



# **SIOS Protection Suite for Linux**

9.6.1 — Last update: 8 June 2022

SIOS TECHNOLOGY CORP.

# Table of Contents

<b>1. LifeKeeper for Linux</b> .....	<b>25</b>
<b>2. LifeKeeper for Linux Release Notes</b> .....	<b>26</b>
<b>3. LifeKeeper for Linux Getting Started Guide</b> .....	<b>60</b>
<b>4. LifeKeeper for Linux Installation Guide</b> .....	<b>75</b>
4.1. Software Packaging .....	76
4.2. Planning Your LifeKeeper Environment .....	78
4.2.1. Mapping Server Configurations .....	79
4.2.2. Storage and Adapter Requirements .....	81
4.2.3. Storage and Adapter Options .....	82
4.3. Setting Up Your LifeKeeper Environment .....	83
4.3.1. Installing the Linux OS and Associated Communication Packages .....	84
4.3.2. Linux Dependencies.....	85
4.3.3. Connecting Servers and Shared Storage .....	89
4.3.4. Configuring Shared Storage .....	90
4.3.5. Verifying Network Configuration.....	91
4.3.6. Creating a Switchable IP Address .....	93
4.3.7. Installing and Setting Up Database Applications .....	94
4.3.8. Configuring GUI Users .....	95
4.3.9. Licensing .....	97
4.3.9.1. Obtaining an Internet HOST ID .....	101
4.4. Installing the Software .....	102
4.5. How to Use Setup Scripts.....	105
4.6. Verifying the LifeKeeper Installation .....	111
4.7. Upgrading LifeKeeper .....	112
<b>5. LifeKeeper for Linux Technical Documentation</b> .....	<b>115</b>
5.1. Introduction .....	116
5.2. Documentation and Training .....	117
5.3. Ikbackup.....	119
5.4. LifeKeeper .....	121
5.4.1. SIOS LifeKeeper for Linux Introduction .....	122
5.4.1.1. Protected Resources .....	124
5.4.1.2. LifeKeeper Core .....	125
5.4.1.3. Configuration Concepts .....	128
5.4.1.3.1. Common Hardware Components .....	129
5.4.1.3.2. System Grouping Arrangements .....	131
5.4.1.3.3. Active – Active Grouping .....	132
5.4.1.3.4. Active – Standby Grouping .....	133
5.4.1.3.5. Intelligent Versus Automatic Switchback.....	134
5.4.1.3.6. Logging With syslog .....	135
5.4.1.3.7. Resource Hierarchies .....	136
5.4.1.3.7.1. Resource Types .....	137

5.4.1.3.7.2. Resource States.....	138
5.4.1.3.7.3. Hierarchy Relationships .....	139
5.4.1.3.7.4. Shared Equivalencies.....	140
5.4.1.3.7.5. Resource Hierarchy Information .....	141
5.4.1.3.7.6. Resource Hierarchy Example .....	142
5.4.1.3.7.7. Detailed Status Display .....	143
5.4.1.3.7.8. Short Status Display.....	149
5.4.1.4. Fault Detection and Recovery Scenarios .....	151
5.4.1.4.1. IP Local Recovery .....	152
5.4.1.4.2. Resource Error Recovery Scenario.....	153
5.4.1.4.3. Server Failure Recovery Scenario .....	155
5.4.2. Installation and Configuration.....	157
5.4.2.1. LifeKeeper Configuration Steps .....	158
5.4.2.1.1. Set Up TTY Connections .....	160
5.4.2.2. LifeKeeper Event Forwarding via SNMP .....	161
5.4.2.2.1. Overview of LifeKeeper Event Forwarding via SNMP.....	162
5.4.2.2.2. Configuring LifeKeeper Event Forwarding.....	166
5.4.2.2.3. SNMP Troubleshooting.....	168
5.4.2.3. LifeKeeper Event Email Notification .....	169
5.4.2.3.1. Overview of LifeKeeper Event Email Notification .....	170
5.4.2.3.2. Configuring LifeKeeper Event Email Notification.....	172
5.4.2.3.3. Email Notification Troubleshooting .....	174
5.4.2.4. Optional Configuration Tasks.....	175
5.4.2.4.1. Confirm Failover and Block Resource Failover Settings.....	176
5.4.2.4.2. Setting Server Shutdown Strategy.....	184
5.4.2.4.3. Tuning the LifeKeeper Heartbeat.....	185
5.4.2.4.4. Using Certificates with the LifeKeeper API.....	188
5.4.2.5. Linux Configuration.....	190
5.4.2.6. Data Replication Configuration .....	194
5.4.2.7. Network Configuration .....	195
5.4.2.8. Application Configuration .....	196
5.4.2.9. Storage and Adapter Configuration.....	197
5.4.2.10. LifeKeeper I/O Fencing Introduction.....	233
5.4.2.10.1. SCSI Reservations .....	234
5.4.2.10.2. Disabling Reservations .....	236
5.4.2.10.2.1. I/O Fencing Chart.....	238
5.4.2.10.3. Quorum/Witness .....	240
5.4.2.10.3.1. Majority Mode.....	245
5.4.2.10.3.2. tcp_remote Mode .....	250
5.4.2.10.3.3. Storage Mode.....	251
5.4.2.10.3.4. Quorum/Witness Cluster Recommendations in AWS.....	259
5.4.2.10.4. STONITH.....	265
5.4.2.10.5. Watchdog .....	270
5.4.2.10.6. I/O Fencing Mechanisms .....	273
5.4.2.10.6.1. Available I/O Fencing Mechanisms (Physical Servers) .....	274

5.4.2.10.6.2. Available I/O Fencing Mechanisms (Virtual Machines in VMware) .....	278
5.4.2.11. Resource Policy Management .....	283
5.4.2.12. Configuring Credentials .....	288
5.4.2.13. Standby Node Health Check .....	290
5.4.2.13.1. Node Monitoring .....	291
5.4.2.13.2. OSU Resource Monitoring .....	292
5.4.3. LifeKeeper Administration Overview.....	293
5.4.3.1. Error Detection and Notification .....	295
5.4.3.2. N-Way Recovery.....	296
5.4.3.3. Administrator Tasks .....	297
5.4.3.3.1. Editing Server Properties.....	298
5.4.3.3.2. Creating a Communication Path .....	299
5.4.3.3.3. Deleting a Communication Path.....	301
5.4.3.3.4. Server Properties – Failover .....	302
5.4.3.3.5. Creating Resource Hierarchies.....	304
5.4.3.3.5.1. Creating a File System Resource Hierarchy .....	306
5.4.3.3.5.2. Creating a Generic Application Resource Hierarchy .....	308
5.4.3.3.5.3. Creating a Raw Device Resource Hierarchy .....	310
5.4.3.3.5.4. Quick Service Protection (QSP) Recovery Kit.....	311
5.4.3.3.6. Editing Resource Properties .....	317
5.4.3.3.7. Editing Resource Priorities .....	318
5.4.3.3.8. Extending Resource Hierarchies.....	320
5.4.3.3.8.1. Extending a File System Resource Hierarchy .....	322
5.4.3.3.8.2. Extending a Generic Application Resource Hierarchy.....	323
5.4.3.3.8.3. Extending a Raw Device Resource Hierarchy.....	324
5.4.3.3.9. Unextending a Hierarchy .....	325
5.4.3.3.10. Creating a Resource Dependency .....	327
5.4.3.3.11. Deleting a Resource Dependency.....	329
5.4.3.3.12. Deleting a Hierarchy from All Servers .....	330
5.4.4. User Guide .....	331
5.4.4.1. Using LifeKeeper for Linux.....	332
5.4.4.1.1. GUI .....	333
5.4.4.1.1.1. GUI Overview – General .....	334
5.4.4.1.1.1.1. LifeKeeper GUI Software Package.....	335
5.4.4.1.1.2. Menus .....	336
5.4.4.1.1.2.1. Resource Context Menu.....	337
5.4.4.1.1.2.2. Server Context Menu .....	339
5.4.4.1.1.2.3. File Menu .....	340
5.4.4.1.1.2.4. Edit Menu – Resource.....	341
5.4.4.1.1.2.5. Edit Menu – Server .....	342
5.4.4.1.1.2.6. View Menu .....	343
5.4.4.1.1.2.7. Help Menu .....	345
5.4.4.1.1.3. Toolbars.....	346
5.4.4.1.1.3.1. GUI Toolbar .....	347
5.4.4.1.1.3.2. Resource Context Toolbar .....	349

5.4.4.1.1.3.3. Server Context Toolbar .....	350
5.4.4.1.1.4. Preparing to Run the GUI .....	351
5.4.4.1.1.4.1. Configuring the LifeKeeper GUI .....	352
5.4.4.1.1.4.2. Starting and Stopping the GUI Server .....	355
5.4.4.1.1.4.3. Java Security Policy .....	357
5.4.4.1.1.4.4. Running the GUI on a LifeKeeper Server .....	360
5.4.4.1.1.4.5. Lifekeeper GUI Overview .....	361
5.4.4.1.2. Status Table .....	362
5.4.4.1.3. Properties Panel .....	363
5.4.4.1.4. Output Panel .....	364
5.4.4.1.5. Message Bar .....	365
5.4.4.1.6. Exiting the GUI .....	366
5.4.4.1.7. Common Tasks .....	367
5.4.4.1.7.1. Starting LifeKeeper .....	368
5.4.4.1.7.2. Stopping LifeKeeper .....	370
5.4.4.1.7.3. Viewing LifeKeeper Processes .....	371
5.4.4.1.7.4. Viewing LifeKeeper GUI Server Processes .....	373
5.4.4.1.7.5. Viewing LifeKeeper Controlling Processes .....	374
5.4.4.1.7.6. Connecting Servers to a Cluster .....	375
5.4.4.1.7.7. Disconnecting from a Cluster .....	376
5.4.4.1.7.8. Viewing Connected Servers .....	377
5.4.4.1.7.9. Viewing the Status of a Server .....	378
5.4.4.1.7.10. Viewing Server Properties .....	379
5.4.4.1.7.11. Viewing Server Log Files .....	380
5.4.4.1.7.12. Viewing Resource Tags and IDs .....	381
5.4.4.1.7.13. Viewing the Status of Resources .....	382
5.4.4.1.7.14. Viewing Resource Properties .....	384
5.4.4.1.7.15. Resource Labels .....	385
5.4.4.1.7.16. Viewing Message History .....	386
5.4.4.1.7.17. Expanding and Collapsing a Resource Hierarchy Tree .....	387
5.4.4.1.7.18. Cluster Connect Dialog .....	389
5.4.4.1.7.19. Cluster Disconnect Dialog .....	390
5.4.4.1.7.20. Resource Properties Dialog .....	391
5.4.4.1.7.21. Server Properties Dialog .....	393
5.4.4.1.8. Operator Tasks .....	397
5.4.4.1.8.1. Bringing a Resource In Service .....	398
5.4.4.1.8.2. Taking a Resource Out of Service .....	399
5.4.4.1.9. Advanced Tasks .....	400
5.4.4.1.9.1. LCD .....	401
5.4.4.1.9.1.1. LCDI Commands .....	402
5.4.4.1.9.1.2. LCD Configuration Data .....	406
5.4.4.1.9.1.3. LCD Directory Structure .....	407
5.4.4.1.9.1.4. LCD Resource Types .....	408
5.4.4.1.9.1.5. LifeKeeper Flags .....	409
5.4.4.1.9.1.6. Resources Subdirectories .....	410

5.4.4.1.9.1.7. Structure of LCD Directory in /opt/LifeKeeper.....	412
5.4.4.1.9.2. LCM .....	413
5.4.4.1.9.2.1. Communication Status Information.....	414
5.4.4.1.9.2.2. LifeKeeper Alarming and Recovery .....	415
5.4.4.1.9.3. LifeKeeper API for Monitoring .....	417
5.4.4.1.10. Maintenance Tasks.....	429
5.4.4.1.10.1. Changing LifeKeeper Configuration Values .....	430
5.4.4.1.10.2. File System Health Monitoring .....	433
5.4.4.1.10.3. Maintaining a LifeKeeper Protected System .....	436
5.4.4.1.10.4. Maintaining a Resource Hierarchy .....	437
5.4.4.1.10.5. Recovering After a Failover .....	438
5.4.4.1.10.6. Removing LifeKeeper .....	439
5.4.4.1.10.7. Running LifeKeeper With a Firewall.....	440
5.4.4.1.10.8. Running the LifeKeeper GUI Through a Firewall.....	442
5.4.4.1.10.9. Transferring Resource Hierarchies .....	444
5.4.4.1.11. Technical Notes.....	445
5.4.4.2. Cluster Example .....	450
5.4.4.3. Dialogs .....	451
5.4.5. Troubleshooting .....	461
5.4.5.1. Solutions .....	463
5.4.5.2. Common Causes of a LifeKeeper Initiated Failover.....	470
5.4.5.3. Known Issues and Restrictions .....	475
5.4.5.3.1. Installation – Known Issues / Restrictions .....	476
5.4.5.3.2. LifeKeeper Core – Known Issues / Restrictions .....	479
5.4.5.3.3. Internet/IP Licensing – Known Issues / Restrictions .....	487
5.4.5.3.4. GUI – Known Issues / Restrictions.....	488
5.4.5.3.5. Data Replication – Known Issues / Restrictions .....	490
5.4.5.3.6. IPv6 – Known Issues / Restrictions.....	493
5.4.5.3.7. Apache – Known Issues / Restrictions .....	495
5.4.5.3.8. Oracle – Known Issues / Restrictions .....	496
5.4.5.3.9. MySQL – Known Issues / Restrictions .....	497
5.4.5.3.10. NAS Recovery Kit – Known Issues / Restrictions .....	498
5.4.5.3.11. NFS Server – Known Issues / Restrictions.....	499
5.4.5.3.12. SAP Recovery Kit – Known Issues / Restrictions .....	501
5.4.5.3.13. LVM – Known Issues / Restrictions.....	503
5.4.5.3.14. Multipath Recovery Kits (DMMP / HDLM / PPATH /NECSPS) Known Issues / Restrictions .....	504
5.4.5.3.15. DMMP – Known Issues / Restrictions .....	505
5.4.5.3.16. DB2 – Known Issues / Restrictions .....	506
5.4.5.3.17. Sybase ASE – Known Issues / Restrictions .....	507
5.4.5.3.18. WebSphere MQ – Known Issues / Restrictions .....	510
5.4.5.3.19. SAP HANA – Known Issues / Restrictions .....	511
5.4.5.3.20. EC2 Recovery Kit Known Issues / Restrictions .....	512
5.4.5.3.21. Known Issues/Restrictions when using LifeKeeper on Oracle Cloud Infrastructure (OCI) .....	513

5.4.5.4. Communication Paths Going Up and Down.....	514
5.4.5.5. Incomplete Resource Created.....	515
5.4.5.6. Incomplete Resource Priority Modification .....	516
5.4.5.7. No Shared Storage Found When Configuring a Hierarchy .....	518
5.4.5.8. Recovering from a LifeKeeper Server Failure .....	520
5.4.5.9. Recovering from a Non-Killable Process .....	521
5.4.5.10. Recovering from a Panic during a Manual Recovery .....	522
5.4.5.11. Recovering Out-of-Service Hierarchies .....	523
5.4.5.12. Resource Tag Name Restrictions.....	524
5.4.5.13. Serial (TTY) Console WARNING.....	525
5.4.5.14. Taking the System to init state S WARNING .....	526
5.4.5.15. Thread is Hung Messages on Shared Storage .....	527
5.5. DataKeeper.....	528
5.5.1. Mirroring with SIOS DataKeeper for Linux.....	529
5.5.2. How SIOS DataKeeper Works.....	531
5.5.3. SIOS DataKeeper Installation and Configuration .....	541
5.5.3.1. Hardware and Software Requirements.....	543
5.5.3.2. General Configuration.....	545
5.5.3.3. DataKeeper for Linux Network Configuration .....	546
5.5.3.4. DataKeeper Events Table .....	547
5.5.3.5. Changing the Data Replication Path .....	548
5.5.3.6. Network Bandwidth Requirements .....	549
5.5.3.6.1. Measuring Rate of Change on a Linux System (Physical or Virtual).....	550
5.5.3.7. WAN Configuration .....	560
5.5.3.8. SIOS DataKeeper for Linux Resource Types .....	561
5.5.3.9. I/O Fencing with DataKeeper Configuration .....	563
5.5.3.10. Resource Configuration Tasks .....	564
5.5.3.10.1. Creating a DataKeeper Resource Hierarchy .....	565
5.5.3.10.1.1. Replicate New File System.....	567
5.5.3.10.1.2. Replicate Existing File System .....	570
5.5.3.10.1.3. DataKeeper Resource .....	572
5.5.3.10.2. Extending Your DataKeeper Hierarchy .....	575
5.5.3.10.3. Unextending Your DataKeeper Hierarchy .....	578
5.5.3.10.4. Deleting a DataKeeper Resource Hierarchy.....	579
5.5.3.10.5. Taking a DataKeeper Resource Out of Service.....	580
5.5.3.10.6. Bringing a DataKeeper Resource In Service .....	581
5.5.3.10.7. Testing Your DataKeeper Resource Hierarchy.....	582
5.5.4. Administering SIOS DataKeeper for Linux.....	583
5.5.4.1. Viewing Mirror Status.....	584
5.5.4.2. GUI Mirror Administration .....	586
5.5.4.2.1. Pause and Resume .....	588
5.5.4.2.2. Set Compression Level.....	589
5.5.4.3. Command Line Mirror Administration .....	590
5.5.4.4. Monitoring Mirror Status via Command Line .....	594
5.5.4.5. Server Failure .....	596

5.5.4.6. Resynchronization .....	597
5.5.4.7. Avoiding Full Resynchronizations .....	598
5.5.4.8. Verify Data Before Resync (Wait to Resync) .....	603
5.5.5. Using LVM with DataKeeper .....	609
5.5.6. Clustering with Fusion-io .....	610
5.5.7. Using External Snapshot Functions for Disks and Devices Protected by DataKeeper.....	613
5.5.8. DataKeeper for Linux Troubleshooting .....	614
5.6. Command Line Interface .....	621
5.6.1. Commands .....	623
5.6.1.1. Iklogmsg .....	627
5.6.1.2. SYS – LifeKeeper Commands Related to the Systems in the LifeKeeper Cluster ...	629
5.6.1.3. NET – Communication Paths Related Commands .....	631
5.6.1.4. FLAG – Commands Related to Internal LifeKeeper Flags .....	633
5.6.1.5. TYP – LifeKeeper Commands Related to Resource Hierarchy Types .....	634
5.6.1.6. APP – LifeKeeper Commands Related to Resource Applications (Group of Related Types) .....	635
5.6.1.7. DEP – LifeKeeper Commands Related to How Resource Applications Relate to Each Other .....	636
5.6.1.8. INS – Commands Related to Individual LifeKeeper Hierarchy Instances .....	643
5.6.1.8.1. Unextend a Hierarchy .....	645
5.6.1.9. Accessing man (Manual) pages .....	646
5.6.2. LKCLI (LifeKeeper Command Line Interface) .....	648
5.6.2.1. LKCLI Subcommands for Each ARK .....	662
5.6.3. Setting up LifeKeeper with Ansible .....	682
5.6.4. LKCLI Guide .....	690
5.6.4.1. LKCLI – Communication Path Creation and Deletion .....	691
5.6.4.2. LKCLI – Resource Creation .....	696
5.6.4.2.1. Creating a File System Resource .....	697
5.6.4.2.2. Creating an IP Resource .....	700
5.6.4.2.3. Creating a PostgreSQL Resource .....	703
5.6.4.2.4. Creating a DataKeeper Resource .....	706
5.6.4.3. LKCLI – Checking Cluster Status .....	710
5.6.4.4. LKCLI – Verifying Switchover Behavior .....	711
5.6.4.5. LKCLI – Maintenance Tasks .....	713
5.6.4.6. LKCLI – Replicate the Existing Cluster Settings .....	718
<b>6. Application Recovery Kits .....</b>	<b>722</b>
6.1. Apache Recovery Kit Administration Guide .....	723
6.1.1. LifeKeeper Documentation and Apache References .....	724
6.1.2. Apache Recovery Kit Requirements .....	725
6.1.3. Configuring Apache Web Server with LifeKeeper .....	726
6.1.3.1. Configuration Definitions and Examples .....	727
6.1.3.1.1. Active/Standby and Active/Active Configurations .....	731
6.1.3.1.2. Configuration Considerations for Apache Web Server .....	732
6.1.4. LifeKeeper Configuration Tasks for Apache .....	736
6.1.4.1. Creating an Apache Web Server Resource Hierarchy .....	737

6.1.4.2. Extending an Apache Web Server Resource Hierarchy.....	739
6.1.4.3. Unextending an Apache Web Server Resource Hierarchy.....	741
6.1.4.4. Deleting an Apache Web Server Resource Hierarchy .....	742
6.1.4.5. Testing an Apache Web Server Resource Hierarchy.....	743
6.1.5. Apache Web Server Troubleshooting .....	744
6.1.5.1. Apache Hierarchy Creation Errors .....	745
6.1.5.2. Apache Extend Hierarchy Errors.....	749
6.1.5.3. Apache Hierarchy Restore, Remove, and Recover Messages and Errors .....	751
6.2. DB2 Recovery Kit Administration Guide .....	755
6.2.1. DB2 Documentation and References .....	756
6.2.2. DB2 Recovery Kit Hardware and Software Requirements .....	757
6.2.3. DB2 Recovery Kit Overview .....	758
6.2.4. Configuring the LifeKeeper for Linux DB2 Recovery Kit.....	759
6.2.4.1. Using DB2 with Raw I/O .....	760
6.2.4.2. Running DB2 .....	761
6.2.4.3. Configuration Considerations for DB2 Single Partition .....	762
6.2.4.4. Configuration Considerations for DB2 Multiple Partition .....	763
6.2.4.4.1. Issues Regarding DB2 EEE or multiple partition ESE and NFS.....	764
6.2.4.4.2. DB2 Configuration Requirements .....	766
6.2.4.5. Configuration Considerations for All DB2 Configurations .....	769
6.2.4.6. DB2 Configuration Examples .....	771
6.2.5. LifeKeeper for Linux DB2 Recovery Kit Configuration Tasks .....	776
6.2.5.1. Creating a DB2 Resource Hierarchy .....	777
6.2.5.2. Deleting a DB2 Resource Hierarchy.....	780
6.2.5.3. Extending Your DB2 Resource Hierarchy .....	782
6.2.5.4. Unextending Your DB2 Resource Hierarchy .....	785
6.2.5.5. Testing Your DB2 Resource Hierarchy.....	786
6.2.6. DB2 Troubleshooting .....	787
6.2.7. Setting Up DB2 to use Raw I/O .....	788
6.3. Recovery Kit for EC2™ Administration Guide.....	791
6.3.1. Recovery Kit for EC2™ Principles of Operation.....	792
6.3.2. Recovery Kit for EC2™ Requirements .....	796
6.3.3. Recovery Kit for EC2™ Configuration .....	798
6.3.3.1. EC2™ Event Table .....	800
6.3.3.2. Adjusting Recovery Kit for EC2™ Tunable Values .....	801
6.3.3.3. Creating an EC2™ Resource Hierarchy.....	802
6.3.3.4. Deleting an EC2™ Resource Hierarchy .....	804
6.3.3.5. Extending Your EC2™ Hierarchy .....	805
6.3.3.6. EC2™ Local Recovery and Configuration .....	807
6.3.3.7. EC2™ Resource Monitoring and Configuration .....	809
6.3.3.8. Unextending Your EC2™ Hierarchy.....	810
6.3.3.9. EC2™ User System Setup.....	811
6.3.4. Recovery Kit for EC2™ Troubleshooting .....	813
6.4. Generic Application Kit for Load Balancer Health Checks.....	814
6.4.1. Configuration Examples .....	815

6.4.2. Basic Behaviors .....	818
6.4.3. Script Specifications.....	819
6.4.4. Script Parameter List .....	821
6.4.5. Creating/Extending a Resource.....	822
6.4.6. Messages List.....	823
6.5. LVM Recovery Kit Administration Guide .....	824
6.5.1. LVM Documentation and References .....	825
6.5.2. LVM Recovery Kit Requirements.....	826
6.5.2.1. LVM Hardware and Software Requirements.....	827
6.5.3. LVM Recovery Kit Overview.....	828
6.5.3.1. LVM Recovery Kit Notes and Restrictions.....	830
6.5.4. LifeKeeper LVM Hierarchy Creation and Administration .....	832
6.5.4.1. LVM Hierarchy Creation Procedures.....	833
6.5.4.2. Using the LVM Recovery Kit with DataKeeper .....	834
6.5.4.3. Volume Group Reconfiguration .....	835
6.5.5. LVM Troubleshooting .....	839
6.6. IP Recovery Kit Administration Guide.....	841
6.6.1. IP Recovery Kit Principles of Operation.....	842
6.6.2. IP Recovery Kit Requirements .....	846
6.6.3. IP Recovery Kit Configuration .....	847
6.6.3.1. IP Interface Selection.....	849
6.6.3.2. IP User System Setup.....	850
6.6.3.3. General IP Planning Considerations .....	851
6.6.3.4. IP Resource Monitoring and Configuration .....	852
6.6.3.5. IP Local Recovery and Configuration .....	853
6.6.3.6. IP Recovery Kit Configuration Examples.....	854
6.6.3.7. Creating an IP Resource Hierarchy.....	861
6.6.3.8. Deleting an IP Resource Hierarchy .....	863
6.6.3.9. Extending Your IP Hierarchy.....	864
6.6.3.10. Unextending Your IP Hierarchy.....	867
6.6.3.11. Testing Your IP Resource Hierarchy .....	868
6.6.3.12. Viewing and Editing IP Configuration Properties .....	869
6.6.3.13. Adjusting IP Recovery Kit Tunable Values .....	881
6.7. MySQL Recovery Kit Administration Guide.....	882
6.7.1. MySQL Recovery Kit Hardware and Software Requirements.....	883
6.7.2. MySQL Recovery Kit Configuration .....	884
6.7.2.1. Configuration Considerations for MySQL .....	885
6.7.2.2. Client Configuration Considerations for MySQL .....	888
6.7.2.3. MySQL Configuration Requirements .....	889
6.7.2.4. MySQL Configuration Examples .....	890
6.7.2.5. Active/Standby MySQL Configuration .....	891
6.7.2.6. Active/Active MySQL Configuration.....	893
6.7.2.7. Multiple Database Server Environment .....	897
6.7.2.8. Using mysqld Groups with LifeKeeper .....	898
6.7.2.9. Using Network Attached Storage .....	902

6.7.2.10. Considerations on MySQL use in systemd Environments.....	906
6.7.3. Installing/Configuring MySQL with LifeKeeper.....	907
6.7.3.1. LifeKeeper Configuration Tasks for MySQL .....	908
6.7.3.2. Creating a MySQL Resource Hierarchy .....	909
6.7.3.3. Deleting a MySQL Resource Hierarchy .....	913
6.7.3.4. Extending Your MySQL Hierarchy.....	915
6.7.3.5. Unextending Your MySQL Hierarchy.....	919
6.7.4. MySQL Administration .....	921
6.7.4.1. Performing a Manual Switchover from the GUI .....	922
6.7.5. MySQL Troubleshooting.....	923
6.8. WebSphere MQ Recovery Kit Administration Guide .....	926
6.8.1. MQ Recovery Kit Abbreviations.....	928
6.8.2. MQ Recovery Kit Requirements .....	929
6.8.2.1. MQ Hardware and Software Requirements .....	930
6.8.2.2. Upgrading an MQ LifeKeeper Cluster .....	932
6.8.3. WebSphere MQ Recovery Kit Overview .....	933
6.8.3.1. MQ Recovery Kit Resource Hierarchies .....	934
6.8.3.2. MQ Recovery Kit Features.....	935
6.8.4. WebSphere MQ Configuration Considerations .....	936
6.8.4.1. MQ Configuration Requirements .....	937
6.8.4.1.1. MQ Supported File System Layouts.....	940
6.8.4.1.1.1. Configuration 1 – /var/mqm on Shared Storage.....	941
6.8.4.1.1.2. Configuration 2 – Direct Mounts .....	942
6.8.4.1.1.3. Configuration 3 – Symbolic Links .....	943
6.8.4.1.2. Configuring WebSphere MQ for use with LifeKeeper .....	944
6.8.4.1.3. MQ Configuration Changes After Resource Creation.....	949
6.8.4.1.3.1. Relocating QMDIR and QMLOGDIR.....	950
6.8.4.1.3.2. Changing the Listener Port.....	951
6.8.4.1.3.3. Changing the IP for the Queue Manager .....	952
6.8.4.1.4. WebSphere MQ Configuration Examples.....	953
6.8.4.1.4.1. Active/Standby Configuration with /var/mqm on Shared Storage .....	954
6.8.4.1.4.2. Active/Standby Configuration with NAS Storage.....	955
6.8.4.1.4.3. Active/Active Configuration with Shared Storage.....	957
6.8.4.1.4.4. Active/Active Configuration with NAS Storage.....	959
6.8.5. LifeKeeper Configuration Tasks for MQ.....	961
6.8.5.1. Creating a WebSphere MQ Resource Hierarchy .....	963
6.8.5.2. Extending a WebSphere MQ Hierarchy.....	965
6.8.5.3. Unextending a WebSphere MQ Hierarchy.....	967
6.8.5.4. Deleting a WebSphere MQ Hierarchy .....	968
6.8.5.5. Testing a WebSphere MQ Resource Hierarchy.....	969
6.8.5.5.1. Testing MQ Client Connectivity.....	971
6.8.5.6. Viewing MQ Resource Properties .....	973
6.8.5.7. Editing MQ Configuration Resource Properties .....	974
6.8.5.7.1. Enable/Disable Listener Protection .....	979
6.8.5.7.2. Changing the LifeKeeper Test Queue Name.....	980

6.8.5.7.3. Changing the Log Level .....	981
6.8.5.7.4. Changing Shutdown Timeout Values .....	982
6.8.5.7.5. Changing the Server Connection Channel .....	984
6.8.5.7.6. Changing the Command Server Protection Configuration .....	985
6.8.5.7.7. Changing LifeKeeper WebSphere MQ Recovery Kit Defaults .....	987
6.8.6. WebSphere MQ Troubleshooting .....	989
6.8.6.1. MQ Error Messages .....	990
6.8.7. Appendix A – Sample mqs.ini Configuration File .....	998
6.8.8. Appendix B – Sample qm.ini Configuration File .....	1000
6.8.9. Appendix C – WebSphere MQ Configuration Sheet .....	1001
6.9. NAS Recovery Kit Administration Guide .....	1004
6.9.1. NAS Documentation and References .....	1005
6.9.2. NAS Recovery Kit Hardware and Software Requirements .....	1006
6.9.3. NAS Recovery Kit Overview .....	1007
6.9.4. Configuring the LifeKeeper for Linux NAS Recovery Kit .....	1009
6.9.4.1. NAS Configuration Considerations .....	1010
6.9.4.2. NAS Configuration Examples .....	1012
6.9.5. LifeKeeper Configuration Tasks for NAS .....	1014
6.9.5.1. Creating a NAS Resource Hierarchy .....	1015
6.9.5.2. Deleting a NAS Resource Hierarchy .....	1018
6.9.5.3. Extending Your NAS Hierarchy .....	1020
6.9.5.4. Unextending Your NAS Hierarchy .....	1023
6.9.5.5. Testing Your NAS Resource Hierarchy .....	1024
6.9.6. NAS Troubleshooting .....	1025
6.9.6.1. NAS Error Messages .....	1026
6.9.6.2. LifeKeeper GUI Related Errors .....	1028
6.10. NFS Server Recovery Kit Administration Guide .....	1029
6.10.1. NFS Server Recovery Kit Overview .....	1030
6.10.2. NFS Server Recovery Kit Requirements .....	1031
6.10.3. NFS Server Recovery Kit Configuration Considerations .....	1032
6.10.3.1. Configuring NFS Server with LifeKeeper .....	1033
6.10.3.2. NFS Specific Configuration Considerations .....	1037
6.10.3.3. NFS Configuration Examples .....	1038
6.10.3.3.1. Active – Standby Configuration .....	1039
6.10.3.3.2. Active – Active Configuration .....	1041
6.10.4. NFS Configuration Tasks .....	1043
6.10.4.1. Creating an NFS Resource Hierarchy .....	1044
6.10.4.2. Deleting an NFS Resource Hierarchy .....	1047
6.10.4.3. Extending Your NFS Hierarchy .....	1049
6.10.4.4. Unextending Your NFS Hierarchy .....	1053
6.10.4.5. Testing Your NFS Hierarchy .....	1055
6.10.5. NFS Troubleshooting .....	1057
6.10.5.1. HA nfs-utils Installation and Configuration .....	1058
6.10.5.2. NFS Hierarchy Creation Errors .....	1059
6.10.5.3. NFS Extend Hierarchy Errors .....	1062

6.10.5.4. NFS Hierarchy Restore, Remove and Recover Messages and Errors .....	1063
6.10.5.5. NFS Hierarchy Delete Messages and Errors .....	1066
6.11. Recovery Kit for Oracle Cloud Infrastructure Administration Guide .....	1067
6.11.1. Principles of Operation .....	1068
6.11.2. Resource Monitoring and Local Recovery .....	1069
6.11.3. Requirements .....	1070
6.11.4. Recovery Kit for Oracle Cloud Infrastructure Notes .....	1072
6.11.5. Restrictions when using LifeKeeper on Oracle Cloud Infrastructure (OCI) .....	1073
6.11.6. Configuration .....	1074
6.11.6.1. Creating a Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy .....	1075
6.11.6.2. Deleting a Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy .....	1077
6.11.6.3. Extending Your OCI Resource Hierarchy .....	1078
6.11.6.4. Unextending Your OCI Hierarchy .....	1080
6.11.6.5. Adjusting Recovery Kit for Oracle Cloud Infrastructure Tunable Values .....	1081
6.11.7. Troubleshooting .....	1082
6.11.7.1. Known Issues / Restrictions .....	1083
6.11.7.2. Error Messages .....	1084
6.12. Oracle Recovery Kit Administration Guide .....	1085
6.12.1. Oracle Recovery Kit Hardware and Software Requirements .....	1086
6.12.2. Configuring Oracle with LifeKeeper .....	1088
6.12.2.1. Specific Configuration Considerations for Oracle .....	1089
6.12.2.2. Configuring the Oracle Net Listener for LifeKeeper Protection .....	1094
6.12.2.3. Configuring Transparent Application Failover with LifeKeeper .....	1098
6.12.2.4. Configuring a Pluggable Database with Oracle Multitenant .....	1100
6.12.2.5. Oracle Configuration Examples .....	1105
6.12.2.5.1. Oracle Configuration Requirements .....	1106
6.12.2.5.2. Oracle Active/Standby Configurations .....	1107
6.12.2.5.3. Oracle Active/Active Configurations .....	1110
6.12.3. LifeKeeper Configuration Tasks for Oracle .....	1114
6.12.3.1. Creating an Oracle Resource Hierarchy .....	1116
6.12.3.2. Deleting an Oracle Resource Hierarchy .....	1118
6.12.3.3. Extending Your Oracle Hierarchy .....	1120
6.12.3.4. Unextending Your Oracle Hierarchy .....	1124
6.12.3.5. Viewing Oracle Configuration Settings .....	1126
6.12.3.6. Changing Username / Password for the Oracle Database Account .....	1127
6.12.3.7. Testing Your Oracle Resource Hierarchy .....	1129
6.12.3.8. Patching Oracle Nodes (SAP/Oracle) with DataKeeper .....	1130
6.12.4. Oracle Troubleshooting .....	1131
6.12.4.1. Oracle Known Issues and Restrictions .....	1132
6.12.4.1.1. Oracle Database Creation Problems .....	1137
6.12.4.1.2. Oracle Database Startup Problems .....	1138
6.12.4.1.3. inqfail error in the LifeKeeper Log .....	1139
6.12.5. Oracle Appendix .....	1140
6.12.5.1. Setting up Oracle to Use Raw I/O .....	1141
6.12.5.1.1. Adding a Tablespace After Creating Hierarchy .....	1144

6.12.5.2. Creating an Oracle Listener for Multiple Resources .....	1145
6.12.5.2.1. Updating the Oracle Listener Protection Level .....	1148
6.12.5.2.2. Updating the Oracle Listener Recovery Level .....	1149
6.12.5.2.3. Updating the Oracle Protected Listener(s) .....	1150
6.12.5.3. Migrating a Pluggable Database .....	1151
6.13. PostgreSQL Recovery Kit Administration Guide .....	1154
6.13.1. PostgreSQL Resource Hierarchy .....	1155
6.13.2. PostgreSQL Hardware and Software Requirements .....	1156
6.13.3. PostgreSQL Configuration Considerations .....	1157
6.13.3.1. Protecting PostgreSQL Best Practices .....	1158
6.13.3.2. Using Mirrored File Systems with DataKeeper .....	1159
6.13.4. PostgreSQL Installation .....	1160
6.13.4.1. Install the PostgreSQL Software .....	1162
6.13.4.2. Creating a PostgreSQL Database .....	1163
6.13.4.3. Install the LifeKeeper Software .....	1165
6.13.4.4. LifeKeeper Tunable Settings for PostgreSQL .....	1166
6.13.4.5. Creating a PostgreSQL Resource Hierarchy .....	1168
6.13.4.6. Deleting a PostgreSQL Resource Hierarchy .....	1170
6.13.4.7. Extending a PostgreSQL Resource Hierarchy .....	1171
6.13.4.8. Unextending a PostgreSQL Resource Hierarchy .....	1173
6.13.4.9. Viewing PostgreSQL Configuration Settings .....	1174
6.13.4.10. Upgrading PostgreSQL .....	1175
6.13.5. PostgreSQL Administration .....	1178
6.13.5.1. Performing a Manual Switchover from the LifeKeeper GUI .....	1179
6.13.5.2. Protecting EnterpriseDB Postgres Plus Advanced Server .....	1180
6.13.5.3. Protecting Symfoware Server/Enterprise Postgres .....	1181
6.13.5.4. Updating Database Administrator User .....	1182
6.13.6. PostgreSQL Troubleshooting .....	1183
6.13.6.1. PostgreSQL General Tips .....	1184
6.13.6.2. PostgreSQL Tunables .....	1186
6.14. Postfix Recovery Kit Administration Guide .....	1188
6.14.1. Postfix Hardware and Software Requirements .....	1190
6.14.1.1. Postfix Recovery Kit Installation .....	1191
6.14.2. Configuring the LifeKeeper for Linux Postfix Recovery Kit .....	1192
6.14.2.1. Postfix Protection Objects .....	1193
6.14.2.2. Postfix Configuration Requirements .....	1194
6.14.2.3. Port and TCP Interface Definition and the Postfix Recovery Kit .....	1196
6.14.2.4. DNS, Postfix and LifeKeeper .....	1197
6.14.2.5. Postfix Configuration Examples .....	1198
6.14.3. Postfix Configuration Validation .....	1202
6.14.3.1. LifeKeeper Configuration Tasks for Postfix .....	1204
6.14.3.1.1. Creating a Postfix Resource Hierarchy .....	1205
6.14.3.1.2. Extending a Postfix Resource Hierarchy .....	1207
6.14.3.1.3. Unextending a Postfix Resource Hierarchy .....	1209
6.14.3.1.4. Deleting a Postfix Resource Hierarchy .....	1210

6.14.3.1.5. Create Dependency with Mailbox Spool Resource .....	1211
6.14.3.1.6. Testing Your Postfix Resource Hierarchy .....	1212
6.14.4. Postfix Troubleshooting .....	1214
6.14.4.1. Postfix Hierarchy Creation Error Messages .....	1215
6.14.4.2. Postfix Hierarchy Extend Error Messages .....	1216
6.14.4.3. Postfix Resource In-Service / Out-of-Service / Health Monitoring Error	
Messages .....	1217
6.15. Recovery Kit for Route 53™ Administration Guide .....	1218
6.15.1. Recovery Kit for Route 53™ Requirements .....	1219
6.15.2. Recovery Kit for Route 53™ Configuration .....	1220
6.15.2.1. Creating a Route53™ Resource Hierarchy .....	1222
6.15.2.2. Deleting a Route53™ Resource Hierarchy .....	1224
6.15.2.3. Extending Your Route53™ Resource Hierarchy .....	1225
6.15.2.4. Unextending Your Route53™ Resource Hierarchy .....	1228
6.15.2.5. Adjusting Recovery Kit for Route 53™ Tunable Values .....	1229
6.15.2.6. Route53™ Resource Monitoring and Recovery .....	1230
6.15.2.7. Route53™ User System Setup .....	1231
6.15.3. Recovery Kit for Route 53™ Troubleshooting .....	1232
6.16. Samba Recovery Kit Administration Guide .....	1233
6.16.1. Samba Recovery Kit Requirements .....	1234
6.16.2. Samba Recovery Kit Installation .....	1235
6.16.3. Samba Recovery Kit Overview .....	1236
6.16.4. Configuring Samba with LifeKeeper .....	1237
6.16.4.1. The Samba Configuration File .....	1238
6.16.4.2. [Global] Section of the Configuration File .....	1239
6.16.4.3. [Homes] Section of the Configuration File .....	1241
6.16.4.4. [Printers] Section of the Configuration File .....	1242
6.16.4.5. Share Definition Sections of the Configuration File .....	1243
6.16.4.6. Running Multiple Instances of Samba .....	1244
6.16.4.7. Samba Configuration Examples .....	1246
6.16.5. Samba Configuration Steps .....	1251
6.16.6. LifeKeeper Configuration Tasks for Samba .....	1253
6.16.6.1. Creating a Samba Resource Hierarchy .....	1254
6.16.6.2. Extending Your Samba Resource Hierarchy .....	1256
6.16.6.3. Unextending Your Samba Resource Hierarchy .....	1258
6.16.6.4. Deleting a Samba Resource Hierarchy .....	1259
6.16.6.5. Testing Your Samba Resource Hierarchy .....	1260
6.16.7. Samba Hierarchy Administration .....	1261
6.16.7.1. Modifying the Samba Configuration File .....	1262
6.16.7.2. Maintaining the smvpasswd File .....	1265
6.16.8. Samba Troubleshooting .....	1266
6.16.8.1. Common Samba Error Messages .....	1267
6.16.8.2. Hierarchy Creation .....	1268
6.16.8.3. Hierarchy Extension .....	1272
6.16.8.4. Restore .....	1273

6.16.8.5. Remove .....	1274
6.16.8.6. Resource Monitoring .....	1275
6.16.8.7. Configuration File Synchronization Utility .....	1276
6.17. SAP Recovery Kit Administration Guide .....	1277
6.17.1. SAP Abbreviations and Definitions .....	1278
6.17.2. LifeKeeper – SAP Icons .....	1280
6.17.3. SAP Recovery Kit Overview .....	1281
6.17.4. LifeKeeper SAP Solution Page.....	1284
6.17.5. SAP Hardware and Software Requirements .....	1286
6.17.6. SAP Configuration Considerations .....	1288
6.17.6.1. ABAP+Java Configuration (ASCS and SCS) .....	1290
6.17.6.2. ABAP SCS (ASCS) .....	1293
6.17.6.3. Java Only Configuration (SCS) .....	1294
6.17.6.4. SAP Directory Structure .....	1296
6.17.6.5. SAP Virtual Server Name.....	1302
6.17.6.6. SAP Health Monitoring.....	1303
6.17.6.7. SAP License .....	1305
6.17.6.8. SAP Automatic Switchback.....	1306
6.17.6.9. Notes – Special Configuration Steps.....	1307
6.17.7. SAP Installation .....	1308
6.17.7.1. Plan Your SAP Configuration .....	1311
6.17.7.2. Installation of the Core Services .....	1313
6.17.7.3. Installation of the Database .....	1314
6.17.7.4. Installation of the Primary Application Server Instance .....	1315
6.17.7.5. Installation of Additional Application Server Instances .....	1316
6.17.7.6. Installation on the Backup Server.....	1317
6.17.7.7. Install LifeKeeper .....	1318
6.17.7.8. Create File Systems and Directory Structure .....	1320
6.17.7.9. Move Data to Shared Disk and LifeKeeper.....	1322
6.17.7.10. Modify ASCS and ERS Instance Profile Settings.....	1329
6.17.7.11. Upgrading from a Previous Version of the SAP Recovery Kit .....	1332
6.17.7.12. SAP IP Resources .....	1334
6.17.7.13. Creating an SAP Resource Hierarchy .....	1335
6.17.7.14. Deleting an SAP Resource Hierarchy.....	1346
6.17.7.15. Common SAP Recovery Kit Tasks .....	1347
6.17.7.16. Setting Up SAP from the Command Line .....	1348
6.17.7.17. Activating the SAP SIOS HA Cluster Connector (SSHCC) .....	1350
6.17.7.18. SAP Test Preparation .....	1352
6.17.7.19. Perform SAP Tests .....	1353
6.17.8. SAP Administration .....	1355
6.17.8.1. NFS Considerations.....	1356
6.17.8.2. SAP Client Reconnect .....	1358
6.17.8.3. Adjusting SAP Recovery Kit Tunable Values.....	1359
6.17.8.4. Separation of SAP and NFS Hierarchies .....	1361
6.17.8.5. Update SAP Protection Level.....	1362

6.17.8.6. Update SAP Recovery Level.....	1364
6.17.8.7. View SAP Properties.....	1365
6.17.8.8. Special Considerations for Oracle.....	1366
6.17.8.9. SSHCC HA Actions.....	1367
6.17.8.10. ERS Resource Types in LifeKeeper.....	1368
6.17.8.11. Upgrading from ENSAv1 to ENSAv2.....	1374
6.17.8.12. Upgrading from ERSv1 to ERSv2.....	1376
6.17.8.13. Automatic Mounting of Critical NFS Shares.....	1380
6.17.9. SAP Troubleshooting.....	1382
6.17.9.1. Disable Autostart in ERS Profile.....	1383
6.17.9.2. ASCS + ERS Restart_Program Parameter.....	1384
6.17.9.3. SAP Hierarchy Remove Errors.....	1385
6.17.9.4. SAP Hierarchy Restore Errors.....	1386
6.17.9.5. SAP Error Messages During Failover or In-Service.....	1388
6.17.9.6. SAP Installation Errors.....	1389
6.17.9.7. Troubleshooting sapinit.....	1390
6.17.9.8. tset Errors Appear in the LifeKeeper Log File.....	1391
6.17.10. Maintenance Mode.....	1393
6.17.10.1. SAP Maintenance Mode.....	1394
6.17.10.2. Custom and Maintenance-Mode Behavior via Policies.....	1398
6.18. SAP HANA Recovery Kit Administration Guide.....	1403
6.18.1. Upgrading from the SAP HANA Gen/App to the SAP HANA Recovery Kit.....	1405
6.18.2. SAP HANA Supported Configurations.....	1408
6.18.3. SAP HANA Recovery Kit Hardware and Software Requirements.....	1409
6.18.4. SAP HANA Recovery Kit Overview.....	1410
6.18.4.1. SAP HANA GUI States.....	1412
6.18.4.2. SAP HANA Resource Hierarchy.....	1415
6.18.4.3. Multitenant Database Support.....	1416
6.18.5. Configuring SAP HANA with LifeKeeper.....	1417
6.18.5.1. Install the SAP HANA Software.....	1418
6.18.5.2. Configure SAP HANA System Replication.....	1419
6.18.5.3. Modify the SAP HANA Instance Profile.....	1420
6.18.5.4. Install the LifeKeeper Software.....	1421
6.18.6. SAP HANA Resource Configuration Tasks.....	1422
6.18.6.1. Creating an SAP HANA Resource Hierarchy.....	1423
6.18.6.2. Extending an SAP HANA Resource Hierarchy.....	1426
6.18.6.3. Unextending an SAP HANA Resource Hierarchy.....	1430
6.18.6.4. Deleting an SAP HANA Resource Hierarchy.....	1432
6.18.6.5. Testing your SAP HANA Resource Hierarchy.....	1434
6.18.7. SAP HANA Resource Hierarchy Administration.....	1449
6.18.7.1. Changing Replication and Operation Modes.....	1453
6.18.7.2. Resolving Split Brain Scenarios.....	1457
6.18.7.3. Takeover with Handshake.....	1460
6.18.7.4. Setting Local and Temporal Recovery Policies for SAP HANA Resources.....	1465
6.18.8. SAP HANA Troubleshooting.....	1469

6.19. SAP MaxDB Recovery Kit Administration Guide .....	1470
6.19.1. SAP MaxDB Recovery Kit Hardware and Software Requirements .....	1472
6.19.2. SAP MaxDB Recovery Kit Overview .....	1473
6.19.2.1. SAP MaxDB Resource Hierarchy .....	1474
6.19.3. SAP MaxDB Configuration Considerations .....	1475
6.19.3.1. Using Raw I/O with SAP MaxDB .....	1476
6.19.3.2. Using SAP MaxDB Mirrored File Systems with DataKeeper .....	1477
6.19.3.3. Using Internal Load Balancer .....	1478
6.19.3.4. SAP MaxDB Active/Standby Considerations .....	1479
6.19.3.4.1. Active/Standby Configuration Example .....	1481
6.19.3.5. SAP MaxDB Active/Active Considerations .....	1482
6.19.3.5.1. Active/Active Configuration Example .....	1483
6.19.4. Configuring SAP MaxDB with LifeKeeper .....	1484
6.19.4.1. Install the SAP MaxDB Software .....	1485
6.19.4.2. Create the SAP MaxDB Database .....	1486
6.19.4.3. Create the User_Key .....	1487
6.19.4.4. Install the LifeKeeper Software .....	1489
6.19.5. SAP MaxDB Resource Configuration Tasks .....	1490
6.19.5.1. Creating an SAP MaxDB Resource Hierarchy .....	1492
6.19.5.2. Extending an SAP MaxDB Resource Hierarchy .....	1494
6.19.5.3. Unextending an SAP MaxDB Resource Hierarchy .....	1496
6.19.5.4. Deleting an SAP MaxDB Resource Hierarchy .....	1497
6.19.5.5. Testing Your SAP MaxDB Resource Hierarchy .....	1498
6.19.6. SAP MaxDB Resource Hierarchy Administration .....	1499
6.19.6.1. Modifying User_Keys .....	1500
6.19.6.2. Modifying OS User .....	1501
6.19.6.3. Updating SAP MaxDB Parameters .....	1502
6.19.7. SAP MaxDB Troubleshooting .....	1503
6.19.7.1. SAP MaxDB Recovery Kit Error Messages .....	1505
6.19.8. Appendix – Creating Device Spaces Using Raw I/O with SAP MaxDB .....	1507
6.19.8.1. Raw I/O Setup Steps .....	1508
6.19.8.2. Adding a Device Space after Creating a Hierarchy .....	1509
6.20. Sybase ASE Recovery Kit Administration Guide .....	1510
6.20.1. Sybase ASE Recovery Kit Overview .....	1511
6.20.2. Sybase ASE Recovery Kit Hardware and Software Requirements .....	1512
6.20.3. Sybase ASE Recovery Kit Configuration Considerations .....	1513
6.20.3.1. Using Raw I/O with Sybase .....	1514
6.20.3.2. Using Sybase ASE Mirrored File Systems with DataKeeper .....	1515
6.20.3.3. Sybase Interfaces File Considerations .....	1516
6.20.3.4. Sybase ASE Software Asset Manager (SySAM) .....	1518
6.20.3.5. Sybase ASE Active/Standby Considerations .....	1519
6.20.3.6. Sybase ASE Active/Active Considerations .....	1521
6.20.3.7. Sybase ASE Monitor Server and Backup Server .....	1523
6.20.3.8. Using Network Attached Storage with Sybase ASE .....	1524
6.20.4. Installing and Configuring Sybase ASE with LifeKeeper .....	1527

6.20.4.1. Install the Sybase ASE Software.....	1529
6.20.4.2. Create the Sybase ASE Servers .....	1530
6.20.4.3. Install the LifeKeeper Software with Sybase .....	1531
6.20.4.4. Creating a Sybase ASE Resource Hierarchy.....	1532
6.20.4.5. Extending a Sybase ASE Resource Hierarchy .....	1534
6.20.4.6. Unextending a Sybase ASE Resource Hierarchy .....	1536
6.20.4.7. Deleting a Sybase ASE Resource Hierarchy .....	1537
6.20.4.8. Testing Your Sybase ASE Resource Hierarchy.....	1538
6.20.5. Sybase ASE Recovery Kit Administration.....	1539
6.20.5.1. Modifying Protection for the Sybase Backup Server.....	1540
6.20.5.2. Modifying Protection for the Sybase Monitor Server.....	1542
6.20.5.3. Updating Sybase ASE Parameters.....	1544
6.20.6. Troubleshooting Sybase ASE Error During Resource Creation.....	1545
6.20.7. Appendix – Creating Device Spaces Using Raw I/O with Sybase ASE .....	1549
6.20.7.1. Requirements for Using Sybase ASE with Raw I/O .....	1550
6.20.7.2. Naming Conventions.....	1551
6.20.7.3. Using Raw I/O with Sybase Setup Steps.....	1552
6.20.7.4. Adding a Device Space after Creating a Sybase Hierarchy .....	1553
6.20.7.5. Creating Links for ASE and OCS .....	1554
6.21. VMDK Shared Storage Recovery Kit Administration Guide .....	1557
6.21.1. VMDK Documentation and References .....	1558
6.21.2. VMDK Hardware and Software Requirements .....	1559
6.21.3. VMDK Recovery Kit Overview .....	1560
6.21.4. Configuring the VMDK Recovery Kit.....	1561
6.21.4.1. VMDK Configuration Considerations .....	1562
6.21.4.2. VMDK Configuration Examples .....	1563
6.21.5. LifeKeeper VMDK Recovery Kit Configuration Tasks .....	1564
6.21.5.1. Register ESXi Host .....	1565
6.21.5.2. Changing the Virtual Machine Option Settings .....	1566
6.21.5.3. Creating a VMDK Resource Hierarchy .....	1568
6.21.5.4. Deleting a VMDK Resource Hierarchy .....	1571
6.21.5.5. Extending Your VMDK Hierarchy .....	1573
6.21.5.6. Unextending Your VMDK Hierarchy .....	1576
6.21.5.7. Testing Your VMDK Resource Hierarchy .....	1577
6.21.5.8. VMDK Maintenance .....	1578
6.21.6. VMDK Troubleshooting .....	1580
6.21.6.1. VMDK Error Messages .....	1581
<b>7. Parameters List .....</b>	<b>1586</b>
7.1. EC2 Parameters List .....	1591
7.2. IP Parameters List.....	1593
7.3. MQ Parameters List .....	1596
7.4. NFS Parameters List .....	1601
7.5. Recovery Kit for Oracle Cloud Infrastructure Parameters List.....	1602
7.6. Oracle Parameters List.....	1603
7.7. PostgreSQL Parameters List .....	1604

7.8. Quorum Parameters List .....	1606
7.9. Route53 Parameters List.....	1609
7.10. SAP Parameters List .....	1611
7.11. DataKeeper Parameters List .....	1613
7.12. Standby Node Health Check Parameters List.....	1617
7.13. SAP HANA Parameters List .....	1620
7.14. SAP MaxDB Parameters List.....	1623
<b>8. Search for an Error Code .....</b>	<b>1624</b>
8.1. Combined Message Catalog .....	1625
8.1.1. DataKeeper Kit Message Catalog .....	1898
8.1.2. DB2 Kit Message Catalog .....	1918
8.1.3. DMMP Kit Message Catalog .....	1932
8.1.4. Recovery Kit for EC2 Message Catalog .....	1947
8.1.5. File System Kit Message Catalog.....	1958
8.1.6. Gen/App Kit Message Catalog .....	1979
8.1.7. IP Kit Message Catalog.....	1990
8.1.8. Recovery Kit for Oracle Cloud Infrastructure Message Catalog .....	1996
8.1.9. Oracle Kit Message Catalog.....	2002
8.1.10. Oracle Listener Kit Message Catalog .....	2033
8.1.11. Oracle PDB Kit Message Catalog.....	2050
8.1.12. SCSI Kit Message Catalog .....	2055
8.1.13. Quick Service Protection Kit Message Catalog.....	2057
8.1.14. GUI Message Catalog .....	2063
8.1.15. SAP Message Catalog .....	2065
8.1.16. SAP HANA Recovery Kit Message Catalog.....	2088
<b>9. LifeKeeper for Linux Support Matrix .....</b>	<b>2112</b>
<b>10. Supported Storage .....</b>	<b>2127</b>
<b>11. Evaluation Guides .....</b>	<b>2148</b>
11.1. DataKeeper for Linux Evaluation Guide.....	2149
11.1.1. DataKeeper for Linux Terms to Know.....	2150
11.1.2. The Evaluation Process .....	2153
11.1.3. Prepare to Install DK for Linux .....	2154
11.1.4. Configure Storage for DK for Linux .....	2158
11.1.5. Install LifeKeeper for Linux.....	2160
11.1.6. Configure the Cluster – DK for Linux.....	2163
11.1.7. Test Your DK for Linux Environment .....	2177
11.2. LifeKeeper Evaluation Guide for Cloud Environments .....	2181
11.2.1. Before Starting an Evaluation of LifeKeeper for Linux .....	2182
11.2.1.1. High Availability, RTO, and RPO.....	2183
11.2.1.2. LifeKeeper for Linux – Integrated Components .....	2184
11.2.1.3. Benefits of LifeKeeper for Linux .....	2185
11.2.1.4. How Workloads Should be Distributed when Migrating to a Cloud Environment .....	2186
11.2.1.5. Public Cloud Platforms and their Network Structure Differences .....	2188

11.2.1.6. How a Client Connects to the Active Node .....	2189
11.2.1.6.1. AWS Route Table Scenario .....	2191
11.2.1.6.2. AWS Elastic IP Scenario .....	2193
11.2.1.6.3. Azure Internal Load Balancer Scenario .....	2194
11.2.1.6.4. Google Cloud Internal Load Balancer Scenario .....	2196
11.2.1.7. How does Data Replication between Nodes Work?.....	2197
11.2.1.8. What is “Split Brain” and How to Avoid It.....	2198
11.2.2. Documentation Style Used in this Guide .....	2200
11.2.2.1. AWS Route53 Scenario .....	2202
11.2.3. Configuring Network Components and Creating Instances .....	2203
11.2.3.1. Network Structure Used in this Tutorial.....	2204
11.2.3.2. Computing Resources Used in this Tutorial .....	2206
11.2.3.3. Creating an Instance in AWS from Scratch .....	2209
11.2.3.3.1. Switching between AWS Services .....	2212
11.2.3.3.2. Deciding on an AWS Region .....	2213
11.2.3.3.3. Creating the VPC.....	2214
11.2.3.3.4. Creating a Subnet.....	2217
11.2.3.3.5. Creating an Internet Gateway and Assigning it to the VPC .....	2222
11.2.3.3.6. Creating the Route Table.....	2224
11.2.3.3.7. Creating a Security Group .....	2228
11.2.3.3.8. Creating the First EC2 Instance.....	2233
11.2.3.3.9. Creating the Second and Third Instances .....	2241
11.2.3.4. Creating an Instance in Azure from Scratch .....	2242
11.2.3.4.1. Switching between Azure Services .....	2245
11.2.3.4.2. Deciding on an Azure Region .....	2246
11.2.3.4.3. Creating the Resource Group .....	2247
11.2.3.4.4. Creating a Virtual Network.....	2249
11.2.3.4.5. Creating a Network Security Group.....	2252
11.2.3.4.6. Creating the First Azure Virtual Machine.....	2258
11.2.3.4.7. Creating the Second and Third Virtual Machines .....	2268
11.2.3.5. Creating an Instance in Google Cloud from Scratch.....	2270
11.2.3.5.1. Switching between Google Cloud Services .....	2273
11.2.3.5.2. Deciding on a Google Cloud Region .....	2274
11.2.3.5.3. Creating the Project.....	2275
11.2.3.5.4. Creating a VPC Network.....	2278
11.2.3.5.5. Creating a New SSH Key.....	2280
11.2.3.5.6. Creating the First Google Cloud VM .....	2282
11.2.3.5.7. Configuring the Firewall Rules .....	2291
11.2.3.5.8. Creating the Second and Third VM .....	2296
11.2.4. Configure Linux Nodes to Run LifeKeeper for Linux .....	2297
11.2.4.1. Connecting to a Linux Node from Windows Client Using ssh.....	2298
11.2.4.2. Set a Hostname for Each Instance.....	2302
11.2.4.3. Disable SELinux .....	2304
11.2.4.4. Disable the Firewall .....	2306
11.2.4.5. Set a Password for the Root User.....	2307

11.2.4.6. Install x11 .....	2308
11.2.5. Install LifeKeeper for Linux.....	2310
11.2.5.1. Install AWS CLI .....	2315
11.2.5.2. Assign Permission to Use EC2 Recovery Kit.....	2316
11.2.5.3. Disable PING Broadcasting .....	2317
11.2.5.4. AWS – Disable Source/Destination Checking.....	2318
11.2.6. Login and Basic Configuration Tasks .....	2319
11.2.6.1. Connecting to a Linux Node with “X11 Forwarding” .....	2321
11.2.6.2. Setup X Window Client Software on Microsoft Windows .....	2323
11.2.6.3. Connecting to the First Node (node-a) .....	2327
11.2.6.4. Connecting to Other Nodes (node-b and node-c).....	2329
11.2.6.5. Define Communication Paths.....	2331
11.2.7. Protecting Our Resources .....	2335
11.2.7.1. Creating an IP Resource.....	2338
11.2.7.2. Switching between Nodes in a Cloud Environment.....	2343
11.2.7.2.1. Creating an AWS EC2 Resource (RouteTable Scenario) .....	2344
11.2.7.2.2. Creating an AWS EC2 Resource (Elastic IP Scenario) .....	2350
11.2.7.2.3. Creating an AWS Route53 Resource.....	2354
11.2.7.2.4. Azure – Using an Internal Load Balancer.....	2365
11.2.7.2.5. Google Cloud – Using an Internal Load Balancer.....	2372
11.2.7.2.6. Responding to Load Balancer Health Checks .....	2383
11.2.7.3. Switch to Standby Node to Confirm Switchover is Working .....	2395
11.2.7.4. How to Create Data Replication of a File System.....	2398
11.2.7.4.1. How to Prepare Disks for Replication on AWS.....	2405
11.2.7.4.2. How to Prepare Disks for Replication on Azure.....	2408
11.2.7.4.3. How to Prepare Disks for Replication on Google Cloud .....	2410
11.2.7.5. How to Protect Other Resources (Databases or Applications).....	2412
11.2.7.5.1. Protecting an Oracle Resource (non-PDB).....	2413
11.2.7.5.1.1. Install Oracle .....	2418
11.2.7.5.1.2. Create an Oracle Database (non-PDB).....	2419
11.2.7.5.1.3. Stop the Oracle Instance .....	2421
11.2.7.5.1.4. Rename /datakeeper/oradata/ORCL .....	2423
11.2.7.5.1.5. Update Config File for Oracle Listener on Both Nodes.....	2425
11.2.7.5.1.6. Start Database and Listener on node-a .....	2427
11.2.7.5.1.7. Configure Oracle LISTENER Resource .....	2428
11.2.7.5.1.8. Configure the Oracle Resource .....	2435
11.2.7.5.1.9. Test Switchover of the Oracle Resource.....	2440
11.2.7.5.2. Protecting MSSQL Using Quick Service Protection.....	2444
11.2.7.5.2.1. Install MSSQL 2017 .....	2448
11.2.7.5.2.2. Relocate Master Database and Log Files to Replicated Storage.....	2449
11.2.7.5.2.3. Rename Folders Under /dataKeeper/MSSQL .....	2451
11.2.7.5.2.4. Configure MSSQL Resource .....	2453
11.2.7.5.2.5. Customize LocalRecovery Parameter on Both Nodes.....	2459
11.2.7.5.2.6. Update Dependency between Resources .....	2461
11.2.7.5.3. Protecting a PostgreSQL Resource .....	2464

11.2.7.5.3.1. Install Postgres 12 on Linux Nodes .....	2468
11.2.7.5.4. Protecting an NFS Resource .....	2470
11.2.7.5.5. Protecting SAP Resources .....	2475
11.2.7.5.5.1. Create ASCS and ERS Virtual IPs.....	2481
11.2.7.5.5.1.1. AWS – Create ASCS and ERS Virtual IPs.....	2482
11.2.7.5.5.1.2. Azure – Create ASCS and ERS Internal Load Balancer .....	2488
11.2.7.5.5.1.3. Google Cloud – Create ASCS and ERS Internal Load	
Balancers .....	2497
11.2.7.5.5.2. Create SAP File Systems .....	2506
11.2.7.5.5.2.1. AWS/Azure – Create SAP Shared and Replicated File Systems	2507
11.2.7.5.5.2.2. Google Cloud – Create SAP Shared and Replicated File	
Systems .....	2518
11.2.7.5.5.3. Install SAP Instances .....	2530
11.2.7.5.5.4. Create LifeKeeper SAP Resources.....	2541
11.2.7.5.5.4.1. AWS/Azure – Create LifeKeeper SAP Resources.....	2542
11.2.7.5.5.4.2. Google Cloud – Create LifeKeeper SAP Resources .....	2548
11.2.7.5.5.5. Test Switchover and Failover .....	2554
11.2.7.5.6. Protecting SAP HANA Resources.....	2565
11.2.7.5.6.1. Create SAP HANA Primary Database Virtual IP .....	2569
11.2.7.5.6.1.1. AWS – Create the SAP HANA Virtual IP .....	2570
11.2.7.5.6.1.2. Azure – Create the SAP HANA Primary Database Load	
Balancer.....	2573
11.2.7.5.6.1.3. Google Cloud – Create the SAP HANA Primary Database Load	
Balancer.....	2579
11.2.7.5.6.2. Attach Disks for SAP HANA File Systems .....	2586
11.2.7.5.6.3. Install SAP HANA and Configure System Replication .....	2588
11.2.7.5.6.4. Create LifeKeeper SAP HANA Resource.....	2594
11.2.7.5.6.4.1. AWS/Azure – Create LifeKeeper SAP HANA Resource.....	2595
11.2.7.5.6.4.2. Google Cloud – Create LifeKeeper SAP HANA Resource .....	2597
11.2.7.5.6.5. Test Switchover and Failover .....	2600
11.2.7.6. Common Tasks.....	2609
11.2.7.6.1. How to Confirm if the Data Storage is Available on a Node.....	2610
11.2.7.6.2. Switchover the Data Storage to the Other Node .....	2611
<b>12. Quick Start Guides .....</b>	<b>2613</b>
12.1. AWS Direct Connect Quick Start Guide .....	2614
12.1.1. AWS Direct Connect Requirements.....	2615
12.1.2. AWS Direct Connect Setup Procedure .....	2617
12.1.2.1. AWS Direct Connect Preparations .....	2618
12.1.2.2. Creating Direct Connect Resources .....	2619
12.1.2.3. Configuring a Route Table .....	2620
12.1.3. Considerations for Settings and Operations in AWS Direct Connect.....	2621
12.2. Microsoft Azure Quick Start Guide.....	2622
12.2.1. Microsoft Azure Overview .....	2623
12.2.2. Configurations.....	2624
12.2.3. LifeKeeper-Specific Configurations in Azure .....	2627

12.2.4. Building a Virtual Machine and Starting the OS .....	2629
12.2.4.1. Creating a Resource Group .....	2630
12.2.4.2. Creating a Virtual Network .....	2631
12.2.4.3. Creating a Virtual Machine .....	2633
12.2.4.3.1. Creating a Cluster Node (Active) .....	2634
12.2.4.3.2. Creating a Cluster Node (Standby) .....	2645
12.2.4.3.3. Creating a Client and Witness Server .....	2648
12.2.4.4. Creating a Load Balancer .....	2651
12.2.4.5. Configuring the OS .....	2658
12.2.5. Building an HA Cluster with LifeKeeper .....	2665
12.2.5.1. Creating a Communication Path .....	2666
12.2.5.2. Configuring Quorum/Witness .....	2667
12.2.5.3. Disable Broadcast Ping .....	2668
12.2.5.4. Creating IP Resources .....	2669
12.2.5.5. Creating a Data Replication Resource Hierarchy .....	2672
12.2.5.6. Creating an Oracle Resource Hierarchy .....	2678
12.2.5.6.1. Installing Oracle DB .....	2682
12.2.5.6.2. Configuring a Listener .....	2686
12.2.5.6.3. Creating the DB .....	2691
12.2.5.6.4. Configuring Oracle .....	2692
12.2.5.6.5. Creating an Oracle Database Listener Resource Hierarchy .....	2694
12.2.5.6.6. Creating an Oracle Resource Hierarchy .....	2699
12.2.5.6.7. Creating an Oracle Pluggable Database Resource Hierarchy .....	2704
12.2.5.6.7.1. Setting Resource Dependencies .....	2708
12.2.5.7. Connectivity Check .....	2710
12.2.6. Availability Zone (High Availability Zone) .....	2713
12.2.6.1. Azure Configuration .....	2715
12.2.6.2. Checking the Client Redirection .....	2719
12.2.7. References and Acknowledgements .....	2721
12.3. MySQL Cluster with Data Replication (“Shared Nothing” Cluster) .....	2722
12.3.1. Terms to Know .....	2723
12.3.2. The Evaluation Process – MySQL Cluster .....	2726
12.3.3. Prepare to Install .....	2727
12.3.4. Configure Storage .....	2731
12.3.5. Install, Configure and Start MySQL .....	2734
12.3.6. Install LifeKeeper for Linux – MySQL Cluster .....	2737
12.3.7. Configure the Cluster .....	2740
12.3.8. Test Your Environment .....	2758
12.4. LifeKeeper for Linux in the AWS Cloud (SAP) .....	2765
12.4.1. Additional Steps to Configure SAP on LifeKeeper .....	2766
12.4.2. ASCS without NFS .....	2771
12.4.2.1. General Setup Steps for ASCS without NFS .....	2772
12.4.2.2. Installing SAP .....	2773
12.4.2.3. Creating the SAP Resource Hierarchy .....	2774
12.4.3. ASCS + ERS with NFS .....	2779

12.4.3.1. General Setup Steps.....	2780
12.4.3.2. Installing SAP .....	2782
12.4.3.3. Setting up NFS .....	2783
12.4.3.4. Creating a Resource Hierarchy .....	2788
12.4.3.5. Creating the SAP Resource Hierarchy .....	2791
12.4.3.6. Create the ERS Resources .....	2796
12.4.3.7. Enforcing ASCS/ERS Avoidance Behavior When Using ENSA2/ERSv2.....	2798
12.4.4. Switchover and Failover Testing .....	2807
12.4.4.1. Additional Resources .....	2809
12.5. Connecting to a LifeKeeper Cluster using AWS Transit Gateway Quick Start Guide .....	2810
12.5.1. AWS VPC Peering Connections Requirements .....	2812
12.5.1.1. LifeKeeper Software Requirements for AWS Environment .....	2814
12.5.2. AWS VPC Peering Setup Procedure .....	2815
12.5.3. Configuring the Route Table .....	2816
12.5.4. Considerations for Settings and Operations in AWS VPC Peering.....	2817
12.5.4.1. Considering the Use of LifeKeeper I-O Fencing – AWS VPC Peering .....	2818
12.5.5. AWS Direct Connect Known Issues and Troubleshooting .....	2819
12.6. Connecting to a LifeKeeper Cluster using AWS VPC Peering Quick Start Guide .....	2820
12.6.1. Connecting to a LifeKeeper Cluster using AWS Requirements .....	2822
12.6.1.1. Peering Requirements for Connecting to a LifeKeeper Cluster using AWS.....	2824
12.6.1.2. Other AWS VPC Requirements.....	2825
12.6.2. Setup Procedure for Connecting to a LifeKeeper Cluster using AWS .....	2826
12.6.3. Related LifeKeeper Resources for AWS VPC Peering.....	2828
12.6.4. Connecting to a LifeKeeper Cluster using AWS Settings and Operations	
Considerations .....	2829
12.6.4.1. Considering the Use of LifeKeeper I-O Fencing when Connecting to a LifeKeeper	
Cluster using AWS .....	2830
12.7. PostgreSQL Cluster with Shared Storage (iSCSI) .....	2831
12.7.1. Terms to Know – PostgreSQL .....	2832
12.7.2. The Evaluation Process – PostgreSQL .....	2835
12.7.3. Prepare to Install – PostgreSQL.....	2836
12.7.4. Configure Storage – PostgreSQL .....	2840
12.7.5. Install, Configure and Start PostgreSQL .....	2842
12.7.6. Install LifeKeeper for Linux – PostgreSQL.....	2844
12.7.7. Configure the Cluster – PostgreSQL .....	2847
12.7.8. Test Your Environment – PostgreSQL.....	2862
12.8. Apache/MySQL Cluster Using Both Shared and Replicated Storage .....	2867
12.8.1. Terms to Know – Apache .....	2868
12.8.2. The Evaluation Process – Apache.....	2871
12.8.3. Prepare to Install – Apache .....	2872
12.8.4. Configure Storage – Apache .....	2876
12.8.5. Install and Configure Apache and PHP .....	2880
12.8.6. Install, Configure, and Start MySQL – Apache .....	2882
12.8.7. Install LifeKeeper for Linux – Apache .....	2885
12.8.8. Configure the Cluster – Apache .....	2888

12.8.9. Test Your Environment – Apache .....	2911
<b>13. LifeKeeper Single Server Protection .....</b>	<b>2922</b>
13.1. LifeKeeper Single Server Protection for Linux Release Notes .....	2923
13.2. LifeKeeper Single Server Protection for Linux Installation Guide .....	2940
13.2.1. LifeKeeper Single Server Protection for Linux Introduction .....	2942
13.2.2. Installing the LifeKeeper Single Server Protection Software .....	2944
13.2.3. How to Use Setup Scripts .....	2946
13.2.4. Upgrading LKSSP .....	2952
13.2.5. Obtaining and Installing the License for LKSSP .....	2953
13.2.6. Resource Policy Management .....	2956
13.2.7. Verifying LifeKeeper Single Server Protection Installation .....	2961
13.3. LifeKeeper Single Server Protection for Linux Technical Documentation .....	2962
13.3.1. Documentation and Training .....	2964
13.3.2. Intergration with VMware HA .....	2966
13.3.3. Administration .....	2967
13.3.3.1. Enabling VMware HA Integration with LifeKeeper Single Server Protection .....	2968
13.3.3.2. Enabled VMware HA Fault Detection and Recovery Scenario .....	2969
13.3.3.2.1. Maintaining a LifeKeeper Single Server Protection Protected System .....	2970
13.3.3.3. LifeKeeper Single Server Protection Heartbeat with VMware HA .....	2971
13.3.3.4. Quick Service Protection (QSP) Recovery Kit .....	2972
13.3.3.5. LifeKeeper API for Monitoring .....	2973
13.3.3.6. Watchdog .....	2974
13.3.3.7. LKCLI (LifeKeeper Command Line Interface) .....	2977
13.3.4. Troubleshooting .....	2979
13.3.4.1. Known Issues and Workarounds .....	2980
13.4. Application Recovery Kits .....	2984
<b>14. Product Support Schedule .....</b>	<b>2985</b>

# 1. LifeKeeper for Linux

---

## 2. LifeKeeper for Linux Release Notes

---

### Version 9.6.1

Released April 20, 2022

#### Important!!

**Read This Document Before Attempting To Install Or Use This Product!**

**This document contains last minute information that must be considered before, during and after installation.**

### Introduction

This release notes document is written for the person who installs, configures and/or administers the LifeKeeper for Linux product. The document contains important information not detailed in the formal LifeKeeper and DataKeeper documentation sets such as package versions and last-minute changes to instructions and procedures as well as a link to the Troubleshooting sections for product restrictions and troubleshooting hints and tips that were discovered through final product testing. It is important that you review this document before installing and configuring LifeKeeper software.

### SIOS Product Descriptions

#### LifeKeeper for Linux

The LifeKeeper product includes fault detection and recovery software that provides high availability for file systems, network addresses, applications and processes running on Linux. LifeKeeper supports the configuration and switchover of a given application across multiple servers. The servers on which the application is configured are assigned priorities to determine the sequence in which the application will move from server to server in the event of multiple failures.

LifeKeeper for Linux provides switchover protection for a range of system resources. Automatic recovery is supported for the following resource types:

- Processes and Applications
- Shared Storage Devices (Including VMWare virtual hard disks)
- [Network Attached Storage Devices](#)
- [LVM Volume Groups and Logical Volumes](#)
- File Systems (ext3, ext4, vxfs, xfs and nfs) **Note:** btrfs is not currently supported by the LifeKeeper for Linux. For detailed information see [LifeKeeper Core – Known Issues / Restrictions](#).

- Communication Resources (TCP/IP)
- Database Applications
  - [Oracle](#)
  - [MySQL](#)
  - [DB2](#)
  - [SAP MaxDB](#)
  - [PostgreSQL](#)
  - [EnterpriseDB Postgres Plus Advanced Server/EDB Postgres Advanced Server](#)
  - [Sybase](#)
- [Web Server Resources](#)
- [Samba Resources](#)
- [DataKeeper for Linux](#)
- [SAP Application Environment Resources](#)
- [WebSphere MQ Resources](#)
- [Postfix Resources](#)

## DataKeeper for Linux

The SIOS DataKeeper product:

- Provides volume-based synchronous and asynchronous data replication.
- Integrates into the LifeKeeper Graphical User Interface for administration and monitoring.
- Automatically resynchronizes data between source and target servers at system recovery.
- Monitors the health of underlying system components and performs local recovery in the event of failure.
- Allows manual resource switchovers and failovers of mirrored volumes.
- Can be easily upgraded to provide high availability clustering and automatic failover and recovery using a license key to enable new functionality.

## LifeKeeper Components

### LifeKeeper Core

LifeKeeper for Linux is bundled for, and only runs on, 64-bit systems (AMD64 and EM64T systems).

### The LifeKeeper Core bundle includes:

- LifeKeeper
- DataKeeper
- [Application Recovery Kits](#)

The LifeKeeper Core Package Cluster includes the following installable packages:

Package	Package Name	Description
<a href="#">LifeKeeper Core</a>	steeleye- lk-9.6.1-7412.x86_64.rpm	The LifeKeeper package provides recovery software for failures associated with core system components such as memory, CPUs, the operating system, the SCSI disk subsystem and file systems.
<a href="#">DataKeeper Core</a>	steeleye- lkDR-9.6.1-7412.x86_64.rpm	The DataKeeper package provides data replication (synchronous or asynchronous mirrors with intent logging).
<a href="#">LifeKeeper GUI</a>	steeleye- lkGUI-9.6.1-7412.x86_64.rpm	The LifeKeeper GUI package provides a graphical user interface for LifeKeeper and DataKeeper administration and status monitoring.
<a href="#">LifeKeeper IP Recovery Kit</a>	steeleye- lkIP-9.6.1-7412.noarch.rpm	The LifeKeeper IP Recovery Kit provides recovery software for automatic switchover of IP addresses.
<a href="#">LifeKeeper Raw I/O Recovery Kit</a>	steeleye- lkRAW-9.6.1-7412.noarch.rpm	The LifeKeeper Raw I/O Recovery Kit provides support for applications that use raw i/o to bypass kernel buffering.
<a href="#">Quick Service Protection</a>	steeleye- lkQSP-9.6.1-7412.noarch.rpm	The Quick Service Protection Recovery Kit provides a simple disaster recovery function for various services.
LifeKeeper Man Pages	steeleye- lkMAN-9.6.1-7412.noarch.rpm	The LifeKeeper Man Page package provides reference manual pages for the LifeKeeper product.

### LifeKeeper Optional Recovery Software

The following optional software provides resource definition and recovery software for the application versions listed. See the [Support Matrix](#) and [Recovery Kit Administration Guides](#) for the requirements for each recovery software.

Package	Package Name	Description
<a href="#">LifeKeeper Apache Web Server Recovery Kit</a>	steeleye- lkAPA-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux Apache Web Server Recovery Kit provides fault resilience for Apache Web Server software in a LifeKeeper environment.

<a href="#">LifeKeeper SAP Recovery Kit</a>	steeleye- lkSAP-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux SAP Recovery Kit provides a mechanism to recover SAP NetWeaver from a failed primary server onto a backup server in a LifeKeeper environment working in conjunction with other LifeKeeper Recovery Kits to provide comprehensive failover protection.
<a href="#">LifeKeeper SAP MaxDB Recovery Kit</a>	steeleye- lkSAPDB-9.6.1-7412.noarch.rpm	The SAP MaxDB Recovery Kit provides fault resilient protection for SAP MaxDB databases in a LifeKeeper for Linux environment.
<a href="#">LifeKeeper DB2 Recovery Kit</a>	steeleye- lkDB2-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux DB2 Recovery Kit provides fault resilient protection for DB2 database instances. LifeKeeper together with the DB2 Universal Database product family afford increased availability to DB2 operating environments by effectively recovering database server failures without significant down-time or human intervention.
<a href="#">LifeKeeper Oracle Recovery Kit</a>	steeleye- lkORA-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux Oracle Recovery Kit provides fault resilience for Oracle software in a LifeKeeper environment furnishing a mechanism to tie the data integrity of Oracle databases to the increased availability provided by LifeKeeper.
<a href="#">LifeKeeper MySQL Recovery Kit</a>	steeleye- lkSQL-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux MySQL Recovery Kit provides an easy way to add LifeKeeper fault-resilient protection for MySQL resources and databases enabling a failure on the primary database server to be recovered on a designated backup server without significant lost time or human intervention.
<a href="#">LifeKeeper PostgreSQL Recovery Kit</a>	steeleye- lkPGSQL-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux PostgreSQL Recovery Kit is an SQL compliant, object-relational database management system (ORDBMS) based on POSTGRES providing a mechanism for protecting PostgreSQL instances within LifeKeeper.
<a href="#">LifeKeeper Sybase ASE Recovery Kit</a>	steeleye- lkSYBASE-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux Sybase ASE Recovery Kit provides LifeKeeper resource protection for the Sybase ASE components Adaptive Server, Monitor Server, and Backup Server.
<a href="#">LifeKeeper Postfix Recovery Kit</a>	steeleye- lkPOSTFIX-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux Postfix Recovery Kit provides a mechanism to recover Postfix from a failed primary server to a backup server in a

		LifeKeeper environment.
<a href="#">LifeKeeper Samba Recovery Kit</a>	steeleye- lkSMB-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux Samba Recovery Kit provides fault resilient protection for Samba file and print shares on a Linux server existing in a heterogeneous network enabling a failure on the primary Samba server to be recovered on a designated backup server without significant lost time or human intervention.
<a href="#">LifeKeeper NFS Server Recovery Kit</a>	steeleye- lkNFS-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux NFS Server Recovery Kit provides fault resilience for Network File System (NFS) software in a LifeKeeper environment enabling a failure on the primary NFS server to be recovered on a designated backup server without significant lost time or human intervention.
<a href="#">LifeKeeper Network Attached Storage Recovery Kit</a>	steeleye- lkNAS-9.6.1-7412.noarch.rpm	<p>The LifeKeeper for Linux Network Attached Storage Recovery Kit provides fault resilience for Network File System (NFS) software in a LifeKeeper environment affording LifeKeeper users the opportunity to employ an exported NFS file system as the storage basis for LifeKeeper hierarchies.</p> <p>NFS over UDP is not supported on Red Hat Enterprise Linux 8 and later.</p> <p>Some environments may require additional configurations. Refer to <a href="#">Specific Configuration Considerations</a> for details.</p>
<a href="#">LifeKeeper Logical Volume Manager (LVM) Recovery Kit</a>	steeleye- lkLVM-9.6.1-7412.noarch.rpm	The LifeKeeper for Linux Logical Volume Manager (LVM) Recovery Kit provides logical volume support for other LifeKeeper recovery kits allowing LifeKeeper-protected applications to take advantage of the benefits offered by the Logical Volume Manager, including simplified storage management and the ability to dynamically re-size volumes as needs change.
LifeKeeper PowerPath Recovery Kit	steeleye- lkPPATH-9.6.1-7412.noarch.rpm	The LifeKeeper PowerPath Recovery Kit protects applications that use EMC PowerPath multipath I/O devices.
LifeKeeper Device Mapper Multipath	steeleye- lkDMMP-9.6.1-7412.noarch.rpm	The LifeKeeper Device Mapper Multipath (DMMP Recovery Kit) protects applications and file

<p>(DMMP) Recovery Kit</p>		<p>systems that use DMMP devices allowing LifeKeeper to operate with and protect these applications and file systems.</p>
<p>Hitachi Dynamic Link Manager Software (HDLM) Recovery Kit</p>	<p>steeleye- lkHDLM-9.6.1-7412.noarch.rpm</p>	<p>The Hitachi Dynamic Link Manager Software (HDLM) Recovery Kit protects applications that use Hitachi Dynamic Link Manager Software devices.</p>
<p>LifeKeeper NEC iStorage StoragePathSavior (NECSPS) Recovery Kit</p>	<p>steeleye- lkSPS-9.6.1-7412.noarch.rpm</p>	<p>The SPS NEC iStorage StoragePathSavior (NECSPS) Recovery Kit protects applications that use NEC iStorage StoragePathSavior v3.3 or later multipath I/O devices.</p>
<p><a href="#">SIOS DataKeeper</a></p>	<p>steeleye- lkDR-9.6.1-7412.x86_64.rpm</p>	<p>SIOS DataKeeper for Linux provides an integrated data mirroring capability for LifeKeeper environments enabling LifeKeeper resources to operate in shared and non-shared storage environments.</p>
<p><a href="#">LifeKeeper WebSphere MQ Recovery Kit</a></p>	<p>steeleye- lkMQS-9.6.1-7412.noarch.rpm</p>	<p>The LifeKeeper for Linux WebSphere MQ Recovery Kit provides fault resilient protection for WebSphere MQ queue managers and queue manager storage locations enabling a failure on a primary WebSphere MQ server or queue manager to be recovered on the primary server or a designated backup server without significant lost time or human intervention.</p>
<p><a href="#">Quorum/Witness Package</a></p>	<p>steeleye- lkQWK-9.6.1-7412.noarch.rpm</p>	<p>The LifeKeeper Quorum/Witness Package allows a node to get a “second opinion” on the status of a failing node acting as an intermediary to determine which servers are part of the cluster. When determining when to fail over, the Witness Server, or Storage Witness, allows resources to be brought in service on a backup server only in cases where it verifies the primary server has failed and is no longer part of the cluster.</p> <p><b>Note:</b> Quorum is required for all 2 node DataKeeper clusters. Each 2 node cluster must use its own quorum/witness node. Shared witness servers are not recommended.</p>
<p><a href="#">Quick Service Protection</a></p>	<p>steeleye- lkQSP-9.6.1-7412.noarch.rpm</p>	<p>LifeKeeper Quick Service Protection supplies functionality to easily protect OS services.</p>

<a href="#">Recovery Kit for EC2™</a>	steeleye- lkECC-9.6.1-7412.noarch.rpm	The Recovery Kit for EC2 provides a mechanism to recover an Elastic IP from a failed primary server to a backup server. It also provides a mechanism to enable the IP Recovery Kit to work in multiple availability zones.
<a href="#">Recovery Kit for Route 53™</a>	steeleye- lkROUTE53-9.6.1-7412.noarch.rpm	Route53 Recovery Kit provides a mechanism for updating Amazon Route 53 DNS information corresponding to a virtual IP address and an actual IP address information of IP resources that are in dependency relation when switching to a failed primary server to a backup server.
<a href="#">VMDK as Shared Storage Recovery Kit</a>	steeleye- lkVMDK-9.6.1-7412.noarch.rpm	With the VMDK as Shared Storage Recovery Kit, VMware virtual hard disks and their file systems used as shared disks can be protected as LifeKeeper resources.
<a href="#">Recovery Kit for Oracle Cloud Infrastructure</a>	steeleye- lkOCIVIP-9.6.1-7412.noarch.rpm	Communication via virtual IP address using the LifeKeeper IP Recovery Kit is allowed in the Oracle Cloud environment.

## New Features of LifeKeeper for Linux Version 9

Product	Feature
<b>New in Version 9.6.1</b>	
LifeKeeper Core	Supports Rocky Linux 8.4 <ul style="list-style-type: none"> <li>Rocky Linux is not supported in the Cloud. (i.e. AWS/Azure/GCP/OCI) .</li> <li>Rocky Linux v8.4 is only supported in LifeKeeper for Linux v9.6.1.</li> <li>LifeKeeper Single Server Protection is not supported.</li> <li>SAP Recovery Kit and SAP HANA Recovery Kit are not supported.</li> </ul>
	Oracle Linux 8.5 is supported. <b>Note:</b> If you are using DataKeeper with RHCK, follow <a href="#">these steps</a> when installing LifeKeeper.
	Red Hat Enterprise Linux 8.5 is supported.
	<a href="#">Bug Fixes</a>
SAP Recovery Kit	The SAP Recovery Kit now supports the use of LifeKeeper-specific critical_nfs_mounts_<tag> files. Any NFS mount entry (e.g., for the /sapmnt/<SID> NFS share) added to this file for a given SAP resource on a given server will be mounted before the resource is brought in-service.
	<a href="#">Bug Fixes</a>
Recovery Kit for Oracle Cloud	Communication via virtual IP address using the LifeKeeper IP Recovery Kit is allowed in the Oracle Cloud environment.

Infrastructure	
NAS Recovery Kit	It is no longer necessary to set “NFS_RPC_PROTOCOL=tcp” when protecting an NFS shared file system with UDP disabled
	<a href="#">Bug Fixes</a>
GUI, CLI, DataKeeper, Route53, Quorum/ Witness, NAS, EC2	<a href="#">Bug Fixes</a>
PostgreSQL Recovery Kit	PowerGres on Linux 13 is now supported. (Certified in May 2022)
<b>New in Version 9.6.0</b>	
LifeKeeper Core	Support SLES 15 SP3 (Supported kernel versions are 5.3.18-59.5 and later)
	Supports Oracle Linux 8.4
	The STONITH feature is now available on Microsoft Azure
Oracle Recovery Kit	Supports Oracle 21c (21.3) running on-premises
Recovery Kit for EC2™	Added the ability to check source/destination checks in the EC2 ARK route scenario and disable it if they are enabled. New permissions (ec2:DescribeNetworkInterfaceAttribute, ec2:ModifyNetworkInterfaceAttribute IAM) are now required.
PostgreSQL Recovery Kit	Supports PostgreSQL 14
	Supports FUJITSU Software Enterprise Postgres 13 (Advanced, Standard, Community)
	Supports EDB Postgres Advanced Server 14.0 (Certified in January 2022)
SAP Recovery Kit	The SAP Recovery Kit is now supported on RHEL 8.4
	The SAP Recovery Kit is now supported on SLES 15 SP3
	The SAP Recovery Kit is now supported on S/4HANA 2021 Platform (SAP kernel 7.85) (Certified in March 2022)
SAP HANA Recovery Kit	Supports RHEL 8.4/SLES15.3
	Supports SAP HANA2 SP6 (Certified in February 2022)
	<b>See video on:</b> <a href="#">Local Recovery Enhancements Added To The HANA Recovery Kit</a>
Bug Fixes	
<b>New in Version 9.5.2</b>	
LifeKeeper Core	Supports <b>Red Hat Enterprise Linux 8.4</b>

	Supports <b>CentOS 8.3</b>
	Supports <b>Oracle Linux 8.3</b>
	Supports <b>Oracle Linux 7 UEK 6</b> <b>Note:</b> The kernel should be updated to 5.4.17-2102.202.5 or higher
	Supports <b>Oracle Linux 8 UEK 6</b> <b>Note:</b> The kernel should be updated to 5.4.17-2102.202.5 or higher
	LKCLI has been enhanced to control the following Recovery Kits. <ul style="list-style-type: none"> <li>• SAP RK</li> <li>• SAP MaxDB RK</li> <li>• MQ RK</li> <li>• DMMP RK</li> <li>• HDLM RK</li> <li>• NEC SPS RK</li> </ul>
	Implemented <a href="#">Standby Node Health Check</a> for the following Recovery Kits: <ul style="list-style-type: none"> <li>• Filesystem RK</li> <li>• NAS RK</li> </ul>
	You can now run lkcli from non-root users who belong to the lk group
	You can now use the “-i” option with lkcli stop as well as with lkstop
	Bug Fixes
SAP	SAP resources can now be created, extended, configured, and cloned with <a href="#">LKCLI</a>
	<b>Video demo</b> of SAP CLI: <a href="#">here</a>
SAP MaxDB	SAP MaxDB resources can now be created, extended, and cloned with <a href="#">LKCLI</a>
SAP HANA	Supports SAP HANA support for RHEL 8.2
	Supports SAP HANA support for SLES 15.2
	Supports <a href="#">Takeover with Handshake</a>
	HANA ARK added adjustable parameters to /etc/default/LifeKeeper with postinstall
	Bug Fixes
MySQL	Supports MariaDB 10.5
NAS	Implemented standby node health check for <b>NAS resources</b>
MQ	Supports WebSphere MQ 9.2 for RHEL 7.9
	Supports WebSphere MQ 9.2 for RHEL 8.3
	MQ resources can now be created, extended, configured, and cloned with <a href="#">LKCLI</a>
	<b>Video demo</b> of MQ CLI: <a href="#">here</a>
DataKeeper, Install, NFS, IP,	Bug Fixes

<p>Oracle, SAP, Quorum/Witness, Filesystem, Generic, Iksupport, DMMP</p>	
<p><b>New in Version 9.5.1</b></p>	
<p>LifeKeeper Core</p>	<p>Supports Red Hat Enterprise Linux 7.9 (Certified in December 2020)  <b>Note:</b> If you are using DataKeeper, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.</p>
	<p>Supports CentOS 7.9 (Certified in December 2020)  <b>Note:</b> If you are using DataKeeper, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.</p>
	<p>Supports Oracle Linux 7.9 (Certified in December 2020)  <b>Note:</b> If you are using DataKeeper, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.</p>
	<p><a href="#">lkstop -i</a> command stops LifeKeeper core but does not stop the protected resources. The user is prompted to confirm (yes/no) that they want to continue.</p>
	<p>LKCLI has been enhanced to control the following Recovery Kits.</p> <ul style="list-style-type: none"> <li>• DB2 RDBMS</li> <li>• RAW</li> <li>• Postfix</li> <li>• VMDK as Shared Storage</li> <li>• SAP ASE</li> <li>• Samba</li> <li>• HULFT / HULFT HUB</li> <li>• OraclePDB for Oracle RDBMS</li> </ul>
	<p>Supports Red Hat Enterprise Linux 8.3 (Certified in January 2021)  <b>Note:</b> If you are using DataKeeper, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.</p>
	<p>Supports Red Hat Enterprise Linux 8.2</p>
	<p>Supports CentOS 8.2</p>
	<p>Supports Oracle Linux 8.2 (excluding UEK6)</p>
	<p>Supports OpenSSL package to 1.1.1g</p>
<p>Supports SLES15 SP2</p>	
<p>Supports cURL package to 7.68.0</p>	
<p>Bug Fixes</p>	
<p>Generic Application Kit for Load Balancer</p>	<p>LifeKeeper for Linux now supports the <a href="#">Generic Application Kit for Load Balancer Health Checks</a> in Azure (January 2021).</p>

Health Checks	
MQ	LifeKeeper for Linux now supports IBM MQ 9.2 (January 2021)
QSP	Apache Tomcat can be protected with the Quick Service Protection (QSP) Recovery Kit
install	Implemented a <a href="#">setup script improvement</a>
	Bug Fixes
PostgreSQL	Supports PostgreSQL 13 (Certified in December 2020)
	Supports EDB Postgres Advanced Server 13.0 (Certified in December 2020)
	Supports FUJITSU Software Enterprise Postgres 12 (Certified in December 2020)
SAP HANA	Supports SAP HANA2 SP5 (Certified in December 2020)
	Supports SAP HANA for RHEL 8.1(Certified in December 2020)
	Supports SAP HANA for SLES 12.5 (Certified in December 2020)
	Bug Fixes
SAP	Supports SAP S/4HANA 2020 (Certified in December 2020)
	Bug Fixes
MaxDB, DB2, Filesystem, Generic Application, IP	Bug Fixes
<b>New in Version 9.5.0</b>	
SAP HANA Recovery Kit	<p>The SAP HANA Recovery Kit, providing high-availability for SAP HANA 2.0 SPS04 clusters, is now available. See the <a href="#">SAP HANA Recovery Kit Administration Guide</a> for details.</p> <ul style="list-style-type: none"> <li>• If you are using LifeKeeper for Linux v9.5 or later you must use the new (built-in) SAP HANA Application Recovery Kit.</li> <li>• The existing SAP HANA gen/app based Recovery Kit is <b>not</b> supported with v9.5.0. Users who wish to upgrade to the LifeKeeper for Linux v9.5.0 <b>must</b> convert their existing SAP HANA gen/app based Recovery Kit to the new SAP HANA Recovery Kit. Refer to <a href="#">Upgrading from the SAP HANA Gen/App to the SAP HANA Recovery Kit</a> for details.</li> <li>• SIOS will continue to support the SAP HANA gen/app based Recovery Kit with the 9.4.x releases until <b>March 31, 2022</b>.</li> <li>• The SAP HANA Recovery Kit does not support HANA v1.</li> </ul>
LifeKeeper Core	<p>Supports Red Hat Enterprise Linux 7.8 (Certified in July 2020)</p> <p><b>Note:</b> If you are using DataKeeper, follow <a href="#">these steps</a> when installing LifeKeeper.</p>

	Supports CentOS 7.8 (Certified in July 2020) <b>Note:</b> If you are using DataKeeper, follow <a href="#">these steps</a> when installing LifeKeeper.
	Supports Oracle Linux 7.8 (Certified in July 2020) <b>Note:</b> If you are using DataKeeper, follow <a href="#">these steps</a> when installing LifeKeeper.
	Supports SUSE Linux Enterprise Server 12 SP5 (Certified in July 2020)
	Support VMWare vSphere 7.0 (Certified in July 2020)
	Supports CentOS 8.0
	Supports Oracle Linux 8.0
	Supports Red Hat Enterprise Linux 8.1
	Supports CentOS 8.1
	Supports Oracle Linux 8.1
	The CLI has been enhanced to allow you to control LifeKeeper through the Command Line Interface. See <a href="#">LKCLI</a> for details.
	Bug Fixes
PostgreSQL	Support PostgreSQL 12
	EDB Postgres Advanced Server v12.0 is supported. (Certified in July 2020)
Oracle	PDBs with Multitenant configurations can now be protected. See <a href="#">Configuring a Pluggable Database with Oracle Multitenant</a> for details.
DataKeeper	DataKeeper online mirrored volumes can now be resized. See <a href="#">Mirror Resize</a> for more information.
	Mirror recovery of data replication resources can now be performed in parallel.
	Added LKDR_CONNECT_NBD_DURING_RESTORE parameter. Refer to <a href="#">DataKeeper parameter list</a> for details.
	Bug Fixes
Filesystem, LVM, NFS, IP, DB2, MaxDB, SAP, Sybase, Sybase ASE, Quorum/ Witness	Bug Fixes
<b>New in Version 9.4.1</b>	

LifeKeeper Core	OpenJDK included with OS is installed. See <a href="#">Configuring the LifeKeeper GUI</a> for details.  (Updated December 2020) Some environments may install OpenJDK that is included with the LifeKeeper installation image. Please refer to <a href="#">Configuring the LifeKeeper GUI</a> for more information.
	Supports SUSE Linux Enterprise Server 15 SP1
	Supports Oracle Linux 7.7
	Supports CentOS 7.7
	Supports AWS Nitro system
	Supports AWS Transit Gateway
	Bug Fixes
DataKeeper	Supports NVMe devices
	Bug Fixes
VMDK as Shared Storage	LifeKeeper for Linux VMDK as Shared Storage Recovery Kit is now available. See the <a href="#">VMDK Shared Storage Recovery Kit Management Guide</a> for details.
PostgreSQL	PowerGres Plus (for Linux) v10 and PowerGres on Linux v11 can be protected with the PostgreSQL Recovery Kit
	Support FUJITSU Software Symfoware Server (Postgres) V12.4 (Certified in March 2020) For the details, refer to the <a href="#">LifeKeeper Optional Recovery Software Requirements, PostgreSQL Recovery Kit Administration Guide &gt; Administration</a> .
Install, IP, MaxDB, EC2	Bug Fixes
<b>New in Version 9.4.0</b>	
LifeKeeper Core	Standby Node Health Check – allows the user to monitor CPU and memory utilization on the standby node and monitor the health of out-of-service (OSU) resources to detect errors on the standby node.
	Oracle Linux 7 Unbreakable Enterprise Kernel Release 5 (UEK R5) is supported.
	Red Hat Enterprise Linux 8 is supported.  <b>Note:</b> Upgrading from one kernel version to another major version such as from RHEL7 to RHEL8 is NOT supported. (i.e DataKeeper resource does NOT work when upgrading from RHEL7 to RHEL8.)
	Red Hat Enterprise Linux 7.7 is supported (Certified in November 2019)

	<p><b>Note:</b> If you are using DataKeeper, follow <a href="#">these steps</a> when installing LifeKeeper.</p>
SAP	<p><a href="#">SAP-certified support of SAP S/4HANA Platform via SAP High Availability Clustering Certification S/4-HA-CLU-1.0</a>  <a href="#">SAP S/4HANA 1809 Platform is now supported.</a>  <a href="#">SAP S/4HANA 1909 Platform is now supported.</a> (Added support in November 2019)</p>
	<p><a href="#">Support for Standalone Enqueue Server 2</a> and <a href="#">Enqueue Replication Server 2</a></p>
	<p><a href="#">SAP Resource UI Enhancements</a></p>
	<p><a href="#">Optimizations for the LifeKeeper SAP ERS Resource</a></p>
MySQL	<p>MariaDB10.3 is supported.</p>
DB2	<p>DB2 11.5 is supported.</p>
PostgreSQL	<p>FUJITSU Software Enterprise Postgres 11 is supported. (Certified in November 2019)</p>
General maintenance	<p>Bug Fixes</p>
<p><b>New in Version 9.3.2</b></p>	
LifeKeeper Core	<p>Red Hat Enterprise Linux 7.6 is supported.</p> <p><b>Note:</b> DataKeeper asynchronous mirrors are not supported because of a Linux kernel bug.</p>
	<p>CentOS 7.6 is supported.</p> <p><b>Note:</b> DataKeeper asynchronous mirrors are not supported because of a Linux kernel bug.</p>
	<p>Oracle Linux Version 7.6 is supported.</p> <p><b>Note:</b> DataKeeper asynchronous mirrors are not supported because of a Linux kernel bug.</p> <p><b>Note:</b> Unbreakable Enterprise Kernel Release 5 (UEK R5) is NOT supported. (i.e. DataKeeper resource does NOT work on UEK R5.)</p>
	<p>Linux Enterprise Server 12 SP4 is supported.</p>
	<p>SUSE Linux Enterprise Server 15 is supported.</p>

	<p><b>Note:</b> Upgrading from one kernel version to another major version such as from SLES 12 to SLES 15 is NOT supported. (i.e DataKeeper resource does NOT work when upgrading from SLES 12 to SLES 15.)</p>
	<p>OpenJDK v10.0.2 is supported.</p> <p>SIOS has tested OpenJDK v10.0.2 downloaded from <a href="https://jdk.java.net/10/">https://jdk.java.net/10/</a> with LifeKeeper for Linux ( v9.4.0. OpenJDK is compatible with LifeKeeper v9.4.0, therefore customers may use this version or any compatible version of OpenJDK with LifeKeeper v9.4.0. If a customer encounters an issue due to the OpenJDK version, SIOS may recommend using a newer version of OpenJDK or the OracleJDK included in LifeKeeper package.</p>
Install	The -s option for saving the current setup configuration has been added to the setup command.
	RHEL7.6 also does not support DataKeeper's asynchronous mode. The warning message is output by setup.
DataKeeper	<p>Wait For Previous Source for multi-target mirrors.</p> <p>For a multi-target mirror DataKeeper keeps track of the last server, aka previous source, that had the mirror in-service. When there is a failover, the bitmap from the previous source is required to keep all of the targets in-sync. DataKeeper will now automatically wait for the previous source to join the cluster before resuming replication to any target. This allows the bitmap from the previous source to be merged so that only partial resyncs are necessary.</p>
	Unnecessary synchronization is avoided in the environment with three or more nodes.
	Add updated messages for "wait for source" in GUI and mirror_status.
PostgreSQL	PostgreSQL 11 is supported.
	EDB Postgres Advanced Server v11 is supported.
	FUJITSU Software Enterprise Postgres 10 is supported. For the details, refer to the <a href="#">LifeKeeper Optional Recovery Software Requirements</a> .
MQ	LifeKeeper for Linux now supports IBM MQ 9.1
	Supports WebSphere MQ 9.2 for RHEL 7.9.
	Supports WebSphere MQ 9.2 for RHEL 8.3.
Oracle	Support Oracle 19c (Certified in August 2019).
SAP	Supports new maintenance mode feature available with SAP kernel 7.49 and above.

General Maintenance	Bug Fixes
<b>New in Version 9.3.1</b>	
LifeKeeper Core	Updated the OpenSSL package to 1.0.2p
	Support Red Hat Enterprise Linux 6.10
	Support CentOS 6.10
	Support Oracle Linux 6 Update 10
MySQL	Support MySQL 8.0
Oracle	Support Oracle 18c (Certified in March 2019)
Install, EC2, Route53	Bug Fix
<b>New in Version 9.3</b>	
LifeKeeper Core	Red Hat Enterprise Linux Version 7.5 is supported. However, DataKeeper asynchronous mirrors are not supported because of a Linux kernel bug.
	CentOS7.5 is supported. However, DataKeeper asynchronous mirrors are not supported because of a Linux kernel bug.
	Oracle Linux Version 7.5 is supported. However, DataKeeper asynchronous mirrors are not supported because of a Linux kernel bug.
	Support VMware vSphere 6.7 (Certified in October 2018)
	Bug Fixes
EC2, Route53	EC2 and Route53 RK now support HTTP Proxy.
	Bug Fixes
Quorum/Witness	Storage QWK is now supported. For details, please click <a href="#">here</a> .
	Bug Fixes
Install	The LifeKeeper for Linux installation process has been upgraded. For details, please click <a href="#">here</a> .
SAP, Oracle, Samba, MQ, Sybase, Filesystem, Generic Application, QSP, SAP MaxDB,	Bug Fix

DataKeeper	
<b>New in Version 9.2.2</b>	
EC2,Route53	<p>IAM Role is now supported.</p> <ul style="list-style-type: none"> <li>Openswan Recovery Kit does not support IAM Role. You may use v9.2.1 in case of Cross Region configuration.</li> </ul>
DataKeeper	<p>Support GUID Partition Table (GPT) to identify protected disks</p> <ul style="list-style-type: none"> <li>The supported disk is SCSI Hard Disk or Xen Virtual Disk(xvd) in case of Linux kernel 2.6.27 or earlier.</li> </ul>
PostgreSQL	Support PostgreSQL 10
	EDB Postgres Advanced Server v10.0 is now supported. (Certified in April 2018)
SAP, NAS, EC2	Bug Fix
<b>New in Version 9.2.1</b>	
LifeKeeper Core	Support Oracle Linux 7.4
	Support CentOS 7.4
	Support SUSE Linux Enterprise Server 12 SP3 <ul style="list-style-type: none"> <li>The kernel should be updated to 4.4.82-6.9.1 for SUSE Linux Enterprise Server 12 SP3</li> </ul>
	The Recovery Kit for EC2, Route 53 Recovery Kit, Openswan Recovery Kit can now be installed from the setup menu. Openswan Recovery Kit is supported only when using with Cross Region configuration
	Bug Fixes
PostgreSQL	Support EDB Postgres Advanced Server 9.6
MQ	Support IBM MQ 9.0
<b>New in Version 9.2</b>	
LifeKeeper Core	Support Red Hat Enterprise Linux 7.4
	SNMP trap can be sent to multiple targets
	Bug fixes
IP	IP resources using real IP(primary IP address configured for NIC) can be created
PostgreSQL	Support PostgreSQL 9.6
	Support FUJITSU Software Enterprise Postgres 9.6 For the details, refer to the <a href="#">LifeKeeper Optional Recovery Software Requirements, PostgreSQL Recovery Kit Administration Guide &gt; Administration</a> .
MQ	Support IBM MQ 9.0 (Certified in December 2017)
MD, SAP, SAP MaxDB, Quorum/	Bug fixes

Witness, Route53, Install	
<b>New in Version 9.1.2</b>	
LifeKeeper Core	SUSE Linux Enterprise Server 12 SP2 is supported.
	CentOS7.3 is supported.
	Red Hat Enterprise Linux Version 6.9 is supported.
	kernel of Oracle Linux Version 7.3 is supported.
	Bug fixes
PostgreSQL	Support PostgreSQL 9.6
	Support FUJITSU Software Enterprise Postgres 9.6 For the details, refer to the <a href="#">LifeKeeper Optional Recovery Software Requirements, PostgreSQL Recovery Kit Administration Guide &gt; Administration.</a>
Oracle	Oracle 12c R2 is supported.
DB2	DB2 11.1 is supported.
IP, QSP, MySQL, NFS	Bug fixes
<b>New in Version 9.1.1</b>	
LifeKeeper Core	SUSE Linux Enterprise Server 12 SP1 support. <ul style="list-style-type: none"> <li>• SLES12.0 is not supported.</li> <li>• Btrfs is not supported.</li> </ul>
	Red Hat Enterprise Linux Version 7.3 support.
	Oracle Linux Version 7.3 support. <ul style="list-style-type: none"> <li>• UEK is not supported.</li> </ul>
	vSphere 6.5 support.
	Bug Fixes
PostgreSQL	PostgreSQL 9.5 support EDB Postgres Advanced Server v9.5 support FUJITSU Software Symfoware Server (Open Interface) V12.2 support FUJITSU Software Symfoware Server (Postgres) V12.3 support FUJITSU Software Enterprise Postgres 9.5 support For the details, refer to the <a href="#">LifeKeeper Optional Recovery Software Requirements, PostgreSQL Recovery Kit Administration Guide &gt; Administration.</a>
Sybase ASE	Sybase ASE 16.0 support.
MySQL	MySQL 5.7 support on RHEL 7.x/CentOS 7.x/OEL 7.x. <ul style="list-style-type: none"> <li>• MySQL 5.7 on other OS is already supported.</li> </ul>
SAP	SAP 7.5 support.

<b>New in Version 9.1.0</b>	
LifeKeeper Core	Red Hat Enterprise Linux 6.8 support (Certified in September 2016). CentOS 6.8, Oracle Linux 6.8 support (Certified in September 2016). <ul style="list-style-type: none"> <li>MD Recovery Kit is not supported on these OS.</li> </ul>
	LifeKeeper API for Monitoring Added API to supply LifeKeeper status and log information.
	Quick Service Protection support Added functionality to easily protect OS services.
	Bug Fixes
<b>New in Version 9.0.2</b>	
LifeKeeper Core	Support of Red Hat Enterprise Linux Version 7.2. <ul style="list-style-type: none"> <li>SQL RK is not supported when running on RHEL 7.x/CentOS 7.x/OEL 7.x.</li> </ul>
	Updated OpenSSL package to version to 1.0.1q
	Bug Fixes
MQ	Added support for Multi-version WebSphere MQ. With this support queue managers for 7.1, 7.5, and 8.x can all be protected on the same cluster node.
	Removed the Recovery Kit restriction that only the mqm user could be used for running MQ commands. With this change any user in the mqm group can be used by the Recovery Kit to run MQ commands.
	Bug Fixes
IP, Filesystem, DMMP, DataKeeper, EC2, PostgreSQL, Power Path, SAP, SAP MaxDB, Oracle	Bug Fixes.
Licensing	Update to a newer version of FlexNet
<b>New in Version 9.0.1</b>	
LifeKeeper Core	Bug Fixes.
DataKeeper	Bug Fixes.
<b>New in Version 9.0</b>	
LifeKeeper Core	Red Hat Enterprise Linux Version 6 Update 7 support. (Certified in October 2015)
	Community ENTerprise Operating System (CentOS) Version 6 Update 7 support. (Certified in October 2015)
	Oracle Linux Version 6 Update 7 support. (Certified in October 2015)

	SUSE LINUX Enterprise Server 11 SP4 support. (Certified in October 2015)
	Chef support
	Added LifeKeeper for Linux <a href="#">Parameters List</a> , document detailing tunable values. Added the <a href="#">lkchkconf command</a> .
	vSphere 6 support
	reiserfs filesystem type is no longer supported.
	Arks supported with Red Hat Enterprise Linux Version 7.0/7.1, Community ENTerprise Operating System (CentOS) Version 7.0/7.1, and Oracle Linux Version 7.0/7.1 are the same as LifeKeeper for Linux v8.4.1. (Arks to be supported are: PostgreSQL, MySQL, Oracle, DB2, Apache, Postfix, DMMP, LVM, NFS, NAS, Samba, MD, EC2, Route53, Openswan)
	Bug Fixes.
DataKeeper	The DK rewind feature is no longer supported in version 9. Prior to upgrading to version 9, you will need to deactivate all rewind configuration settings, and perform any necessary archival of data.
	Bug Fixes.
GUI	JRE 8u51 support. (JRE 7 is no longer supported.)
	Chrome Browser is no longer supported.
	Bug Fixes.

## Bug Fixes

The following is a list of the latest bug fixes and enhancements.

Bug	Description
PL-3157	Fixed an issue where lkbackup could output incorrect error.
PL-3402	An error is now displayed when autofs is detected when creating a filesystem resource.
PL-4168	Fixed an issue where restore would succeed but the root resource of the hierarchy would be OSF when resource failover takes a significant amount of time.
PL-4782	Fixed an issue where “Automatic” switchback would automatically be changed to “Intelligent” during Oracle resource creation.
PL-5617	Fixed an issue where an abnormality of a NAS is not detected. Note that following NAS formats are not supported by this fix. -The export part contains ":" -The host part is IPv6
PL-5679	Fixed the refresh button on the main frame that was grayed out and could not be selected when logging into the GUI as a user belonging to lkoper or lkguest.

PL-8042	An error is now displayed and "lkcli dependency create" will now fail when: - the destination node of an extended resource is not found. - the parent resource is in service but the child is not in service.
PL-8132	An error is now displayed and "lkcli dependency create" will now fail when resources have different extensions.
PL-8934	Fixed an issue where the target node would become QUORUM_LOST if the COMM_DOWN and COMM_UP events occurred in succession.
PL-8936	Enhanced the configuration check of the qwk_storage_init command.
PL-9246	Fixed an issue where the insufficient permissions warning was not displayed when creating Route53 resources.
PL-9251	Fixed an issue where a help message about the Host Name isn't displayed when creating a Route53 resource.
PL-9252	Fixed an issue where duplicated warning messages are displayed when a Route53 resource is brought 'In Service' after the target A Record is removed.
PL-9253	Fixed an issue where duplicated warning messages are displayed when a Route53 resource is brought 'In Service' after removing the 'ListResourceRecordSets' privilege from the attached IAM policy.
PL-9460	The message for error 135802 has been changed to be more understandable.
PL-9786	The SAP Recovery Kit now performs an additional check during resource creation to verify that the virtual host name associated to the SAP instance is not a match to the fully qualified or short physical local host name of the server. The results of this check may be ignored by setting the parameter SAP_IGNORE_HOSTNAME_CHECK=1 in /etc/default/LifeKeeper.
PL-10574	Fixed an issue where exclusive control of qwk_storage was not working.
PL-10894	Fixed an issue where the 'remove' script could be called depending on the timing of when lkstop was executed on the target side and when lkstop -f was executed to stop the resource while it was running on the source side.
PL-11129	Fixed an issue where NO_PROXY configuration was not detected.
PL-11316	Fixed an issue where certain resource-specific files would not be created for an SAP resource when the tag name contained forward slashes.
PL-11317	Fixed an issue where the GUI mirror status for a DataKeeper resource was not updated from "Wait to Resync" when a bitmap merge failed.
PL-11542	Enhanced the configuration check of the qwk_storage_init command.
PL-11547	Added additional troubleshooting information when the aws command fails while attempting to bring an EC2 resource in-service.
PL-11771	Fixed a timing issue that could lead to the fuser command hanging.
PL-12204	Fixed a problem in which Sybase resources were displayed as ISP even though ASE was not running.

PL-12800	Fixed an issue that caused communication between the SAP HA Cluster Connector Library and LifeKeeper to fail on recent SAP kernel patch levels.
PL-13254	Fixed an issue in the LifeKeeper GUI where the properties panel for an SAP ERS resource created in LifeKeeper 9.3.2 or earlier did not display the instance status on both cluster nodes.

## Hotfixes and Add-on Support Packages

The patches are located [here](#).

Hotfix/Add-on Support Package	Bug Description	Patch Description
steeleye-lkHOTFIX-Gen-LB-PL-7172-9.5.1-7154.x86_64.rpm	Support for load balancers in Microsoft Azure and Google Cloud Platform	The Generic Application Recovery Kit for Load Balancer Health Checks provides a mechanism to receive and respond to a TCP health check probe for target instances of load balancers in Microsoft Azure (Azure) and Google Cloud Platform (GCP) environments.
steeleye-lkHOTFIX-core-PL-7770-9.6.1-7412.noarch.rpm	This patch fixes a problem related to the lklicmgr command not properly recognizing licenses for the SAP HANA ARK and showing them as unknown.	This patch avoids erroneous unknown status messages for installed HANA ARK licenses.
steeleye-lkHOTFIX-NFS-PL-12327-9.6.1-7412.noarch.rpm	NFS resource hierarchy switchover hangs when exportfs hangs.	This patch will wait up to 30 seconds for the exportfs command to complete before aborting the action. A notification will be provided to the user that the process appears to have hung and may require a system reboot to clear the hang.
steeleye-lkHOTFIX-DR-PL-13245-9.6.1-7412.x86_64.rpm	Suppress message for mirrors without a unique identifier.	This patch provides an updated DRBase.pm file that limits the Emergency message to one quickCheck cycle for the resource per LifeKeeper start. A unique identifier is still required

		to protect the data from loss or corruption.
steeleye- lkHOTFIX-SAP-PL-13460-9.6.1-7412.x86_64.rpm	For ERS resources created prior to LifeKeeper 9.4.0, the resource properties panel in the LifeKeeper GUI does not show the process name ("ENREPSERVER") in the protected instance.	This patch fixes an issue where the process name is not displayed in the properties panel for an ERS resource created prior to LifeKeeper 9.4.0.

## Discontinued Features

Feature	Description
<b>Discontinued in Version 9.6.1</b>	
	None
<b>Discontinued in Version 9.6.0</b>	
Using the Client/Web Browser to access the LifeKeeper GUI is no longer supported.	
<b>Discontinued in Version 9.5.2</b>	
Core	Red Hat Enterprise Linux 6 is no longer supported.
	Oracle Linux 6 is no longer supported.
	CentOS 6 is no longer supported.
Software RAID	LifeKeeper Software RAID (md) Recovery Kit is no longer supported.
Oracle	Oracle virtual machine (OVM) latest 3.4.6 is no longer supported.
PostgreSQL	PostgreSQL 9.5 is no longer supported.
	EnterpriseDB Postgres Plus Advanced Server 9.5 is no longer supported.
VMware	VMware 5.5 and 6.0 are no longer supported.
<b>Discontinued in Version 9.5.1</b>	
LifeKeeper Core	The configuration using Chef is no longer supported.
<b>Discontinued in Version 9.5.0</b>	
LifeKeeper Core	System log management using syslog-ng is no longer supported. Please use rsyslog.

	SUSE Linux Enterprise Server (SLES) 11.0 to SP4 is no longer supported.
DataKeeper	Environments that use DEVNAME for disk identification using DataKeeper for Linux (DK resources) are no longer supported. Please use a GPT partition (GUID Partition Table).
Oracle	Oracle Database Enterprise Edition 11g R2 is no longer supported.
	Oracle Database Standard Edition 11g R2 is no longer supported.
	Oracle Database Standard Edition One 11g R2 is no longer supported.
MySQL	MariaDB 5.5, 10.0 is no longer supported.
PostgreSQL	PostgreSQL 9.4 is no longer supported.
	EnterpriseDB Postgres Plus Advanced Server 9.4 is no longer supported.
<b>Discontinued in Version 9.4.1</b>	
	None
<b>Discontinued in Version 9.4.0</b>	
DataKeeper	Multi-Site Cluster Feature

## System Requirements

### LifeKeeper Product Requirements

LifeKeeper for Linux is supported on any Linux platform that satisfies the minimum requirements included in the [Linux Configuration Table](#). Also refer to the [LifeKeeper Support Matrix](#) for supported operating systems, applications and virtualization.

 **Note:** LifeKeeper on a Linux server will not inter-operate with LifeKeeper for Windows.

Description	Requirement
Linux Operating System	See the <a href="#">Linux Configuration Table</a> for specific operating system information.
Virtual Environments	<p>The guest operating system running on the virtual machine must be one of the supported versions listed in the <a href="#">Linux Configuration Table</a>. The following virtual environment is an example where LifeKeeper for Linux is deployed. Please refer to the <a href="#">Support Matrix</a> for detailed versions of supported virtualization environments.</p> <ul style="list-style-type: none"> <li>• KVM</li> <li>• Oracle VM Server for x86</li> <li>• VMware vSphere v6.5, v6.7 and v7.0</li> </ul>

	<ul style="list-style-type: none"> <li>• Amazon EC2</li> <li>• Microsoft Azure</li> <li>• Nutanix Acropolis Hypervisor</li> <li>• Google Cloud™</li> <li>• Oracle Cloud Infrastructure (<b>See Note3</b>)</li> </ul> <p>SAN configuration is supported for vSphere 6.5 or later except RDM which is not supported by VMWare.</p> <p>Fibre channel SAN and shared SCSI cluster configurations are not supported with LifeKeeper for Linux running in a KVM and Oracle VM Server for x86 virtual machine.</p> <p><b>Note1:</b> Some Amazon EC2 configurations have issues when the Shutdown Strategy is set to “Do not Switchover Resources”. For detailed information, see Troubleshooting &gt; <a href="#">Known Issues and Restrictions</a>.</p> <p><b>Note2:</b> On SLES v12 or later running on AWS or Azure, the dynamic change of the virtual IP address by the cloud network plug-in may affect the operation of the LifeKeeper cluster. For detailed information, see <a href="#">LifeKeeper Core – Known Issues / Restrictions</a>.</p> <p><b>Note3:</b> Refer to <a href="#">Support Configuration</a> for the supported configuration and the restrictions.</p>
Memory	<p>The LifeKeeper for Linux minimum memory requirement is the same as the OS minimum requirement. System memory should be sized for the applications that will be running on the LifeKeeper protected system as well. Refer to <a href="#">Application Configuration</a> for further information.</p>
Disk Space	<p>The LifeKeeper Package Cluster requires the following disk space:</p> <p>/opt – approx 100MB (depending on kits installed)</p> <p>/ – approx 110MB</p>
Java Runtime Environment	<ul style="list-style-type: none"> <li>• OpenJDK 1.8, 10 or later</li> </ul>

## LifeKeeper Optional Recovery Software Requirements

The following table shows the software requirements for the optional LifeKeeper recovery software.

See [Application Configuration](#) for additional requirements and/or restrictions that may apply to applications under LifeKeeper protection.

Product	Requirement(s)
<a href="#">LifeKeeper Apache Web Server Recovery Kit</a>	Apache Web Server v2.4
<a href="#">SAP Recovery Kit</a>	SAP NetWeaver 7.0 including Enhancement Package 1,2 and 3 SAP NetWeaver 7.3 including Enhancement Package 1 SAP NetWeaver 7.4 SAP NetWeaver 7.5 SAP NetWeaver AS for ABAP 7.51 innovation package
<a href="#">LifeKeeper SAP MaxDB Recovery Kit</a>	SAP MaxDB v7.9  LifeKeeper Core Package Cluster
<a href="#">LifeKeeper Postfix Recovery Kit</a>	Postfix software provided with the supported Linux distributions installed and configured on each server. The same version of Postfix should be installed on each server.  LifeKeeper Core Package Cluster
<a href="#">LifeKeeper Oracle Recovery Kit</a>	Oracle Database Enterprise Edition v12c, v12c R2, v18c, v19c (excluding ASM) and v21c (excluding ASM)  Oracle Database Standard Edition 2 (SE2) v12c, v12c R2, v18c, v19c (excluding ASM) and v21c (excluding ASM)
<a href="#">LifeKeeper DB2 Recovery Kit</a>	IBM Db2 Universal Database v10.5, v11.1  IBM Db2 Enterprise Server Edition (ESE) v10.5, v11.1 and v11.5  IBM Db2 Workgroup Server Edition (WSE) v10.5, v11.1 and v11.5  IBM Db2 Express Edition v10.5, v11.1 and v11.5  LifeKeeper NFS Server Recovery Kit v5.1 or later (for DB2 EEE and DB2 ESE with multiple partitions only)
<a href="#">LifeKeeper MySQL Recovery Kit</a>	MySQL and MySQL Enterprise v5.7 and v8.0  MariaDB v10.3, v10.4, and v10.5
<a href="#">LifeKeeper PostgreSQL Recovery Kit</a>	PostgreSQL v9.6, v10, v11, v12, v13 and v14  EnterpriseDB Postgres Plus Advanced Server/EDB Postgres Advanced Server v9.6, v10.0, v11.0, v12.0, v13.0 and v14.0  PowerGres Plus (for Linux) v10

PowerGres on Linux v11, v13

The following edition of FUJITSU Software Symfoware Server.

Symfoware Server V12.2

- Symfoware Server (Open Interface) V12.2 Enterprise Edition
- Symfoware Server (Open Interface) V12.2 Standard Edition
- Symfoware Server V12.3
- Symfoware Server (Postgres) V12.3 Enterprise Edition
- Symfoware Server (Postgres) V12.3 Standard Edition
- Symfoware Server (Postgres) V12.3 Lite Edition br>
- Symfoware Server V12.4
- Symfoware Server (Postgres) V12.4 Enterprise Edition
- Symfoware Server (Postgres) V12.4 Standard Edition

The following edition of FUJITSU Software Enterprise Postgres 9.5

- FUJITSU Software Enterprise Postgres 9.5 Advanced Edition
- FUJITSU Software Enterprise Postgres 9.5 Standard Edition

The following edition of FUJITSU Software Enterprise Postgres 9.6

- FUJITSU Software Enterprise Postgres 9.6 Standard Edition

The following edition of FUJITSU Software Enterprise Postgres 10

- FUJITSU Software Enterprise Postgres 10 Advanced Edition
- FUJITSU Software Enterprise Postgres 10 Standard Edition
- FUJITSU Software Enterprise Postgres 10 Community Edition

The following editions of FUJITSU Software Enterprise Postgres 11

- FUJITSU Software Enterprise Postgres 11 Advanced Edition
- FUJITSU Software Enterprise Postgres 11 Standard Edition

	<ul style="list-style-type: none"> <li>• FUJITSU Software Enterprise Postgres 11 Community Edition</li> </ul> <p>The following editions of FUJITSU Software Enterprise Postgres 12</p> <ul style="list-style-type: none"> <li>• FUJITSU Software Enterprise Postgres 12 Advanced Edition</li> <li>• FUJITSU Software Enterprise Postgres 12 Standard Edition</li> <li>• FUJITSU Software Enterprise Postgres 12 Community Edition</li> </ul>
<a href="#">LifeKeeper Sybase ASE Recovery Kit</a>	Sybase ASE 15.7 and 16.0
<a href="#">LifeKeeper Samba Recovery Kit</a>	Standard Samba file services provided with the supported Linux distributions
<a href="#">LifeKeeper NFS Server Recovery Kit</a>	<p>Linux kernel version 2.6 or later</p> <p>The NFS Server and client packages must be installed on SLES systems.</p> <p>NFSv2 is not supported on Red Hat Enterprise Linux 7 or later, CentOS 7 or later, Oracle Linux 7 or later.</p> <p>NFS over UDP is not supported on Red Hat Enterprise Linux 8 and later.</p> <p>Some environments may require additional configurations. Refer to <a href="#">NFS Specific Configuration Considerations</a>.</p>
<a href="#">LifeKeeper Network Attached Storage Recovery Kit</a>	NFS version of Mounted NFS file systems from an NFS server or Network Attached Storage (NAS) device v2, v3 and v4
<a href="#">LifeKeeper Logical Volume Manager (LVM) Recovery Kit</a>	Linux Logical Volume Manager (LVM) Version 1 or 2 volume groups and logical volumes
EMC PowerPath	<p>PowerPath for Linux v5.3 or later</p> <p>The sg3_utils package must be installed.</p>
Device Mapper Multipath (DMMP)	<p>The device-mapper-multipath package attached to the operating system.</p> <p>The sg3_utils package must be installed.</p>
Hitachi Dynamic Link Manager Software (HDLM)	<p>Please see <a href="#">Hitachi Dynamic Link Manager Software Multipath I/O Configurations and Linux Distribution Requirements</a>.</p> <p>The sg3_utils package must be installed.</p>

NEC iStorage Storage Path Savior (NECSPS)	<p>iStorage StoragePathSavior for Linux v3.3 or later</p> <p>For the supported Linux kernel and distribution, please refer to the support information of StoragePathSystem for Linux.</p> <p>The sg3_utils package must be installed on Red Hat and SLES.</p>
<a href="#">WebSphere MQ Resources</a>	<p>IBM MQ v8.0, v9.0, v9.1, v9.2</p> <p>See <a href="#">Known Issues and Restrictions &gt; Installation</a>.</p>
<a href="#">Quorum/Witness Package</a>	<p>All nodes which will participate in a quorum/witness mode cluster, including witness-only nodes, should be installed with the Quorum/Witness Server Support Package for LifeKeeper.</p>

## Open Source Packages

The following open source packages are included in the LifeKeeper installation image.

Name	Version	License Type and Version
curl-7.68.0-1	7.68.0	MIT
libcurl-7.68.0-1	7.68.0	MIT
gnutls-2.8.6-3.1	2.8.6	GPLv3+ and LGPLv2+
gnutls-utils-2.8.6-3.1	2.8.6	GPLv3+
libgcrypt-1.5.0-2.1	1.5.0	LGPLv2+
libgpg-error-1.10-2.1	1.10	LGPLv2+
libxml2-2.7.8-7.1	2.7.8	MIT
libxml2-static-2.7.8-7.1	2.7.8	MIT
lighttpd-1.4.41-2	1.4.41	BSD
lighttpd-fastcgi-1.4.41-2	1.4.41	BSD
openssl-1.1.1g-1	1.1.1g	BSDish
openssl-perl-1.1.1g-1	1.1.1g	BSDish
pcre-4.5-2.1	4.5	distributable
pdksh-5.2.14-780.7.1	5.2.14	GPL; distributable
perl-5.8.8-8.2	5.8.8	Artistic or GPL
perl-addons-5.8.8-26	5.8.8	Various
powercli-11.5.0-2	11.5.0	various license

powershell-6.2.3-2	6.2.3	MIT
readline-4.3-14.1	4.3	GPL
runit-2.0.0-4.11	2.0.0	BSD
util-linux-2.31.1-2	2.31.1	GPLv2 and GPLv2+ and LGPLv2+ and BSD with advertising and Public Domain
Perl Config::IniFiles (CPAN module)	2.27	GPL/Artistic (Same as Perl)
openjdk-12.0.2	12.0.2	GPLv2+
kconfig-frontends	4.11.0	GPLv2
balance	3.54	GPL
mdadm	3.2.6	GPL v2
nbd-client	1.0	GPL
nbd-server	1.3	GPL
HADR-CentOS-2.6.32-all	2.6.32	GPLv2
HADR-CentOS-3.10.0-514.el7	3.10.0	GPLv2
HADR-CentOS-3.10.0-693.el7	3.10.0	GPLv2
HADR-CentOS-3.10.0-862.el7	3.10.0	GPLv2
HADR-CentOS-3.10.0-all	3.10.0	GPLv2
HADR-CentOS-4.18.0-147.el8	4.18.0	GPLv2
HADR-CentOS-4.18.0-193.el8	4.18.0	GPLv2
HADR-CentOS-4.18.0-240.el8	4.18.0	GPLv2
HADR-CentOS-4.18.0-305.el8	4.18.0	GPLv2
HADR-CentOS-4.18.0-all	4.18.0	GPLv2
HADR-OEL-2.6.32-all	2.6.32	GPLv2
HADR-OEL-3.10.0-514.el7	3.10.0	GPLv2
HADR-OEL-3.10.0-693.el7	3.10.0	GPLv2
HADR-OEL-3.10.0-862.el7	3.10.0	GPLv2
HADR-OEL-3.10.0-all	3.10.0	GPLv2
HADR-OEL-4.18.0-147.el8	4.18.0	GPLv2
HADR-OEL-4.18.0-193.el8	4.18.0	GPLv2
HADR-OEL-4.18.0-240.el8	4.18.0	GPLv2
HADR-OEL-4.18.0-305.el8	4.18.0	GPLv2

HADR-OEL-4.18.0-all	4.18.0	GPLv2
HADR-RHAS-2.6.32-71.el6	2.6.32	GPLv2
HADR-RHAS-2.6.32-all	2.6.32	GPLv2
HADR-RHAS-3.10.0-514.el7	3.10.0	GPLv2
HADR-RHAS-3.10.0-693.el7	3.10.0	GPLv2
HADR-RHAS-3.10.0-862.el7	3.10.0	GPLv2
HADR-RHAS-3.10.0-all	3.10.0	GPLv2
HADR-RHAS-4.18.0-147.el8	4.18.0	GPLv2
HADR-RHAS-4.18.0-193.el8	4.18.0	GPLv2
HADR-RHAS-4.18.0-240.el8	4.18.0	GPLv2
HADR-RHAS-4.18.0-305.el8	4.18.0	GPLv2
HADR-RHAS-4.18.0-348.el8	4.18.0	GPLv2
HADR-RHAS-4.18.0-all	4.18.0	GPLv2
HADR-OEL.UEK-4.14.35-2025.400.8.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.400.9.1.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.400.9.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.401.4.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.402.2.1.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.403.3.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.404.1.1.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.404.1.2.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2025.405.3.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.500.10.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.500.9.1.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.500.9.3.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.501.1.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.501.2.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.502.4.1.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.502.4.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.502.5.el7uek	4.14.35	GPLv2
HADR-OEL.UEK-4.14.35-2047.503.1.1.el7uek	4.14.35	GPLv2

HADR-OEL.UEK-4.14.35-2047.503.1.el7uek	4.14.35	GPLv2
HADR-SuSE-4.12.14-122.20.1	4.12.14	GPLv2
HADR-SuSE-4.12.14-197.37.1	4.12.14	GPLv2
HADR-SuSE-4.12.14-95.51.1	4.12.14	GPLv2
HADR-SuSE-5.3.18-all	5.3.18	GPLv2

## Installation and Configuration

See the [LifeKeeper for Linux Installation Guide](#) for complete installation and configuration information.

## Upgrades

LifeKeeper can be upgraded to Version 9.6.1 from either LifeKeeper Version 9.4.x or Version 9.5.x. If upgrading from a version other than 9.4.x or 9.5.x, the older version will need to be uninstalled and SIOS Protection Suite for Linux will have to be reinstalled. An alternative to uninstalling the older version would be to upgrade from the older version to 9.4.x or 9.5.x, then perform the upgrade to 9.6.x.

## Storage and Adapter Options

For a list of the disk array storage models and adapters currently supported by LifeKeeper in shared storage configurations as well as their type of certification, see the [Storage and Adapter Options](#) topic. Details about driver versions and other configuration requirements for these arrays and adapters are listed in the [Storage and Adapter Configuration](#) topic.

## Technical Notes

We strongly recommend that you read the [Technical Notes](#) section concerning configuration and operational issues related to your LifeKeeper environment.

## Known Issues

See [Known Issues and Restrictions](#) in the [Troubleshooting](#) section of [LifeKeeper for Linux Technical Documentation](#) and the [DataKeeper Troubleshooting](#) section..

### Trademarks:

- “Amazon Web Services, the “Powered by Amazon Web Services” logo, “AWS”, “Amazon EC2”, “EC2”, “Amazon Elastic Compute Cloud”, “Amazon Virtual Private Cloud”, and “Amazon VPC” are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries.
  - Google Cloud, BigQuery and Google Compute Engine are trademarks of Google LLC.
  - “Oracle Cloud” is a registered trademark of Oracle Corporation and its affiliates.
- Trademark symbols such as ® and ™ may be omitted from system names and product names in this document.

# Important Notice: Kernel issue with SIOS DataKeeper for Linux

## **IMPORTANT NOTICE:** Patch Available, Immediate Installation Required

We are pleased to announce that we have completed all Engineering (Development and Quality Assurance) related activities for the issue that impacted the following kernel and versions:

Distribution	Affected Kernels	Supported LifeKeeper for Linux versions for the Distribution at Left
RHEL/CentOS/OEL 8.2	All	LifeKeeper for Linux v9.5.1
RHEL 8.3	All	LifeKeeper for Linux-L v9.5.1
SLES 12 SP4	>= 4.12.14-95.51	LifeKeeper for Linux-L v9.3.2 – v9.5.1
SLES 12 SP5	4.12.14-122.20 – 4.12.14-122.74.0	LifeKeeper for Linux-L v9.5.0 – v9.5.1
SLES 15 SP1	>= 4.12.14-197.37.1	LifeKeeper for Linux-L v9.4.1 – v9.5.1
SLES 15 SP2	5.3.18-22.2 – 5.3.18-24.67.1	LifeKeeper for Linux-L v9.5.1
OEL 7.x UEK 5	4.14.35-2025.400.8 – 4.14.35-2047.504.1	LifeKeeper for Linux-L v9.4.0 – v9.5.1

We recommend that all customers running LifeKeeper for Linux v9.5.1 or earlier with DataKeeper replication on the operating systems listed above apply the LifeKeeper for Linux patch as soon as possible. The patch is located at: [http://ftp.us.sios.com/pickup/HOTFIX-PL-9146-raid1\\_data\\_integrity\\_patch](http://ftp.us.sios.com/pickup/HOTFIX-PL-9146-raid1_data_integrity_patch). Patched versions of the raid1 kernel module are distributed natively with v9.5.2 and are automatically installed on affected systems when installing or upgrading to v9.5.2. Please see the [readme](#) for the required steps to upgrade LifeKeeper for Linux on systems running the patch.

The steps required to successfully install the patch are outlined in the readme file located in the patch folder. [http://ftp.us.sios.com/pickup/HOTFIX-PL-9146-raid1\\_data\\_integrity\\_patch/readme.txt.html](http://ftp.us.sios.com/pickup/HOTFIX-PL-9146-raid1_data_integrity_patch/readme.txt.html)

The readme file will assist you with application of the patch by describing the pre-conditions, steps, and post-conditions and by providing recommended steps for verifying your data after the patch has been applied. This patch addresses the issue in the impacted raid1 kernel module immediately.

After you have applied the patch, we recommend that you do not update the kernel to a different version that does not include our patched raid1 kernel module. Refer to your Operating System documentation for restricting the automatic kernel updates. Also, you may refer to [Solution 995 in our Self Service portal](#) for restricting kernel updates with RHEL (log into the [Customer Portal](#) first to access). We will continue to provide updates for the final resolution.

If a kernel upgrade is required after applying the patch (due to security issues or regulations), please follow the steps in the “Performing a planned kernel upgrade after the patch has been installed” section of the [readme file](#).

If you have additional questions or would like assistance applying the patch, please contact your [SIOS Sales Representative](#) or [SIOS Technical Support](#).

We are committed to providing you with the highest quality of products and services and we are always available to answer any questions you may have.

# 3. LifeKeeper for Linux Getting Started Guide

---

This document will guide you through the installation of the LifeKeeper for Linux and assumes the user has basic knowledge of the Linux operating system. Please refer to the [LifeKeeper for Linux product documentation](#) for more information.

## Pre-Installation Requirements

Before installing LifeKeeper for Linux, please check the following:

- [LifeKeeper for Linux Release Notes](#) -The Release Notes include supported platforms, operating systems, applications, and storage. They also include the latest features and Bug Fixes.
- **TCP/IP Connection and Name Resolution** – In order to use the GUI function, both cluster nodes need to be able to resolve the name. Use the DNS service or `/etc/hosts` for name resolution. Also, localhost needs to be resolved to 127.0.0.1.
- **Firewall** – The following ports are used:
  - Communication Path (TCP): 7365/tcp
  - Communication of a GUI Server: 81/tcp, 82/tcp
  - RMI Communication between the GUI Server and Client: all the ports after 1024/tcp
  - Synchronization of DataKeeper (when using DataKeeper): “10001+<mirror number>+<256 \* i>”

### More Firewall Information

- The port used for communication with the GUI server and a client needs to be open on the cluster node where LifeKeeper is installed and on all systems where the GUI client runs.
- The ports used by DataKeeper can be calculated using the formula above. The value of i starts at 0 and uses an unused port when found. For example, in an environment where a DataKeeper resource with mirror number 0 exists, if port 10001 is being used by another application, port 10257 will be used.
- For communication between the GUI server and a client, Java RMI (Remote Method Invocation) randomly uses ports 1024 and above. When applying access control etc. to a cluster system, packet filtering needs to be performed considering these ports. If this specification is an issue from a security standpoint, you can use ssh X forwarding. Please refer to the [Technical Documentation](#) for the setting details.
- **Check the SELinux Setting** – When the SELinux setting is enabled, LifeKeeper for Linux cannot be installed. Please refer to the OS distribution documentation on how to disable SELinux. It is not recommended to use SELinux permissive mode unless it is required in an SAP environment. Please make sure that the application to be run on the cluster supports a permissive mode. SELinux permissive mode has been tested for following ARKs: SAP / SAP MaxDB / Sybase / Oracle / DB2 / NFS / DataKeeper / NAS / EC2 / IP / FileSystem / MQ  
Refer to [Linux Dependencies](#) for required packages.
  - Install the appropriate package provided by your distribution.
  - The sg3\_utils package is required for environments using recovery kits for Multipath such as the DMMP Recovery Kit and the PowerPath Recovery Kit. This is not required for

environments where recovery kits for Multipath are not used.

- **Check [Known Issues](#)** – Please make sure that there are no known issues for your environment.

## Installing LifeKeeper for Linux

 **Note:** These installation instructions assume that you are familiar with the Linux operating system installed on your servers.

Time required for deployment: Approx. 2-3 hours. Actual times may vary depending on individual customer environment.

Install the LifeKeeper software on each server in the LifeKeeper configuration.

Packages that LifeKeeper is dependent on are installed automatically because the LifeKeeper installation setup script uses package manager tools (yum or zypper) to ensure installation of all dependent packages.

 **IMPORTANT:** A functional yum or zypper configuration is required for the successful installation of LifeKeeper. A non-functional configuration can result in an installation failure (see [Installation Known Issues](#) for more information). Additionally, the package manager repo or rpm database must not be locked as that could cause the install to hang. If the dependent packages cannot be installed automatically via the package manager, refer to [Linux Dependencies](#) and install the necessary packages in advance.

The LifeKeeper for Linux image file (sps.img) provides a set of installation scripts designed to perform the user interactive system setup tasks that are necessary when installing LifeKeeper on your system.

A licensing utilities package is also installed providing utilities for obtaining and displaying the Host ID or Entitlement ID of your server. Host IDs and/or Entitlement IDs are used to obtain valid licenses for running LifeKeeper.

 **IMPORTANT:**

- Installing LifeKeeper on your shared storage is not supported. Each server should have its own copy installed on its local disk.
- All LifeKeeper packages are installed in the directory `/opt/LifeKeeper`.

## Obtaining and Installing the License

LifeKeeper for Linux requires a unique license for each server. The license is a run-time license, which means that you can install LifeKeeper without it, but the license must be installed before you can successfully start and run the product.

**Note:** If using newer hardware with RHEL 6.1, please see the IP Licensing [Known Issues](#) in the LifeKeeper for Linux [Troubleshooting](#) Section.

The Installation script installs the Licensing Utilities package which obtains and displays all of the available Host IDs for your server during the initial install of your LifeKeeper Software. Once your

licenses have been installed the utility will return the Entitlement ID if it is available or the Host IDs if it is not.

**Note:** Host IDs, if displayed will always be based on the MAC address of the NICs.

The new licenses obtained from the [SIOS Technology Corp. Licensing Operations Portal](#) will contain your Entitlement ID and will be locked to a specific node or IP address in the cluster. The Entitlement ID (Authorization Code) which was provided with your LifeKeeper Software, is used to obtain the permanent license required to run the LifeKeeper Software. The process is illustrated below.



**Note:** Each software package requires a license for each server.

**Perform the following steps to obtain and install your license(s) for each server in the LifeKeeper cluster:**

1. **Ensure you have your LifeKeeper Entitlement ID** (Authorization Code). You should have received an email with your software containing the Entitlement ID needed to obtain the license.
2. **Obtain your licenses from the SIOS Technology Corp. Licensing Operations Portal.**
  - a. Using the system that has internet access, log in to the [SIOS Technology Corp. Licensing Operations Portal](#).
  - b. From the **Activation & Entitlements** dropdown list select **List Entitlements**. **Note:** If changing password, use the **Profile** button in the upper right corner of the display.
  - c. Find your **Entitlement ID** and select each **Activation ID** associated with that Entitlement ID by checking the box to the left of the line item.
  - d. From the **Action** dropdown list select **Activate**.
  - e. Define the required fields and select **Next**.
  - f. Click on the **Green Plus Sign** to add a new host.
  - g. Select and Define the required fields and click **Okay**. (**Note:** Internet = IP address, Ethernet = MAC address)

- h. Check the box to the left of the **Host ID** or **IP address** and select **Generate**. The **Fulfillment ID** will display on the **License Summary** screen.
- i. Select **Complete**.
- j. Check the box to the left of the **Fulfillment ID** and select the **Email** from the View dropdown list.
- k. Enter a valid email address to send the license to and select **Send**.
- l. Retrieve the email(s).
- m. Copy the file(s) to a temporary directory on each node. **Make sure that the licenses match the MAC address**. This path and filename(s) will be used during the 'Install License Key' portion of the setup script.

 **NOTE:** To install the license outside of the 'Setup' script, copy the license file(s) to `/var/LifeKeeper/license` on each system, or run `/opt/LifeKeeper/bin/lkkeyins` and specify the filename (including full path) to the file.

## How to Install / Upgrade LifeKeeper Using the Setup Script

To install or upgrade LifeKeeper, follow the steps below.

### Interactive Mode

1. After logging in as the root user, use the following command to mount the sps.img file:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

Where PATH is the path to the image  
IMAGE\_NAME is the name of the image  
MOUNT\_POINT is the path to mount location

2. Change to the directory where sps.img is mounted and enter the following:

```
./setup
```

3. The script collects information about the system environment and determines what you need to do to install LifeKeeper.

If the system requirements for installation or upgrade are not satisfied, then an error message is displayed and the installation / upgrade is cancelled.

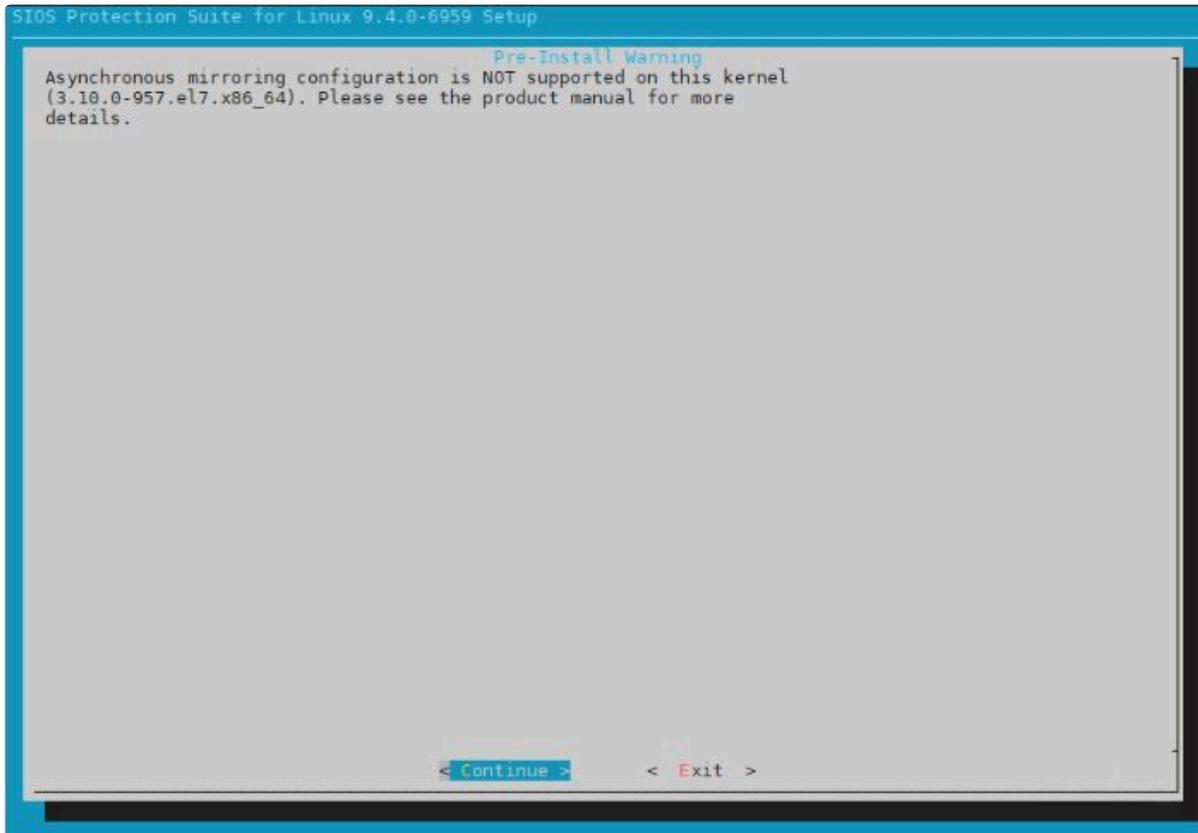
Also, if some restrictions arise or a configuration change is required, a warning message will be displayed requiring the user to decide whether to continue or abort the installation.

4. Select the LifeKeeper features and Application Recovery Kits (ARKs) to install via the main dialog

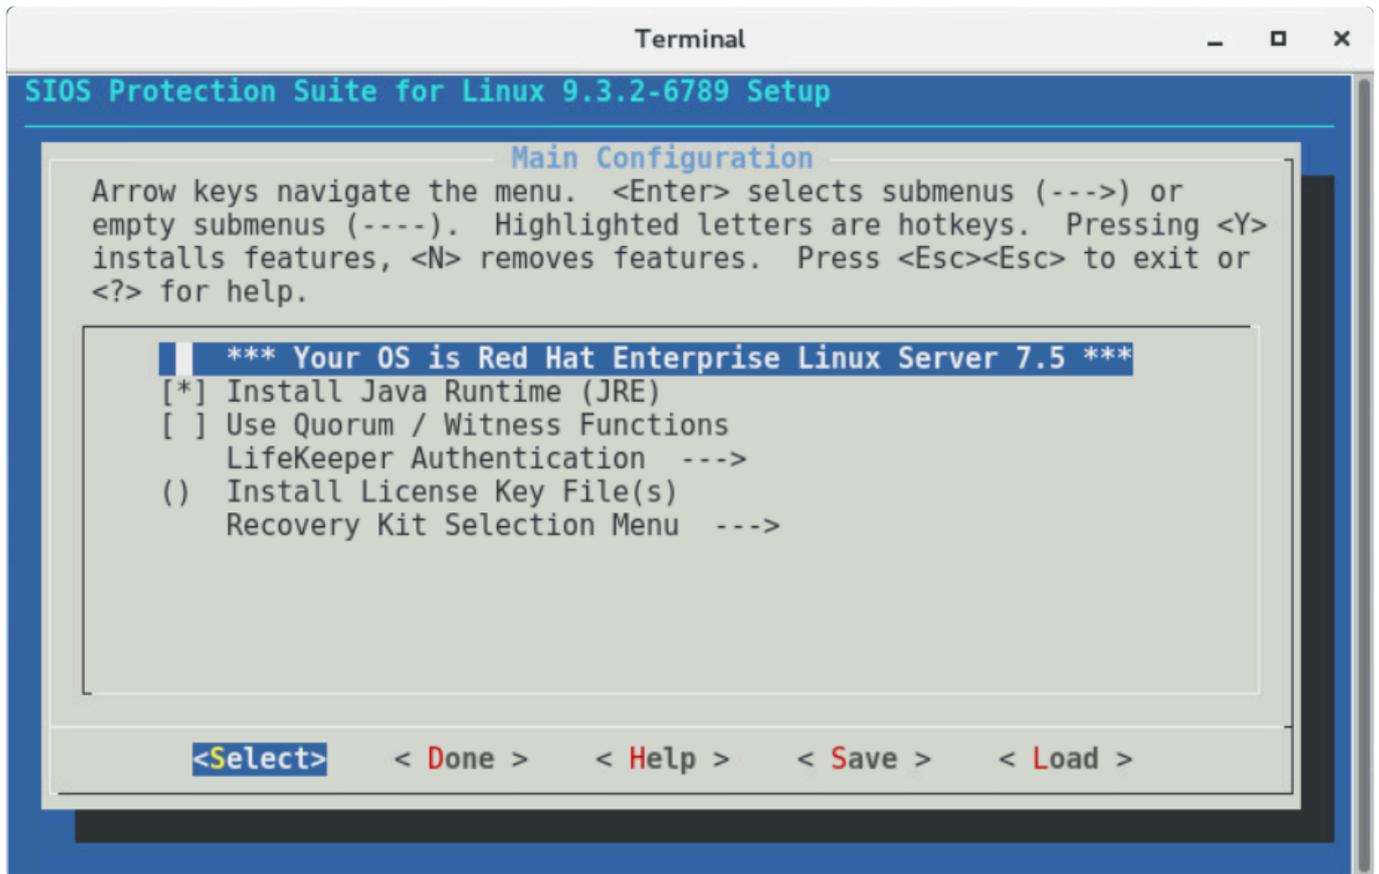
screen.

## How to Use the Dialog Screen

If the kernel version is not supported for asynchronous mirroring the following dialog will appear.



The dialog screen is displayed below.



Use the following keys to navigate the menu.

↑ ↓	Navigate between menu items
← →	Navigate between the menu buttons at the bottom of the screen
ENTER	Open the selected sub menu
Y / N / SPACE	Turn on, turn off or invert the selected item

The menu buttons at the bottom of the screen are used for the following operations.

Select	Opens a sub menu dialog screen
Done	Closes the current screen and returns to the previous screen. Selecting this button on the main screen completes the configuration.
Help	Displays help for the highlighted item
Save	Saves the current settings in a configuration file. The saved configuration file can be used for non-interactive installations.
Load	Loads a saved configuration file

The “Save” and “Load” menu buttons display a dialog screen asking for a configuration file name for use in saving the current configuration or for loading a saved configuration. If you want to change the default file name provided, move to the file name field using the [TAB] key, and enter a new name. **Note:** The “Save” operation will prompt for confirmation before overwriting a file with the same name.

The items listed below are configurable during installation. During an upgrade only items that can be configured are listed. Using the hotkey <Z> will show those items that will remain unchanged during the upgrade.

- **Install Java Runtime (JRE)**

Install the Java runtime environment used by the LifeKeeper GUI.

- **Restart NFS Service**

When configuring High Availability NFS, restarting the NFS services is required. When this is selected, the services are restarted automatically after the configuration is completed.

**Note:** If you do not want to restart the NFS services automatically, a restart will need to be done to pick up the configuration changes before using the NFS Recovery Kit.

- **Use Quorum / Witness Functions**



Note: Quorum is required for all 2 node DataKeeper clusters. Each 2 node cluster must use its own quorum/witness node. Shared witness servers are not recommended.

Use Quorum / Witness for I/O fencing. For details, please refer to [Quorum/Witness](#) in the technical documentation.

The Quorum/Witness Server Support Package for LifeKeeper will need to be installed on every node in the cluster that uses quorum/witness functionality, including a witness-only node. The only configuration requirement for the witness node is to [create appropriate comm paths](#). When using a quorum mode with tcp\_remote, LifeKeeper does not need to be installed on the host which was set as QUORUM\_MODE in /etc/default/LifeKeeper configuration file.

The general process for setting up quorum/witness functionality will involve the following steps:

1. Set up the server and make sure that it can communicate with other servers.
2. Install LifeKeeper on the server. During the installation, enable “Use Quorum / Witness functions” with the setup command and install the quorum/witness package as well.
3. Create appropriate communication paths between the nodes including witness-only nodes.
4. [Configure quorum/witness](#).

When the above steps are completed, the quorum/witness functions will be activated in the cluster and quorum checking and witness checking will be performed before failovers are allowed.

- **LifeKeeper Authentication**

Specify the users allowed to log in to the LifeKeeper for Linux GUI along with their privilege levels. Multiple user accounts can be specified by separating them with blanks. For details, refer to [GUI User Settings](#).

- **Install License Key File(s)**

Install the licenses required to start LifeKeeper for Linux by entering the path name of the license file to install. Multiple files can be specified by separating them with spaces.

- **Recovery Kit Selection**

Select the Application Recovery Kits to install.

Application Recovery Kits are broken into several categories based on common functionality.

Please refer to [Categories for Application Recovery Kits](#) for details.

- **LifeKeeper Startup After Install**

When selected, LifeKeeper for Linux will be started when the installation is completed.

 **Note:** Because the LifeKeeper Data Replication package may install kernel modules for some of the supported OS distributions, a re-install of LifeKeeper may be required when the kernel is upgraded. This applies to OS distributions for RedHat, CentOS and Oracle Enterprise Linux (non-UEK kernels only) running kernel versions 3.10.0-514 or later.

## Categories for Application Recovery Kits

Category	Description
Application Suite	A group of recovery kits that protect applications such as SAP and IBM MQ.
Networking	A group of recovery kits that protect network services in the cloud such as EC2 and Route53.
Database	A group of recovery kits that protect database applications, including, but not limited to, Oracle, PostgreSQL, and MaxDB.
File Sharing	A group of recovery kits that protect file sharing services such as NFS and Samba.
Mail Server	A group of recovery kits that protect email services such as Postfix.
Storage	A group of recovery kits that protect data storage methods, including, but not limited to, DataKeeper (replication), Device Mapper (DM) Multipath (DMMP), and Network Attached Storage (NAS).
Web Server	A group of recovery kits that protect web services such as Apache.

5. Once all the required LifeKeeper features and ARKs have been selected, select <Done> to begin the installation.

If any notifications are output when the installation completes, please take the necessary actions to correct them.

## Creating a Cluster

✿ LifeKeeper must be installed on all systems before creating a cluster.

To create a cluster system first you need set up a “communication path” between the nodes that make up a HA cluster. Then create “resources” to define what to protect.

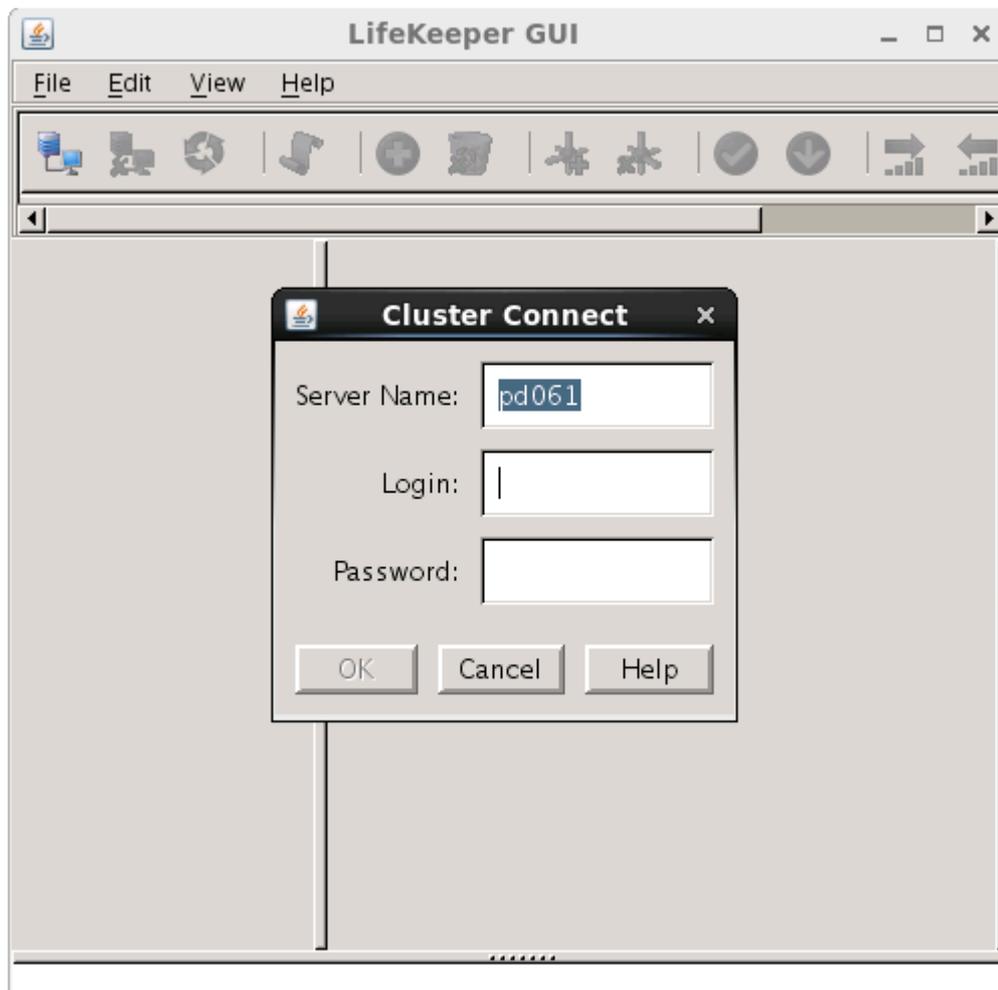
### Connecting with LifeKeeper GUI Client

#### Configure LifeKeeper using the GUIs.

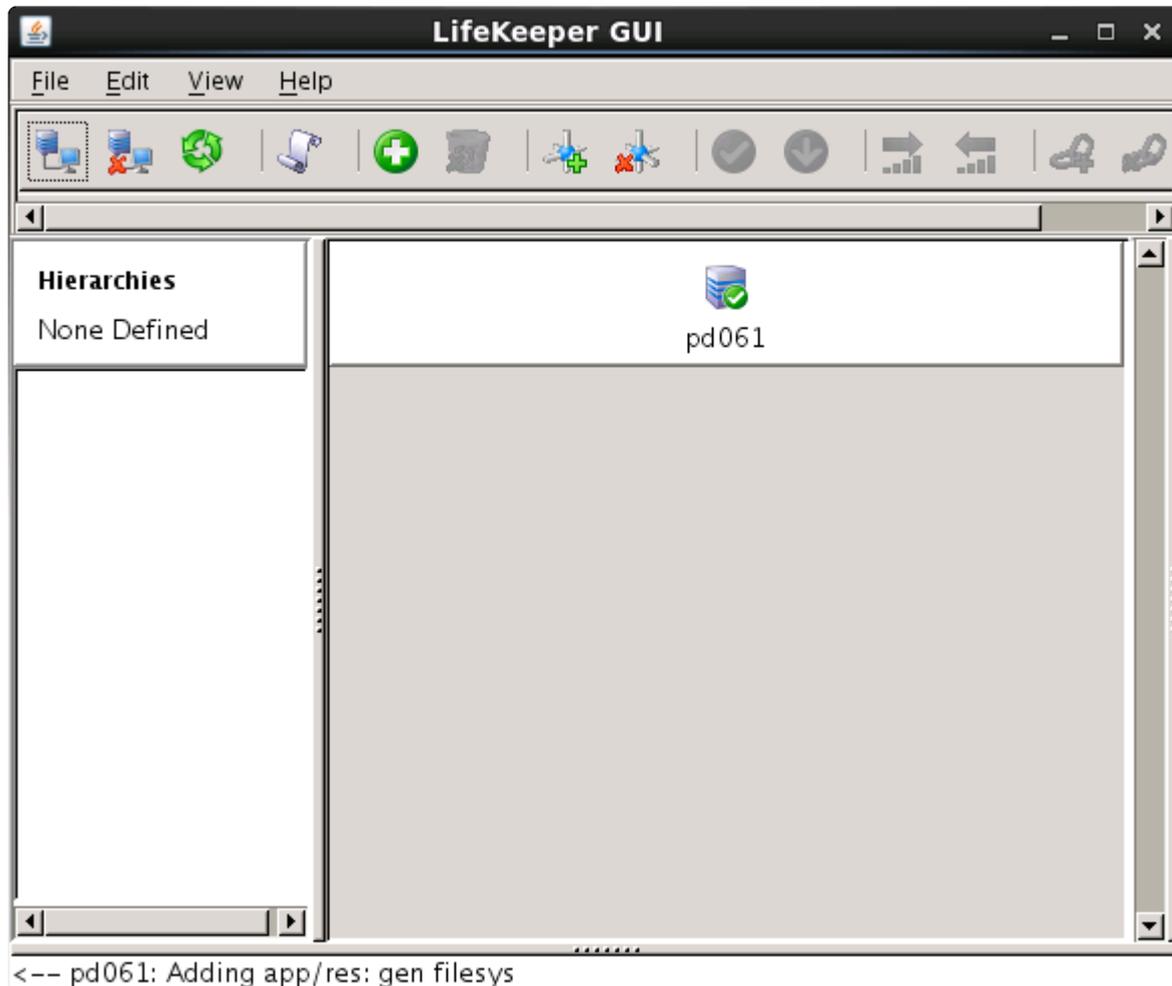
The GUI client is started by lkGUIapp command. After starting LifeKeeper, start the LifeKeeper GUI client with the following command.

```
# lkGUIapp
java version "1.8.0_51"
Java(TM) SE Runtime Environment (build 1.8.0_51-b16)
Java HotSpot(TM) 64-Bit Server VM (build 25.51-b03, mixed mode)
Setting up secure random number generator
Random number setup completed
█
```

After executing the command, the GUI client is started and the login screen is launched. Server Name is the name of the server you are running. For login username and password, enter the LifeKeeper admin user name and password. By default, the operating system super user (root) and its password are used for the admin user.



After successfully logging in the following screen is displayed.



## Creating a Communication Path

To create a communication path between a pair of servers, you must define the path individually on both servers. LifeKeeper allows you to create both TCP (TCP/IP) and TTY communication paths between a pair of servers. Only one TTY path can be created between a given pair. However, you can create multiple TCP communication paths between a pair of servers by specifying the local and remote addresses that are to be the end-points of the path. A priority value is used to tell LifeKeeper the order in which TCP paths to a given remote server should be used.

**IMPORTANT:** Using a single communication path can potentially compromise the ability of servers in a cluster to communicate with one another. If a single comm path is used and the comm path fails, LifeKeeper hierarchies may come in service on multiple servers simultaneously. This is known as “false failover”. Additionally, heavy network traffic on a TCP comm path can result in unexpected behavior, including false failovers and LifeKeeper initialization problems.

1. On the global toolbar, click the **Create Comm Path** button.
2. A dialog entitled **Create Comm Path** will appear. For each of the options that follow, click **Help** for an explanation of each choice.
3. Select the **Local Server** from the list box and click **Next**.

4. Select one or more **Remote Servers** in the list box. If a remote server is not listed in the list box (i.e. it is not yet connected to the cluster), you may enter it using **Add**. You must make sure that the network addresses for both the local and remote servers are resolvable (for example, with DNS or added to the `/etc/hosts` file). Click **Next**.
5. Select either **TCP** or **TTY** for **Device Type** and click **Next**.
6. Select one or more **Local IP Addresses** if the **Device Type** was set for **TCP**. Select the **Local TTY Device** if the **Device Type** was set to **TTY**. Click **Next**.
7. Select the **Remote IP Address** if the **Device Type** was set for **TCP**. Select the **Remote TTY Device** if the **Device Type** was set to **TTY**. Click **Next**.
8. Enter or select the **Priority** for this comm path if the **Device Type** was set for **TCP**. Enter or select the **Baud Rate** for this Comm Path if the **Device Type** was set to **TTY**. Click **Next**.
9. Click **Create**. A message should be displayed indicating the network connection is successfully created. Click **Next**.
10. If you selected multiple Local IP Addresses or multiple Remote Servers and the **Device Type** was set for **TCP**, then you will be taken back to Step 6 to continue with the next Comm Path. If you selected multiple Remote Servers and the **Device Type** was set for **TTY**, then you will be taken back to Step 5 to continue with the next Comm Path.
11. Click **Done** when presented with the concluding message.

You can verify the comm path by viewing the [Server Properties Dialog](#) or by entering the command `lcdstatus -q`. See the `LCD` man page for information on using `lcdstatus`. You should see an **ALIVE** status.

In addition, check the server icon in the right pane of the GUI. If this is the first comm path that has been created, the server icon shows a yellow heartbeat, indicating that one comm path is **ALIVE**, but there is no redundant comm path. 

The server icon will display a green heartbeat when there are at least two comm paths **ALIVE**. 

## Creating Resource Hierarchies

Create resources for the services and applications you want to protect.

1. On the global toolbar, click on the **Create Resource Hierarchy** button.
2. A dialog entitled Create Resource Hierarchy will appear with a list of all recognized recovery kits installed within the cluster. Select the **Recovery Kit** that builds resource hierarchies to protect your application and click **Next**.
3. Select the **Switchback Type** and click **Next**.

4. Select the **Server** and click **Next**. **Note:** If you began from the server context menu, the server will be determined automatically from the server icon that you clicked on, and this step will be skipped.
5. Continue through the succeeding dialogs, entering whatever data is needed for the type of resource hierarchy that you are creating.

## Recovery Kit Options

Each optional recovery kit that you install adds entries to the Select Recovery Kit list; for example, you may see Oracle, Apache, and NFS Recovery Kits. Refer to the Administration Guide that accompanies each recovery kit for directions on creating the required resource hierarchies.

 Please refer to the procedure for creating each resource for the Recovery Kits in the [Application Recovery Kit Documentation](#). There you will find setup requirements for each Recovery Kit.

## Creating a File System Resource Hierarchy

1. On the global toolbar, click on the **Create Resource Hierarchy** button.
2. A dialog entitled *Create Resource Wizard* will appear with a **Recovery Kit** list. Select *File System Resource* and click **Next**.
3. Select the **Switchback Type** and click **Next**.
4. Select the **Server** and click **Next**. **Note:** *If you began from the server context menu, the server will be determined automatically from the server icon that you clicked on, and this step will be skipped.*
5. The *Create gen/filesys Resource* dialog will now appear. Select the **Mount Point** for the file system resource hierarchy and click **Next**. The selected mount point will be checked to see that it is shared with another server in the cluster by checking each storage kit to see if it recognizes the mounted device as shared. If no storage kit recognizes the mounted device, then an error dialog will be presented:

**<file system>** is not a shared file system

Selecting **OK** will return to the *Create gen/filesys Resource* dialog.

### Notes:

- In order for a mount point to appear in the choice list, the mount point must be currently mounted. If an entry for the mount point exists in the `/etc/fstab` file, LifeKeeper will remove this entry during the creation and extension of the hierarchy. It is advisable to make a backup of `/etc/fstab` prior to using the NAS Recovery Kit, especially if you have complex mount settings. You can direct that entries are re-populated back into `/etc/fstab` on deletion by setting the `/etc/default/LifeKeeper` tunable `REPLACEFSTAB=true|TRUE`.

- Many of these resources (SIOS DataKeeper, LVM, Device Mapper Multipath, etc.) require LifeKeeper recovery kits on each server in the cluster in order for the file system resource to be created. If these kits are not properly installed, then the file system will not appear to be shared in the cluster.
6. LifeKeeper creates a default **Root Tag** for the file system resource hierarchy. (This is the label used for this resource in the status display). You can select this root tag or create your own, then click **Next**.
  7. Click **Create Instance**. A window will display a message indicating the status of the instance creation.
  8. Click **Next**. A window will display a message that the file system hierarchy has been created successfully.
  9. At this point, you can click **Continue** to move on to [extending the file system resource hierarchy](#), or you can click **Cancel** to return to the GUI. If you click **Cancel**, you will receive a warning message that your hierarchy exists on only one server, and it is not protected at this point.

## Frequently Used Commands

- Starting the LifeKeeper GUI client
 

```
# /opt/LifeKeeper/bin/lkGUIapp
```
- Starting LifeKeeper
 

```
# /opt/LifeKeeper/bin/lkstart
```
- Stopping LifeKeeper (stopping resources)
 

```
# /opt/LifeKeeper/bin/lkstop
```
- Stopping LifeKeeper (without stopping resources)
 

```
# /opt/LifeKeeper/bin/lkstop -f
```
- Checking a status of LifeKeeper
 

Specify “-e” option to display the simple status

```
# /opt/LifeKeeper/bin/lcdstatus (or lcdstatus -e)
```
- Checking a LifeKeeper log
 

Refer to `/var/log/lifekeeper.log`. If you want to check the log output in real time, you can also use the tail command as follows.

```
# tail -f /var/log/lifekeeper.log
```
- Collect LifeKeeper Configuration Information and Logs together
 

```
# /opt/LifeKeeper/bin/lksupport
```
- Backup/Restore of LifeKeeper Configuration Information
 

Taking a backup of the LifeKeeper configuration information

```
# /opt/LifeKeeper/bin/lkbackup -c
```

- Restoring the LifeKeeper configuration information

```
# /opt/LifeKeeper/bin/lkbackup -x -f archive..tar.gz
```

## Support for LifeKeeper for Linux

Contact SIOS Technology Corp. Support at [support@us.sios.com](mailto:support@us.sios.com)

You can also contact SIOS Technology Corp. Support at:

- 1-877-457-5113 (Toll Free)
- 1-803-808-4270 (International)

Email: [support@us.sios.com](mailto:support@us.sios.com)



In order to begin our investigation, we will need the 'lksupport' logs. These are critical in diagnosing the issue/status of the cluster and should be included whenever you contact Support. Run: ***/opt/LifeKeeper/bin>lksupport*** on each node in the cluster. The lksupport command will create a .tar file for each node under the directory: */tmp/lksupport*. For faster diagnosis send all of the logs for each node in the cluster when contacting support.

## 4. LifeKeeper for Linux Installation Guide

---

The LifeKeeper for Linux Installation Guide contains information on how to plan and install your LifeKeeper environment. In addition to providing the necessary steps for setting up your server, storage device and network components, it includes details for configuring your LifeKeeper graphical user interface (GUI).

Once you have completed the steps in this guide, you will be ready to configure your LifeKeeper and DataKeeper resources. The [LifeKeeper for Linux Technical Documentation](#) provides the information needed to complete your LifeKeeper configuration.

### System Requirements

For a complete list of hardware and software requirements and versions, see the [LifeKeeper for Linux Release Notes](#).

Also, before installing LifeKeeper, be sure that you have completed the planning and hardware configuration tasks described in this document.

### Technical Notes

Refer to the [Technical Notes](#) and [Troubleshooting](#) sections of the LifeKeeper for Linux Technical Documentation for information detailing troubleshooting issues, restrictions, etc., pertaining to this software.

## 4.1. Software Packaging

---

The LifeKeeper for Linux software, including [Optional LifeKeeper Recovery Kits](#), is contained within a single image file (sps.img).

### LifeKeeper for Linux Installation Image File

The LifeKeeper for Linux image file (sps.img) provides a set of installation scripts designed to perform the user interactive system setup tasks that are necessary when installing LifeKeeper on your system (see [Interactive Way](#) for more information). A non-user interactive install can be performed as well (see [Non-interactive Way](#) for more information).

The LifeKeeper installation process is broken down into 3 steps:

- Collection
- Selection (user interactive only)
- Installation and Configuration

The first step of the process is the Collection phase and is responsible for collecting information about the system, such as the Linux distribution being used, to ensure the system meets the requirements for a successful install. Step 2 of the process is the Selection phase and is responsible for interacting with the user via a menu based selection process to determine what LifeKeeper packages to install and the configurations required to support those selections. The third and final step is the Installation and Configuration phase. This step is responsible for installing the LifeKeeper Core Package Cluster and Optional Recovery Software, and configuring the system for LifeKeeper. This step also installs any required OS supporting packages that are not already on the system.

The LifeKeeper for Linux image file includes a core package cluster containing the following software packages:

### LifeKeeper Core Package Cluster

- LifeKeeper (**steeleye-ik**). The LifeKeeper core packages provide recovery software for core system components, such as memory, CPUs, the operating system, the SCSI disk subsystem and file systems.
- LifeKeeper GUI (**steeleye-ikGUI**). The LifeKeeper GUI package provides a graphical user interface for LifeKeeper administration and monitoring.
- DataKeeper (**steeleye-ikDR**). The DataKeeper package provides data replication (synchronous or asynchronous mirrors) with intent logging.
- IP Recovery Kit (**steeleye-ikIP**). The LifeKeeper IP Recovery Kit provides switchover software for automatic recovery of IP addresses.
- Raw I/O Recovery Kit (**steeleye-ikRAW**). The LifeKeeper Raw I/O Recovery Kit provides support for applications that use raw i/o to bypass kernel buffering.

- **Man Pages (steeleye-ikMAN)**. The LifeKeeper Man Page package provides reference manual pages for the LifeKeeper product.

## Optional Recovery Software

Recovery kits are also released with the LifeKeeper Core software. During the installation, you will be presented with a complete, up-to-date, selectable list of available recovery kits. For information regarding these recovery kits, see the [Application Recovery Kits](#) section of the LifeKeeper for Linux Technical Documentation.

## 4.2. Planning Your LifeKeeper Environment

---

The following topics will assist in defining the LifeKeeper for Linux cluster environment.

---

[Mapping Server Configurations](#)

[Storage and Adapter Requirements](#)

[Storage and Adapter Options](#)

## 4.2.1. Mapping Server Configurations

Document your server configuration using the following guidelines:

1. Determine the server names, processor types, memory and other I/O devices for your configuration. When you specify a backup server, you should ensure that the server you select has the capacity to perform the processing should a failure occur on the primary server.
2. Determine your communications connection requirements.

**Important:** Potentially, clustered configurations have two types of communications requirements: cluster requirements and user requirements.

- **Cluster** – A LifeKeeper cluster requires at least two communication paths (also called “comm paths” or “heartbeats”) between servers. This redundancy helps avoid “split-brain” scenarios due to communication failures. Two separate LAN-based (TCP) comm paths using dual independent subnets are recommended, and at least one of these should be configured as a private network. Using a combination of TCP and TTY is also supported. A TTY comm path uses an RS-232 null-modem connection between the servers’ serial ports.

Note that using only one comm path can potentially compromise the ability of systems in a LifeKeeper cluster to communicate with each other. If a single comm path is used and the comm path fails, then LifeKeeper hierarchies may try to come into service on multiple systems simultaneously. This is known as a false failover or a “split-brain” scenario. In the “split-brain” scenario, each server believes it is in control of the application and thus may try to access and write data to the shared storage device. To resolve the split-brain scenario, LifeKeeper may cause servers to be powered off or rebooted or leave hierarchies out-of-service to assure data integrity on all shared data. Additionally, heavy network traffic on a TCP comm path can result in unexpected behavior, including false failovers and the failure of LifeKeeper to initialize properly.

- **User** – We recommend that you provide alternate LAN connections for user traffic – that is, a separate LAN connection than the one used for the cluster heartbeat. However, if two TCP comm paths are configured (as recommended), one of those comm paths can share the network address with other incoming and outgoing traffic to the server.

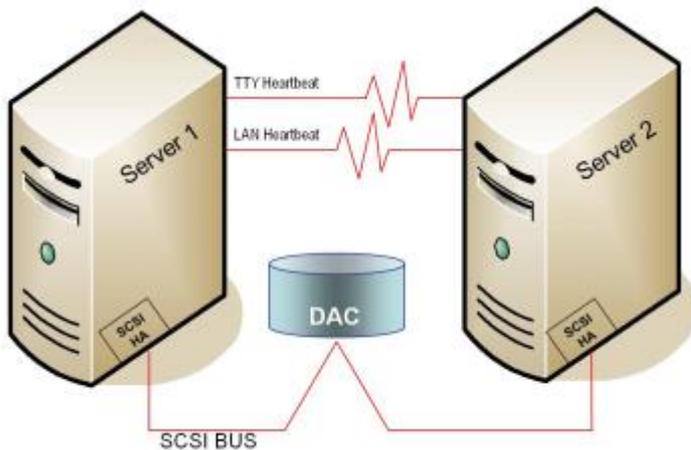
 **Note:** To help ensure that resources are brought into service only when necessary, you may elect to utilize the [Quorum/Witness Server Support Package for LifeKeeper](#).

3. Identify and understand your shared resource access requirements. If you are planning to use LifeKeeper in a data replication (mirroring) environment, see the [SIOS DataKeeper Administration Guide](#). If you are using LifeKeeper in a network attached storage (NAS) environment, see the [LifeKeeper Network Attached Storage Recovery Kit Administration Guide](#). Clusters that use shared storage can utilize shared SCSI buses, Fibre Channel loops, or iSCSI. Because LifeKeeper locks resources to one server, you must ensure that only one server requires access to all locked resources at any given time. LifeKeeper device locking is done at the Logical Unit (LUN) level. For active/active configurations, each hierarchy must access its own unique LUN. All hierarchies

accessing a common LUN must be active (in-service) on the same server.

- Determine your shared memory requirements. Remember to take into account the shared memory requirements of third-party applications as well as those of LifeKeeper when configuring shared memory and semaphore parameters. See [Tuning](#) in [Technical Notes](#) for LifeKeeper's shared memory requirements.

## Sample Configuration Map for LifeKeeper Pair



This sample configuration map depicts a pair of LifeKeeper servers sharing a disk array subsystem where, normally, Server 1 runs the application(s) and Server 2 is the backup or secondary server. In this case, there is no contention for disk resources because one server at a time reserves the entire disk storage space of the disk array. The disk array controller is labeled “DAC,” and the SCSI host adapters (parallel SCSI, Fibre Channel, etc.) are labeled “SCSI HA.”

A pair of servers is the simplest LifeKeeper configuration. When you plan a cluster consisting of more than two servers, your map is even more critical to ensure that you have the appropriate connections between and among servers. For example, in a multi-directional failover configuration, it is possible to define communications paths within LifeKeeper when the physical connections do not exist. Each server must have a physical communication path to every other server in the cluster in order to provide cascading failover capability.

## 4.2.2. Storage and Adapter Requirements

---

Determine your storage and host adapter requirements using the following guidelines:

**Storage Devices** – Based on your application's data storage requirements, you will need to determine the type and number of data storage devices required by your configuration. Your shared files should reside on a disk array subsystem (Redundant Array of Inexpensive Disks, or RAID). LifeKeeper supports a number of hardware RAID peripherals for use in LifeKeeper configurations. See [Supported Storage](#) for a list of the supported peripherals.

Consider the following issues when planning the configuration of your storage devices:

- LifeKeeper manages resources at the physical disk or Logical Unit (LUN) level, making the resources on each physical disk or LUN available to only one server in the configuration at a time. As a result, it is a good idea to plan disk allocations before you begin to configure LifeKeeper. For example, each hierarchy in active/active configurations must access its own unique LUN, so a minimum of two LUNs is required for a two-node active/active configuration.
- Some model-specific issues and hardware configuration details are listed in [Supported Storage](#).

**Adapters** – Based upon the type of configuration and the number of peripherals, determine the types and number of SCSI or Fibre Channel Host Adapters required. It is important that any adapter you choose be supported by LifeKeeper, as well as by your Linux distribution so that there is a driver available. Refer to Supported Adapter Models for a list of supported host adapters. For reference purposes, you should add the host adapter specifications to your configuration map.

## 4.2.3. Storage and Adapter Options

---

For a list of the disk array storage models currently supported by LifeKeeper in shared storage configurations, see the [Supported Storage](#). Refer to [Storage and Adapter Configuration](#) for details about driver versions and other configuration requirements for these arrays and adapters.

Note that a supported disk array and adapter are not required in LifeKeeper configurations involving non-shared storage with IP failover only or when using SIOS Data Replication or Network Attached Storage.

SIOS Technology Corp. does not specifically certify fibre channel hubs and switches, because there are no known LifeKeeper-specific restrictions or requirements on these devices. Unless otherwise noted for a given array in [Storage and Adapter Configuration](#), LifeKeeper recommends the hubs and switches that the disk array vendor supports.

## 4.3. Setting Up Your LifeKeeper Environment

---

Now that the requirements have been determined and LifeKeeper configuration has been mapped, components of this LifeKeeper environment can be set up.

\* Although it is possible to perform some setup tasks in a different sequence, this list is provided in the recommended sequence.

[Installing the Linux OS and Associated Communications Packages](#)

[Linux Dependencies](#)

[Connecting Servers and Shared Storage](#)

[Configuring Shared Storage](#)

[Verifying Network Configuration](#)

[Creating Switchable IP Address](#)

[Installing and Setting Up Database Applications](#)

[Configuring GUI Users](#)

## 4.3.1. Installing the Linux OS and Associated Communication Packages

---

Before attempting to install the LifeKeeper for Linux software, you must first ensure that your Linux operating system is successfully installed and operational. Please see the Linux installation instructions provided with your distribution of Linux for complete installation details.

### Notes:

- Refer to the [Linux Dependencies](#) topic for further dependencies that may be necessary for the required packages.
- It is possible to install Linux *after* connecting and configuring your shared storage, but it may be simpler to have Linux installed and running before introducing new peripheral devices.
- The LifeKeeper for Linux Installation Image File provides a set of installation scripts designed to perform user-interactive system setup tasks and installation tasks for installing LifeKeeper on your system.

## 4.3.2. Linux Dependencies

Successful completion of the installation of LifeKeeper for Linux requires the installation of a number of prerequisite packages. To prevent script failures, these packages should be installed prior to attempting to run the installation setup script.

The prerequisite packages are broken down into the following three groups:

- [General Package Dependencies](#)
- [Optional Recovery Kit Package Dependencies](#)

Depending on the operating system version and the packages installed based on the operating system type selected (minimal, default, etc.), additional dependent packages may be required.

 **Note:** The dependencies are based on the versions of the OS supported by LifeKeeper. Refer to the [Support Matrix](#) for details.

 **Note:** You may want to consider using a repository-based package manager such as **yum** or **zypper** that is designed to automatically resolve dependencies by searching in predefined software repositories thereby easing the installation of these required packages. To facilitate the installation of dependent packages, the LifeKeeper for Linux installer uses **yum** or **zypper** to install the LifeKeeper for Linux packages. Therefore, it is highly recommended that an OS package repository be setup and configured. This will negate the need to install the OS dependent packages listed below before attempting to install LifeKeeper for Linux.

### rpm Install Example

```
rpm -ivh <package(s)>
```

### yum Install Example

```
yum install <package(s)>
```

### Zypper Install Example

```
zypper install <package(s)>
```

## yum/Zypper Package Lists

The following list of rpm packages, for each distribution listed and installed with the corresponding package installer, is the minimum list of packages that will resolve all the required dependencies for LifeKeeper for Linux:

## Red Hat Enterprise Linux, CentOS and Oracle Linux

```
yum install libXtst libstdc++ bzip2-libs pam zlib patch redhat-lsb
ncurses-libs
```

## SLES

```
zypper install libstdc++ bzip2 pam pam-modules zlib lsb libncurses5
```

## General Package Dependencies

The following packages are always required to successfully install LifeKeeper for Linux. The package architecture version of the installed package should always match the operating system architecture (x86 or x86\_64):

### Red Hat Enterprise Linux, CentOS and Oracle Linux

- bzip2
- iproute
- iputils
- patch (version 2.5 or later)
- redhat-lsb
- ncurses-libs

**Note:** Some or all of these packages may already be installed depending on the selections made during the install of the operating system.

### SLES

- bzip2
- iproute2
- iptables
- iputils
- insserv
- patch (version 2.5 or later)
- lsb-release
- libncurses5
- libXtst6 (SLES15 only)
- libXi6 (SLES15 only)

**Note:** Some or all of these packages may already be installed depending on the selections made during the install of the operating system.

## Dependency with syslog daemon

LifeKeeper logs use the syslog daemon. LifeKeeper supports rsyslog. Before installing LifeKeeper the rsyslog daemon must be installed and activated.

**Note:** In the distributions using systemd such as RHEL7 or SLES12, journald administrates the log collectively. Because LifeKeeper does log output using the syslog daemon, the syslog daemon also must be operating in these environments. Therefore, set up syslog daemon to operate when using LifeKeeper.

**Note:** journald records the log output to a temporary file system (tmpfs) mount on /run/log/journal by default. Thus, the system log is not saved at the time of the OS shutdown. Change the setup to let the journald log perpetuate.

**Note:** To let the journald log perpetuate, set up “Storage=persistent” in /etc/systemd/journald.conf, or, create the /var/log/journal directory with the set up “Storage=auto” (default). After changing the set up, restart systemd-journald.service.

## Optional Recovery Kit Package Dependencies

Additionally, some of the LifeKeeper for Linux optional Application Recovery Kits (ARKs) require supporting packages to be installed.

If NFS exports are to be protected via the LifeKeeper for Linux NFS Application Recovery Kit, then the following dependent packages are required:

- nfs-utils (Red Hat Enterprise Linux, CentOS, Oracle Linux)
- nfs-client (SLES)
- nfs-kernel-server (SLES)

If multipath devices are to be protected via Device Mapper Multipath (DMMP), Hitachi Dynamic Link Manager Software (HDLM), Power Path or NEC iStorage StoragePathSavior (NECSPS), then the following dependent packages are required:

- sg3\_utils (All multipath kits)
- sg3\_utils-libs (All multipath kits)
- HDLM (Hitachi Dynamic Link Manager Software Kit)
- EMCpower.LINUX (Power Path Kit)
- sps (NEC iStorage StoragePathSavior Kit **4.2.0** or prior)
- sps-utils and sps-driver (NEC iStorage StoragePathSavior Kit **4.2.1** or later)

If Websphere MQSeries queue managers are to be protected via the LifeKeeper for Linux Websphere MQ/MQSeries Application Recovery Kit, then the following dependent Websphere MQ packages are required:

- MQSeriesServer
- MQSeriesSamples
- MQSeriesClient
- MQSeriesRuntime
- MQSeriesSDK

If Software RAID devices are to be protected via the LifeKeeper for Linux Software RAID (md) Recovery Kit, then the following dependent package is required:

- mdadm

## 4.3.3. Connecting Servers and Shared Storage

---

If you are planning to use LifeKeeper in a non-shared storage environment, then you may skip this information. If you are using LifeKeeper in a data replication (mirroring) environment, see the [DataKeeper](#) section of this documentation. If you are using LifeKeeper in a network attached storage environment, see [LifeKeeper Network Attached Storage Recovery Kit Administration Guide](#).

Once Linux is installed, you should set the host adapter and shared peripheral addressing. Refer to the documentation accompanying your adapter and storage device for specific details.

## 4.3.4. Configuring Shared Storage

---

LifeKeeper configurations may use the facilities of shared Small Computer System Interface (SCSI) host adapters and shared disk hardware to switch resources from a failed server to a designated backup server. A Fibre Channel Storage Area Network (SAN) may also be used to switch resources from a failed server to a designated backup server.

Perform the following tasks before creating disk-based application resource hierarchies that enable LifeKeeper to provide failover protection.

1. Partition disks and LUNs. Because all disks placed under LifeKeeper protection must be partitioned, your shared disk arrays must now be configured into logical units, or LUNs. Use your disk array management software to perform this configuration. You should refer to your disk array software documentation for detailed instructions.

**Note:** Remember that LifeKeeper locks **its disks** at the LUN level. Therefore, one LUN may be adequate in an Active/Standby configuration. But, if you are using an Active/Active configuration, then you must configure at least two separate LUNs, so that each hierarchy can access its **own unique LUN**.

2. Verify that both servers recognize the shared disks (for example, using the **gdisk** command). If Linux does not recognize the LUNs you have created, then LifeKeeper will not either.
3. Create file systems on your shared disks from the system you plan to use as the primary server in your LifeKeeper hierarchy. Refer to the Linux documentation for complete instructions on the administration of file systems.

## 4.3.5. Verifying Network Configuration

---

It is important to ensure that your network is configured and working properly before you install LifeKeeper. There are several tasks you should do at this point to verify your network operation:

1. If your server installation has a firewall enabled, you will either need to accommodate the LifeKeeper ports or disable the firewall. Please refer to [Running LifeKeeper With a Firewall](#).
2. From each server, ping the local server, and ping the other server(s) in the cluster. If the ping fails, then do the necessary troubleshooting and perform corrective actions before continuing.
3. If your server has more than one network adapter, you should configure the adapters to be on different subnets. If the adapters are on the same subnet, TCP/IP cannot effectively utilize the second adapter.
4. Ensure that *localhost* is resolvable by each server in the cluster. If DNS is not implemented, edit the */etc/hosts* file and add an entry for the *localhost* name. This entry can list either the IP address for the local server, or it can list the default entry (127.0.0.1). If *localhost* is not resolvable, the LifeKeeper GUI may not work.
5. If DNS is implemented, verify the configuration to ensure the servers in your LifeKeeper cluster can be resolved using DNS.
6. Ensure each server's hostname is correct and will not change after LifeKeeper is installed. If you later decide to change the hostname of a LifeKeeper system, you should follow these steps *on all servers in the cluster*.

- a. Stop LifeKeeper on all servers in the cluster using the command:

```
/opt/LifeKeeper/bin/lkstop -f
```

- b. Change the server's hostname using the Linux **hostname** command.
- c. Before continuing, you should ensure that the new hostname is resolvable by each server in the cluster (see the previous bullets).
- d. Run the following command on every server in the cluster to update LifeKeeper's hostname. (Refer to *lk\_chg\_value(1M)* for details.)

```
/opt/LifeKeeper/bin/lk_chg_value -o oldhostname -n newhostname
```

- e. Start LifeKeeper using the command:

```
/opt/LifeKeeper/bin/lkstart
```

LifeKeeper for Linux v7.x supports VLAN interface for Communication Paths and IP resources. The type of VLAN interface can be chosen as described below.

## VLAN Interface Support Matrix

- not supported \ x supported

### LK Linux v7.1 or Prior Version

VLAN_NAME_TYPE	CommPath	IP Resource
DEV_PLUS_VID (eth0.0100)	-	x
DEV_PLUS_VID_NO_PAD (eth0.100)	-	x
VLAN_PLUS_VID (vlan0100)	x	x
VLAN_PLUS_VID_NO_PAD (vlan100)	x	x

### LK Linux v7.2 or Later Version

VLAN_NAME_TYPE	CommPath	IP Resource
DEV_PLUS_VID (eth0.0100)	x	x
DEV_PLUS_VID_NO_PAD (eth0.100)	x	x
VLAN_PLUS_VID (vlan0100)	x	x
VLAN_PLUS_VID_NO_PAD (vlan100)	x	x
<b>Note:</b> The NIC name can be anything (i.e. eth0, ens192...)		

## 4.3.6. Creating a Switchable IP Address

---

A switchable IP address is a “virtual” IP address that can be switched between servers. It is separate from the IP address associated with the network interface card of each server. Applications under LifeKeeper protection are associated with the switchable IP address. Then, if there is a failure on the primary server, that IP address “switches” to the backup server.

If you plan to configure resource hierarchies for switchable IP addresses, you must do the following on each server in the cluster:

- Verify that the computer name is correct and will not be changed.
- Verify that the switchable IP addresses are unique using the ping command.
- Edit the `/etc/hosts` file to add an entry for each switchable IP address.

Refer to the [LifeKeeper for Linux IP Recovery Kit Technical Documentation](#) for additional information.

## 4.3.7. Installing and Setting Up Database Applications

---

If your environment includes a protected database application such as Oracle, DB2 or MySQL, you should install the application using the documentation provided with the database. Ensure that the database is on a shared file system and that the configuration files are on a shared file system. The executables may either be on each local or a shared file system.

Although it is possible to install your application *after* LifeKeeper is installed, you should test the application to ensure it is configured and operating properly before placing it under LifeKeeper protection. Please reference the specific [LifeKeeper database recovery kit documentation](#) for additional installation and setup considerations.

## 4.3.8. Configuring GUI Users

### GUI Authentication with PAM

LifeKeeper for Linux now leverages the Pluggable Authentication Module (PAM) provided in the Linux Standard Base (LSB). LifeKeeper no longer uses its private password file once located in `/opt/LifeKeeper/website/passwd`. Instead, users are identified and authenticated against the system's PAM configuration. Privilege levels are determined from group membership as provided through PAM.

In order to access the GUI, a user must be a member in one of the three LifeKeeper groups: `lkadmin`, `lkoper` or `lkguest`. Membership in these groups should be set by the system administrator using whatever technique is appropriate for the type of user account database that is being used throughout the cluster.

These three LifeKeeper groups provide three different sets of permissions (see [Permissions Table](#)).

1. Users with **Administrator** permission (`lkadmin`) throughout a cluster can perform all possible actions through the GUI.
2. Users with **Operator** permission (`lkoper`) on a server can view LifeKeeper configuration and status information and can bring resources into service and take them out of service on that server.
3. Users with **Guest** permission (`lkguest`) on a server can view LifeKeeper configuration and status information on that server.

During installation of the GUI package, the *root user* on the system is automatically added to the `lkadmin` group in the system's local group database allowing *root* to perform all LifeKeeper tasks on that server via the GUI application or web client. If you plan to allow users other than *root* to use LifeKeeper GUI clients, then these LifeKeeper GUI users will need to be configured by adding them to the appropriate group.

If PAM is configured to use a non-local database such as NIS, LDAP or AD, then the system administrator must ensure that the accounts are correctly configured in those databases. The groups listed above must exist and users who are allowed to log into the LifeKeeper GUI must be a member of one of these groups. These groups should be created in the remote database only and they should be removed from the local `/etc/group` file.

If any system in the cluster is using an LK GUI password other than the system's 'root' password, the LK GUI login will fail. Once the root passwords are the same on each system in the cluster, the LK GUI login for 'root' will succeed.



**Note:** To avoid confusion and maintain consistency if leveraging more complex PAM configurations such as LDAP, NIS or AD, it is recommended that all user and LifeKeeper group accounts exist prior to installing or upgrading LifeKeeper.

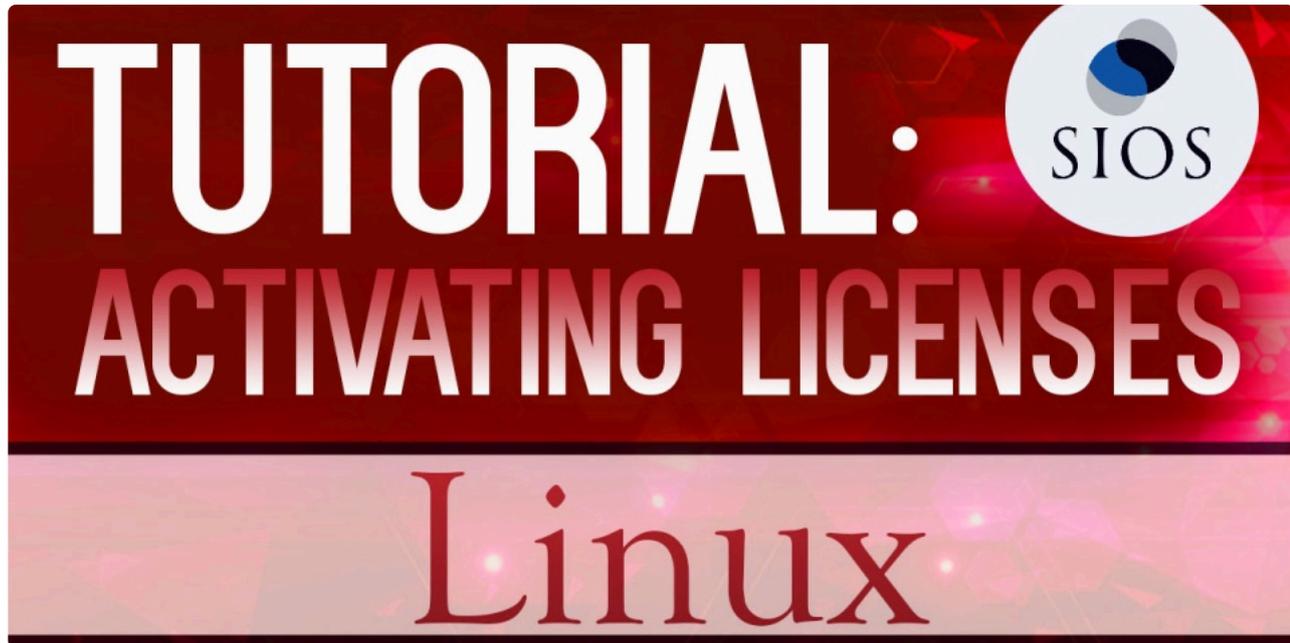
The best practice is to always grant permissions on a cluster-wide basis. It is possible to grant permissions on a single-server basis, but that is confusing to users and makes it impossible to perform administrative tasks.

## 4.3.9. Licensing

---

### Obtaining and Installing the License

LifeKeeper for Linux requires a unique license for each server. The license is a run-time license, which means that you can install LifeKeeper without it, but the license must be installed before you can successfully start and run the product.



<https://fast.wistia.net/embed/iframe/76bo8epo9g>



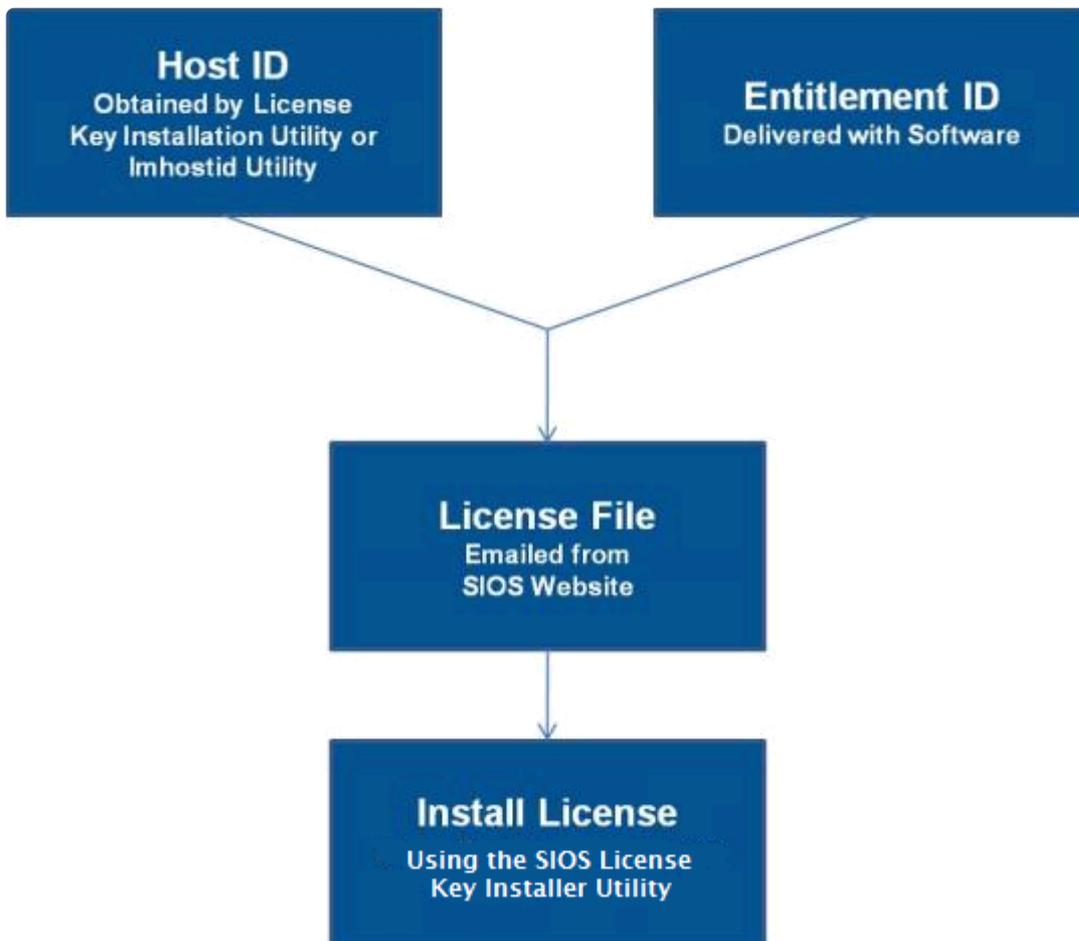
<https://fast.wistia.net/embed/iframe/mtl4wubwfk>

The Installation script installs the Licensing Utilities package which obtains and displays all of the

available Host IDs for your server during the initial install of your LifeKeeper Software. Once your licenses have been installed the utility will return the Entitlement ID if it is available or the Host IDs if it is not.

✿ **Note:** Host IDs, if displayed will always be based on the MAC address of the NICs.

Starting with v8.2.0 any new licenses obtained from the [SIOS Technology Corp. Licensing Operations Portal](#) will contain your Entitlement ID and will be locked to a specific node or IP address in the cluster. The Entitlement ID (Authorization Code) which was provided with your LifeKeeper Software, is used to obtain the permanent license required to run the LifeKeeper Software. The process is illustrated below.



✿ **Note:** Each software package requires a license for each server.

Perform the following steps to obtain and install your licenses for each server:

1. Get your **Host ID**. At the end of the LifeKeeper installation, make note of the **Host ID** displayed by the **License Key Installer** utility.
2. Write the **Host IDs** in a notebook or save them in a file. If saved in a file, copy that file to a system with internet access. Otherwise, take your notebook with you to the system with internet access.

3. Ensure you have your LifeKeeper for Linux **Entitlement ID** (Authorization Code). You should have received an email with your software containing the Entitlement ID needed to obtain the license.
4. Obtain your licenses from the [SIOS Technology Corp. Licensing Operations Portal](#).
  - a. Using the system that has internet access, navigate to the [SIOS Technology Corp. Licensing Operations Portal](#) and log in entering your **User Name** and **Password** (or register if you do not already have an account).

 **Note:** New users must enter the Entitlement ID that is included in the delivery email..

- b. From the **Activation and Entitlements** dropdown select **List Entitlements**.
- c. Check the box to the left of the product line item(s) that you wish to license.
- d. From the **Action** dropdown select **Activate** and enter the requested information (including your system HOSTNAME) then select **Next**.
- e. Click on the **Gray Plus Sign** to choose an already defined host or create a new host by selecting the **Green Plus Sign**.
- f. Select **ANY** for the Node Locked Host choice if it is available, otherwise select **ETHERNET MAC ADDRESS** and enter the Host ID (MAC address), click **OK** then click **Generate**.

 **Note:** The Host ID is 12 characters with no spaces, no colons, no dashes, and no separators.

- g. Check the box to the left of the **Fulfillment ID** and select **Complete**.
  - h. From the **License Support** dropdown select **List Licenses**. Check the box to the left of the **Fulfillment ID** and select **Email** from the **View** dropdown.
  - i. Enter a valid email address to send the license to and select **Send**.
  - j. Retrieve the email(s).
  - k. Copy the file(s) to the appropriate system(s).
5. Install your license(s).
    - On each system, copy the license file(s) to /var/LifeKeeper/license. Run /opt/LifeKeeper/bin/lkkeyins and specify the filename (including full path) to the file.
  6. Repeat on all additional servers. You must install a license on the other LifeKeeper server(s) using the unique Host ID for each server.

7. Restart LifeKeeper for Linux.

## Primary Network Interface Change May Require a License Rehost

The Host ID used by the License Key Installer utility is obtained from the LifeKeeper for Linux server's primary network interface card (NIC). LifeKeeper for Linux will check for a valid license each time it starts. If your LifeKeeper for Linux server should require a NIC replacement in the future that would cause the Host ID to change, then the next time LifeKeeper for Linux is stopped, a License Rehost must be performed before starting it again. Log in to the [SIOS Technology Corp. Licensing Operations Portal](#) and select **License Support, List Licenses, Action, Rehost**.

 **Note:** A rehost can be performed four times per six-month period (per Activation ID) without contacting support.

## 4.3.9.1. Obtaining an Internet HOST ID

Use `lmutil` to obtain your machine's Internet Host ID. The Internet Host ID is normally the primary IP address of the primary network interface in the system. Internet Host IDs can be used as an alternative to Ethernet (or MAC) Host IDs and may be preferable in virtual environments where MAC addresses can change due to VM cloning.

1. Type the following command:

```
# /opt/LifeKeeper/bin/lmutil lmhostid -internet -n
```

2. Record the ID returned by the program.

### Example:

```
# /opt/LifeKeeper/bin/lmutil lmhostid -internet -n
```

```
"INTERNET=172.17.100.161"
```



**Note:** This info must match the information contained in the permanent license key obtained from SIOS Technology Corp.

## 4.4. Installing the Software

---

This document will guide you through the installation of the LifeKeeper for Linux and assumes the user has basic knowledge of the Linux operating system. Please refer to the [LifeKeeper for Linux product documentation](#) for more information.

### Pre-Installation Requirements

Before installing LifeKeeper for Linux, please check the following:

- [LifeKeeper for Linux Release Notes](#) -The Release Notes include supported platforms, operating systems, applications, and storage. They also include the latest features and Bug Fixes.
- **TCP/IP Connection and Name Resolution** – In order to use the GUI function, both cluster nodes need to be able to resolve the name. Use the DNS service or */etc/hosts* for name resolution. Also, localhost needs to be resolved to 127.0.0.1.
- **Firewall** – The following ports are used:
  - Communication Path (TCP): 7365/tcp
  - Communication of a GUI Server: 81/tcp、 82/tcp
  - RMI Communication between the GUI Server and Client: all the ports after 1024/tcp
  - Synchronization of DataKeeper (when using DataKeeper): “10001+<mirror number>+<256 \* i>”

#### More Firewall Information

- The port used for communication with the GUI server and a client needs to be open on the cluster node where LifeKeeper is installed and on all systems where the GUI client runs.
- The ports used by DataKeeper can be calculated using the formula above. The value of i starts at 0 and uses an unused port when found. For example, in an environment where a DataKeeper resource with mirror number 0 exists, if port 10001 is being used by another application, port 10257 will be used.
- For communication between the GUI server and a client, Java RMI (Remote Method Invocation) randomly uses ports 1024 and above. When applying access control etc. to a cluster system, packet filtering needs to be performed considering these ports. If this specification is an issue from a security standpoint, you can use ssh X forwarding. Please refer to the [Technical Documentation](#) for the setting details.
- Add the following to the port numbers you are using: *WebGUI server process and policy setting with the `lkpolicy` command : 778(SSL) /tcp*
- **Check the SELinux Setting** – When the SELinux setting is enabled, LifeKeeper for Linux may not be able to be installed depending on the mode.
  - enforcing mode – LifeKeeper for Linux cannot be installed
  - permissive mode – LifeKeeper for Linux can be installed (not recommended except in some ARK environments)
    - It is not recommended to use SELinux permissive mode unless it is required in an SAP environment. Please make sure that the application to be run on the cluster supports permissive mode. SELinux permissive mode has been tested for following ARKs: SAP / SAP MaxDB / Sybase / Oracle / DB2 / NFS / DataKeeper / NAS / EC2 / IP / FileSystem / MQ. Refer to [Linux Dependencies](#) for required packages.

- Install the appropriate package provided by your distribution.
  - disabled mode – LifeKeeper for Linux can be installed
    - Please refer to the OS distribution documentation on how to disable SELinux.
  - The sg3\_utils package is required for environments using recovery kits for Multipath such as the DMMP Recovery Kit and the PowerPath Recovery Kit. This is not required for environments where recovery kits for Multipath are not used.
- **Check [Known Issues](#)** – Please make sure that there are no known issues for your environment.

## Installing LifeKeeper for Linux

Install the LifeKeeper software on each server in the LifeKeeper configuration.

Packages that LifeKeeper is dependent on are installed automatically because the LifeKeeper installation setup script uses package manager tools (yum or zypper) to ensure installation of all dependent packages.

**!** **IMPORTANT:** A functional yum or zypper configuration is required for the successful installation of LifeKeeper. A non-functional configuration can result in an installation failure (see [Installation Known Issues](#) for more information). Additionally, the package manager repo or rpm database must not be locked as that could cause the install to hang. If the dependent packages cannot be installed automatically via the package manager, refer to [Linux Dependencies](#) and install the necessary packages in advance.

The LifeKeeper for Linux image file (sps.img) provides a set of installation scripts designed to perform the user interactive system setup tasks that are necessary when installing LifeKeeper on your system (see [Interactive Mode](#) for more information). A non-user interactive install can be performed as well (see [Non-interactive Mode](#) for more information).

A licensing utilities package is also installed providing utilities for obtaining and displaying the Host ID or Entitlement ID of your server. Host IDs and/or Entitlement IDs are used to obtain valid licenses for running LifeKeeper. Refer to [Licensing](#) for information on how to obtain and install your licenses.

**\* Note:** These installation instructions assume that you are familiar with the Linux operating system installed on your servers.

**!** **IMPORTANT:**

- Installing LifeKeeper on your shared storage is not supported. Each server should have its own copy installed on its local disk.
- All LifeKeeper packages are installed in the directory /opt/LifeKeeper.

LifeKeeper will be installed through the command line regardless of the Linux distribution you are operating under.

- Please refer to [How to Use Setup Scripts](#) for the installation activities.
- For upgrade installations, see [Upgrading LifeKeeper](#).

## 4.5. How to Use Setup Scripts

---

To install or upgrade LifeKeeper, follow the steps below.

### How the Setup Scripts Works

1. Interactive installation

Configure and install LifeKeeper from the menu.

If you save the configuration information at this time, it can be used for the non-interactive installation described below.

2. Non-interactive installation

Install LifeKeeper using the saved configuration information.

Since no inquiry to the user occurs, you can perform this using a building tool (e.g. Ansible).

### How in Install / Upgrade LifeKeeper Using the Setup Script

#### Interactive Installation

1. After logging in as the root user, use the following command to mount the sps.img file:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

Where PATH is the path to the image

IMAGE\_NAME is the name of the image

MOUNT\_POINT is the path to mount location

2. Change to the directory where sps.img is mounted and enter the following:

```
./setup [-s <response_file>]
```

When the `-s` option is specified, you can save your configuration information in the `response_file`, which is used for a non-interactive installation.

3. The script collects information about the system environment and determines what you need to do to install LifeKeeper.

If the system requirements for installation or upgrade are not satisfied, then an error message is displayed and the installation / upgrade is cancelled.

Also, if some restrictions arise or a configuration change is required, a warning message will be displayed requiring the user to decide whether to continue or abort the installation.

4. Select the LifeKeeper features and Application Recovery Kits (ARKs) to install via the main dialog screen.

Please refer to [How to Use the Dialog Screen](#).

5. Once all the required LifeKeeper features and ARKs have been selected, select <Done> to begin the installation.

If any notifications are output when the installation completes, please take the necessary actions to correct them.

## Non-interactive Installation

1. After logging in as root user, use the following command to mount the sps.img file:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

PATH is the path to the image

IMAGE\_NAME is the name of the image

MOUNT\_POINT is the path to mount location

2. After copying the configuration file to the system where you want to install LifeKeeper, run the following command:

```
./setup -f <response_file> -q y
```

The “-q y” option gives the answer that the warning has been noted.

 **'sh setup'** (bourne shell) cannot be used. Use **bash (** instead.

 If you use a configuration file saved with the `-s` option for a non-interactive installation, the system on which the file is used must be configured the same way as the system on which the file was generated. If the systems have too many differences the non-interactive installation may fail. The configuration file created with the `create_response_file` script has no such restrictions.

## Creating the Configuration Information

The configuration information file used for non-interactive installation can be created during setup with `setup -s <response_file>` or created in advance with the `create_response_file` script.

1. After logging in as root user, run the following command to mount the lkssp.img file:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

PATH is the path to the image

IMAGE\_NAME is the name of the image

MOUNT\_POINT is the path to mount location

2. Change to the directory where sps.img is mounted and enter the following:

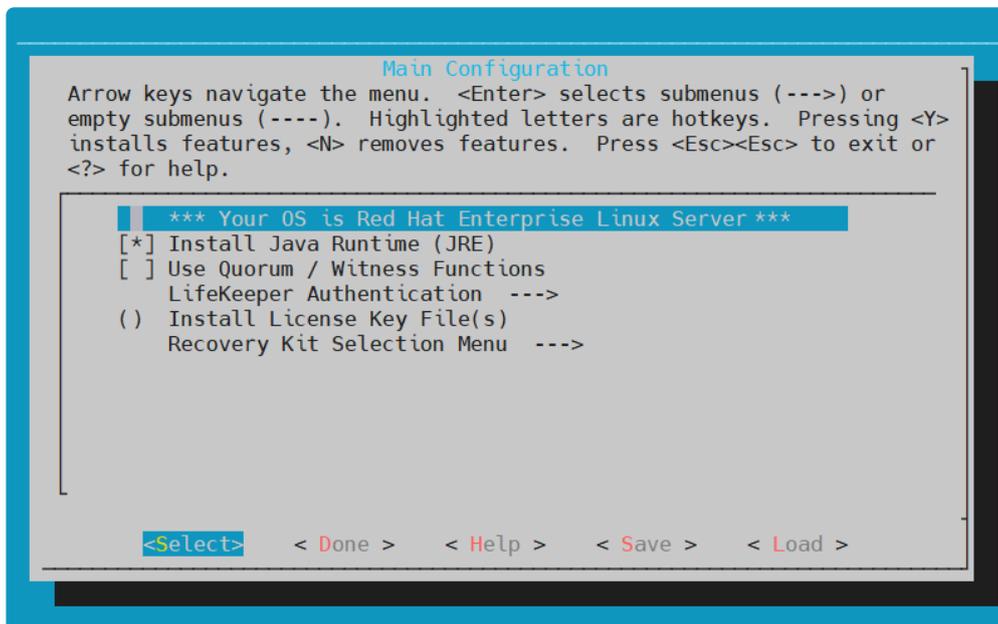
```
./create_response_file <response_file>
```

3. Select the LifeKeeper features and Application Recovery Kits (ARKs) to install via the main dialog screen. Please refer to [How to Use the Dialog Screen](#).
4. Select the LifeKeeper features and Application Recovery Kits (ARKs) to install, select **Done** to save the configuration to response\_file and exit the script. The response\_file is copied to the destination system.

Repeat steps 2 through 4 to change the saved configuration information.

## How to Use the Dialog Screen

The dialog screen is displayed below.



Use the following keys to navigate the menu.

↑ ↓	Navigate between menu items
← →	Navigate between the menu buttons at the bottom of the screen
ENTER	Open the selected sub menu
Y / N / ENTER / SPACE	Turn on, turn off or invert the selected item

The menu buttons at the bottom of the screen are used for the following operations.

Select	Opens a sub menu dialog screen
Done	Closes the current screen and returns to the previous screen. Selecting this button on the main screen completes the configuration.
Help	Displays help for the highlighted item

Save	Saves the current settings in a configuration file. The saved configuration file can be used for non-interactive installations. The full path to the file where the configuration information is to be saved should be specified. (Disabled for create_response_file)
Load	Loads a saved configuration file (Disabled for create_response_file)

The “Save” and “Load” menu buttons display a dialog screen asking for a configuration file name for use in saving the current configuration or for loading a saved configuration. If you want to change the default file name provided, move to the file name field using the [TAB] key, and enter a new name. **Note:** The “Save” operation will prompt for confirmation before overwriting a file with the same name.

The items listed below are configurable during installation. During an upgrade only items that can be configured are listed. Using the hotkey <Z> will show those items that will remain unchanged during the upgrade.

- **Install Java Runtime (JRE)**

Install the Java runtime environment used by the LifeKeeper GUI.

- **Restart NFS Service**

When configuring High Availability NFS, restarting the NFS services is required. When this is selected, the services are restarted automatically after the configuration is completed.

**Note:** If you do not want to restart the NFS services automatically, a restart will need to be done to pick up the configuration changes before using the NFS Recovery Kit.

- **Use Quorum / Witness Functions**

Use Quorum / Witness for I/O fencing. For details, please refer to [Quorum/Witness](#) in the technical documentation.

- **LifeKeeper Authentication**

Specify the users allowed to log in to the LifeKeeper for Linux GUI along with their privilege levels. Multiple user accounts can be specified by separating them with blanks. For details, refer to [GUI User Settings](#).

- **Install License Key File(s)**

Install the licenses required to start LifeKeeper for Linux by entering the path name of the license file to install. Multiple files can be specified by separating them with spaces.

Please refer to [Licensing](#) for details.

- **Recovery Kit Selection**

Select the Application Recovery Kits to install.

Application Recovery Kits are broken into several categories based on common functionality.

Please refer to [Categories for Application Recovery Kits](#) for details.

- **LifeKeeper Startup After Install**

When selected, LifeKeeper for Linux will be started when the installation is completed.

 **Note:** Because the LifeKeeper for Linux Data Replication package may install kernel modules for some of the supported OS distributions, a reinstall of LifeKeeper for Linux may be required when the kernel is upgraded. This applies to OS distributions for RedHat, CentOS and Oracle Enterprise Linux (including UEK kernels) running kernel versions 3.10.0-514 or later, and for SUSE Linux Enterprise Server running kernel version 4.12.14-95 or later.

## Adding / Removing Application Recovery Kits

To add Application Recovery Kits after completing an installation, simply execute `setup`, select the Recovery Kit in the Recovery Kit Selection, followed by the Application Recovery Kit Category and then select the desired kit. If you deselect an Application Recovery Kit which is no longer necessary, that kit will be removed. However, since the kit cannot be removed for the resources in use, delete the resources in advance.

## Repair Installation

To repair a LifeKeeper for Linux installation run `setup` with the “—force” option. A repair installation will update the installation replacing any lost or corrupted files.

## Setup Script Options

The setup script can be executed with the following options:

- `-f <response_file>`

Install non-interactively. `<response_file>` contains the configuration information to use during the installation.

- `-s <response_file>`

Save a configuration file containing your menu selections. This file can then be used with the “-f” option to install the same LifeKeeper configuration to another system. For example, run:

```
setup -s <response_file>
```

Select the necessary packages and options and complete setup.

Then run:

```
setup -f <response_file> -q y
```

to run a silent installation of LifeKeeper (on another system) with the same options that were selected the first time setup was run.

- `-force`

Forcibly reinstall LifeKeeper for Linux.

- `-q <y/n>`

Specifies the response to any confirmation questions that may arise during non-interactive installation.

## Categories for Application Recovery Kits

Category	Description
Application Suite	A group of recovery kits that protect applications such as SAP and IBM MQ.
Networking	A group of recovery kits that protect network services in the cloud such as EC2 and Route53.
Database	A group of recovery kits that protect database applications, including, but not limited to, Oracle, PostgreSQL, and MaxDB.
File Sharing	A group of recovery kits that protect file sharing services such as NFS and Samba.
Mail Server	A group of recovery kits that protect email services such as Postfix.
Storage	A group of recovery kits that protect data storage methods, including, but not limited to, DataKeeper (replication), Device Mapper (DM) Multipath (DMMP), and Network Attached Storage (NAS).
Web Server	A group of recovery kits that protect web services such as Apache.

## 4.6. Verifying the LifeKeeper Installation

---

You can verify that the LifeKeeper packages were installed correctly by entering the following at the command line:

```
rpm -V <package name>
```

 **Note:** If the package is installed correctly, no output will be displayed by this command.

To perform a query from the command line, type

```
rpm -qi <package name>
```

 **Note:** The expected output for this command is the package information.

## 4.7. Upgrading LifeKeeper

- \* Beginning in v9.5.0 SIOS has released the new SAP HANA Application Recovery Kit. SIOS will continue to support the SAP HANA gen/app based Recovery Kit with the 9.4.x releases until **March 31, 2022**.
  - The existing SAP HANA gen/app based Recovery Kit is **not** supported with v9.5.0.
  - If you are using LifeKeeper for Linux v9.5 or later you must use the new (built-in) SAP HANA Application Recovery Kit.
  - Users who wish to upgrade to the LifeKeeper for Linux v9.5.0 **must** convert their existing SAP HANA gen/app based Recovery Kit to the new [SAP HANA Recovery Kit](#). Refer to [Upgrading from the SAP HANA Gen/App to the SAP HANA Recovery Kit](#) for details.

LifeKeeper for Linux may be upgraded to future releases while preserving existing resource hierarchies. Review this information carefully to ensure that you minimize application downtime.

- \* **Note:** LifeKeeper can be upgraded to the current version from up to **two versions back**. (An example of a two version upgrade is from **9.1.x -> 9.3.x**) If upgrading from a version previous to that, the older version will need to be uninstalled, and LifeKeeper for Linux will have to be reinstalled. An alternative to uninstalling the older version would be to upgrade from the older version to one of the two acceptable versions, then perform the upgrade to the current version.

- \* **Note:** If using [lkbackup](#) during your upgrade, refer to [the known issues of lkbackup](#) for further information.

### Upgrading LifeKeeper

1. While upgrading LifeKeeper in a cluster, switch all applications away from the server to be upgraded now. Do this manually or by setting the LifeKeeper shutdown strategy to “**Switchover**” which causes the applications to be switched when LifeKeeper is stopped or the server is shut down. Refer to the [Setting Server Shutdown Strategy](#) for more information.
2. Upgrade LifeKeeper by referring to [How to Use Setup Scripts](#).
3. Switch all applications back to the upgraded server.
4. Repeat this procedure for each server in the LifeKeeper cluster to be upgraded. For clusters containing a dedicated Witness/Quorum node (a node with no resource instances) no switching of applications is required prior to upgrading LifeKeeper.

**!** **CAUTION:** The same version and release of LifeKeeper must be installed on all systems in a cluster. In general, different versions and/or releases of LifeKeeper are not

compatible. For situations other than rolling upgrades, LifeKeeper should not be started when a different version or release is resident and running on another system in the cluster.

## Upgrading the OS / Kernel on a node with LifeKeeper

### NOTES:

When upgrading the OS, make sure the currently installed version of LifeKeeper supports the upgraded version of the OS. If it is not supported, LifeKeeper will need to be upgraded as well provided a version of LifeKeeper has been released that supports the new OS version. If no version of LifeKeeper has been released that supports the new OS version you may not be able to upgrade the OS. Refer to the [Supported Operating Systems](#).

Before upgrading the OS, it is recommended that the LifeKeeper configuration be backed up via the [lkbackup command](#).

 **Note:** When using [lkbackup](#), refer to [the known issues of lkbackup](#).

1. When upgrading the cluster, all the resource hierarchies and thus the applications they protect, must be switched from the server to be upgraded to a standby node in the cluster. This can be done manually, or, by setting the LifeKeeper Shutdown Strategy to “Switchover”. By setting the Shutdown Strategy to “Switchover”, the resource hierarchies are switched over to a standby node when LifeKeeper stops or the servers are shut down.
2. Stop LifeKeeper.
3. Upgrade the OS / Kernel (*A message will indicate whether a **reboot** is necessary*)
  - a. **Note:** After an upgrade, [see scenarios](#) below for when a reboot is **required**.
4. Upgrade LifeKeeper if required to support the new OS / Kernel. If you do not upgrade LifeKeeper, you must still run the LifeKeeper setup again to update the settings corresponding to the new OS / Kernel
5. Start up LifeKeeper.
6. Switch all the resource hierarchies to the upgraded server.
7. Execute these steps for all the nodes in the LifeKeeper cluster. For clusters containing a dedicated Witness/Quorum node (a node with no resource instances) no switching of applications is required prior to updating LifeKeeper.
8. Restart the LifeKeeper GUI (via `/opt/LifeKeeper/bin/lkGUIapp`) if it was open during the upgrade.

 **Note:** All nodes in the cluster must be running the same version of the OS and the same version of LifeKeeper to be considered supported. Only during the upgrade process can the nodes differ in the OS and LifeKeeper versions as this would be considered a temporary condition.

### Scenarios Requiring a Reboot After Installing/Upgrading

1. The system **must be rebooted** if the LifeKeeper setup script is unable to reload any required kernel modules. In this case, the setup script will display a warning message resembling either of the following:
  - `“Unable to reload modules after adding LifeKeeper for Linux specific configuration information in /etc/modprobe.d. Please reboot your system after setup completes to ensure these modules load correctly and allow LifeKeeper for Linux to function properly.”`
  - `“Updated some modules for DataKeeper. Reboot the system to use the new module.”`
2. The system **must be rebooted** after disabling secure boot. If secure boot is enabled when DataKeeper is installed, the setup script will display a warning message such as:
  - `“Secure Boot cannot be enabled in a DataKeeper environment. Please take one of the following actions: a) Disable Secure Boot [recommended] or b) Disable signature verification (mokutil --disable-validation)”`.

## New or Deprecated Mount Options After a Kernel Upgrade

When upgrading the Linux kernel, it is possible that some existing file system mount options may be deprecated in the new kernel or that the new kernel may add new default mount options to existing mounts. For example, the “nobarrier” mount option was deprecated in RedHat Enterprise Linux 8, and some kernel versions have added new default mount options such as “logbufs=8” and “logsize=32k”.

If a LifeKeeper-protected file system resource contains mount options which become deprecated after a kernel upgrade, the deprecated options should be removed from the list of mount options for the LifeKeeper resource on every server in the cluster. See the [Modifying Mount Options for a LifeKeeper File System Resource](#) section for more details.

If new default mount options are added by the kernel to an existing LifeKeeper-protected mount point after a kernel upgrade, then the new options should be added to the list of mount options for the LifeKeeper resource on every server in the cluster. See the [Modifying Mount Options for a LifeKeeper File System Resource](#) section for more details.

# 5. LifeKeeper for Linux Technical Documentation

---

LifeKeeper for Linux integrates high availability clustering with innovative data replication functionality in a single, enterprise-class solution.

## LifeKeeper for Linux Integrated Components

**SIOS LifeKeeper** provides a complete fault-resilient software solution to provide high availability for your servers' file systems, applications, and processes. LifeKeeper does not require any customized, fault-tolerant hardware. LifeKeeper simply requires two or more systems to be grouped in a network, and site-specific configuration data is then created to provide automatic fault detection and recovery.

In the case of a failure, LifeKeeper migrates protected resources from the failed server to a designated back-up server. Users experience a brief interruption during the actual switchover; however, LifeKeeper restores operations on the back-up server without operator intervention.

**SIOS DataKeeper** provides an integrated data mirroring capability for LifeKeeper environments. This feature enables LifeKeeper resources to operate in shared and non-shared storage environments.

---

[Documentation and Training](#)

# 5.1. Introduction

---

## About LifeKeeper for Linux

LifeKeeper for Linux integrates high availability clustering with innovative data replication functionality in a single, enterprise-class solution.

### LifeKeeper for Linux Integrated Components

**SIOS LifeKeeper** provides a complete fault-resilient software solution to provide high availability for your servers' file systems, applications, and processes. LifeKeeper does not require any customized, fault-tolerant hardware. LifeKeeper simply requires two or more systems to be grouped in a network, and site-specific configuration data is then created to provide automatic fault detection and recovery.

In the case of a failure, LifeKeeper migrates protected resources from the failed server to a designated back-up server. Users experience a brief interruption during the actual switchover; however, LifeKeeper restores operations on the back-up server without operator intervention.

**SIOS DataKeeper** provides an integrated data mirroring capability for LifeKeeper environments. This feature enables LifeKeeper resources to operate in shared and non-shared storage environments.

[Documentation and Training](#)

## 5.2. Documentation and Training

### Documentation

A complete reference providing instructions for installing, configuring, administering and troubleshooting LifeKeeper for Linux is available in the [LifeKeeper for Linux Technical Documentation](#). The following sections cover every aspect of LifeKeeper for Linux:

Section	Description
<a href="#">Introduction</a>	Provides an introduction to the LifeKeeper for Linux product, including software packaging and configuration concepts.
<a href="#">LifeKeeper for Linux Installation Guide</a>	Provides useful information for planning and setting up your LifeKeeper environment, installing and licensing LifeKeeper and configuring the LifeKeeper graphical user interface (GUI).
<a href="#">Configuration</a>	Contains detailed information and instructions for configuring the LifeKeeper software on each server in your cluster.
<a href="#">Administration</a>	Discusses server-level tasks such as editing server properties and creating resources and resource-level tasks such as editing, extending or deleting resources.
<a href="#">User's Guide</a>	Contains detailed information on the <a href="#">LifeKeeper GUI</a> , including the many tasks that can be performed within the LifeKeeper GUI. Also includes a <a href="#">Technical Notes</a> section along with many more <a href="#">Advanced Tasks</a> .
<a href="#">DataKeeper</a>	Contains planning and installation instructions as well as administration, configuration and user information for SIOS DataKeeper for Linux.
<a href="#">Troubleshooting</a>	Describes known issues and restrictions and suggests solutions to problems that may be encountered during installation, configuration and/or use of SIOS LifeKeeper for Linux.
<a href="#">Recovery Kits</a>	Contains planning and installation instructions as well as administration, configuration and user information for the Optional Recovery Kits that allow LifeKeeper to manage and control specific applications.
<a href="#">Error Code Search</a>	Provides a listing of all messages that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received.

### Training

LifeKeeper training is available through SIOS Technology Corp. or through your reseller. Contact your sales representative for more information.

## Technical Support

As a SIOS Technology Corp. customer with a valid Support contract, you are entitled to access the [SIOS Technology Corp. Support Self-Service Portal](#).

The [SIOS Technology Corp. Support Self-Service Portal](#) offers you the following capabilities:

- Search our **Solution Knowledge Base** to find solutions to problems and answers to questions
- Always on 24/7 service with the SIOS Technology Corp. Support team to:
  - **Log a Case** to report new incidents.
  - **View Cases** to see all of your open and closed incidents.
  - **Review Top Solutions** providing information on the most popular problem resolutions being viewed by our customers.

Contact SIOS Technology Corp. Support at [support@us.sios.com](mailto:support@us.sios.com) to set up and activate your Self-Service Portal account.

You can also contact SIOS Technology Corp. Support at:

1-877-457-5113 (Toll Free)

1-803-808-4270 (International)

Email: [support@us.sios.com](mailto:support@us.sios.com)

## 5.3. Ikbbackup

### Creating a Backup using Ikbbackup

**!** Before performing a LifeKeeper or OS upgrade, create a backup of your LifeKeeper hierarchies on all nodes.

To create a backup run the following command:

```
/opt/LifeKeeper/bin/lkbackup -c
```

The backup will be created in a file called:

```
/opt/LifeKeeper/config/archive.<date-time-stamp>.tar.gz
```

### Automatic Ikbbackup

**\*** An automatic **lkbackup** runs in versions 7.5 and above. The backup file is saved with the name: `/opt/LifeKeeper/config/auto-backup.<x>.tgz` (where <x> is a sequential number)

**To change the time of the automatic Ikbbackup:**

1. Go to `/etc/crontab`

The entry looks like this:

```
0 3 * * * root /opt/LifeKeeper/bin/backupadm -c (3 represents the hour the backup will run)
```

2. Change the hour to the desired time

### Restoring a Backup

**\*** Restoring from a backup is limited to current versions only. Backups should not be restored across versions.

To restore a backup run the following commands:

1. `/opt/LifeKeeper/bin/lkcli stop`
2. `/opt/LifeKeeper/bin/lkbackup -x -f <file-name>`

3. `/opt/LifeKeeper/bin/lkcli start`

 Restore the backup on all systems in the cluster.

The use of `lkbackup` with DataKeeper resources requires a [full resync](#).

## 5.4. LifeKeeper

---

The following LifeKeeper product documentation is available from the SIOS Technology Corp. website:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [Optional Recovery Kit Documentation](#)

## 5.4.1. SIOS LifeKeeper for Linux Introduction

---

SIOS LifeKeeper for Linux provides high availability clustering for up to 32 nodes with many supported storage configurations, including shared storage (Fiber Channel SAN, iSCSI), network attached storage (NAS), host-based replication, integration with array-based SAN replication including HP Continuous Access and VMware virtual hard disk (VMDK).

---

[Protected Resources](#)

[LifeKeeper Core](#)

[Configuration Concepts](#)

[Common Hardware Components](#)

[System Grouping Arrangements](#)

[Active – Active Grouping](#)

[Active – Standby Grouping](#)

[Intelligent vs Automatic Switchback](#)

[Logging With syslog](#)

[Resource Hierarchies](#)

[Resource Types](#)

[Resource States](#)

[Hierarchy Relationships](#)

[Shared Equivalencies](#)

[Resource Hierarchy Information](#)

[Resource Hierarchy Example](#)

[Detailed Status Display](#)

[Short Status Display](#)

[Fault Detection Recovery Scenarios](#)

[IP Local Recovery](#)

[Resource Error Recovery Scenario](#)

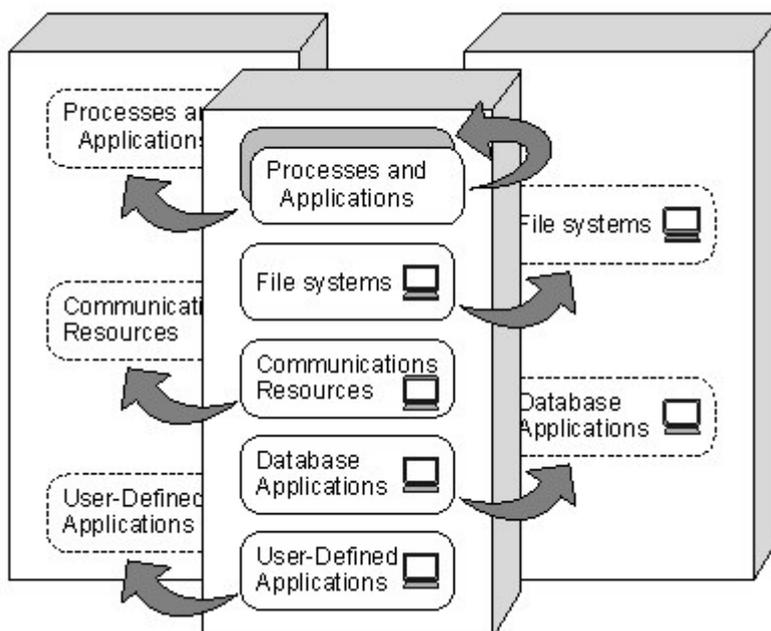
[Server Failure Recovery Scenario](#)

## 5.4.1.1. Protected Resources

The LifeKeeper family of products includes software that allows you to provide failover protection for a range of system resources. The following figure demonstrates LifeKeeper's flexibility and identifies the resource types you can specify for automatic recovery:

- **File systems.** LifeKeeper allows for the definition and failover of file systems, such as ext3, ext4, NFS, vxfs or xfs.
- **Communications resources.** LifeKeeper provides communications Recovery Kits for communications resources, such as TCP/IP.
- **Infrastructure resources.** LifeKeeper provides optional Recovery Kits for Linux infrastructure services, such as NFS, Samba, LVM, WebSphere MQ, and software RAID (md).
- **Web Server resources.** LifeKeeper provides an optional Recovery Kit for Apache Web Server resources.
- **Databases and other applications.** LifeKeeper provides optional Recovery Kits for major RDBMS products such as Oracle, MySQL and PostgreSQL, Sybase, SAP MaxDB, [SAP HANA DB](#), and for enterprise applications such as SAP.
- **Cloud resources.** LifeKeeper provides communications Recovery Kits for communications resources, such as EC2 EIP, and Route53.

LifeKeeper supports [N-Way Recovery](#) for a range of resource types.



## 5.4.1.2. LifeKeeper Core

---

LifeKeeper Core is composed of four major components:

- LifeKeeper Core Software
- File System, Generic Application, Raw I/O and IP Recovery Kit Software
- LifeKeeper GUI Software
- LifeKeeper Man Pages

### LifeKeeper Core Software

- The LifeKeeper Core Software consists of the following components:
  - [LifeKeeper Configuration Database \(LCD\)](#) – The LCD stores information about the LifeKeeper-protected resources. This includes information on resource instances, dependencies, shared equivalencies, recovery direction, and LifeKeeper operational flags. The data is cached in shared memory and stored in files so that the data can be remembered over system boots.
  - [LCD Interface \(LCDI\)](#) – The LCDI queries the configuration database (LCD) to satisfy requests for data or modifications to data stored in the LCD. The LCDI may also be used by the Application Recovery Kit to obtain resource state or description information.
  - [LifeKeeper Communications Manager \(LCM\)](#) – The LCM is used to determine the status of servers in the cluster and for LifeKeeper inter-process communication (local and remote). Loss of LCM communication across all communication paths on a server in the cluster indicates the server has failed.
  - [LifeKeeper Alarm Interface](#) – The LifeKeeper Alarm Interface provides the infrastructure for triggering an event. The sendevent program is called by application daemons when a failure is detected in a LifeKeeper-protected resource. The sendevent program communicates with the LCD to determine if recovery scripts are available.
  - LifeKeeper Recovery Action and Control Interface (LRACI) – The LRACI determines the appropriate recovery script to execute for a resource and invokes the appropriate restore / remove scripts for the resource.

### File System, Generic Application, IP and RAW I/O Recovery Kit Software

The LifeKeeper Core provides protection of specific resources on a server. These resources are:

- File Systems – LifeKeeper allows for the definition and failover of file systems on shared storage devices. A file system can be created on a disk that is accessible by two servers via a shared

SCSI bus. A LifeKeeper file system resource is created on the first server and then extended to the second server. [File System Health Monitoring](#) detects disk full and improperly mounted (or unmounted) file system conditions. Depending on the condition detected, the Recovery Kit may log a warning message, attempt a local recovery, or failover the file system resource to the backup server.

Specific help topics related to the File System Recovery Kit include [Creating](#) and [Extending](#) a File System Resource Hierarchy and [File System Health Monitoring](#).

- Generic Applications – The Generic Application Recovery Kit allows protection of a generic or user-defined application that has no predefined Recovery Kit to define the resource type. This kit allows a user to define monitoring and recovery scripts that are customized for a specific application.

Specific help topics related to the Generic Application Recovery Kit include [Creating](#) and [Extending](#) a Generic Application Resource Hierarchy.

- IP Addresses – The IP Recovery Kit provides a mechanism to recover a “switchable” IP address from a failed primary server to one or more backup servers in a LifeKeeper environment. A switchable IP address is a virtual IP address that can switch between servers and is separate from the IP address associated with the network interface card of each server. Applications under LifeKeeper protection are associated with the switchable IP address, so if there is a failure on the primary server, the switchable IP address becomes associated with the backup server. The resource under LifeKeeper protection is the switchable IP address.

Refer to the [IP Recovery Kit Technical Documentation](#) included with the Recovery Kit for specific product, configuration and administration information.

- RAW I/O – The RAW I/O Recovery Kit provides support for raw I/O devices for applications that prefer to bypass kernel buffering. The RAW I/O Recovery Kit allows for the definition and failover of raw devices bound to shared storage devices. The raw device must be configured on the primary node prior to resource creation. Once the raw resource hierarchy is [created](#), it can be [extended](#) to additional servers.
- Quick Service Protection (QSP) – QSP Recovery Kit provides a mechanism to simply protect OS services. Resources can be created for services that can be started/stopped with OS service commands. Generic Applications can provide the same protection, but QSP doesn't require code development. Also, you can create dependencies to start/stop services with applications protected by other resources.

However, QuickCheck of QSP only performs a simple check (using service command's “status”) and does not ensure the provision of the services and running of the processes. If complicated start/stop processing or robust check is required, please consider the use of Generic Applications.

For other topics regarding QSP, please see [Creating/extending QSP resources](#).

## LifeKeeper GUI Software

The LifeKeeper GUI is a client / server application developed using Java technology that provides a graphical administration interface to LifeKeeper and its configuration data. The LifeKeeper GUI client is implemented as a [stand-alone Java application](#).

## LifeKeeper Man Pages

The LifeKeeper Core reference manual pages for the LifeKeeper product.

## 5.4.1.3. Configuration Concepts

---

LifeKeeper works on the basis of resource hierarchies you define for groups of two or more servers. The following topics introduce the LifeKeeper failover configuration concepts:

---

[Common Hardware Components](#)

[System Grouping Arrangements](#)

[Active – Active Grouping](#)

[Active – Standby Grouping](#)

[Intelligent vs Automatic Switchback](#)

[Logging With syslog](#)

[Resource Hierarchies](#)

## 5.4.1.3.1. Common Hardware Components

All LifeKeeper configurations share these common components:

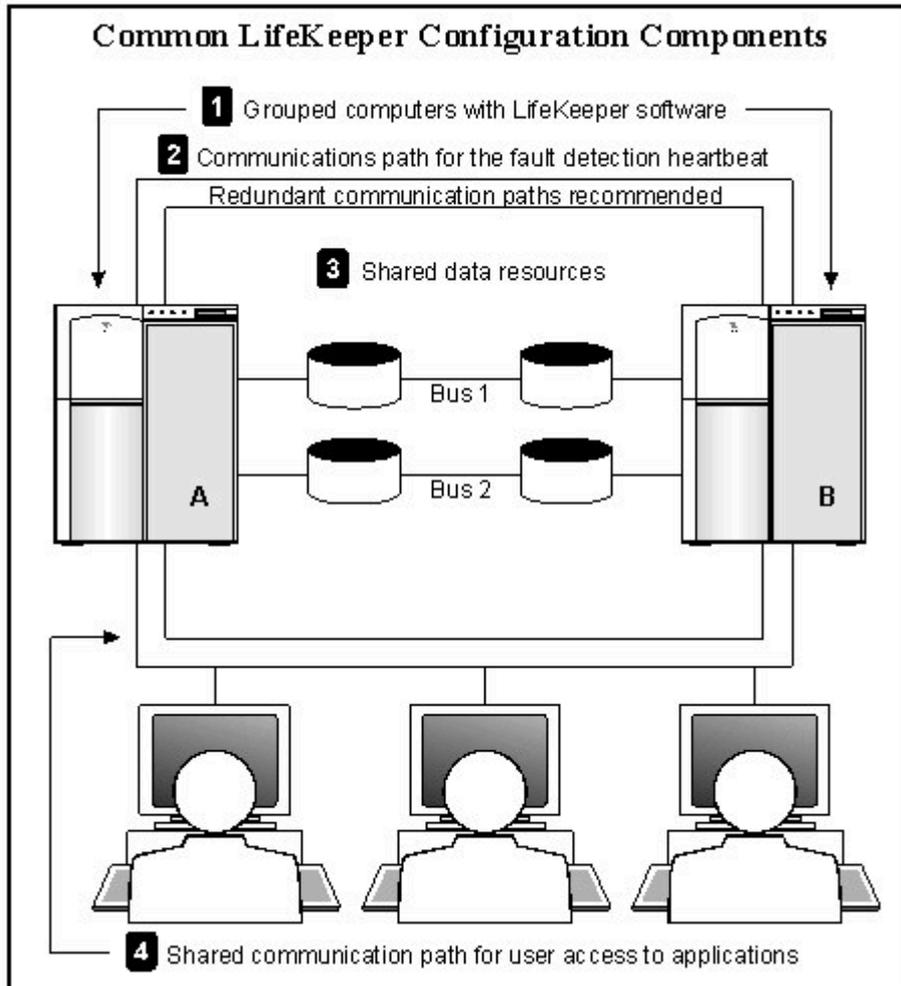
1. **Server groups.** The basis for the fault resilience provided by LifeKeeper is the grouping of two or more servers into a cluster. The servers can be any supported platform running a supported distribution of Linux. LifeKeeper gives you the flexibility to configure servers in multiple overlapping groups, but, for any given recoverable resource, the critical factor is the linking of a group of servers with defined roles or priorities for that resource. The priority of a server for a given resource is used to determine which server will recover that resource should there be a failure on the server where it is currently running. The highest possible priority value is one (1). The server with the highest priority value (normally 1) for a given resource is typically referred to as the primary server for that resource; any other servers are defined as backup servers for that resource.
2. **Communications paths.** The LifeKeeper heartbeat, a periodic message between servers in a LifeKeeper cluster, is a key fault detection facility. All servers within the cluster require redundant heartbeat communications paths (comm paths) to avoid system panics due to simple communications failures. Two separate LAN-based (TCP) comm paths using dual independent subnets are recommended (at least one of these should be configured as a private network); however, using a combination of TCP and TTY comm paths is supported. A TCP comm path can also be used for other system communications.

 **Note:** In a **cloud environment**, the internal configuration of the network is not open to the public, so it is difficult to physically prepare two LAN lines with different routes. Since it is expected that the physical network is basically redundant on the cloud side, operational reliability can be ensured even if there is **only one communication path**.

**Note:** A TTY comm path is used by LifeKeeper only for detecting whether other servers in the cluster are alive. The LifeKeeper GUI uses TCP/IP for communicating status information about protected resources; if there are two TCP comm paths configured, LifeKeeper uses the comm path on the public network for communicating resource status. Therefore if the network used by the LifeKeeper GUI is down, the GUI will show hierarchies on other servers in an UNKNOWN state, even if the TTY (or other TCP) comm path is operational.

3. **Shared data resources.** In shared storage configurations, servers in the LifeKeeper cluster share access to the same set of disks. In the case of a failure of the primary server, LifeKeeper automatically manages the unlocking of the disks from the failed server and the locking of the disks to the next available back-up server.
4. **Shared communication.** LifeKeeper can automatically manage switching of communications resources, such as TCP/IP addresses, allowing users to connect to the application regardless of where the application is currently active.

# Components Common to All LifeKeeper Configurations



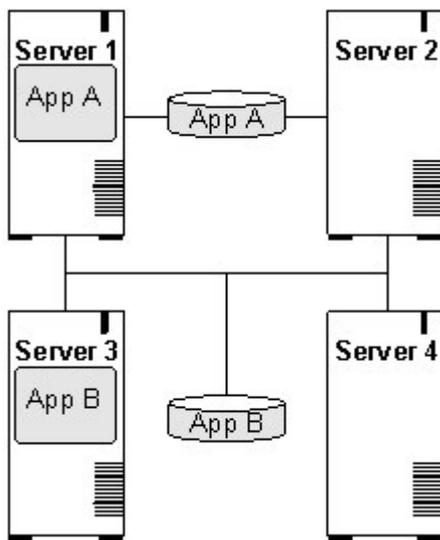
## 5.4.1.3.2. System Grouping Arrangements

A resource hierarchy is defined on a cluster of LifeKeeper servers. For a given hierarchy, each server is assigned a priority, with one (1) being the highest possible priority. The primary, or highest priority, server is the computer you want to use for the normal operation of those resources. The server having the second highest priority is the backup server to which you want LifeKeeper to switch those resources should the primary server fail.

In an [active/active group](#), all servers are active processors, but they also serve as the backup server for resource hierarchies on other servers. In an [active/standby group](#), the primary server is processing and any one of the backup servers can be configured to stand by in case of a failure on the primary server. The standby systems can be smaller, lower-performance systems, but they must have the processing capability to assure resource availability should the primary server fail.

Your physical connections and access to the shared resources determine your grouping options. To be grouped, servers must have communications and heartbeat paths installed and operational, and all servers must have access to the disk resources through a shared SCSI or Fibre Channel interface. For example, in the following diagram, there is only one grouping option for the resource *AppA* on Server 1. Server 2 is the only other server in the configuration that has shared access to the *AppA* database.

The resource *AppB* on Server 3, however, could be configured for a group including any one of the other three servers, because the shared SCSI bus in this example provides all four servers in the configuration access to the *AppB* database.



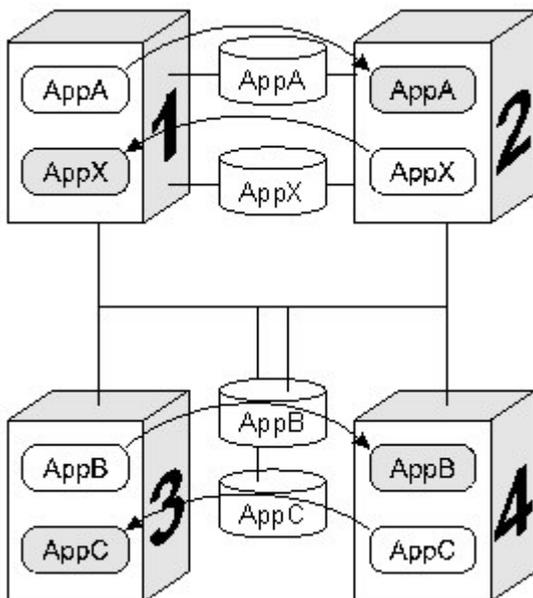
## 5.4.1.3.3. Active – Active Grouping

In an active/active pair configuration, all servers are active processors; they also serve as the backup server for resource hierarchies on other servers.

For example, the configuration example below shows two active/active pairs of servers. Server 1 is processing *AppA*, but also serves as the backup server for *AppX* running on Server 2. The reverse is also true. Server 2 is processing *AppX*, but also serves as the backup server for *AppA* running on Server 1. Servers 3 and 4 have the same type of active/active relationships.

Although the configurations on Servers 1 and 2 and the configurations on Servers 3 and 4 are similar, there is a critical difference. For the *AppA* and *AppX* applications, Servers 1 and 2 are the only servers available for grouping. They are the only servers that have access to the shared resources.

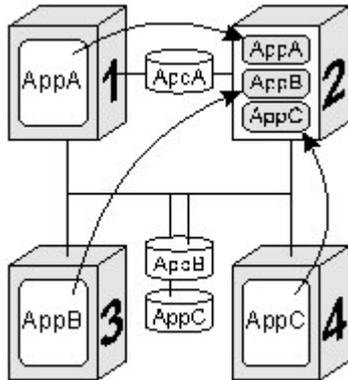
*AppB* and *AppC*, however, have several grouping options because all four servers have access to the *AppB* and *AppC* shared resources. *AppB* and *AppC* could also be configured to failover to Server1 and/or Server2 as a third or even fourth backup system.



\* **Note:** Because LifeKeeper applies locks at the disk level, only one of the four systems connected to the *AppB* and *AppC* disk resources can have access to them at any time. Therefore, when Server 3 is actively processing *AppB*, those disk resources are no longer available to Servers 1, 2, and 4, even though they have physical connections.

## 5.4.1.3.4. Active – Standby Grouping

In an active/standby pair configuration, the primary server is processing, and the back-up servers are standing by in case of a failure on the primary server. The standby systems can be smaller, lower-performance systems, but they must have the processing capability to assure resource availability should the primary server fail.



A standby server can provide backup for more than one active server. For example in the figure above, Server 2 is the standby server in three active/standby resource pairs. The LifeKeeper resource definitions specify the following active/standby paired relationships:

- *AppA* on *Server1* fails over to *Server2*.
- *AppB* on *Server3* fails over to *Server2*.
- *AppC* on *Server4* fails over to *Server2*.

Be aware of these three critical configuration concepts when you are considering configurations with multiple active/standby groups:

- **Disk ownership.** Different active applications cannot use disk partitions on the same shared disk or LUN from different servers. LifeKeeper applies locks at the disk or LUN level. When the SCSI locks are applied, only one system on the shared SCSI bus can access partitions on the disk or LUN. This requires that applications accessing different partitions on the same disk be active on the same server. In the example, Server 3 has ownership of the *AppB* disk resources and Server 4 owns the *AppC* resources.
- **Processing capacity.** Although it is unlikely that Servers 1, 3 and 4 would fail at the same time, you must take care when designating a standby server to support multiple resource relationships so that the standby server can handle all critical processing should multiple faults occur.
- **LifeKeeper administration.** In the example, Server 2 provides backup for three other servers. In general it is not desirable to administer the LifeKeeper database on the different logical groups simultaneously. You should first create the resources between the spare and one active system, then between the spare and another active system, and so on.

## 5.4.1.3.5. Intelligent Versus Automatic Switchback

---

By default, the switchback setting of a resource is *intelligent*. This means that once the failover occurs for that resource from *Server A* to *Server B*, the resource remains on *Server B* until another failure or until an administrator *intelligently* switches the resource to another server. Thus, the resource continues to run on *Server B* even after *Server A* returns to service. *Server A* now serves as a backup for the resource.

In some situations, it may be desirable for a resource to switch back automatically to the original failed server when that server recovers. LifeKeeper offers an *automatic switchback* option as an alternative to the default *intelligent switchback* behavior described above. This option can be configured for individual resource hierarchies on individual servers. If *automatic switchback* is selected for a resource hierarchy on a given server and that server fails, the resource hierarchy is failed over to a backup system; when the failed server recovers, the hierarchy is automatically switched back to the original server.

### Notes:

- For automatic switchback, switch back will take place automatically after the primary server comes back online and LifeKeeper communications path is re-established.
- LifeKeeper never performs an *automatic switchback* from a higher priority server to a lower priority server.

## 5.4.1.3.6. Logging With syslog

Beginning with LifeKeeper v8.0, logging is done through the standard syslog facility. LifeKeeper supports rsyslog, which is an extension of the original syslog protocol. During package installation, `syslog` will be configured to use the “local6” facility for all LifeKeeper log messages (if “local6” is already in use another local should be used). The `syslog` configuration file `/etc/rsyslog.conf` is modified to include LifeKeeper-specific routing sending all LifeKeeper log messages to `/var/log/lifekeeper.log` (the original configuration file will be backed up with the same name ending in “~”).

**!** **Important:** DO NOT edit LifeKeeper’s unique setup steps manually or upgrading and uninstalling may not function correctly.

The facility can be changed after installation by using the `lklogconfig` tool located in `/opt/LifeKeeper/bin`. For example, changing the facility to local5, run the following command.

```
lkstop -f
lklogconfig --action=update --facility=local5
lkstart
```

See the `lklogconfig(8)` manpage on a system with LifeKeeper installed for more details on this tool.

If a generic resource script puts a message into `/opt/LifeKeeper/out/log` directly, LifeKeeper will send a log message with the error severity level as the default into `/var/log/lifekeeper.log`. The severity level can be changed to information level by adding the following parameter into `/etc/default/LifeKeeper`, **LOGMGR\_LOGLEVEL=LK\_INFO**.

**\* Note:** When LifeKeeper is removed from a server, the LifeKeeper-specific `syslog` configuration will be removed.

**\* Note:** The LifeKeeper configuration files that control the log rotation can be found in `/etc/logrotate.d` as `lifekeeper` for the `lifekeeper.log` file and `lifekeeper-err` for `lifekeeper.err` log file. By default, log rotation will only occur when the size of the log file reaches 100MB.

## 5.4.1.3.7. Resource Hierarchies

---

The LifeKeeper GUI enables you to create a resource hierarchy on one server, then extend that hierarchy to one or more backup servers. LifeKeeper then automatically builds the designated hierarchies on all servers specified. LifeKeeper maintains hierarchy information in a database on each server. If you use the command line interface, you must explicitly define the hierarchy on each server.

After you create the resource hierarchy, LifeKeeper manages the stopping and starting of the resources within the hierarchy. The related topics below provide background for hierarchy definition tasks.

---

[Resource Types](#)

[Resource States](#)

[Hierarchy Relationships](#)

[Shared Equivalencies](#)

[Resource Hierarchy Information](#)

[Resource Hierarchy Example](#)

[Detailed Status Display](#)

[Short Status Display](#)

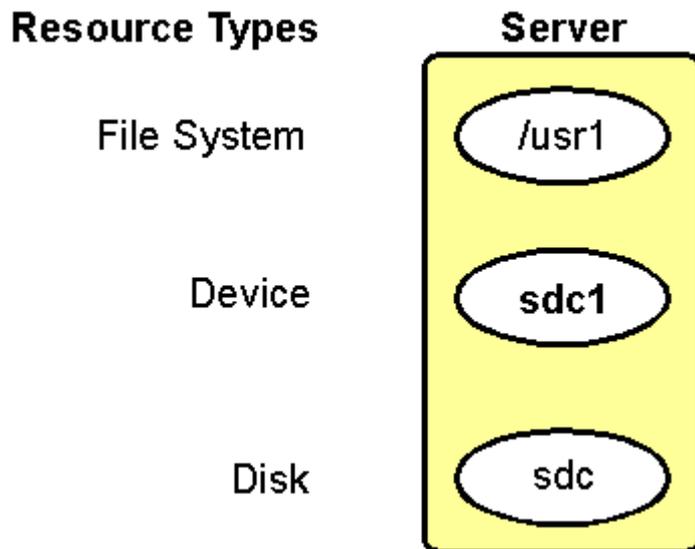
## 5.4.1.3.7.1. Resource Types

---

A resource can be either a hardware or software entity, categorized by resource type. LifeKeeper supplies file system and SCSI resource types, and the recovery kits provide communications, RDBMS and other application resource types.

For example, a hierarchy for a protected file system includes instances for resources of the following types:

- **filesystem**. Linux file system resource objects identified by their mount point.
- **device**. SCSI disk partitions and virtual disks, identified by their device file names, for example *sdc1*.
- **disk**. SCSI disks or RAID system logical units, identified by SCSI device name, for example *sd*.

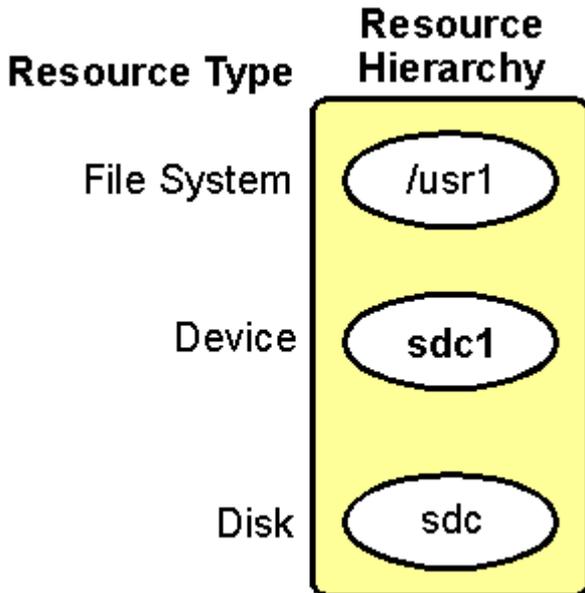


## 5.4.1.3.7.2. Resource States

State	Meaning
In-Service, Protected (ISP)	Resource is operational. LifeKeeper local recovery operates normally. LifeKeeper inter-server recovery and failure recovery is operational.
In-Service, Unprotected (ISU)	Resource is operational. However, no local recovery or failure recovery will occur because the LifeKeeper protection is not operational.  <b>Note:</b> When the file system protected by the file system resource(filesys) has reached at least 90% (the default threshold) of its capacity, the resource status will be changed to ISU to alert the user. In this case, the monitoring process is continued. For the monitoring of file system resources and its capacity, the ISU state is used differently from other resource types. Once the file system capacity drops below the threshold, the resource state will return to ISP.
Out-of-Service, Failed (OSF)	Resource has gone out-of-service because of a failure in the resource. Recovery has not been completed or has failed. LifeKeeper alarming is not operational for this resource.
Out-of-Service, Unimpaired (OSU)	Resource is out-of-service but available to take over a resource from another server.
Illegal (Undefined) State (ILLSTATE)	This state appears in situations where no state has been set for a resource instance. Under normal circumstances, this invalid state does not last long: a transition into one of the other states is expected. This state will occur if switchover occurs before all LifeKeeper information tables have been updated (for example, when LifeKeeper is first started up).

## 5.4.1.3.7.3. Hierarchy Relationships

LifeKeeper allows you to create relationships between resource instances. The primary relationship is a dependency, for example one resource instance depends on another resource instance for its operation . The combination of resource instances and dependencies is the resource hierarchy.



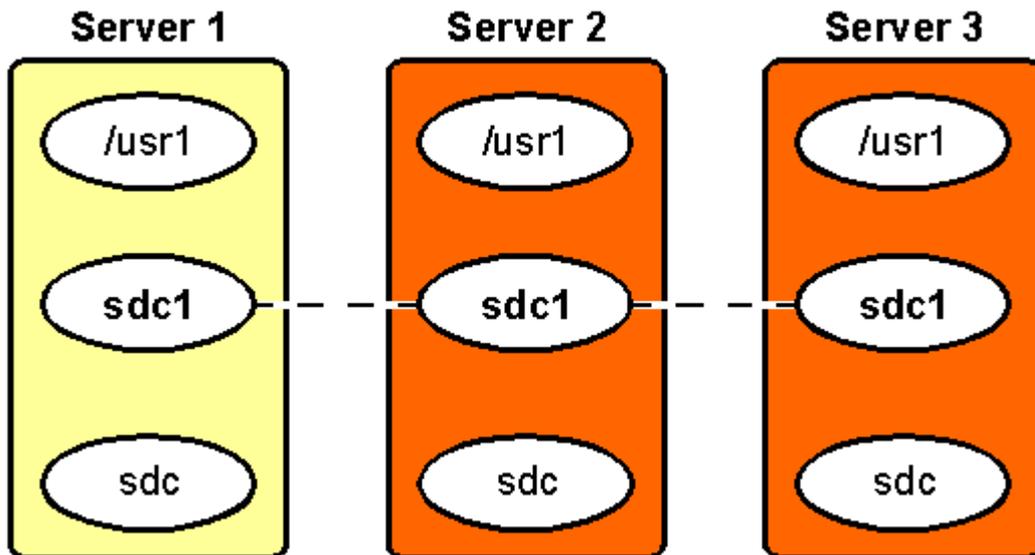
For example, since */usr1* depends on its operation upon the disk subsystem, you can create an ordered hierarchy relationship between */usr1* and those instances representing the disk subsystem.

The dependency relationships specified by the resource hierarchy tell LifeKeeper the appropriate order for bringing resource instances in service and out-of-service. In the example resource hierarchy, LifeKeeper cannot bring the */usr1* resource into service until it successfully brings into service first the *disk* and *device* instances.

## 5.4.1.3.7.4. Shared Equivalencies

When you create and extend a LifeKeeper resource hierarchy, the hierarchy exists on *both* the primary and the secondary servers. Most resource instances can be active on only one server at a time. For such resources, LifeKeeper defines a second kind of relationship called a shared equivalency that ensures that when the resource is *in-service* on one server, it is *out-of-service* on the other servers on which it is defined.

In the example below, the shared equivalency between the disk partition resource instances on each server is represented. Each resource instance will have a similar equivalency in this example.



## 5.4.1.3.7.5. Resource Hierarchy Information

The resource status of each resource is displayed in the [Detailed Status Display](#) and the [Short Status Display](#). The LifeKeeper tag names of root resources are displayed beginning in the left-most position of the TAG column, with tag names of resources within the hierarchy indented appropriately to indicate dependency relationships between resources.

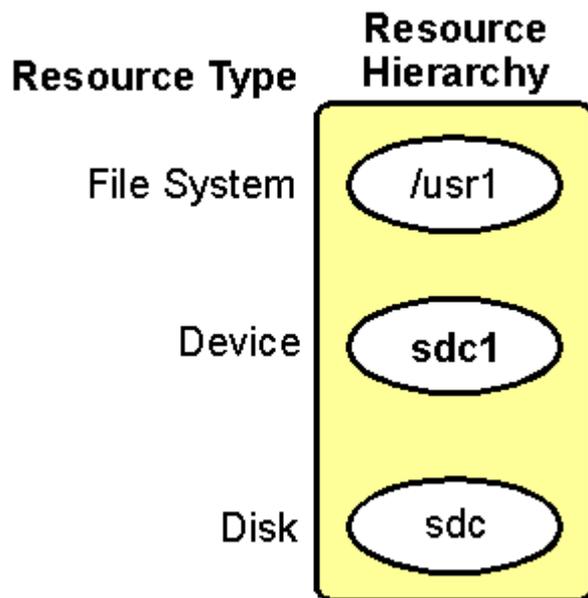
The following sample is from the resource hierarchy section of a short status display (the device and disk ID's are shortened to fit in the display area):

LOCAL	TAG	ID	STATE	PRIO	PRIMARY
svr1	app3910-on-svr1	app4238	ISP	1	svr2
svr1	filesys4083	/jrl1	ISP	1	svr2
svr1	device2126	000...300-1	ISP	1	svr2
svr1	disk2083	000...300	ISP	1	svr2

See [Resource Hierarchy Example](#) for an illustration of a hierarchy. For more information, see the Resource Hierarchy Information section of [Detailed Status Display](#) and [Short Status Display](#).

## 5.4.1.3.7.6. Resource Hierarchy Example

---



## 5.4.1.3.7.7. Detailed Status Display

---

This topic describes the categories of information provided in the detailed status display as shown in the following example of output from the **lcdstatus** command. For information on how to display this information, see the LCD(1M) man page. At the command line, you can enter either **man lcdstatus** or **man LCD**. For status information available in the LifeKeeper GUI, see [Viewing the Status of a Server](#) or [Viewing the Status of Resources](#).

Example of detailed status display:

### [Resource Hierarchy Information](#)

```
Resource hierarchies for machine "wileecoyote":

ROOT of RESOURCE HIERARCHY

apache-home.fred: id=apache-home.fred app=webserver type=apache
state=ISP

initialize=(AUTORES_ISP) automatic restore to IN-SERVICE by LifeKeeper

info=/home/fred /usr/sbin/httpd

reason=restore action has succeeded

depends on resources:
ipeth0-172.17.104.25,ipeth0-172.17.106.10,ipeth0-172.17.106.105

Local priority = 1

SHARED equivalency with "apache-home.fred" on "roadrunner", priority =
10

FAILOVER ALLOWED

ipeth0-172.17.104.25: id=IP-172.17.104.25 app=comm type=ip state=ISP

initialize=(AUTORES_ISP) automatic restore to IN-SERVICE by LifeKeeper

info=wileecoyote eth0 172.17.104.25 fffffc00

reason=restore action has succeeded

these resources are dependent: apache-home.fred
```

Local priority = 1

SHARED equivalency with "ipeth0-172.17.104.25" on "roadrunner",  
priority = 10

FAILOVER ALLOWED

ipeth0-172.17.106.10: id=IP-172.17.106.10 app=comm type=ip state=ISP

initialize=(AUTORES\_ISP) automatic restore to IN-SERVICE by LifeKeeper

info=wileecoyote eth0 172.17.106.10 fffffc00

reason=restore action has succeeded

these resources are dependent: apache-home.fred

Local priority = 1

SHARED equivalency with "ipeth0-172.17.106.10" on "roadrunner",  
priority = 10

FAILOVER ALLOWED

ipeth0-172.17.106.105: id=IP-172.17.106.105 app=comm type=ip state=ISP

initialize=(AUTORES\_ISP) automatic restore to IN-SERVICE by LifeKeeper

info=wileecoyote eth0 172.17.106.105 fffffc00

reason=restore action has succeeded

These resources are dependent: apache-home.fred

Local priority = 1

SHARED equivalency with "ipeth0-172.17.106.105" on "roadrunner",  
priority = 10

FAILOVER ALLOWED

### Communication Status Information

The following LifeKeeper servers are known:

```
machine=wileecoyote state=ALIVE
```

```
machine=roadrunner state=DEAD (eventslcm detected failure at Wed Jun 7
15:45:14 EDT 2000)
```

The following LifeKeeper network connections exist:

```
to machine=roadrunner type=TCP addresses=192.168.1.1/192.168.105.19
state="DEAD" priority=2 #comm_downs=0
```

### [LifeKeeper Flags](#)

The following LifeKeeper flags are on:

```
shutdown_switchover
```

### [Shutdown Strategy](#)

The shutdown strategy is set to: switchover.

## Resource Hierarchy Information

LifeKeeper displays the resource status beginning with the root resource. The display includes information about all resource dependencies.

Elements common to multiple resources appear only once under the first root resource. The first line for each resource description displays the resource tag name followed by a colon (:), for example:

`device13557: .` These are the information elements that may be used to describe the resources in the hierarchy:

- **id.** Unique resource identifier string used by LifeKeeper.
- **app.** Identifies the type of application, for example the sample resource is a *webserver* application.
- **type.** Indicates the resource class type, for example the sample resource is an *Apache* application.
- **state.** Current state of the resource:
  - ISP—In-service locally and protected.
  - ISU—In-service, unprotected.
  - OSF—Out-of-service, failed.
  - OSU—Out-of-service, unimpaired.
- **initialize.** Specifies the way the resource is to be initialized, for example LifeKeeper restores the application resource, but the host adapter initializes without LifeKeeper.

- **info.** Contains object-specific information used by the object's `remove` and `restore` scripts.
- **reason.** If present, describes the reason the resource is in its current state. For example, an application might be in the OSU state because it is in-service (ISP or ISU) on another server. Shared resources can be active on only one of the grouped servers at a time.
- **depends on resources.** If present, lists the tag names of the resources on which this resource depends.
- **these resources are dependent.** If present, indicates the tag names of all parent resources that are directly dependent on this object.
- **Local priority.** Indicates the failover priority value of the targeted server, for this resource.
- **SHARED equivalency.** Indicates the resource tag and server name of any remote resources with which this resource has a defined equivalency, along with the failover priority value of the remote server, for that resource.
- **FAILOVER ALLOWED.** If present, indicates that LifeKeeper is operational on the remote server identified in the equivalency on the line above, and the application is protected against failure. `FAILOVER INHIBITED` means that the application is not protected due to either the shutting down of LifeKeeper or the stopping of the remote server.

## Communication Status Information

This section of the status display lists the servers known to LifeKeeper and their current state, followed by information about each communications path.

These are the communications information elements you can see on the status display:

- **State.** Status of communications path. These are the possible communications state values:
  - ALIVE — Functioning normally
  - DEAD — No longer functioning normally
- **priority.** The assigned priority value for the communications path. This item is displayed only for TCP paths.
- **#comm\_downs.** The number of times the port has failed and caused a failover. The path failure causes a failover only if no other communications paths are marked "ALIVE" at the time of the failure.

In addition, the status display can provide any of the following statistics maintained only for TTY communications paths:

- **wrpid.** Each TTY communications path has unique reader and writer processes. The `wrpid` field contains the process ID for the writer process. The writer process sleeps until one of two conditions occurs:
  - Heartbeat timer expires, causing the writer process to send a message.

- Local process requests the writer process to transmit a LifeKeeper maintenance message to the other server. The writer process transmits the message, using its associated TTY port, to the reader process on that port on the other system.
- **rdpid**. Each TTY communications path has unique reader and writer processes. The `rdpid` field contains the process ID for the reader process. The reader process sleeps until one of two conditions occurs:
  - Heartbeat timer expires and the reader process must determine whether the predefined heartbeat intervals have expired. If so, the reader process marks the communications path in the DEAD state, which initiates a failover event if there are no other communications paths marked ALIVE.
  - Remote system writer process transmits a LifeKeeper maintenance message, causing the reader process to perform the protocol necessary to receive the message.
- **#NAKs**. Number of times the writer process received a negative acknowledgment (NAK). A `NAK` message means that the reader process on the other system did not accept a message packet sent by the writer process, and the writer process had to re-transmit the message packet. The `#NAKs` statistic can accumulate over a long period of time due to line noise. If, however, you see the numbers increasing rapidly, you should perform diagnostic procedures on the communications subsystem.
- **#chksumerr**. Number of mismatches in the check sum message between the servers. This statistic can accumulate over a long period of time due to line noise. If, however, you see the numbers increasing rapidly, you should perform diagnostic procedures on the communications subsystem.
- **#incmpltmes**. Number of times the incoming message packet did not match the expected size. A high number of mismatches may indicate that you should perform diagnostic procedures on the hardware port associated with the communications path.
- **#noreply**. Number of times the writer process timed out while waiting for an acknowledgment and had to re-transmit the message. Lack of acknowledgment may indicate an overloaded server or it can signal a server failure.
- **#pacresent**. Number of times the reader process received the same packet. This can happen when the writer process on the sending server times out and resends the same message.
- **#pacoutseq**. Number of times the reader received packets out of sequence. High numbers in this field can indicate lost message packets and may indicate that you should perform diagnostic procedures on the communications subsystem.
- **#maxretrys**. Metric that increments for a particular message when the maximum retransmission count is exceeded (for `NAK` and `noreply` messages). If you see a high number in the `#maxretrys` field, you should perform diagnostic procedures on the communications subsystem.

## LifeKeeper Flags

Near the end of the detailed status display, LifeKeeper provides a list of the flags set for the system. A common type is a Lock LCD flag used to ensure that other processes wait until the process lock completes its action. The following is the standard LCD lock format:

```
!action!processID!time!machine:id.
```

These are examples of general LCD lock flags:

- `!action!02833!701236710!server1:filesys` – The creation of a file system hierarchy produces a flag in this format in the status display. The `filesys` designation can be a different resource type for other application resource hierarchies, or `app` for generic or user-defined applications.
- Other typical flags include `!nofailover!machine`, `!notarmode!machine`, and `shutdown_switchover`. The `!nofailover!machine` and `!notarmode!machine` flags are internal, transient flags created and deleted by LifeKeeper, which control aspects of server failover. The `shutdown_switchover` flag indicates that the shutdown strategy for this server has been set to `switchover` such that a shutdown of the server will cause a switchover to occur. See the LCDI-flag(1M) for more detailed information on the possible flags.

## Shutdown Strategy

The last item on the detailed status display identifies the LifeKeeper shutdown strategy selected for this system. See [Setting Server Shutdown Strategy](#) for more information.

## 5.4.1.3.7.8. Short Status Display

This topic describes the categories of information provided in the short status display as shown in the following example of output from the **lcdstatus -e** command. For information on how to display this information, see the `LCD` man page. At the command line, you can enter either **man lcdstatus** or **man LCD**. For status information available in the LifeKeeper GUI, see [Viewing the Status of a Server](#) or [Viewing the Status of Resources](#).

Example of Short Status Display (Example of output from the `lcdstatus -e` command):

### [Resource Hierarchy Information](#)

BACKUP	TAG	ID	STATE	PRIO	PRIMARY
svr1	appfs3910-on-svr1	appfs4238	ISP	1	svr2
svr1	filesys4083	/jrl1	ISP	1	svr2
svr1	device2126	000...300-1	ISP	1	svr2
svr1	disk2083	000...300	ISP	1	svr2

### [Communication Status Information](#)

MACHINE	NETWORK	ADDRESSES/DEVICE	STATE	PRIO
svr1	TCP	100.10.1.20/100.11.1.21	ALIVE	1
svr1	TTY	/dev/ttyS0	ALIVE	--

## Resource Hierarchy Information

LifeKeeper displays the resource status of each resource. The LifeKeeper tag names of root resources are displayed beginning in the left-most position of the **TAG** column, with tag names of resources within the hierarchy indented appropriately to indicate dependency relationships between resources.

The **BACKUP** column indicates the next system in the failover priority order, after the system for which the status display pertains. If the target system is the lowest priority system for a given resource, the **BACKUP** column for that resource contains dashes (for example, -----).

- **TAG column.** Contains the root tag for the resource.
- **ID column.** Contains each resource's identifier string.
- **STATE column.** Contains the current state of each resource, as described in [Resource States](#).
- **PRIO column.** Contains the failover priority value of the local server, for each resource.
- **PRIMARY column.** Contains the name of the server with the highest priority, for each resource.

## Communication Status Information

This section of the display lists each communications path that has been defined on the target system. For each path, the following information is provided.

- **MACHINE.** Remote server name for the communications path.
- **NETWORK.** The type of communications path (TCP or TTY)
- **ADDRESSES/DEVICE.** The pair of IP addresses or device name for the communications path
- **STATE.** The state of the communications path (ALIVE or DEAD)
- **PRIOR.** For TCP paths, the assigned priority of the path. For TTY paths, this column will contain dashes (----), since TTY paths do not have an assigned priority.

## 5.4.1.4. Fault Detection and Recovery Scenarios

---

To demonstrate how the various LifeKeeper components work together to provide fault detection and recovery, see the following topics that illustrate and describe three types of recovery scenarios:

---

[IP Local Recovery](#)

[Resource Error Recovery Scenario](#)

[Server Failure Recovery Scenario](#)

## 5.4.1.4.1. IP Local Recovery

---

### Local Recovery Scenario

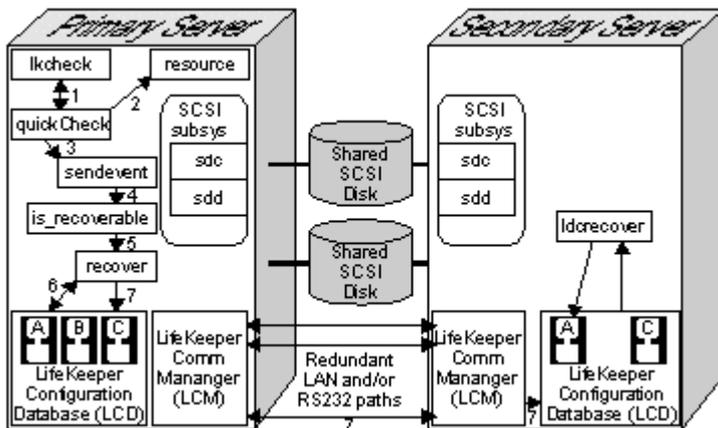
When a failure of an IP address is detected by the IP Recovery Kit, the resulting failure triggers the execution of the **IP local recovery script** which is responsible for running remove and restore functions.

\* SIOS recommends the use of **bonded interfaces** via the standard Linux NIC bonding mechanism in any **LifeKeeper release where a backup interface** is required.

## 5.4.1.4.2. Resource Error Recovery Scenario

LifeKeeper provides a real-time daemon monitor, **lkcheck**, to check the status and health of LifeKeeper-protected resources. For each in-service resource, **lkcheck** periodically calls the quickCheck script for that resource type. The **quickCheck** script performs a quick health check of the resource, and if the resource is determined to be in a failed state, the **quickCheck** script calls the event notification mechanism, **sendevent**.

The following figure illustrates the recovery process tasks when **lkcheck** initiates the process:



1. **lkcheck** runs. By default, the **lkcheck** process runs once every two minutes. When **lkcheck** runs, it invokes the appropriate **quickCheck** script for each in-service resource on the system.
2. **quickCheck** script checks resource. The nature of the checks performed by the **quickCheck** script is unique to each resource type. Typically, the script simply verifies that the resource is available to perform its intended task by imitating a client of the resource and verifying that it receives the expected response.
3. **quickCheck** script invokes **sendevent**. If the **quickCheck** script determines that the resource is in a failed state, it initiates an event of the appropriate class and type by calling **sendevent**.
4. Recovery instruction search. The system event notification mechanism, **sendevent**, first attempts to determine if the LCD has a resource and/or recovery for the event type or component. To make this determination, the **is\_recoverable** process scans the resource hierarchy in LCD for a resource instance that corresponds to the event (in this example, the filesystem name).

The action in the next step depends upon whether the scan finds resource-level recovery instructions:

- Not found. If resource recovery instructions are not found, **is\_recoverable** returns to **sendevent** and **sendevent** continues with basic event notification.
- Found. If the scan finds the resource, **is\_recoverable** forks the **recover** process into the background. The **is\_recoverable** process returns and **sendevent** continues with basic event notification, passing an advisory flag “-A” to the basic alarming event response scripts, indicating that LifeKeeper is performing recovery.

5. Recover process initiated. Assuming that recovery continues, `is_recoverable` initiates the recover process which first attempts local recovery.
6. Local recovery attempt. If the instance was found, the recover process attempts local recovery by accessing the resource hierarchy in LCD to search the hierarchy tree for a resource that knows how to respond to the event. For each resource type, it looks for a recovery subdirectory containing a subdirectory named for the event class, which in turn contains a recovery script for the event type.

The recover process runs the recovery script associated with the resource that is farthest above the failing resource in the resource hierarchy. If the recovery script succeeds, recovery halts. If the script fails, recover runs the script associated with the next resource, continuing until a recovery script succeeds or until recover attempts the recovery script associated with the failed instance.

*If local recovery succeeds, the recovery process halts.*

7. Inter-server recovery begins. If local recovery fails, the event then escalates to inter-server recovery.
8. Recovery continues. Since local recovery fails, the recover process marks the failed instance to the *Out-of-Service-FAILED* (OSF) state and marks all resources that depend upon the failed resource to the *Out-of-Service-UNIMPAIRED* (OSU) state. The recover process then determines whether the failing resource or a resource that depends upon the failing resource has any shared equivalencies with a resource on any other systems, and selects the one to the highest priority alive server. Only one equivalent resource can be active at a time.

*If no equivalency exists, the recover process halts.*

If a shared equivalency is found and selected, LifeKeeper initiates inter-server recovery. The recover process sends a message through the LCM to the LCD process on the selected backup system containing the shared equivalent resource. This means that LifeKeeper would attempt inter-server recovery.

9. **lcdrecover** process coordinates transfer. The LCD process on the backup server forks the process **lcdrecover** to coordinate the transfer of the equivalent resource.
10. Activation on backup server. The **lcdrecover** process finds the equivalent resource and determines whether it depends upon any resources that are not in-service. **lcdrecover** runs the restore script (part of the resource recovery action scripts) for each required resource, placing the resources in-service.

The act of restoring a resource on a backup server may result in the need for more shared resources to be transferred from the primary system. Messages pass to and from the primary system, indicating resources that need to be removed from service on the primary server and then brought into service on the selected backup server to provide full functionality of the critical applications. This activity continues, until no new shared resources are needed and all necessary resource instances on the backup are restored.

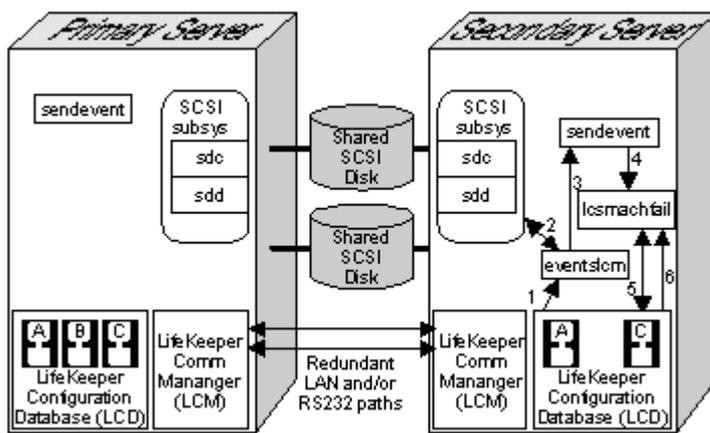
## 5.4.1.4.3. Server Failure Recovery Scenario

The LifeKeeper Communications Manager ([LCM](#)) has two functions:

- Messaging. The LCM serves as a conduit through which LifeKeeper sends messages during recovery, configuration, or when running an audit.
- Failure detection. The LCM also plays a role in detecting whether or not a server has failed.

LifeKeeper has a built-in heartbeat signal that periodically notifies each server in the configuration that its paired server is operating. If a server fails to receive the heartbeat message through one of the communications paths, LifeKeeper marks that path DEAD.

The following figure illustrates the recovery tasks when the LCM heartbeat mechanism detects a server failure.



The following steps describe the recovery scenario, illustrated above, if LifeKeeper marks all communications connections to a server DEAD.

1. LCM activates **eventslcm**. When LifeKeeper marks all communications paths dead, the LCM initiates the **eventslcm** process.

Only one activity stops the **eventslcm** process:

- Communication path alive. If one of the communications paths begins sending the heartbeat signal again, the LCM stops the **eventslcm** process.

It is critical that you configure two or more physically independent, redundant communication paths between each pair of servers to prevent failovers and possible system panics due to communication failures.

2. Message to **sendevent**. **eventslcm** sends the system failure alarm by calling **sendevent** with the event type *machfail*.
3. **sendevent** initiates failover recovery. The **sendevent** program determines that LifeKeeper can handle the system failure event and executes the LifeKeeper failover recovery process

**lcdmachfail.**

4. **lcmachfail** checks. The **lcmachfail** process first checks to ensure that the non-responding server was not shut down. Failovers are inhibited if the other system was shut down gracefully before system failure. Then **lcmachfail** determines all resources that have a shared equivalency with the failed system. This is the commit point for the recovery.
5. **lcmachfail** restores resources. **lcmachfail** determines all resources on the backup server that have shared equivalencies with the failed primary server. It also determines whether the backup server is the highest priority alive server for which a given resource is configured. All backup servers perform this check, so that only one server will attempt to recover a given hierarchy. For each equivalent resource that passes this check, **lcmachfail** invokes the associated restore program. Then, **lcmachfail** also restores each resource dependent on a restored resource, until it brings the entire hierarchy into service on the backup server.

## 5.4.2. Installation and Configuration

---

### LifeKeeper for Linux Installation

For complete installation instructions on installing the LifeKeeper for Linux software, see the [LifeKeeper for Linux Installation Guide](#). Refer to the [LifeKeeper for Linux Release Notes](#) for additional information.

### LifeKeeper for Linux Configuration

Once the LifeKeeper environment has been installed, the LifeKeeper software can be configured on each server in the cluster. Follow the steps in **LifeKeeper Configuration Steps** below which contains links to topics with additional details.

---

[Configuration Steps](#)

[Event Forwarding via SNMP](#)

[Event Email Notification](#)

[Optional Configuration Tasks](#)

[Linux Configuration](#)

[Data Replication Configuration](#)

[Network Configuration](#)

[Application Configuration](#)

[Storage and Adapter Configuration](#)

[Fencing](#)

[Resource Policy Management](#)

[Configuring Credentials](#)

## 5.4.2.1. LifeKeeper Configuration Steps

---

If you have installed your LifeKeeper environment as described in the LifeKeeper for Linux Installation Guide, you should be ready to start and configure the LifeKeeper for Linux software on each server in your cluster.

Follow the steps below which contain links to topics with additional details. Perform these tasks on each server in the cluster.

1. Start LifeKeeper by typing the following command as root:

```
$LKROOT/bin/lkcli start
```

This command starts all LifeKeeper daemon processes on the server being administered if they are not currently running.

For additional information on starting and stopping LifeKeeper, see [Starting LifeKeeper](#) and [Stopping LifeKeeper](#).

2. [Set Up TTY Communications Connections](#). If you plan to use a TTY communications (comm) path for a LifeKeeper heartbeat, you need to set up the physical connection for that heartbeat.
3. Configure the GUI. There are multiple tasks involved with configuring the GUI. Start with the [LifeKeeper GUI – Overview](#) topic within [Preparing to Run the GUI](#). Then for detailed instructions, follow the browse sequence throughout [Preparing to Run the GUI](#).

**Note:** The first time you run the LifeKeeper GUI, you will see a QuickStart button which opens a window with instructions and links to help you step through the configuration of your LifeKeeper resources. Subsequently, you can access this QuickStart Configuration Assistant under the [Help menu](#).

4. [Create Communication Paths](#). Before you can activate LifeKeeper protection, you must create the communications path (heartbeat) definitions within LifeKeeper.
5. Perform any of the following optional configuration tasks:
  - [Set the Server Shutdown Strategy](#)
  - [Configure the manual failover confirmation option](#)
  - [Tune the LifeKeeper heartbeat](#)
  - [Configure SNMP Event Forwarding via SNMP](#)
  - [Configure Event Email Notification](#)
  - If you plan to use [STONITH](#) devices in your cluster, create the scripts to control the STONITH devices and place them in the appropriate LifeKeeper events directory.

6. LifeKeeper for Linux is now ready to protect your applications. The next step depends on whether you will be using one of the optional LifeKeeper Recovery Kits:
- If you are using a LifeKeeper Recovery Kit, refer to the Documentation associated with the kit for instructions on creating and extending your resource hierarchies.
  - If you are using an application that does not have an associated Recovery Kit, then you have two options:
    - If it is a simple application, you should carefully plan how to create an interface between your application and LifeKeeper. You may decide to protect it using the [Generic Application Recovery Kit](#) included with the LifeKeeper core.
    - Services provided by the operating system can easily be protected by using the [Quick Service Protection \(QSP\) Recovery Kit](#) include with the LifeKeeper Core. However, please be aware that quickCheck will only perform a simple check of the service state.

## 5.4.2.1.1. Set Up TTY Connections

---

If you plan to use a TTY communications (comm) path for a LifeKeeper heartbeat, you need to set up the physical connection for that heartbeat. Remember that multiple communication paths are required to avoid false failover due to a simple communications failure. Two or more LAN-based (TCP) comm paths should also be used.

Connect the TTY cable to the serial ports of each server to be used for the serial heartbeat.

1. Test the serial path using the following command:

```
/opt/LifeKeeper/bin/portio -r -p port -b baud
```

where:

- `baud` is the baud rate selected for the path (normally 9600)
- `port` is the serial port being tested on Server 1, for example `/dev/ttyS0`.
- Server 1 is now waiting for input from Server 2.

2. Run command `portio` on Server 2. On the second system in the pair, type the following command:

```
echo Helloworld | /opt/LifeKeeper/bin/portio -p port -b baud
```

where:

- `baud` is the same baud rate selected for Server 1.
- `port` is the serial port being tested on Server 2, for example `/dev/ttyS0`.

3. View the console. If the communications path is operational, the software writes “Helloworld” on the console on Server 1. If you do not see that information, perform diagnostic and correction operations before continuing with LifeKeeper configuration.

## 5.4.2.2. LifeKeeper Event Forwarding via SNMP

---

[Overview](#)

[Configuration](#)

[Troubleshooting](#)

✿ See the following documentation for kit specific traps:

[DataKeeper Events Table](#)

[EC2 Event Table](#)

## 5.4.2.2.1. Overview of LifeKeeper Event Forwarding via SNMP

The Simple Network Management Protocol (SNMP) defines a device-independent framework for managing networks. Devices on the network are described by MIB (Management Information Base) variables that are supplied by the vendor of the device. An SNMP agent runs on each node of the network, and interacts with a Network Manager node. The Network Manager can query the agent to get or set the values of its MIB variables, there by monitoring or controlling the agent's node. The agent can also asynchronously generate messages called traps to notify the manager of exceptional events. There are a number of applications available for monitoring and managing networks using the Simple Network Management Protocol (SNMP).

LifeKeeper has an event notification mechanism for registering applications that wish to be notified of specific events or alarms (see the `sendevent(5)` man page). LifeKeeper can be easily enabled to send SNMP trap notification of key LifeKeeper events to a third party network management console wishing to monitor LifeKeeper activity.

The remote management console receiving SNMP traps must first be configured through the administration software of that system; LifeKeeper provides no external SNMP configuration. The remote management server is typically located outside of the LifeKeeper cluster (i.e., it is not a LifeKeeper node).

### LifeKeeper Events Table

The following table contains the list of LifeKeeper events and associated trap numbers. The entire Object ID (OID) consists of a prefix followed by a specific trap number in the following format:

```
prefix.0.specific trap number
```

The prefix is **.1.3.6.1.4.1.7359**, which expands to **iso.org.dod.internet.private.enterprises.7359** in the MIB tree. (7359 is SteelEye's [SIOS Technology] enterprise number, followed by 1 for LifeKeeper.) For example, the LifeKeeper Startup Complete event generates the OID: **.1.3.6.1.4.1.7359.1.0.100**.

LifeKeeper Event/Description	Trap #	Object ID
<b>LifeKeeper Startup Complete</b> Sent from a node when LifeKeeper is started on that node	100	.1.3.6.1.4.1.7359.1.0.100
<b>LifeKeeper Shutdown Initiated</b> Sent from a node beginning LifeKeeper shutdown	101	.1.3.6.1.4.1.7359.1.0.101

<p><b>LifeKeeper Shutdown Complete</b></p> <p>Sent from a node completing LifeKeeper shutdown</p>	102	.1.3.6.1.4.1.7359.1.0.102
<p><b>LifeKeeper Manual Switchover Initiated on Server</b></p> <p>Sent from the node from which a manual switchover was requested</p>	110	.1.3.6.1.4.1.7359.1.0.110
<p><b>LifeKeeper Manual Switchover Complete – recovered list</b></p> <p>Sent from the node where the manual switchover was completed</p>	111	.1.3.6.1.4.1.7359.1.0.111
<p><b>LifeKeeper Manual Switchover Complete – failed list</b></p> <p>Sent from each node within the cluster where the manual switchover failed</p>	112	.1.3.6.1.4.1.7359.1.0.112
<p><b>LifeKeeper Node Failure Detected for Server</b></p> <p>Sent from each node within the cluster when a node in that cluster fails</p>	120	.1.3.6.1.4.1.7359.1.0.120
<p><b>LifeKeeper Node Recovery Complete for Server – recovered list</b></p> <p>Sent from each node within the cluster that has recovered resources from the failed node</p>	121	.1.3.6.1.4.1.7359.1.0.121
<p><b>LifeKeeper Node Recovery Complete for Server – failed list</b></p> <p>Sent from each node within the cluster that has failed to recover resources from the failed node</p>	122	.1.3.6.1.4.1.7359.1.0.122
<p><b>LifeKeeper Resource Recovery Initiated</b></p> <p>Sent from a node recovering a resource; a 131 or 132 trap always follows to indicate whether the recovery was completed or failed.</p>	130	.1.3.6.1.4.1.7359.1.0.130

<b>LifeKeeper Resource Recovery Failed</b>  Sent from the node in trap 130 when the resource being recovered fails to come into service	131*	.1.3.6.1.4.1.7359.1.0.131
<b>LifeKeeper Resource Recovery Complete</b>  Sent from the node in trap 130 when the recovery of the resource is completed	132	.1.3.6.1.4.1.7359.1.0.132
<b>LifeKeeper Communications Path Up</b>  A communications path to a node has become operational	140	.1.3.6.1.4.1.7359.1.0.140
<b>LifeKeeper Communications Path Down</b>  A communications path to a node has gone down	141	.1.3.6.1.4.1.7359.1.0.141
<b>LifeKeeper &lt;Node Monitoring&gt; Failure</b>  Sent from a node where a failure was detected with Node Monitoring of the Standby Node Health Check. Detected failure is described in <Node Monitoring>.	190	.1.3.6.1.4.1.7359.1.0.190
<b>LifeKeeper &lt;OSUquickCheck&gt; Failure</b>  Sent from a node where a failure was detected with OSU Resource Monitoring of the Standby Node Health Check. Tag name of the resource where the failure was detected is described in <OSUquickCheck>.	200	.1.3.6.1.4.1.7359.1.0.200
<b>The following variables are used to “carry” additional information in the trap PDU:</b>		
Trap message	all	.1.3.6.1.4.1.7359.1.1
Resource Tag	130	.1.3.6.1.4.1.7359.1.2
Resource Tag	131	.1.3.6.1.4.1.7359.1.2
Resource Tag	132	.1.3.6.1.4.1.7359.1.2
List of recovered resources	111	.1.3.6.1.4.1.7359.1.3
List of recovered resources	121	.1.3.6.1.4.1.7359.1.3

List of failed resources	112	.1.3.6.1.4.1.7359.1.4
List of failed resources	122	.1.3.6.1.4.1.7359.1.4

- This trap may appear multiple times if recovery fails on multiple backup servers.

## 5.4.2.2.2. Configuring LifeKeeper Event Forwarding

---

### Prerequisites

The SNMP event forwarding feature is included as part of the LifeKeeper Core functionality and does not require additional LifeKeeper packages to be installed. Since LifeKeeper uses the `snmptrap` utility to generate the traps, the `snmptrap` command is required to be installed on the node that will generate SNMP notification.

The `snmptrap` utility is provided by the following packages:

RHEL 5 or later and supported operating systems – `net-snmp-utils`

SLES 11 or later – `net-snmp`

In older versions of the `snmp` implementation (prior to 4.1) where the `defCommunity` directive is not supported, the traps will be sent using the “public” community string.

It is not necessary to have an SNMP agent `snmpd` running on the LifeKeeper node.

The configuration of a trap handler on the network management console and its response to trap messages is beyond the scope of this LifeKeeper feature. See the documentation associated with your system management tool for related instructions.

### Configuration Tasks

The following tasks must be performed to set up LifeKeeper SNMP Event Forwarding. All but the last task must be repeated on each node in the LifeKeeper cluster that will be generating SNMP trap messages.

1. Ensure that the `snmptrap` utility is available as noted above.
2. Specify the network management node to which the SNMP traps will be sent. This can be done either by command line or by editing the `/etc/default/LifeKeeper` file. You must specify the IP address rather than domain name to avoid DNS issues.
  - By command line, use the `lk_configsnmp` (see the `lk_configsnmp(1m)` man page for details). This utility will only accept IP addresses.
  - Or, edit the defaults file `/etc/default/LifeKeeper` to add the IP address. Find the entry `LK_TRAP_MGR=` and insert one or more IP addresses (separated by commas) to the right of “=” (no white space before or after “=” or commas).
3. If you are using an older version of the `snmp` implementation that does not support the `defCommunity` directive, skip this step. Traps will be sent using the “public” community string. Otherwise, do the following:

Specify a default community in `/usr/share/snmp/snmp.conf`. If this file does not exist, create it using sufficiently secure permissions. Add the directive “`defCommunity`” with a value. This specifies the SNMP version 2c community string to use when sending traps. For example, add a line like this:

```
defCommunity myCommunityString
```

Refer to the `snmp.conf` man page (delivered with the `snmp` package) for more information about this configuration file.

4. Perform whatever configuration steps are needed on the remote management console to detect and respond to the incoming trap OIDs from LifeKeeper events. If the management node is a Linux server, the minimum that you would need to do to begin verification of this feature would be to start the `snmptrapd` daemon with the `-f -Lo` option (print the messages to `stdout`).

## Verifying the Configuration

To verify that the configuration is working, initiate a LifeKeeper action (for example, start or stop LifeKeeper, or bring a resource in-service manually using the LifeKeeper GUI). Verify that the trap message was received at the management console. If a trap is not received, inspect the appropriate log files on the management system, and follow the normal troubleshooting practices provided with the management software. The LifeKeeper log can be inspected to determine if there was a problem sending the trap message. See [SNMP Troubleshooting](#) for more information.

## Disabling SNMP Event Forwarding

To disable the generation of SNMP traps by LifeKeeper, simply remove the assignment of an IP address from the `LK_TRAP_MGR` environment variable in the file `/etc/default/LifeKeeper`. This can be accomplished using the `lk_configsnmp` utility from the command line with the “`disable`” option (see the `lk_configsnmp(1M)` page for an example). Or, edit `/etc/default/LifeKeeper` and change the entry for `LK_TRAP_MGR` to `LK_TRAP_MGR=` (or remove the line entirely). This must be done on each node that should be disabled from sending trap messages.

## 5.4.2.2.3. SNMP Troubleshooting

---

Following are some possible problems and solutions related to SNMP Event Forwarding. For specific error messages, see the [LifeKeeper Message Catalog](#).

**Problem:** No SNMP trap messages are sent from LifeKeeper.

**Solution:** Verify that the `snmptrap` utility is installed on the system (it is usually located in `/usr/bin`). If it is not installed, install the appropriate `snmp` package (see [Prerequisites](#)). If it is installed in some other location, edit the `PATH` variable in the file `/etc/default/LifeKeeper` and add the appropriate path.

**Problem:** No SNMP error messages are logged and SNMP trap messages do not appear to be sent from a LifeKeeper server.

**Solution:** Check to see if `LK_TRAP_MGR` is set to the IP address of the network management server that will receive the traps. From the command line, use the `lk_configsnmp` utility with the “query” option to verify the setting (See the `lk_configsnmp(1M)` man page for an example.) Or, search for the entry for `LK_TRAP_MGR` in the file `/etc/default/LifeKeeper`. This variable must be set on each LifeKeeper node that will generate SNMP trap messages.

## 5.4.2.3. LifeKeeper Event Email Notification

---

[Overview](#)

[Configuration](#)

[Troubleshooting](#)

## 5.4.2.3.1. Overview of LifeKeeper Event Email Notification

LifeKeeper Event Email Notification is a mechanism by which one or more users may receive email notices when certain events occur in a LifeKeeper cluster. LifeKeeper has an event notification mechanism for registering applications that wish to be notified of specific events or alarms (see the `sendevent(5)` man page). LifeKeeper can be easily enabled to send email notification of key LifeKeeper events to a selected set of users wishing to monitor LifeKeeper activity.

By default, LifeKeeper Event Email Notification is disabled. Enabling this feature requires setting the `LK_NOTIFY_ALIAS` environment variable defined in `/etc/default/LifeKeeper`. The `LK_NOTIFY_ALIAS` environment variable can be set to a single email address or alias, or it can contain multiple addresses or aliases separated by commas. To set `LK_NOTIFY_ALIAS` either run `lk_confignotifyalias` (See the `lk_confignotifyalias(1M)` man page for an example) from the command line and supply the address or list of addresses that should receive email when an event occurs or edit the defaults file `/etc/default/LifeKeeper` to add the email address or address list. Search for the entry `LK_NOTIFY_ALIAS=` and insert the address or address list separated by commas. Repeat this action on all nodes in the cluster that need to send email for the selected LifeKeeper events.

To disable Email Notification, either run `lk_confignotifyalias` (See the `lk_confignotifyalias(1M)` man page for an example) with the `—disable` argument or edit the defaults file `/etc/default/LifeKeeper` and remove the setting of `LK_NOTIFY_ALIAS` (change the line to `LK_NOTIFY_ALIAS=`).

### LifeKeeper Events Generating Email

The following LifeKeeper events will generate email notices when `LK_NOTIFY_ALIAS` is set.

LifeKeeper Event	Event Description
LifeKeeper Startup Complete	Sent from a node when LifeKeeper is started on that node.
LifeKeeper Shutdown Initiated	Sent from a node beginning LifeKeeper shutdown.
LifeKeeper Shutdown Complete	Sent from a node completing LifeKeeper shutdown.
LifeKeeper Manual Switchover Initiated on Server	Sent from the node from which a manual switchover was requested.
LifeKeeper Manual Switchover Complete – recovered list	Sent from the node where the manual switchover was completed listing the resource successfully recovered.

LifeKeeper Manual Switchover Complete – failed list	Sent from the node where the manual switchover was completed listing the resource that failed to successfully switchover.
LifeKeeper Node Failure Detected	Sent from each node within the cluster when a node in that cluster fails.
LifeKeeper Node Recovery Complete for Server – recovered list	Sent from each node within the cluster that has recovered resources from the failed node listing the resource successfully recovered.
LifeKeeper Node Recovery Complete for Server – failed list	Sent from each node within the cluster that has failed to recover resources from the failed node listing the resource that failed to successfully recover.
LifeKeeper Resource Recovery Initiated	Sent from a node recovering a resource; a “Resource Recovery Complete” or “Resource Recovery Failed” message always follows to indicate whether the recovery was completed or failed.
LifeKeeper Resource Recovery Complete	Sent from the node that issued a “Resource Recovery Initiated” message when the recovery of the resource is completed listing the resource successfully recovered.
LifeKeeper Resource Recovery Failed	Sent from the node that issued a “Resource Recovery Initiated” message if the resource fails to come into service listing the resource that failed to successfully recover.
LifeKeeper Communications Path Up	A communications path to a node has become operational.
LifeKeeper Communications Path Down	A communications path to a node has gone down.
LifeKeeper <Node Monitoring> Failure Detected	Sent from a node where a failure was detected with Node Monitoring of the Standby Node Health Check. Detected failure is described in <Node Monitoring>.
LifeKeeper <OSUquickCheck> Failure Detected	Sent from a node where a failure was detected with OSU resource monitoring of the Standby Node Health Check. Tag name of the resource where the failure was detected is described in <OSUquickCheck>.

## 5.4.2.3.2. Configuring LifeKeeper Event Email Notification

---

### Prerequisites

The Event Email Notification feature is included as part of the LifeKeeper core functionality and does not require additional LifeKeeper packages to be installed. It does require that email software be installed and configured on each LifeKeeper node that will generate email notification of LifeKeeper events. LifeKeeper uses the mail utility, usually installed by the mailx package to send notifications.

The configuration of email is beyond the scope of this LifeKeeper feature. By default, LifeKeeper Event Email Notification is disabled.

### Configuration Tasks

The following tasks must be performed to set up LifeKeeper Event Email Notification.

1. Ensure that the mail utility is available as noted above.
2. Identify the user or users that will receive email notices of LifeKeeper events and set `LK_NOTIFY_ALIAS` in the LifeKeeper defaults file `/etc/default/LifeKeeper`. This can be done either from the command line or by editing the file `/etc/default/LifeKeeper` and specifying the email address or alias or the list of email addresses or aliases that should receive notification.
  - From the command line, use the `lk_confignotifyalias` utility (see the `lk_confignotifyalias(1M)` man page for details). This utility will only accept email addresses or aliases separated by commas.
  - Or, edit the defaults file `/etc/default/LifeKeeper` to add the email address or alias. Search for the entry `LK_NOTIFY_ALIAS=` and insert the email address or alias (single or list separated by commas) to the right of the `=` (no white space around the `=`).

### Verifying the Configuration

To verify that the configuration is working, initiate a LifeKeeper action (for example, [start](#) or [stop](#) LifeKeeper or bring a resource in-service manually using the LifeKeeper GUI). Verify that an email message was received by the users specified in `LK_NOTIFY_ALIAS` in the file `/etc/default/LifeKeeper` and a message was logged in the LifeKeeper log file. If an email message has not been received, follow your normal debugging procedures for email failures. The LifeKeeper log can be inspected to determine if there was a problem sending the email message. See [Email Notification Troubleshooting](#) for more information.

## Disabling Event Email Notification

To disable the generation of email notices by LifeKeeper, simply remove the assignment of an email address or alias from the `LK_NOTIFY_ALIAS` environment variable in the file `/etc/default/LifeKeeper`. This can be accomplished using the `lk_confignotifyalias` utility from the command line with the “*—disable*” option (see the `lk_confignotifyalias(1M)` page for an example). Or, edit `/etc/default/LifeKeeper` and change the entry for `LK_NOTIFY_ALIAS` to `LK_NOTIFY_ALIAS=`. This must be done on each node that should be disabled from sending email messages.

## 5.4.2.3.3. Email Notification Troubleshooting

Following are some possible problems and solutions related to email notification of LifeKeeper events. For specific error messages, see the [LifeKeeper Message Catalog](#).

**Problem:** No email messages are received from LifeKeeper.

**Solution:** Verify that the mail utility is installed on the system (it is usually located in `/bin/mail`). If it is not installed, install the `mailx` package. If it is installed in some other location, edit the `PATH` variable in the file `/etc/default/LifeKeeper` and add the path to the mail utility.

**Problem:** No email messages are received from LifeKeeper.

**Solution:** Check the email configuration and ensure email messages have not be queued for delivery indicating a possible email configuration problem. Also ensure that the email address or addresses specified in `LK_NOTIFY_ALIAS` are valid and are separated by a comma.

**Problem:** The log file has a “mail returned” error message.

**Solution:** There was some problem invoking or sending mail for a LifeKeeper event, such as a “node failure”, as the mail command return the error X. Verify the mail configuration and that `LK_NOTIFY_ALIAS` contains a valid email address or list of addresses separated by a comma and ensure that email can be sent to those addresses by sending email from the command line using the email recipient format defined in `LK_NOTIFY_ALIAS`.

**Problem:** No messages, success or failure, are logged and the user or users designated to receive email have not received any mail when a LifeKeeper Event has occurred, such as a node failure.

**Solution:** Check to see if `LK_NOTIFY_ALIAS` is, in fact, set to an email address or list of addresses separated by commas. From the command line, use the `lk_confignotifyalias` utility with the “`—query`” option to verify the setting (See the `lk_confignotifyalias(1M)` man page for an example.) Or, search for the entry `LK_NOTIFY_ALIAS` in the file `/etc/default/LifeKeeper`. This variable must be set on each LifeKeeper node that will generate email notification messages. Also, see the [Overview of LifeKeeper Event Email Notification](#) to see if the LifeKeeper event generates an email message (not all events generate email messages).

## 5.4.2.4. Optional Configuration Tasks

---

[Confirm Failover and Block Resource Failover Settings](#)

[Setting Server Shutdown Strategy](#)

[Tuning the LifeKeeper Heartbeat](#)

[Using Custom Certificates](#)

## 5.4.2.4.1. Confirm Failover and Block Resource Failover Settings

---

Normally, LifeKeeper will automatically switch operations to a backup node when a node failure or a resource failure occurs. However, depending on the environment, requiring manual confirmation by a system administrator may be desirable, instead of an automatic failover recovery initiated by LifeKeeper. In these cases, the **Confirm Failover** or **Block Resource Failover** settings are available. By using these functions, automatic failover can be blocked and a time to wait for failover can be set when a resource failure or a node failure occurs.

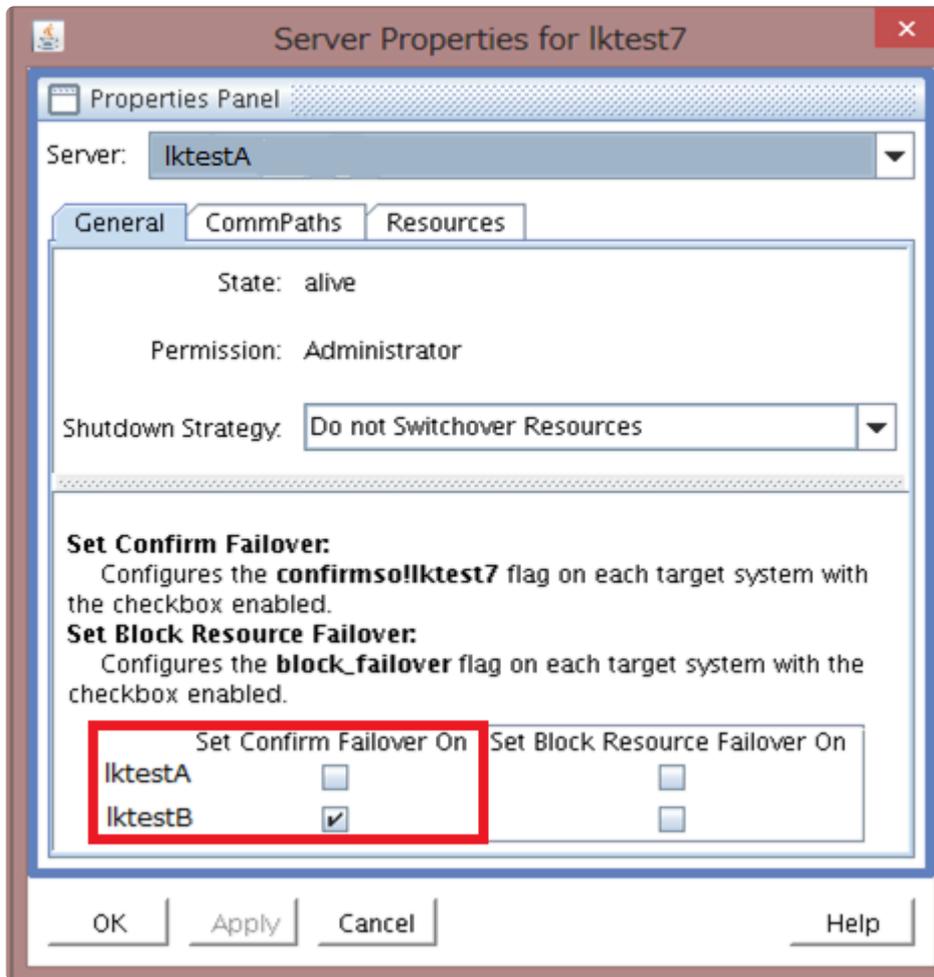
Set **Confirm Failover** or **Block Resource Failover** in your LifeKeeper environment after carefully reading the descriptions, examples, and considerations below. These settings are available from the **Server Properties** dialog of the GUI and via the command line of LifeKeeper.

### Set Confirm Failover On

When a failover occurs because a node in the LifeKeeper cluster fails (**Note:** a node failure is identified by a failure of all LifeKeeper communication paths to that system), the time to wait before LifeKeeper switches resources to a backup node can be set with the **Confirm Failover** setting (see the discussion on the CONFIRMSOTO variable later in this document). Also, a user can decide whether to automatically switch to the backup node or not after the time to wait expires (see the discussion of the CONFIRMSODEF variable later in this document).

 **Note:** Set **Confirm Failover On** actions are only available when a node failure occurs. It is not available for resource failures where one or more communications paths are still active.

To enable the **Confirm Failover** setting via the GUI, use the General tab of the **Server Properties** dialog. An example of the General tab for Server Properties is shown below. The part outlined in red on the screen addresses the **Confirm Failover** setting.



**Note:** This setting is only available for users with Administrator permission for LifeKeeper.

In this example, the setting is seen from the host named lktestA. The part outlined in red on the screen is used for this setting **Confirm Failover**. The node names for the HA cluster are displayed vertically. In this example, the standby node for lktestA is lktestB.

The screen shows the configuration status for server lktestA, with the checkbox for lktestB set. In this case, the confirm failover flag is created on lktestB. When a failover from lktestA to lktestB is executed, the confirmation process for executing a failover occurs on lktestB. This process includes checking the default action to take based on the CONFIRMSODEF variable setting (see the discussion later in this document) and how long to wait before taking that action based on the CONFIRMSOTO variable setting.

The **Confirm Failover** flag creation status can be checked via the command line. When the checkbox for lktestB is set on the host named lktestA, the **Confirm Failover** flag is created on lktestB. (**NOTE:** In this example, the flag is not created on lktestA, only on lktestB.) An example of the command line output is below.

```
[root@lktestB~]# /opt/LifeKeeper/bin/flg_list
```

```
confirmso!lktestA
```

The “confirmso!lktestA” output is the result of the flg\_list command, and indicates that the **Confirm Failover** flag is set on node lktestB to confirm lktestA failures.

When failover occurs with the confirmso flag, the following messages are recorded in the LifeKeeper log file.

```
INFO:lcd.recover:::004113:
```

```
chk_man_interv: Flag confirmso!hostname is set, issuing confirmso event and waiting for switchover instruction.
```

```
NOTIFY:event.confirmso:::010464:
```

```
LifeKeeper: FAILOVER RECOVERY OF MACHINE lktestA requires manual confirmation! Execute '/opt/LifeKeeper/bin/lk_confirmso -y -s lktestA ' to allow this failover, or execute '/opt/LifeKeeper/bin/lk_confirmso -n -s lktestA' to prevent it. If no instruction is provided, LifeKeeper will timeout in 600 seconds and the failover will be allowed to proceed.
```

Execute one of the following commands to confirm the failover:

**To proceed with the failover:**

```
# /opt/LifeKeeper/bin/lk_confirmso -y -s hostname
```

**To block the failover:**

```
# /opt/LifeKeeper/bin/lk_confirmso -n -s hostname
```

The host name that is specified when executing the command is the host name listed in **Confirm Failover** flag which for this example would be lktestA. Execute the command by referring to the example commands provided in the Log output.

In the case where the set time to wait is exceeded, the default failover action is executed (allow failover or block failover). The default failover action is determined by the CONFIRMSODEF variable (see discussion later in this document).

The following message is output to the LifeKeeper log when the timeout expires.

```
lcdrecover[xxxx]: INFO:lcd.recover:::004408:chk_man_interv: Timed out waiting for instruction, using default CONFIRMSODEF value 0.
```

The LifeKeeper operation when the time to wait is exceeded is controlled by the setting of the variable “CONFIRMSODEF” which is set in the /etc/default/LifeKeeper file with a value of “1” or “0”. A value of “0” is set by default, and this indicates that the failover will proceed when the time to wait is exceeded. If the value is set to a “1”, the failover is blocked when the time to wait is exceeded.

The time to wait for confirmation of a failover can be changed by adjusting the value of the CONFIRMSOTO variable in the /etc/default/LifeKeeper file. The value of the variable specifies the

number of seconds to wait for a manual confirmation from the user before proceeding or blocking the failover as determined by the value of the “CONFIRMSODEF” variable (see above).

Restarting LifeKeeper or rebooting the OS **is not** required for changes to these variables to take effect. If the value of CONFIRMSOTO is set to 0 seconds, then the operation based on the CONFIRMSODEF setting will occur immediately.

## When to Select [Confirm Failover] Setting

This setting is used for Disaster Recovery or WAN configurations in the environment which the communication paths are not redundant.

- Open the **Properties** page from one server and then select the server that you want the **Confirm Failover flag** to be set on.

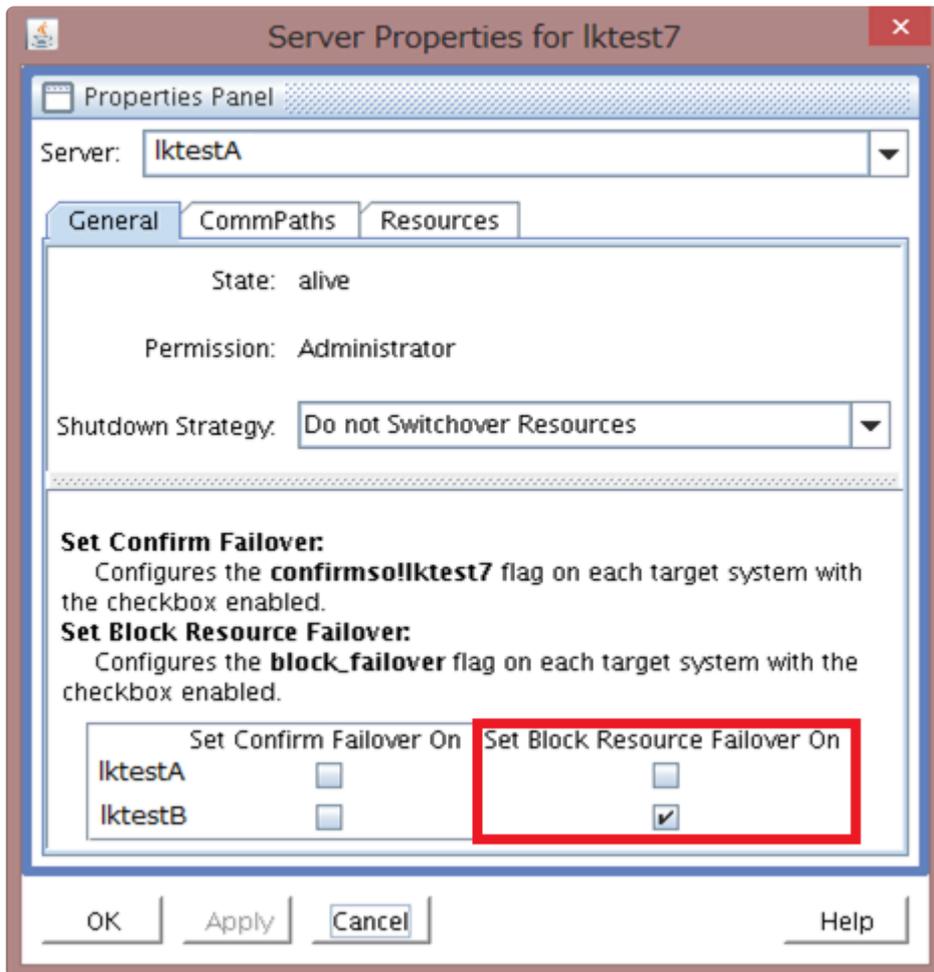
## Block Resource Failover On

The **Block Resource Failover On** setting blocks all resource transfers due to a resource failure from the given system.

 **Note:** The **Block Resource Failover On** setting has no effect on the failover processing when a node failure occurs. This setting only blocks a failover attempt when a local resource recovery fails and attempts to transfer the resource to another node in the cluster.

By default, the recovery of resource failures in a local system (local recovery) is performed when a resource failure is detected. When the local recovery has failed or is not enabled, a failover is initiated to the next highest priority standby node defined for the resource. The Block Resource Failover On setting will prevent this failover attempt.

To enable the **Block Resource Failover On** setting by the GUI, use the General tab of the Server Properties. An example of the General tab for Server Properties is below. The part outlined in red on the screen addresses the setting **Block Resource Failover On**.



**Note:** This setting is only available for users with Administrator permission for LifeKeeper.

In this example, the setting is seen from the host named lktestA. The part outlined in red on the screen is used for setting **Block Resource Failover**. The node names for the HA cluster are displayed vertically here. In this example, the standby node for lktestA is lktestB.

In this case, the Block Resource Failover flag is created on lktestB. The “block\_failover” flag can be verified on the command line by executing the flg\_list command. An example of the output is below.

```
[root@lktestB~]# /opt/LifeKeeper/bin/flg_list
```

```
block_failover
```

When the block\_failover flag is set, the failover to other node (lktestA) is blocked when a resource failure occurs on lktestB.

**The block\_failover flag prevents resource failovers from occurring on the node where the flag is set.** The following log message is output to the LifeKeeper log when the failover is blocked by this setting.

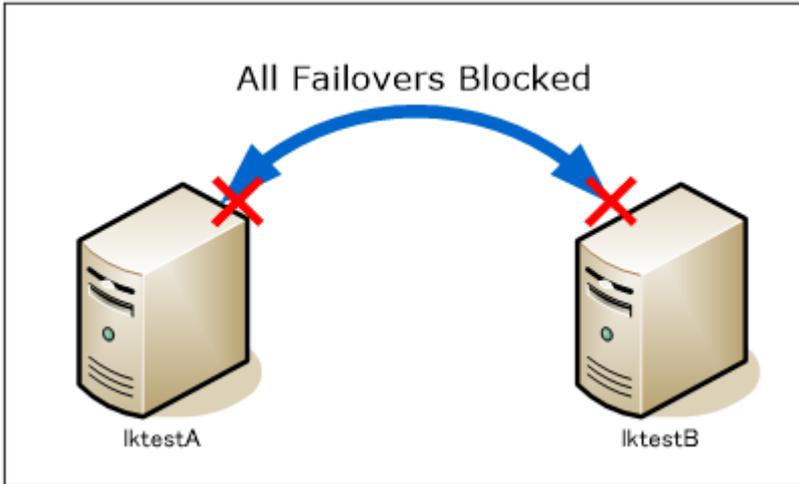
```
ERROR:lcd.recover:::004787:Failover is blocked by current settings. MANUAL INTERVENTION IS
```

REQUIRED

## Configuration examples

Some configuration examples are described below.

### Block All Automatic Failovers



In this example, the failover is blocked when a node failure or a resource failure is detected on either IktestA or IktestB. Use the Confirm failover and Block Resource failover settings for this. LKCLI allows you to configure the setting with just one command. Refer to [lkcli server block-all-failovers](#) for more information. The configuration example with the GUI is as follows.

1. Select IktestA and view **Server Properties**. On the General tab, check the “**Set Confirm Failover On**” box for IktestB and the “**Set Block Resource Failover On**” box for both IktestA and IktestB. The setting status in the GUI is below.

The configuration of Server Properties for IktestA as displayed in the GUI once set.

	Set Confirm Failover On	Set Block Resource Failover On
IktestA	(Not checked)	✓
IktestB	✓	✓

\*When viewing the Server Properties in the GUI the node name can be found near the top of the properties panel display.

2. Select IktestB and view **Server Properties**.

On General tab, check “**Set Confirm Failover On**” box for IktestA. The “**Set Block Resource Failover On**” property will already be set based on the actions taken in step 1.

The configuration of Server Properties for IktestB as displayed in the GUI once set.

	Set Confirm Failover On	Set Block Resource Failover On
lktestB	(Not checked)	✓
lktestA	✓	✓

\*When viewing the Server Properties in the GUI the node name can be found near the top of the properties panel display.

After completing these steps, confirm that the “confirmso!hostname” and “ block\_failover” flags are set for each node using the flg\_list command. For the confirmso flag, verify that the host name for which failover confirmation is to be performed is listed as part of the contents of the flag name (to block failover from lktestB on lktestA, lktestA should list lktestB in the contents of the confirmso flag name on lktestA, see the table below).

	Confirm Failover Flag	Block Resource Failover Flag
lktestA	confirmso!lktestB	block_failover
lktestB	confirmso!lktestA	block_failover

- Set the values for “CONFIRMSOTO” and “CONFIRMSODEF” in /etc/default/LifeKeeper on each node. (Restarting LifeKeeper or rebooting the OS is not required.)

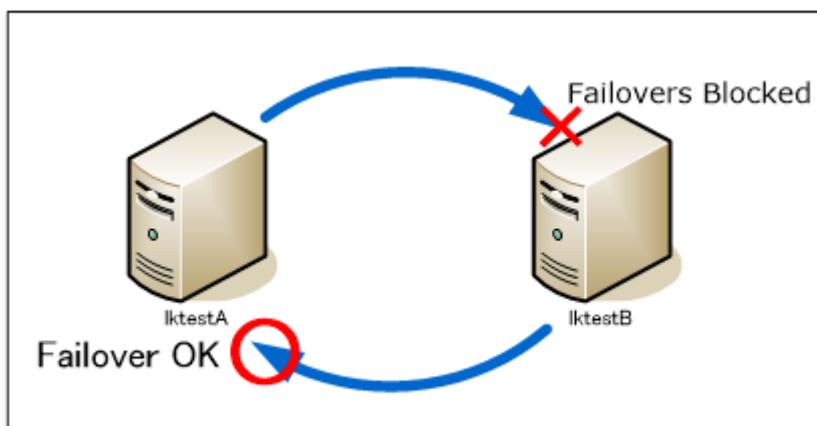
CONFIRMSODEF=1

CONFIRMSOTO=0

When setting the time to wait value, it is specified in seconds via CONFIRMSOTO. For the default action to be taken on failover, the CONFIRMSODEF setting must be either 0 (failover is executed) or 1 (failover is blocked).

With the above settings any node failure will be immediately blocked without any operator intervention.

### Block Failovers in One Direction



In this example, the failover to lktestB is blocked when a node failure or a resource failure is detected on

IktestA. On the contrary, the failover to IktestA is allowed when a node failure or a resource failure is detected on IktestB.

1. Select IktestA and view **Server Properties**.
2. On the **General** tab, check “**Set Confirm Failover On**” box for IktestB and the “**Set Block Resource Failover On**” for IktestA.

The configuration of Server Properties for IktestA as displayed in the GUI once set.

	<b>Set Confirm Failover On</b>	<b>Set Block Resource Failover On</b>
IktestA	(Not checked)	✓
IktestB	✓	(Not checked)

3. Select IktestB and view **Server Properties**.

In the **General** tab, the “**Set Block Resource Failover On**” box for IktestA should already be set (from the action taken on IktestA).

The configuration of Server Properties for IktestB as displayed in the GUI once set.

	<b>Set Confirm Failover On</b>	<b>Set Block Resource Failover On</b>
IktestB	(Not checked)	(Not checked)
IktestA	(Not checked)	✓

\*In the GUI, the local host name is listed first.

For this configuration, verify that the “confirmso!IktestA” flag is set on IktestB (no confirmso flag should be set on IktestA) and the block failover flag is set on IktestA.

	<b>Confirm Failover Flag</b>	<b>Block Resource Failover Flag</b>
IktestA	N/A	block_failover
IktestB	confirmso!IktestA	N/A

4. Set the values for “CONFIRMSODEF” and “CONFIRMSOTO” in /etc/default/LifeKeeper on IktestB.

CONFIRMSODEF=1

CONFIRMSOTO=0

For this configuration resource and machine failovers from IktestA to IktestB are blocked. Resource and machine failovers from IktestB to IktestA are allowed.

## 5.4.2.4.2. Setting Server Shutdown Strategy

The Shutdown Strategy is a LifeKeeper configuration option that governs whether or not resources are switched over to a backup server when a server is shut down. The options are:

Do Not Switch Over Resources (default)	LifeKeeper will not bring resources in service on a backup server during an orderly shutdown.
Switch Over Resources	LifeKeeper will bring resources in service on a backup server during an orderly shutdown.

The Shutdown Strategy is set by default to “Do Not Switch Over Resources.” You should decide which strategy you want to use on each server in the cluster, and if you wish, change the Shutdown Strategy to “Switch Over Resources”.

For each server in the cluster:

1. On the [Edit Menu](#), point to **Server** and then click **Properties**.
2. Select the server to be modified.
3. On the [General Tab](#) of the **Server Properties** dialog, select the **Shutdown Strategy**.

 **Note:** The LifeKeeper process must be running during an orderly shutdown for the Shutdown Strategy to have an effect.

## 5.4.2.4.3. Tuning the LifeKeeper Heartbeat

### Overview of the Tunable Heartbeat

The LifeKeeper heartbeat is the signal sent between LifeKeeper servers over the communications path(s) to ensure each server is “alive”. There are two aspects of the heartbeat that determine how quickly LifeKeeper detects a failure:

- **Interval:** the time interval between heartbeats signal sent (unit is second). Failing to receive the LCM signal, which includes heartbeat signal, from another server within the interval time is determined as a missed heartbeat.
- **Number of Heartbeats:** the consecutive number of heartbeats by which the communications path is determined as dead, triggering a failover.

The heartbeat values are specified by two tunables in the LifeKeeper defaults file `/etc/default/LifeKeeper`. These tunables can be changed if you wish LifeKeeper to detect a server failure sooner than it would using the default values:

- LCMHBEATTIME (interval)
- LCMNUMHBEATS (number of heartbeats)

The following table summarizes the defaults and minimum values for the tunables over both TCP and TTY heartbeats. The interval for a TTY communications path cannot be set below 2 seconds because of the slower nature of the medium.

Tunable	Default Value	Minimum Value
LCMHBEATTIME	5	1 (TCP) 2 (TTY)
LCMNUMHBEATS	3	2 (TCP or TTY)

**!** The values for both tunables **MUST** be the **SAME** on all servers in the cluster.

### Example

Consider a LifeKeeper cluster in which both intervals are set to the default values. LifeKeeper sends a heartbeat between servers every 5 seconds. If a communications problem causes the heartbeat to skip two beats, but it resumes on third heartbeat, LifeKeeper takes no action. However, if the communications path remains dead for 3 beats, LifeKeeper will label that communications path as dead, but will initiate a failover only if the redundant communications path is also dead.

## Configuring the Heartbeat

You must manually edit file `/etc/default/LifeKeeper` to add the tunable and its associated value. Normally, the defaults file contains no entry for these tunables; you simply append the following lines with the desired value as follows:

```
LCMHBEATTIME=x
```

```
LCMNUMHBEATS=y
```

If you assign the value to a number below the minimum value, LifeKeeper will ignore that value and use the minimum value instead.

## Configuration Considerations

- If you wish to set the interval at less than 5 seconds, then you should ensure that the communications path is configured on a private network, since values lower than 5 seconds create a high risk of false failovers due to network interruptions.
- Testing has shown that setting the number of heartbeats to less than 2 creates a high risk of false failovers. This is why the value has been restricted to 2 or higher.
- In order to avoid false failovers, both the interval and heartbeat count values must be the same on all servers in the cluster. For this reason, LifeKeeper must be stopped on both servers before modifying these values. After starting LifeKeeper, you can use the command `/opt/LifeKeeper/bin/lkstop -f` to edit the heartbeat settings while the application is protected. This command stops LifeKeeper but does not stop the protected application.
- LifeKeeper does not impose an upper limit for the `LCMHBEATTIME` and `LCMNUMHBEATS` values. But setting these values at a very high number can effectively disable LifeKeeper's ability to detect a failure. For instance, setting both values to 25 would instruct LifeKeeper to wait 625 seconds (over 10 minutes) to detect a server failure, which may be enough time for the server to re-boot and re-join the cluster.



**Note:** If you are using both TTY and TCP communications paths, the value for each tunable applies to both communications paths. The only exception is if the interval value is below 2, which is the minimum for a TTY communications path.

For example, suppose you specify the lowest values allowed by LifeKeeper in order to detect failure as quickly as possible:

```
LCMHBEATTIME=1
```

```
LCMNUMHBEATS=2
```

LifeKeeper will use a 1 second interval for the TCP communications path, and a 2 second interval for

TTY. In the case of a server failure, LifeKeeper will detect the TCP failure first because its interval is shorter (2 heartbeats that are 1 second apart), but then will do nothing until it detects the TTY failure, which will be after 2 heartbeats that are 2 seconds apart.

## 5.4.2.4.4. Using Certificates with the LifeKeeper API

The LifeKeeper API uses SSL/TLS to communicate between different systems. By default, the product is installed with default certificates that provide some assurance of identity between nodes. This document explains how to replace these default certificates with certificates created by your own Certificate Authority (CA).

\* **Note:** Currently, this API is only partially used and is reserved for internal use only but may be opened up to customer and third party usage in a future release.

\* **Note:** Normal LifeKeeper communication does not use these certificates.

### How Certificates Are Used

In cases where SSL/TLS is used for communications between LifeKeeper servers to protect the data being transferred, a certificate is provided by systems to identify themselves. The systems also use a CA certificate to verify the certificate that is presented to them over the SSL connection.

Three certificates are involved:

- `/opt/LifeKeeper/etc/certs/LK4LinuxValidNode.pem` (server certificate)
- `/opt/LifeKeeper/etc/certs/LK4LinuxValidClient.pem` (client certificate)
- `/opt/LifeKeeper/etc/certs/LKCA.pem` (certificate authority)

The first two certificates must be signed by the CA certificate to satisfy the verification performed by the servers. Note that the common name of the certificates is not verified, only that the certificates are signed by the CA.

### Using Your Own Certificates

In some installations, it may be necessary to replace the default certificates with certificates that are created by an organization's internal or commercial CA. If this is necessary, replace the three certificates listed above with new certificates *using the same certificate file names*. These certificates are of the PEM type. The `LK4LinuxValidNode.pem` and `LK4LinuxValidClient.pem` each contain both their respective key and certificate. The `LK4LinuxValidNode.pem` certificate is a *server* type certificate. `LK4LinuxValidClient.pem` is a *client* type certificate.

If the default certificates are replaced, LifeKeeper will need to be restarted to reflect the changes. If the certificates are misconfigured, `SIOS-lighttpd` daemon will not start successfully and errors will be

received in the LifeKeeper log file. If problems arise, refer to this log file to see the full command that should be run.

## Updating Expired Certificates

- \* Certificates in older versions prior to 9.5.2 **will expire** in September 2021 and should be replaced with updated certificates. These updated certificates can be found via the [SIOS Support Portal](#). When replacing certificates, **please replace all the certificates listed above with new certificates using the same certificate file name.**

## 5.4.2.5. Linux Configuration

<p><b>Operating System</b></p>	<p>The default operating system must be installed to ensure that all required packages are installed. The minimal operating system install does not contain all of the required packages, and therefore, cannot be used with LifeKeeper.</p>														
<p><b>Kernel updates</b></p>	<p>In order to provide the highest level of availability for a LifeKeeper cluster, the kernel version used on a system is very important. The table below lists each supported distribution and version with the kernel that has passed LifeKeeper certification testing.</p> <p><b>Note:</b> When upgrading the kernel, you may need to rerun the setup script (<code>./setup</code>) using the installation image. If DataKeeper fails to start after upgrading the kernel, run the setup script to install the appropriate kernel module.</p> <table border="1" data-bbox="323 828 1455 2067"> <thead> <tr> <th data-bbox="323 828 802 936">Distribution/Version</th> <th data-bbox="802 828 1038 936">Supported Version</th> <th data-bbox="1038 828 1455 936">Supported Kernels</th> </tr> </thead> <tbody> <tr> <td data-bbox="323 936 802 1547"> <p>Red Hat Enterprise Linux for AMD64/EM64T</p> <p>(*DataKeeper asynchronous mirrors are not supported on some kernels on RHEL 7.4-7.6. Click <a href="#">here for full details.</a>)</p> <p>(Some kernel versions do not support asynchronous mode. Please see <a href="#">Known Issues and Restrictions</a> for details)</p> </td> <td data-bbox="802 936 1038 1547"> <p>7</p> <p>7.1</p> <p>7.2</p> <p>7.3</p> <p>7.4*</p> <p>7.5*</p> <p>7.6*</p> <p>7.7</p> <p>7.8</p> <p>7.9</p> </td> <td data-bbox="1038 936 1455 1547"> <p>3.10.0-123.el7</p> <p>3.10.0-229.el7</p> <p>3.10.0-327.el7</p> <p>3.10.0-514.el7</p> <p>3.10.0-693.el7</p> <p>3.10.0-862.el7</p> <p>3.10.0-957.el7</p> <p>3.10.0-1062.el7</p> <p>3.10.0-1127.el7</p> <p>3.10.0-1160.el7</p> </td> </tr> <tr> <td data-bbox="323 1547 802 1895"> <p>Red Hat Enterprise Linux for AMD64/EM64T</p> <p>(Upgrading from RHEL 7 to RHEL 8 is not supported)</p> </td> <td data-bbox="802 1547 1038 1895"> <p>8.0</p> <p>8.1</p> <p>8.2</p> <p>8.3</p> <p>8.4</p> <p>8.5</p> </td> <td data-bbox="1038 1547 1455 1895"> <p>4.18.0-80.el8.x86_64</p> <p>4.18.0-147.el8.x86_64</p> <p>4.18.0-193.el8.x86_64</p> <p>4.18.0-240.el8.x86_64</p> <p>4.18.0-305.el8.x86_64</p> <p>4.18.0-348.el8.x86_64</p> </td> </tr> <tr> <td data-bbox="323 1895 802 2067"> <p>SUSE Linux Enterprise Server 12 for x86_64</p> </td> <td data-bbox="802 1895 1038 2067"> <p>12 SP1</p> <p>12 SP2</p> </td> <td data-bbox="1038 1895 1455 2067"> <p>3.12.49-11.1</p> <p>4.4.21-69.1</p> </td> </tr> </tbody> </table>			Distribution/Version	Supported Version	Supported Kernels	<p>Red Hat Enterprise Linux for AMD64/EM64T</p> <p>(*DataKeeper asynchronous mirrors are not supported on some kernels on RHEL 7.4-7.6. Click <a href="#">here for full details.</a>)</p> <p>(Some kernel versions do not support asynchronous mode. Please see <a href="#">Known Issues and Restrictions</a> for details)</p>	<p>7</p> <p>7.1</p> <p>7.2</p> <p>7.3</p> <p>7.4*</p> <p>7.5*</p> <p>7.6*</p> <p>7.7</p> <p>7.8</p> <p>7.9</p>	<p>3.10.0-123.el7</p> <p>3.10.0-229.el7</p> <p>3.10.0-327.el7</p> <p>3.10.0-514.el7</p> <p>3.10.0-693.el7</p> <p>3.10.0-862.el7</p> <p>3.10.0-957.el7</p> <p>3.10.0-1062.el7</p> <p>3.10.0-1127.el7</p> <p>3.10.0-1160.el7</p>	<p>Red Hat Enterprise Linux for AMD64/EM64T</p> <p>(Upgrading from RHEL 7 to RHEL 8 is not supported)</p>	<p>8.0</p> <p>8.1</p> <p>8.2</p> <p>8.3</p> <p>8.4</p> <p>8.5</p>	<p>4.18.0-80.el8.x86_64</p> <p>4.18.0-147.el8.x86_64</p> <p>4.18.0-193.el8.x86_64</p> <p>4.18.0-240.el8.x86_64</p> <p>4.18.0-305.el8.x86_64</p> <p>4.18.0-348.el8.x86_64</p>	<p>SUSE Linux Enterprise Server 12 for x86_64</p>	<p>12 SP1</p> <p>12 SP2</p>	<p>3.12.49-11.1</p> <p>4.4.21-69.1</p>
Distribution/Version	Supported Version	Supported Kernels													
<p>Red Hat Enterprise Linux for AMD64/EM64T</p> <p>(*DataKeeper asynchronous mirrors are not supported on some kernels on RHEL 7.4-7.6. Click <a href="#">here for full details.</a>)</p> <p>(Some kernel versions do not support asynchronous mode. Please see <a href="#">Known Issues and Restrictions</a> for details)</p>	<p>7</p> <p>7.1</p> <p>7.2</p> <p>7.3</p> <p>7.4*</p> <p>7.5*</p> <p>7.6*</p> <p>7.7</p> <p>7.8</p> <p>7.9</p>	<p>3.10.0-123.el7</p> <p>3.10.0-229.el7</p> <p>3.10.0-327.el7</p> <p>3.10.0-514.el7</p> <p>3.10.0-693.el7</p> <p>3.10.0-862.el7</p> <p>3.10.0-957.el7</p> <p>3.10.0-1062.el7</p> <p>3.10.0-1127.el7</p> <p>3.10.0-1160.el7</p>													
<p>Red Hat Enterprise Linux for AMD64/EM64T</p> <p>(Upgrading from RHEL 7 to RHEL 8 is not supported)</p>	<p>8.0</p> <p>8.1</p> <p>8.2</p> <p>8.3</p> <p>8.4</p> <p>8.5</p>	<p>4.18.0-80.el8.x86_64</p> <p>4.18.0-147.el8.x86_64</p> <p>4.18.0-193.el8.x86_64</p> <p>4.18.0-240.el8.x86_64</p> <p>4.18.0-305.el8.x86_64</p> <p>4.18.0-348.el8.x86_64</p>													
<p>SUSE Linux Enterprise Server 12 for x86_64</p>	<p>12 SP1</p> <p>12 SP2</p>	<p>3.12.49-11.1</p> <p>4.4.21-69.1</p>													

Distribution/Version	Supported Version	Supported Kernels
<p>(The kernel should be updated to 4.4.82-6.9.1 for SP3.)</p>	<p>12 SP3* 12 SP4 12 SP5 * SLES12.0 is not supported.</p>	<p>4.4.82-6.9.1 4.12.14-94.41.1 4.12.14-120.1</p>
<p>SUSE Linux Enterprise Server 15 for x86_64</p> <p>(*DataKeeper cannot use disks with an odd sector size.) (*Upgrading from version SLES12 to SLES15 is not supported.) (*For SP3, upgrade the kernel to 5.3.18-59.5)</p>	<p>15* 15 SP1 15 SP2 15 SP3*</p>	<p>4.12.14-23.1 4.12.14-195.1 5.3.18-14.1 5.3.18-59.5</p>
<p>Oracle Linux</p> <p><a href="#">(*DataKeeper asynchronous mirrors are not supported on some kernels on OEL 7.4-7.6. Click here for full details.)</a></p> <p>(**The kernel should be updated to 5.4.17-2102.202.5 for UEK6.)</p>	<p>7 7.1 7.2 7.3 7.4* 7.5* 7.6* 7.7 7.8 7.9 UEK R3 UEK R4 UEK R5 UEK R6**</p>	<p>3.10.0-123.el7 3.10.0-229.el7 3.10.0-327.el7 3.10.0-514.el7 3.10.0-693.el7 3.10.0-862.el7 3.10.0-957.el7 3.10.0-1062.el7 3.10.0-1127.el7 3.10.0-1160.el7 3.8.13-16.2.1.el7uek 4.1.12-37.3.1.el7uek 4.14.35-1818.3.3.el7uek 5.4.17-2102.202.5.el7uek</p>
<p>Oracle Linux</p> <p>(Upgrading from OEL 7 to OEL 8 is not supported.)</p>	<p>8.0 8.1 8.2 8.3 8.4 8.5</p>	<p>4.18.0-80.el8 4.18.0-147.el8.x86_64 4.18.0-193.el8.x86_64 4.18.0-240.el8.x86_64 4.18.0-305.el8.x86_64 4.18.0-348.el8.x86_64</p>

	Distribution/Version	Supported Version	Supported Kernels
	(**The kernel should be updated to 5.4.17-2102.202.5 for UEK6.	UEK R6**	5.4.17-2102.202.5.el8uek
	CentOS  <a href="#">(*DataKeeper asynchronous mirrors are not supported on some kernels on CentOS 7.4-7.6. Click here for full details.)</a>	7 7.1 7.2 7.3 7.4* 7.5* 7.6* 7.7 7.8 7.9	3.10.0-123.el7 3.10.0-229.el7 3.10.0-327.el7 3.10.0-514.el7 3.10.0-693.el7 3.10.0-862.el7 3.10.0-957.el7 3.10.0-1062.el7 3.10.0-1127.el7 3.10.0-1160.el7
	CentOS  (Upgrading from CentOS 7 to CentOS 8 is not supported.)	8.0 8.1 8.2 8.3	4.18.0-80.el8 4.18.0-147.el8.x86_64 4.18.0-193.el8.x86_64 4.18.0-240.el8.x86_64
	<p><b>Note:</b> This list of supported distributions and kernels is for LifeKeeper only. You should also determine and adhere to the supported distributions and kernels for your server and storage hardware as specified by the manufacturer.</p>		
<p><b>LUN support</b></p>	<p>The Linux SCSI driver has several parameters that control which devices will be probed for Logical Units (LUNs):</p> <ul style="list-style-type: none"> <li>List of devices that <b>do not</b> support LUNs – this list of devices are known to NOT support LUNs, so the SCSI driver will not allow the probing of these devices for LUNs.</li> <li>List of devices that <b>do</b> support LUNs – this list of devices is known to support LUNs well, so always probe for LUNs.</li> <li>Probe all LUNs on each SCSI device – if a device is not found on either list, whether to probe or not. This parameter is configured by make config in the SCSI module section.</li> </ul>		

	<p>While most distributions (including SUSE) have the Probe all LUNs setting enabled by default, Red Hat has the setting disabled by default. External RAID controllers that are typically used in LifeKeeper configurations to protect data are frequently configured with multiple LUNs (Logical Units). To enable LUN support, this field must be selected and the kernel remade.</p> <p>To enable Probe all LUNs without rebuilding the kernel or modules, set the variable <code>max_scsi_luns</code> to 255 (which will cause the scan for up to 255 LUNs). To set the <code>max_scsi_luns</code> on a kernel where the scsi driver is a module (e.g. Red Hat), add the following entry to <code>/etc/modules.conf</code>, rebuild the initial ramdisk and reboot loading that ramdisk:</p> <pre>options scsi_mod max_scsi_luns=255</pre> <p>To set the <code>max_scsi_luns</code> on a kernel where the scsi driver is compiled into the kernel (e.g. SUSE), add the following entry to <code>/etc/lilo.conf</code>:</p> <pre>append="max_scsi_luns=255"</pre> <p><b>Note:</b> For some devices, scanning for 255 LUNs can have an adverse effect on boot performance (in particular devices with the <code>BLIST_SPARSELUN</code> defined). The Dell PV650F is an array where this has been experienced. To avoid this performance problem, set the <code>max_scsi_luns</code> to the maximum number of LUNs you have configured on your arrays such as 16 or 32. For example,</p> <pre>append="max_scsi_luns=16"</pre>
<p><b>Testing environment of channel bonding, network teaming</b></p>	<p>In LifeKeeper, we performed tests in the environment using the channel bonding or network teaming with the following settings:</p> <ul style="list-style-type: none"> <li>• Bonding policy in channel bonding             <ul style="list-style-type: none"> <li>+ balance-rr</li> <li>+ active-backup</li> </ul> </li> <li>• Runner in network teaming             <ul style="list-style-type: none"> <li>+ round-robin</li> <li>+ active-backup</li> </ul> </li> </ul>

## 5.4.2.6. Data Replication Configuration

Item	Description
<b>SIOS DataKeeper Feature/ Distribution Matrix</b>	SIOS DataKeeper uses Linux kernel functionality. Therefore, Linux kernel versions should be 2.6 and higher and it should be 2.6.27 and higher (backported to RHEL 5.4 or later for RHEL 5 series) when you use the bitmap merge operation as the minimum requirements. Refer to the <a href="#">Support Matrix</a> for available OS
<b>SIOS DataKeeper Documentation</b>	The documentation for <a href="#">SIOS DataKeeper</a> is located within the LifeKeeper Technical Documentation on the SIOS Technology Corp. Website.

## 5.4.2.7. Network Configuration

Item	Description
<b>IP Recovery Kit impact on routing table</b>	<p>LifeKeeper-protected IP addresses are implemented on Linux as logical interfaces. When a logical interface is configured on Linux, a route to the subnet associated with the logical interface is automatically added to the routing table, even if a route to that subnet already exists (for example, through the physical interface). This additional route to the subnet could possibly result in multiple routing-table entries to the same subnet.</p> <p>If an application is inspecting and attempting to verify the address from which incoming connections are made, the multiple routing-table entries could cause problems for such applications on other systems (non-LifeKeeper installed) to which the LifeKeeper system may be connecting. The multiple routing table entries can make it appear that the connection was made from the IP address associated with the logical interface rather than the physical interface.</p>
<b>IP subnet mask</b>	<p>For IP configurations under LifeKeeper protection, if the LifeKeeper-protected IP address is intended to be on the same subnet as the IP address of the physical interface on which it is aliased, the subnet mask of the two addresses must be the same. Incorrect settings of the subnet mask may result in connection delays and failures between the LifeKeeper GUI client and server.</p>
<b>EEpro100 driver initialization</b>	<p>The Intel e100 driver should be installed to resolve initialization problems with the eepr100 driver on systems with Intel Ethernet Interfaces. With the eepr100 driver, the following errors may occur when the interface is started at boot time and repeat continuously until the interface is shut down.</p> <p>eth0: card reports no Rx buffers</p> <p>eth0: card reports no resources</p>

## 5.4.2.8. Application Configuration

Item	Description
<b>Database Initialization Files</b>	The initialization files for databases need to be either on a shared device and symbolically linked to specified locations in the local file system or kept on separate systems and manually updated on both systems when changes need to be implemented.
<b>Localized Oracle Mount Points</b>	Localized Oracle environments are different depending on whether you connect as <i>internal</i> or as <i>sysdba</i> . A database on a localized mount point must be created with "connect / as sysdba" if it is to be put under LifeKeeper protection.
<b>Apache Updates</b>	<p>Upgrading a LifeKeeper protected Apache application as part of upgrading the Linux operating system requires that the default server instance be disabled on start up.</p> <p>If your configuration file (<i>httpd.conf</i>) is in the default directory (<i>/etc/httpd/conf</i>), the Red Hat upgrade will overwrite the config file. Therefore, you should make a copy of the file before upgrading and restore the file after upgrading.</p> <p>Also, see the Specific Configuration Considerations for Apache Web Server section in the <a href="#">Apache Web Server Recovery Kit Administration Guide</a>.</p>

## 5.4.2.9. Storage and Adapter Configuration

Item	Description
<p>Multipath I/O and Redundant Controllers</p>	<p>There are several multipath I/O solutions either already available or currently being developed for the Linux environment. SIOS Technology Corp. is actively working with a number of server vendors, storage vendors, adapter vendors and driver maintainers to enable LifeKeeper to work with their multipath I/O solutions. LifeKeeper's use of SCSI reservations to protect data integrity presents some special requirements that frequently are not met by the initial implementation of these solutions.</p> <p>Refer to the technical notes below for supported disk arrays to determine if a given array is supported with multiple paths and with a particular multipath solution. Unless an array is specifically listed as being supported by LifeKeeper with multiple paths and with a particular multipath solution, it must be assumed that it is not.</p>
<p>Heavy I/O in Multipath Configurations</p>	<p>In multipath configurations, performing heavy I/O while paths are being manipulated can cause a system to temporarily appear to be unresponsive. When the multipath software moves the access of a LUN from one path to another, it must also move any outstanding I/Os to the new path. The rerouting of the I/Os can cause a delay in the response times for these I/Os. If additional I/Os continue to be issued during this time, they will be queued in the system and can cause a system to run out of memory available to any process. Under very heavy I/O loads, these delays and low memory conditions can cause the system to be unresponsive such that LifeKeeper may detect a server as down and initiate a failover.</p> <p>There are many factors that will affect the frequency at which this issue may be seen.</p> <ul style="list-style-type: none"> <li>• The speed of the processor will affect how fast I/Os can be queued. A faster processor may cause the failure to be seen more frequently.</li> <li>• The amount of system memory will affect how many I/Os can be queued before the system becomes unresponsive. A system with more memory may cause the failure to be seen less frequently.</li> <li>• The number of LUNs in use will affect the amount of I/O that can be queued.</li> <li>• Characteristics of the I/O activity will affect the volume of I/O queued. In test cases where the problem has been seen, the test was writing an unlimited amount of data to the disk. Most applications will both read and write data. As the reads are blocked waiting on the failover, writes will also be throttled, decreasing the I/O rate such that the failure may be seen less frequently.</li> </ul>

	<p>For example, during testing of the IBM DS4000 multipath configuration with RDAC, when the I/O throughput to the DS4000 was greater than 190 MB per second and path failures were simulated, LifeKeeper would (falsely) detect a failed server approximately one time out of twelve. The servers used in this test were IBM x345 servers with dual Xeon 2.8GHz processors and 2 GB of memory connected to a DS4400 with 8 volumes (LUNs) per server in use. To avoid the failovers, the LifeKeeper parameter LCMNUMHBEATS (in <code>/etc/default/LifeKeeper</code>) was increased to 16. The change to this parameter results in LifeKeeper waiting approximately 80 seconds before determining that an unresponsive system is dead, rather than the default wait time of approximately 15 seconds.</p>
<p>Special Considerations for Switchovers with Large Storage Configurations</p>	<p>With some large storage configurations (for example, multiple logical volume groups with 10 or more LUNs in each volume group), LifeKeeper may not be able to complete a sendevent within the default timeout of 300 seconds when a failure is detected. This results in the switchover to the backup system failing. All resources are not brought in-service and an error message is logged in the LifeKeeper log.</p> <p>The recommendation with large storage configurations is to change SCSIERROR from “event” to “halt” in the <code>/etc/default/LifeKeeper</code> file. This will cause LifeKeeper to perform a “halt” on a SCSI error. LifeKeeper will then perform a successful failover to the backup system.</p>
<p>HP 3PAR StoreServ 7200 FC</p>	<p>The HP 3PAR StoreServ 7200 was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>HP 3PAR StoreServ 7200 (Firmware (HP 3PAR OS) version 3.1.2) using QLogic QMH2572 8Gb FC HBA for HP BladeSystem c-Class (Firmware version 5.06.02 (90d5)), driver version 8.03.07.05.06.2-k (RHEL bundled) with DMMP (device-mapper-1.02.66-6, device-mapper-multipath-0.4.9-46.el6).</p> <p>The test was performed with LifeKeeper for Linux v8.1.1 using RHEL 6.2 (x86_64).</p> <p><b>Note:</b> 3PAR StoreServ 7200 returns a reservation conflict with the default path checker. To avoid this conflict, set the following parameter in “<code>/etc/default/LifeKeeper</code>”:</p> <p><b>DMMP_REGISTRATION_TYPE=hba</b></p>
<p>HP 3PAR StoreServ 7400 FC</p>	<p>The HP 3PAR StoreServ 7400 was tested by a SIOS Technology Corp. partner</p>

	<p>with the following configurations:</p> <p>HP 3PAR StoreServ 7400 (Firmware (HP 3PAR OS) version 3.1.2) with HP DL380p Gen8 with Emulex LightPulse Fibre Channel SCSI HBA (driver version 8.3.5.45.4p) with DMMP (device-mapper-1.02.66-6, device-mapper-multipath-0.4.9-46.el6).</p> <p>The test was performed with LifeKeeper for Linux v8.1.1 using RHEL 6.2 (x86_64).</p> <p><b>Note:</b> 3PAR StoreServ 7400 returns a reservation conflict with the default path checker. To avoid this conflict, set the following parameter in “<code>/etc/default/LifeKeeper</code>”:</p> <p><b>DMMP_REGISTRATION_TYPE=hba</b></p> <p>And user friendly device mapping are not supported. Set the following parameter in “<code>multipath.conf</code>”</p> <p><b>“user_friendly_names no”</b></p>
<p>HP 3PAR StoreServ 7400 iSCSI (multipath configuration using the DMMP Recovery Kit)</p>	<p>The HP 3PAR StoreServ 7400 was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>HP 3PAR StoreServ 7400 (Firmware (HP 3PAR OS) version 3.1.3) using HP Ethernet 10Gb 2-port 560SFP+ Adapter (Networkdriver ixgbe-3.22.0.2) with iSCSI (iscsi-initiator-utils-6.2.0.873-10.el6.x86_64), DMMP (device-mapper-1.02.79-8.el6, device-mapper-multipath-0.4.9-72.el6).</p> <p><b>Note:</b> 3PAR StoreServ 7400 iSCSI returns a reservation conflict. To avoid this conflict, set the following parameter in “<code>/etc/default/LifeKeeper</code>”:</p> <p><b>DMMP_REGISTER_IGNORE=TRUE</b></p>
<p>HP 3PAR StoreServ 7400 iSCSI(using Quorum/Witness Kit)</p>	<p>The HP 3PAR StoreServ 7400 was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>iSCSI (iscsi-initiator-utils-6.2.0.872-21.el6.x86_64), DMMP (device-mapper-multipath-0.4.9-41, device-mapper-1.02.62-3), nx_nic v4.0.588.</p> <p>DMMP with the DMMP Recovery Kit on RHEL 6.1 — must be used with the combination of Quorum/Witness Server Kit and STONITH. To disable SCSI reservation, set RESERVATIONS=none in “<code>/etc/default/LifeKeeper</code>”.</p>

	<p>Server must have interface based on IPMI 2.0.</p>
<p>HP 3PAR StoreServ 10800 FC</p>	<p>The HP 3PAR StoreServ 10800 FC was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>Firmware (HP 3PAR OS) version 3.1.2 with HP DL380p Gen8 with Emulex LightPulse Fibre Channel HBA (driver version 8.3.5.45.4p) with DMMP (device-mapper-1.02.66-6, device-mapper-multipath-0.4.9-46 el6). The test was performed with SPS for Linux v8.1.2 using RHEL 6.2 (x86_64).</p> <p><b>Note:</b> 3PAR StoreServ 10800 FC returns a reservation conflict with the default path checker. To avoid this conflict, set the following parameter in <code>"/etc/default/LifeKeeper"</code>:</p> <p><b>DMMP_REGISTRATION_TYPE=hba</b></p> <p>And user friendly device mapping are not supported. Set the following parameter in "multipath.conf"</p> <p><b>"user_friendly_names no"</b></p>
<p>HP MSA1040/2040fc</p>	<p>The HP MSA 2040 Storage FC was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>HP MSA 2040 Storage FC (Firmware GL101R002) using HP SN1000Q 16Gb 2P FC HBA QW972A (Firmware version 6.07.02, driver version 8.04.00.12.06.0-k2 (RHEL bundled)) with DMMP (device-mapper-1.02.74-10, device-mapper-multipath-0.4.9-56).</p> <p>The test was performed with LifeKeeper for Linux v8.1.2 using RHEL 6.3 (X86_64).</p>
<p>HP P9500/XP</p>	<p>Certified by Hewlett-Packard Company using SIOS LifeKeeper for Linux v7.2 or later. Model tested was the HP P9500/XP and has been qualified to work with LifeKeeper on the following:</p> <ul style="list-style-type: none"> <li>• Red Hat Enterprise for 32-bit, x64 (64-bit; Opteron and Intel EMT64) RHEL 5.3, RHEL 5.4, RHEL 5.5</li> <li>• SuSE Enterprise Server for 32-bit, x64 (64-bit; Opteron and Intel EMT64) SLES 10 SP3, SLES 11, SLES 11 SP1</li> <li>• Native or Inbox Clustering Solutions RHCS and SLE HA</li> </ul>

<p>HP StoreVirtual 4330 iSCSI (multipath configuration using the DMMP Recovery Kit)</p>	<p>The HP StoreVirtual 4330 was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>HP StoreVirtual 4330 (Firmware HP LeftHand OS 10.5) using HP Ethernet 1Gb 4-port 331FLR (Networkdriver tg3-3.125g) with iSCSI (iscsi-initiator-utils-6.2.0.872-41.el6.x86_64), DMMP (device-mapper-1.02.74-10.el6,device-mapper-multipath-0.4.9-56.el6).</p>
<p>StoreVirtual (LeftHand) series OS (SAN/iQ) version 11.00 iSCSI (multipath configuration using the DMMP Recovery Kit)</p>	<p>OS (SAN/iQ) version 11.00 is supported in HP StoreVirtual (LeftHand) storage. All StoreVirtual series are supported, including StoreVirtual VSA as the virtual storage appliance. This storage was tested with the following configurations:</p> <p>StoreVirtual VSA + RHEL 6.4(x86_64) + DMMP.</p>
<p>HP StoreVirtual 4730 iSCSI (multipath configuration using the DMMP Recovery Kit)</p>	<p>The HP StoreVirtual 4730 was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>HP StoreVirtual 4730 (Firmware HP LeftHand OS 11.5) using HP FlexFabric 10Gb 2-port 536FLB Adapter (Networkdriver bnx2x-1.710.40) with iSCSI (iscsi-initiator-utils-6.2.0.873-10.el6.x86_64), DMMP (device-mapper-1.02.79-8.el6,device-mapper-multipath-0.4.9-72.el6).</p>
<p>HP StoreVirtual LeftHand OS version 11.5 iSCSI (multipath configuration using the DMMP Recovery Kit)</p>	<p>LeftHand OS version 11.5 is supported in HP StoreVirtual (LeftHand) storage. All StoreVirtual series are supported, including StoreVirtual VSA as the virtual storage appliance. This storage was tested with the following configurations:</p> <p>StoreVirtual 4730(11.5.00.0673.0) + RHEL 6.5(x86_64) + DMMP (device-mapper-1.02.79-8.el6.x86_64, device-mapper-multipath-0.4.9-72.el6.x86_64)</p>
<p>HP StoreVirtual 4330 iSCSI (using Quorum/Witness Kit)</p>	<p>The HP StoreVirtual 4330 was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>HP StoreVirtual 4330 (Firmware HP LeftHand OS 10.5) using iSCSI (iscsi-initiator-utils-6.2.0.872-41.el6.x86_64), bonding(version: 3.6.0),tg3(version: 3.125g)</p> <p>To disable SCSI reservation, set RESERVATIONS=none in <i>"/etc/default/LifeKeeper"</i>.</p>

<p>IBM San Volume Controller (SVC)</p>	<p>Certified by partner testing in a single path configuration. Certified by SIOS Technology Corp. in multipath configurations using the Device Mapper Multipath Recovery Kit.</p>
<p>IBM Storwize V7000 iSCSI</p>	<p>The IBM Storwize V7000 (Firmware Version 6.3.0.1) has been certified by partner testing using iSCSI (iscsi-initiator-utils-6.2.0.872-34.el6.x86_64) with DMMP (device-mapper-1.02.66-6.el6, device-mapper-multipath-0.4.9-46.el6). The test was performed with LifeKeeper for Linux v7.5 using RHEL 6.2.</p> <p><b>Restriction:</b> IBM Storwize V7000 must be used in combination with the Quorum/Witness Server Kit and STONITH. Disable SCSI reservation by setting the following in <code>/etc/default/LifeKeeper</code>:</p> <ul style="list-style-type: none"> <li>• RESERVATIONS=none</li> </ul>
<p>IBM Storwize V7000 FC</p>	<p>The IBM Storwize V7000 FC has been certified by partner testing in multipath configurations on Red Hat Enterprise Linux Server Release 6.2 (Tikanga), HBA: QLE2562 DMMP: 0.4.9-46.</p>
<p>IBM Storwize V3700 FC</p>	<p>The IBM Storwize V3700 FC has been certified by partner testing in multipath configurations on Red Hat Enterprise Linux Server Release 6.5 (Santiago), HBA: QLE2560 DMMP: 0.4.9-72.</p>
<p>IBM XIV Storage System</p>	<p>Certified by partner testing in only multipath configuration on Red Hat Enterprise Linux Release 5.6, HBA: NEC N8190-127 Single CH 4Gbps (Emulex LPe1150 equivalent), XIV Host Attachment Kit: Version 1.7.0.</p> <p><b>Note:</b> If you have to create over 32 LUNs on IBM XIV Storage System with LifeKeeper, please contact your IBM sales representative for details.</p>
<p>Dell EqualLogic PS4000/4100/4110/6000/6010/6100/6110/6500/6510</p>	<p>The Dell EqualLogic was tested by a SIOS Technology Corp. partner with the following configurations: Dell EqualLogic PS4000/4100/4110/6000/6010/6100/6110/6500/6510 using DMMP with the DMMP Recovery Kit with RHEL 5.3 with iscsi-initiator-utils-6.2.0.868-0.18.el5. With a large number of luns (over 20), change the REMOTETIMEOUT setting in <code>/etc/default/LifeKeeper</code> to REMOTETIMEOUT=600.</p>
<p>Fujitsu ETERNUS DX60 S2 / DX80 S2 / DX90 S2 iSCSI ETERNUS DX410</p>	<p>When using LifeKeeper DMMP ARK for multipath configuration it is necessary to set the following parameters to <code>/etc/multipath.conf</code>.</p> <pre>prio alua path_grouping_policy group_by_prio</pre>

S2 / DX440 S2 iSCSI	
ETERNUS DX8100 S2/DX8700 S2 iSCSI	
ETERNUS DX100 S3/DX200 S3 iSCSI	
ETERNUS DX500 S3/DX600 S3 iSCSI	
ETERNUS DX200F iSCSI	
ETERNUS DX60 S3 iSCSI	
ETERNUS AF250 / AF650 iSCSI	failback immediate
ETERNUS DX60 S4 / DX100 S4 / DX200 S4 iSCSI	no_path_retry 10  Path_checker tur
ETERNUS DX500 S4 / DX600 S4 / DX8900S4 iSCSI	
ETERNUS AF250 S2 / AF650 S2 iSCSI	
ETERNUS DX60 S5 / DX100 S5 / DX200 S5 iSCSI	
ETERNUS DX500 S5 / DX600 S5 / DX900S5 iSCSI	
ETERNUS AF150 S3 / AF250 S3 / AF650 S3 iSCSI	

<p>Fujitsu</p> <p>ETERNUS DX60 S2 / DX80 S2 / DX90 S2</p> <ul style="list-style-type: none"> <li>• FC, single path and multipath configurations</li> </ul> <p>ETERNUS DX410 S2 / DX440 S2</p> <ul style="list-style-type: none"> <li>• FC, single path and multipath configurations</li> </ul> <p>ETERNUS DX8100 S2/DX8700 S2</p> <ul style="list-style-type: none"> <li>• FC, single path and multipath configurations</li> </ul> <p>ETERNUS DX100 S3/DX200 S3/ DX500 S3/ DX600S3</p> <ul style="list-style-type: none"> <li>• FC, single path and multipath configurations</li> </ul> <p>ETERNUS DX200F</p> <ul style="list-style-type: none"> <li>• FC, single path and multipath configurations</li> </ul> <p>ETERNUS DX60 S3</p> <ul style="list-style-type: none"> <li>• FC, single path and</li> </ul>	<p>When using LifeKeeper DMMP ARK for multipath configuration it is necessary to set the following parameters to /etc/multipath.conf.</p> <pre>prio alua path_grouping_policy group_by_prio failback immediate no_path_retry 10 Path_checker tur</pre> <p>When using ETERNUS Multipath Driver for multipath configuration, it is no need to set parameters to any configure file.</p>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

<p>multipath configurations</p> <p>ETERNUS DX8700 S3 / DX8900 S3</p> <ul style="list-style-type: none"> <li>• FC, single path and multipath configurations</li> </ul> <p>ETERNUS AF250 / AF650</p> <p>ETERNUS DX60 S4 / DX100 S4 / DX200 S4</p> <p>ETERNUS DX500 S4 / DX600 S4 / DX8900S4</p> <p>ETERNUS AF250 S2 / AF650 S2</p> <p>ETERNUS DX60 S5 / DX100 S5 / DX200 S5</p> <p>ETERNUS DX500 S5 / DX600 S5 / DX900S5</p> <p>ETERNUS AF150 S3 / AF250 S3 / AF650 S3</p> <ul style="list-style-type: none"> <li>• iSCSI, single path and multipath configurations</li> </ul>	
<p>NEC iStorage M10e iSCSI (Multipath configuration using the SPS Recovery Kit)</p>	<p>The NEC iStorage M10e iSCSI was tested by a SIOS Technology Corp. partner with the following configurations:</p> <p>NEC iStorage M10e iSCSI + 1GbE NIC + iSCSI (iscsi-initiator-utils-6.2.0.873-10.el6.x86_64),SPS (sps-utils-5.3.0-0.el6,sps-driver-</p>

	E-5.3.0-2.6.32.431.el6)
<p>NEC iStorage Storage Path Savior Multipath I/O</p>	<p><b>Protecting Applications and File Systems That Use Multipath Devices:</b> In order for SPS to configure and protect applications or file systems that use SPS devices, the SPS recovery kit must be installed.</p> <p>Once the SPS kit is installed, simply creating an application hierarchy that uses one or more of the multipath device nodes will automatically incorporate the new resource types provided by the SPS kit.</p> <p><b>Multipath Device Nodes:</b> To use the SPS kit, any file systems and raw devices must be mounted or configured on the multipath device nodes (/dev/dd*) rather than on the native /dev/sd* device nodes.</p> <p><b>Use of SCSI-3 Persistent Reservations:</b> The SPS kit uses SCSI-3 persistent reservations with a "Write Exclusive" reservation type. This means that devices reserved by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other nodes will be unable to write to the device. Note that this does not mean that you can expect to be able to mount file systems on those other nodes for ongoing read-only access.</p> <p>LifeKeeper uses the sg_persist utility to issue and monitor persistent reservations. If necessary, LifeKeeper will install the sg_persist(8) utility.</p> <p><b>Tested Environment:</b> The SPS kit has been tested and certified with the NEC iStorage disk array using Emulex HBAs and Emulex lpfc driver. This kit is expected to work equally well with other NEC iStorage D, S and M supported by SPS.</p> <p><b>[Tested Emulex HBA]</b></p> <p>iStorage D-10 =====</p> <p>LP952 LP9802 LP1050 LP1150 =====</p> <p>iStorage M100 =====</p> <p>LPe1150 LPe11002 LPe1250</p>

	<p>LPe12002 LPe1105 LPe1205 =====</p> <p><b>Multipath Software Requirements:</b> The SPS kit has been tested with SPS for Linux 3.3.001. There are no known dependencies on the version of the SPS package installed.</p> <p><b>Installation Requirements:</b> SPS software must be installed prior to installing the SPS recovery kit.</p> <p><b>Adding or Repairing SPS Paths:</b> When LifeKeeper brings an SPS resource into service, it establishes a persistent reservation registered to each path that was active at that time. If new paths are added after the initial reservation, or if failed paths are repaired and SPS automatically reactivates them, those paths will not be registered as a part of the reservation until the next LifeKeeper quickCheck execution for the SPS resource. If SPS allows any writes to that path prior to that point in time, reservation conflicts that occur will be logged to the system message file. The SPS driver will retry these IOs on the registered path resulting in no observable failures to the application. Once quickCheck registers the path, subsequent writes will be successful.</p>
<p>Pure Storage FA-400 Series FC (Multipath configuration using the DMMP Recovery Kit)</p>	<p>By partner testing in multipath configuration of FC connection using the DMMP Recovery Kit.</p>
<p>QLogic Drivers</p>	<p>For other supported fibre channel arrays with QLogic adapters, use the qla2200 or qla2300 driver, version 6.03.00 or later.</p>
<p>Emulex Drivers</p>	<p>For the supported Emulex fibre channel HBAs, you must use the lpfc driver v8.0.16 or later.</p>
<p>Adaptec 29xx Drivers</p>	<p>For supported SCSI arrays with Adaptec 29xx adapters, use the aic7xxx driver, version 6.2.0 or later, provided with the OS distribution.</p>

## HP Multipath I/O Configurations

Item	Description
<p>Multipath Cluster Installation Using</p>	<p>For a fresh installation of a multiple path cluster that uses Secure Path, perform these steps:</p>

Secure Path	<ol style="list-style-type: none"> <li>1. Install the OS of choice on each server.</li> <li>2. Install the clustering hardware: FCA2214 adapters, storage, switches and cables.</li> <li>3. Install the HP Platform Kit.</li> <li>4. Install the HP Secure Path software. This will require a reboot of the system. Verify that Secure Path has properly configured both paths to the storage. See Secure Path documentation for further details.</li> <li>5. Install LifeKeeper.</li> </ol>
Secure Path Persistent Device Nodes	<p>Secure Path supports “persistent” device nodes that are in the form of /dev/spdev/spXX where XX is the device name. These nodes are symbolic links to the specific SCSI device nodes /dev/sdXX. LifeKeeper v4.3.0 or later will recognize these devices as if they were the “normal” SCSI device nodes /dev/sdXX. LifeKeeper maintains its own device name persistence, both across reboots and across cluster nodes, by directly detecting if a device is /dev/sda1 or /dev/sdq1, and then directly using the correct device node.</p> <p><b>Note:</b> Support for symbolic links to SCSI device nodes was added in LifeKeeper v4.3.0.</p>
Active/Passive Controllers and Controller Switchovers	<p>The MSA1000 implements multipathing by having one controller active with the other controller in standby mode. When there is a problem with either the active controller or the path to the active controller, the standby controller is activated to take over operations. When a controller is activated, it takes some time for the controller to become ready. Depending on the number of LUNs configured on the array, this can take 30 to 90 seconds. During this time, IOs to the storage will be blocked until they can be rerouted to the newly activated controller.</p>
Single Path on Boot Up Does Not Cause Notification	<p>If a server can access only a single path to the storage when the system is loaded, there will be no notification of this problem. This can happen if a system is rebooted where there is a physical path failure as noted above, but transient path failures have also been observed. It is advised that any time a system is loaded, the administrator should check that all paths to the storage are properly configured, and if not, take actions to either repair any hardware problems or reload the system to resolve a transient problem.</p>

## Hitachi Multipath I/O Configurations

Item	Description
Protecting Applications and File Systems That Use Multipath Devices	<p>In order for LifeKeeper to configure and protect applications or file systems that use devices, the HDLM Kit must be installed.</p> <p>Once the HDLM Kit is installed, simply creating an application hierarchy that uses one or more of the new resource types will automatically incorporate the new resource types provided by the HDLM Kit.</p>
Multipath Device Nodes	<p>To use the HDLM Kit, any file systems and raw devices must be mounted or configured on the multipath device nodes rather than on the native /dev/sd* device nodes.</p>

<p>Use of SCSI-3 Persistent Reservations</p>	<p>The HDLM Kit uses SCSI-3 persistent reservations with a “Write Exclusive” reservation type. This by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other to the device. Note that this does not mean that you can expect to be able to mount file systems o ongoing read-only access.</p> <p>LifeKeeper uses the sg_persist utility to issue and monitor persistent reservations. If necessary, L sg_persist(8) utility.</p>
<p>Hardware Requirements</p>	<p>The HDLM Kit has been tested and certified with the Hitachi SANRISE AMS1000 disk array using 8.02.00-k5-rhel5.2-04 driver and Silkworm3800 FC switch. This kit is expected to work equally well arrays. The HDLM Kit has also been certified with the SANRISE AMS series, SANRISE USP and HBA driver must be supported by HDLM.</p> <p>BR1200 is certified by Hitachi Data Systems. Both single path and multipath configuration require BR1200 configuration using the RDAC driver is supported, and the BR1200 configuration using H supported.</p>
<p>Multipath Software Requirements</p>	<p>The HDLM kit has been tested with HDLM for Linux as follows:</p> <p>05-80, 05-81, 05-90, 05-91, 05-92, 05-93, 05-94, 6.0.0, 6.0.1, 6.1.0, 6.1.1, 6.1.2, 6.2.0, 6.2.1, 6.3. 6.5.2, 6.6.0, 6.6.2, 7.2.0, 7.2.1, 7.3.0, 7.3.1, 7.4.0, 7.4.1, 7.5.0, 7.6.0, 7.6.1, 8.0.0, 8.0.1, 8.1.0, 8. 8.2.1, 8.4.0, 8.5.0, 8.5.1, 8.5.2, 8.5.3, 8.6.0, 8.6.1, 8.6.2, 8.6.4, 8.6.5, 8.7.0, 8.7.1, 8.7.2, 8.7.3, 8.7 8.8.1</p> <p>There are no known dependencies on the version of the HDLM package installed.</p> <p><b>Note:</b> The product name changed to “Hitachi Dynamic Link Manager Software (HDLM)” for HDLM than 6.0.0 (05-9x) are named “Hitachi Command Dynamic Link Manager (HDLM)”.</p> <p><b>Note:</b> HDLM version 6.2.1 or later is not supported by HDLM Recovery Kit v6.4.0-2. If you need to you can use HDLM Recovery Kit v7.2.0-1 or later with LK Core v7.3 or later.</p> <p><b>Note:</b> If using LVM with HDLM, the version supported by HDLM is necessary. Also, a filter setting lvm.conf to ensure that the system does not detect the /dev/sd* corresponding to the /dev/sddlm*. please see “LVM Configuration” in the HDLM manual.</p>
<p>Linux Distribution Requirements</p>	<p>Linux Distribution Requirements</p> <p>HDLM is supported in the following distributions:</p> <p>RHEL 4 (AS/ES) (x86 or x86_64) Update 1, 2, 3, 4, Update 4 Security Fix (*2), 4.5,4.5 Security Fix Fix(*8),4.7,4.7 Security Fix(*9), 4.8,4.8 Security Fix(*12) (x86_64(*1) non-English version)</p>

RHEL 5, 5.1, 5.1 Security Fix(\*5), 5.2, 5.2 Security Fix(\*6), 5.3, 5.3 Security Fix(\*10),5.4 , 5.4 Security Fix(\*13), 5.6, 5.6 Security Fix(\*14), 5.7 (x86/x86\_64(\*1) non-English version), 5.8,5.9,5.10,5.11 (x86\_64(\*1) non-English version)

RHEL 6, 6.1, 6.2 , 6.3, 6.4, 6.5, 6.6, 6.7, 6.8, 6.9, 6.10(x86/x86\_64(\*1) non-English version)(\*15)

RHEL 7, 7.1, 7.2, 7.3, 7.3, 7.4, 7.5, 7.6, 7.7, 7.8, 7.9 (x86\_64(\*1) non-English version)

RHEL 8.1, 8.2, 8.3 (x86\_64(\*1) non-English version)

(\*1) AMD Opteron(Single Core , Dual Core) or Intel EM64T architecture CPU with x86\_64 kernel.

(\*2) The following kernels are supported

x86 : 2.6.9-42.0.3.EL , 2.6.9-42.0.3.ELsmp , 2.6.9-42.0.3.ELhugemem

x86\_64 : 2.6.9-42.0.3.EL , 2.6.9-42.0.3.ELsmp , 2.6.9-42.0.3.ELlargesmp

(\*3) Hitachi does not support RHEL4 U2 environment

(\*4) The following kernels are supported

x86 : 2.6.9-55.0.12.EL , 2.6.9-55.0.12.ELsmp , 2.6.9-55.0.12.ELhugememx

x86\_64 : 2.6.9-55.0.12.EL , 2.6.9-55.0.12.ELsmp , 2.6.9-55.0.12.ELlargesmp

(\*5) The following kernels are supported

x86 : 2.6.18-53.1.13.el5 , 2.6.18-53.1.13.el5PAE , 2.6.18-53.1.21.el5 , 2.6.18-53.1.21.el5PAE

x86\_64 : 2.6.18-53.1.13.el5 , 2.6.18-53.1.21.el5

(\*6) The following kernels are supported

x86 : 2.6.18-92.1.6.el5 , 2.6.18-92.1.6.el5PAE , 2.6.18-92.1.13.el5 , 2.6.18-92.1.13.el5PAE , 2.6.18-92.1.22.el5

2.6.18-92.1.22.el5PAE

x86\_64 : 2.6.18-92.1.6.el5 , 2.6.18-92.1.13.el5 , 2.6.18-92.1.22.el5

(\*7) The following kernels are supported

x86 : 2.6.9-34.0.2.EL , 2.6.9-34.0.2.ELsmp , 2.6.9-34.0.2.ELhugemem

x86\_64 : 2.6.9-34.0.2.EL , 2.6.9-34.0.2.ELsmp , 2.6.9-34.0.2.ELlargesmp

(\*8) The following kernels are supported

x86 : 2.6.9-67.0.7.EL , 2.6.9-67.0.7.ELsmp , 2.6.9-67.0.7.ELhugemem , 2.6.9-67.0.22.EL , 2.6.9-67.0.22.ELhugemem

2.6.9-67.0.22.ELhugemem

x86\_64 : 2.6.9-67.0.7.EL , 2.6.9-67.0.7.ELsmp , 2.6.9-67.0.7.ELlargesmp

2.6.9-67.0.22.EL , 2.6.9-67.0.22.ELsmp , 2.6.9-67.0.22.ELlargesmp

(\*9) The following kernels are supported

x86 : 2.6.9-78.0.1.EL , 2.6.9-78.0.1.ELsmp , 2.6.9-78.0.1.ELhugemem , 2.6.9-78.0.5.EL , 2.6.9-78.0.5.ELhugemem , 2.6.9-78.0.8.EL , 2.6.9-78.0.8.ELsmp , 2.6.9-78.0.8.ELhugemem, 2.6.9-78.0.17.ELsmp , 2.6.9-78.0.17.ELhugemem , 2.6.9-78.0.22.EL , 2.6.9-78.0.22.ELsmp , 2.6.9-78.0.22.ELhugemem

2.6.9-78.0.22.ELhugemem

x86\_64 : 2.6.9-78.0.1.EL , 2.6.9-78.0.1.ELsmp , 2.6.9-78.0.1.ELlargesmp, 2.6.9-78.0.5.EL , 2.6.9-78.0.5.ELhugemem , 2.6.9-78.0.8.EL , 2.6.9-78.0.8.ELsmp , 2.6.9-78.0.8.ELhugemem, 2.6.9-78.0.17.ELsmp , 2.6.9-78.0.17.ELhugemem , 2.6.9-78.0.22.EL , 2.6.9-78.0.22.ELsmp , 2.6.9-78.0.22.ELhugemem

2.6.9-78.0.22.ELhugemem

2.6.9-78.0.5.ELlargesmp , 2.6.9-78.0.8.EL , 2.6.9-78.0.8.ELsmp , 2.6.9-78.0.8.ELlargesmp , 2.6.9-78.0.17.ELsmp , 2.6.9-78.0.17.ELlargesmp , 2.6.9-78.0.22.EL , 2.6.9-78.0.22.ELsmp , 2.6.9-78.0.22.ELlargesmp

(\*10) The following kernels are supported

x86 : 2.6.18-128.1.10.el5 , 2.6.18-128.1.10.el5PAE , 2.6.18-128.1.14.el5, 2.6.18-128.1.14.el5PAE  
2.6.18-128.7.1.el5PAE

x86\_64 : 2.6.18-128.1.10.el5 , 2.6.18-128.1.14.el5

(\*11) The following kernels are supported

x86 : 2.6.18-164.9.1.el5 , 2.6.18-164.9.1.el5PAE , 2.6.18-164.11.1.el5 , 2.6.18-164.11.1.el5PAE

x86\_64 : 2.6.18-164.9.1.el5 , 2.6.18-164.11.1.el5

(\*12) The following kernels are supported

x86 : 2.6.9-89.0.20.EL , 2.6.9-89.0.20.ELsmp , 2.6.9-89.0.20.Elhugemem

x86\_64 : 2.6.9-89.0.20.EL , 2.6.9-89.0.20.ELsmp , 2.6.9-89.0.20.Ellargesmp

(\*13) The following kernels are supported

x86 : 2.6.18-194.11.1.el5, 2.6.18-194.11.1.el5PAE, 2.6.18-194.11.3.el5, 2.6.18-194.11.3.el5PAE,  
2.6.18-194.17.1.el5PAE, 2.6.18-194.32.1.el5, 2.6.18-194.32.1.el5PAE

x86\_64 : 2.6.18-194.11.1.el5, 2.6.18-194.11.3.el5, 2.6.18-194.17.1.el5, 2.6.18-194.32.1.el5

(\*14) The following kernels are supported

x86 :

2.6.18-238.1.1.el5,2.6.18-238.1.1.el5PAE,2.6.18-238.9.1.el5,2.6.18-238.9.1.el5PAE,2.6.18-238.19.1.el5

x86\_64 : 2.6.18-238.1.1.el5,2.6.18-238.9.1.el5,2.6.18-238.19.1.el5

(\*15) The following kernels are supported

x86 : 2.6.32-71.el6.i686, 2.6.32-131.0.15.el6.i686, 2.6.32-220.el6.i686 , 2.6.32-279.el6.i686

x86\_64 : 2.6.32-71.el6.x86\_64, 2.6.32-131.0.15.el6.x86\_64, 2.6.32-220.el6.x86\_64 , 2.6.32-279.el6.x86\_64

(\*16) The following kernels are supported

x86 : 2.6.18-274.12.1.el5 , 2.6.18-274.12.1.el5PAE , 2.6.18-274.18.1.el5 , 2.6.18-274.18.1.el5PAE

x86\_64 : 2.6.18-274.12.1.el5 , 2.6.18-274.18.1.el5

(\*17) The following kernels are supported

x86 : 2.6.18-308.8.2.el5 , 2.6.18-308.8.2.el5PAE

x86\_64 : 2.6.18-308.8.2.el5

(\*18) The following kernels are supported

x86 : 2.6.32-220.4.2.el6.i686, 2.6.32-220.17.1.el6.i686, 2.6.32-220.23.1.el6.i686, 2.6.32-220.31.1.el6.i686,  
2.6.32-220.45.1.el6.i686, 2.6.32-220.77.1.el6.x86\_64

x86\_64 : 2.6.32-220.4.2.el6.x86\_64, 2.6.32-220.17.1.el6.x86\_64, 2.6.32-220.23.1.el6.x86\_64, 2.6.32-220.31.1.el6.x86\_64,  
2.6.32-220.45.1.el6.x86\_64, 2.6.32-220.48.1.el6.x86\_64 , 2.6.32-220.64.1.el6.x86\_64 , 2.6.32-220.72.2.el6.x86\_64 ,

2.6.32-220.72.2.el6.x86\_64 , 2.6.32-220.73.1.el6.x86\_64 , 2.6.32-220.75.1.el6.x86\_64, 2.6.32-220.77.1.el6.x86\_64

2.6.32-220.77.1.el6.x86\_64 , 2.6.32-220.73.1.el6.x86\_64 , 2.6.32-220.75.1.el6.x86\_64, 2.6.32-220.77.1.el6.x86\_64

(\*19) The following kernels are supported

x86 : 2.6.32-279.19.1.el6.i686  
 x86\_64 : 2.6.32-279.19.1.el6.x86\_64

(\*20) The following kernels are supported

x86 : 2.6.32-358.6.2.el6.i686, 2.6.32-358.11.1.el6.i686, 2.6.32-358.14.1.el6.i686, 2.6.32-358.23.2.el6.i686,  
 x86\_64 : 2.6.32-358.6.2.el6.x86\_64, 2.6.32-358.11.1.el6.x86\_64, 2.6.32-358.14.1.el6.x86\_64, 2.6.32-358.23.2.el6.x86\_64,  
 2.6.32-358.28.1.el6.x86\_64, 2.6.32-358.87.1.el6.x86\_64

(\*21) The following kernels are supported

x86 : 2.6.32-431.1.2.el6.i686, 2.6.32-431.3.1.el6.i686, 2.6.32-431.5.1.el6.i686, 2.6.32-431.17.1.el6.i686,  
 2.6.32-431.20.3.el6.i686, 2.6.32-431.23.3.el6.i686, 2.6.32-431.29.2.el6.i686, 2.6.32-431.72.1.el6.i686,  
 x86\_64 : 2.6.32-431.1.2.el6.x86\_64, 2.6.32-431.3.1.el6.x86\_64, 2.6.32-431.5.1.el6.x86\_64, 2.6.32-431.17.1.el6.x86\_64,  
 2.6.32-431.20.3.el6.x86\_64, 2.6.32-431.23.3.el6.x86\_64, 2.6.32-431.29.2.el6.x86\_64, 2.6.32-431.72.1.el6.x86\_64,  
 2.6.32-431.77.1.el6.x86\_64, 2.6.32-431.87.1.el6.x86\_64

(\*22) The following kernels are supported

x86 : 2.6.32-504.3.3.el6.i686 , 2.6.32-504.12.2.el6.i686 , 2.6.32-504.30.3.el6.i686  
 x86\_64 : 2.6.32-504.3.3.el6.x86\_64 , 2.6.32-504.12.2.el6.x86\_64 , 2.6.32-504.16.2.el6.x86\_64 , 2.6.32-504.30.3.el6.x86\_64 ,  
 2.6.32-504.40.1.el6.x86\_64 , 2.6.32-504.43.1.el6.x86\_64, 2.6.32-504.66.1.el6.x86\_64

(\*23) The following kernels are supported

x86 : 2.6.18-348.1.1.el5, 2.6.18-348.1.1.el5PAE, 2.6.18-348.6.1.el5, 2.6.18-348.6.1.el5PAE, 2.6.18-348.18.1.el5,  
 2.6.18-348.18.1.el5PAE  
 x86\_64 : 2.6.18-348.1.1.el5, 2.6.18-348.6.1.el5, 2.6.18-348.18.1.el5

(\*24) The following kernels are supported

x86\_64:3.10.0-123.13.2.el7.x86\_64, 3.10.0-123.20.1.el7.x86\_64

(\*25) The following kernels are supported

x86\_64:3.10.0-229.4.2.el7.x86\_64,3.10.0-229.20.1.el7.x86\_64, 3.10.0-229.34.1.el7.x86\_64

(\*26) The following kernels are supported

x86\_64:3.10.0-327.4.4.el7.x86\_64, 3.10.0-327.4.5.el7.x86\_64, 3.10.0-327.10.1.el7.x86\_64, 3.10.0-327.10.2.el7.x86\_64,  
 3.10.0-327.22.2.el7.x86\_64, 3.10.0-327.36.1.el7.x86\_64, 3.10.0-327.36.3.el7.x86\_64, 3.10.0-327.36.4.el7.x86\_64,  
 3.10.0-327.46.1.el7.x86\_64, 3.10.0-327.49.2.el7.x86\_64 , 3.10.0-327.55.2.el7.x86\_64 , 3.10.0-327.55.3.el7.x86\_64,  
 3.10.0-327.58.1.el7.x86\_64, 3.10.0-327.62.1.el7.x86\_64 , 3.10.0-327.62.4.el7.x86\_64, 3.10.0-327.62.5.el7.x86\_64,  
 3.10.0-327.93.1.el7.x86\_64, 3.10.0-327.96.1.el7.x86\_64

(\*27) The following kernels are supported

x86: 2.6.32-573.8.1.el6.i686, 2.6.32-573.12.1.el6.i686, 2.6.32-573.18.1.el6.i686, 2.6.32-573.53.1.el6.i686,  
 x86\_64: 2.6.32-573.8.1.el6.x86\_64 , 2.6.32-573.12.1.el6.x86\_64, 2.6.32-573.18.1.el6.x86\_64, 2.6.32-573.53.1.el6.x86\_64

(\*28) The following kernels are supported

x86:2.6.32-642.1.1.el6.i686, 2.6.32-642.6.2.el6.i686, 2.6.32-642.13.1.el6.i686  
 x86\_64:2.6.32-642.1.1.el6.x86\_64 , 2.6.32-642.6.1.el6.x86\_64,2.6.32-642.6.2.el6.x86\_64 , 2.6.32-642.13.1.el6.x86\_64,  
 2.6.32-642.15.1.el6.x86\_64

(\*29) The following kernels are supported

x86 : 2.6.18-416.el5 , 2.6.18-416.el5PAE, 2.6.18-419.el5 , 2.6.18-419.el5PAE, 2.6.18-426.el5 , 2.6.18-426.el5PAE  
 x86\_64 : 2.6.18-416.el5, 2.6.18-419.el5, 2.6.18-426.el5

(\*30)The following kernels are supported

x86\_64 : 3.10.0-514.6.1.el7.x86\_64 , 3.10.0-514.10.2.el7.x86\_64 , 3.10.0-514.16.1.el7.x86\_64 , 3.10.0-514.26.2.el7.x86\_64 , 3.10.0-514.36.5.el7.x86\_64 , 3.10.0-514.44.1.el7.x86\_64, 3.10.0-514.52.2.el7.x86\_64

(\*31)The following kernels are supported

x86 : 2.6.32-696.3.2.el6.i686 2.6.32-696.6.3.el6.i686 , 2.6.32-696.10.3.el6.i686 , 2.6.32-696.18.7.el6.i686 , 2.6.32-696.20.1.el6.i686 , 2.6.32-696.23.1.el6.i686  
 x86\_64 : 2.6.32-696.3.2.el6.x86\_64 , 2.6.32-696.10.3.el6.x86\_64 , 2.6.32-696.18.7.el6.x86\_64 , 2.6.32-696.20.1.el6.x86\_64 , 2.6.32-696.23.1.el6.x86\_64

(\*32)The following kernels are supported

x86\_64 : 3.10.0-693.1.1.el7.x86\_64, 3.10.0-693.5.2.el7.x86\_64, 3.10.0-693.11.1.el7.x86\_64 , 3.10.0-693.13.1.el7.x86\_64 , 3.10.0-693.17.1.el7.x86\_64 , 3.10.0-693.21.1.el7.x86\_64, 3.10.0-693.23.1.el7.x86\_64 , 3.10.0-693.27.1.el7.x86\_64 , 3.10.0-693.31.1.el7.x86\_64

(\*33)The following kernels are supported

x86\_64 3.10.0-862.3.2.el7.x86\_64, 3.10.0-862.14.4.el7.x86\_64

(\*34)The following kernels are supported

x86 : 2.6.32-754.3.5.el6.i686 , 2.6.32-754.15.3.el6.i686, 2.6.32-754.24.3.el6.i686  
 x86\_64 2.6.32-754.3.5.el6.x86\_64 , 2.6.32-754.15.3.el6.x86\_64, 2.6.32-754.24.3.el6.x86\_64

(\*35)The following kernels are supported

x86\_64 : 3.10.0-957.10.1.el7.x86\_64 , 3.10.0-957.12.2.el7.x86\_64,  
 3.10.0-957.21.3.el7.x86\_64 , 3.10.0-957.27.2.el7.x86\_64

(\*36)The following kernels are supported

x86\_64 3.10.0 1062 .1 .1 .el 7.x86\_64, 3.10.0 1062 9 .1 .el 7.x86\_64, 3.10.0-1062.18.1.el7.x86\_64, 3.10.0-1062.22.1.el7.x86\_64

(\*37)The following kernels are supported

x86\_64 : 4.18.0-147.5.1.el8\_1.x86\_64, x86\_64 : 4.18.0-147.8.1.el8\_1.x86\_64

(\*38)The following kernels are supported x86\_64 : 4.18.0-193.28.1.el8\_2.x86\_64

(\*39) Only iSCSI environments are supported

(\*40)The following kernels are supported

x86\_64 : 4.18.0-240.22.1.el8\_3.x86\_64

(\*41) The following kernels are supported.

x86\_64 3.10. 0 1 160 .1 5.2 .el 7.x86\_6 4

Installation Requirements	HDLM software must be installed prior to installing the HDLM recovery kit. Also, customers wanting to SCSI devices to HDLM devices must run the Installation setup script after configuring the HDLM environment. If the HDLM driver is not installed.
Adding or Repairing HDLM Paths	When LifeKeeper brings an HDLM resource into service, it establishes a persistent reservation register at that time. If new paths are added after the initial reservation, or if failed paths are repaired and HDLM registers them, those paths will not be registered as a part of the reservation until the next LifeKeeper quickCheck of the resource. If HDLM allows any writes to that path prior to that point in time, reservation conflicts that occur will be recorded in a message file. The HDLM driver will retry these IOs on the registered path resulting in no observable failure. When quickCheck registers the path, subsequent writes will be successful. The status will be changed to "Offline(R)". If there is a reservation conflict. If the status is "Offline(E)", customers will need to manually change the status to "Online" using the "ckchk" command.
Additional settings for RHEL7.x	If Red Hat Enterprise Linux 7.0 or later are used with HDLM Recovery Kit, you must add "HDLM_DLMMGR=default/LifeKeeper. HDLM_DLMMGR=.dlmmgr_exe" to the /etc/defaults/ldmrc file.

		OS version / Architecture										
		RHEL4										
		U1-U4	U3 Security Fix(*7)	U4 Security Fix(*2)	4.5	4.5 Security Fix(*4)	4.6	4.6 Security Fix(*8)	4.7	4.7 Security Fix(*9)	4.8	4.8 Security Fix(*12)
		x86/x86_64										
HDLM	05-80 05-81 05-90	X										
	05-91 05-92	X		X								
	05-93	X(*3)		X	X							
	05-94	X(*3)		X	X	X	X	X				
	6.0.0	X(*3)		X	X	X	X	X	X	X	X	X
	6.0.1	X(*3)		X	X	X	X	X	X	X	X	X
	6.1.0	X(*3)		X	X	X	X	X	X	X	X	X
	6.1.1	X(*3)	X	X	X	X	X	X	X	X	X	X
	6.1.2	X(*3)	X	X	X	X	X	X	X	X	X	X
	6.2.0	X(*3)	X	X	X	X	X	X	X	X	X	X
	6.2.1	X(*3)	X	X	X	X	X	X	X	X	X	X
	6.3.0	X(*3)	X	X	X	X	X	X	X	X	X	X

	6.4.0	X(*3)	X	X	X	X	X	X	X	X	X	X	
	6.4.1	X(*3)	X	X	X	X	X	X	X	X	X	X	
	6.5.0	X(*3)	X	X	X	X	X	X	X	X	X	X	
	6.5.1	X(*3)	X	X	X	X	X	X	X	X	X	X	
	6.5.2	X(*3)	X	X	X	X	X	X	X	X	X	X	
	6.6.0	X(*3)	X	X	X	X	X	X	X	X	X	X	
	6.6.2	X(*3)	X	X	X	X	X	X	X	X	X	X	
	7.2.0	X(*3)	X	X	X	X	X	X	X	X	X	X	
	7.2.1	X(*3)	X	X	X	X	X	X	X	X	X	X	
	7.3.0 or later	X(*3)	X	X	X	X	X	X	X	X	X	X	
LifeKeeper	v6.0	X	X	X									
	v6.0(v6.0.1-2 or later)												
	v6.1	X	X	X									
	(v6.1.0-5 or later)												
	v6.2	X	X	X	X	X	X	X					
	(v6.2.0-5 or later)												
	v6.2	X	X	X	X	X	X	X					
	(v6.2.2-1or later)												
	v6.3	X	X	X	X	X	X	X					
	(v6.3.2-1or later)												
	v6.4	X	X	X	X	X	X	X	X	X			
	(v6.4.0-10 or later)												
	v7.0	X	X	X	X	X	X	X	X	X	X	X	X
	(v7.0.0-5 or later)												
V 7.1	X	X	X	X	X	X	X	X	X	X	X	X	
(v7.1.0-8 or later)													

	V7.2												
	(v7.2.0-10 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	V 7.3												
	(v7.3.0-21 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	V 7.4												
	(v7.4.0-63 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	V 7.5	RHEL4 is not supported in v7.5 or later version of LK.											
(v7.5.0-3640 or later)													
HDLM ARK	6.0.1-2	X	X	X	X	X	X	X					
	6.1.0-4	X	X	X	X	X	X	X					
	6.2.2-3	X	X	X	X	X	X	X					
	6.2.3-1	X	X	X	X	X	X	X	X	X	X	X	X
	6.4.0-2	X	X	X	X	X	X	X	X	X	X	X	X
	7.0.0-1	X	X	X	X	X	X	X	X	X	X	X	X
	7.2.0-1	X	X	X	X	X	X	X	X	X	X	X	X
X = supported blank = not supported													

		OS version / Architecture											
		RHEL5											
		No Updates	5.1	5.1 Security Fix (*5)	5.2	5.2 Security Fix (*6)	5.3	5.3 Security Fix(*10)	5.4	5.4 Security Fix(*11)	5.5	5.5 Security Fix(*13)	5.6
		x86/x86_64											
HDLM	05-80 05-81 05-90												
	05-91 05-92												
	05-93	X											
	05-94	X	X										

	6.0.0	X	X	X	X	X							
	6.0.1	X	X	X	X	X							
	6.1.0	X	X	X	X	X							
	6.1.1	X	X	X	X	X							
	6.1.2	X	X	X	X	X	X	X	X	X	X	X	X
	6.2.0	X	X	X	X	X	X	X	X	X	X	X	X
	6.2.1	X	X	X	X	X	X	X	X	X	X	X	X
	6.3.0	X	X	X	X	X	X	X	X	X	X	X	X
	6.4.0	X	X	X	X	X	X	X	X	X	X	X	X
	6.4.1	X	X	X	X	X	X	X	X	X	X	X	X
	6.5.0	X	X	X	X	X	X	X	X	X	X	X	X
	6.5.1	X	X	X	X	X	X	X	X	X	X	X	X
	6.5.2	X	X	X	X	X	X	X	X	X	X	X	X
	6.6.0	X	X	X	X	X	X	X	X	X	X	X	X
	6.6.2	X	X	X	X	X	X	X	X	X	X	X	X
	7.2.0	X	X	X	X	X	X	X	X	X	X	X	X
	7.2.1	X	X	X	X	X	X	X	X	X	X	X	X
	7.3.0	X	X	X	X	X	X	X	X	X	X	X	X
	7.3.1	X	X	X	X	X	X	X	X	X	X	X	X
	7.4.0	X	X	X	X	X	X	X	X	X	X	X	X
	7.4.1	X	X	X	X	X	X	X	X	X	X	X	X
	7.5.0	X	X	X	X	X	X	X	X	X	X	X	X
	7.6.0	X	X	X	X	X	X	X	X	X	X	X	X
	7.6.1	X	X	X	X	X	X	X	X	X	X	X	X
	8.0.0	X	X	X	X	X	X	X	X	X	X	X	X
	8.0.1	X	X	X	X	X	X	X	X	X	X	X	X
	8.1.0	X	X	X	X	X	X	X	X	X	X	X	X
	8.1.1	X	X	X	X	X	X	X	X	X	X	X	X
	8.1.2	X	X	X	X	X	X	X	X	X	X	X	X
	8.1.3	X	X	X	X	X	X	X	X	X	X	X	X
	8.1.4	X	X	X	X	X	X	X	X	X	X	X	X

	8.2.0	X	X	X	X	X	X	X	X	X	X	X	X
	8.4.0	X	X	X	X	X	X	X	X	X	X	X	X
	8.5.0	X	X	X	X	X	X	X	X	X	X	X	X
	8.5.1	X	X	X	X	X	X	X	X	X	X	X	X
	8.5.2	X	X	X	X	X	X	X	X	X	X	X	X
	8.5.3	X	X	X	X	X	X	X	X	X	X	X	X
	8.5.4	X	X	X	X	X	X	X	X	X	X	X	X
	8.6.0	X	X	X	X	X	X	X	X	X	X	X	X
	8.6.1	X	X	X	X	X	X	X	X	X	X	X	X
	8.6.2	X	X	X	X	X	X	X	X	X	X	X	X
	8.6.4	X	X	X	X	X	X	X	X	X	X	X	X
	8.6.5	X	X	X	X	X	X	X	X	X	X	X	X
	8.7.0	X	X	X	X	X	X	X	X	X	X	X	X
	8.7.1	X	X	X	X	X	X	X	X	X	X	X	X
	8.7.2	X	X	X	X	X	X	X	X	X	X	X	X
8.7.3	X	X	X	X	X	X	X	X	X	X	X	X	
LifeKeeper	v6.0 (v6.0.1-2 or later)												
	v6.1 (v6.1.0-5 or later)	X	X										
	v6.2 (v6.2.0-5 or later)	X	X										
	v6.2 (v6.2.2-1 or later)	X	X	X									
	v6.3 (v6.3.2-1 or later)	X	X	X	X	X							
	v6.4 (v6.4.0-10 or later)	X	X	X	X	X	X	X					
	v7.0	X	X	X	X	X	X	X	X	X			

(v7.0.0-5 or later)													
v7.1 (v7.1.0-8 or later)	X	X	X	X	X	X	X	X	X	X	X	X	
v7.2 (v7.2.0-10 or later)	X	X	X	X	X	X	X	X	X	X	X	X	
v7.3 (v7.3.0-21 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v7.4 (v7.4.0-63 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v7.5 (v7.5.0-3640 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.0 (v8.0.0-510 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.1 (v8.1.1-5620 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.2 (v8.2.0-6213 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.2.1 (v8.2.1-6353 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.3.0 (v8.3.0-6389 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.3.1 (v8.3.1-6397 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.3.2 (v8.3.2-6405 or later)	X	X	X	X	X	X	X	X	X	X	X	X	X
v8.4.0 (v8.4.0-6427	X	X	X	X	X	X	X	X	X	X	X	X	X

	or later)												
	v8.4.1 (v8.4.1-6449 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.0.0 (v9.0.0.0-6488 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.0.1 (v9.0.1-6492 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.0.2 (v9.0.2-6213 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.1.0 (v9.1.0-6538 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.1.1 (v9.1.1-6594 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.2.0 (v9.2.0-6629 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.2.1 (v9.2.1.0-6653 or later)	X	X	X	X	X	X	X	X	X	X	X	X
	v9.2.2 (v9.2.2-6679 or later)	X	X	X	X	X	X	X	X	X	X	X	X
HDLM ARK	6.0.1-2												
	6.1.0-4	X	X										
	6.2.2-3	X	X	X									
	6.2.3-1	X	X	X	X	X							
	6.4.0-2	X	X	X	X	X	X	X					
	7.0.0-1	X	X	X	X	X	X	X	X	X	X	X	
	7.2.0-1	X	X	X	X	X	X	X	X	X	X	X	X
	8.1.1-5620	X	X	X	X	X	X	X	X	X	X	X	X
	8.2.0-6213	X	X	X	X	X	X	X	X	X	X	X	X

	8.2.1-6213	X	X	X	X	X	X	X	X	X	X	X	X
	8.3.0-6213	X	X	X	X	X	X	X	X	X	X	X	X
	8.4.0-6213	X	X	X	X	X	X	X	X	X	X	X	X
	8.4.1-6213	X	X	X	X	X	X	X	X	X	X	X	X
		X = supported blank = not supported											

**Note:** RHEL5 is not supported in v9.3 or later version of LK.

		OS version / Architecture											
		RHEL6											
		6	6.1	6.2 Security Fix(*18)	6.3 Security Fix(*19)	6.4 Security Fix(*20)	6.5 Security Fix(*21)	6.6 Security Fix(*22)	6.7 Security Fix(*27)	6.8 Security Fix(*28)	6.9 Security Fix(*31)	6	6
		x86/x86_64											
HDLM	6.5.0												
	6.5.1												
	6.5.2	X											
	6.6.0	X											
	6.6.2	X											
	6.6.2-01	X	X										
	7.2.0	X	X	X									
	7.2.1	X	X	X									
	7.3.0	X	X	X									
	7.3.1	X	X	X									
	7.4.0	X	X	X	X	X	X	X	X	X			
	7.4.1	X	X	X	X	X	X	X	X	X			
	7.5.0	X	X	X	X	X	X	X	X	X			
	7.6.0	X	X	X	X	X	X	X	X	X			
	7.6.1	X	X	X	X	X	X	X	X	X			
	8.0.0	X	X	X	X	X	X	X	X	X	X		
8.0.1	X	X	X	X	X	X	X	X	X	X			
8.1.0	X	X	X	X	X	X	X	X	X	X			
8.1.1	X	X	X	X	X	X	X	X	X	X			

	8.1.2	X	X	X	X	X	X	X	X	X	X	
	8.1.3	X	X	X	X	X	X	X	X	X	X	
	8.1.4	X	X	X	X	X	X	X	X	X	X	
	8.2.0	X	X	X	X	X	X	X	X	X	X	
	8.2.1	X	X	X	X	X	X	X	X	X	X	
	8.4.0	X	X	X	X	X	X	X	X	X	X	
	8.5.0	X	X	X	X	X	X	X	X	X	X	
	8.5.1	X	X	X	X	X	X	X	X	X	X	
	8.5.2	X	X	X	X	X	X	X	X	X	X	
	8.5.3	X	X	X	X	X	X	X	X	X	X	
	8.5.4	X	X	X	X	X	X	X	X	X	X	
	8.6.0	X	X	X	X	X	X	X	X	X	X	
	8.6.1	X	X	X	X	X	X	X	X	X	X	
	8.6.2	X	X	X	X	X	X	X	X	X	X	
	8.6.4	X	X	X	X	X	X	X	X	X	X	
	8.6.5	X	X	X	X	X	X	X	X	X	X	
	8.7.0	X	X	X	X	X	X	X	X	X	X	
	8.7.1	X	X	X	X	X	X	X	X	X	X	
	8.7.2	X	X	X	X	X	X	X	X	X	X	
	8.7.3	X	X	X	X	X	X	X	X	X	X	
	8.7.4	X	X	X	X	X	X	X	X	X	X	
	8.7.6	X	X	X	X	X	X	X	X	X	X	
	8.7.7	X	X	X	X	X	X	X	X	X	X	
	8.7.8	X	X	X	X	X	X	X	X	X	X	
8.8.0	X	X	X	X	X	X	X	X	X	X		
LifeKeeper	v7.0											
	(v7.0.0-5 or later)											
	V 7.1											
	(v7.1.0-8 or later)											
	V7.2											

(v7.2.0-10 or later)											
V 7.3											
(v7.3.0-21 or later)	X										
V 7.4											
(v7.4.0-63 or later)	X										
V 7.5											
(v7.5.0-3640 or later)	X	X	X	X							
v8.0											
(v8.0.0-510 or later)	X	X	X	X							
v8.1											
(v8.1.1-5620 or later)	X	X	X	X							
v8.1.2											
(v8.1.2-5795 or later)	X	X	X	X	X						
v8.2.0											
(v8.2.0-6213 or later)	X	X	X	X	X						
v8.2.1											
(v8.2.1-6353 or later)	X	X	X	X	X	X					
v8.3.0											
(v8.3.0-6389 or later)	X	X	X	X	X	X					
v8.3.1											
(v8.3.1-6397 or later)	X	X	X	X	X	X					
v8.3.2											
(v8.3.2-6405 or later)	X	X	X	X	X	X	X	X	X		

v8.4.0												
(v8.4.0-6427 or later)	X	X	X	X	X	X	X	X	X	X		
v8.4.1												
(v8.4.1-6449 or later)	X	X	X	X	X	X	X	X	X	X		
v9.0.0												
(v9.0.0-6488 or later)	X	X	X	X	X	X	X	X	X	X		
v9.0.1												
(v9.0.1-6492 or later)	X	X	X	X	X	X	X	X	X	X		
v9.0.2												
(v9.0.2-6513 or later)	X	X	X	X	X	X	X	X	X	X		
v9.1.0												
(v9.1.0-6538 or later)	X	X	X	X	X	X	X	X	X	X		
v9.1.1												
(v9.1.1-6594 or later)	X	X	X	X	X	X	X	X	X	X		
v9.1.2												
(v9.1.2-6609 or later)	X	X	X	X	X	X	X	X	X	X	X	
v9.2.0												
(v9.2.0-6629 or later)	X	X	X	X	X	X	X	X	X	X	X	
v9.2.1												
(v9.2.1-6653 or later)	X	X	X	X	X	X	X	X	X	X	X	
v9.2.2												
(v9.2.2-6679 or later)	X	X	X	X	X	X	X	X	X	X	X	
v9.3												
(v9.3.0-6738)	X	X	X	X	X	X	X	X	X	X	X	

	or later)											
	v9.3.1											
	(v9.3.1-6750 or later)	X	X	X	X	X	X	X	X	X	X	
	v9.3.2											
	(v9.3.2-6863 or later)	X	X	X	X	X	X	X	X	X	X	
	v9.4.0											
	(v9.4.0-6959 or later)	X	X	X	X	X	X	X	X	X	X	
	v9.4.1											
	(v9.4.1-6983 or later)	X	X	X	X	X	X	X	X	X	X	
	v9.5.0											
(v9.5.0-7075 or later)	X	X	X	X	X	X	X	X	X	X		
HDLM ARK	7.0.0-1											
	7.2.0-1	X	X	X	X							
	8.1.1-5620	X	X	X	X							
	8.1.2-5795	X	X	X	X	X						
	8.2.0-6213	X	X	X	X	X						
	8.2.1-6353	X	X	X	X	X	X	X	X	X		
	8.3.0-6389	X	X	X	X	X	X	X	X	X		
	8.3.1-6397	X	X	X	X	X	X	X	X	X		
	8.3.2-6405	X	X	X	X	X	X	X	X	X		
	8.4.0-6427	X	X	X	X	X	X	X	X	X		
	8.4.1-6449	X	X	X	X	X	X	X	X	X		
	9.0.0-6488	X	X	X	X	X	X	X	X	X		
	9.0.1-6492	X	X	X	X	X	X	X	X	X		
	9.0.2-6513	X	X	X	X	X	X	X	X	X		
	9.1.0-6538	X	X	X	X	X	X	X	X	X		
	9.1.1-6594	X	X	X	X	X	X	X	X	X		
9.1.2-6609	X	X	X	X	X	X	X	X	X	X		

	9.2.0-6629	X	X	X	X	X	X	X	X	X	X
	9.2.1-6653	X	X	X	X	X	X	X	X	X	X
	9.2.2-6679	X	X	X	X	X	X	X	X	X	X
	9.3.0-6738	X	X	X	X	X	X	X	X	X	X
	9.3.1-6750	X	X	X	X	X	X	X	X	X	X
	9.3.2-6863	X	X	X	X	X	X	X	X	X	X
	9.4.0-6959	X	X	X	X	X	X	X	X	X	X
	9.4.1-6983	X	X	X	X	X	X	X	X	X	X
	9.5.0-7075	X	X	X	X	X	X	X	X	X	X
	X = supported blank = not supported										

		RHEL7										
		7.0 Security Fix(*24)	7.1 Security Fix(*25)	7.2 Security Fix(*26)	7.3 Security Fix(*30)	7.4 Security Fix(*32)	7.5 Security Fix(*33)	7.6 Security Fix(*35)	7.7 Security Fix(*36)	7.8 Security Fix	7.9 Security Fix	
		x86/x86_64										
HDLM	8.0.1	X	X									
	8.1.0	X	X									
	8.1.1	X	X									
	8.1.2	X	X									
	8.1.3	X	X									
	8.1.4	X	X									
	8.2.0	X	X									
	8.2.1	X	X									
	8.4.0	X	X	X								
	8.5.0	X	X	X								
	8.5.1	X	X	X	X	X						
	8.5.2	X	X	X	X	X						
	8.5.3	X	X	X	X	X						
	8.5.4	X	X	X	X	X						

	8.6.0	X	X	X	X	X					
	8.6.1	X	X	X	X	X	X				
	8.6.2	X	X	X	X	X	X	X (Note: 3)			
	8.6.4	X	X	X	X	X	X	X			
	8.6.5	X	X	X	X	X	X	X			
	8.7.0	X	X	X	X	X	X	X	X		
	8.7.1	X	X	X	X	X	X	X	X		
	8.7.2	X	X	X	X	X	X	X	X		
	8.7.3	X	X	X	X	X	X	X	X	X	X
	8.7.4	X	X	X	X	X	X	X	X	X	X
	8.7.6	X	X	X	X	X	X	X	X	X	X
	8.7.7	X	X	X	X	X	X	X	X	X	X
	8.7.8	X	X	X	X	X	X	X	X	X	X
	8.8.0	X	X	X	X	X	X	X	X	X	X
	8.8.1	X	X	X	X	X	X	X	X	X	X
LifeKeeper	v9.0.0 (v9.0.0-6488 or later)	X	X								
	v9.0.1 (v9.0.1-6492 or later)	X	X								
	v9.0.2 (v9.0.2-6513 or later)	X	X	X							
	v9.1.0 (v9.1.0-6538 or later)	X	X	X							
	v9.1.1 (v9.1.1-6594 or later)	X	X	X	X						
	v9.1.2 (v9.1.2-6609 or later)	X	X	X	X						
	v9.2.0 (v9.2.0-6629	X	X	X	X	X					

	or later)										
	v9.2.1 (v9.2.1-6653 or later)	X	X	X	X	X					
	v9.2.2 (v9.2.2-6679 or later)	X	X	X	X	X					
	v9.3 (v9.3.0-6738 or later)	X	X	X	X	X	X				
	v9.3.1 (v9.3.1-6750or later)	X	X	X	X	X	X				
	v9.3.2 (v9.3.2-6863 or later)	X	X	X	X	X	X	X			
	v9.4.0 (v9.4.0-6959 or later)	X	X	X	X	X	X	X	X		
	v9.4.1 (v9.4.1-6983 or later)	X	X	X	X	X	X	X	X		
	v9.5.0 (v9.5.0-7075 or later)	X	X	X	X	X	X	X	X	X	
	v9.5.1 (v9.5.1-7154 or later)	X	X	X	X	X	X	X	X	X	X
	v9.5.2 (v9.5.2-7301 or later)	X	X	X	X	X	X	X	X	X	X
HDLM ARK	9.0.0-6488	X	X								
	9.0.1-6492	X	X								
	9.0.2-6513	X	X	X							
	9.1.0-6538	X	X	X							
	9.1.1-6594	X	X	X	X						
	9.1.2-6609	X	X	X	X						
	9.2.0-6629	X	X	X	X	X					

	9.2.1-6653	X	X	X	X	X					
	9.2.2-6679	X	X	X	X	X					
	9.3.0-6738	X	X	X	X	X	X				
	9.3.1-6750	X	X	X	X	X	X				
	9.3.2-6863	X	X	X	X	X	X	X			
	9.4.0-6959	X	X	X	X	X	X	X	X		
	9.4.1-6983	X	X	X	X	X	X	X	X		
	9.5.0-7075	X	X	X	X	X	X	X	X	X	
	9.5.1-7154	X	X	X	X	X	X	X	X	X	X
	9.5.2-7301	X	X	X	X	X	X	X	X	X	X

X = supported blank = not supported

**Note: 1** If you are running the system with LifeKeeper v9.0.x on RHEL7/7.1/7.2, you need to apply the Bug7205's patch.

**Note: 2** The Raw device configuration is not supported on RHEL7/7.1/7.2/7.3/7.4/7.5/7.6/7.7/7.8.

**Note: 3** Supported with HDLM 8.6.2-02 or later.

		OS version / Architecture									
		RHEL8									
		8.1 Security Fix(*37)	8.2 Security Fix(*38)(*39)	8.3 Security Fix(*40)							
		x86/x86_64									
HDLM	8.7.2	X									
	8.7.3	X									
	8.7.4	X	X		X						
	8.7.6	X	X		X						
	8.7.7	X	X		X						
	8.7.8	X	X		X						
	8.8.0	X	X		X						
LifeKeeper	v9.5.1 (9.5.1-7154 以降)	X	X		X						

HDLM ARK	9.5.1-7154	X	X	X									
X = supported, blank = not supported.													

**Note:** The Raw device configuration is not supported on RHEL 8.1, RHEL 8.2 and RHEL 8.3.

## Device Mapper Multipath I/O Configurations

Protecting Applications and File Systems That Use Device Mapper Multipath Devices	<p>In order for LifeKeeper to operate with and protect applications or file systems that use Device Mapper Multipath devices, the Device Mapper Multipath (DMMP) Recovery Kit must be installed.</p> <p>Once the DMMP Kit is installed, simply creating an application hierarchy that uses one or more of the multipath device nodes will automatically incorporate the new resource types provided by the DMMP Kit.</p>
Multipath Device Nodes	<p>To use the DMMP Kit, any file systems and raw devices must be mounted or configured on the multipath device nodes rather than on the native <code>/dev/sd*</code> device nodes. The supported multipath device nodes to address the full disk are <code>/dev/dm-#</code>, <code>/dev/mapper/&lt;uuid&gt;</code>, <code>/dev/mapper/&lt;user_friendly_name&gt;</code> and <code>/dev/mpath/&lt;uuid&gt;</code>. To address the partitions of a disk, use the device nodes for each partition created in the <code>/dev/mapper</code> directory.</p>
Use of SCSI-3 Persistent Reservations	<p>The Device Mapper Multipath Recovery Kit uses SCSI-3 persistent reservations with a "Write Exclusive" reservation type. This means that devices reserved by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other nodes will be unable to write to the device. Note that <b>this does not mean</b> that you can expect to be able to mount file systems on those other nodes for ongoing read-only access.</p> <p>LifeKeeper uses the <code>sg_persist</code> utility to issue and monitor persistent reservations. If necessary, LifeKeeper will install the <code>sg_persist(8)</code> utility.</p> <p>SCSI-3 Persistent Reservations must be enabled on a per LUN basis when using EMC Symmetrix (including VMAX) arrays with multipathing software and LifeKeeper. This applies to both DMMP and PowerPath.</p>
Hardware Requirements	<p>The Device Mapper Multipath Kit has been tested by SIOS Technology Corp. with the EMC CLARiiON CX300, the HP EVA 8000, HP MSA1500, HP P2000, the IBM SAN Volume Controller (SVC), the IBM DS8100, the IBM DS6800, the IBM ESS, the DataCore SANsymphony, and the HDS 9980V. Check with your storage vendor to determine their support for Device Mapper Multipath.</p> <p>Enabling support for the use of reservations on the CX300 and the VNX Series requires that the hardware handler be notified to honor reservations. Set the following parameter in <code>/etc/multipath.conf</code> for this array:</p> <pre>hardware_handler    "3 emc 0 1"</pre>

	<p>The HP MSA1500 returns a reservation conflict with the default path checker setting (tur). This will cause the standby node to mark all paths as failed. To avoid this condition, set the following parameter in <i>/etc/multipath.conf</i> for this array:</p> <pre>path_checker    readsector0</pre> <p>The HP 3PAR F400 returns a reservation conflict with the default path checker. To avoid this conflict, set (add) the following parameter in <i>/etc/default/LifeKeeper</i> for this array:</p> <pre>DMMP_REGISTRATION_TYPE=hba</pre> <p>For the HDS 9980V the following settings are required:</p> <ul style="list-style-type: none"> <li>• Host mode: 00</li> <li>• System option: 254 (must be enabled; global HDS setting affecting all servers)</li> <li>• Device emulation: OPEN-V</li> </ul> <p>Refer to the HDS documentation “Suse Linux Device Mapper Multipath for HDS Storage” or “Red Hat Linux Device Mapper Multipath for HDS Storage” v1.15 or later for details on configuring DMMP for HDS. This documentation also provides a compatible multipath.conf file.</p> <p>For the EVA storage with firmware version 6 or higher, DMMP Recovery Kit v6.1.2-3 or later is required. Earlier versions of the DMMP Recovery Kit are supported with the EVA storage with firmware versions prior to version 6.</p>
<p>Multipath Software Requirements</p>	<p>For SUSE, multipath-tools-0.4.5-0.14 or later is required.</p> <ul style="list-style-type: none"> <li>• For Red Hat, device-mapper-multipath-0.4.5-12.0.RHEL4 or later is required.</li> <li>• It is advised to run the latest set of multipath tools available from the vendor. The feature content and the stability of this multipath product are improving at a very fast rate.</li> </ul>
<p>Linux Distribution Requirements</p>	<p>Some storage vendors such as IBM have not certified DMMP with SLES 11 at this time.</p> <ul style="list-style-type: none"> <li>• SIOS Technology Corp. is currently investigating reported issues with DMMP, SLES 11, and EMCs CLARiiON and Symmetrix arrays.</li> </ul>
<p>Transient path failures</p>	<p>While running IO tests on Device Mapper Multipath devices, it is not uncommon for actions on the SAN, for example, a server rebooting, to cause paths to temporarily be reported as failed. In most cases, this will simply cause one path to fail leaving other paths to send IOs down resulting in no observable failures other than a small performance impact. In some cases, multiple paths can be reported as failed leaving no paths working. This can cause an application, such as a file system or database, to see IO errors. There has been much improvement in Device Mapper Multipath and the vendor support to eliminate these failures. However, at times, these can still be seen. To avoid these situations, consider these actions:</p>

1. Verify that the multipath configuration is set correctly per the instructions of the disk array vendor.
2. Check the setting of the “failback” feature. This feature determines how quickly a path is reactivated after failing and being repaired. A setting of “immediate” indicates to resume use of a path as soon as it comes back online. A setting of an integer indicates the number of seconds after a path comes back online to resume using it. A setting of 10 to 15 generally provides sufficient settle time to avoid thrashing on the SAN.
3. Check the setting of the “no\_path\_retry” feature. This feature determines what Device Mapper Multipath should do if all paths fail. We recommend a setting of 10 to 15. This allows some ability to “ride out” temporary events where all paths fail while still providing a reasonable recovery time. The LifeKeeper DMMP kit will monitor IOs to the storage and if they are not responded to within four minutes LifeKeeper will switch the resources to the standby server. **NOTE:** LifeKeeper does not recommend setting “no\_path\_retry” to “queue” since this will result in IOs that are not easily killed. The only mechanism found to kill them is on newer versions of DM, the settings of the device can be changed:

```
/sbin/dmsetup message -u `DMid` 0 fail_if_no_path
```

This will temporarily change the setting for no\_path\_retry to fail causing any outstanding IOs to fail. However, multipathd can reset no\_path\_retry to the default at times. When the setting is changed to fail\_if\_no\_path to flush failed IOs, it should then be reset to its default prior to accessing the device (manually or via LifeKeeper).

If “no\_path\_retry” is set to “queue” and a failure occurs, LifeKeeper will switch the resources over to the standby server. However, LifeKeeper will not kill these IOs. The recommended method to clear these IOs is through a reboot but can also be done by an administrator using the dmsetup command above. If the IOs are not cleared, then data corruption can occur if/when the resources are taken out of service on the other server thereby releasing the locks and allowing the “old” IOs to be issued.

## **5.4.2.10. LifeKeeper I/O Fencing Introduction**

I/O fencing is the locking away of data from a malfunctioning node preventing uncoordinated access to shared storage. In an environment where multiple servers can access the same data, it is essential that all writes are performed in a controlled manner to avoid data corruption. Problems can arise when the failure detection mechanism breaks down because the symptoms of this breakdown can mimic a failed node. For example, in a two-node cluster, if the connection between the two nodes fails, each node would “think” the other has failed, causing both to attempt to take control of the data resulting in data corruption. I/O fencing removes this data corruption risk by blocking access to data from specific nodes.

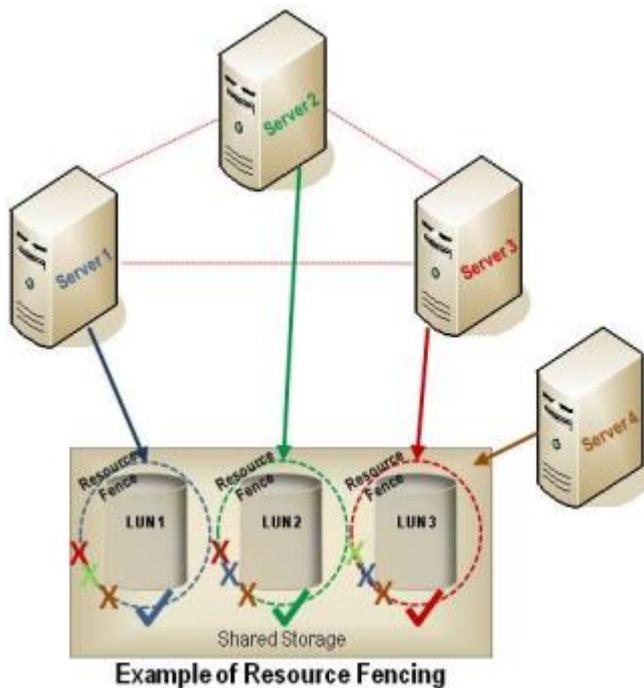
For DataKeeper, please refer to [DataKeeper I-O Fencing Introduction](#).

## 5.4.2.10.1. SCSI Reservations

### Storage Fencing Using SCSI Reservations

While LifeKeeper for Linux supports both resource fencing and node fencing, its primary fencing mechanism is storage fencing through SCSI reservations. This fence, which provides the highest level of data protection for shared storage, allows for maximum flexibility and maximum security providing very granular locking to the LUN level. The underlying shared resource (LUN) is the primary quorum device in this architecture. Quorum can be defined as exclusive access to shared storage, meaning this shared storage can only be accessed by one server at a time. The server who has quorum (exclusive access) owns the role of “primary.” The establishment of quorum (who gets this exclusive access) is determined by the “quorum device.”

As stated above, with reservations enabled, the quorum device is the shared resource. The shared resource establishes quorum by determining who owns the reservation on it. This allows a cluster to continue to operate down to a single server as long as that single server can access the LUN.



SCSI reservations protect the shared user data so that only the system designated by LifeKeeper can modify the data. No other system in the cluster or outside the cluster is allowed to modify that data. SCSI reservations also allow the application being protected by LifeKeeper to safely access the shared user data when there are multiple server failures in the cluster. A majority quorum of servers is not required; the only requirement is establishing ownership of the shared data.

Adding quorum/witness capabilities provides for the establishment of quorum membership. Without this membership, split-brain situations could result in multiple servers, even all servers, killing each other. Watchdog added to configurations with reservations enabled provides a mechanism to recover from partially hung servers. In cases where a hung server goes undetected by LifeKeeper, watchdog will begin recovery. Also, in the case where a server is hung and not able to detect that the reservation has

been stolen, watchdog can reboot the server to begin its recovery.

## Alternative Methods for I/O Fencing

In addition to resource fencing using SCSI reservations, LifeKeeper for Linux also supports disabling reservations. Regardless of whether reservations are enabled or disabled, there are two issues to be aware of:

- Access to the storage must be controlled by LifeKeeper.
- Great care must be taken to ensure that the storage is not accessed unintentionally such as by mounting file systems manually, fsck manually, etc.

If these two rules are followed and reservations are enabled, LifeKeeper will prevent most errors from occurring. With reservations disabled (alone), there is no protection. Therefore, other options must be explored in order to provide this protection. The following sections discuss these different fencing options and alternatives that help LifeKeeper provide a reliable configuration even without reservations.

## 5.4.2.10.2. Disabling Reservations

While reservations provide the highest level of data protection for shared storage, in some cases, the use of reservations is not available and must be disabled within LifeKeeper. With reservations disabled, the storage no longer acts as an arbitrator in cases where multiple systems attempt to access the storage, intentionally or unintentionally.

Consideration should be given to the use of other methods to fence the storage through cluster membership which is needed to handle system hangs, system busy situations and any situation where a server can appear to not be alive.

The key to a reliable configuration without reservations is to “know” that when a failover occurs, the “other” server has been powered off or power cycled. There are four fencing options that help accomplish this, allowing LifeKeeper to provide a very reliable configuration, even without SCSI reservations. These include the following:

- [STONITH](#) (Shoot the Other Node in the Head) using a highly reliable interconnect, i.e. serial connection between server and STONITH device. STONITH is the technique to physically disable or power-off a server when it is no longer considered part of the cluster. LifeKeeper supports the ability to power off servers during a failover event thereby insuring safe access to the shared data. This option provides reliability similar to reservations but is limited to two nodes physically located together.
- [Quorum/Witness](#) – Quorum/witness servers are used to confirm membership in the cluster, especially when the cluster servers are at different locations. While this option can handle split-brain, it, alone, is not recommended due to the fact that it does not handle system hangs.
- [Watchdog](#) – Watchdog monitors the health of a server. If a problem is detected, the server with the problem is rebooted or powered down. This option can recover from a server hang; however, it does not handle split-brain; therefore this option alone is also not recommended.
- `CONFIRM_SO` – This option requires that automatic failover be turned off, so while very reliable (depending upon the knowledge of the administrator), it is not as available.

While none of these alternative fencing methods alone are likely to be adequate, when used in combination, a very reliable configuration can be obtained.

## Non-Shared Storage

If planning to use LifeKeeper in a non-shared storage environment, the risk of data corruption that exists with shared storage is not an issue; therefore, reservations are not necessary. However, partial or full resyncs and merging of data may be required. To optimize reliability and availability, the above options should be considered with non-shared storage as well.



**Note:** For further information comparing the reliability and availability of the different

options, see the [I/O Fencing Comparison Chart](#).

It is important to note that no option will provide complete data protection, but the following combination will provide almost the same level of protection as reservations.

## Configuring I/O Fencing Without Reservations

To configure a cluster to support node fencing, complete the following steps:

1. Stop LifeKeeper.
2. Disable the use of SCSI reservations within LifeKeeper. This is accomplished by editing the LifeKeeper defaults file, `/etc/default/LifeKeeper`, on all nodes in the cluster. Add or modify the `Reservations` variable to be "none", e.g. `RESERVATIONS="none"`. (**Note:** This option should only be used when reservations are not available.)
3. Obtain and configure a STONITH device or devices to provide I/O fencing. Note that for this configuration, STONITH devices should be configured to do a system "poweroff" command rather than a "reboot". Take care to avoid bringing a device hierarchy in service on both nodes simultaneously via a manual operation when LifeKeeper communications have been disrupted for some reason.
4. If desired, obtain and configure a quorum/witness server(s). For complete instructions and information on configuring and using a witness server, see [Quorum/Witness Server Support Package](#) topic.



**Note:** The quorum/witness server should reside at a site apart from the other servers in the cluster to provide the greatest degree of protection in the event of a site failure.

5. If desired, configure watchdog. For more information, see the [Watchdog](#) topic.

## 5.4.2.10.2.1. I/O Fencing Chart

	Split-Brain	Hung Server
<b>Reservations On</b>		
Alone		
Quorum/Witness		
Watchdog		
Watchdog & Quorum/Witness		
STONITH (serial)		
<b>Reservations Off</b>		
Nothing		
STONITH (serial)		
CONFIRM_SO*		
Quorum/Witness		
Watchdog		
<b>Non-Shared Storage</b>		
Default Features		
Quorum/Witness		
CONFIRM_SO*		
Watchdog		
STONITH (serial)		





\* While `CONFIRM_SO` is highly reliable (depending upon the knowledge of the administrator), it has lower availability due to the fact that automatic failover is turned off.

## 5.4.2.10.3. Quorum/Witness

### Quorum/Witness Server Support Package for LifeKeeper

#### Feature Summary

The Quorum/Witness Server Support Package for LifeKeeper (steeleye-lkQWK, hereinafter “Quorum/Witness Package”) combined with the existing failover process of the LifeKeeper core allows system failover to occur with a greater degree of confidence in situations where total network failure could be common. This effectively means that local site failovers and failovers to nodes across a WAN can be done while greatly reducing the risk of [split-brain](#) situations.

In a distributed system that takes network partitioning into account, there is a concept called quorum to obtain consensus across the cluster. A node having quorum is a node that can obtain consensus of all the clusters and is allowed to bring resources in service. On the other hand, a node not having quorum is a node that cannot obtain consensus of all the clusters and it is not allowed to bring resources in service. This will prevent split brain from happening. To check whether a node has quorum is called quorum check. It is expressed as “quorum check succeeded” if it has quorum, and “quorum check failed” if it does not have quorum.

In case of a communication failure, using one node where failure occurred and another multiple nodes (or other devices) will allow a node to get a “second opinion” on the status of the failing node. The node to get a “second opinion” is called a witness node (or a witness device), and getting a “second opinion” is called witness checking. When determining when to fail over, the witness node (the witness device) allows resources to be brought in service on a backup server only in cases where it verifies the primary server has failed and is no longer part of the cluster. This will prevent failovers from happening due to simple communication failures between nodes when those failures don’t affect the overall access to, and performance of, the in-service node. During actual operation, the witness node (the witness device) will be consulted when LifeKeeper is started or the failed communication path is restored. Witness checking can only be performed for nodes having quorum.

✿ SIOS recommends using one quorum/witness server per cluster for a more reliable cluster, when you are not using shared storage.

✿ The Witness node system requirements are the same as the [LifeKeeper’s system requirements](#). Since the Witness node does not participate in replication or host any protected resources, it only needs to meet the minimum requirements.

#### Package Installation and Configuration

The Quorum/Witness Server Support Package for LifeKeeper will need to be installed on every node in the cluster that uses quorum/witness functionality, including a witness-only node. The only configuration requirement for the witness node is to [create appropriate comm paths](#). When using a quorum mode with

`tcp_remote`, LifeKeeper does not need to be installed on the host which was set as `QUORUM_MODE` in `/etc/default/LifeKeeper` configuration file.

The general process for setting up quorum/witness functionality will involve the following steps:

1. Set up the server and make sure that it can communicate with other servers.
2. Install LifeKeeper on the server. During the installation, enable “Use Quorum / Witness functions” with the setup command and install the quorum/witness package as well.
3. Create appropriate communication paths between the nodes including witness-only nodes.
4. [Configure quorum/witness](#).

When the above steps are completed, the quorum/witness functions will be activated in the cluster and quorum checking and witness checking will be performed before failovers are allowed.

See the Configurable Components section below for additional configuration options.

 **Note:** Any node that has the quorum/witness package installed can participate in quorum/witness functionality. The witness-only nodes will have communication paths with all the other nodes and will not host any protected resources.

## Configurable Components

The quorum/witness package contains two configurable modes: quorum and witness. By default, installing the quorum/witness package will enable both quorum and witness modes.

The behavior of these modes can be customized via the `/etc/default/LifeKeeper` configuration file, and the quorum and witness modes can be individually adjusted. The package installs default settings into the configuration file when it is installed, *majority* being the default quorum mode and *remote\_verify* being the default witness mode. An example is shown below:

```
QUORUM_MODE=majority
WITNESS_MODE=remote_verify
```

## Available Quorum Modes

Four quorum checking modes are available which can be set via the `QUORUM_MODE` setting in `/etc/default/LifeKeeper`.

QUORUM_MODE	Description
<i>majority</i> (default)	With majority as the quorum mode setting quorum checks occur via LifeKeeper for Linux communication paths. A node has quorum when it is able to communicate with the majority of the nodes in the cluster. This quorum mode is available on clusters with

	three or more nodes. A witness – only node needs to be added when using a two-node configuration. See “ <a href="#">majority mode</a> ” for details.
<i>tcp_remote</i>	Checks the connection to the TCP/IP service on the specified port for the host independent from the communication path. It is determined that the node has quorum when it is able to communicate with the majority of the nodes in the cluster. A host for connection checking is required separately. See “ <a href="#">tcp_remote mode</a> ” for details.
<i>storage</i>	With storage as the quorum mode setting quorum checks occur using a “shared storage” device. A node has quorum when it is able to access the shared storage device and update its own quorum object. A cluster is considered to have quorum consensus with this mode when each node is able to access the shared storage device. This mode can be used for 2, 3 or 4 node clusters. Shared storage devices can be a block storage device shared to all nodes, a file accessed via a NFS share on all nodes or an Amazon S3 storage object. The “shared storage” used for this solution must be obtained separately. See “ <a href="#">storage mode</a> ” for details. When <i>storage</i> is selected for the quorum mode, the witness mode, which is described later, must also be set to <i>storage</i> .
<i>none/off</i>	Quorum checking is disabled. With this configuration, quorum checking is always determined to be successful.

## Available Witness Modes

Three witness modes are available which can be set via the WITNESS\_MODE setting in `/etc/default/LifeKeeper`.

WITNESS_MODE	Description
<i>remote_verify</i> (default)	Consults all the other nodes in the cluster about their view of the status of a node which appears to be failing. If any node determines that there is no failure, witness checking determines that there is no failure. If all the nodes determine that there is failure, witness checking determines that the node is failing.
<i>storage</i>	A witness mode where shared storage is used as a witness device. The shared storage device is used to “share” status information between nodes in the cluster. Each node updates its own information and reads the other nodes information. If a node detects that information for another node is not being updated then that node will be considered failed. See “ <a href="#">storage mode</a> ” for details. When storage is selected for the witness mode, then storage must be selected for quorum mode. See above.
<i>none/off</i>	In this mode, witness checking is disabled. With this setting, it is always determined that there is no failure.

 **Note:** It would be unnecessary for witness checks to ever be performed by servers acting as dedicated quorum/witness nodes that do not host resources; therefore, this setting should be set to *none/off* on these servers.

## Supported Combinations of a Quorum Mode and Witness Mode

LifeKeeper supports the following combinations.

		QUORUM_MODE			
		<i>majority</i>	<i>tcp_remote</i>	<i>storage</i>	<i>none/off</i>
WITNESS_MODE	<i>remote_verify</i>	Supported 3 or more nodes	Supported 3 or more nodes	Not supported	Supported 3 or more nodes
	<i>storage</i>	Not Supported	Not Supported	Supported Between 2 and 4 nodes	Not supported
	<i>none/off</i>	Supported 3 or more nodes	Supported 2 or more nodes	Not supported	Supported

## Available Actions When Quorum is Lost

The quorum/witness package offers three different options for how the system should react if quorum is lost — “*fastboot*”, “*fastkill*” and “*osu*”. These options can be selected via the QUORUM\_LOSS\_ACTION setting in `/etc/default/LifeKeeper`. All three options take the system’s resources out of service; however, they each allow a different behavior.

Mode	Description
<i>fastboot</i>	<p>The system will be <b>immediately</b> rebooted when a loss of quorum is detected (from a communication path failure). Although this is an aggressive option, it ensures that the system will be disconnected from any external resources right away. In many cases, such as with storage-level replication, this immediate release of resources is desired.</p> <p>Two important notes on this option are:</p> <ol style="list-style-type: none"> <li>1. The system performs an <b>immediate</b> hard reboot without first performing any shut-down procedure; no shutdown tasks are performed (disk syncing, etc.).</li> <li>2. The system will come back up performing normal startup routines, including negotiating storage and resource access, etc.</li> </ol>
<i>fastkill</i> (default)	<p>The fastkill option is very similar to the fastboot option, but instead of a hard reboot, the system will <b>immediately</b> halt when quorum is lost. As with the fastboot option, no tasks are performed (disk syncing, etc.), and the system will then need to be manually started and will come back up performing normal startup routines, including negotiating storage and resource access, etc.</p>
<i>osu</i>	<p>This is the least aggressive option, leaving the system operational but taking resources out of service on the system where quorum is lost. In some cluster configurations, this is all that is</p>

	needed, but it may not be strong enough or fast enough in others.
--	-------------------------------------------------------------------

## 5.4.2.10.3.1. Majority Mode

Quorum checking is performed via LifeKeeper for Linux communication paths. A node has quorum when it is able to communicate with the majority of the nodes in the cluster. This quorum mode is available on clusters with three or more nodes. A node dedicated for witness checking needs to be added when using a two-node configuration.

 **Note:** Due to requirements for node majority, it is recommended that clusters always be configured with an odd number of nodes (the count includes the quorum node).

### Stopping the cluster with quorum:

1. Stop the target
2. Stop the source
3. Stop the witness

### Starting the cluster with quorum:

1. Start the witness
2. Start the target
3. Start the source

## Majority Mode Configuration

Set QUORUM\_MODE to majority in `/etc/default/LifeKeeper`. No other setting is required for this mode.

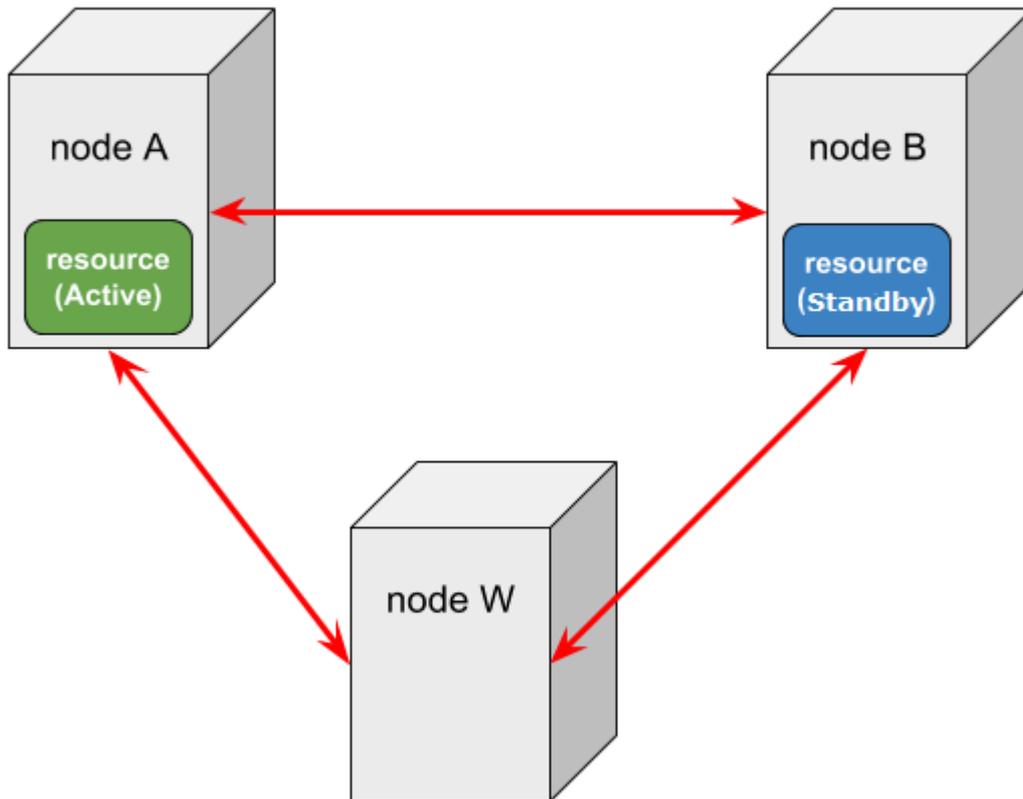
## Available Witness mode settings for Majority Mode

The following witness modes are available for majority mode. For details on each mode, please refer to "[Available Witness Mode](#)".

- *remote\_verify*
- *none/off*

## Expected Behaviors for Majority Mode (Assuming Default Modes)

The scenarios listed below shows the LifeKeeper for Linux behavior of a three-node cluster with Node A (resources are in-service), Node B (resources are on stand-by), and Node W (a witness-only node without protected resources).



The following three events may change the resource status on a node failure.

- **COMM\_DOWN** event  
An event called when all the communication paths between nodes are disconnected.
- **COMM\_UP** event  
An event called when communication paths are recovered from a **COMM\_DOWN** state.
- **LCM\_AVAIL** event  
An event called after [LCM](#) initialization is completed and it is called only once when starting LifeKeeper. Once this state has been reached heartbeat transmission to other nodes in the cluster begins over the established communication paths. It also ready to receive heartbeat requests from other nodes cluster. **LCM\_AVAIL** is always processed before processing a **COMM\_UP** event.

## Scenario 1

### A communication path fails between Node A and B

In this case, the following will happen:

1. Both Node A and Node B will begin processing **COMM\_DOWN** events, though not necessarily at exactly the same time.
2. Both nodes will perform the quorum check and determine that they still have quorum (since both Node A and B can see Node W and they have communication with two of the three known nodes,

they think that they are in the majority).

3. Each will consult the other nodes with whom they can still communicate about the true status of the server with whom they've lost communications (witness checking). In this scenario, this means that Node A will consult Node W about Node B's status and Node B will also consult Node W about Node A's status.
4. Node A and Node B will both determine that the other is still alive by having consulted Node W and no failover processing will occur. Resources will be left in service on Node A.

## Scenario 2

### **A communication path fails between Node A and Node W**

Since all nodes can and will act as witness nodes when the quorum/witness package is installed, this scenario is the same as the previous. In this case, Node A and Node W will determine that the other is still alive by consulting with Node B.

## Scenario 3

### **Node A fails and stops**

In this case, Node B will do the following:

1. Begin processing the COMM\_DOWN event from Node A.
2. Determine that it can still communicate with Node W and thus has quorum.
3. Verify via Node W that Node A really appears to be lost and, begin the usual failover activity.
4. Node B will continue processing the event and bring the protected resources in service.

### **With resources being in-service on Node B, Node A is powered on and establishes communications with the other nodes**

In this case, Node A will process an LCM\_AVAIL event. Node A will determine that it has quorum and not bring resources in service because they are currently in service on Node B. Next, a COMM\_UP event will be processed between Node A and Node B and also between Node A and Node W (processed twice at Node A). Each node will determine that it has quorum during the COMM\_UP events and will not bring resources in service because they are currently in service on Node B.

### **With resources being in-service on Node B, Node A is powered on and cannot establish communications to the other nodes**

In this case, Node A will process an LCM\_AVAIL event and Node B and Node W will do nothing since they can't communicate with Node A. Node A will determine that it does not have quorum since it can only communicate with one of the three nodes (Node A itself). Because it does not have quorum, Node A will not bring resources in service.

## Scenario 4

### **A failure occurs with the network for Node A (Node A is running without communications to other nodes)**

In this case, Node A will do the following:

1. Begin processing a COMM\_DOWN event from Node B (processing of a COMM\_DOWN event from Node W is started almost simultaneously).
2. Determine that it cannot communicate with Node B or Node W and thus does not have quorum.
3. LifeKeeper takes action based on the QUORUM\_LOSS\_ACTION (see [Quorum/Witness](#) for more details).

Node B will do the following:

1. Begin processing a COMM\_DOWN event from Node A.
2. Determine that it can still communicate with Node W and thus has quorum.
3. Verify via Node W that Node A really appears to be lost (witness checking) and, begin the usual failover activity.
4. Node B will now have the protected resources in service.

### **With resources being in-service at Node B, communication resumes for Node A**

In this case, Node B will process a COMM\_UP event, determine that it has quorum (all three of the nodes are visible) and that it has the resources in service. Node A will process a COMM\_UP event, determine that it also has quorum and that the resources are in service on Node B. Node A will not bring resources in service at this time.

## Scenario 5

### **All three nodes lose communications with each other**

In this case, Node A will do the following:

1. Begin processing COMM\_DOWN events between node B. (Processing of a COMM\_DOWN event from Node W is started almost simultaneously).
2. Determine that it cannot communicate with Node B or Node W and thus does not have quorum.
3. LifeKeeper takes action based on the QUORUM\_LOSS\_ACTION (see [Quorum/Witness](#) for more details).

Node B will do the following:

1. Begin processing a COMM\_DOWN event between Node A.
2. Determine that it cannot communicate with Node A or Node W and thus does not have quorum.
3. Since it does not have the resources in service, no QUORUM\_LOSS\_ACTION will occur.

If all the communication paths are recovered, Node A will bring the resources in service. The following requirements should be met for this behavior.

- As initialization behavior, AUTORES\_ISP is set for the resources on Node A.
- The Resource Priority value is the highest on Node A.

## 5.4.2.10.3.2. tcp\_remote Mode

In this setting Quorum is determined by checking the ability to connect to TCP/IP services on remote hosts. Connections are done via TCP/IP to a specific port on a host and are done independent of any of the defined communication paths. Being able to connect to the majority of the specified hosts will determine if the node has quorum. Host and port combinations are defined via the QUORUM\_HOSTS setting discussed below. This mode is also available with a two-node cluster, however, three nodes are required when using a witness node for *remote\_verify*.

\* **Note:** Due to majority-based quorum, it is recommended that the hosts always be specified with an odd number of nodes.

\* **Note:** Due to the inherent flexibility and complexity of this mode, it should be used with caution by someone experienced with both LifeKeeper and the particular network/cluster configuration involved.

### tcp\_remote Mode Configuration

Set QUORUM\_MODE> to *tcp\_remote* in `/etc/default/LifeKeeper`. The following configuration settings are required when using *tcp\_remote*:

- QUORUM\_HOSTS – This is a comma delimited list of host:port values used to define the hosts and ports to connect to when checking for quorum.
- QUORUM\_TIMEOUT\_SECS – This is the time allowed for TCP/IP connections to complete. It defaults to 20 seconds.

See “[Quorum Parameter List](#)” for more information.

### Available Witness Mode setting with tcp\_remote mode

The following witness mode settings are available with *tcp\_remote*. Refer to “[Available Witness Mode](#)” for more details on each mode.

- *remote\_verify*
- *none/off*

## 5.4.2.10.3.3. Storage Mode

With this mode each node writes information about itself to a shared storage device on a regular basis and periodically reads the information written by the other nodes. A cluster is considered to have quorum consensus when each node is able to access the shared storage device and update its quorum object as well as see that the quorum objects for all other nodes are being updated. The node information located on the shared storage device is called a quorum (QWK) object or QWK object for short. QWK objects are required for every node configured in the cluster.

Quorum checking determines that a node has quorum when it has access to the shared storage device. Witness checking accesses the QWK objects for the other nodes to determine that node's current state. During a check it is verifying that updates to the QWK objects of the other nodes are still occurring on a regular basis. If no updates have occurred on a particular node after a certain period of time, the node will be considered in a failed state. During this time the checking node will update its own QWK object. Witness checking is performed when quorum checking is performed.

When “*storage*” is selected for quorum mode, “*storage*” must be selected for witness mode.

This quorum mode setting can be used for a two-node, three node, or four node cluster. The shared storage used for storing QWK objects for all the nodes must be configured separately. If a node loses access to the shared storage, it affects bringing resources in service. Select a shared storage device which is always accessible from all the nodes.

 **Note:** Using Storage for the Quorum Mode requires a storage device that can be accessed by all nodes in the cluster. The storage solution is to be used for quorum / witness functionality and must not be protected by LifeKeeper for Linux. For supported storage solutions see the topic on Available Share Storage.

 **Note:** In order to use this mode, initialization of the QWK object is required after configuring (See “[Storage Mode Configuration](#)”). In addition, reinitialization is necessary to add/delete nodes in the cluster or change the configuration after initial configuration.

 **Note:** This mode cannot be used if the names of the nodes in the cluster are similar such that the only difference is in the use of ‘-’ and ‘.’. For example a cluster with nodes named lifekeeper-sios and lifekeeper.sios would not be allowed but a cluster with nodes named lifekeeper-sios and lifekeeper.sios2 would be acceptable.

## Available Shared Storage

The purpose of the quorum/witness function is to avoid a split brain scenario. Therefore, correctly configuring the storage quorum mode choice is critical to ensure all nodes in the cluster can see all the QWK objects. This is accomplished by placing all the QWK objects in the same type of shared storage:

block devices, regular files, EFS or S3 objects.

The available shared storage choices are shown below. Specify the type of shared storage being used via the QWK\_STORAGE\_TYPE setting in the /etc/default/LifeKeeper configuration file.

QWK_STORAGE_TYPE	QWK Object Location
block	<p>When using physical storage, RDM (physical compatibility), iSCSI (in-VM initiator) for shared storage, allocate one QWK object in one of the following ways:</p> <p style="text-align: center;">(a) <b>1 QWK object = 1 partition</b></p> <p style="text-align: center;">(b) <b>1 QWK object = 1 LU</b></p> <p>In the case of (a), since multiple hosts will write to one LU, align the offset with 4K (sector size of the storage device) within the LU of the partition. Also, do not mix partitions used for other purposes.</p> <p>In the case of (b), do not create partitions within LU.</p> <p>No file system needs to be created for either (a) or (b).</p>
	<p>When using VMDK for shared storage, allocate one QWK object as follows:</p> <p style="text-align: center;"><b>1 QWK object = 1 VMDK</b></p> <p>Do not create partitions. Also, no file system needs to be created.</p> <p>Set the provisioning option for VMDK as follows:</p> <p style="text-align: center;">thick (eager zeroed)</p>
file (Supports NFS and EFS)	<p>When using either EFS or NFS for shared storage, allocate one QWK object as follows:</p> <p style="text-align: center;"><b>1 QWK object = 1 regular file system in the NFS or EFS file system</b></p> <p><b>Note:</b> The following options are not relevant for EFS, however the filesystem should be placed in /etc/fstab. EFS has it's own set of parameters.</p> <p>Set the export option for the NFS server (not applicable for EFS) as follows:</p> <p style="text-align: center;">rw,no_root_squash,sync,no_wdelay</p> <p>Set the mount option for the NFS server as follows:</p>

	<p style="text-align: center;">soft,timeo=20,retrans=1,noac</p> <p>Configure <code>/etc/fstab</code> to mount automatically after rebooting the OS.</p> <p><a href="#">For EFS configuration, please reference this AWS article to set up EFS.</a></p>
aws_s3	<p>When using Amazon Simple Storage Service (S3) for shared storage, allocate one QWK object as follows:</p> <p style="text-align: center;"><b>1 QWK object = 1 S3 object</b></p> <p>Use S3 in a region different from the region where LifeKeeper is running. Also, due to the Amazon S3 Data Consistency Model, the old data may be returned if the request is made right after updating the QWK objects; therefore, two QWK objects can be specified on one node when using S3(this is only available with S3).</p> <p>All of the nodes configured in the cluster need to satisfy the following requirements:</p> <ul style="list-style-type: none"> <li>• AWS Command Line Interface (AWS CLI) is installed and available to the root user. See “<a href="#">Installing the AWS Command Line Interface</a>”.</li> <li>• Ability to access the endpoint in Amazon S3 (<a href="#">the AWS region and endpoint</a>) with the HTTP and HTTPS protocols.</li> <li>• Ability to access the S3 object as the root user by properly configuring <a href="#">the IAM role for EC2</a> and the <a href="#">AWS CLI</a></li> </ul> <p><b>Note:</b> If the path name for the AWS CLI executable files are not already specified as a part of the “PATH” parameter in the LifeKeeper defaults file <code>/etc/default/LifeKeeper</code>, you must append the path to the AWS CLI executables for LifeKeeper to function correctly when using S3 objects.</p>

The size of 1 QWK object is 4096 bytes.

Quorum witness checking performs a read and/or /write to its own QWK object and will only read the QWK objects of other nodes. Set the access rights appropriately (be careful of permission restrictions such as granting Persistent Reservation to the shared storage).

 **Note:** 4K native disk of vSphere cannot be used as QWK object’s shared storage.

## Storage Mode Configuration

QUORUM\_MODE and WITNESS\_MODE should be configured as storage in the `/etc/default/LifeKeeper` configuration file. The following configuration parameters are also available when using storage:

- QWK\_STORAGE\_TYPE – Specifies the type of shared storage being used.
- QWK\_STORAGE\_HBEATTIME – Specifies the interval in seconds between reading and writing the QWK objects. This setting must be greater than or equal to the LCMHBEATTIME default setting.
- QWK\_STORAGE\_NUMHBEATS – Specifies the number of consecutive heartbeat checks that when missed indicates the target node has failed. A missed heartbeat occurs when the QWK object has not been updated since the last check. This setting must be greater than or equal to the LCMNUMHBEATS default setting.

**Note:** Based on the added traffic and no traffic time comparisons, you can **tune the number of heartbeats and the time** mentioned above. Defaults are 6 (minimum of 5, maximum of 10) seconds for heart beat time and 4 (minimum of 3) missed heart beats.

In the `/etc/default/LifeKeeper` file, **SIOS recommends** editing the QWK\_STORAGE\_NUMHBEATS value, changing it to 9.

```
QWK_STORAGE_NUMHBEATS=9
```

- QWK\_STORAGE\_OBJECT\_ – Specifies the path to the QWK object for each node in the cluster. Entries for all nodes in the cluster are required.
- HTTP\_PROXY, HTTPS\_PROXY, NO\_PROXY – Set this parameter when using HTTP proxy for accessing the service endpoint. The value set here will be passed to AWS CLI.

See the [“Quorum Parameter List”](#) for more information.

## How to use Storage Mode

Initialization is required in order to use the storage quorum mode. The initialization steps for all the nodes in the cluster are as follows.

1. Set up all the nodes and make sure that they can communicate with each other.
2. On all the nodes run the LifeKeeper for Linux setup and enable “Use Quorum/Witness Functions” to install the Quorum/Witness package.
3. Create communication paths between all the nodes.
4. Configure the quorum setting in the `/etc/default/LifeKeeper` configuration file on all nodes.

 In order to use storage quorum, **SIOS recommends** that you increase the **LCMHEARTBEATS** to allow for a longer time before the path is marked as failed. This will change the timeout period from the default of 15 seconds to 45 seconds.

**Edit the** `/etc/default/LifeKeeper` file and change **LCMNUMHBEATS to 9:**

```
LCMNUMHBEATS=9
```

 **Verify that all comm paths are up and ALIVE before running** `qwk_storage_init`.

5. Run the `qwk_storage_init` command on all nodes. This command will wait until the initialization of the QWK objects on all nodes is complete. Quorum/Witness functions will become available in the storage mode once the init completes on all nodes.

Reinitialization is necessary to add/delete cluster nodes after initial configuration, or when quorum parameters are changed in the `/etc/default/LifeKeeperconfiguration` file. Please reinitialize according to the following steps.

1. Execute the `qwk_storage_exit` command on all nodes.
2. Delete communication paths between the node that is being deleted and all the other nodes. Create communication paths between the node that is being added and all the other nodes.
3. Modify the quorum parameters in the `/etc/default/LifeKeeper` configuration file on all nodes.
4. Execute the `qwk_storage_init` command on all nodes.

## Troubleshooting

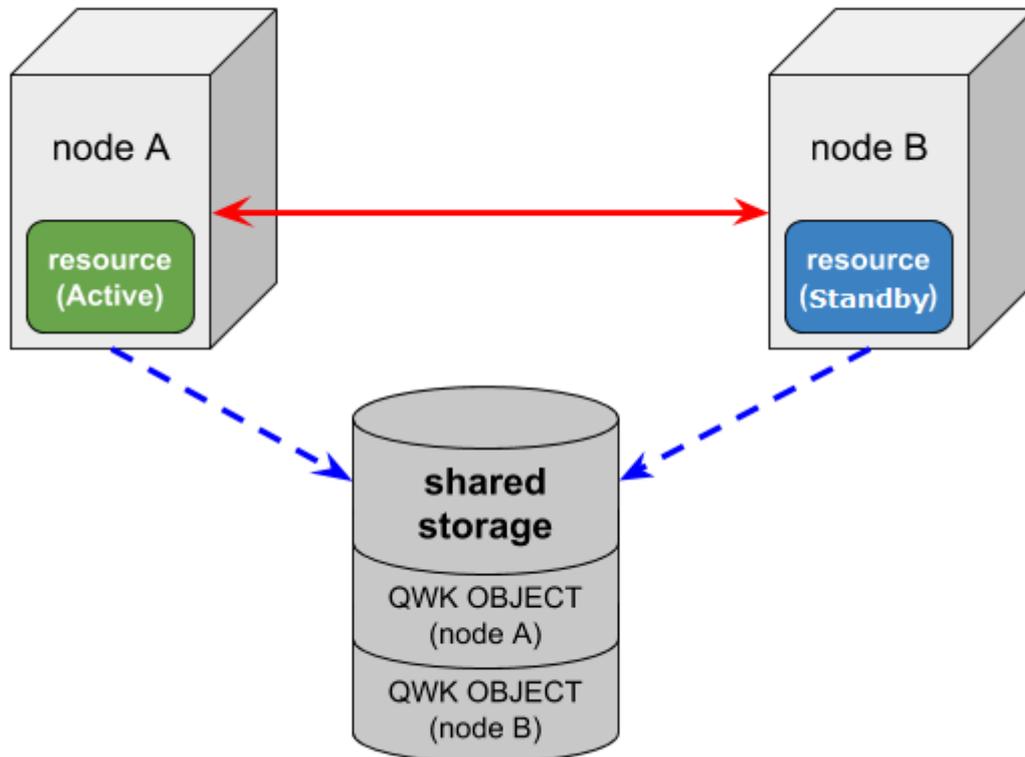
The frequent logging of message ID 135802 in the `lifekeeper.log` indicates that the periodic reading and writing of the QWK objects is overloaded and causing a delay in processing. If the number of nodes in the cluster is large and S3 is used for the shared storage (especially when used in two regions), the load from reading and writing of the QWK objects can be high.

Follow these steps to avoid the frequent logging of message ID 135802 in the `lifekeeper.log`.

- Increase the value of `QWK_STORAGE_HBEATTIME` (the time required for a failure to be detected and for a failover to begin will increase)
- If the shared storage choice is S3, use only in regions with the lowest network latency.
- Increase the throughput (for Amazon EC2 instances, change the instance type, etc.)

## Expected Behaviors for Storage Mode (Assuming Default Modes)

Behavior of a two-node cluster; Node A (resources are in-service) and Node B (resources are on stand-by), is shown below.



The following three events may change the resource status on a node failure:

- **COMM\_DOWN** event  
An event called when all the communication paths between the nodes are disconnected.
- **COMM\_UP** event  
An event called when the communication paths are recovered from a COMM\_DOWN state.
- **LCM\_AVAIL** event  
An event called after [LCM](#) initialization is completed and it is called only once when starting LifeKeeper. Once this state has been reached heartbeat, transmission to other nodes in the cluster begins over the established communication paths. It is also ready to receive heartbeat requests from other nodes in the cluster. LCM\_AVAIL will always processed before processing a COMM\_UP event.

### Scenario 1

**The communication paths fail between Node A and Node B (Both Node A and Node B can access**

**the shared storage)**

In this case, the following will happen:

1. Both Node A and Node B will begin processing a COMM\_DOWN event, though not necessarily at exactly the same time.
2. Both nodes will perform the quorum check and determine that they still have quorum (both A and B can access the shared storage).
3. Each node will check the QWK object for the node with whom it has lost communication to see if it is still being updated on a regular basis. Both nodes will find that the other's QWK object is being updated on a regular as both nodes are still running witness checks.
4. It will be determined, via the witness checking on each node, that the other is still alive so no failover processing will take place. Resources will be left in service at Node A.

## Scenario 2

**Node A fails and stops**

In this case, Server B will do the following:

1. Begin processing a COMM\_DOWN event from Node A.
2. Determine that it can still access the shared storage and thus has quorum.
3. Check to see that updates to the QWK object for Node A have stopped (witness checking).
4. Verify via witness checking that Node A really appears to be lost and begins the usual failover activity. Node B will continue processing and bring the protected resources in service.

**With resources being in-service on Node B, Node A is powered on and establishes communications with the other nodes and is able to access the QWK shared storage**

In this case, Node A will process a LCM\_AVAIL event. Node A will determine that it has quorum and not bring resources in service because they are currently in service on Node B. Next, a COMM\_UP event will be processed between Node A and Node B.

Each node will determine that it has quorum during the COMM\_UP events and Node A will not bring resources in service because they are currently in service on Node B.

**With resources being in-service on Node B, Node A is powered on and cannot establish communications to the other nodes but is able to access the QWK shared storage**

In this case, Node A will process a LCM\_AVAIL event. Node A will determine that it has quorum since it can access the shared storage for the QWK objects. It will then perform witness checks to determine the status for Node B since the communication to Node B is down. Since Node B is running and has been

updating its QWK object, Node A detects this and does not bring resources in service. Node B will do nothing since it can't communicate with Node A and already has the resource in-service.

### Scenario 3

**A failure occurs with the network for Node A (Node A is running without communication paths to the other nodes and does not have access to the QWK objects on shared storage)**

In this case, Node A will do the following:

1. Begin processing a COMM\_DOWN event from Node B.
2. Determine that it cannot access the shared storage and thus does not have quorum.
3. **Immediately** force-quit ("*fastkill*", default behavior of QUORUM\_LOSS\_ACTION).

Also, in this case, Node B will do the following:

1. Begin processing a COMM\_DOWN event from Node A.
2. Determine that it can still access the shared storage and thus has quorum.
3. Verify that the updating for the QWK objects for Node A has stopped (witness checking).
4. Verify via witness checking that Node A really appears to be lost and, begin the usual failover activity. Node B will now have the protected resources in service.

**With resources being in-service on Node B, Node A is powered on and establishes communications with the other nodes and is able to access the QWK shared storage**

Same as scenario 2.

**With the resources being in-service on Node B, Node A powered-on but is not able to access the QWK shared storage**

In this case, Node A will process an LCM\_AVAIL event. Node A will determine that it does not have quorum and will not bring resources in service.

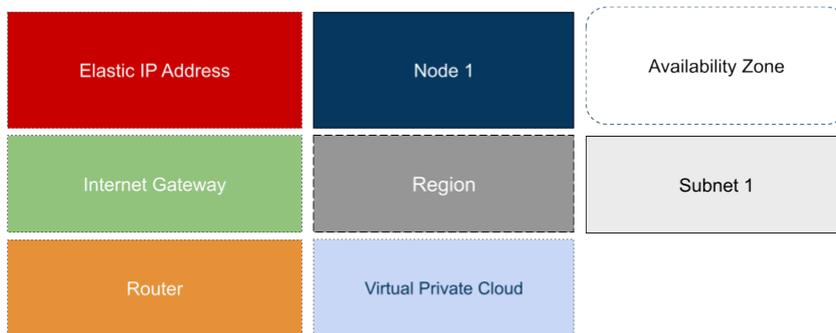
If the communication paths to Node B are available, then a COMM\_UP event will be processed. However, because Node A does not have quorum, it will not bring resources in service.

## 5.4.2.10.3.4. Quorum/Witness Cluster Recommendations in AWS

---

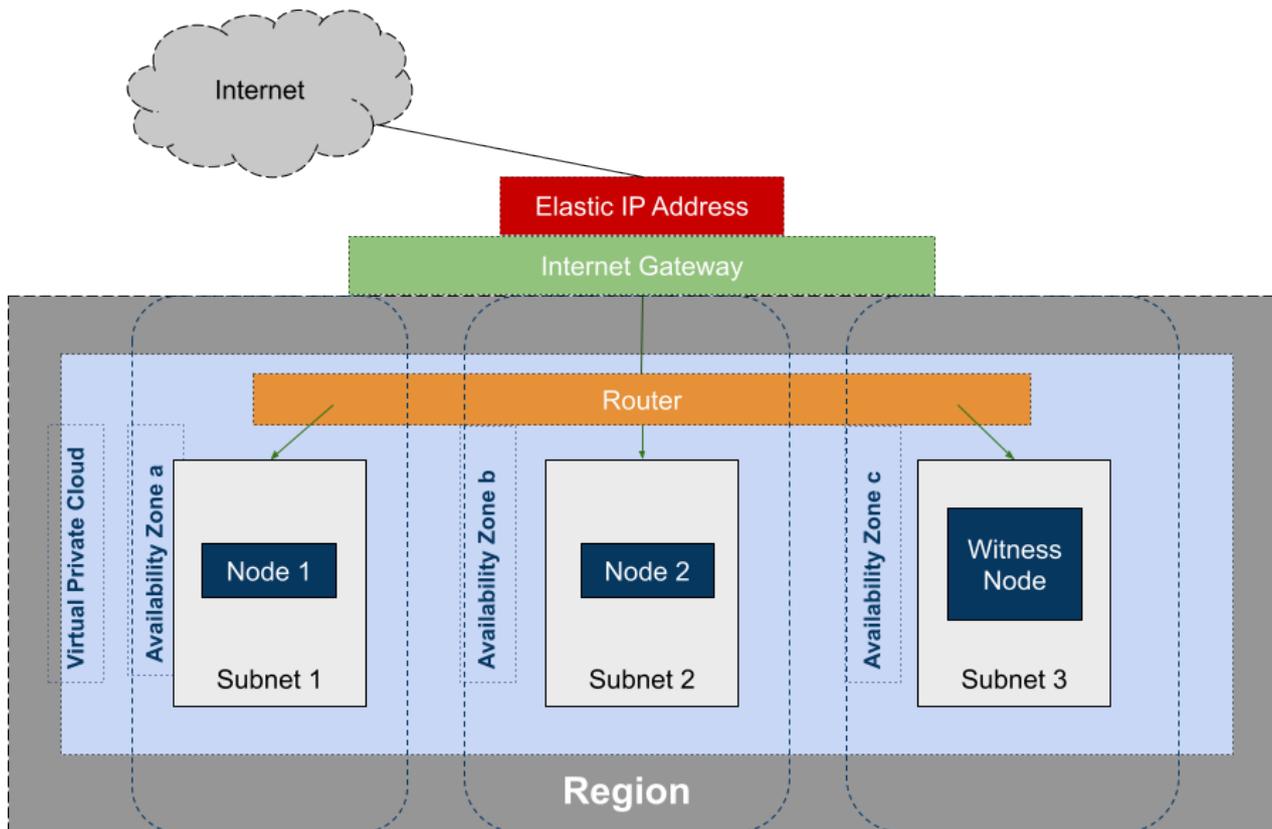
On this page you will find cluster configurations that provide ways you can help setup your **quorum/witness** alongside existing nodes within an AWS cloud environment.

### Key for diagrams below:



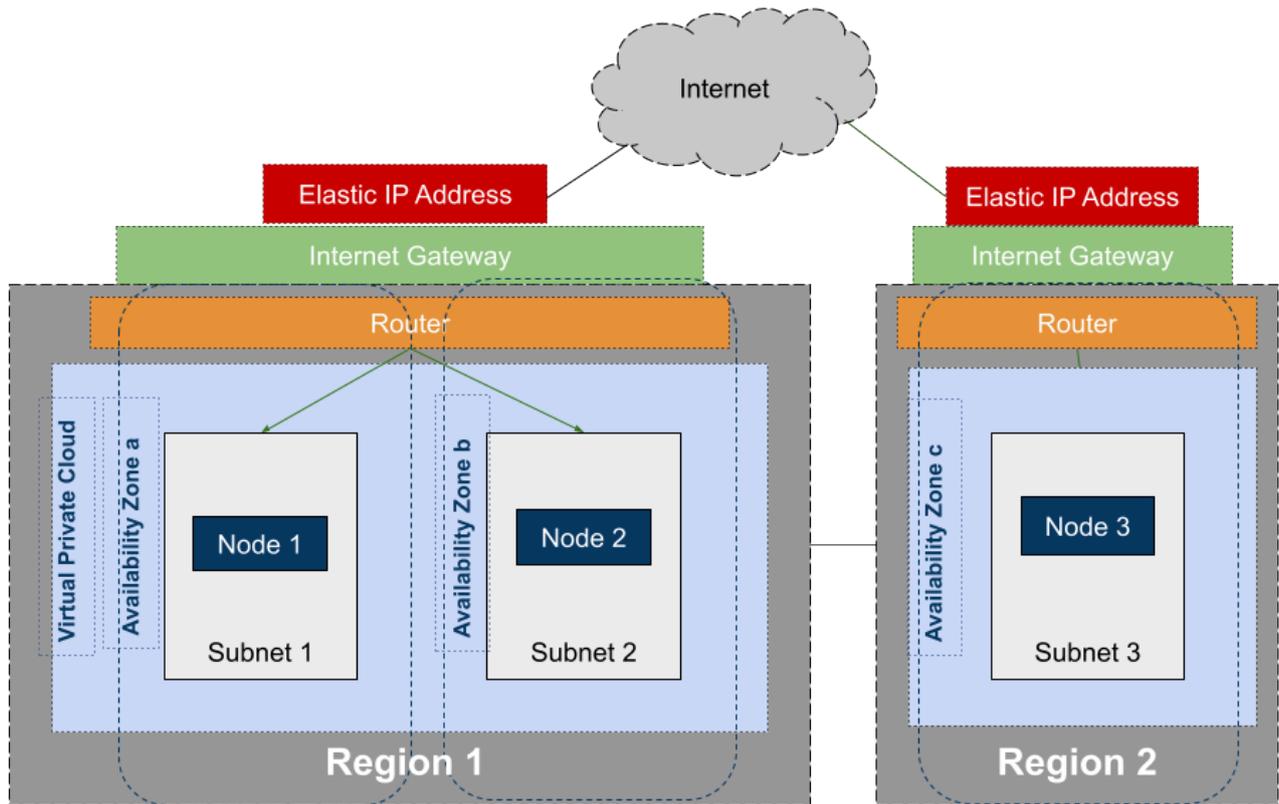
### 2 Node, Single Region Deployment

If both nodes all reside in the same region, then the witness node should reside in same region. **If any availability zone in the region fails, you still have one node and a witness node.** If the entire region fails you have no nodes anyway, so failover becomes irrelevant.



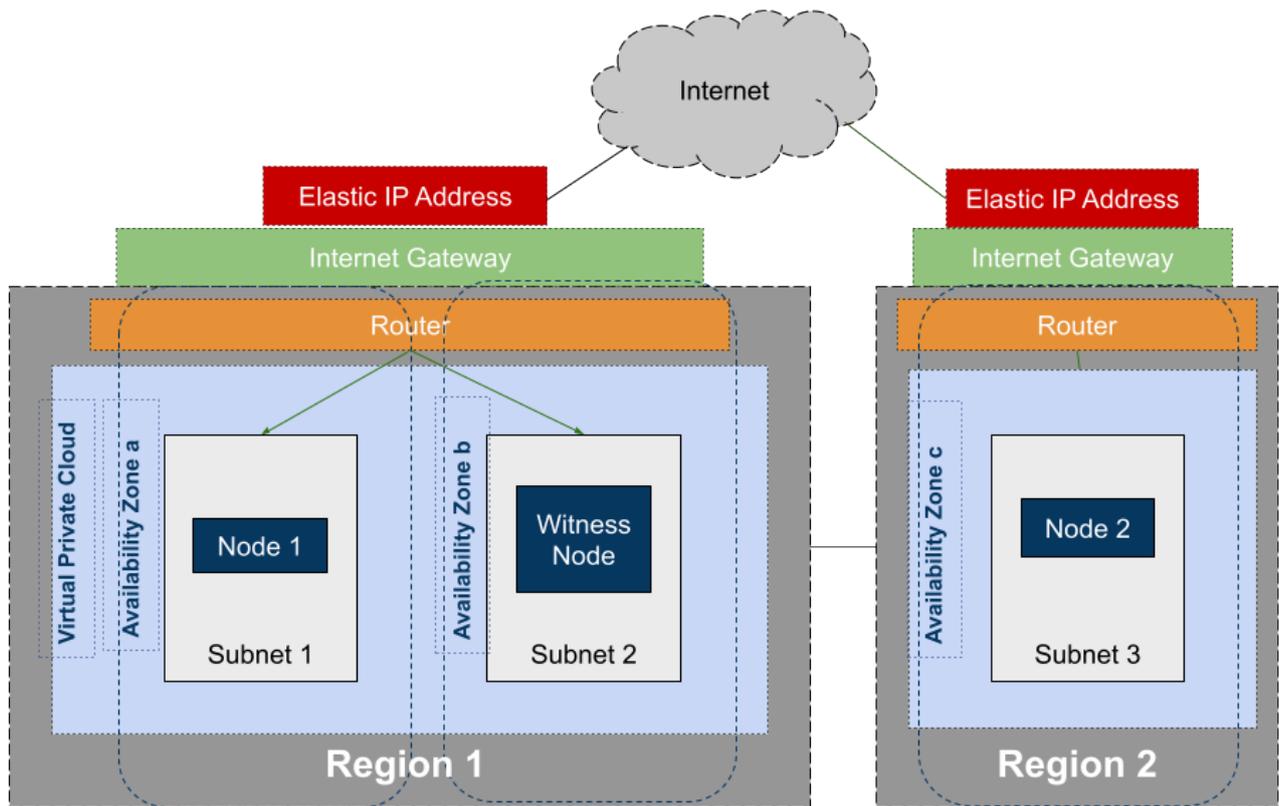
### 3 Node, 2 Region Deployment

In a 3-node cluster where two nodes live in one region and the 3rd lives in a different region then a witness node **does not** need to be added to the cluster. This is because we already have an **odd number** of votes. Failover to the DR region will always be a manual process in the event of the failover of the entirety of region 1.

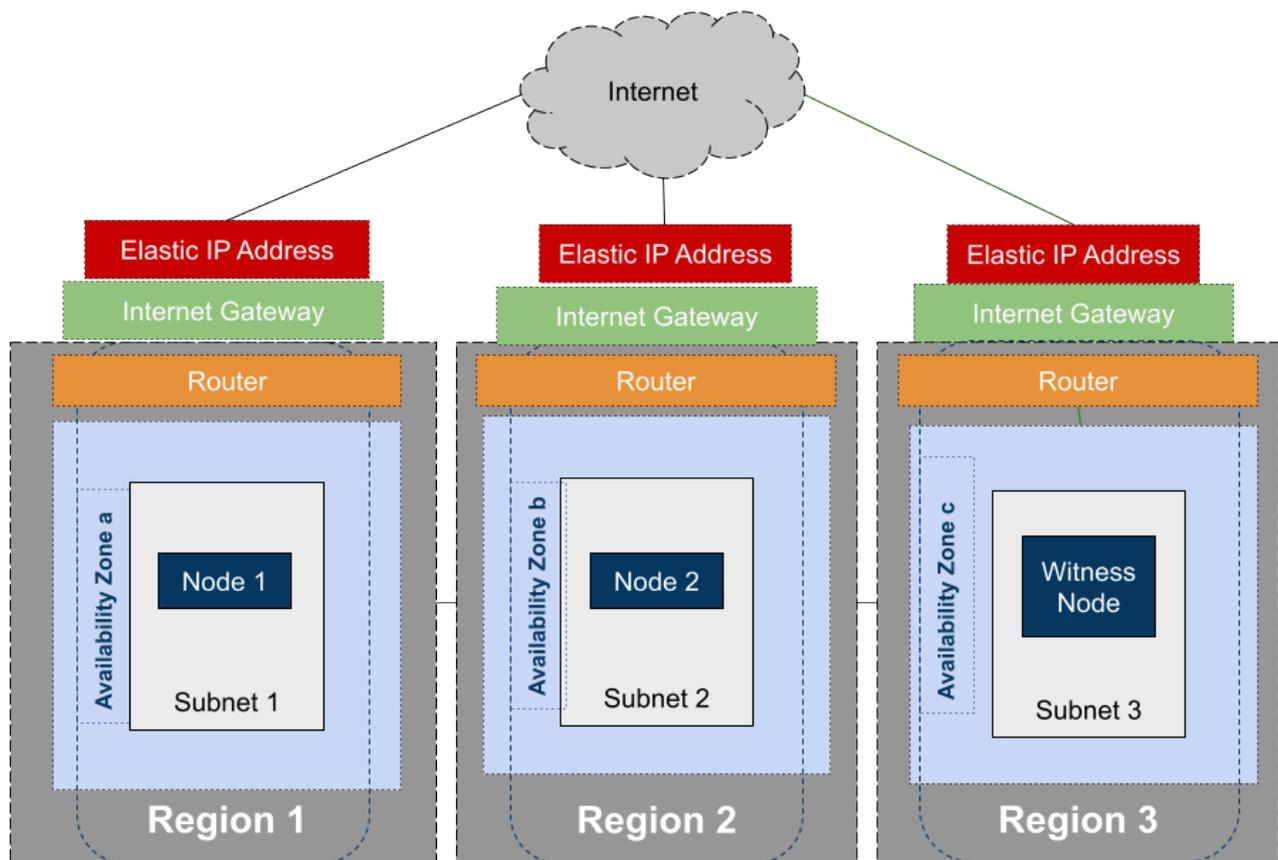


## 2 Node w/Witness, Multi Region Deployment

In a 2-node cluster where one node lives in region 1 and node 2 lives in region 2, the witness should reside in region 3 if you have it, but region 1 if you don't have it. The reason to not put it in region 2 is that a failure of the network between region 1 and region 2 will cause a failover to region 2 unnecessarily. Having a 3rd region eliminates that possibility. If the 3rd region is not practical, then putting it in the primary region is preferable to eliminate false failover. However, in the event of a complete failure of region 1, you will have to bring the secondary server online manually. **When this event occurs, please be sure to log the time/date of the failure and specifically what was done to bring the server online.**

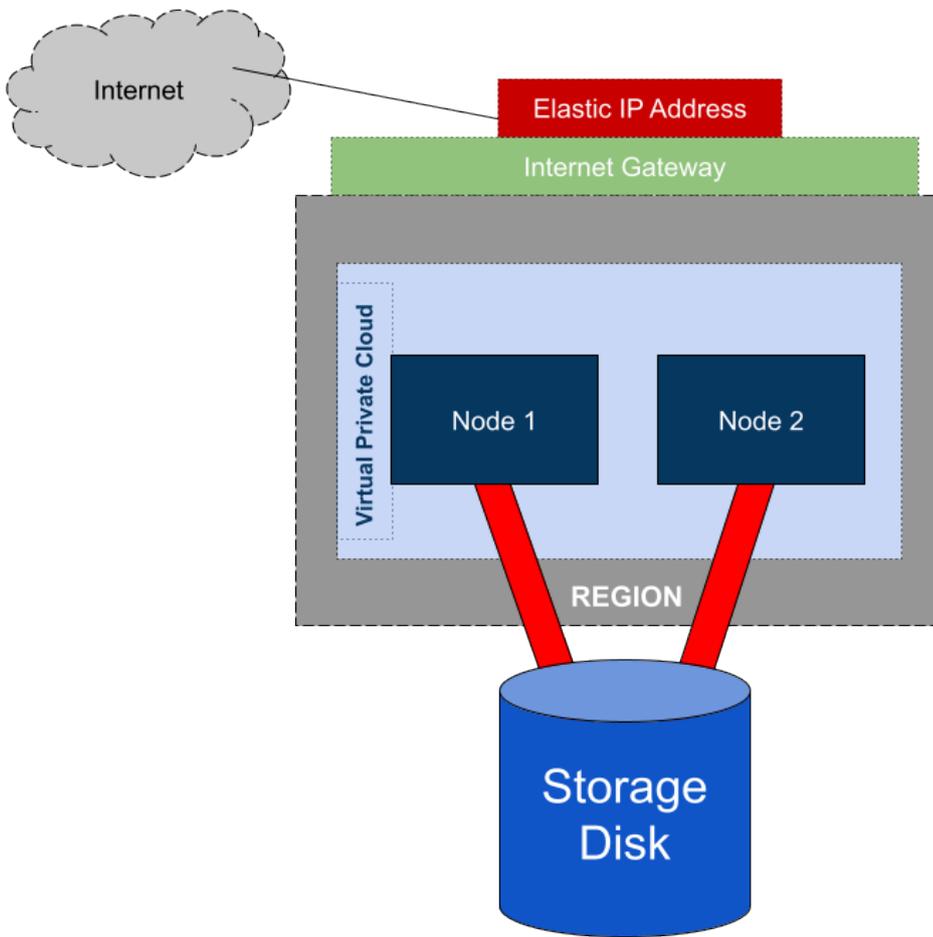


## Alternative Method



## 2 nodes and storage disk

Amazon S3 buckets are accessible globally. **NFS or EFS** could be set up in the same region.



\* **Note:** In addition to block storage we can use **NFS, EFS and S3 buckets,**

## 5.4.2.10.4. STONITH

**STONITH** (Shoot The Other Node in the Head) is a fencing technique for remotely powering down a node in a cluster. LifeKeeper can provide STONITH capabilities by using external power switch controls, IPMI-enabled motherboard controls, hypervisor-provided power capabilities and cloud vendor tools to power off the other nodes in a cluster. Each STONITH method allows the cluster software to *power off* a cluster node that appears to have died thus ensuring that the unhealthy node cannot access or corrupt any shared data.

\* **Note:** Power off is recommended over reboot to avoid fence loops (i.e. two machines have lost communication but can still STONITH each other, taking turns powering each other off and rebooting).

### STONITH using IPMI

IPMI (Intelligent Platform Management Interface) defines a set of common interfaces to a computer system which can be used to monitor system health and manage the system. Used with STONITH, it allows the cluster software to instruct the switch via a serial or network connection to power off a cluster node that appears to have died thus ensuring that the unhealthy node cannot access or corrupt any shared data.

#### Package Requirements

- IPMI tools package on **each server** in the cluster(e.g. ipmitool-1.8.11-6.el6.x86\_64.rpm)

#### Configure Baseboard Management Controller (BMC)

Using BIOS or the ipmitool command on **each server** in the cluster (example using ipmitool):

- **Use Static IP:** ipmitool lan set 1 ipsrc static
  - **Add IP address:** ipmitool lan set 1 ipaddr 192.168.0.1
  - **Set Sub netmask:** ipmitool lan set 1 netmask 255.0.0.0
  - **Set User name:** ipmitool user set name 1 root
  - **Set Password:** ipmitool user set password 1 secret
  - **Add administrator privilege level to the user:** ipmitool user priv 1 4
  - **Enable network access to the user:** ipmitool user enable 1
- (For detailed information, see the ipmitool man page.)

#### STONITH Installation

Install the LifeKeeper STONITH script on each server where LifeKeeper is installed and communication paths are configured to all servers by running the following command:

```
# /opt/LifeKeeper/samples/STONITH/stonith-install
```

## Update /opt/LifeKeeper/config/stonith.conf

Add entries to the stonith.conf file for each server in the cluster.

```
# LifeKeeper STONITH configuration
#
# Example: <host> ipmitool -I <interface> -H <ip> -U root -P secret power off
minute-maid ipmitool -I lanplus -H 192.168.0.1 -U root -P secret power off
kool-aid ipmitool -I lanplus -H 192.168.0.2 -U root -P secret power off
```

## STONITH in VMware vSphere Environments

vCLI (vSphere Command-Line Interface) is a command-line interface supported by VMware for managing your virtual infrastructure including the ESXi hosts and virtual machines. You can choose the vCLI command best suited for your needs and apply it for your LifeKeeper STONITH usage between VMware virtual machines.

### Package Requirements

#### STONITH Server

VMware vSphere SDK Package or VMware vSphere CLI (vSphere CLI is included in the same installation package as the vSphere SDK).

#### Each Monitored Virtual Machine

VMware Tools

## Configuration

vSphere CLI commands run on top of vSphere SDK for Perl.

- vCLI-esxcli
  - esxcli —server=10.0.0.1 —username=root —password=secret vms vm kill —type='hard' —world-id=1234567
- vCLI-vmware\_cmd
  - vmware-cmd -H 10.0.0.1 -U root -P secret <vm\_id> stop hard

### Determining <vm\_id>

vSphere CLI commands run on top of vSphere SDK for Perl. <vm\_id> is used as an identifier of the VM. This variable should point to the VM's configuration file for the VM being configured.

1. Get the list of available hosts:

```
vmware-cmd -H <vmware host> -l
```

## 2. Sample output:

```

/vmfs/volumes/4e08c1b9-d741c09c-1d3e-0019b9cb28be/lampserver/
lampserver.vmx
/vmfs/volumes/4e1e1386-0b862fae-a859-0019b9cb28bc/oracle10/oracle.vmx
/vmfs/volumes/4e08c1b9-d741c09c-1d3e-0019b9cb28be/lampserver02/
lampserver02.vmx

```

## 3. The command referencing the first VM in the output that is in bold:

```

vmware-cmd -H 10.0.0.1 -U root -P secret
/vmfs/volumes/4e08c1b9-d741c09c-1d3e-0019b9cb28be/l
ampserver/lampserver.vmx stop hard

```

## STONITH Installation

Install the LifeKeeper STONITH script on **each server** where LifeKeeper is installed and communication paths are configured to all servers by running the following command:

```
# /opt/LifeKeeper/samples/STONITH/stonith-install
```

## Update /opt/LifeKeeper/config/stonith.conf

The entries for the 3 hosts listed in the output above for the stonith.conf file (all other entries should be commented out or deleted):

```

# LifeKeeper STONITH configuration
#
# Example: vmware-cmd -H 10.0.0.1 -U root -P secret stop hard
lampserver vmware-cmd -H 10.0.0.1 -U root -P secret
/vmfs/volumes/4e08c1b9-d741c09c-1d3e-0019b9cb28be/lampserver/lampserver.vmx
stop hard
oracle vmware-cmd -H 10.0.0.1 -U root -P secret
/vmfs/volumes/4e1e1386-0b862fae-a859-0019b9cb28bc/oracle10/oracle.vmx stop hard
lampserver02 vmware-cmd -H 10.0.0.1 -U root -P secret /
vmfs/volumes/4e08c1b9-d741c09c-1d3e-0019b9cb28be/lampserver02/lampserver02.vmx
stop hard

```

## STONITH in Microsoft Azure Environments

The Azure CLI is a command line interface supported by Microsoft to manage Azure resources such as virtual machines. Used with STONITH, it allows the cluster software to power off a cluster node that appears to have died thus ensuring that the unhealthy node cannot access or corrupt any shared data.

## Requirements

### Package

Azure CLI – [install Linux Azure CLI](#) on **each server** in the cluster.

### Custom Role

The virtual machines on Microsoft Azure used for Azure Fencing and the custom roles assigned to users must have at least powerOff permissions on the virtual machines.

```
Microsoft.Compute/*/read
Microsoft.Compute/virtualMachines/powerOff/action
Microsoft.Compute/virtualMachines/start/action
```

## Pre-checking

On **each server** run the following command to verify that authentication on Microsoft Azure is working.

```
# az vm show --resource-group <group name> --name <vm name>
```

## STONITH Installation

Install the LifeKeeper for Microsoft Azure STONITH script on **each server** where LifeKeeper is installed and communication paths are configured to all servers by running the following command:

```
# /opt/LifeKeeper/samples/STONITH/azure-stonith-install
```

Since the above command works interactively, enter the group name and check the virtual machine name on Microsoft Azure for the cluster node displayed.

Example output from the command

```
STONITH script install...
Please enter the Resource Group name in Azure: rg-Group1

Please enter the System name in Azure[vm-HostA]:
Enable Stonith on node vm-HostA [Yes]:
s
Please enter the System name in Azure[vm-HostB]:
Enable Stonith on node vm-HostB [Yes]:
Configuration file /opt/LifeKeeper/config/stonith.conf was saved.
```

## Update /opt/LifeKeeper/config/stonith.conf

After the installation is completed, the settings to power off the virtual machine will be added to the following file.

/opt/LifeKeeper/config/stonith.conf

```
# LifeKeeper STONITH configuration
#
# Example: <host> az vm restart -g <resource group> -n <node name>
vm-HostA az vm stop -g rg-Group1 -n vm-HostA --skip-shutdown
vm-HostB az vm stop -g rg-Group1 -n vm-HostB --skip-shutdown
```

 **Note:** 'stop' is recommended over restart to avoid fence loops (i.e. two machines have lost communication but can still STONITH each other, taking turns powering each other off and rebooting). 'skip-shutdown' is recommended to quickly power off the node.

## Expected Behaviors

When LifeKeeper detects a communication failure with a node, that node will be powered off and a failover will occur. Once the issue is repaired, the node will have to be manually powered on.

## 5.4.2.10.5. Watchdog

Watchdog is a method of monitoring a server to ensure that if the server is not working properly, corrective action (reboot) will be taken so that it does not cause problems. Watchdog can be implemented using special watchdog hardware or using a software-only option.

\* **Note:** This configuration has only been tested with Red Hat Enterprise Linux Version 7. No other operating systems have been tested; therefore, no others are supported at this time.

### Components

- Watchdog timer – software driver or an external hardware component
- Watchdog daemon – rpm available through the Linux distribution
- LifeKeeper core daemon – installed with the LifeKeeper installation
- Health check script – Script to check the status of LifeKeeper core



LifeKeeper Interoperability with Watchdog

Read the next section carefully. The daemon is designed to recover from errors and will reset the system if not configured carefully. Planning and care should be given to how this is installed and configured. This section is not intended to explain and configure watchdog, but only to explain and configure how LifeKeeper interoperates in such a configuration.

### Configuration

The following steps should be carried out by an administrator with root user privileges. The administrator should already be familiar with some of the risks and issues with watchdog.

The health check script (LifeKeeper monitoring script) is the component that ties the LifeKeeper

configuration with the watchdog configuration (`/opt/LifeKeeper/samples/watchdog/LifeKeeper-watchdog`). This script can monitor the basic parts of LifeKeeper core components.

1. If watchdog has been previously configured, enter the following command to stop it. If not, go to Step 2.

```
systemctl stop watchdog
```

2. Edit the watchdog configuration file (`/etc/watchdog.conf`) supplied during the installation of watchdog software.

- Modify test-binary:

```
test-binary = /opt/LifeKeeper/samples/watchdog/LifeKeeper-watchdog
```

- Modify test-timeout:

```
test-timeout = 5
```

- Modify interval:

```
interval = 7
```

The interval value should be less than LifeKeeper communication path timeout (15 seconds), so a good number for the interval is generally half of this value.

3. Make sure LifeKeeper has been started. If not, please refer to the [Starting LifeKeeper](#) topic.
4. Start watchdog by entering the following command:

```
systemctl start watchdog
```

5. To start watchdog automatically on future restarts, enter the following command:

```
systemctl enable watchdog
```

 **Note:** Configuring watchdog may cause some unexpected reboots from time to time. This is the general nature of how watchdog works. If processes are not responding correctly, the watchdog feature will assume that LifeKeeper (or the operating system) is hung, and it will reboot the system (without warning).

## Uninstall

Care should be taken when uninstalling LifeKeeper. The above steps should be done in reverse order as listed below.

**! WARNING:** IF UNINSTALLING LIFEKEEPER BY REMOVING THE RPM PACKAGES THAT MAKE UP LIFEKEEPER, **TURN OFF WATCHDOG FIRST!** In Step 2 above, the watchdog config file was modified to call on the LifeKeeper-watchdog script; therefore, if watchdog is not turned off first, it will call on that script that is no longer there. An error will occur when this script is not found which will trigger a reboot. This will continue until watchdog is turned off.

1. Stop watchdog by entering the following command:

```
systemctl stop watchdog
```

2. Edit the watchdog configuration file (`/etc/watchdog.conf`) supplied during the installation of watchdog software.

- Modify test-binary and interval by commenting out those entries (add # at the beginning of each line):

```
#test-binary =
```

```
#interval =
```

**Note:** If interval was used previously for other functions, it can be left as-is

3. Uninstall LifeKeeper. See the [Removing LifeKeeper](#) topic.
4. Watchdog can now be started again. If only used by LifeKeeper, watchdog can be permanently disabled by entering the following command:

```
systemctl disable watchdog
```

## 5.4.2.10.6. I/O Fencing Mechanisms

---

LifeKeeper for Linux provides various fencing mechanisms. Depending on the server and storage configuration, available fencing mechanisms and allowed combinations of these may differ.

Refer to the information linked below for fencing mechanisms available in each server configuration. Details of storage configuration are described on the server configuration pages.

- For physical servers see [Available I/O Fencing Mechanisms \(Physical Servers\)](#)
- For virtual machines in VMware see [Available I/O Fencing Mechanisms \(Virtual Machines in VMware\)](#)

### I/O Fencing Mechanism Summary

- SCSI Fencing with SCSI-2 Reservations – By issuing a SCSI-2 reservation to the storage from the active node, the protected logical unit (LU) is locked, preventing simultaneous access from other nodes. When a communication failure is detected, a standby node will disable LU locks from other nodes, and then lock the protected LU, preventing simultaneous access from other nodes.
- SCSI Fencing with SCSI-3 Reservations – By issuing a SCSI-3 reservation to the storage from the active node, the protected logical unit (LU) is locked, preventing simultaneous access from other nodes. When a communication failure is detected, a standby node will disable LU locks from other nodes, and then lock the protected LU, preventing simultaneous access from other nodes.
- IPMI STONITH – When a communication failure is detected, the nodes will issue IPMI “power off” commands to the other nodes, preventing the protected service from starting on multiple nodes. This mechanism is available only for physical servers with IPMI interfaces.
- VMware STONITH – When a communication failure is detected, the nodes will issue “power off” commands to the other nodes via the VMware host or vCenter APIs, preventing the protected service from starting on multiple nodes. This mechanism is available only for virtual machines running in VMware environments.
- Quorum/Witness – When a communication failure is detected, an arbitrator will determine whether failover to a given node would effectively achieve service continuation. LifeKeeper normally prevents simultaneous access to LUs from multiple nodes with the reservation commands described above, but Quorum/Witness is mandatory in environments where reservations are not used. There are three modes of arbitration: dedicated Witness node, independent host, or shared storage (see [Quorum/Witness](#) for details).

## 5.4.2.10.6.1. Available I/O Fencing Mechanisms (Physical Servers)

This page describes the combinations of fencing mechanisms that can be used with various storage configurations in a physical server environment.

### Shared Disk Configuration (Single or Multipath Configuration Using SCSI Reservations)

This configuration corresponds to the case where exclusive control is enforced by SCSI reservations in a certified shared storage system, as listed in the [Supported Storage List](#). When using storage that requires a multipath driver for SCSI-3 reservations and multipath control, a multipath kit that supports the multipath driver is required.

#### Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used.

##### Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations					
SCSI-2 Reservations	●	●	●	? <sup>1</sup>	? <sup>1</sup>	? <sup>1</sup>
SCSI-3 Reservations	? <sup>1</sup>	? <sup>1</sup>	? <sup>1</sup>	●	●	●
IPMI STONITH	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>
Quorum/Witness (tcp_remote)	○	? <sup>3</sup>	? <sup>3</sup>	○	? <sup>3</sup>	? <sup>3</sup>
Quorum/Witness (majority)	? <sup>3</sup>	○	? <sup>3</sup>	? <sup>3</sup>	○	? <sup>3</sup>
Quorum/Witness (storage)	? <sup>3</sup>	? <sup>3</sup>	○	? <sup>3</sup>	? <sup>3</sup>	○

<sup>1</sup>SCSI-2 and SCSI-3 reservations cannot coexist on a single shared disk.

<sup>2</sup> When SCSI Reservations and STONITH are used together, the functions of each may conflict and an unexpected system outage may occur.

<sup>3</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

### Mechanisms Not Available with this Configuration

- Fencing mechanisms for virtual and cloud environments cannot be used with this configuration.

## Shared Disk Configuration (Single or Multipath Configuration Not Using SCSI Reservations)

This configuration is applicable when SCSI reservations cannot be used with a shared storage system connected with a method including SCSI/FC/iSCSI/SAS (but excluding NAS). When using storage that requires a multipath driver to control multiple paths, a multipath kit that supports the multipath driver is required.

### Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used.

#### Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations		
<b>IPMI STONITH</b>	○	○	○
<b>Quorum/Witness (tcp_remote)</b>	●	? <sup>1</sup>	? <sup>1</sup>
<b>Quorum/Witness (majority)</b>	? <sup>1</sup>	●	? <sup>1</sup>
<b>Quorum/Witness (storage)</b>	? <sup>1</sup>	? <sup>1</sup>	●

<sup>1</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

### Mechanisms Not Available with this Configuration

- SCSI Fencing cannot be used with storage systems that are not certified to properly support SCSI reservations.
- Fencing mechanisms for virtual and cloud environments cannot be used with this configuration.

## Data Replication Configuration

This configuration is applicable when local storage connected to each node is replicated between nodes using DataKeeper.

### Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used. Fencing is not mandatory in this configuration; however, it is recommended to use fencing to enhance data protection.

#### Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations		
<b>IPMI STONITH</b>	○	○	○
<b>Quorum/Witness (tcp_remote)</b>	○	? <sup>1</sup>	? <sup>1</sup>
<b>Quorum/Witness (majority)</b>	? <sup>1</sup>	○	? <sup>1</sup>
<b>Quorum/Witness (storage)</b>	? <sup>1</sup>	? <sup>1</sup>	○

<sup>1</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

### Mechanisms Not Available with this Configuration

- SCSI Fencing cannot be used with this configuration.
- Fencing mechanisms for virtual and cloud environments cannot be used with this configuration.

## Network Attached Storage (NAS) Configuration

This configuration is applicable when using NAS storage connected using Network File System (NFS).

### Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used. Fencing is not mandatory in this configuration; however, it is recommended to use fencing to enhance data protection.

## Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations		
IPMI STONITH	○	○	○
Quorum/Witness (tcp_remote)	○	? <sup>1</sup>	? <sup>1</sup>
Quorum/Witness (majority)	? <sup>1</sup>	○	? <sup>1</sup>
Quorum/Witness (storage)	? <sup>1</sup>	? <sup>1</sup>	○

<sup>1</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

## Mechanisms Not Available with this Configuration

- SCSI Fencing cannot be used with this configuration.
- Fencing mechanisms for virtual and cloud environments cannot be used with this configuration.

## 5.4.2.10.6.2. Available I/O Fencing Mechanisms (Virtual Machines in VMware)

This page describes the combinations of fencing mechanisms that can be used with various storage configurations in a VMware virtual server environment.

### Shared Disk Configuration (Single or Multipath Configuration Using SCSI Reservations)

This configuration corresponds to the case where exclusive control is enforced by SCSI reservations in a certified shared storage system, as listed in the [Supported Storage List](#). When using storage that requires a multipath driver for SCSI-3 reservations and multipath control, a multipath kit that supports the multipath driver is required.

#### Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used.

##### Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations					
SCSI-2 Reservations	●	●	●	? <sup>1</sup>	? <sup>1</sup>	? <sup>1</sup>
SCSI-3 Reservations	? <sup>1</sup>	? <sup>1</sup>	? <sup>1</sup>	●	●	●
VMware STONITH	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>	? <sup>2</sup>
Quorum/Witness (tcp_remote)	○	? <sup>3</sup>	? <sup>3</sup>	○	? <sup>3</sup>	? <sup>3</sup>
Quorum/Witness (majority)	? <sup>3</sup>	○	? <sup>3</sup>	? <sup>3</sup>	○	? <sup>3</sup>
Quorum/Witness (storage)	? <sup>3</sup>	? <sup>3</sup>	○	? <sup>3</sup>	? <sup>3</sup>	○

<sup>1</sup>SCSI-2 and SCSI-3 reservations cannot coexist on a single shared disk.

<sup>2</sup> When SCSI Reservations and STONITH are used together, the functions of each may conflict and an unexpected system outage may occur.

<sup>3</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

### Mechanisms Not Available with this Configuration

- Fencing mechanisms for virtual environments other than VMware cannot be used with this configuration.
- Fencing mechanisms for cloud and physical environments cannot be used with this configuration.

## Shared Disk Configuration (Single or Multipath Configuration Not Using SCSI Reservations)

This configuration is applicable when SCSI reservations cannot be used with a shared storage system connected with a method including SCSI/FC/iSCSI/SAS (but excluding NAS). When using storage that requires a multipath driver to control multiple paths, a multipath kit that supports the multipath driver is required.

### Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used.

#### Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations		
<b>VMware STONITH</b>	○	○	○
<b>Quorum/Witness (tcp_remote)</b>	●	? <sup>1</sup>	? <sup>1</sup>
<b>Quorum/Witness (majority)</b>	? <sup>1</sup>	●	? <sup>1</sup>
<b>Quorum/Witness (storage)</b>	? <sup>1</sup>	? <sup>1</sup>	●

<sup>1</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

### Mechanisms Not Available with this Configuration

- SCSI Fencing cannot be used with storage systems that are not certified to properly support SCSI reservations.
- Fencing mechanisms for virtual environments other than VMware cannot be used with this

configuration.

- Fencing mechanisms for cloud and physical environments cannot be used with this configuration.

## Data Replication Configuration

This configuration is applicable when local storage connected to each node is replicated between nodes using DataKeeper.

### Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used. Fencing is not mandatory in this configuration; however, it is recommended to use fencing to enhance data protection.

#### Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations		
VMware STONITH	○	○	○
Quorum/Witness (tcp_remote)	○	? <sup>1</sup>	? <sup>1</sup>
Quorum/Witness (majority)	? <sup>1</sup>	○	? <sup>1</sup>
Quorum/Witness (storage)	? <sup>1</sup>	? <sup>1</sup>	○

<sup>1</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

### Mechanisms Not Available with this Configuration

- SCSI Fencing cannot be used with this configuration.
- Fencing mechanisms for virtual environments other than VMware cannot be used with this configuration.
- Fencing mechanisms for cloud and physical environments cannot be used with this configuration.

## Network Attached Storage (NAS) Configuration

This configuration is applicable when using NAS storage connected using Network File System (NFS).

## Available Fencing Mechanisms

The following table shows which fencing mechanisms are available and the allowed combinations, if multiple fencing mechanisms are used. Fencing is not mandatory in this configuration; however, it is recommended to use fencing to enhance data protection.

### Symbol Definitions

- – Required
- – Available as an option
- ? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations		
VMware STONITH	○	○	○
Quorum/Witness (tcp_remote)	○	? <sup>1</sup>	? <sup>1</sup>
Quorum/Witness (majority)	? <sup>1</sup>	○	? <sup>1</sup>
Quorum/Witness (storage)	? <sup>1</sup>	? <sup>1</sup>	○

<sup>1</sup> Multiple Quorum/Witness modes cannot be used together in a single cluster.

## Mechanisms Not Available with this Configuration

- SCSI Fencing cannot be used with this configuration.
- Fencing mechanisms for virtual environments other than VMware cannot be used with this configuration.
- Fencing mechanisms for cloud and physical environments cannot be used with this configuration.

## VMDK as Shared Storage Configuration

This configuration is applicable when using a VMware virtual hard disk configured with the VMDK as the Shared Storage method.

### Available Functions and Appropriate Combinations

The table below shows the fencing functions available with this configuration and the appropriate combination patterns.

### Symbol Definitions

- – Required

○ – Available as an option

? – Not available

For the functions listed as “Available as an option” in the table, use of the functions in the combination is not mandatory. Please refer to the note below the table for the functions that are “Not available” in the table.

	Allowed Combinations		
<b>VMware STONITH</b>	? <sup>1</sup>	? <sup>1</sup>	? <sup>1</sup>
<b>Quorum/Witness (tcp_remote)</b>	○	? <sup>2</sup>	? <sup>2</sup>
<b>Quorum/Witness (majority)</b>	? <sup>2</sup>	○	? <sup>2</sup>
<b>Quorum/Witness (storage)</b>	? <sup>2</sup>	? <sup>2</sup>	○

<sup>1</sup>VMDK as Shared Storage and VMware STONITH cannot coexist because service may stop due to a conflict.

<sup>2</sup> Due to the LifeKeeper mechanism, Quorum Witness modes cannot coexist for a single cluster.

### Unavailable Function with this Configuration

- The fencing function that requires the use of a shared disk cannot be used with this configuration.

## 5.4.2.11. Resource Policy Management

---

Resource Policy Management in LifeKeeper for Linux provides behavior management of resource local recovery and failover. Resource policies are managed with the **lkpolicy** command line tool (CLI).

### LifeKeeper

LifeKeeper is designed to monitor individual applications and groups of related applications, periodically performing local recoveries or notifications when protected applications fail. Related applications, by example, are hierarchies where the primary application depends on lower-level storage or network resources. When an application or resource failure occurs, the default behavior is:

1. **Local Recovery:** First, attempt **local** recovery of the resource or application. An attempt will be made to restore the resource or application on the local server without external intervention. If local recovery is successful, then LifeKeeper will not perform any additional action.
2. **Failover:** Second, if a local recovery attempt *fails* to restore the resource or application (or the *recovery kit* monitoring the resource has no support for local recovery), then a **failover** will be initiated. The failover action attempts to bring the application (and all dependent resources) into service on another server within the cluster.

Please see [LifeKeeper Fault Detection and Recovery Scenarios](#) for more detailed information about our recovery behavior.

### Custom and Maintenance-Mode Behavior via Policies

LifeKeeper Version 7.5 and later supports the ability to set additional policies that modify the default recovery behavior. There are four policies that can be set for individual resources (see the section below about precautions regarding individual resource policies) or for an entire server. ***The recommended approach is to alter policies at the server level.***

The available policies are:

#### Standard Policies

- **Failover** This policy setting can be used to turn on/off resource failover. (**Note:** In order for reservations to be handled correctly, **Failover** cannot be turned off for individual scsi resources.)
- **LocalRecovery** – LifeKeeper, by default, will attempt to recover protected resources by restarting the individual resource or the entire protected application prior to performing a failover. This policy setting can be used to turn on/off local recovery.
- **TemporalRecovery** – Normally, LifeKeeper will perform local recovery of a failed resource. If local recovery fails, LifeKeeper will perform a resource hierarchy failover to another node. If the local recovery succeeds, failover will not be performed.

There may be cases where the local recovery succeeds, but due to some irregularity in the server, the local recovery is re-attempted within a short time; resulting in multiple, consecutive local recovery attempts. This may degrade availability for the affected application.

To prevent this repetitive local recovery/failure cycle, you may set a temporal recovery policy. The temporal recovery policy allows an administrator to limit the number of local recovery attempts (successful or not) within a defined time period.

*Example:* If a user sets the policy definition to limit the resource to three local recovery attempts in a 30-minute time period, LifeKeeper will fail over when a third local recovery attempt occurs within the 30-minute period.

Defined temporal recovery policies may be turned *on* or *off*. When a temporal recovery policy is *off*, temporal recovery processing will continue to be done and notifications will appear in the log when the policy *would* have fired; however, no actions will be taken.

 **Note:** It is possible to disable failover and/or local recovery with a temporal recovery policy also in place. This state is illogical as the temporal recovery policy will **never** be acted upon if failover or local recovery are disabled.

## Meta Policies

The “meta” policies are the ones that can affect more than one other policy at the same time. These policies are usually used as shortcuts for getting certain system behaviors that would otherwise require setting multiple standard policies.

- **NotificationOnly** – This mode allows administrators to put LifeKeeper in a “monitoring only” state. **Both** local recovery **and** failover **of a resource (or all resources in the case of a server-wide policy) are affected**. The user interface will indicate a **Failure** state if a failure is detected; *but no recovery or failover action will be taken*. **Note:** The administrator will need to correct the problem that caused the failure manually and then bring the affected resource(s) back in service to continue normal LifeKeeper operations.

## Important Considerations for Resource-Level Policies

Resource level policies are policies that apply to a specific resource only, as opposed to an entire resource hierarchy or server.

*Example:*

app

- IP

- file system

In the above resource hierarchy, app depends on both IP and file system. A policy can be set to *disable* local recovery or failover of a specific resource. This means that, for example, if the IP resource's local recovery fails and a policy was set to disable failover of the IP resource, then the IP resource will not fail over or cause a failover of the other resources. However, if the file system resource's local recovery fails and the file system resource policy does not have failover disabled, then the entire hierarchy will fail over.



**Note:** It is important to remember that resource level policies apply *only* to the specific resource for which they are set.

This is a simple example. Complex hierarchies can be configured, so care must be taken when setting resource-level policies.



**See known issue.** (“Resources removed in the wrong order during failover”)

## The Ikpolicy Tool

The `lkpolicy` tool is the command-line tool that allows management (querying, setting, removing) of policies on servers running LifeKeeper for Linux. `lkpolicy` supports setting/modifying policies, removing policies and viewing all available policies and their current settings. In addition, defined policies can be set on or off, preserving resource/server settings while affecting recovery behavior.

The general usage is :

```
lkpolicy [-list-policies | -get-policies | -set-policy | -remove-policy] <name
value pair data...>
```

The `<name value pair data...>` differ depending on the operation and the policy being manipulated, particularly when setting policies. *For example:* Most on/off type policies only require `-on` or `-off` switch, but the temporal policy requires additional values to describe the threshold values.

## Example Ikpolicy Usage

### Authenticating With Local and Remote Servers

The `lkpolicy` tool communicates with LifeKeeper servers via an API that the servers expose. This API requires authentication from clients like the `lkpolicy` tool. The first time the `lkpolicy` tool is asked to access a LifeKeeper server, if the credentials for that server are not known, it will ask the user for credentials for that server. These credentials are in the form of a username and password and:

1. Clients must have LifeKeeper admin rights. This means the username must be in the `lkadmin` group according to the operating system's authentication configuration (via `pam`). It is **not** necessary to run as **root**, but the root user can be used since it is in the appropriate group by

default.

- The credentials will be stored in the *credential store* so they do not have to be entered manually each time the tool is used to access this server.

See [Configuring Credentials for LifeKeeper](#) for more information on the credential store and its management with the credstore utility.

An example session with `lkpolicy` might look like this:

```
[root@thor49 ~]# lkpolicy -l -d v6test4
Please enter your credentials for the system 'v6test4'.
Username: root
Password:
Confirm password:
Failover
LocalRecovery
TemporalRecovery
NotificationOnly
[root@thor49 ~]# lkpolicy -l -d v6test4
Failover
LocalRecovery
TemporalRecovery
NotificationOnly
[root@thor49 ~]#
```

## Listing Policies

```
lkpolicy --list-policy-types
```

## Showing Current Policies

```
lkpolicy --get-policies
```

```
lkpolicy --get-policies tag=\*
```

```
lkpolicy --get-policies --verbose tag=mysql\* # all resources starting with mysql
```

```
lkpolicy --get-policies tag=mytagonly
```

## Setting Policies

```
lkpolicy --set-policy Failover --off
```

```
lkpolicy --set-policy Failover --on tag=myresource
```

```
lkpolicy --set-policy Failover --on tag=\*
```

```
lkpolicy —set-policy LocalRecovery —off tag=myresource
```

```
lkpolicy —set-policy NotificationOnly —on
```

```
lkpolicy —set-policy TemporalRecovery —on recoverylimit=5 period=15
```

```
lkpolicy —set-policy TemporalRecovery —on —force recoverylimit=5 period=10
```

## Removing Policies

```
lkpolicy —remove-policy Failover tag=steve
```

 **Note:** *NotificationOnly is a policy alias. Enabling NotificationOnly is the equivalent of disabling the corresponding LocalRecovery and Failover policies.*

## 5.4.2.12. Configuring Credentials

Credentials for communicating with other systems are managed via a *credential store*. This store can be managed, as needed, by the `/opt/LifeKeeper/bin/credstore` utility. This utility allows server access credentials to be set, changed and removed – on a per server basis.

### Adding or Changing Credentials

Adding and changing credentials are handled in the same way. A typical example of adding or changing credentials for a server, `server.mydomain.com`, would look like this:

```
/opt/LifeKeeper/bin/credstore -k server.mydomain.com myuser
```

In this case, *myuser* is the username used to access `server.mydomain.com` and the password will be asked for via a prompt with confirmation (like *passwd*).

 **Note:** The key name used to store LifeKeeper server credentials must match *exactly* the hostname used in commands such as `lkpolicy`. If the hostname used in the command is an FQDN, then the credential key must also be the FQDN. If the hostname is a short name, then the key must also be the short name.

You may wish to set up a **default** key in the credential store. The **default** credentials will be used for authentication when no specific server key exists. To add or change the **default** key, run:

```
/opt/LifeKeeper/bin/credstore -k default myuser
```

### Listing Stored Credentials

The currently stored credentials can be listed by the following command:

```
/opt/LifeKeeper/bin/credstore -l
```

This will list the *keys* stored in the credential store and, in this case, the key indicates the server for which the credentials are used. (This command will not actually list the credentials, only the key names, since the credentials themselves may be sensitive.)

### Removing Credentials for a Server

Credentials for a given server can be removed with the following command:

```
/opt/LifeKeeper/bin/credstore -d -k myserver.mydomain.com
```

In this case, the credentials for the server `myserver.mydomain.com` will be removed from the store.

## Additional Information

More information on the credstore utility can be found by running:

```
/opt/LifeKeeper/bin/credstore -man
```

This will show the entire man/help page for the command.

## 5.4.2.13. Standby Node Health Check

---

### Overview

The Standby Node Health Check feature allows you to monitor CPU and memory utilization on the standby node and monitor the health of out-of-service resources to detect errors on the standby node. This allows for issues to be resolved in advance, reducing the risk of an unsuccessful failover, if a failure occurs on the active node. This monitoring is performed at the same interval as the normal LifeKeeper resource monitoring (*/etc/default/LifeKeeper* setting LKCHECKINTERVAL).

The Standby Node Health Check performs the following two functions:

#### Node Monitoring

If all resources on a node are out of service, LifeKeeper considers it a standby node and calls the node monitoring script. The node monitoring script monitors CPU and memory utilization. If it determines that the node cannot be switched to successfully (due to high CPU or memory load), it sends this information to the administrator by email or SNMP event forwarding. See [Node Monitoring](#) for details.

#### Out-of-Service (OSU) Resource Monitoring

For each out-of-service (OSU) resource, *lkcheck* periodically calls the *OSUquickCheck* script. The *OSUquickCheck* script performs a quick health check for the resource. If it determines that the resource cannot start successfully, it changes the state of the resource to OSF and sends this information to the administrator by email or SNMP event forwarding. See [OSU Resource Monitoring](#) for details.

## Installation and Configuration

There is no special installation required.

### Setting up Standby Node Health Check

1. Configure email notification and event forwarding via SNMP2.
2. Configure Standby Node Health Check (Set the SNHC settings in the */etc/default/LifeKeeper* configuration file. See [Standby Node Health Check Parameters List](#) for details.)
3. If LifeKeeper is already started, restart the *lkcheck* process in order to reflect the configuration. Run the following command to restart the *lkcheck* process:

```
killall lkcheck
```

Once the above steps are completed, the Standby Node Health Check is activated on that node.

## 5.4.2.13.1. Node Monitoring

If all resources on a node are out of service, LifeKeeper considers it a standby node and calls the node monitoring script. The node monitoring script monitors CPU and memory utilization. If it determines that the node cannot be switched to successfully (due to high CPU or memory load), it sends this information to the administrator by email or SNMP event forwarding. This monitoring is performed at the same interval as the normal LifeKeeper resource monitoring (*/etc/default/LifeKeeper* setting LKCHECKINTERVAL).

### Monitored Resources

The following can be monitored with Node Monitoring:

Resource Name	Monitoring Details
<i>CPU Utilization</i>	Check CPU Utilization in <i>/proc/stat</i> file
<i>Memory Utilization</i>	Check Memory Utilization in <i>/proc/meminfo</i> file

### Node Monitoring Configuration

Set the SNHC\_CPUCHECK and SNHC\_MEMCHECK settings in the */etc/default/LifeKeeper* configuration file. You will also need to configure the following settings. See [Standby Node Health Check Parameters List](#) for details.

- SNHC\_CPUCHECK\_THRESHOLD
- SNHC\_CPUCHECK\_TIME
- SNHC\_MEMCHECK\_THRESHOLD
- SNHC\_MEMCHECK\_TIME

## 5.4.2.13.2. OSU Resource Monitoring

For each out-of-service (OSU) resource, *lkcheck* periodically calls the *OSUquickCheck* script for the resource. The *OSUquickCheck* script performs a quick health check for the resource. If it determines that the resource cannot start successfully, it changes the state of the resource to OSF and sends this information to the administrator by email or SNMP event forwarding. This monitoring is performed at the same interval as the normal LifeKeeper resource monitoring (*/etc/default/LifeKeeper* setting *LKCHECKINTERVAL*).

### Monitored Resources

The following can be monitored with OSU Resource Monitoring:

Resource Name	Monitoring Details
<i>IP Resource</i>	Verify the NIC link is up (disable with <i>/etc/default/LifeKeeper</i> setting <i>IP_NOLINKCHECK=1</i> ).  Also, verify network reachability (if a ping list is configured).
<i>Disk or DMMP resource(s)</i>	Verify that the paths to the monitored disk are functional by using commands for each resource.
NAS Resource	Verify that NFS access is available for the NFS server. Refer to <a href="#">NAS Configuration Considerations</a> for information on the timeout value for NFS access.

### OSU Resource Monitoring Configuration

Set the *SNHC\_IPCHECK* and *SNHC\_DISKCHECK* settings in the */etc/default/LifeKeeper* configuration file. You may also need to configure the following setting. See [Standby Node Health Check Parameters List](#) for details.

- *SNHC\_IPCHECK\_SLEEPTIME*

### Recovery from Failure

If an error is detected during OSU resource monitoring, the state of the corresponding resource is changed to OSF (out of service with failure). When the status is changed, OSU resource monitoring is no longer performed for the resource. After checking the details of the notified failure and addressing it, you should change the resource state to OSU. The state can be changed from OSF to OSU using the following command:

```
/opt/LifeKeeper/lkadm/bin/retstate <resource tag>
```

## 5.4.3. LifeKeeper Administration Overview

---

LifeKeeper does not require administration during operation. LifeKeeper works automatically to monitor protected resources and to perform the specified recovery actions if a fault should occur. You use the LifeKeeper GUI in these cases:

- **Resource and hierarchy definition.** LifeKeeper provides these interface options:
  - LifeKeeper GUI
  - LifeKeeper command line interface
- **Resource monitoring.** The LifeKeeper GUI provides access to resource status information and to the LifeKeeper logs.
- **Manual intervention.** You may need to stop servers or specific resources for maintenance or other administrative actions. The LifeKeeper GUI provides menu functions that allow you to bring specific resources in and out of service. Once applications have been placed under LifeKeeper protection, they should be started and stopped only through these LifeKeeper interfaces. Starting and stopping LifeKeeper is done through the command line only.

See [GUI Tasks](#) and [Maintenance Tasks](#) for detailed instructions on performing LifeKeeper administration, configuration and maintenance operations.



**Note:** All LifeKeeper executable scripts and programs run via the command line require super user authority.

A super user granted permissions by running the “su” or “sudo” command is able to execute LifeKeeper commands. However, SIOS Technology Corp. has tested executing LifeKeeper commands via the root user only.

---

[Error Detection and Notification](#)

[N-Way Recovery](#)

[Administrator Tasks](#)

[Editing Server Properties](#)

[Creating a Communication Path](#)

[Deleting a Communication Path](#)

[Server Properties – Failover](#)

[Creating Resource Hierarchies](#)

[Creating a File System Resource Hierarchy](#)

[Creating a Generic Application Resource Hierarchy](#)

[Creating a Raw Device Resource Hierarchy](#)

[QSP Resource Hierarchy](#)

[Editing Resource Properties](#)

[Editing Resource Priorities](#)

[Extending Resource Hierarchies](#)

[Extending a File System Resource Hierarchy](#)

[Extending a Generic Application Resource Hierarchy](#)

[Extending a Raw Device Resource Hierarchy](#)

[Unextending a Hierarchy](#)

[Creating a Resource Dependency](#)

[Deleting a Resource Dependency](#)

[Deleting a Hierarchy from All Servers](#)

## 5.4.3.1. Error Detection and Notification

---

The ability to provide detection and alarming for problems within an application is critical to building the best total fault resilient solution. Since every specific application varies on the mechanism and format of failures, no one set of generic mechanisms can be supplied. In general, however, many application configurations can rely on the Core system error detection provided within LifeKeeper. Two common fault situations are used to demonstrate the power of LifeKeeper's core facilities in the topics [Resource Error Recovery Scenario](#) and [Server Failure Recovery Scenario](#).

LifeKeeper also provides a complete environment for defining errors, alarms, and events that can trigger recovery procedures. This interfacing usually requires pattern match definitions for the system error log (`/var/log/messages`), or custom-built application specific monitor processes.

## 5.4.3.2. N-Way Recovery

---

N-Way recovery allows different resources to fail over to different backup servers in a cluster.

Return to [Protected Resources](#)

## 5.4.3.3. Administrator Tasks

---

[Editing Server Properties](#)

[Creating a Communication Path](#)

[Deleting a Communication Path](#)

[Server Properties – Failover](#)

[Creating Resource Hierarchies](#)

[Editing Resource Properties](#)

[Editing Resource Priorities](#)

[Extending Resource Hierarchies](#)

[Unextending a Hierarchy](#)

[Creating a Resource Dependency](#)

[Deleting a Resource Dependency](#)

[Deleting a Hierarchy from All Servers](#)

## 5.4.3.3.1. Editing Server Properties

---

1. To edit the properties of a server, bring up the Server Properties dialog just as you would for [viewing server properties](#).
2. If you are logged into that server with the appropriate permissions, the following items will be editable.
  - [Shutdown Strategy](#)
  - [Failover Confirmation](#)
3. Once you have made changes, the **Apply** button will be enabled. Clicking this button will apply your changes without closing the window.
4. When you are finished, click **OK** to save any changes and close the window, or **Cancel** to close the window without applying changes.

## 5.4.3.3.2. Creating a Communication Path

Before configuring a LifeKeeper communication path between servers, verify the hardware and software setup. For more information, see the [LifeKeeper for Linux Release Notes](#).

To create a communication path between a pair of servers, you must define the path individually on both servers. LifeKeeper allows you to create both TCP (TCP/IP) and TTY communication paths between a pair of servers. Only one TTY path can be created between a given pair. However, you can create multiple TCP communication paths between a pair of servers by specifying the local and remote addresses that are to be the end-points of the path. A priority value is used to tell LifeKeeper the order in which TCP paths to a given remote server should be used.

 **IMPORTANT:** Using a single communication path can potentially compromise the ability of servers in a cluster to communicate with one another. If a single comm path is used and the comm path fails, LifeKeeper hierarchies may come in service on multiple servers simultaneously. This is known as “false failover”. Additionally, heavy network traffic on a TCP comm path can result in unexpected behavior, including false failovers and LifeKeeper initialization problems.

1. There are four ways to begin.
  - Right-click on a server icon, then click **Create Comm Path** when the [server context menu](#) appears.
  - On the [global toolbar](#), click the **Create Comm Path** button.
  - On the [server context toolbar](#), if displayed, click the **Create Comm Path** button.
  - On the [Edit menu](#), select **Server**, then **Create Comm Path**.
2. A dialog entitled **Create Comm Path** will appear. For each of the options that follow, click **Help** for an explanation of each choice.
3. Select the **Local Server** from the list box and click **Next**.
4. Select one or more **Remote Servers** in the list box. If a remote server is not listed in the list box (i.e. it is not yet connected to the cluster), you may enter it using **Add**. You must make sure that the network addresses for both the local and remote servers are resolvable (for example, with DNS or added to the `/etc/hosts` file). Click **Next**.
5. Select either **TCP** or **TTY** for **Device Type** and click **Next**.
6. Select one or more **Local IP Addresses** if the **Device Type** was set for **TCP**. Select the **Local TTY Device** if the **Device Type** was set to **TTY**. Click **Next**.
7. Select the **Remote IP Address** if the **Device Type** was set for **TCP**. Select the **Remote TTY Device** if the **Device Type** was set to **TTY**. Click **Next**.

8. Enter or select the **Priority** for this comm path if the **Device Type** was set for **TCP**. Enter or select the **Baud Rate** for this Comm Path if the **Device Type** was set to **TTY**. Click **Next**.
9. Click **Create**. A message should be displayed indicating the network connection is successfully created. Click **Next**.
10. If you selected multiple Local IP Addresses or multiple Remote Servers and the **Device Type** was set for **TCP**, then you will be taken back to Step 6 to continue with the next Comm Path. If you selected multiple Remote Servers and the **Device Type** was set for **TTY**, then you will be taken back to Step 5 to continue with the next Comm Path.
11. Click **Done** when presented with the concluding message.

You can verify the comm path by viewing the [Server Properties Dialog](#) or by entering the command `lcdstatus -q`. See the `LCD` man page for information on using `lcdstatus`. You should see an **ALIVE** status.

In addition, check the server icon in the right pane of the GUI. If this is the first comm path that has been created, the server icon shows a yellow heartbeat, indicating that one comm path is **ALIVE**, but there is no redundant comm path. 

The server icon will display a green heartbeat when there are at least two comm paths **ALIVE**. 

 **IMPORTANT:** When using IPv6 addresses to create a comm path, statically assigned addresses should be used instead of auto-configured/stateless addresses as the latter may change over time which will cause the comm path to fail.

If the comm path does not activate after a few minutes, verify that the paired server's computer name is correct. If using TTY comm paths, verify that the cable connection between the two servers is correct and is not loose. Use the `portio(1M)` command if necessary to verify the operation of the TTY connection.

## 5.4.3.3.3. Deleting a Communication Path

---

1. There are four ways to begin.
  - Right-click on a server icon, then click **Delete Comm Path** when the [server context menu](#) appears.
  - On the [global toolbar](#), click the **Delete Comm Path** button.
  - On the [server context toolbar](#), if displayed, click the **Delete Comm Path** button.
  - On the [Edit menu](#), select **Server**, then **Delete Comm Path**.
2. A dialog entitled Delete Comm Path will appear. For each of the options that follow, click **Help** for an explanation of each choice.
3. Select **Local Server** from the list and click **Next**. This dialog will only appear if the delete is selected using the **Delete Comm Path** button on the [global toolbar](#) or via the [Edit menu](#) selecting **Server**.
4. Select the communications path(s) that you want to delete and click **Next**.
5. Click **Delete Comm Path(s)**. If the output panel is enabled, the dialog closes, and the results of the commands to delete the communications path(s) are shown in the [output panel](#). If not, the dialog remains up to show these results, and you click **Done** to finish when all results have been displayed. A message should be displayed indicating the network connection is successfully removed
6. Click **Done** to close the dialog and return to the GUI status display.

## 5.4.3.3.4. Server Properties – Failover

---

In the event that the primary server has attempted and failed local recovery, or failed completely, most server administrators will want LifeKeeper to automatically restore the protected resource(s) to a backup server. This is the default LifeKeeper behavior. However, some administrators may not want the protected resource(s) to automatically go in-service at a recovery site. For example, if LifeKeeper is installed in a WAN environment where the network connection between the servers may not be reliable in a disaster recovery situation.

Automatic failover is enabled by default for all protected resources. To disable automatic failover for protected resources or to prevent automatic failover to a backup server, use the **Failover** section located on the **General** tab of Server Properties to configure as follows:

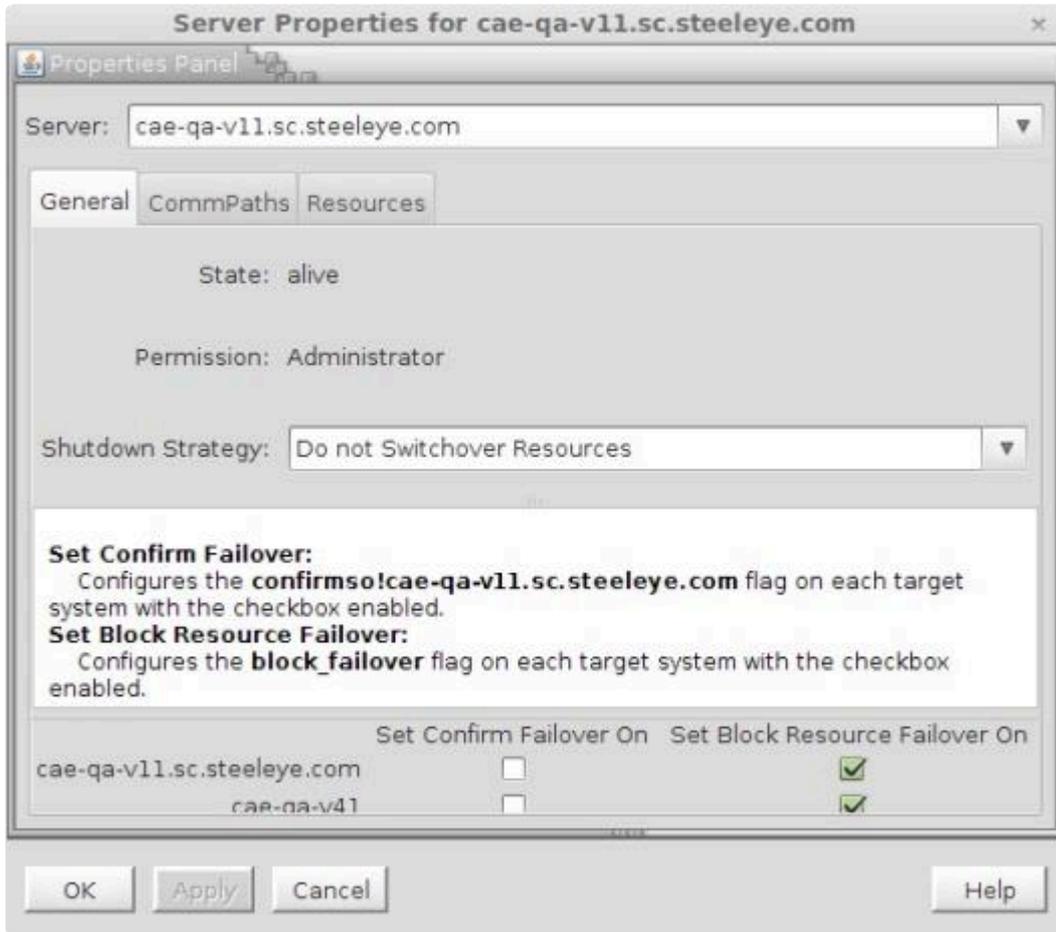
For each server in the cluster:

1. Bring up the **Server Properties** dialog just as you would for [viewing server properties](#).
2. Select the **General** tab. In the **Failover** section of the Server Properties dialog, check the server to disable system and resource failover capabilities. By default, all failover capabilities of LifeKeeper are enabled.

In the **Set Confirm Failover On** column, select the server to be disqualified as a backup server for a complete failure of the local server.

In the **Set Block Resource Failover On** column, select the server to be disqualified as a backup server for any failed resource hierarchy on this local server. Resource failovers cannot be disabled without first disabling system failover capabilities.

To commit your selections, press the **Apply** button.



Refer to [\[Confirm Failover\]](#) and [\[Block Resource Failover\] Settings](#) for configuration details.

## 5.4.3.3.5. Creating Resource Hierarchies

1. There are four ways to begin creating a resource hierarchy.

✿ A Quorum node cannot be used to create hierarchies since the node does not have all of the kit licenses.

- Right-click on a **server icon** to bring up the [server context menu](#), then click on **Create Resource Hierarchy**.
  - On the [global toolbar](#), click on the **Create Resource Hierarchy** button.
  - On the [server context toolbar](#), if displayed, click on the **Create Resource Hierarchy** button.
  - On the [Edit menu](#), select **Server**, then click on **Create Resource Hierarchy**.
2. A dialog entitled Create Resource Wizard will appear with a list of all recognized recovery kits installed within the cluster. Select the **Recovery Kit** that builds resource hierarchies to protect your application and click **Next**.
  3. Select the **Switchback Type** and click **Next**.
  4. Select the **Server** and click **Next**. **Note:** If you began from the server context menu, the server will be determined automatically from the server icon that you clicked on, and this step will be skipped.
  5. Continue through the succeeding dialogs, entering whatever data is needed for the type of resource hierarchy that you are creating.

## LifeKeeper Application Resource Hierarchies

If you install LifeKeeper without any recovery kits, the Select Recovery Kit list includes options for File System or Generic Application by default. The Generic Application option may be used for applications that have no associated recovery kits.

If you install the Raw I/O or IP Recovery Kits (both of which are Core Recovery Kits that are packaged separately and included on the LifeKeeper Core media), the Select Recovery Kit list will provide additional options for these Recovery Kits.

See the following topics describing these available options:

[Creating a File System Resource Hierarchy](#)

[Creating a Generic Application Resource Hierarchy](#)

[Creating a Raw Device Resource Hierarchy](#)

See the [IP Recovery Kit Technical Documentation](#) for more information.

## Recovery Kit Options

Each optional recovery kit that you install adds entries to the Select Recovery Kit list; for example, you may see Oracle, Apache, and NFS Recovery Kits. Refer to the Administration Guide that accompanies each recovery kit for directions on creating the required resource hierarchies.

 **Note:** If you wish to create a File System or other application resource hierarchy that is built on a logical volume, you must first have the [Logical Volume Manager \(LVM\) Recovery Kit](#) installed.

## 5.4.3.3.5.1. Creating a File System Resource Hierarchy

---

Use this option to protect a file system on storage that is directly accessible by two or more servers using the SCSI protocol (over a SCSI bus or iSCSI) or with network attached storage (requires the [NAS Recovery Kit](#)). To create a replicated file system hierarchy using SIOS DataKeeper, see [Creating a DataKeeper Resource Hierarchy](#).

1. There are four ways to begin creating a file system resource hierarchy.
  - Right-click on a server icon to bring up the [server context menu](#), then click on **Create Resource Hierarchy**.
  - On the [global toolbar](#), click on the **Create Resource Hierarchy** button.
  - On the [server context toolbar](#), if displayed, click on the **Create Resource Hierarchy** button.
  - On the [Edit menu](#), select **Server**, then click on **Create Resource Hierarchy**.
2. A dialog entitled *Create Resource Wizard* will appear with a **Recovery Kit** list. Select *File System Resource* and click **Next**.
3. Select the **Switchback Type** and click **Next**.
4. Select the **Server** and click **Next**. **Note:** *If you began from the server context menu, the server will be determined automatically from the server icon that you clicked on, and this step will be skipped.*
5. The Create *gen/filesys Resource* dialog will now appear. Select the **Mount Point** for the file system resource hierarchy and click **Next**. The selected mount point will be checked to see that it is shared with another server in the cluster by checking each storage kit to see if it recognizes the mounted device as shared. If no storage kit recognizes the mounted device, then an error dialog will be presented:

**<file system>** is not a shared file system

Selecting **OK** will return to the *Create gen/filesys Resource* dialog.

### Notes:

- In order for a mount point to appear in the choice list, the mount point must be currently mounted. If an entry for the mount point exists in the `/etc/fstab` file, LifeKeeper will remove this entry during the creation and extension of the hierarchy. It is advisable to make a backup of `/etc/fstab` prior to using the NAS Recovery Kit, especially if you have complex mount settings. You can direct that entries are re-populated back into `/etc/fstab` on deletion by setting the `/etc/default/LifeKeeper` tunable `REPLACEFSTAB=true|TRUE`.

- Many of these resources (SIOS DataKeeper, LVM, Device Mapper Multipath, etc.) require LifeKeeper recovery kits on each server in the cluster in order for the file system resource to be created. If these kits are not properly installed, then the file system will not appear to be shared in the cluster.

6. LifeKeeper creates a default **Root Tag** for the file system resource hierarchy. (This is the label used for this resource in the status display). You can select this root tag or create your own, then click **Next**.
7. Click **Create Instance**. A window will display a message indicating the status of the instance creation.
8. Click **Next**. A window will display a message that the file system hierarchy has been created successfully.
9. At this point, you can click **Continue** to move on to [extending the file system resource hierarchy](#), or you can click **Cancel** to return to the GUI. If you click **Cancel**, you will receive a warning message that your hierarchy exists on only one server, and it is not protected at this point.

## 5.4.3.3.5.2. Creating a Generic Application Resource Hierarchy

Use this option to protect a user-defined application that has no associated recovery kit. Templates are provided for the user supplied scripts referenced below in `$LKROOT/1kadm/subsys/gen/app/templates`. Copy these templates to another directory before customizing them for the application that you wish to protect and testing them.

 **Note:** For applications depending upon other resources such as a file system, disk partition, or IP address, create each of these resources separately, and use Create Dependency to create the appropriate dependencies.

1. There are four ways to begin creating a generic application resource hierarchy.
  - Right-click on a server icon to bring up the [server context menu](#), then click on **Create Resource Hierarchy**.
  - On the [global toolbar](#), click on the **Create Resource Hierarchy** button.
  - On the [server context toolbar](#), if displayed, click on the **Create Resource Hierarchy** button.
  - On the [Edit menu](#), select **Server**, then click on **Create Resource Hierarchy**.

2. A dialog entitled *Create Resource Wizard* will appear with a **Recovery Kit** list. Select **Generic Application** and click **Next**.

3. Select the **Switchback Type** and click **Next**

4. Select the **Server** and click **Next**

**Note:** *If you began from the server context menu, the server will be determined automatically from the server icon that you clicked on, and this step will be skipped.*

5. On the next dialog, enter the path to the **Restore Script** for the application and click **Next**. This is the command that starts the application. A template restore script, `restore.template`, is provided in the templates directory. The restore script must not impact applications that are already started.

6. Enter the path to the **Remove Script** for the application and click **Next**. This is the command that stops the application. A template remove script, `remove.template`, is provided in the templates directory.

7. Enter the path to the **quickCheck Script** for the application and click **Next**. This is the command that monitors the application. A template quickCheck script, `quickCheck.template`, is provided in the templates directory.

8. Enter the path to the **Local Recovery Script** for the application and click **Next**. This is the command that attempts to restore a failed application on the local server. A template recover script, `recover.template`, is provided in the `templates` directory.
9. Enter any **Application Information** and click **Next**. This is optional information about the application that may be needed by the `restore`, `remove`, `recover`, and `quickCheck` scripts.
10. Select either **Yes** or **No** for **Bring Resource In Service**, and click **Next**. Selecting **No** will cause the resource state to be set to `OSU` following the `create`; selecting **Yes** will cause the previously provided `restore` script to be executed. For applications depending upon other resources such as a file system, disk partition, or IP address, select **No** if you have not already created the appropriate dependent resources.
11. Enter the **Root Tag**, which is a unique name for the resource instance. (This is the label you will see for this resource in the status display.)
12. Click **Create Instance** to start the creation process. A window will display a message indicating the status of the instance creation.
13. Click **Next**. A window will display a message that the hierarchy has been created successfully.
14. At this point, you can click **Continue** to move on to [extending the generic application resource hierarchy](#), or you can click **Cancel** to return to the GUI. If you click **Cancel**, you will receive a warning that your hierarchy exists on only one server, and it is not protected at this point.

**Note:** The scripts which are provided when resource hierarchy is created, such as `restore`, `remove`, `quickCheck`, are located in each directory under `LKROOT/subsys/gen/resource/app/`.

- `restore` – `actions/!restore/<tag name>`
- `remove` – `actions/!remove/<tag name>`
- `quickCheck` – `actions/!quickCheck/<tag name>`
- `recover` – `recovery/!recover/<tag name>`

## 5.4.3.3.5.3. Creating a Raw Device Resource Hierarchy

Use this option to protect a raw device resource. For example, if you create additional table space on a raw device that needs to be added to an existing database hierarchy, you would use this option to create a raw device resource.

 **Note:** LifeKeeper locks shared disk partition resources at the disk logical unit (or LUN) level to one system in a cluster at a time.

1. There are four ways to begin creating a raw device resource hierarchy.
  - Right-click on a server icon to bring up the [server context menu](#), then click on **Create Resource Hierarchy**.
  - On the [global toolbar](#), click on the **Create Resource Hierarchy** button.
  - On the [server context toolbar](#), if displayed, click on the **Create Resource Hierarchy** button.
  - On the [Edit menu](#), select **Server**, then click on **Create Resource Hierarchy**.
2. A dialog entitled Create Resource Wizard will appear with a **Recovery Kit** list. Select Raw Device and click **Next**.
3. Select the **Switchback Type** and click **Next**.
4. Select the **Server** and click **Next**.

**Note:** *If you began from the server context menu, the server will be determined automatically from the server icon that you clicked on, and this step will be skipped.*
5. Select the **Raw Partition** on a shared storage device where this resource will reside, and click **Next**.
6. Enter the **Root Tag**, which is a unique name for the resource instance. (This is the label you will see for this resource in the status display.)
7. Click **Create Instance** to start the creation process. A window titled Creating scsi/raw resource will display text indicating what is happening during creation.
8. Click **Next**. A window will display a message that the hierarchy has been created successfully.
9. At this point, you can click **Continue** to move on the [extending the raw resource hierarchy](#), or you can click **Cancel** to return to the GUI. If you click **Cancel**, you will receive a message warning that your hierarchy exists on only one server, and it is not protected at this point

## 5.4.3.3.5.4. Quick Service Protection (QSP) Recovery Kit

---

The QSP Recovery Kit provides a simplified method to protect the OS service. With the QSP Recovery Kit, users can easily create a LifeKeeper resource instance to protect an OS service provided that service can be started and stopped by the OS service command or the systemctl command (start/stop). The service can also be protected via the Generic Application Recovery Kit but the use of that kit requires code development, whereas the QSP Recovery Kit does not. Also, by creating a dependency relationship, protected services can be started and stopped in conjunction with the application that requires the service.

The QSP Recovery Kit quickCheck can only perform simple health checks (using the “status” action of the service command). QSP doesn’t guarantee that the service is provided or the process is functioning. If complicated starting and/or stopping is necessary, or more robust health checking operations are necessary, using a Generic Application is recommended.

### Requirements

The service to be protected by the QSP Recovery Kit needs to meet the following requirements.

- It must support start and stop actions via the OS service command or the systemctl command. Also, it must return 0 when the start and stop action succeeds.
- To perform health checks the service must support the status action via the OS service command or the systemctl command. If it does not support the status action then quickCheck health check operations must be disabled. Also, it must return 0 when the status action succeeds.
- The name of the service to be protected must not exceed 256 characters in length and can contain only alphanumeric characters.

 **Note:** The compatible service command may be used to control protected resources even in a systemd environment.

The service to be protected by the QSP resource must be running (started) before attempting a resource create. Please notice that some services which are already supplied with a dedicated Recovery Kit are not the target of QSP (hereinafter referred as “the Services not targeted by QSP protection”) and cannot be protected by the QSP Recovery Kit.

### Create the QSP Resource Hierarchy

This option is used to protect OS services via the QSP Recovery Kit.

1. There are 4 methods to start the creation of a QSP resource instance.

- Right-click on a server icon to bring up the [server context menu](#), then click on [Create Resource Hierarchy].
  - On the [global toolbar](#), click on the [Create Resource Hierarchy] button.
  - On the [server context toolbar](#), if displayed, click on the [Create Resource Hierarchy] button.
  - On the [Edit menu](#), select [Server], then click on [Create Resource Hierarchy].
2. A dialogue box titled [Create Resource Wizard] is displayed. In the [Recovery Kit] drop down is a list of available resource types to create. Select **Quick Service Protection** and click [Next].
  3. Select [Switchback Type] and click [Next].
  4. Select [Server] and click [Next].

**Note:** If the create was started via the server context menu, this step is skipped because the server is detected based on the start context (defaults to the name of the server on which the create process started).

5. The next dialog box contains a drop down of the available services that can be protected. Select the [Service Name] to be protected and click [Next].

**Note:** The list may not show the service if it is not running. In this case, click **Cancel** to discontinue the process, and start the service. Once the service is running restart the create process. The list will not show the Services not targeted by QSP protection.

6. In the next dialog box the quickCheck action is configured. To enable the quickCheck monitoring function, select [enable]. To disable it, select [disable]. Click [Next] to continue. The quickCheck action can be changed at any time.

**Note:** If the selected service does not support the “status” action via the OS service command, set the quickCheck action to “disabled” because the QSP Recovery Kit cannot monitor the service state.

7. Input the [Resource Tag] This is a unique name for the resource instance. (This is the label that uniquely identifies the resource instance and is used whenever displaying LifeKeeper protected resource instances in UI.)
8. Click [Create Instance] to start the creation process. The status of the resource instance creation is displayed in the status window.
9. Click [Next] to display the resource extension dialog. Click [Next] to begin the extension process or click [Cancel] to go back to the GUI. When [Cancel] is clicked, an alert is indicating that the hierarchy exists on only one server and protection by LifeKeeper is not available at this time.

## Extending the QSP Resource Hierarchy

This function, as explained in the section [Extending Resource Hierarchies](#), starts automatically after finishing the Create QSP Resource Hierarchy (URL) process or from right clicking on an existing QSP resource and selecting [Extend Resource Hierarchy]. After finishing the pre-extend process, complete the following steps.

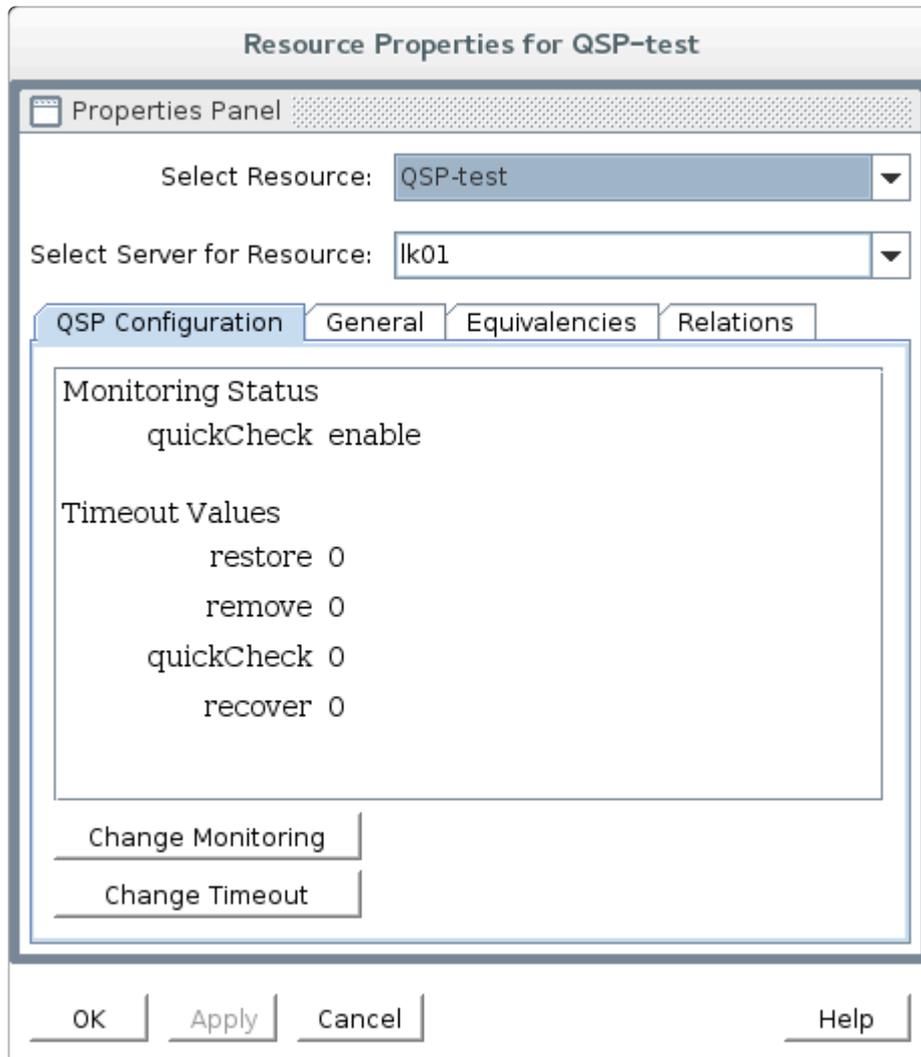
1. Select [Resource Tag] provided by LifeKeeper or input a unique tag for the resource hierarchy on the target server.
2. Click [Extend] to start the extension process. The status of the extension process is displayed in the dialogue box. When it is finished it will show a message indicating the hierarchy is correctly extended. If the hierarchy is to be extended to another server, click [Next Server], otherwise click [Finish] to complete the extension. If [Next Server] is selected, the extension operation is repeated.
3. When [Finish] is clicked the integrity of the hierarchy is checked. If any problems are detected the extension is reversed. To complete the verification and close the dialog box click [Done].

## QSP Resource Configuration

The following parameters are unique to each QSP resource instance and are available for modification.

Set Up Items		Default Value	Description
Monitoring	quickCheck	Specified when creating the resource	Set to enable to check the status of the service or to disable / skip the monitoring function
Time Out	restore	0	Specify the restore timeout (unit: seconds). If set to 0, no timeout occurs when restoring the resource instance.
	remove	0	Specify the remove timeout (unit: seconds). If set to 0, no timeout occurs when removing the resource instance.
	quickCheck	0	Specify the quickCheck timeout (unit: seconds). If set to 0, no time out occurs when performing health checking of the resource instance.
	recover	0	Specify the recover timeout (unit: seconds). If set to 0, no timeout occurs during recovery of the resource instance.

Checking / changing of the set value is possible from the **QSP Configuration** tab by [Display Resource Properties](#) and must be performed on each node in the hierarchy. If the quickCheck function is disabled, quickCheck and recover of timeouts are not displayed and thus cannot be changed.



## How to Change the Monitoring Function

1. Display the [QSP Configuration] tab of the resource properties and click [Change quickCheck]
2. Select [enable] to enable quickCheck, or [disable] to disable it.
3. Clicking [Change] starts the change process and displays the change process message.
4. Finish by clicking [Done].

**Note:** Modification of these values is a per node operation. If the same change is needed on another node, the process must be repeated on that node.

**Change monitoring for QSP-test**

Enable or Disable monitoring

enable

enable

disable

Select "enable" or "disable" for quickCheck. If "enable" is selected, LifeKeeper will provide monitoring for using the service command.

<BackChangeCancelHelp

## How to Change the Timeout Value

1. Display the [QSP Configuration] tab of resource properties and click [Change Timeout].
2. Select the timeout action to be changed (restore, remove, quickCheck or recover) and click **Next**.

**Note:** [quickCheck] and [recover] timeouts are not displayed in the list if the monitoring function is disabled.

3. Input the timeout value in seconds.

**Note:** Input decimal numbers only. Non numerical characters are invalid.

4. Clicking [Change] starts the timeout change process and displays the change process messages.
5. Finish by clicking [Done].

**Note:** Modification of these values is a per node operation. If the same change is needed on another node, the process must be repeated on that node.

### Change action timeout(s) for QSP-test

Please Select an Action

restore	▼
restore	
remove	
quickCheck	
recover	

Select the action name for the timeout that will be updated for **QSP-test**.

[<Back](#)   [Next>](#)   [Cancel](#)   [Help](#)

### Change action timeout(s) for QSP-test

Timeout for restore in seconds

0	▼

Enter the new timeout value for the restore action.

[<Back](#)   [Change](#)   [Cancel](#)   [Help](#)

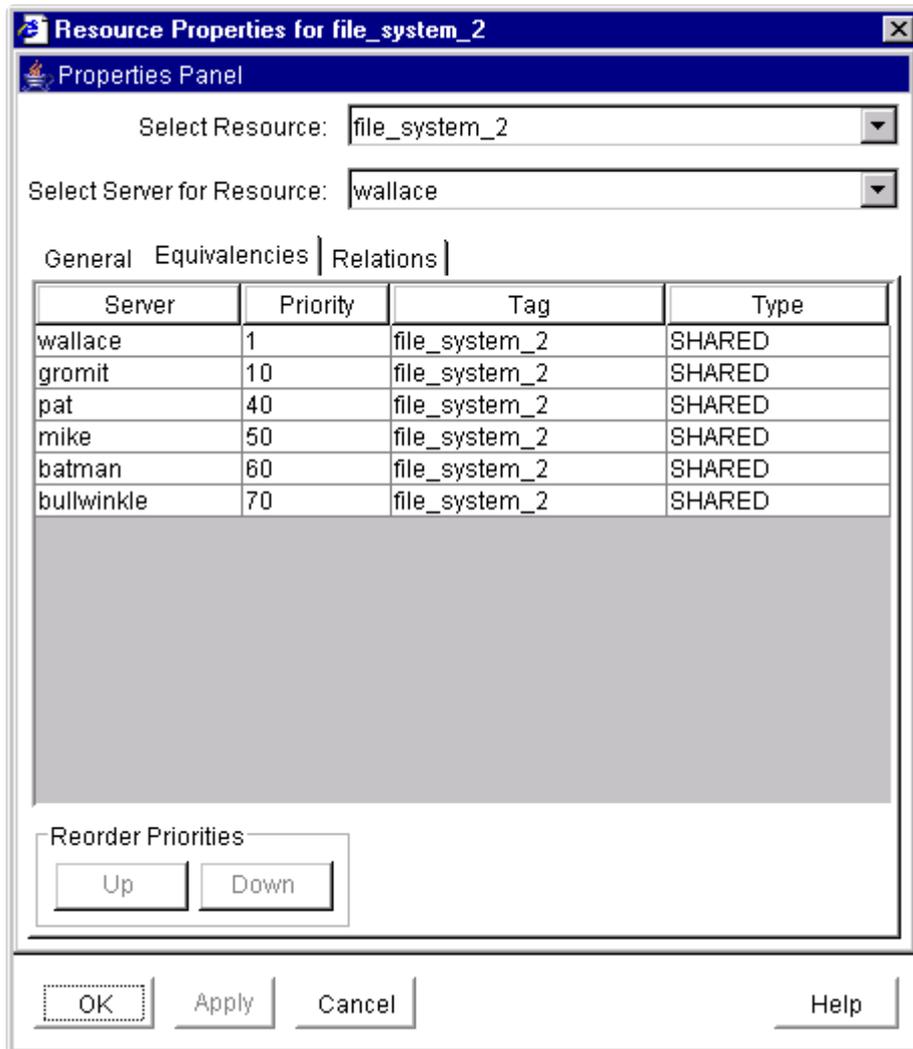
## 5.4.3.3.6. Editing Resource Properties

---

1. To edit the properties of a resource, bring up the Resource Properties dialog just as you would for [viewing resource properties](#).
2. If you are logged into that server with the appropriate permissions, the following items will be editable.
  - Switchback
  - Resource Configuration (only for resources with specialized configuration settings)
  - [Resource Priorities](#)
3. Once you have made changes, the **Apply** button will be enabled. Clicking this button will apply your changes without closing the window.
4. When you are finished, click **OK** to save any changes and close the window, or **Cancel** to close the window without applying changes.

## 5.4.3.3.7. Editing Resource Priorities

You can edit or reorder the priorities of servers on which a resource hierarchy has been defined. First, bring up the Resource Properties dialog just as you would for [viewing resource properties](#). The Resource Properties dialog displays the priority for a particular resource on a server in the Equivalencies Tab as shown below.



There are two ways to modify the priorities:

- ◦ Reorder the priorities by moving an equivalency with the **Up/Down** buttons ,or
- ◦ Edit the priority values directly.

### Using the Up and Down Buttons

1. Select an equivalency by clicking on a row in the Equivalencies table. The **Up** and/or **Down** buttons will become enabled, depending on which equivalency you have selected. The **Up** button is enabled unless you have selected the highest priority server. The **Down** button is enabled unless you have selected the lowest priority server.

2. Click **Up** or **Down** to move the equivalency in the priority list.

The numerical priorities column will not change, but the equivalency will move up or down in the list.

## Editing the Priority Values

1. Select a priority by clicking on a priority value in the Priority column of the Equivalencies table. A box appears around the priority value, and the value is highlighted.
2. Enter the desired priority and press **Enter**.

 **Note:** Valid server priorities are 1 to 999.

After you have edited the priority, the Equivalencies table will be re-sorted.

## Applying Your Changes

Once you have the desired priority order in the Equivalencies table, click **Apply** (or **OK**) to commit your changes. The **Apply** button applies any changes that have been made. The **OK** button applies any changes that have been made and then closes the window. The **Cancel** button closes the window without saving any changes made since **Apply** was last clicked.

## 5.4.3.3.8. Extending Resource Hierarchies

---

The LifeKeeper **Extend Resource Hierarchy** option copies an existing hierarchy from one server and creates a similar hierarchy on another LifeKeeper server. Once a hierarchy is extended to other servers, cascading failover is available for that resource. The server where the existing hierarchy currently resides is referred to as the template server. The server where the new extended hierarchy will be placed is referred to as the target server.

The target server must be capable of supporting the extended hierarchy and it must be able to communicate with equivalent hierarchies on other remote servers (via active LifeKeeper communications paths). This means that all recovery kits associated with resources in the existing hierarchy must already be installed on the target server, as well as every other server where the hierarchy currently resides.

1. There are five ways to extend a resource hierarchy through the GUI.
  - [Create](#) a new resource hierarchy. When the dialog tells you that the hierarchy has been created, click on the **Continue** button to start extending your new hierarchy via the Pre-Extend Wizard.
  - Right-click on a global or server-specific resource icon to bring up the [resource context menu](#), then click on Extend Resource Hierarchy to extend the selected resource via the Pre-Extend Wizard.
  - On the [global toolbar](#), click on the **Extend Resource Hierarchy** button. When the Pre-Extend Wizard dialog appears, select a **Template Server** and a **Tag to Extend**, clicking on **Next** after each choice.
  - On the [resource context toolbar](#), if displayed, click on the **Extend Resource Hierarchy** button to bring up the Pre-Extend Wizard.
  - On the [Edit menu](#), select **Resource**, then click on **Extend Resource Hierarchy**. When the Pre-Extend Wizard dialog appears, select a **Template Server** and a **Tag to Extend**, clicking on **Next** after each choice.
2. Either select the default **Target Server** or enter one from the list of choices, then click **Next**.
3. Select the **Switchback Type**, then click **Next**.
4. Either select the default or enter your own **Template Priority**, then click **Next**.
5. Either select or enter your own **Target Priority**, then click **Next**.
6. The dialog will then display the pre-extend checks that occur next. If these tests succeed, LifeKeeper goes on to perform any steps that are needed for the specific type of resource that you are extending.

The **Accept Defaults** button which is available for the **Extend Resource Hierarchy** option is intended for the user who is familiar with the **LifeKeeper Extend Resource Hierarchy** defaults, and wants to quickly extend a LifeKeeper resource hierarchy without being prompted for input or confirmation. Users

who prefer to extend a LifeKeeper resource hierarchy using the interactive, step-by-step interface of the GUI dialogs should use the **Next** button.

\* **Note:** ALL roots in a multi-root hierarchy must be extended together, they may not be extended as single root hierarchies.

\* **Note:** For command line instructions, see [Extending the SAP Resource from the Command Line](#).

## 5.4.3.3.8.1. Extending a File System Resource Hierarchy

---

This operation can be started automatically after you have finished [creating a file system resource hierarchy](#), or from an existing file system resource, as described in the section on [extending resource hierarchies](#). Then complete the steps below, which are specific to file system resources.

1. The *Extend gen/filesys Resource Hierarchy* dialog box appears. Select the **Mount Point** for the file system hierarchy, then click **Next**.
2. Select the **Root Tag** that LifeKeeper offers, or enter your own tag for the resource hierarchy on the target server, then click **Next**.
3. The dialog displays the status of the extend operation, which should finish with a message saying that the hierarchy has been successfully extended. Click **Next Server** if you want to extend the same resource hierarchy to a different server. This will repeat the extend operation. Or click **Finish** to complete this operation.
4. The dialog then displays verification information as the extended hierarchy is validated. When this is finished, the **Done** button will be enabled. Click **Done** to finish.

## 5.4.3.3.8.2. Extending a Generic Application Resource Hierarchy

---

This operation can be started automatically after you have finished [creating a generic application resource hierarchy](#), or from an existing generic application resource, as described in the section on [extending resource hierarchies](#). Then complete the steps below, which are specific to generic application resources.

1. Select the **Root Tag** that LifeKeeper offers, or enter your own tag for the resource hierarchy on the target server, then click **Next**.
2. Enter any **Application Information** next (optional), then click **Next**.
3. The dialog displays the status of the extend operation, which should finish with a message saying that the hierarchy has been successfully extended. Click **Next Server** if you want to extend the same resource hierarchy to a different server. This will repeat the extend operation. Or click **Finish** to complete this operation.
4. The dialog then displays verification information as the extended hierarchy is validated. When this is finished, the **Done** button will be enabled. Click **Done** to finish.

## 5.4.3.3.8.3. Extending a Raw Device Resource Hierarchy

---

This operation can be started automatically after you have finished [creating a raw device resource hierarchy](#), or from an existing raw device resource, as described in the section on [extending resource hierarchies](#). Then complete the steps below, which are specific to raw device resources.

1. Select the **Root Tag** that LifeKeeper offers, or enter your own tag for the resource hierarchy on the target server, then click **Next**.
2. The dialog displays the status of the extend operation, which should finish with a message saying that the hierarchy has been successfully extended. Click **Next Server** if you want to extend the same resource hierarchy to a different server. This will repeat the extend operation. Or click **Finish** to complete this operation.
3. The dialog then displays verification information as the extended hierarchy is validated. When this is finished, the **Done** button will be enabled. Click **Done** to finish.

## 5.4.3.3.9. Unextending a Hierarchy

---

While most Recovery Kits create their dependencies during the original resource hierarchy creation task, under certain circumstances, you may want to create new or additional resource dependencies or delete existing ones. An example might be that you wish to change an existing IP dependency to another IP address. Instead of deleting the entire resource hierarchy and creating a new one, you can delete the existing IP dependency and create a new dependency with a different IP address.

1. There are four possible ways to begin.

- ° Right-click on the icon for the parent server-specific resource under the server, or the parent global resource, to which you want to add a parent-child dependency. When the [resource context menu](#) appears, click **Create Dependency**.

**Note:** If you right-clicked on a server-specific resource in the right pane, the value of the **Server** will be that server. If you right-clicked on a global resource in the left pane, the value of the **Server** will be the server where the resource has the highest priority.

- ° On the [global toolbar](#), click the **Create Dependency** button. When the dialog comes up, select the server in the **Server** list from which you want to begin creating the resource dependency, and click **Next**. On the next dialog, select the parent resource from the **Parent Resource Tag** list, and click **Next** again.

- ° On the [resource context toolbar](#), if displayed, click the **Create Dependency** button.

- ° On the [Edit menu](#), point to **Resource** and then click **Create Dependency**. When the dialog comes up, select the server in the **Server** list from which you want to begin creating the resource dependency, and click **Next**. On the next dialog, select the parent resource from the **Parent Resource Tag** list, and click **Next** again.

2. Select a **Child Resource Tag** from the drop down box of existing and valid resources on the server. The dialog will display all the resources available on the server with the following exceptions:

- ° The parent resource, its ancestors, and its children.

- ° A resource that has not been extended to the same servers as the parent resource.

- ° A resource that does not have the same relative priority as the parent resource.

- ° Any resource that is not in-service on the same server as the parent, if the parent resource is in-service.

Click **Next** to proceed to the next dialog.

3. The dialog will then confirm that you have selected the appropriate parent and child resource tags for your dependency creation. Click **Create Dependency** to create the dependency on all servers in the cluster to which the parent has been extended.
4. If the [output panel](#) is enabled, the dialog closes, and the results of the commands to create the dependency are shown in the output panel. If not, the dialog remains up to show these results, and you click **Done** to finish when all results have been displayed.

## 5.4.3.3.10. Creating a Resource Dependency

While most Recovery Kits create their dependencies during the original resource hierarchy creation task, under certain circumstances, you may want to create new or additional resource dependencies or delete existing ones. An example might be that you wish to change an existing IP dependency to another IP address. Instead of deleting the entire resource hierarchy and creating a new one, you can delete the existing IP dependency and create a new dependency with a different IP address.

1. There are four possible ways to begin.

- Right-click on the icon for the parent server-specific resource under the server, or the parent global resource, to which you want to add a parent-child dependency. When the [resource context menu](#) appears, click **Create Dependency**.

**Note:** If you right-clicked on a server-specific resource in the right pane, the value of the **Server** will be that server. If you right-clicked on a global resource in the left pane, the value of the **Server** will be the server where the resource has the highest priority.

- On the [global toolbar](#), click the **Create Dependency** button. When the dialog comes up, select the server in the **Server** list from which you want to begin creating the resource dependency, and click **Next**. On the next dialog, select the parent resource from the **Parent Resource Tag** list, and click **Next** again.
- On the [resource context toolbar](#), if displayed, click the **Create Dependency** button.
- On the [Edit menu](#), point to **Resource** and then click **Create Dependency**. When the dialog comes up, select the server in the **Server** list from which you want to begin creating the resource dependency, and click **Next**.
- On the next dialog, select the parent resource from the **Parent Resource Tag** list, and click **Next** again.

2. Select a **Child Resource Tag** from the drop down box of existing and valid resources on the server. The dialog will display all the resources available on the server with the following exceptions:

- The parent resource, its ancestors, and its children.
- A resource that has not been extended to the same servers as the parent resource.
- A resource that does not have the same relative priority as the parent resource.
- Any resource that is not in-service on the same server as the parent, if the parent resource is in-service.

Click **Next** to proceed to the next dialog.

3. The dialog will then confirm that you have selected the appropriate parent and child resource tags for your dependency creation. Click **Create Dependency** to create the dependency on all servers in the cluster to which the parent has been extended.
4. If the [output panel](#) is enabled, the dialog closes, and the results of the commands to create the dependency are shown in the output panel. If not, the dialog remains up to show these results, and you click **Done** to finish when all results have been displayed.

## 5.4.3.3.11. Deleting a Resource Dependency

1. There are four possible ways to begin.
  - Right-click on the icon for the parent server-specific resource under the server, or the parent global resource, from which you want to delete a parent-child dependency. When the [resource context menu](#) appears, click **Delete Dependency**.
  - On the [global toolbar](#), click the **Delete Dependency** button. When the dialog comes up, select the server in the **Server** list from which you want to begin deleting the resource dependency, and click **Next**. On the next dialog, select the parent resource from the **Parent Resource Tag** list, and click **Next** again.
  - On the [resource context toolbar](#), if displayed, click the **Delete Dependency** button.
  - On the [Edit menu](#), point to **Resource** and then click **Delete Dependency**. When the dialog comes up, select the server in the **Server** list from which you want to begin deleting the resource dependency, and click **Next**. On the next dialog, select the parent resource from the **Parent Resource Tag** list, and click **Next** again.
2. Select the **Child Resource Tag** from the drop down box. This should be the tag name of the child in the dependency that you want to delete. Click **Next** to proceed to the next dialog box.
3. The dialog then confirms that you have selected the appropriate parent and child resource tags for your dependency deletion. Click **Delete Dependency** to delete the dependency on all servers in the cluster.
4. If the [output panel](#) is enabled, the dialog closes, and the results of the commands to delete the dependency are shown in the output panel. If not, the dialog remains up to show these results, and you click Done to finish when all results have been displayed.

## 5.4.3.3.12. Deleting a Hierarchy from All Servers

---

1. There are five possible ways to begin.
  - Right-click on the icon for a resource in the hierarchy that you want to delete under the server where you want the deletion to begin. When the [resource context menu](#) appears, click **Delete Resource Hierarchy**.

 **Note:** Deleting a resource hierarchy before bringing the resource out-of-service on all nodes may prevent the use of some system resources. To mitigate this situation we suggest bringing the resource out-of-service on all nodes prior to deleting the resource hierarchy.

- Right-click on the icon for a global resource in the hierarchy that you want to delete. When the [resource context menu](#) appears, click **Delete Resource Hierarchy**. When the dialog comes up, select the server in the **Target Server** list from which you want to begin deleting the resource hierarchy, and click **Next**.
  - On the [global toolbar](#), click the **Delete Resource Hierarchy** button. When the dialog comes up, select the server in the **Target Server** list from which you want to begin deleting the resource hierarchy, and click **Next**. On the next dialog, select a resource in the hierarchy that you want to delete from the **Hierarchy to Delete** list, and click **Next** again.
  - On the [resource context toolbar](#) in the [properties panel](#), if displayed, click the **Delete Resource Hierarchy** button.
  - On the [Edit menu](#), point to **Resource** and then click **Delete Resource Hierarchy**. When the dialog comes up, select the server in the **Target Server** list from which you want to begin deleting the resource hierarchy, and click **Next**. On the next dialog, select a resource in the hierarchy that you want to delete from the **Hierarchy to Delete** list, and click **Next** again.
2. The dialog will display a message verifying the hierarchy you have specified for deletion. Click **Delete** to perform the action.
  3. If the [output panel](#) is enabled, the dialog closes, and the results of the commands to delete the hierarchy are shown in the output panel. If not, the dialog remains up to show these results, and you click **Done** to finish when all results have been displayed.

## 5.4.4. User Guide

The [User Guide](#) is a complete, searchable resource containing detailed information on the many tasks that can be performed within the LifeKeeper GUI. Click [User Guide](#) to access this documentation.

The tasks that can be performed through the GUI can be grouped into three areas:

[Common Tasks](#) – These are basic tasks that can be performed by any user such as connecting to a cluster, viewing server or resource properties, viewing log files and changing GUI settings.

[Operator Tasks](#) – These are more advanced tasks that require Operator permission, such as bringing resources in and out of service.

[Administrator Tasks](#) – These are tasks that require Administrator permission. They include server-level tasks such as editing server properties, creating resources, creating or deleting comm paths and resource-level tasks such as editing, extending, or deleting resources.

The table below lists the default tasks that are available for each user permission. Additional tasks may be available for specific resource types, and these will be described in the associated resource kit documentation.

Task	Permission		
	Guest	Operator	Administrator
View servers and resources	X	X	X
Connect to and disconnect from servers	X	X	X
View server properties and logs	X	X	X
Modify server properties			X
Create resource hierarchies			X
Create and delete comm paths			X
View resource properties	X	X	X
Modify resource properties			X
Take resources into and out of service		X	X
Extend and unextend resource hierarchies			X
Create and delete resource dependencies			X
Delete resource hierarchies			X

## 5.4.4.1. Using LifeKeeper for Linux

---

The following topics provide detailed information on the LifeKeeper graphical user interface (GUI) as well as the many tasks that can be performed within the LifeKeeper GUI.

---

[GUI](#)

[Status Table](#)

[Properties Panel](#)

[Output Panel](#)

[Message Bar](#)

[Exiting the GUI](#)

[Common Tasks](#)

[Operator Tasks](#)

[Advanced Tasks](#)

[Maintenance Tasks](#)

[Technical Notes](#)

## 5.4.4.1.1. GUI

---

The GUI components should have already been installed as part of the LifeKeeper Core installation.

---

### [GUI Overview – General](#)

#### [LifeKeeper GUI Software Package](#)

### [Menus](#)

#### [Resource Context Menu](#)

#### [Server Context Menu](#)

#### [File Menu](#)

#### [Edit Menu – Resource](#)

#### [Edit Menu – Server](#)

#### [View Menu](#)

#### [Help Menu](#)

### [Toolbars](#)

#### [GUI Toolbar](#)

#### [Resource Context Toolbar](#)

#### [Server Context Toolbar](#)

### [Preparing to Run the GUI](#)

#### [Overview](#)

#### [Configuration](#)

#### [Starting and Stopping the GUI Server](#)

#### [Java Security Policy](#)

#### [Running the GUI on a LifeKeeper Server](#)

## 5.4.4.1.1.1. GUI Overview – General

---

The GUI allows users working on any machine to administer, operate or monitor servers and resources in any cluster as long as they have the required group memberships on the cluster machines. (For details, see [Configuring GUI Users](#). The GUI Server and Client components are described below.

### GUI Server

The GUI server communicates with GUI clients using Hypertext Transfer Protocol (HTTP) and Remote Method Invocation (RMI). By default, the GUI server is initialized during LifeKeeper startup, but this can also be configured — see [Starting/Stopping the GUI Server](#).

### GUI Client

The GUI client can be run as an [application](#).

The client includes the following components:

- The [status table](#) on the upper left displays the high level status of connected servers and their resources.
- The [properties panel](#) on the upper right displays detailed information about the most recently selected status table object.
- The [output panel](#) on the bottom displays command output.
- The [message bar](#) at the very bottom of the window displays processing status messages.
- The context (in the properties panel) and [global toolbars](#) provide fast access to frequently used tasks.
- The context (popup) and [global menus](#) provide access to all tasks.

### Exiting GUI Clients

Select **Exit** from the [File Menu](#) to disconnect from all servers and close the client.

## 5.4.4.1.1.1. LifeKeeper GUI Software Package

---

The LifeKeeper GUI is included in the **steeleye-1kGUI** software package which is bundled with the LifeKeeper Core Package Cluster. The **steeleye-1kGUI** package:

- Installs the LifeKeeper GUI Client in Java archive format.
- Installs the LifeKeeper GUI Server.
- Installs the LifeKeeper administration web server.

**Note:** The LifeKeeper administration web server is configured to use Port 81, which should be different from any public web server.

- Installs a Java policy file in `/opt/LifeKeeper/htdocs/` which contains the minimum permissions required to run the LifeKeeper GUI. The LifeKeeper GUI application uses the `java.policy` file in this location for access control.
- Prepares LifeKeeper for GUI administration.

Before continuing, you should ensure that the LifeKeeper GUI package has been installed on the LifeKeeper server(s). You can enter the command `rpm -qi steeleye-1kGUI` to verify that this package is installed. You should see output including the package name **steeleye-1kGUI** if the GUI package is installed.

## 5.4.4.1.1.2. Menus

---

### LifeKeeper for Linux Menus

---

[Resource Context Menu](#)

[Server Context Menu](#)

[File Menu](#)

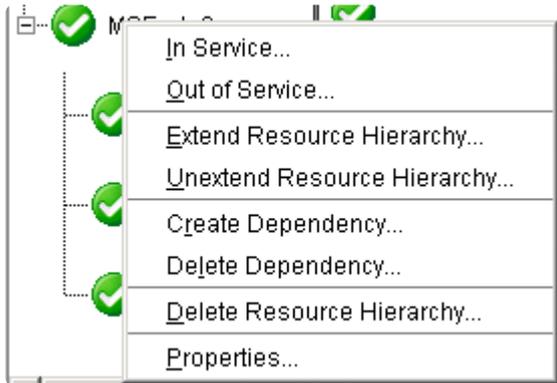
[Edit Menu – Resource](#)

[Edit Menu – Server](#)

[View Menu](#)

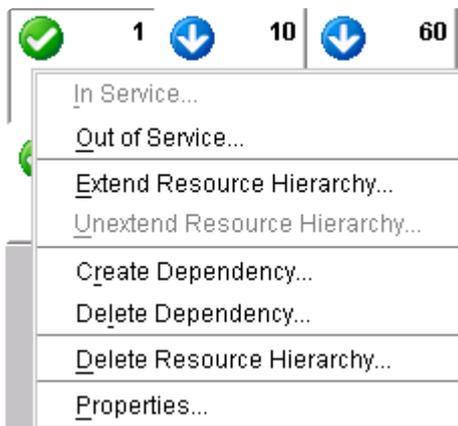
[Help Menu](#)

## 5.4.4.1.1.2.1. Resource Context Menu



The Resource Context Menu appears when you right-click on a global (cluster-wide) resource, as shown above, or a server-specific resource instance, as shown below, in the [status table](#). The default resource context menu is described here, but this menu might be customized for specific resource types, in which case the menu will be described in the appropriate resource kit documentation.

The actions are invoked for the resource that you select. If you select a resource instance on a specific server, the action is invoked for that server while if you select a global (cluster-wide) resource, you will need to select the server.



[In Service.](#) Bring a resource hierarchy into service.

[Out of Service.](#) Take a resource hierarchy out of service.

[Extend Resource Hierarchy.](#) Copy a resource hierarchy to another server for failover support.

[Unextend Resource Hierarchy.](#) Remove an extended resource hierarchy from a single server.

[Create Dependency.](#) Create a parent/child relationship between two resources.

[Delete Dependency.](#) Remove a parent/child relationship between two resources.

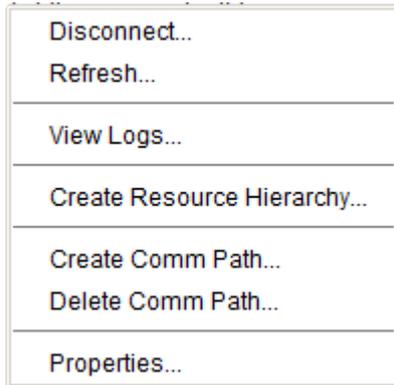
[Delete Resource Hierarchy.](#) Remove a resource hierarchy from all servers in the LifeKeeper cluster.

[Properties.](#) Display the [Resource Properties Dialog.](#)

## 5.4.4.1.1.2.2. Server Context Menu

---

The Server Context Menu appears when you right-click on a server icon in the [status table](#). This menu is the same as the Edit Menu's Server submenu except that the actions are always invoked on the server that you initially selected.



[Disconnect.](#) Disconnect from a cluster.

Refresh. Refresh GUI.

[View Logs.](#) View LifeKeeper log messages on connected servers.

[Create Resource Hierarchy.](#) Create a resource hierarchy.

[Create Comm Path.](#) Create a communication path between servers.

[Delete Comm Path.](#) Remove communication paths from a server.

[Properties.](#) Display the [Server Properties Dialog](#).

## 5.4.4.1.1.2.3. File Menu

---

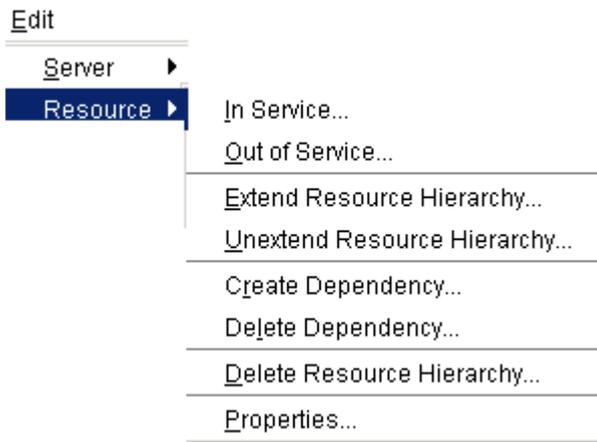


**Connect:** Connect to a LifeKeeper cluster. Connection to each server in the LifeKeeper cluster requires login authentication on that server.

**Exit:** Disconnect from all servers and close the GUI window.

## 5.4.4.1.1.2.4. Edit Menu – Resource

---



[In Service.](#) Bring a resource hierarchy into service.

[Out of Service.](#) Take a resource hierarchy out of service.

[Extend Resource Hierarchy.](#) Copy a resource hierarchy to another server for failover support.

[Unextend Resource Hierarchy.](#) Remove an extended resource hierarchy from a single server.

[Create Dependency.](#) Create a parent/child relationship between two resources.

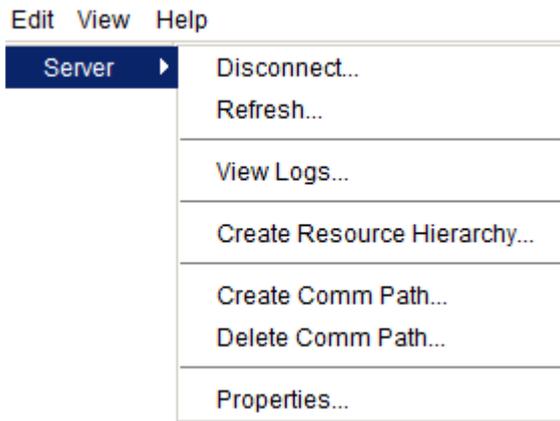
[Delete Dependency.](#) Remove a parent/child relationship between two resources.

[Delete Resource Hierarchy.](#) Remove a resource hierarchy from all servers in the LifeKeeper cluster.

[Properties.](#) Display the [Resource Properties Dialog](#).

## 5.4.4.1.1.2.5. Edit Menu – Server

---



[Disconnect.](#) Disconnect from a cluster.

Refresh. Refresh GUI.

[View Logs.](#) View LifeKeeper log messages on connected servers.

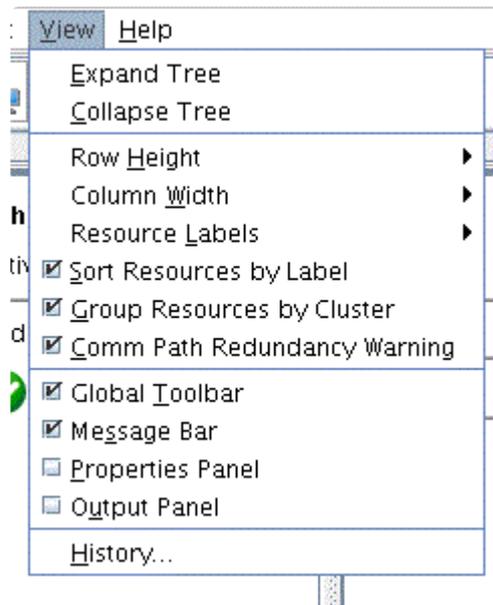
[Create Resource Hierarchy.](#) Create a resource hierarchy.

[Create Comm Path.](#) Create a communication path between servers.

[Delete Comm Path.](#) Remove communication paths from a server.

[Properties.](#) Display the [Server Properties Dialog](#).

## 5.4.4.1.1.2.6. View Menu



[Expand Tree](#). Expand the entire resource hierarchy tree.

[Collapse Tree](#). Collapse the entire resource hierarchy tree.

**Row Height**. Modify the row viewing size of the resources in the resource hierarchy tree and table. Select small, medium or large row height depending upon the number of resources displayed.

**Column Width**. Modify the column with viewing size of the resources in the resource hierarchy tree and table. Select fill available space, large, medium or small depending upon the resource displayed.

[Resource Labels](#). This option group allows you to specify whether resources are viewed in the resource hierarchy tree by their tag name or ID.

**Sort Resources by Label**. will sort resources by resource label only.

**Group Resources by Cluster**. will sort by server cluster and resource label such that resources belonging in the same cluster of servers will be grouped together.

**Comm Path Redundancy Warning**. specifies the representation of comm path status in the server status graphic.

- ◦ If selected, the display will show a server warning graphic if the comm paths between a set of servers are not configured with a redundant comm path.
- ◦ If not selected, the display will ignore a lack of redundant comm paths between a pair of servers but will still present server warning graphic if there are comm path failures.

[Global Toolbar](#). Display this component if the checkbox is selected.

[Message Bar](#). Display this component if the checkbox is selected.

[Properties Panel](#). Display this component if the checkbox is selected.

[Output Panel](#). Display this component if the checkbox is selected.

[History](#). Display the newest messages that have appeared in the Message Bar in the LifeKeeper GUI Message History dialog box (up to 1000 lines).

## 5.4.4.1.1.2.7. Help Menu

---

Help

Technical Documentation

About..

**Technical Documentation:** Displays the landing page of the SIOS Technology Corp. Technical Documentation.

**About:** Displays LifeKeeper GUI version information.

## 5.4.4.1.1.3. Toolbars

---

### LifeKeeper for Linux Toolbars

---

[GUI Toolbar](#)

[Resource Context Toolbar](#)

[Server Context Toolbar](#)

## 5.4.4.1.1.3.1. GUI Toolbar

This toolbar is a combination of the default [server](#) and [resource](#) context toolbars which are displayed on the [properties panel](#) except that you must select a server and possibly a resource when you invoke actions from this toolbar.



	<a href="#">Connect</a> . Connect to a LifeKeeper cluster.
	<a href="#">Disconnect</a> . Disconnect from a LifeKeeper cluster.
	<a href="#">Refresh</a> . Refresh GUI.
	<a href="#">View Logs</a> . View LifeKeeper log messages on connected servers.
	<a href="#">Create Resource Hierarchy</a> . Create a resource hierarchy.
	<a href="#">Delete Resource Hierarchy</a> . Remove a resource hierarchy from all servers in the LifeKeeper cluster.
	<a href="#">Create Comm Path</a> . Create a communication path between servers.
	<a href="#">Delete Comm Path</a> . Remove communication paths from a server.
	<a href="#">In Service</a> . Bring a resource hierarchy into service.
	<a href="#">Out of Service</a> . Take a resource hierarchy out of service.
	<a href="#">Extend Resource Hierarchy</a> . Copy a resource hierarchy to another server for failover support.
	<a href="#">Unextend Resource Hierarchy</a> . Remove an extended resource hierarchy from a single server.

	<a href="#">Create Dependency</a> . Create a parent/child relationship between two resources.
	<a href="#">Delete Dependency</a> . Remove a parent/child relationship between two resources.
	The Multi-Site feature has been discontinued.

## 5.4.4.1.1.3.2. Resource Context Toolbar

The resource context toolbar is displayed in the [properties panel](#) when you select a server-specific resource instance in the [status table](#).

The actions are invoked for the server and the resource that you select. Actions that are not available for selection for a resource will be grayed out.



	<p><a href="#">In Service</a>. Bring a resource hierarchy into service.</p>
	<p><a href="#">Out of Service</a>. Take a resource hierarchy out of service.</p>
	<p><a href="#">Extend Resource Hierarchy</a>. Copy a resource hierarchy to another server for failover support.</p>
	<p><a href="#">Unextend Resource Hierarchy</a>. Remove an extended resource hierarchy from a single server.</p>
	<p><a href="#">Add Dependency</a>. Create a parent/child relationship between two resources.</p>
	<p><a href="#">Remove Dependency</a>. Remove a parent/child relationship between two resources.</p>
	<p><a href="#">Delete Resource Hierarchy</a>. Remove a resource hierarchy from all servers.</p>

## 5.4.4.1.1.3.3. Server Context Toolbar

The server context toolbar is displayed in the [properties panel](#) when you select a server in the [status table](#). The actions are invoked for the server that you select.



	<a href="#">Disconnect</a> . Disconnect from a LifeKeeper cluster.
	Refresh. Refresh GUI.
	<a href="#">View Logs</a> . View LifeKeeper log messages on connected servers.
	<a href="#">Create Resource Hierarchy</a> . Create a resource hierarchy.
	<a href="#">Delete Resource Hierarchy</a> . Remove a resource hierarchy from all servers in the LifeKeeper cluster.
	<a href="#">Create Comm Path</a> . Create a communication path between servers.
	<a href="#">Delete Comm Path</a> . Remove communication paths from a server.

## 5.4.4.1.1.4. Preparing to Run the GUI

---

[Overview](#)

[Configuration](#)

[Starting and Stopping the GUI Server](#)

[Java Security Policy](#)

[Running the GUI on a LifeKeeper Server](#)

# 5.4.4.1.1.4.1. Configuring the LifeKeeper GUI

## Installing the LifeKeeper Server for GUI Administration

Perform the following steps for each LifeKeeper server. Each step contains references or links for more detailed instructions.

1. You must install the Java Runtime Environment (JRE) or Java Software Development Kit (JDK) on each server. See the [LifeKeeper for Linux Release Notes](#) for the required Java version.
2. Start the LifeKeeper GUI Server on each server (see [Starting/Stopping the GUI Server](#)). Note: Once the GUI Server has been started following an initial installation, starting and stopping LifeKeeper will start and stop all LifeKeeper daemon processes including the GUI Server.
3. If you plan to allow users other than root to use the GUI, then you need to [Configure GUI Users](#).

## Running the GUI

You can run the LifeKeeper GUI on the LifeKeeper server in the cluster.

See [Running the GUI on the LifeKeeper Server](#) for information on configuring and running the GUI on a server in your LifeKeeper cluster.

## GUI Configuration

Item	Description
<b>GUI Client and Server Communication</b>	The LifeKeeper GUI client and server use Java Remote Method Invocation (RMI) to communicate. For RMI to work correctly, the client and server must use resolvable hostnames or IP addresses. If DNS is not implemented (or names are not resolvable using other name lookup mechanisms), edit the /etc/hosts file on each client and server to include the names and addresses of all other LifeKeeper servers.
<b>GUI Server Java Platform</b>	The LifeKeeper GUI server requires that the Java Runtime Environment (JRE) – Java virtual machine, the Java platform core classes and supporting files – be installed. The LifeKeeper GUI supports OpenJDK. In the following environments, the setup script that is executed during installation installs OpenJDK that is included with the OS. If the Linux distributor does not provide OpenJDK, install the OpenJDK package included in the LifeKeeper installation image. See the <a href="#">Release Notes</a> for supported OpenJDK versions. <ul style="list-style-type: none"> <li>• RedHat Enterprise Linux/CentOS/Oracle Linux 7.1 or later</li> <li>• RedHat Enterprise Linux/CentOS/Oracle Linux 8 or later</li> <li>• SUSE Linux Enterprise Server 12 or later (excluding SLES15 and SLES15 SP1)</li> </ul>

	<p><b>Note:</b> When installing LifeKeeper, set the JRE path used by the GUI to LifeKeeper PATH default file <code>/etc/default/LifeKeeper</code>. Edit this PATH if you want to change the JRE version. If LifeKeeper is running when you edit this file, you should stop and restart the LifeKeeper GUI server to reflect the change.</p> <p><b>Note:</b> There is memory management inconsistency between the OpenJDKs that are included with LifeKeeper and the OS as for the combination of LifeKeeper and the OS version. In order to avoid that, OpenJDK that is included with the LifeKeeper installation image needs to be installed under the <code>/opt/LifeKeeper/lib64/java</code>. Memory management inconsistency has already been fixed but when LifeKeeper is upgraded, JRE that is installed on <code>/opt/LifeKeeper/lib64/java</code> will be used. If you want to use JRE that is included with the OS, install JRE and then follow the steps above to change the <code>/etc/default/LifeKeeper</code> PATH. <code>/opt/LifeKeeper/lib64/java</code> can be deleted since it is unnecessary.</p> <ul style="list-style-type: none"> <li>• LifeKeeper             <ul style="list-style-type: none"> <li>◦ v9.4.1 – v9.5.1</li> </ul> </li> <li>• OS             <ul style="list-style-type: none"> <li>◦ RedHat Enterprise Linux/CentOS/Oracle Linux 8.0 – 8.2</li> <li>◦ SUSE Linux Enterprise Server 15.0 – SP2</li> </ul> </li> </ul>
<p><b>Uninstall Java Runtime Environment</b></p>	<ul style="list-style-type: none"> <li>• Environment where the OpenJDK package included with the LifeKeeper installation image is installed: The OpenJDK package will be uninstalled when uninstalling LifeKeeper.</li> <li>• Environment where OpenJDK provided by Linux distributor is installed: The OpenJDK package is not uninstalled when uninstalling LifeKeeper. If necessary, uninstall it manually.</li> </ul>
<p><b>Java Remote Object Registry Server Port</b></p>	<p>The LifeKeeper GUI server uses port 82 for the Java remote object registry on each LifeKeeper server. This should allow servers to support RMI calls from clients behind typical firewalls.</p>
<p><b>LifeKeeper Administration Web Server</b></p>	<p>The LifeKeeper GUI server requires an administration web server for client browser communication. Currently, the LifeKeeper GUI server is using a private copy of the lighttpd web server for its administration web server. This web server is installed and configured by the steeleye-lighttpd package and uses port 81 to avoid a conflict with other web servers.</p>
<p><b>GUI Client Network Access</b></p>	<p>LifeKeeper GUI clients require network access to all hosts in the LifeKeeper cluster. When running the LifeKeeper GUI client in a browser, you will have to lower the security level to allow network access for applets. Be careful not to visit other sites with security set to low values (e.g., change the security settings only for intranet or trusted sites).</p>

## GUI Limitations

Item	Description
GUI	The LifeKeeper for Linux client may only be used to administer LifeKeeper on Linux

<b>Interoperability Restriction</b>	servers. The LifeKeeper for Linux GUI will not interoperate with LifeKeeper for Windows.
-----------------------------------------	------------------------------------------------------------------------------------------

## 5.4.4.1.1.4.2. Starting and Stopping the GUI Server

---

### To Start the LifeKeeper GUI Server

If the LifeKeeper GUI Server is not running, type the following command as root:

```
/opt/LifeKeeper/bin/lkGUIserver start
```

This command starts all LifeKeeper GUI Server daemon processes on the server being administered if they are not currently running. A message similar to the following is displayed.

```
# Installing GUI Log
# LK GUI Server Startup at:
# Mon May 8 14:14:46 EDT 2006
# LifeKeeper GUI Server Startup completed at:
# Mon May 8 14:14:46 EDT 2006
```

Once the LifeKeeper GUI Server is started, all subsequent starts of LifeKeeper will automatically start the LifeKeeper GUI Server processes.

### Troubleshooting

The LifeKeeper GUI uses Ports 81 and 82 on each server for its administration web server and Java remote object registry, respectively. If another application is using the same ports, the LifeKeeper GUI will not function properly. These values may be changed by editing the following entries in the LifeKeeper default file */etc/default/LifeKeeper*.

```
GUI_WEB_PORT=81 GUI_RMI_PORT=82
```

 **Note:** These port values are initialized in the GUI server at start time. If you alter them, you will need to stop and restart the *steeleye-lighttpd* process. These values must be the same across all clusters to which you connect.

### To Stop the LifeKeeper GUI Server

If the LifeKeeper GUI Server is running, type the following command as *root*:

```
/opt/LifeKeeper/bin/lkGUIserver stop
```

This command halts all LifeKeeper GUI Server daemon processes on the server being administered if they are currently running. The following messages are displayed.

```
# LifeKeeper GUI Server Shutdown at:
```

```
# Fri May 19 15:37:27 EDT 2006
# LifeKeeper GUI Server Shutdown Completed at:
# Fri May 19 15:37:28 EDT 2006
```

## LifeKeeper GUI Server Processes

To verify that the LifeKeeper GUI Server is running, type the following command:

```
ps -ef | grep runGuiSer
```

You should see output similar to the following:

```
root 2805 1 0 08:24 ? 00:00:00 sh/opt/LifeKeeper/bin/runGuiSer
```

To see a list of the other GUI Server daemon processes currently running, type the following command:

```
ps -ef | grep S_LK
```

You should see output similar to the following:

```
root 30228 30145 0 11:20 ? 00:00:00 java -Xint -Xss3M
-DS_LK=true -Djava.rmi.server.hostname=thor48 ...
```

## 5.4.4.1.1.4.3. Java Security Policy

---

The LifeKeeper GUI uses policy-based access control. When the GUI client is loaded, it is assigned permissions based on the security policy currently in effect. The policy, which specifies permissions that are available for code from various signers/locations, is initialized from an externally configurable policy file.

There is, by default, a single system-wide policy file and an optional user policy file. The system policy file, which is meant to grant system-wide code permissions, is loaded first, and then the user policy file is added to it. In addition to these policy files, the LifeKeeper GUI policy file may also be loaded if the LifeKeeper GUI is invoked as an application.

### Location of Policy Files

The system policy file is by default at:

*<JAVA.HOME>/lib/security/java.policy (Linux)*

*<JAVA.HOME>\lib\security\java.policy (Windows)*

**Note:** JAVA.HOME refers to the value of the system property named “JAVA.HOME”, which specifies the directory into which the JRE or JDK was installed.

The user policy file starts with `.` and is by default at:

*<USER.HOME>\.java.policy*

**Note:** USER.HOME refers to the value of the system property named “user.home”, which specifies the user’s home directory. For example, the home directory on a Windows NT workstation for a user named Paul might be “paul.000”.

For Windows systems, the user.home property value defaults to:

*C:\WINNT\Profiles\<>USER> (on multi-user Windows NT systems)*

*C:\WINDOWS\Profiles\<>USER> (on multi-user Windows 95/98 systems)*

*C:\WINDOWS (on single-user Windows 95/98 systems)*

The LifeKeeper GUI policy file is by default at:

*/opt/LifeKeeper/htdoc/java.policy (Linux)*

### Policy File Creation and Management

By default, the LifeKeeper GUI policy file is used when the LifeKeeper GUI is invoked as an application.

If you are running the LifeKeeper GUI as an applet, you will need to create a user policy file in your home directory if one does not already exist. The user policy file should specify the minimum permissions required to run the LifeKeeper GUI, which are provided in the “Sample Policy File” section later in this topic.

A policy file can be created and maintained via a simple text editor, or via the graphical Policy Tool utility included with the Java Runtime Environment (JRE) or Java Development Kit (JDK). Using the Policy Tool saves typing and eliminates the need for you to know the required syntax of policy files. For information about using the Policy Tool, see the Policy Tool documentation at <http://docs.oracle.com/javase/8/docs/technotes/tools/>.

The **simplest way to create a user policy file** with the minimum permissions required to run the LifeKeeper GUI is to copy the LifeKeeper GUI policy file located in `/opt/LifeKeeper/htdocs/java.policy` to your home directory and rename it `.java.policy` (note the leading dot before the filename which is required). On a Windows system, you can copy the LifeKeeper GUI policy file by opening the file `http://<server name>*.81/java.policy` (where `<server name>` is the host name of a LifeKeeper server) and saving it as `.java.policy` in your home directory. If you need to determine the correct location for a user policy file, enable the Java Console using the Java Control Panel and start the LifeKeeper GUI as an applet. The home directory path for the user policy file will be displayed in the Java Console.

## Granting Permissions in Policy Files

A permission represents access to a system resource. In order for a resource access to be allowed for an applet, the corresponding permission must be explicitly granted to the code attempting the access. A permission typically has a name (referred to as a “target name”) and, in some cases, a comma-separated list of one or more actions. For example, the following code creates a `FilePermission` object representing read access to the file named `abc` in the `/tmp` directory:

```
perm = new java.io.FilePermission("/tmp/abc", "read");
```

In this, the target name is `/tmp/abc` and the action string is `read`.

A policy file specifies what permissions are allowed for code from specified code sources. An example policy file entry granting code from the `/home/sysadmin` directory read access to the file `/tmp/abc` is:

```
grant codeBase "file:/home/sysadmin/" { permission java.io.FilePermission "/tmp/abc",
    "read"; };
```

## Sample Policy File

The following sample policy file includes the minimum permissions required to run the LifeKeeper GUI. This policy file is installed in `/opt/LifeKeeper/htdocs/java.policy` by the LifeKeeper GUI package.

```
/*
 * Permissions needed by the LifeKeeper GUI. You may want to
 * restrict this by codebase. However, if you do this, remember
 * that the recovery kits can have an arbitrary jar component ** with an
```

```
arbitrary codebase, so you'll need to alter the grant
* to cover these as well.
*/
grant {

/*
* Need to be able to do this to all machines in the
* LifeKeeper cluster. You may restrict the network
* specification accordingly.
*/
permission java.net.SocketPermission "*" , "accept,connect,resolve";
/*
* We use URLClassLoaders to get remote properties files and
* jar pieces.
*/
permission java.lang.RuntimePermission "createClassLoader";
/*
* The following are needed only for the GUI to run as an
* application (the default RMI security manager is more
* restrictive than the one a browser installs for its
* applets.
*/
permission java.util.PropertyPermission "*" , "read";
permission java.awt.AWTPermission "*" ;
permission java.io.FilePermission "<<ALL FILES>>" , "read,execute";

};
```

## 5.4.4.1.1.4.4. Running the GUI on a LifeKeeper Server

---

The simplest way to run the LifeKeeper GUI is as an application on a LifeKeeper server. By doing so you are, in effect, running the GUI client and server on the same system.

1. After configuring the LifeKeeper server for GUI Administration, you can run the GUI as an application on the server by entering the following command as root:

```
/opt/LifeKeeper/bin/lkGUIapp
```

2. The lkGUIapp script sets the appropriate environment variables and starts the application. As the application is loading, an application identity dialog or splash screen for LifeKeeper appears.
3. After the application is loaded, the LifeKeeper GUI appears and the Cluster Connect dialog is automatically displayed. Enter the Server Name you wish to connect to, followed by the login and password.
4. Once a connection to the cluster is established, the GUI window displays a visual representation and status of the resources protected by the connected servers. The GUI menus and toolbar buttons provide administration functions.

## 5.4.4.1.1.4.5. Lifekeeper GUI Overview

---

The LifeKeeper GUI uses Java technology to provide a graphical status interface to LifeKeeper and its configuration data. The LifeKeeper GUI allows users working on any machine to administer, operate, or monitor servers and resources in any cluster, as long as they have the required group memberships on the cluster machines. For details, see [[Configuring GUI Users](#)]. The LifeKeeper GUI Server are described below.

### GUI Server

The LifeKeeper GUI server is initialized on each server in a LifeKeeper cluster at system startup. It communicates with the LifeKeeper core software via the Java Native Interface (JNI), .

## 5.4.4.1.2. Status Table

---

The status table provides a visual representation of the status of connected servers and their resources. It shows:

- the state of each server in the top row
- the global (cross-server) state and the parent-child relationships of each resource in the left-most column
- the state of each resource on each server in the remaining cells

The states of the servers and resources are shown using graphics, text and color. An empty table cell under a server indicates that a particular resource has not been defined on that server.

If you select a server or a resource instance in the status table, detailed state information and a context-sensitive toolbar for that item are shown in the [properties panel](#). You can also pop up the appropriate [server context menu](#) or [resource context menu](#) for any item by right-clicking on that cell.

The status table is split into two sections. The relative sizes of the left and right sections can be modified by moving the divider between them. The status table can also be collapsed to show only the highest level items in the hierarchy trees. [Collapsing or expanding resource items](#) in the tree causes the hierarchies listed in the table to also expand and collapse.

## 5.4.4.1.3. Properties Panel

---

The properties panel displays the properties of the server or resource that is selected in the status table. The properties panel has the same functionality as the [server properties dialog](#) or the [resource properties dialog](#), plus a context-sensitive toolbar to provide fast access to commonly used commands. The caption at the top of this panel is **server\_name** if a server is selected, or **server\_name: resource\_name** if a resource is selected.

The context-sensitive toolbars displayed in the properties panel are the [server context toolbar](#) and the [resource context toolbar](#). Server or resource toolbars may also be customized. For more information on customized toolbars, see the corresponding [application recovery kit documentation](#).

The buttons at the bottom of the properties panel function as follows:

The **Apply** button applies any changes that have been made to editable properties on the panel. This button is only enabled if you have changed an editable property.

The **Reset** button queries the server for the current values of all properties, clearing any changes that you may have made. This button is always enabled.

The **Help** button displays context-sensitive help for the properties panel. This button is always enabled.

You increase or decrease the size of the properties panel by sliding the separator at the left of the panel to the left or right. If you want to open or close this panel, use the **Properties Panel checkbox** on the [View Menu](#).

## 5.4.4.1.4. Output Panel

---

The output panel collects output from commands issued by the LifeKeeper GUI client. When a command begins to run, a time stamped label is added to the output panel, and all of the output from that command is added under this label. If you are running multiple commands at the same time (typically on different servers), the output from each command is sent to the corresponding section making it easy to see the results of each.

You increase or decrease the size of the output panel by sliding the separator at the top of the panel up or down. If you want to open or close this panel, use the **Output Panel checkbox** on the [View Menu](#). When the output panel is closed, the dialog that initiates each command will stay up, the output will be displayed on that dialog until you dismiss it and you will not be able to review the output from any command after you have closed that dialog. After the output panel is reopened, the LifeKeeper GUI will return to its default behavior.

## 5.4.4.1.5. Message Bar

---

The message bar appears beneath the status window. It is used for displaying messages in a single text line. Message such as “Connecting to Server X” or “Failure to connect to Server X” might be displayed.

To hide the message bar, clear the **Message Bar** checkbox in the [View Menu](#).

To display the message bar, select the **Message Bar** checkbox in the View Menu.

To see a history of messages displayed in the message bar, see [Viewing Message History](#).

## 5.4.4.1.6. Exiting the GUI

---

Select **Exit** from the [File Menu](#) to disconnect from all servers and close the GUI window.

## 5.4.4.1.7. Common Tasks

---

The following are basic tasks that can be performed by any user.

---

[Starting LifeKeeper](#)

[Stopping LifeKeeper](#)

[Viewing LifeKeeper Processes](#)

[Viewing LifeKeeper GUI Server Processes](#)

[Viewing LifeKeeper Controlling Processes](#)

[Connecting Servers to a Cluster](#)

[Disconnecting from a Cluster](#)

[Viewing Connected Servers](#)

[Viewing the Status of a Server](#)

[Viewing Server Properties](#)

[Viewing Server Log Files](#)

[Viewing Resource Tags and IDs](#)

[Viewing the Status of Resources](#)

[Viewing Resource Properties](#)

[Resource Labels](#)

[Viewing Message History](#)

[Expanding and Collapsing a Resource Hierarchy Tree](#)

[Cluster Connect Dialog](#)

[Cluster Disconnect Dialog](#)

[Resource Properties Dialog](#)

[Server Properties Dialog](#)

## 5.4.4.1.7.1. Starting LifeKeeper

All LifeKeeper software is installed in the directory `/opt/LifeKeeper`.

When you have completed all of the [verification tasks](#), you are ready to start LifeKeeper on both servers. This section provides information for starting the LifeKeeper server daemon processes. The LifeKeeper GUI application is launched using a separate command and is described in [Configuring the LifeKeeper GUI](#). LifeKeeper provides a [command line interface](#) that starts and stops the LifeKeeper daemon processes. These daemon processes must be running before you start the LifeKeeper GUI.

### Starting LifeKeeper Server Processes

! When LifeKeeper starts and establishes communication with the other servers in a cluster, it **will not** allow data replication (*DataKeeper*) resources to come in-service (*manual or automatic*) until communication is established with all servers in the cluster. The data replication resources will be marked OSF if an attempt is made to *in-service* them before communication is established with all servers. (*Please refer to the system's log files for additional information.*)

If LifeKeeper is not currently running on your system, type the following command as the user root on all servers:

```
/opt/LifeKeeper/bin/lkstart
```

When executing this command, the LifeKeeper service will be started and LifeKeeper will be set to start automatically at system startup.

Following the delay of a few seconds, an informational message is displayed.

\* **Note:** If you receive an error message referencing the **LifeKeeper Distribution Enabling Package** when you start LifeKeeper, you should install / re-install the [LifeKeeper Installation Image File](#).

See the `LCDD` help page by entering `man LCD` at the command line for details on the `lkstart` command.

To start only the LifeKeeper process without enabling automatic startup, execute the following command:

```
service lifekeeper start (or, systemctl start lifekeeper)
```

### Enabling Automatic LifeKeeper Restart

While the above command will start LifeKeeper, it will need to be performed each time the system is rebooted. If you would like LifeKeeper to start automatically when server boots up, type the following command:

```
chkconfig lifekeeper on (or, systemctl enable lifekeeper)
```

See the `chkconfig` man page for further information.

## 5.4.4.1.7.2. Stopping LifeKeeper

If you need to stop LifeKeeper, type the following command as root to stop it:

 **Note:** The `lkstop` process will continue even when a resource remove fails and will not result in a failed `lkstop`.

- `/opt/LifeKeeper/bin/lkstop`

This command will shut down LifeKeeper on the local system if it is currently running. It will first remove all protected resources from service on the local system then shut down the LifeKeeper daemons.

- `/opt/LifeKeeper/bin/lkstop -f`

This command will skip the section that removes resources from service. The resources will remain running on the local system but will no longer be protected by LifeKeeper.

 See [Commands](#) for additional options and best practices in usage.

 When '`lkstop -f`' is used to stop LifeKeeper, in-service resources are left configured and running, including data replication resources. **It is important that the hierarchy/resources are allowed to be brought back in-service on the correct server after an '`lkstop -f`'.** LifeKeeper will automatically bring all resources in-service once communication is established with all servers. Data replication resources may temporarily be marked OSF while LifeKeeper is waiting for all servers to restart and rejoin the cluster. *(Please refer to the system's **log files** for additional information pertaining to that node.)*

### Disabling Automatic LifeKeeper Restart

If you do not want LifeKeeper to automatically restart when the system is restarted, type the following command:

```
chkconfig lifekeeper off
```

*or*

```
systemctl disable lifekeeper
```

See the `chkconfig` (or `systemctl`) man page for further information.

## 5.4.4.1.7.3. Viewing LifeKeeper Processes

To see a list of all LifeKeeper core daemon processes currently running, type the following command:

```
ps -ef | grep LifeKeeper | grep -w bin | grep -v lklogmsg
```

An example of the output is provided below:

```
root 11663 11662 0 14:03 pts/0 00:00:00 /bin/bash /etc/redhat-lsb/
lsb_start_daemon /opt/LifeKeeper/sbin/runsvdir -P /opt/LifeKeeper/etc/
service log: runit just
started.....

root 11666 11663 0 14:03 pts/0 00:00:00 /bin/bash -c ulimit -S -c 0
>/dev/null 2> &1 ; /opt/LifeKeeper/sbin/runsvdir -P /opt/LifeKeeper/etc/
service log: runit just
started.....

root 11880 11873 0 14:03 ? 00:00:00 /opt/LifeKeeper/bin/lk_logmgr -l/opt/
LifeKeeper/out -d/etc/default/LifeKeeper

root 12240 11877 0 14:04 ? 00:00:00 /opt/LifeKeeper/bin/lcm

root 12247 11879 0 14:04 ? 00:00:00 /opt/LifeKeeper/bin/ttymonlcm

root 12250 11876 0 14:04 ? 00:00:00 /opt/LifeKeeper/bin/lcd

root 12307 11874 0 14:04 ? 00:00:00 /opt/LifeKeeper/bin/lkcheck

root 12311 11875 0 14:04 ? 00:00:00 /opt/LifeKeeper/bin/lkscsid

root 12325 11871 0 14:04 ? 00:00:00 /opt/LifeKeeper/bin/lkvmhad

root 12335 12330 0 14:04 ? 00:00:00 /opt/LifeKeeper/bin/perl /opt/
LifeKeeper/htdoc/cgi-bin/DoRequest.fcgi
```

The run state of LifeKeeper can be determined via the following command:

```
/opt/LifeKeeper/bin/lktest
```

If LifeKeeper is running it will output something similar to the following:

```
F S UID PID PPID C CLS PRI NI SZ STIME TIME CMD

4 S root 12240 11877 0 TS 39 -20 6209 14:04 00:00:00 lcm
```

```
4 S root 12247 11879 0 TS 39 -20 30643 14:04 00:00:00 ttymonlcm
```

```
4 S root 12250 11876 0 TS 29 -10 9575 14:04 00:00:00 lcd
```

If LifeKeeper is not running, then nothing is output and the command exists with a 1.

 **Note:** There are additional LifeKeeper processes running that start, stop, and monitor the LifeKeeper core daemon processed along with those required for the Graphical User Interface (GUI). See [Viewing LifeKeeper Controlling Processes](#) and [Viewing LifeKeeper GUI Server Processes](#) for a list of the processes. Additionally, most LifeKeeper processes have a child lklogmsg to capture and log any unexpected output.

## 5.4.4.1.7.4. Viewing LifeKeeper GUI Server Processes

---

To verify that the LifeKeeper GUI Server is running, type the following command:

```
ps -ef | grep runGuiSer
```

You should see output similar to the following:

```
root 2805 1 0 08:24 ? 00:00:00 sh /opt/LifeKeeper/bin/runGuiServer
```

To see a list of the other GUI Server daemon processes currently running, type the following command:

```
ps -efw | grep S_LK
```

You should see output similar to the following:

```
root 819 764 0 Oct16 ? 00:00:00 java -Xint -Xss3M -DS_LK=true  
-Djava.rmi.server.hostname=wake -Dcom.steeleye.LifeKeeper.rmiPort=82  
-Dcom.steeleye.LifeKeeper.LKROOT=/opt/LifeKeeper  
-DGUI_RMI_REGISTRY=internal -DGUI_WEB_PORT=81  
com.steeleye.LifeKeeper.beans.S_LK
```

To verify that the LifeKeeper GUI Server Administration Web Server is running type the following command:

```
ps -ef|grep steeleye-light | egrep -v "lklogmsg|runsv"
```

You should see output similar to the following:

```
root 12330 11872 0 14:04 ? 00:00:00 /opt/LifeKeeper/sbin/steeleye-  
lighttpd -D -f/opt/LifeKeeper/etc/lighttpd/lighttpd.conf
```

## 5.4.4.1.7.5. Viewing LifeKeeper Controlling Processes

---

To verify that the LifeKeeper controlling processes are running, type the following command:

```
ps -ef | grep runsv
```

You should see output similar to the following:

```
root 29093 1 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsvdir -P /opt/
LifeKeeper/etc/service

root 29097 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv lcd

root 29098 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv lcm

root 29099 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv lk_logmgr

root 29100 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv lkcheck

root 29101 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv lkscsid

root 29102 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv lkvmhad

root 29103 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv steeleye-
lighttpd

root 29104 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv ttymonlcm

root 29105 29093 0 11:35 ? 00:00:00 /opt/LifeKeeper/sbin/runsv
lkguiserver

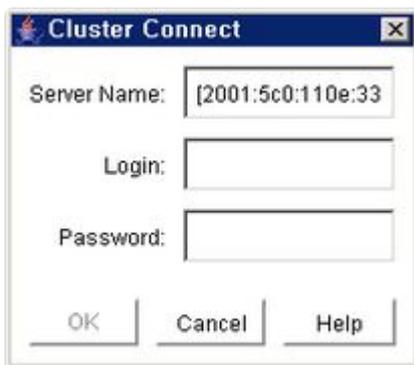
root 29465 2894 0 11:36 pts/0 00:00:00 grep --color=auto runsv
```

These processes start, stop, and monitor LifeKeeper core daemon processes and must be running to start LifeKeeper. These processes are configured by default to start when the system boots and this behavior should not be altered.

## 5.4.4.1.7.6. Connecting Servers to a Cluster

1. There are two possible ways to begin.
  - On the [Global Toolbar](#), click the **Connect** button.
  - On the [File Menu](#), click **Connect**.
2. In the **Server Name** field of the [Cluster Connect dialog](#), enter the name of a server within the cluster to which you want to connect.

\* **Note:** If using an **IPv6** address, this address will need to be enclosed in brackets [ ]. This will allow a connection to be established through a machine's IPv6 address. Alternatively, a name can be assigned to the address, and that name can then be used to connect.



3. In the **Login** and **Password** fields, enter the login name and password of a user with LifeKeeper authorization on the specified server.
4. Click **OK**.

If the GUI successfully connects to the specified server, it will continue to connect to (and add to the status display) all known servers in the cluster until no new servers are found.

\* **Note:** If the initial login name and password fails to authenticate the client on a server in the cluster, the user is prompted to enter another login name and password for that server. If “**Cancel**” is selected from the [Password dialog](#), connection to that server is aborted and the GUI continues connecting to the rest of the cluster.

## 5.4.4.1.7.7. Disconnecting from a Cluster

---

This task disconnects your GUI client from all servers in the cluster, and it does so through the server you select.

1. There are three possible ways to begin.
  - On the [Global Toolbar](#), click the **Disconnect** button.
  - On the [Edit Menu](#), select **Server** and then click **Disconnect**.
  - On the [Server Context Toolbar](#), if displayed, click the **Disconnect** button.
2. In the **Select Server in Cluster** list of the [Cluster Disconnect Dialog](#), select the name of a server in the cluster from which you want to disconnect.
3. Click **OK**. A **Confirmation** dialog listing all the servers in the cluster is displayed.
4. Click **OK** in the **Confirmation** dialog to confirm that you want to disconnect from all servers in the cluster.

After disconnecting from the cluster, all servers in that cluster are removed from the GUI status display.

## 5.4.4.1.7.8. Viewing Connected Servers

---

The state of a server can be determined by looking at the graphic representation of the server in the table's header as shown below. See [Viewing the Status of a Server](#) for an explanation of the server states indicated visually by the server icon.

---

					
wallace	gromit	pat	mike	batman	bullwinkle

---

## 5.4.4.1.7.9. Viewing the Status of a Server

The state of a server can be determined by looking at the graphic representation of the server in the table's header as shown below.



Server State	Visual state	What it Means
ALIVE		<p>Client has valid connection to the server.</p> <p>Comm paths originating from this server to an ALIVE remote server are ALIVE.</p> <p>Comm paths which may be marked DEAD and which target a DEAD server are ignored because the DEAD server will be reflected in its own graphic.</p>
ALIVE		<p>Client has valid connection to the server.</p> <p>One or more comm paths from this server to a given remote server are marked as DEAD.</p> <p>No redundant comm path exists from this server to a given remote server.</p>
DEAD		Reported as DEAD by other servers in the cluster.
UNKNOWN		Network connection was lost. Last known LifeKeeper state is ALIVE.

## 5.4.4.1.7.10. Viewing Server Properties

---

1. There are two possible ways to begin.
  - Right-click on the icon for the server for which you want to view the properties. When the [Server Context Menu](#) appears, click **Properties**. Server properties will also be displayed in the [Properties Panel](#) if it is enabled when clicking on the server.
  - On the [Edit Menu](#), point to **Server** and then click **Properties**. When the dialog comes up, select the server for which you want to view the properties from the Server list.
2. If you want to view properties for a different server, select that server from the dialog's **Server** list.
3. When you are finished, click **OK** to close the window.

## 5.4.4.1.7.11. Viewing Server Log Files

---

1. There are four ways to begin.
  - Right-click on a server icon to display the [Server Context Menu](#), then click **View Log** to bring up the LifeKeeper Log Viewer Dialog.
  - On the [Global Toolbar](#), click the **View Log** button, then select the server that you want to view from the Server list in the LifeKeeper Log Viewer Dialog.
  - On the [Server Context Toolbar](#), if displayed, click the **View Log** button.
  - On the [Edit Menu](#), point to **Server**, click **View Log**, then select the server that you want to view from the Server list in the **LifeKeeper Log Viewer Dialog**.
2. If you started from the **Global Toolbar** or the **Edit Menu** and you want to view logs for a different server, select that server from the **Server** list in the LifeKeeper Log Viewer Dialog. This feature is not available if you selected **View Logs** from the **Server Context Menu** or **Server Context Toolbar**.
3. When you are finished, click **OK** to close the **Log Viewer** dialog.

## 5.4.4.1.7.12. Viewing Resource Tags and IDs

A resource's tag and ID can be viewed quickly by positioning the cursor over a resource icon in the status window and clicking the left mouse button once (single-click). The resource tag and ID of the server having the lowest priority number are displayed in the message bar. To display the resource tag and ID for a resource on a specific server, single-click the appropriate resource instance cell in the table.

Messages displayed in the message bar look similar to the following:

```
Resource Tag = ipdnet0-153.98.87.73, Resource ID = IP-153.98.87.73
```

Under certain circumstances, the GUI may not be able to determine the resource ID, in which case only the resource tag is displayed in the message bar.

## 5.4.4.1.7.13. Viewing the Status of Resources

The status or state of a resource is displayed in two formats: **Global Resource Status** (across all servers), and the **Server Resource Status** (on a single server). The global resource status is shown in the **Resource Hierarchy Tree** in the left pane of the status window. The server resource status is found in the table cell where the resource row intersects with the server column.

### Server Resource Status

The following figure shows servers with resource statuses of active, standby and unknown.

- All resources on “wallace” are active
- All resources on “gromit”, “pat”, “mike” and “batman” are standby
- All resources on “bullwinkle” are unknown

	 wallace	 gromit	 pat	 mike	 batman	 bullwinkle
 1	 10	 20	 30	 40	 50	
Active	StandBy	StandBy	StandBy	StandBy	Unknown	
 1	 10	 20	 30	 40	 50	
Active	StandBy	StandBy	StandBy	StandBy	Unknown	

Server Resource State	Visual State	What it Means
ALIVE		Resource is operational on this server and protected. (ISP)
Degraded		Resource is operational on this server, but not protected by a backup resource. (ISU)
Standby		Server can take over operation of the resource. (OSU)
Failed		Problem with resource detected on this server. For example, an attempt to bring the resource in-service failed. (OSF)
Unknown		Resource is operational on this server, but not protected by a backup resource. (ISU)
	Empty	Server does not have the resource defined.

	panel	
--	-------	--

## Global Resource Status



Visual State	Description	What it Means / Causes
ALIVE 	Normal	Resource is active (ISP) and all backups are active.
	Warning	Resource is active (ISP). One or more backups are marked as unknown or failed (OSF)
	Failed. Resource is not active on any servers (OSF).	Resource has been taken out-of-service for normal reasons.  Resource has stopped running by unconventional means.  Recovery has not been completed or has failed.
	Unknown. Could not determine state from available information.	More than one server is claiming to be active.  Lost connection to server.  All server resource instances are in an unknown state.

## 5.4.4.1.7.14. Viewing Resource Properties

---

1. There are three possible ways to begin.
  - Right-click on the icon for the resource/server combination for which you want to view the properties. When the [Resource Context Menu](#) appears, click **Properties**. Resource properties will also be displayed in the [Properties Panel](#) if it is enabled.
  - Right-click on the icon for the global resource for which you want to view the properties. When the [Resource Context Menu](#) appears, click **Properties**. When the dialog comes up, select the server for which you want to view that resource from the **Server** list.
  - On the [Edit Menu](#), point to **Resource** and then click **Properties**. When the dialog comes up, select the resource for which you want to view properties from the **Resource** list, and the server for which you want to view that resource from the **Server** list.
2. If you want to view properties for a different resource, select that resource from the **Resource** list.
3. If you want to view resource properties for a different server, select that server from the **Server** list.
4. When you are finished, click **OK** to close the window.

## 5.4.4.1.7.15. Resource Labels

---

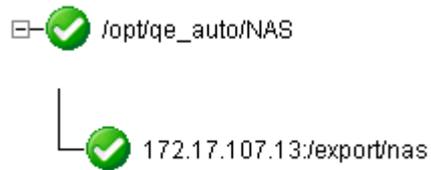
This option group allows you to specify whether resources are viewed in the resource hierarchy tree by their tag name or ID.

 **Note:** The resource tag/ID shown in the resource hierarchy tree belongs to the server having the lowest priority number. If you wish to see the tag/ID for a resource on a specific server, left-click the resource instance cell in the table and its tag/ID will be displayed in the message bar.

### By tag name:



### By ID:



## 5.4.4.1.7.16. Viewing Message History

---

1. On the [View Menu](#), click **History**. The LifeKeeper GUI Message **History** dialog is displayed.
2. If you want to clear all messages from the history, click **Clear**.
3. Click **OK** to close the dialog.

The **Message History** dialog displays the most recent messages from the message bar. The history list can display a maximum of 1000 lines. When the maximum number of lines is exceeded, the new messages will “push out” the oldest messages.

These messages represent only the actions between the client and the server and are displayed in chronological order, the most recent messages appearing at the top of the list.

### Reading the Message History

<- indicates that the message is incoming from a server and typically has a format of:

```
<-“server name”：“action”
```

```
<-“server name”：“app res”：“action”
```

```
<-“server name”：“res instance”：“action”
```

-> indicates that the message is outgoing from a client and typically has a format of:

```
->“server name”：“action”
```

```
->“server name”：“app res”：“action”
```

```
->“server name”：“res instance”：“action”
```

The **Clear** button clears the history but does not close the dialog.

The **OK** button closes the dialog without clearing the history.

## 5.4.4.1.7.17. Expanding and Collapsing a Resource Hierarchy Tree

 <p>The screenshot shows a tree structure with three nodes, each with a green checkmark icon to its right. The root node is 'file_system_2' with a square icon containing a minus sign to its left. It has two children: 'device-nfs15957' and 'nfs-/opt/qe_auto/NFS/export1'. The 'device-nfs15957' node has a square icon containing a minus sign to its left. The 'nfs-/opt/qe_auto/NFS/export1' node has a square icon containing a plus sign to its left.</p>	<p>In this segment of the tree, the resource <i>file_system_2</i> is expanded and the resource <i>nfs-/opt/qe_auto/NFS/export1</i> is collapsed.</p> <p> appears to the left of a resource icon if it is expanded.</p> <p> appears if it is collapsed.</p>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

To **expand** a resource hierarchy tree,

- Click the  or
- Double-click the resource icon to the right of a .

To **expand all** resource hierarchy trees,

- On the **View Menu**, click **Expand Tree** or
- Double-click the **Resource Hierarchy Tree** button in the column header in the left pane of the **Status** window.

 **Note:** The resource tag/ID shown in the resource hierarchy tree belongs to the server having the lowest priority number. If you wish to see the tag/ID for a resource on a specific server, left-click the resource instance cell in the table and its tag/ID will be displayed in the message bar.

To **collapse** a resource hierarchy tree,

- click the  or
- double-click the resource icon to the right of a .

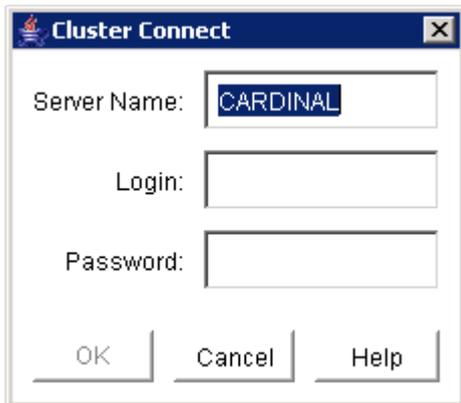
To **collapse all** resource hierarchy trees,

- On the **View Menu**, click **Collapse Tree** or
- Double-click the **Resource Hierarchy Tree** button in the column header in the left pane of the **Status** window

 **Note:** The “9” and “0” keys are defined as hot/accelerator keys to facilitate quickly expanding or collapsing all resource hierarchy trees.

## 5.4.4.1.7.18. Cluster Connect Dialog

---



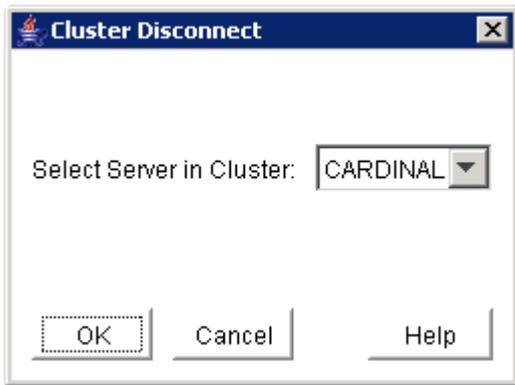
**Server Name.** The name of the server to which you want to connect.

**Login.** The login name of a user with LifeKeeper authorization on the server to which you want to connect.

**Password.** The password that authorizes the specified login on the server to which you want to connect.

## 5.4.4.1.7.19. Cluster Disconnect Dialog

---



### Select Server in Cluster

A drop-down list box containing the names of connected servers will appear. From the list, select a server from the cluster from which you want to disconnect. All servers in the cluster to be disconnected are noted in the confirmation dialog.

## 5.4.4.1.7.20. Resource Properties Dialog

The Resource Properties dialog is available from the [Edit menu](#) or from a [resource context menu](#). This dialog displays the properties for a particular resource on a server. When accessed from the Edit menu, you can select the resource and the server. When accessed from a resource context menu, you can select the server.

### General Tab

- **Tag.** The name of a resource instance, unique to a system, that identifies the resource to an administrator.
- **ID.** A character string associated with a resource instance, unique among all instances of the resource type, that identifies some internal characteristics of the resource instance to the application software associated with it.
- **Switchback.** (editable if user has Administrator permission) The setting that governs the recovery behavior of the server where the resource was in service when it failed. If the setting is intelligent, the server acts as a possible backup for the given resource. If the setting is automatic, the server actively attempts to re-acquire the resource, providing the following conditions are met:
  - The resource hierarchy must have been in service on the server when it left the cluster.
  - If it is in service at all, then the resource must currently be in service on a server with a lower priority.

 **Note:** Checks for automatic switchback are made only when LifeKeeper starts or when a new server is added to the cluster; they are not performed during normal cluster operation.

- **State.** Current state of the resource instance:
  - *Active* – In-service locally and protected.
  - *Warning* – In-service locally, but local recovery will not be attempted.
  - *Failed* – Out-of-service, failed.
  - *Standby* – Out-of-service, unimpaired.
  - *ILLSTATE* – A resource state has not been initialized properly by the resource initialization process which is run as part of the LifeKeeper startup sequence. Resources in this state are not under LifeKeeper protection.
  - *UNKNOWN* – Resource state could not be determined. The GUI server may not be available.
- **Reason.** If present, describes the reason the resource is in its current state, that is, the reason for the last state change. For example the application on galahad is in the OSU state because the shared primary resource ordbfsaa-on-tristan on tristan is in ISP or ISU state. Shared resources can be active on only one of the grouped systems at a time.
- **Initialization.** The setting that determines resource initialization behavior at boot time, for example, AUTORES\_ISP, INIT\_ISP, or INIT\_OSU.

## Relations Tab

- **Parent.** Identifies the tag names of the resources that are directly dependent on this resource.
- **Child.** Identifies the tag names of all resources on which this resource depends.
- **Root.** Tag name of the resource in this resource hierarchy that has no parent.

## Equivalencies Tab

- **Server.** The name of the server on which the resource has a defined equivalency.
- **Priority.** (editable if the user has Administrator permission). The failover priority value of the targeted server, for this resource.
- **Tag.** The tag name of this resource on the equivalent server.
- **Type.** The type of equivalency (SHARED, COMMON, COMPOSITE).
- **Reorder Priorities.** (available if the user has Administrator permission) Up/Down buttons let you to re-order the priority of the selected equivalency.

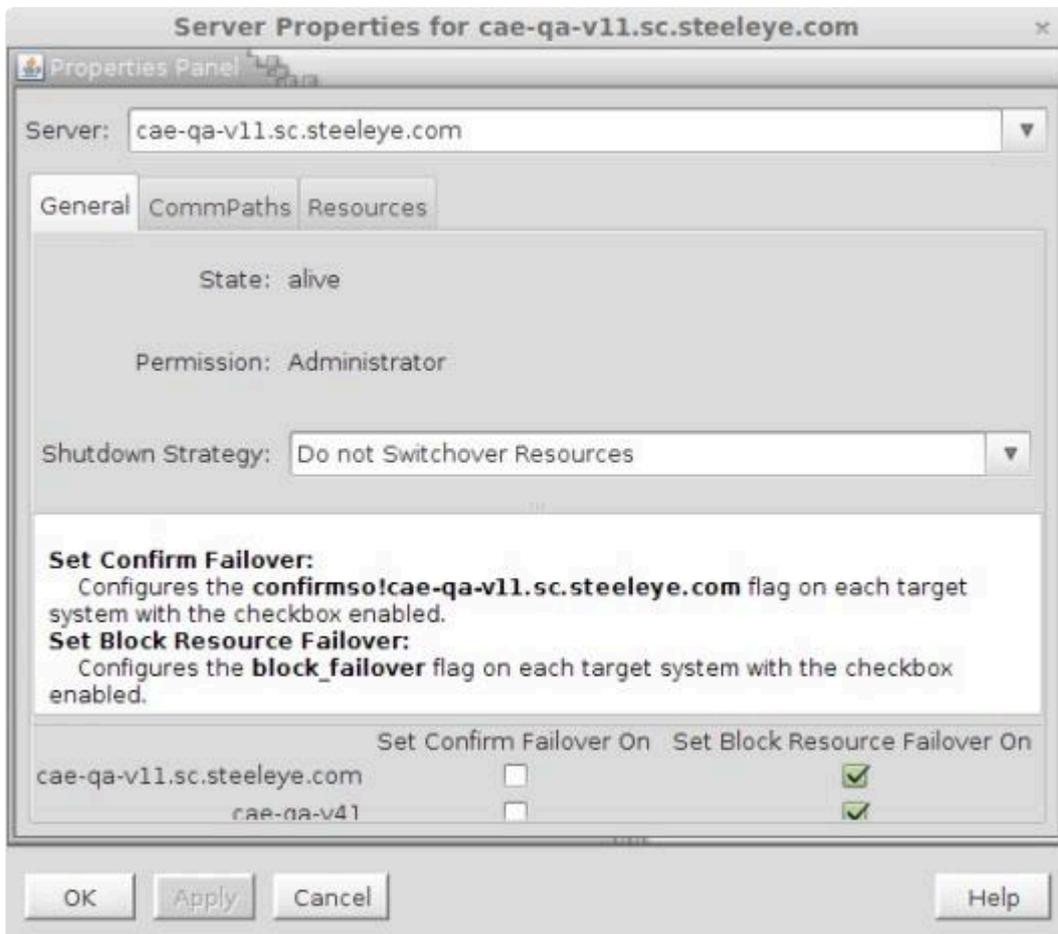
The OK button applies any changes that have been made and then closes the window. The Apply button applies any changes that have been made. The Cancel button, closes the window without saving any changes made since Apply was last clicked.

## 5.4.4.1.7.21. Server Properties Dialog

The Server Properties dialog is available from a server context menu or from the [Edit menu](#). This dialog displays the properties for a particular server. The properties for the server will also be displayed in the [properties panel](#) if it is enabled.

The three tabs of this dialog are described below. The OK button applies any changes that have been made and then closes the window. The Apply button applies any changes that have been made. The Cancel button closes the window without saving any changes made since Apply was last clicked.

### General Tab



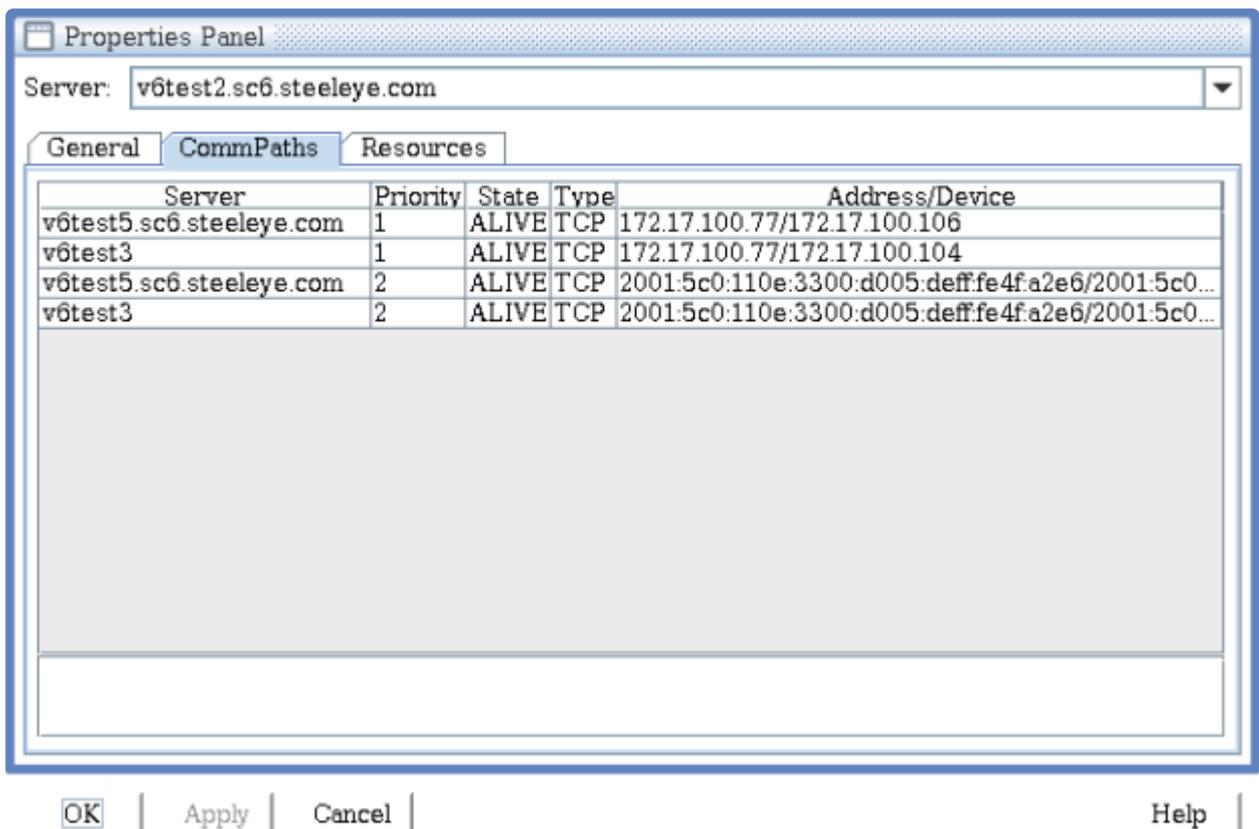
- **Name.** Name of the selected server.
- **State.** Current state of the server. These are the possible server state values:
  - *ALIVE* – server is available.
  - *DEAD* – server is unavailable.
  - *UNKNOWN* – state could not be determined. The GUI server may not be available.
- **Permission.** The permission level of the user currently logged into that server. These are the possible permission values:
  - *Administrator* – the user can perform any LifeKeeper task.
  - *Operator* – the user can monitor LifeKeeper resource and server status, and can bring

resources in service and take them out of service.

- *Guest* – the user can monitor LifeKeeper resource and server status.

- **Shutdown Strategy.** (editable if the user has Administrator permission) The setting that governs whether or not resources are switched over to a backup server in the cluster when a server is shutdown. The setting “*Switchover Resources*” indicates that resources will be brought in service on a backup server in the cluster. The setting “*Do not Switchover Resources*” indicates that resources will not be brought in service on another server in the cluster.
- **Failover Strategy.** The setting allows you to require the confirmation of failovers from specific systems in the LifeKeeper cluster. It is only available to LifeKeeper administrators. Operators and guests will not be able to see it. By default, all failovers proceed automatically with no user intervention. However, once the confirm failover flag is set, failovers from the designated system will require confirmation by executing the command: `lk_confirmso -y system`. The failover may be blocked by executing the command: `lk_confirmso -n system`. The system will take a pre-programmed default action unless one of these commands is executed within a specified interval. Two flags in the `/etc/default/LifeKeeper` file govern this automatic action.
  - CONFIRMSODEF (This specifies the default action. If set to “0”, the default action is to proceed with failover. If set to “1”, the default action is to block failover.)
  - CONFIRMSOTO (This is set to the time in seconds that LifeKeeper should wait before taking the default action.)

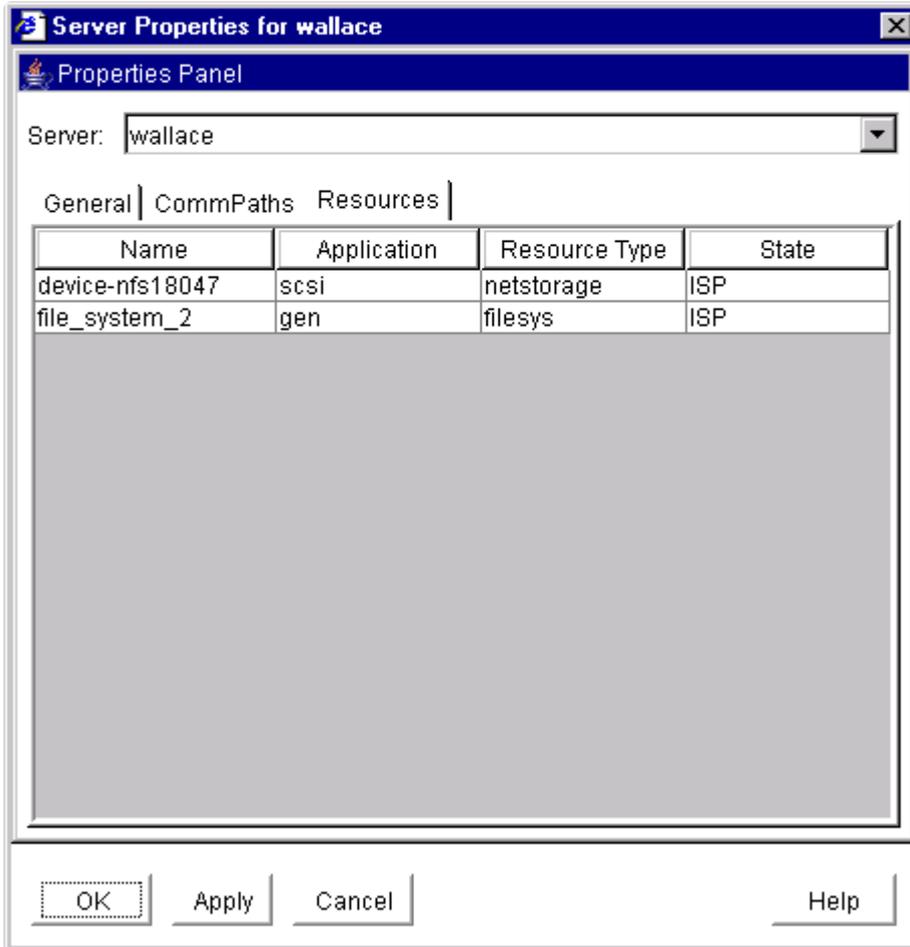
## CommPaths Tab



- **Server.** The server name of the other server the communication path is connected to in the LifeKeeper cluster.

- **Priority.** The priority determines the order by which communication paths between two servers will be used. Priority 1 is the highest and priority 99 is the lowest.
- **State.** State of the communications path in the LifeKeeper Configuration Database (LCD). These are the possible communications path state values:
  - *ALIVE* – functioning normally.
  - *DEAD* – no longer functioning normally.
  - *UNKNOWN* – state could not be determined. The GUI server may not be available.
- **Type.** The type of communications path, TCP (TCP/IP) or TTY, between the server in the list and the server specified in the Server field.
- **Address/Device.** The IP address or device name that this communications path uses.
- **Comm Path Status.** Summary communications path status determined by the GUI based on the state of the communications paths in the LifeKeeper Configuration Database ([LCD](#)). These are the possible communications path status values displayed below the detailed text in the lower panel:
  - *NORMAL* – all comm paths functioning normally.
  - *FAILED* – all comm paths to a given server are dead.
  - *UNKNOWN* – comm path status could not be determined. The GUI server may not be available.
  - *WARNING* – one or more comm paths to a given server are dead.
  - *DEGRADED* – one or more redundant comm paths to a given server are dead.
  - *NONE DEFINED* – no comm paths defined.

## Resources Tab



- **Name.** The tag name of a resource instance on the selected server.
- **Application.** The application name of a resource type (gen, scsi, ...)
- **Resource Type.** The resource type, a class of hardware, software, or system entities providing a service (for example, app, filesystem, nfs, device, disk,...)
- **State.** The current state of a resource instance:
  - *ISP* – In-service locally and protected.
  - *ISU* – In-service locally, but local recovery will not be attempted.
  - *OSF* – Out-of-service, failed.
  - *OSU* – Out-of-service, unimpaired.
  - *ILLSTATE* – Resource state has not been initialized properly by the resource initialization process which is run as part of the LifeKeeper startup sequence. Resources in this state are not under LifeKeeper protection.
  - *UNKNOWN* – Resource state could not be determined. The GUI server may not be available.

## 5.4.4.1.8. Operator Tasks

---

The following topics are more advanced tasks that require Operator permission.

---

[Bringing a Resource In Service](#)

[Taking a Resource Out of Service](#)

## 5.4.4.1.8.1. Bringing a Resource In Service

\* **Note:** LifeKeeper puts resources in-service from the bottom of the hierarchy and works its way to the top level resource. When putting all of the resources in one hierarchy in-service, select the top (parent) resource in the hierarchy.

1. There are five possible ways to begin.
  - Right-click on the icon for the resource/server combination that you want to bring into service. When the [Resource Context Menu](#) appears, click **In Service**.
  - Right-click on the icon for the global resource that you want to bring into service. When the **Resource Context Menu** appears, click **In Service**. When the dialog comes up, select the server on which to perform the In Service from the **Server list** and click **Next**.
  - On the [Global Toolbar](#), click the **In Service** button. When the dialog comes up, select the server on which to perform the In Service from the **Server list** and click **Next**. On the next dialog, select one or more resources that you want to bring into service from the **Resource(s)** list and click **Next** again.
  - On the [Resource Context Toolbar](#), if displayed, click the **In Service** button.
  - On the [Edit Menu](#), point to **Resource** and then click **In Service**. When the dialog comes up, select the server on which to perform the **In Service** from the **Server list**, and click **Next**. On the next dialog, select one or more resources that you want to bring into service from the Resource(s) list and click **Next** again.
2. A dialog appears confirming the server and resource(s) that you have selected to bring into service. This dialog will include a warning that if you are bringing a dependent child resource into service, it will also bring the parent resource into service. Click **In Service** to bring the resource(s) into service along with any dependent child resources.
3. If the [Output Panel](#) is enabled, the dialog closes and the results of the commands to bring the resource(s) in service are shown in the **output panel**. If not, the dialog remains up to show these results and you click **Done** to finish when all results have been displayed. Any additional dependent (child) resources that were brought into service are noted in the dialog or **output panel**.
4. Errors that occur while bringing a resource in service are logged in the LifeKeeper log of the server on which you want to bring the resource into service.

## 5.4.4.1.8.2. Taking a Resource Out of Service

\* **Note:** LifeKeeper takes resources out of service from the top of the hierarchy and works its way down to the other resources. When taking resources out of service if you want the entire hierarchy (parent with the child resources) to be taken out of service, take the lowest level resource out of service.

1. There are four possible ways to begin.
  - Right-click on the icon for the global resource or resource/server combination that you want to take out of service. When the [Resource Context Menu](#) appears, click **Out of Service**.
  - On the [Global Toolbar](#), click the **Out of Service** button. When the [Out of Service](#) dialog comes up, select one or more resources that you want to take out of service from the Resource(s) list, and click **Next**.
  - On the [Resource Context Toolbar](#), if displayed, click the **Out of Service** button.
  - On the [Edit Menu](#), point to **Resource** and then click **Out of Service**. When the **Out of Service** dialog comes up, select one or more resources that you want to take out of service from the **Resource(s)** list, and click **Next**.
2. An **Out of Service** dialog appears confirming the selected resource(s) to be taken out of service. This dialog will include a warning that if you are taking a dependent child resource out of service, it will also take the parent resource(s) out of service. Click **Out of Service** to proceed to the next dialog box.
3. If the [Output Panel](#) is enabled, the dialog closes, and the results of the commands to take the resource(s) out of service are shown in the output panel. If not, the dialog remains up to show these results, and you click **Done** to finish when all results have been displayed.
4. Errors that occur while taking a resource out of service are logged in the LifeKeeper log of the server on which you want to take the resource out of service.

## 5.4.4.1.9. Advanced Tasks

---

[LCD](#)

[LCM](#)

[LifeKeeper API for Monitoring](#)

## 5.4.4.1.9.1. LCD

---

### LifeKeeper Configuration Database

The LifeKeeper Configuration Database (LCD) maintains the object-oriented resource hierarchy information and stores recovery direction information for all resource types known to LifeKeeper. The data is cached within system shared memory and stored in files so that configuration data is retained over system restarts. The LCD also contains state information and specific details about resource instances required for recovery.

See the following related topics for information on the LCD directory structure, types of data stored, resource types available and use of application scripts.

#### Related Topics

---

[LCDI Commands](#)

[LCD Configuration Data](#)

[LCD Directory Structure](#)

[LCD Resource Types](#)

[LifeKeeper Flags](#)

[Resources Subdirectories](#)

[Structure of LCD Directory in /opt/LifeKeeper](#)

## 5.4.4.1.9.1.1. LCDI Commands

---

### Steps to Create Resources by Defining Your Own Recovery Kit

**Note:** Use the GUI to create resources if using the existing Recovery Kit.

LifeKeeper provides two mechanisms for defining an application resource hierarchy:

- LifeKeeper GUI
- LifeKeeper Configuration Database Interface (LCDI) commands

The LCDI is a set of interface commands provided by LifeKeeper that you can use to create and customize resource hierarchy configurations to meet your application needs. You use the command interface when an application depends upon multiple resources (such as two or more file systems).

For a description of the commands, see the LCDI manual pages. This topic provides a development scenario that demonstrates the way you can use both the GUI and command functions to create a resource hierarchy.

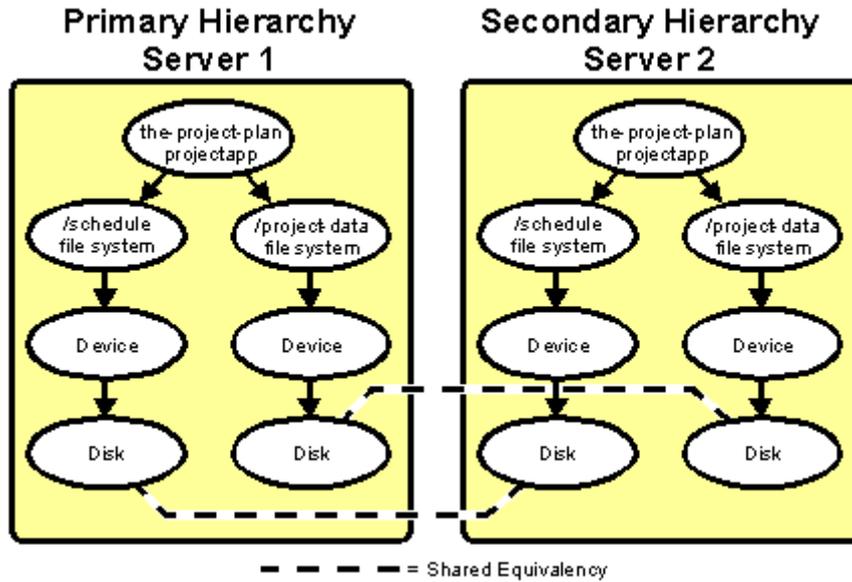
### Scenario Situation

The example application, ProjectPlan, has data stored in SCSI file systems shared by Servers 1 and 2. Server 1 will be the primary hierarchy for the application. The application has two file systems: `/project-data` and `/schedule`. The first step in the hierarchy definition is to determine the dependencies.

The example application has these dependencies:

- **Shared file systems.** The application depends upon its file systems: `/project-data` and `/schedule`.
- **SCSI disk subsystem.** The file systems in turn depend upon the SCSI disk subsystem, which includes the device, disk and host adapter resources.

As a result, the task is to create a hierarchy that looks like the following diagram.



## Hierarchy Definition

These are the tasks required to construct the example application hierarchy:

1. **Create file system resources.** The LifeKeeper GUI provides menus to create file system resources. See [Creating File System Resource Hierarchies](#).

At the end of this definition task, the LCD has two filesys resources defined as follows:

ID	Tag	Server
/project-data	project-data-on-Server1	Server1
/project-data	project-data-from-Server1	Server2
/schedule	schedule-on-Server1	Server1
/schedule	schedule-from-Server1	Server2

**Note:** LifeKeeper does not place any significance on the tag names used; they are simply labels. The tag names shown are the LifeKeeper defaults.

2. **Define resources.** The example requires the following definitions:

Application:	projectapp
Resource Type:	plan
Instance ID:	1yrplan
Tag:	the-project-plan

**Note:** Although you can create much of the definition using the LifeKeeper GUI, the rest of this example demonstrates the command interface.

3. **Create directories.** On each system, you create the necessary application recovery directories under the directory `/opt/LifeKeeper/subsys` with the command:

```
mkdir -p /opt/LifeKeeper/subsys/projectapp/Resources/plan/actions
```

4. **Define application.** The following commands create the application named `projectapp`:

```
app_create -d Server1 -a projectapp
```

```
app_create -d Server2 -a projectapp
```

5. **Define the resource type.** The following commands create the resource type named `plan`:

```
typ_create -d Server1 -a projectapp -r plan
```

```
typ_create -d Server2 -a projectapp -r plan
```

6. **Install recovery scripts.** Copy your restore and remove scripts to the following directory on each server:

```
/opt/LifeKeeper/subsys/projectapp/Resources/plan/actions
```

7. **Define instance.** The following commands define an instance of resource type `plan` with the id `1yrplan`:

```
ins_create -d Server1 -a projectapp -r plan -I\
```

```
AUTORES_ISP -t the-project-plan -i 1yrplan
```

```
ins_create -d Server2 -a projectapp -r plan -I\
```

```
SEC_ISP -t the-project-plan -i 1yrplan
```

The `-I AUTORES_ISP` instruction for the instance created on Server1 tells LifeKeeper to automatically bring the resource in service when LifeKeeper is restarted. In this case, the resource's restore script is run and, if successful, the resource is placed in the ISP state. This operation is not performed if the paired resource is already in service.

The `-I SEC_ISP` instruction for the instance created on Server2 tells LifeKeeper that this resource instance should not be brought into service when LifeKeeper is restarted. Instead, Server2 will serve as the backup for the resource on Server1, and the local resource will be brought in service upon failure of the primary resource or server.

8. **Define dependencies.** The following commands define the dependencies between the application and the file systems:

```
dep_create -d Server1 -p the-project-plan -c project-data-on-System1
```

```
dep_create -d Server2 -p the-project-plan -c project-data-from-  
Server1
```

```
dep_create -d Server1 -p the-project-plan -c schedule-on-Server1
```

```
dep_create -d Server2 -p the-project-plan -cschedule-from-Server1
```

9. **Execute lcdsync.** Execute the following lcdsync commands to inform LifeKeeper to update its copy of the configuration:

```
lcdsync -d Server1
```

```
lcdsync -d Server2
```

10. **Set the resources to “In Service”.** Access LifeKeeper GUI on the primary server, select **[Edit] > [Resource] > [In-Service]** and set the resource “In Service”, or execute the following command on the primary server:

```
perform_action -t the-project-plan -a restore
```

11. **Create equivalency.** Create equivalency of resources registered for each node with the following commands to switch resources:

```
eqv_create -d Server1 -t the-project-plan -p 1 -S Server2 -o the-  
project-plan -r 10 -e SHARED
```

```
eqv_create -d Server2 -t the-project-plan -p 10 -S Server1 -o the-  
project-plan -r 1 -e SHARED
```

## 5.4.4.1.9.1.2. LCD Configuration Data

---

LCD stores the following related types of data:

- Dependency Information
- Resource Status Information
- Inter-Server Equivalency Information

### Dependency Information

For each defined resource, LifeKeeper maintains a list of dependencies and a list of dependents (resources depending on a resource.) For information, see the `LCDI_relationship` (1M) and `LCDI_instances` (1M) manual pages.

### Resource Status Information

LCD maintains status information in memory for each resource instance. The [resource states](#) recognized by **LCD** are **ISP**, **ISU**, **OSF**, **OSU** and **ILLSTATE**. Resources may change from one state to another when a system event occurs or when an administrator takes certain actions. When a resource changes states, the status change is reflected in the LCD on the local server as well as in the database of the backup servers for that resource.

### Inter-Server Equivalency Information

Relationships may exist between resources on various servers. A [shared equivalency](#) is a relationship between two resources on different servers that represents the same physical entity. When two servers have a resource with a shared equivalency relationship, LifeKeeper attempts to ensure in its actions that only one of the two servers has the resource instance in the in-service, protected [ISP] state at any one time. Both servers can have the resource instance in an out-of-service state [**OSU** or **OSF**], but for data integrity reasons, only one server can have the resource in service at any given time.

Disks on a Small Computer System Interface (SCSI) bus are one example of equivalent resources. With the SCSI locking (or reserve) mechanism, only one server can own the lock for a disk device at any point in time. This lock ownership feature guarantees that two or more servers cannot access the same disk resource at the same time.

Furthermore, the dependency relationships within a hierarchy guarantee that all resources that depend upon the disk, such as a file system, are in service on only one server at a time.

## 5.4.4.1.9.1.3. LCD Directory Structure

---

Major subdirectories under */opt/LifeKeeper*:

- **config.** LifeKeeper configuration files, including shared equivalencies.
- **bin.** LifeKeeper executable programs, such as `is_recoverable`. See [Fault Detection and Recovery Scenarios](#) for descriptions.
- **subsys.** Resources and types. LifeKeeper provides resource and type definitions for the shared SCSI disk subsystem in `scsi` and for the generic application menu functions in `gen`. When you define an application interface, you create directories under `subsys`.
- **events.** Alarming events. See [LifeKeeper Alarming and Recovery](#) for further information.

The structure of the LCD directory in */opt/LifeKeeper* is shown in the topic [Structure of LCD Directory in /opt/LifeKeeper](#).

## 5.4.4.1.9.1.4. LCD Resource Types

---

The LCD is maintained in both shared memory and in the `/opt/LifeKeeper` directory. As highlighted on the [directory structure diagram](#), `subsys` contains two application resource sets you can use to define your application interface:

- **gen** – generic application and file system information
- **scsi** – recovery information specific to the SCSI

These subdirectories are discussed in [Resources Subdirectories](#).

## 5.4.4.1.9.1.5. LifeKeeper Flags

---

Near the end of the [detailed status display](#), LifeKeeper provides a list of the flags set for the system. A common type is a Lock LCD flag used to ensure that other processes wait until the process lock completes its action. The following is the standard LCD lock format:

```
!action!processID!time!machine:id.
```

These are examples of general LCD lock flags:

- **!action!02833!701236710!<servername>:filesys**. The creation of a filesystem hierarchy produces a flag in this format in the status display. The *filesys* designation can be a different resource type for other application resource hierarchies or *app* for generic or user-defined applications.
- Other typical flags include **!nofailover!machine** and **shutdown\_switchover**. The **!nofailover!machine** flag is an internal, transient flag created and deleted by LifeKeeper which controls aspects of server failover. The **shutdown\_switchover** flag indicates that the shutdown strategy for this server has been set to switchover such that a shutdown of the server will cause a switchover to occur. See `LCDI-flag(1M)` for more detailed information on the possible flags.

## 5.4.4.1.9.1.6. Resources Subdirectories

The **scsi** and **gen** directories each contain a resources subdirectory. The content of those directories provides a list of the resource types provided by LifeKeeper:

**scsi resource types.** You find these resource types in the `/opt/LifeKeeper/subsys/scsi/resources` directory. Note that there may be additional directories depending upon your configuration.

- **device** —disk partitions or virtual disk devices
- **disk** —physical disks or LUNs
- **hostadp** —host adapters

**gen resource types.** You find these resource types in the `/opt/LifeKeeper/subsys/gen/resources` directory:

- **filesystems** —file systems
- **app** —generic or user-defined applications that may depend upon additional resources

Each resource type directory contains one or more of the following:

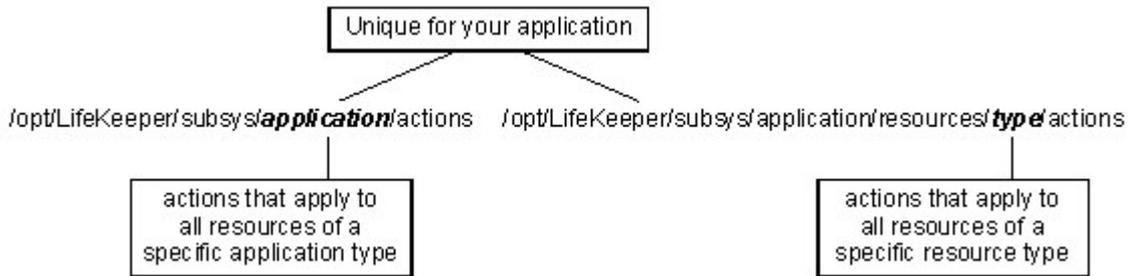
- **instances.** This file reflects the permanent information saved in the LCD about resource instances. It contains descriptive information for the resource instances associated with this resource type.

**! WARNING:** Do not modify the instances file (or any LCD file) directly. To create or manipulate resource instances, use only the LifeKeeper GUI functions or the LifeKeeper LCDI\_instances commands: `ins_create`, `ins_remove`, `ins_gettag`, `ins_setas`, `ins_setinfo`, `ins_setinit`, `ins_setstate` and `ins_list`. Refer to the LCDI\_instances (1M) manual pages for explanations of these commands.

- **recovery.** This optional directory contains the programs used to attempt the local recovery of a resource for which a failure has been detected. The recovery directory contains directories that correspond to event classes passed to `sendevent`. The names of the directories must match the class parameter (-C) passed to the `sendevent` program. (See [LifeKeeper Alarming and Recovery](#).)

In each subdirectory, the application can place recovery programs that service event types of the corresponding event class. The name of these programs must match the string passed to `sendevent` with the -E parameter. This optional directory may not exist for many applications.

- **actions.** This directory contains the set of recovery action programs that act only on resource instances of the specific resource type. If, for your application, any actions apply to all resource types within an application, place them in an **actions** subdirectory under the application directory rather than under the **resource type** directory.



Recovery direction software is used to modify or recover a resource instance. Two actions, **remove** and **restore**, must exist in the **actions** directory for each resource type.

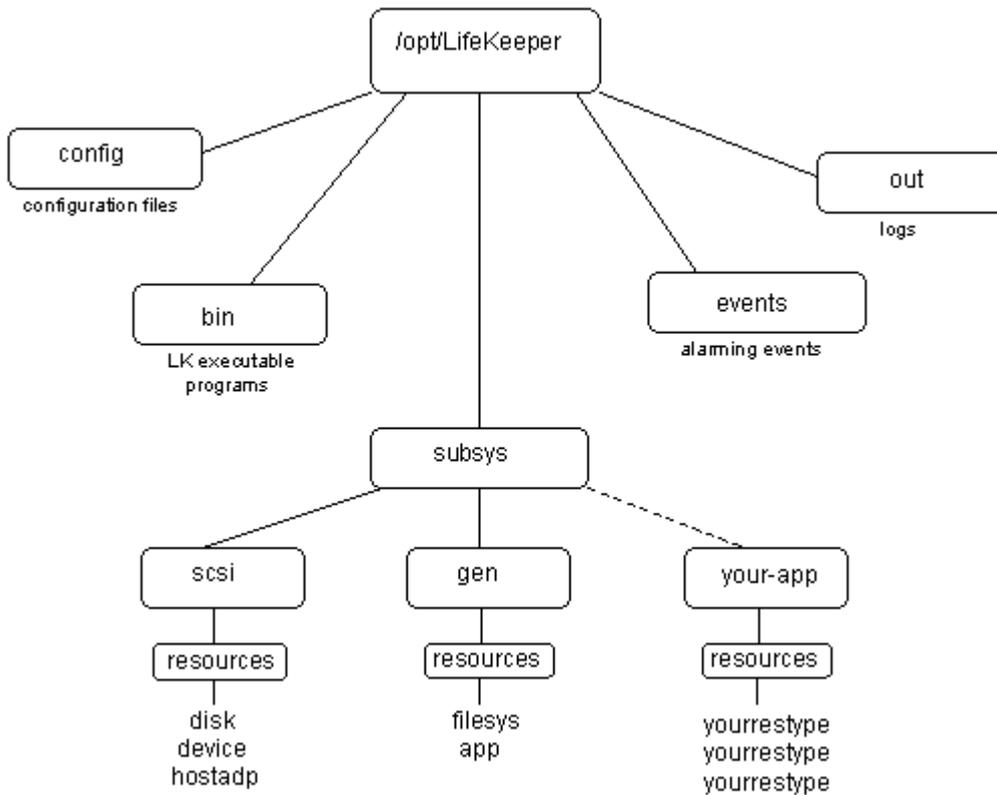
## Resource Actions

The **actions** directory for a resource type contains the programs (most often shell scripts) that describe specific application functions. Two actions are required for every resource type—restore and remove.

The remove and restore programs should perform symmetrically opposite functions; that is, they undo the effect of one another. These scripts should never be run manually. They should only be run by executing the LifeKeeper Recovery Action and Control Interface (LRACI) `perform_action` shell program described in the `LRACI-perform_action (1M)` manual page.

## 5.4.4.1.9.1.7. Structure of LCD Directory in /opt/LifeKeeper

The following diagram shows the directory structure of */opt/LifeKeeper*.



## 5.4.4.1.9.2. LCM

---

The LifeKeeper Communications Manager (LCM) provides reliable communication between processes on one or more LifeKeeper servers. This process can use redundant communication paths between systems so that failure of a single communication path does not cause failure of LifeKeeper or its protected resources. The LCM supports a variety of communication alternatives including RS-232 (TTY) and TCP/IP connections.

The LCM provides the following:

- **LifeKeeper Heartbeat.** Periodic communication with other connected LifeKeeper systems to determine if the other systems are still functioning. LifeKeeper can detect any total system failure that is not detected by another means by recognizing the absence of the heartbeat signal.
- **Administration Services.** The administration functions of LifeKeeper use the LCM facilities to perform remote administration. This facility is used for single-point administration, configuration verification and sanity checking of administrative actions.
- **Configuration and Status Communication.** The LifeKeeper configuration database (LCD) tracks resource status, availability and configuration through the LCM facilities. These
- **Failover Recovery.** If a resource fails on a system, the LCM notifies LifeKeeper to recover the resource on a backup system.

In addition to the LifeKeeper services provided by the LCM, inter-system application communication is possible through a set of shell commands for reliable communication. These commands include `snd_msg`, `rcv_msg`, and `can_talk`. These commands are described in the `LCMI_mailboxes (1M)` manual pages. The LCM runs as a real-time process on the system assuring that critical communications such as system heartbeat will be transmitted.

## Related Topics

---

[Communication Status Information](#)

[Alarming and Recovery](#)

## 5.4.4.1.9.2.1. Communication Status Information

---

The communications status information section of the status display lists the servers known to LifeKeeper and their current state followed by information about each communication path.

The following sample is from the communication status section of a short status display:

```
MACHINE NETWORK ADDRESSES/DEVICE STATE PRIO
tristan TCP 100.10.100.100/100.10.100.200 ALIVE 1
tristan TTY /dev/ttyS0 ALIVE -
```

For more information, see the communication status information section of the topics [Detailed Status Display](#) and the [Short Status Display](#).

## 5.4.4.1.9.2.2. LifeKeeper Alarming and Recovery

---

LifeKeeper error detection and notification is based on the event alarming mechanism, `sendevent`. The key concept of the **sendevent** mechanism is that independent applications can register to receive alarms for critical components. Neither the alarm initiation component nor the receiving application(s) need to be modified to know the existence of the other applications. Application-specific errors can trigger LifeKeeper recovery mechanisms via the **sendevent** facility.

This section discusses topics related to alarming including alarm classes, alarm processing and alarm directory layout and then provides a processing scenario that demonstrates the alarming concepts.

### Alarm Classes

The `/opt/LifeKeeper/events` directory lists a set of alarm classes. These classes correspond to particular sub-components of the system that produces events (for example, `filesys`). For each alarm class, subdirectories contain the set of potential alarms (for example, `badmount` and `diskfull`). You can register an application to receive these alarms by placing shell scripts or programs in the appropriate directories.

LifeKeeper uses a basic alarming notification facility. With this alarming functionality, all applications registered for an event have their handling programs executed asynchronously by `sendevent` when the appropriate alarm occurs. With LifeKeeper present, the **sendevent** process first determines if the LifeKeeper resource objects can handle the class and event. If LifeKeeper finds a class/event match, it executes the appropriate recover scenario.

Defining additional scripts for the **sendevent** alarming functionality is optional. When you define LifeKeeper resources, LifeKeeper provides the basic alarming functionality described in the processing scenarios later in this chapter.

 **Note:** Local recovery for a resource instance is the attempt by an application under control of LifeKeeper to return interrupted resource services to the end-user on the same system that generated the event. Inter-server recovery allows an application to migrate to a backup system. This type of recovery is tried after local recovery fails or is not possible.

### Alarm Processing

Applications or processes that detect an event which may require LifeKeeper attention can report the event by executing the **sendevent** program, passing the following arguments: respective error class, error name and failing instance. Refer to the `sendevent(5)` manual pages for required specifics and optional parameters and syntax.

## Alarm Directory Layout

The `/opt/LifeKeeper/events` directory has two types of content:

- **LifeKeeper supplied classes.** LifeKeeper provides two alarm classes listed under the `events` directory: `lifekeeper` and `filesys`. An example of an alarm event includes `diskfull`. The alarm classes correspond to the strings that are passed with the `-C` option to the `sendevent` command and the alarm events correspond to the strings that are passed with the `-E` option. The `lifekeeper` alarm class is used internally by LifeKeeper for event reporting within the LifeKeeper subsystem.
- **Application-specific classes.** The other subdirectories in the `events` directory are added when specific applications require alarm class definitions. Applications register to receive these alarms by placing shell scripts or binary programs in the directories. These programs are named after the application package to which they belong.

## 5.4.4.1.9.3. LifeKeeper API for Monitoring

### Introduction

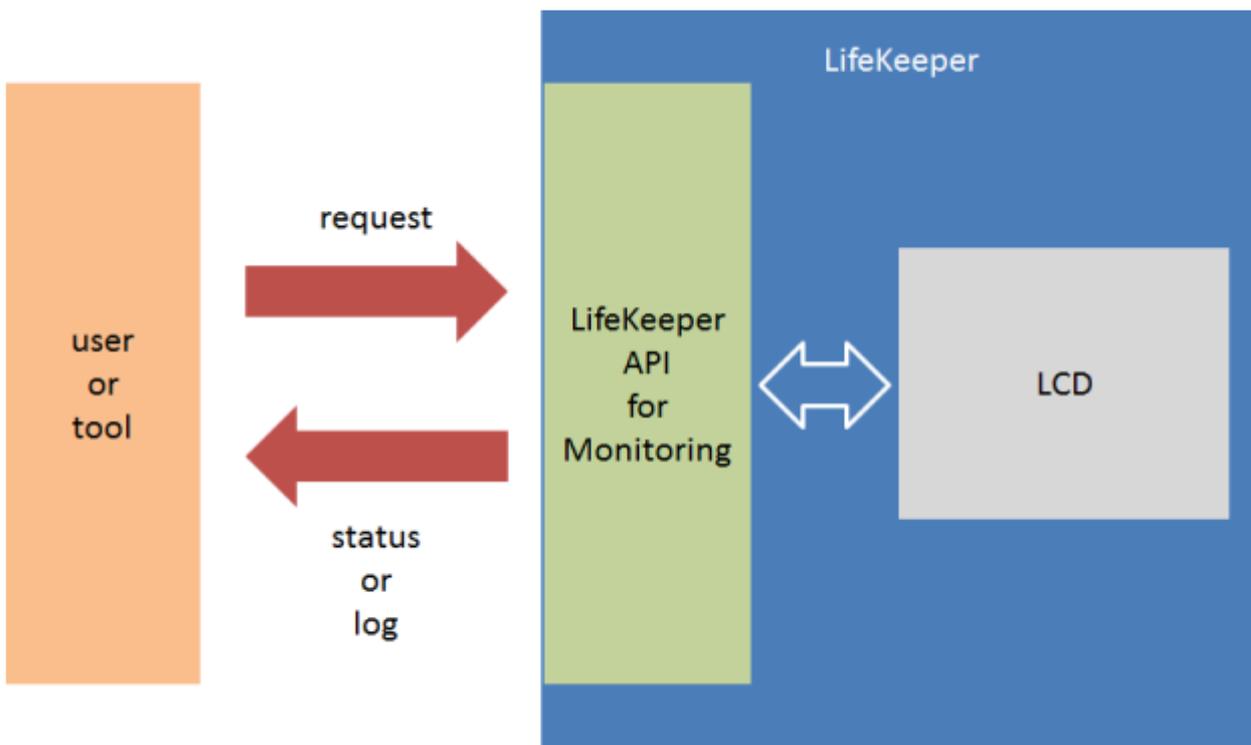
The LifeKeeper API for Monitoring can obtain the operational status of LifeKeeper nodes and their protected resources by making status inquiries to the available nodes in the LifeKeeper cluster.

### Summary

This document describes the LifeKeeper API for Monitoring (hereinafter referred to as the API) and is targeted for developers who manage the resource protected by LifeKeeper. By using the API, the information supplied by the `lcdstatus` command is obtained through CGI script and the `lighttpd` module. By using this API, users can determine the current status of the LifeKeeper nodes and resources without logging-in to LifeKeeper servers. The API can supply the following information.

- LifeKeeper node status is the node alive and processing or down
- Communication path status between nodes in the cluster, are communication path(s) up or down
- Status of protected resources

To get the detailed status of any abnormal condition requires logging-in to LifeKeeper GUI or checking the LifeKeeper log as necessary.



## Information to be supplied with this API

The following information is supplied through this API when the user makes an inquiry to an active LifeKeeper node. The information supplied is about the specific LifeKeeper server to which the inquiry was directed even if the cluster consists of multiple servers.

- Status

- ° Operating status of each server

- Node name
    - Operational status (ALIVE/DEAD)

- ° Operational status of communication path(s)

- Node name
    - Operational status (ALIVE/DEAD)
    - Address / device name

- ° Status of protected resources

- Node Name
    - Tag
    - Status (ISP, OSU, OSF, ...)
    - Dependency setting
    - Mirror information for Data Replication resources (available only if status is ISP)

- Tag
- Mirror status (Sync, Paused, ...)
- Replication status (75%, 100%, ...)
- Log

◦ /var/log/lifekeeper.log \*Not supported if log file path is changed

- Up to 1000 lines (when data output format is HTML)
- All (when data output format is plain text)

◦ /var/log/lifekeeper.err \*Not supported if log file path is changed

- Up to 1000 lines (when data output format is HTML)
- All (when data output format is plain text)

## Communication Format

The API uses HTTP to obtain the requested information. To obtain information, the user sends a HTTP GET request to the CGI scripts via lighttpd on the specific server.

## Data Format

The following 3 data formats are available.

- JSON
  - To be used by an external tool to analyze the status information returned
  - Status checking is possible

- Log output is not available
  
- HTML
  - To be used to visually check via a browser
  
  - Status checking is possible
  
  - Log information is available up to 1000 lines
  
- plain text
  - Used for regular log checking
  
  - For logging purpose only and not for checking the status
  
  - All contents of /var/log/lifekeeper.log and /var/log/lifekeeper.err are available

Available JSON format and HTML format from the status in the following figure.



Figure 2. Example of active LifeKeeper configuration

```
{  
  "resource": [  
    {  
      "replication": {},  
      "child": [  
        {  
          "tag": "datarep-data"  
        }  
      ],  
      "server": {  
        "status": "ISP",  
        "name": "lk01"  
      },  
      "tag": "/data"  
    },  
    {  
      "replication": {  
        "percent": "100%",  
        "mirror": "Fully Operational"  
      },  
      "child": [],  
      "server": {  
        "status": "ISP",
```

```
        "name": "lk01"

    },

    "tag": "datarep-data"

},

{

    "replication": {},

    "child": [],

    "server": {

        "status": "ISP",

        "name": "lk01"

    },

    "tag": "ip-10.125.139.118"

}

],

"compath": [

    {

        "status": "ALIVE",

        "server": [

            {

                "name": "lk01",

                "term": "192.168.139.18"

            },

            {

                "name": "lk02",
```

```
        "term": "192.168.139.19"
    }
]
},
{
    "status": "ALIVE",
    "server": [
        {
            "name": "lk01",
            "term": "172.20.139.18"
        },
        {
            "name": "lk02",
            "term": "172.20.139.19"
        }
    ]
}
],
"server": [
    {
        "status": "ALIVE",
        "name": "lk01"
    },
    ],
```

```

    {
        "status": "ALIVE",
        "name": "lk02"
    }
]

```

```

}

```

Figure 3. Status output example using the JSON data format

#### RESOURCES

tag	lk01
/data	ISP
datarep-data	ISP
ip-10.125.139.118	ISP

#### DATA REPLICATIONS

tag	nodes	mirror status	replication status
datarep-data	lk01 -> lk02	Fully Operational	100%

#### COMMUNICATION PATHS

communication path	status
192.168.139.18/192.168.139.19	ALIVE
172.20.139.18/172.20.139.19	ALIVE

Figure 4. Status output using the HTML format

## How to use

### Activate the API

The API is disabled by default. To activate, requires modification of `/etc/default/LifeKeeper` set the `LKAPI_MONITORING` configuration parameter to true. Setting of the configuration parameter only activates the API on that node and therefore must be set on each node on which the API will be used. Setting of this configuration parameter does not require a restart of LifeKeeper..

```
LKAPI_MONITORING=true
```

### Port number

The API uses port 779 by default. To change the port number, the user needs to set the following in `/etc/`

default/LifeKeeper.

LKAPI\_WEB\_PORT=<port number>

### Usage examples

To obtain information a request is made to a server with an active LifeKeeper API configuration. Basic example using curl.

```
curl http://<IPADDR>:779/Monitoring.cgi
```

If no arguments are given, the current status is obtained using the JSON data format.

Request for log information using HTML data format.

```
curl http://<IPADDR>:779/Monitoring.cgi?format=html&show=log
```

The list of available arguments can be found in the table below.

List 1. Arguments

Name	Explanation	Value	Comments
show	Specify the target information	status, log, log-err	show=status is the default
format	Specify data format	json, html, plain	format=json is the default. If the format is json an error will be displayed if show=log or show=log-err is set.

### Security

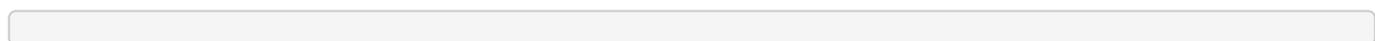
All the users requesting information via the API must be authorized to get LifeKeeper status information.

For this reason, user security settings can limit the users who can get the status by, configuring SSL, and encrypting the information.

#### Basic Authentication

To obtain the information via the API, Basic Authentication is required. To setup the authentication requires modification to the lighttpd configuration file (modify the part in red) plus a restart of the lighttpd module.

After modification execute the command `"/opt/LifeKeeper/sbin/sv restart steeleye-lighttpd"` and reboot lighttpd to restart lighttpd using the new configuration.



```
server.modules = ( : "mod_auth", # uncommenting
```

```
/opt/LifeKeeper/lib64/steeleye-lighttpd/include_server_bind.pl
```

```
print qq/\$SERVER["socket"] == "$addr:$port" {\n/;
print qq/    server.document-root = "\/opt\/LifeKeeper\/api"\n/;
print qq/    auth.backend = "htpasswd"\n/;
print qq/    auth.backend.htpasswd.userfile = "\/opt\/LifeKeeper\/etc\/lighttp
d\/lighttpd.user.htpasswd"\n/;
print qq/    auth.require = ( "\/" =>\n/;
print qq/        (\n/;
print qq/                "method" => "basic",\n/;
print qq/                "realm"  => "LifeKeeperAPI",\n/;
print qq/                "require" => "valid-user"\n/;
print qq/        )\n/;
print qq/    )\n/;
print qq/ }\n/;
```

To create htpasswd file:

```
htpasswd -c /opt/LifeKeeper/etc/lighttpd/lighttpd.user.htpasswd
<USERNAME>
```

## SSL/TLS set up

SSL/TLS is available for the communication via this API. The lighttpd modifications for SSL/TLS is shown in the example below. After modification execute the command `"/opt/LifeKeeper/sbin/sv restart steeleye-lighttpd"` and reboot lighttpd to restart lighttpd with the new configuration.

In addition, if it is no longer necessary to set the default port used by this API due to this support, disable it by referring to the following "Disable port 779".

```
/opt/LifeKeeper/etc/lighttpd/include_ssl_port.pl
```

```
configAPI("0.0.0.0", 443);
if(socket($sock, AF_INET6, SOCK_STREAM, 0)) {
    configAPI("[::]", 443);
}
sub configAPI {
    my $addr = shift;
    my $port = shift;

    print qq/\$SERVER["socket"] == "$addr:$port" {\n/;
    print qq/    server.document-root = "\/opt\/LifeKeeper\/api"\n/;
    print qq/    ssl.engine = "enable"\n/;
    print qq/    ssl.pemfile = "\/opt\/LifeKeeper\/etc\/certs\/LK4LinuxValidNo
de.pem"\n/;
```

```

print qq/      ssl.use-sslv2 = "disable"\n/;
print qq/      ssl.use-sslv3 = "disable"\n/;
print qq/ } \n/;
}

```

### Modification to support SSL/TLS + Basic authentication

Using SSL/TLS, modification example to set up Basic authentication is below. After modification, execute the command “/opt/LifeKeeper/sbin/sv restart steeleye-lighttpd” and restart lighttpd to reflect the modified set up.

/opt/LifeKeeper/lib64/steeleye-lighttpd/include\_server\_bind.pl

```

#lkapi_config("0.0.0.0", $lkapi_port);
#if(socket($sock, AF_INET6, SOCK_STREAM, 0)) {
#    lkapi_config("[:,]", $lkapi_port);
#}

```

### IP address access limitation

The lighttpd configuration can also be setup to limit IP addresses that can be used to access data via the API. The lighttpd configuration to limit access is shown in Figure 9. The example will reject the connections from IP address other than 192.168.10.1. After modification execute the command “/opt/LifeKeeper/sbin/sv restart steeleye-lighttpd” to restart lighttpd with the new configuration.

/opt/LifeKeeper/etc/lighttpd/conf.d/lkapi\_user.conf

```

$HTTP["remoteip"] != "192.168.10.1" {
url.access-deny = ( "" )
}

```

## Error

Errors can occur during the usage of this API when enabled. Should this occur, the summary of the error is output. Error example when JSON format is shown below.

HTTP status code returned by lighttpd is not described here.

```

{
"error" : {
    id : -1,
    message : "Failed to get LCD status"
}
}

```

Similar message is output in the case the output format is HTML.

## 5.4.4.1.10. Maintenance Tasks

---

The following are tasks for maintaining LifeKeeper.

---

[Changing Configuration Values](#)

[File System Health Monitoring](#)

[Maintaining Protected System](#)

[Maintaining a Resource Hierarchy](#)

[Recovering After a Failover](#)

[Removing LifeKeeper](#)

[Running With a Firewall](#)

[Running GUI Through a Firewall](#)

[Transferring Resource Hierarchies](#)

## 5.4.4.1.10.1. Changing LifeKeeper Configuration Values

---

There are a number of values in LifeKeeper that may need to be changed after LifeKeeper has been configured and set up. Examples of values that may be modified include the unname of LifeKeeper servers, comm path ip addresses, ip resource addresses and tag names. To change these values, carefully follow the instructions below.

1. Stop LifeKeeper on all servers in the cluster using the command:

```
$LKROOT/bin/lkcli stop
```

There is no need to delete comm paths or unextend resource hierarchies from any of the servers.

2. If you are changing the unname of a LifeKeeper server, change the server's hostname using the Linux `hostname(1)` command.
3. Before continuing, ensure that any new host names are resolvable by all of the servers in the cluster. If you are changing comm path addresses, check that the new addresses are configured and working (the **ping** and **telnet** utilities can be used to verify this).
4. If more than one LifeKeeper value is to be changed, old and new values should be specified in a file on each server in the cluster in the following format:

```
old_value1=new_value1
```

```
....
```

```
old_value9=new_value9
```

5. Verify that the changes to be made do not have any unexpected side effects by examining the output of running the **lk\_chg\_value** command on **all** servers in the cluster. If there is more than one value to change, run the command:

```
$LKROOT/bin/lk_chg_value -Mvf file_name
```

where **file\_name** is the name of the file created in Step 4.

If there is only one value to change, run the command:

```
$LKROOT/bin/lk_chg_value -Mvo old_value -n new_value
```

The **-M** option specifies that no modifications should be made to any LifeKeeper files.

6. Modify LifeKeeper files by running the `lk_chg_value` command without the **-M** option on all servers in the cluster. If there is more than one value to change, run the command:

```
$LKROOT/bin/lk_chg_value -vf file_name
```

where **file\_name** is the name of the file created in Step 4.

If there is only one value to change, run the command:

```
$LKROOT/bin/lk_chg_value -vo old_value -n new_value
```

- Restart LifeKeeper using the command:

```
$LKROOT/bin/lkcli start
```

If the cluster is being viewed using the LifeKeeper GUI, it may be necessary to close and restart the GUI.

### Example:

*Server1* and *Server2* are the LifeKeeper server unames in a two-node cluster. *Server1* has a comm path with address 172.17.100.48. *Server2* has an ip resource with address 172.17.100.220 which is extended to *Server1*. To change the following values for *Server1*:

Value	Old	New
uname	Server1	Newserver1
comm path address	172.17.100.48	172.17.105.49
IP resource address	172.17.100.220	172.17.100.221

The following steps should be performed to make these changes.

- Stop LifeKeeper on both *Server1* and *Server2* using the command:

```
$LKROOT/bin/lkcli stop
```

- Change the uname of # *Server1* to *Newserver1* using the command:

```
hostname Newserver1
```

- Create the file, */tmp/subs*, with the content below, on both *Newserver1* and *Server2*:

```
Server1=Newserver1
```

```
172.17.100.48=172.17.105.49
```

```
172.17.100.220=172.17.100.221
```

- Verify that the changes specified will not have any unexpected side effects by examining the output of running the following command on both servers:

```
$LKROOT/bin/lk_chg_value -Mvf /tmp/subs
```

5. Modify the LifeKeeper files by running the `lk_chg_value` command without the `-M` option on both servers:

```
$LKROOT/bin/lk_chg_value -vf /tmp/subs
```

6. Restart LifeKeeper on both servers using the command:

```
$LKROOT/bin/lkcli start
```

 **Note:** To see the changes `lk_chg_value` will make without modifying any LifeKeeper files, use the `-M` option. To see the files `lk_chg_value` is examining, use `-v`. To not modify tag names, use the `-T` option. To not modify resource ids, use the `-I` option.

## 5.4.4.1.10.2. File System Health Monitoring

The File System Health Monitoring feature detects conditions that could cause LifeKeeper protected applications that depend on the file system to fail. Monitoring occurs on active/in-service resources (i.e. file systems) only. The two conditions that are monitored are:

- A full (or almost full) file system, and
- An improperly mounted (or unmounted) file system.

When either of these two conditions is detected, one of several actions might be taken.

- A warning message can be logged and email sent to a system administrator.
- Local recovery of the resource can be attempted.
- The resource can be failed over to a backup server.

### Condition Definitions

#### Full or Almost Full File System

A “disk full” condition can be detected, but cannot be resolved by performing a local recovery or failover – administrator intervention is required. A message will be logged by default. Additional notification functionality is available. For example, an email can be sent to a system administrator, or another application can be invoked to send a warning message by some other means. To enable notification for the full/almost full disk conditions a basic event notification script name `notify` has been provided in the directory `/opt/LifeKeeper/events/filesys/diskfull`. Simply add the functionality required to send email or execute another application.

In addition to a “disk full” condition, a “disk almost full” condition can be detected and a warning message logged in the LifeKeeper log.

The “disk full” threshold is:

```
FILESYSFULLERROR=95
```

The “disk almost full” threshold is:

```
FILESYSFULLWARN=90
```

The default values are 90% and 95% as shown, but are configurable via tunables in the `/etc/default/LifeKeeper` file. The meanings of these two thresholds are as follows:

`FILESYSFULLWARN` – When a file system reaches this percentage full, a message will be displayed in the LifeKeeper log.

`FILESYSFULLERROR` – When a file system reaches this percentage full, a message will be displayed in the LifeKeeper log as well as the system log. The file system notify script will also be called.

## Unmounted or Improperly Mounted File System

LifeKeeper checks the `/etc/mtab` file to determine whether a LifeKeeper protected file system that is in service is actually mounted. In addition, the mount options are checked against the stored mount options in the `filesys` resource information field to ensure that they match the original mount options used at the time the hierarchy was created.

If an unmounted or improperly mounted file system is detected, local recovery is invoked and will attempt to remount the file system with the correct mount options.

If the remount fails, failover will be attempted to resolve the condition. The following is a list of common causes for remount failure which would lead to a failover:

- corrupted file system (fsck failure)
- failure to create mount point directory
- mount point is busy
- mount failure
- LifeKeeper internal error

## New or Deprecated Mount Options After a Kernel Upgrade

When upgrading the Linux kernel, it is possible that some existing file system mount options may be deprecated in the new kernel or that the new kernel may add new default mount options to existing mounts. For example, the “noBarrier” mount option was deprecated in RedHat Enterprise Linux 8, and some kernel versions have added new default mount options such as “logbufs=8” and “logsize=32k”.

If a LifeKeeper-protected file system resource contains mount options which become deprecated after a kernel upgrade, the deprecated options should be removed from the list of mount options for the LifeKeeper resource on every server in the cluster. See the [Modifying Mount Options for a LifeKeeper File System Resource](#) section for more details.

If new default mount options are added by the kernel to an existing LifeKeeper-protected mount point after a kernel upgrade, then the new options should be added to the list of mount options for the LifeKeeper resource on every server in the cluster. See the [Modifying Mount Options for a LifeKeeper File System Resource](#) section for more details.

# Modifying Mount Options for a LifeKeeper File System Resource

To modify the mount options used by LifeKeeper when mounting a protected file system:

1. Take the file system resource out of service in LifeKeeper. This will unmount the protected file system.
2. Update the mount options on each server in the cluster.
  - a. Using the LifeKeeper GUI
    - i. Right-click on the file system resource on each server where you would like to change the mount options and select “Change Mount Options”.
    - ii. In the resulting dialog, modify the mount options by providing a comma-separated list of options to be used when LifeKeeper mounts the file system. Once the desired mount options have been entered, click “Set Value”. Click “Finish” to exit the confirmation dialog.
  - b. Using the [LifeKeeper Command Line Interface](#)
    - i. 

```
lkcli resource config fs --tag <tag> --mountopts "comma-separated list"
```
    - ii. This command must be run on each server in the cluster.
3. Bring the file system resource in-service in LifeKeeper. This will mount the protected file system.

## 5.4.4.1.10.3. Maintaining a LifeKeeper Protected System

---

When performing shutdown and maintenance on a LifeKeeper-protected server, you must put that system's resource hierarchies in service on the backup server before performing maintenance. This process stops all activity for shared disks on the system needing maintenance.

Perform these actions in the order specified, where *Server A* is the primary system in need of maintenance and *Server B* is the backup server:

1. **Bring hierarchies in service on Server B.** On the backup, *Server B*, use the LifeKeeper GUI to bring in service any resource hierarchies that are currently in service on *Server A*. This will unmount any file systems currently mounted on *Server A* that reside on the shared disks under LifeKeeper protection. See [Bringing a Resource In Service](#) for instructions.
2. **Stop LifeKeeper on Server A.** Use the LifeKeeper command `/opt/LifeKeeper/bin/lkstop -f` to stop LifeKeeper. Your resources are now unprotected.
3. **Shut down Linux and power down Server A.** Shut down the Linux operating system on *Server A*, then power off the server.
4. **Perform maintenance.** Perform the necessary maintenance on *Server A*.
5. **Power on Server A and restart Linux.** Power on *Server A*, then reboot the Linux operating system.
6. **Start LifeKeeper on Server A.** Use the LifeKeeper command `/opt/LifeKeeper/bin/lkstart` to start LifeKeeper. Your resources are now protected.
7. **Bring hierarchies back in-service on Server A, if desired.** On *Server A*, use the LifeKeeper GUI to bring in service all resource hierarchies that were switched over to *Server B*.

## 5.4.4.1.10.4. Maintaining a Resource Hierarchy

---

You can perform maintenance on a resource hierarchy while maintaining LifeKeeper protection of all other hierarchies on the system. This involves taking the hierarchy in need of maintenance out of service and then bringing it back in-service after you complete the maintenance tasks.

To perform maintenance on a resource hierarchy:

1. **Take the hierarchy out of service.** Use the LifeKeeper GUI to take as much of the resource hierarchy out of service as you need to perform the maintenance. See [Taking a Resource Out of Service](#) for instructions.
2. **Perform maintenance.** Perform the necessary maintenance on the resource hierarchy.
3. **Restore the hierarchy.** Use the LifeKeeper GUI to bring the resource hierarchy back in service. See [Bringing a Resource In Service](#) for instructions.

## 5.4.4.1.10.5. Recovering After a Failover

---

After LifeKeeper performs a failover recovery from a primary server (*Server A*) to a backup server (*Server B*), perform the following steps:

1. **Review logs.** When LifeKeeper on *Server B* performs a failover recovery from *Server A*, status messages are displayed during the failover.

The exact output depends upon the configuration. Some messages on failure to mount or unmount are expected and do not suggest failure of recovery. These messages, as well as any errors that occur while bringing the resource in service on *Server B*, are logged in the LifeKeeper log.

2. **Perform maintenance.** Determine and fix the cause of the failure on *Server A*. *Server A* may need to be powered down to perform maintenance.
3. **Reboot *Server A*, if necessary.** Once maintenance is complete, reboot *Server A* if necessary.
4. **Start LifeKeeper, if necessary.** If LifeKeeper is not running on *Server A*, use the command `$LKROOT/bin/lkcli start` to start LifeKeeper.
5. **Move application back to *Server A*.** At a convenient time, use the LifeKeeper GUI to bring the application back into service on *Server A*. See [Bringing a Resource In Service](#) for instructions. Note that this step may be unnecessary if the application on *Server A* was configured for **Automatic Switchback**.

## 5.4.4.1.10.6. Removing LifeKeeper

---

You can uninstall LifeKeeper in a Linux environment via the command line by entering the following command.

```
/opt/LifeKeeper/bin/rmlk
```

This command uninstalls all the LifeKeeper packages and removes the directory `/opt/LifeKeeper` from the system. The command can be run with or without LifeKeeper running on the system. If the command is run with LifeKeeper running on all nodes in the cluster, then hierarchies are unextended and any comm paths are removed. This ensures that all remnants of the node on which the command was run are removed from the remaining nodes which effectively removes the node from the cluster. If LifeKeeper is not running at the time the command is executed, the other nodes will have remnants remaining which may impact the running system. You may delete these by running the following command on the other nodes in the cluster (without LifeKeeper running), it will complete the uninstall.

```
/opt/LifeKeeper/bin/rmlk -l
```

Optional argument:

- `-l` which remove all LifeKeeper licenses

**Use this command carefully.**

 **Note:** The periodic backup of the LifeKeeper configuration via the command `lkbackup` automatically archives the results in `/opt/LifeKeeper/config/`. Because the `rmlk` command will remove the `/opt/LifeKeeper` directory you may wish to back-up the archives before running the command.

## 5.4.4.1.10.7. Running LifeKeeper With a Firewall

---

LifeKeeper for Linux can work with a firewall in place on the same server if you address the following network access requirements.

 **Note:** If you wish to simply disable your firewall, see [Disabling a Firewall](#) below.

### LifeKeeper Communication Paths

Communication paths are established between pairs of servers within the LifeKeeper cluster using specific IP addresses. Although TCP Port 7365 is used by default on the remote side of each connection as it is being created, the TCP port on the initiating side of the connection is arbitrary. The recommended approach is to configure the firewall on each LifeKeeper server to allow both incoming and outgoing traffic for each specific pair of local and remote IP addresses in the communication paths known to that system.

### LifeKeeper GUI Connections

The LifeKeeper GUI uses a number of specific TCP ports, including Ports 81 and 82 as the default initial connection ports. The GUI also uses Remote Method Invocation (RMI), which uses Ports 1024 and above to send and receive objects. All of these ports must be open in the firewall on each LifeKeeper server to at least those external systems on which the GUI client will be run.

### LifeKeeper IP Address Resources

The firewall should be configured to allow access to any IP address resources in your LifeKeeper hierarchies from those client systems that need to access the application associated with the IP address. Remember that the IP address resource can move from one server to another in the LifeKeeper cluster; therefore, the firewalls on all of the LifeKeeper servers must be configured properly.

LifeKeeper also uses a broadcast ping test to periodically check the health of an IP address resource. This test involves sending a broadcast ping packet from the virtual IP address and waiting for the first response from any other system on the local subnet. To prevent this test from failing, the firewall on each LifeKeeper server should be configured to allow the following types of network activity.

- Outgoing Internet Control Message Protocol (ICMP) packets from the virtual IP address (so that the active LifeKeeper server can send broadcast pings)
- Incoming ICMP packets from the virtual IP address (so that other LifeKeeper servers can receive broadcast pings)
- Outgoing ICMP reply packets from any local address (so that other LifeKeeper servers can respond to broadcast pings)

- Incoming ICMP reply packets to the virtual IP address (so that the active LifeKeeper server can receive broadcast ping replies)

## LifeKeeper Data Replication

When using LifeKeeper Data Replication, the firewall should be configured to allow access to any of the ports used by nbd for replication. The ports used by nbd can be calculated using the following formula:

$$10001 + \text{<mirror number>} + \text{<256 * i>}$$

where *i* starts at zero and is incremented until the formula calculates a port number that is not in use. In use constitutes any port found defined in `/etc/services`, found in the output of `netstat -an --inet --inet6`, or already defined as in use by another LifeKeeper Data Replication resource.

**For example:** If the mirror number for the LifeKeeper Data Replication resource is 0, then the formula would initially calculate the port to use as 10001, but that number is defined in `/etc/services` on some Linux distributions as the SCP Configuration port. In this case, *i* is incremented by 1 resulting in Port Number 10257, which is not in `/etc/services` on these Linux distributions.

## Other Inter-node Communications

Each LifeKeeper server communicates using SSL connection on port 778. You can change this port using the configuration variable `API_SSL_PORT` in `/etc/default/LifeKeeper`.

## Disabling a Firewall

To disable the firewall, please follow the procedure described in the manual of your distribution.

## 5.4.4.1.10.8. Running the LifeKeeper GUI Through a Firewall

---

In some situations, a LifeKeeper cluster is placed behind a corporate firewall and administrators wish to run the LifeKeeper GUI from a remote system outside the firewall.

LifeKeeper uses Remote Method Invocation (RMI) to communicate between the GUI server and client. The RMI client must be able to make connections in each direction. Because the RMI client uses dynamic ports, you can not use preferential ports for the client.

One solution is to use ssh to tunnel through the firewall as follows:

1. Make sure your IT department has opened the secure shell port on the corporate firewall sufficiently to allow you to get behind the firewall. Often the machine IT allows you to get to is not actually a machine in your cluster but an intermediate one from which you can get into the cluster. This machine must be a Unix or Linux machine.
2. Make sure both the intermediate machine and the LifeKeeper server are running sshd (the secure shell daemon) and that X11 port forwarding is enabled (this is usually the line `X11Forwarding yes` in `/etc/ssh/sshd_config`, but if you are unsure, have your IT do this for you).
3. From your Unix client in X, tunnel to the intermediate machine using:

```
ssh -X -C <intermediate machine>
```

The `-C` means `compress the traffic` and is often useful when coming in over slower internet links.

4. From the intermediate machine, tunnel to the LifeKeeper server using:

```
ssh -X <LifeKeeper server>
```

You should not need to compress this time since the intermediate machine should have a reasonably high bandwidth connection to the LifeKeeper server.

5. If all has gone well, when you issue the command:

```
echo $DISPLAY
```

it should be set to something like `localhost:10.0`. If it is not set, it is likely that X11 forwarding is disabled in one of the sshd config files.

6. Verify that you can pop up a simple `xterm` from the LifeKeeper server by issuing the command:

```
/usr/X11R6/bin/xterm
```

7. If the `xterm` appears, you're ready to run **lkGUIapp** on the LifeKeeper server using the following command:

**`/opt/LifeKeeper/bin/lkGUIapp`**

8. Wait (and wait some more). Java uses a lot of graphics operations which take time to propagate over a slow link (even with compression), but the GUI console should eventually appear.

## 5.4.4.1.10.9. Transferring Resource Hierarchies

---

When you need to perform routine maintenance or other tasks on a LifeKeeper Server, you can use the LifeKeeper GUI to move in-service resources to another server. To transfer in-service resource hierarchies from *Server A* to *Server B*, use the GUI to bring the hierarchies into service on *Server B*. Repeat until all of *Server A*'s resources have been placed in-service on their respective backup servers. See [Bringing a Resource In Service](#) for instructions.

When all of *Server A*'s resources are active on their backup server(s), you can shut down *Server A* without affecting application processing. For the maintenance period, however, the resources may not have LifeKeeper protection depending on the number of servers in the cluster.

## 5.4.4.1.11. Technical Notes

\* We strongly recommend that you read the following technical notes concerning configuration and operational issues related to your LifeKeeper environment.

### LifeKeeper Features

Item	Description
<b>Licensing</b>	LifeKeeper requires unique runtime license keys for each server. This applies to both physical and virtual servers. A license key is required for the LifeKeeper core software, as well as for each separately packaged LifeKeeper recovery kit. The installation script installs a License Utilities package that obtains and displays the Host ID of your server during the initial install of LifeKeeper. Once your licenses have been installed the utility will return the Entitlement ID if it is available or the Host ID if it is not. The Host IDs, along with the Activation ID(s) provided with your software, are used to obtain license keys from the <b>SIOS Technology Corp. website</b> .
<b>Large Cluster Support</b>	LifeKeeper supports large cluster configurations, up to 32 servers. There are many factors other than LifeKeeper, however, that can affect the number of servers supported in a cluster. This includes items such as the storage interconnect and operating system or storage software limitations. Refer to the vendor-specific hardware and software configuration information to determine the maximum supported cluster size.
<b>Internationalization and Localization</b>	LifeKeeper for Linux v5.2 and later does support wide/multi-byte characters in resource and tag names but does not include native language message support. The LifeKeeper GUI can be localized by creating locale-specific versions of the Java property files, although currently only the English version is fully localized. However, many of the messages displayed by the LifeKeeper GUI come from the LifeKeeper core, so localization of the GUI will provide only a partial solution for users until the core software is fully localized.  See also <b>Language Environment Effects</b> in <a href="#">Known Issues and Restrictions</a> for additional information.
<b>LifeKeeper MIB File</b>	LifeKeeper can be configured to issue SNMP traps describing the events that are occurring within the LifeKeeper cluster. See the <code>lk_configsnmp(8)</code> man page for more information about configuring this capability. The MIB file describing the LifeKeeper traps can be found at <code>/opt/LifeKeeper/include/LifeKeeper-MIB.txt</code> .
<b>Watchdog</b>	LifeKeeper supports the watchdog feature. The feature was tested by SIOS Technology Corp. on Red Hat EL 5.5 64-bit, and Red Hat EL 6 + softdog.
<b>STONITH</b>	LifeKeeper supports the STONITH feature. This feature was tested by SIOS Technology Corp. on SLES 11 on IBM x3550 x86_64 architecture and RHEL5.5 64-bit.

<b>XFS File System</b>	The XFS file system does not use the fsck utility to check and fix a file system but instead relies on mount to replay the log. If there is a concern that there may be a consistency problem, the system administrator should unmount the file system by taking it out of service and run <code>xfs_check(8)</code> and <code>xfs_repair(8)</code> to resolve any issues.
<b>IPv6</b>	SIOS has migrated to the use of the <code>ip</code> command and away from the <code>ifconfig</code> command (for more information, see <a href="#">IPv6 Known Issue</a> ).

## Tuning

Item	Description												
<b>IPC Semaphores and IPC Shared Memory</b>	<p>LifeKeeper requires Inter-Process Communication (IPC) semaphores and IPC shared memory. The default Red Hat values for the following Linux kernel options are located in <code>/usr/include/linux/sem.h</code> and should be sufficient to support most LifeKeeper configurations.</p> <p><b>Note:</b> The required values are the minimum values for LifeKeeper. These values are for LifeKeeper only and should be adjusted based on other application semaphores requirements but should never fall below the LifeKeeper required values.</p> <table border="1" data-bbox="475 1032 1034 1317"> <thead> <tr> <th>Option</th> <th>Required</th> <th>Default Red Hat 7</th> </tr> </thead> <tbody> <tr> <td>SEMOPM</td> <td>14</td> <td>3</td> </tr> <tr> <td>SEMMNI</td> <td>25</td> <td>128</td> </tr> <tr> <td>SEMMSL</td> <td>20</td> <td>32000</td> </tr> </tbody> </table>	Option	Required	Default Red Hat 7	SEMOPM	14	3	SEMMNI	25	128	SEMMSL	20	32000
Option	Required	Default Red Hat 7											
SEMOPM	14	3											
SEMMNI	25	128											
SEMMSL	20	32000											
<b>System File Table</b>	<p>LifeKeeper requires that system resources be available in order to failover successfully to a backup system. For example, if the system file table is full, LifeKeeper may be unable to start new processes and perform a recovery. In kernels with enterprise patches, including those supported by LifeKeeper, <code>file-max</code>, the maximum number of open files in the system, is configured by default to 1/10 of the system memory size, which should be sufficient to support most LifeKeeper configurations. Configuring the <code>file-max</code> value lower than the default could result in unexpected LifeKeeper failures.</p> <p>The value of <code>file-max</code> may be obtained using the following command:</p> <pre>cat /proc/sys/fs/file-nr</pre> <p>This will return three numbers. The first represents the high watermark of file table entries (i.e. the maximum the system has seen so far); the second represents the current number of file table entries, and the third is the <code>file-max</code> value.</p>												

	<p>To adjust file-max, add (or alter) the “<i>fs,file-max</i>” value in <i>/etc/sysctl.conf</i> (see <i>sysctl.conf(5)</i> for the format) and then run</p> <pre>sysctl -p</pre> <p>to update the system. The value in <i>/etc/sysctl.conf</i> will persist across reboots.</p>
--	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## LifeKeeper Operations

Item	Description
<b>Kernel Debugger (kdb)</b>	Before using the Kernel Debugger ( <b>kdb</b> ) on a LifeKeeper protected server, you should first either shut off LifeKeeper on that server or switch any LifeKeeper protected resources over to the backup server. Use <code>lkcli stop -s</code> with the LifeKeeper SCSI Reservation Daemons ( <b>lkscsid</b> ) enabled (they are enabled by default) can also lead to unexpected panics.
<b>System Panic on Locked Shared Devices</b>	LifeKeeper uses a lock to protect shared data from being accessed by other servers on a shared SCSI device. If LifeKeeper cannot access a device as a result of another server taking the lock on a device, then a critical error has occurred and quick action should be taken or data can be corrupted. When this condition is detected, LifeKeeper enables a feature that will cause the system to panic.  If LifeKeeper stops the LifeKeeper daemons without removing resource(s) such as <code>lkcli stop -s</code> and shared devices still reserved, then the LifeKeeper locking mechanism may trigger a kernel panic when another server recovers the resource(s). All resources must be placed out-of-service before stopping LifeKeeper in this manner.
<b>nolock Option</b>	When using storage applications with locking and following recommendations for the NFS mount options, LifeKeeper requires the additional nolock option be set, e.g. <code>rw,nolock,bg,hard,nointr,tcp,nfsvers=3,timeo=600,rsz=32768,wsz=32768,actimeo=120</code>
<b>Recovering Out-of-Service Hierarchies</b>	As part of the recovery following the failure of a LifeKeeper server, resource hierarchies that were configured on the failed server but which were not <i>in-service</i> anywhere at the time of the server failure are recovered on the highest priority alive server at the time of the failure. This is the case no matter where the <i>out-of-service</i> hierarchy was last in service, including the failed server, the recovering server, or some other server in the cluster.
<b>Coexistence with Linux Firewalls</b>	The firewall is enabled upon installation. After installation is complete, the firewall should be modified to allow LifeKeeper traffic.  LifeKeeper will function if a host firewall is enabled. However, unless absolutely necessary, it is recommended that the firewall be disabled and that the LifeKeeper protected resources reside behind another shielding firewall.  If LifeKeeper must coexist on firewall enabled hosts, then the specific ports that LifeKeeper is using must be opened. Please note that LifeKeeper uses specific ports for communication paths, GUI, IP and Replication. Refer to <a href="#">Running LifeKeeper with a Firewall</a> for details.

	To disable or modify the firewall please refer to the documentation for your OS distribution.
<b>Coexistence with SELinux</b>	Disable SELinux. To Disable SELinux, please refer to the documentation for your OS distribution. AppArmor (for distributions that use this security model) may be enabled.
<b>Suid Mount Option</b>	The suid mount option is the default when mounting as <i>root</i> and is not written to the <i>/etc/mtab</i> by the <i>mount</i> command. The suid mount option is not needed in LifeKeeper environments.

## Server Configuration

Item	Description
<b>BIOS Updates</b>	The latest available BIOS should always be installed on all LifeKeeper servers.

## LifeKeeper Version 8.2.0 and Later GUI Requirement

64-bit versions of any PAM related packages will be required for the LifeKeeper GUI Client to successfully authenticate users.

## Confirm Failover and Block Resource Failover Settings

Make sure you review and understand the following descriptions, examples and considerations before setting the **Confirm Failover** or **Block Resource Failover** in your LifeKeeper environment. These settings are available from the command line or on the **Properties** panel in the LifeKeeper GUI.

### Confirm Failover On:

**Definition** – Enables manual failover confirmation from *System A* to *System B* (where *System A* is the server whose properties are being displayed in the [Properties Panel](#) and *System B* is the system to the left of the checkbox). If this option is set on a system, it will require a manual confirmation by a system administrator before allowing LifeKeeper to perform a failover recovery of a system that it detects as failed.

Use the `lk_confirmso` command to confirm the failover. By default, the administrator has 10 minutes to run this command. This time can be changed by modifying the **CONFIRMSOTO** setting in `/etc/default/LifeKeeper`. If the administrator does not run the `lk_confirmso` command within the time allowed, the failover will either proceed or be blocked. By default, the failover will proceed. This behavior can be changed by modifying the **CONFIRMSODEF** setting in `/etc/default/LifeKeeper`.

**Example:** If you wish to block automatic failovers completely, then you should set the **Confirm Failover On** option in the **Properties** panel and also set **CONFIRMSODEF** to **1** (block failover) and **CONFIRMSOTO** to **0** (do not wait to decide on the failover action).

### When to select this setting:

This setting is used in most Disaster Recovery and other WAN configurations where the configuration does not include redundant heartbeat communications paths.

Open the **Properties** page from one server and then select the server that you want the **Confirm Failover flag** to be set on.

### Set Block Resource Failover On:

**Definition** – By default, all resource failures will result in a recover event that will attempt to recover the failed resource on the local system. If local recovery fails or is not enabled, then LifeKeeper transfers the resource hierarchy to the next highest priority system for which the resource is defined. However, if this setting is selected on a designated system(s), all resource transfers due to a resource failure will be blocked from the given system.

When the setting is enabled, the following message is logged:

Local recovery failure, failover blocked, MANUAL INTERVENTION REQUIRED

## 5.4.4.2. Cluster Example

### Expanded Multicluster Example

Hierarchies	 pat	 mike	 wallace	 gromit	 batman	 bullwinkle
 Backup Not StandB						
 file_system_2	 40 St...	 50 St...	 1 Acti...	 10 St...	 60 St...	 70 U...
 device-nfs180	 40 St...	 50 St...	 1 Acti...	 10 St...	 60 St...	 70 U...

## 5.4.4.3. Dialogs

---

[In Service Dialog](#)

[Out-of-Service Resource Properties – EquivalenciesResource Dialog](#)

[Password Dialog](#)

[Resource Properties – Equivalencies](#)

[Resource Properties – General](#)

[Resource Properties – Relations](#)

[Server Properties – Commpath](#)

[Server Properties – General](#)

[Server Properties – Resource](#)

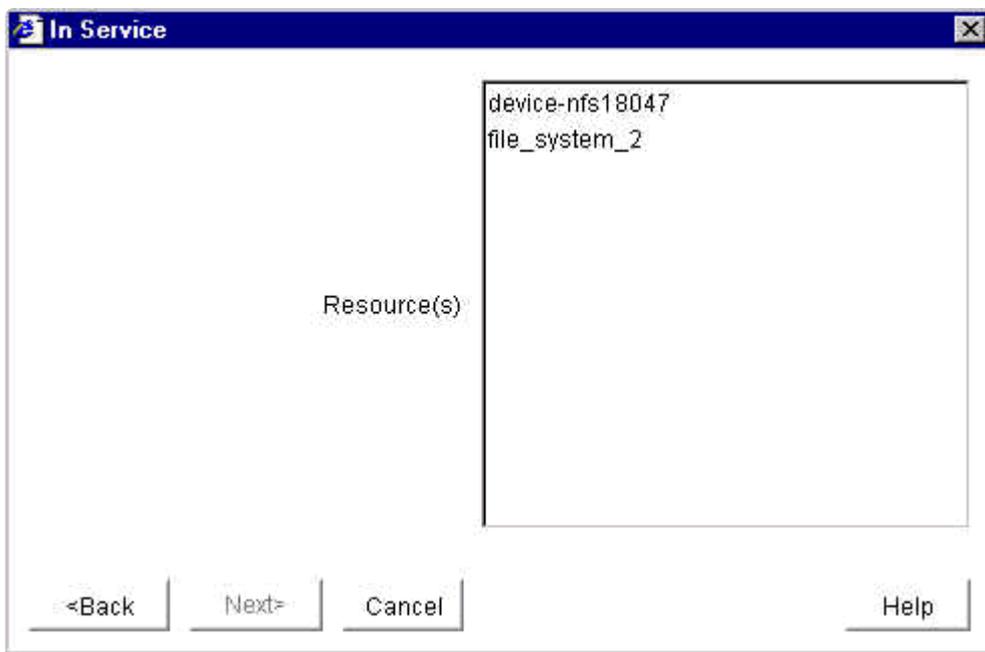
# In Service Dialog

**Select a Server.** The first dialog provides a drop-down list box containing the names of servers in your LifeKeeper cluster. Select the **Server** where you want to bring the resource instance into service. Click on the **Next** button to proceed to the next dialog.



\* **Note:** If you initiated the In Service task by right-clicking from the right pane on a server-specific resource, this dialog and the next will not appear since you will have already specified the server and the resource that you want to bring into service.

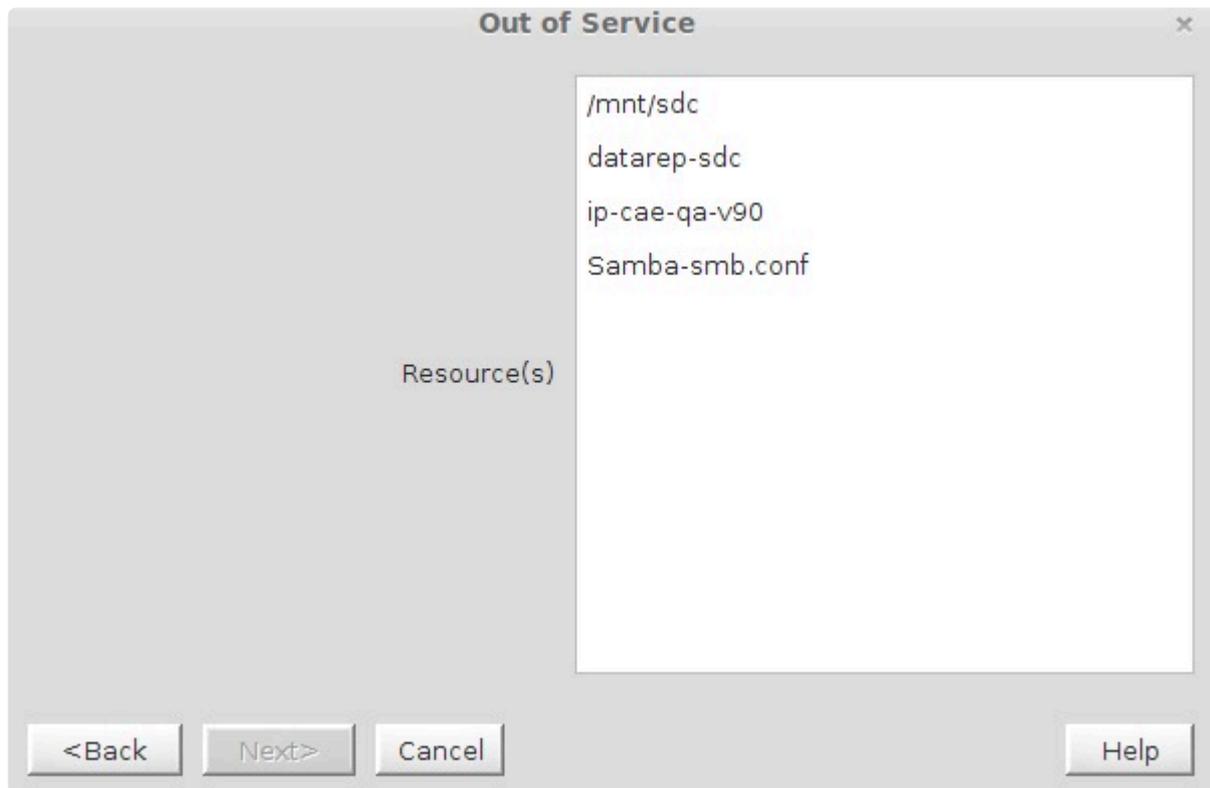
**Select a Resource.** The second dialog provides a drop-down list box containing the names of all the available resources on the server you selected in the previous dialog. Select the resource that you want to bring into service.



\* **Note:** If you initiated the In Service task by right-clicking from the left pane on a global resource, this dialog will not appear since you will have already identified the resource that you want to bring into service.

# Out-of-Service Resource Dialog

**Select a Resource.** This dialog provides a list containing the names of all resources that are in service in your LifeKeeper cluster. Select the resource that you want to take out of service.

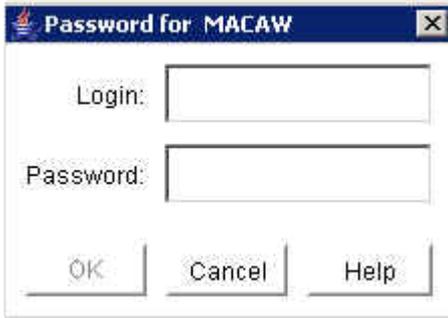


\* **Note:** If you initiated the Out-of-Service task by right-clicking from the left pane on an in-service global resource or from the right pane on an in-service server-specific resource instance, this dialog will not appear since you will have already specified the resource that you want to take out of service.

# Password Dialog

---

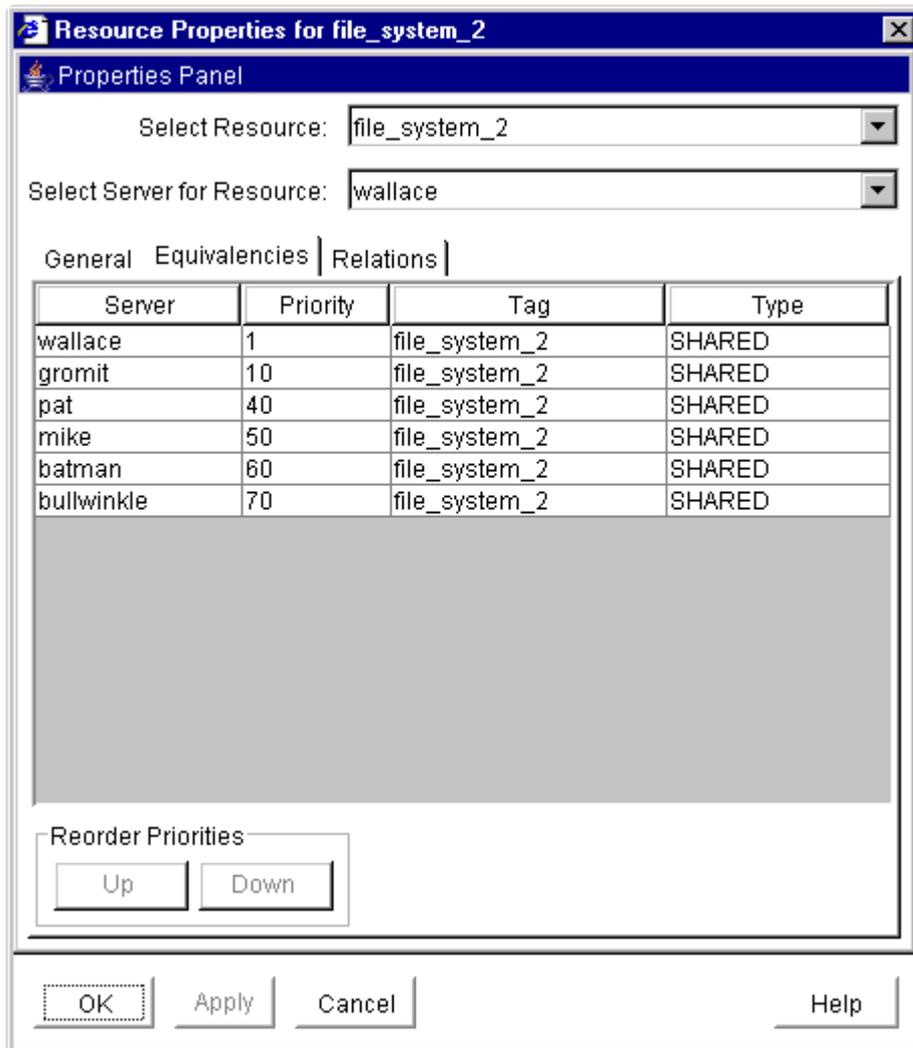
This dialog is displayed if the initial login name or password entered in the [Cluster Connect dialog](#) is invalid.



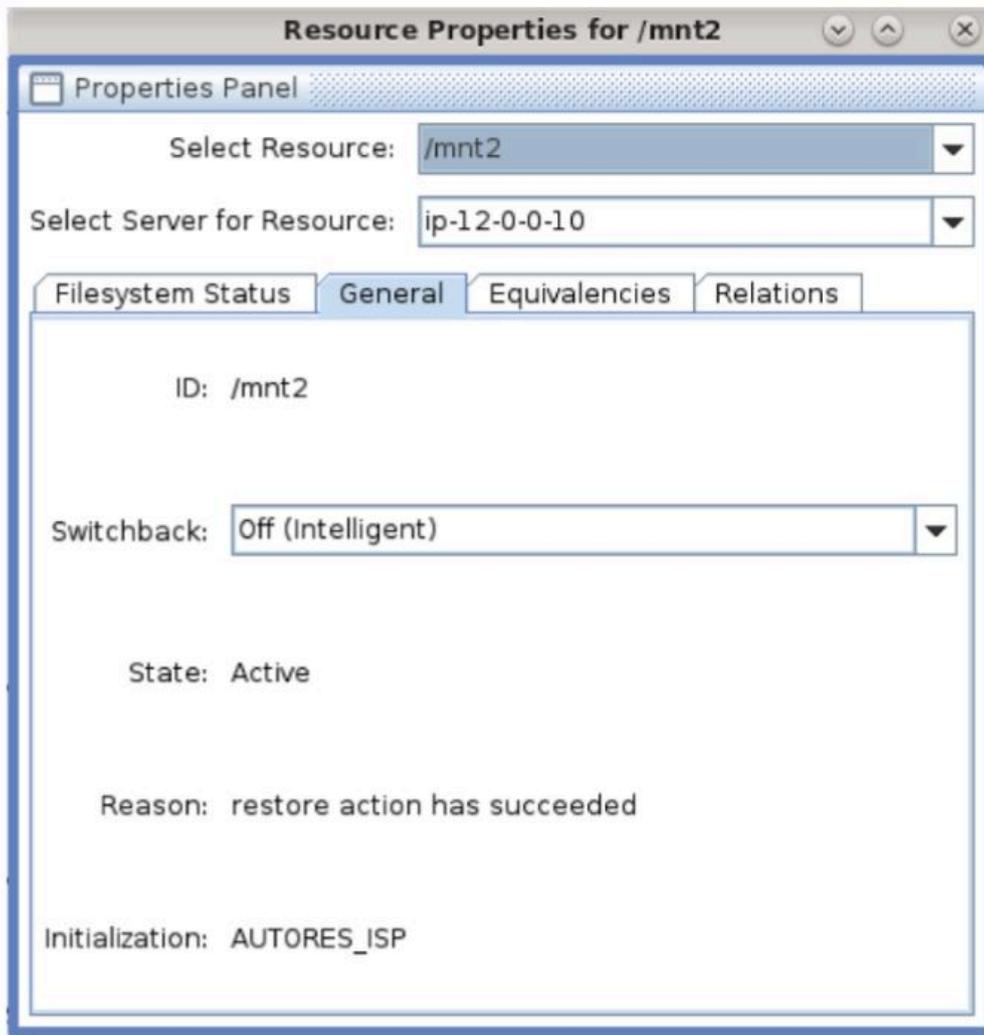
**Login.** The login name of a user with LifeKeeper authorization on the specified server.

**Password.** The password that authorizes the specified login on the server.

# Resource Properties – Equivalencies



# Resource Properties – General



# Resource Properties – Relations

Resource Properties for file\_system\_2

Properties Panel

Select Resource: file\_system\_2

Select Server for Resource: wallace

General | Equivalencies | Relations |

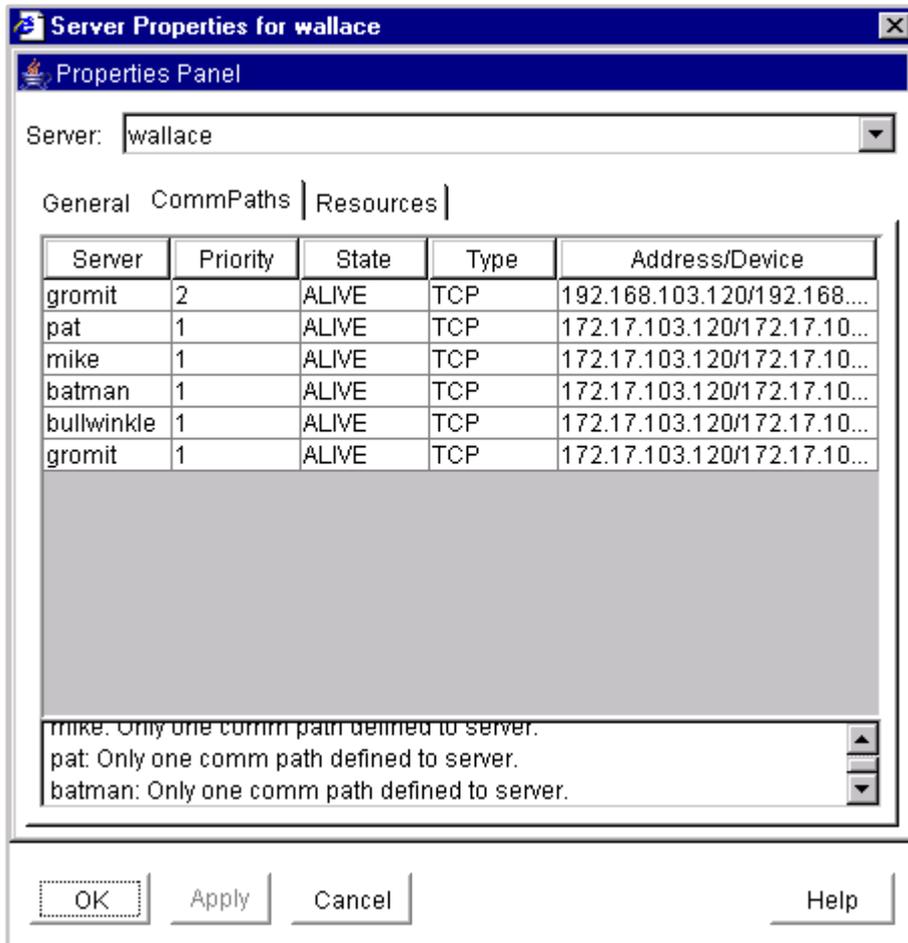
Root:

Parent:

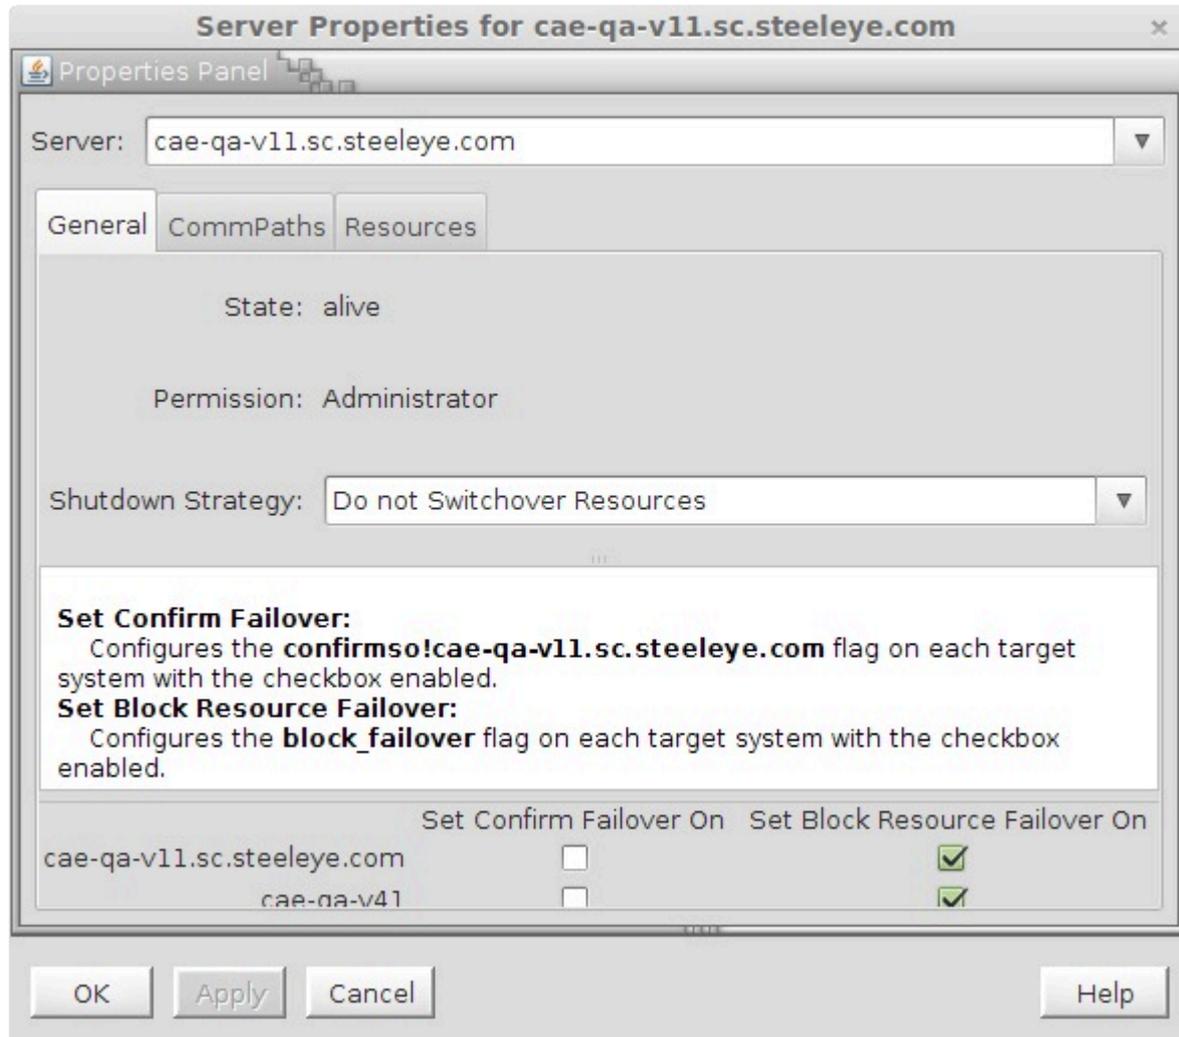
Child: device-nfs18047

OK Apply Cancel Help

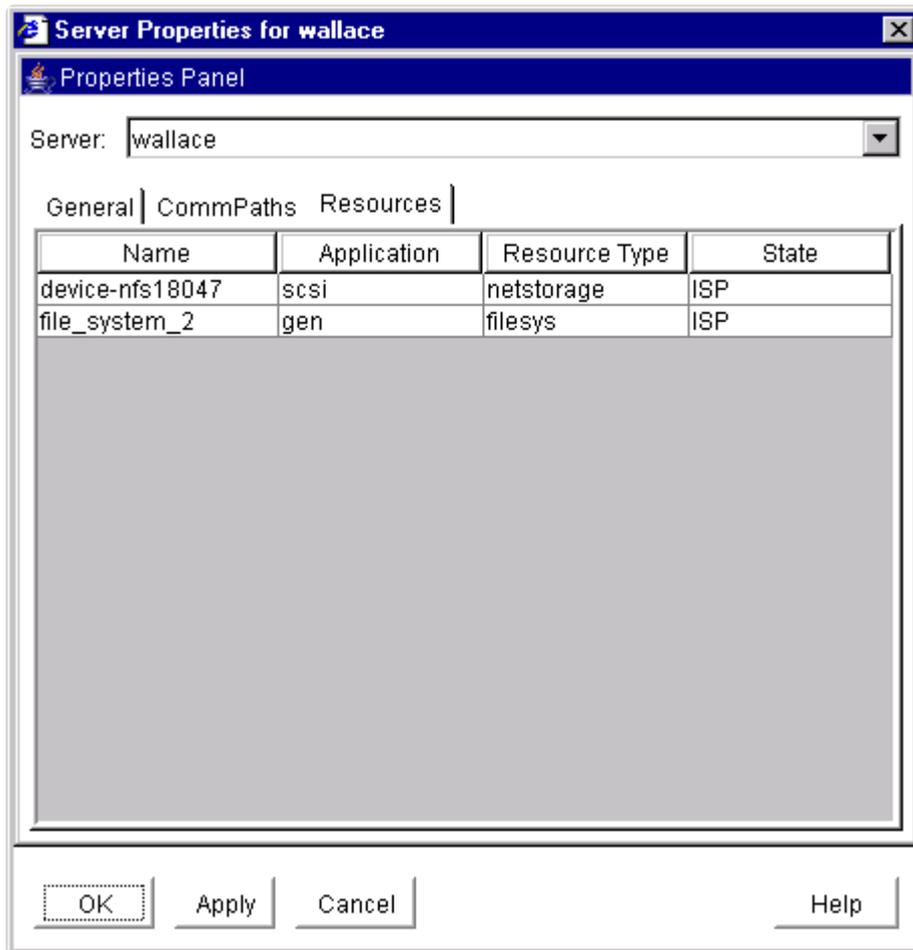
# Server Properties – Commpath



# Server Properties – General



# Server Properties – Resource



## 5.4.5. Troubleshooting

---

The [Message Catalog](#) provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the following individual Message Catalogs:

[Core Message Catalog](#)

[DB2 Kit Message Catalog](#)

[DMMP Kit Message Catalog](#)

[Recovery Kit for EC2 Message Catalog](#)

[File System Kit Message Catalog](#)

[Gen/App Kit Message Catalog](#)

[IP Kit Message Catalog](#)

[Oracle Listener Kit Message Catalog](#)

[Oracle Kit Message Catalog](#)

[SCSI Kit Message Catalog](#)

[DataKeeper Kit Message Catalog](#)

[Quick Service Protection Kit Message Catalog](#)

[GUI Message Catalog](#)

In addition to utilizing the Message Catalog described above, the following topics detail troubleshooting issues, restrictions, etc., that may be encountered:

---

[Common Causes of Failover](#)

[Known Issues and Restrictions](#)

[GUI Troubleshooting](#)

[Communication Paths Going Up and Down](#)

[Incomplete Resource Created](#)

[Incomplete Resource Priority Modification](#)

[No Shared Storage Found when Configuring a Hierarchy](#)

[Recovering from a Server Failure](#)

[Recovering from a Non-Killable Process](#)

[Recovering from a Panic During a Manual Recovery](#)

[Recovering Out-of-Service Hierarchies](#)

[Resource Tag Name Restrictions](#)

[Serial \(TTY\) Console Warning](#)

[Taking the System to INIT State S WARNING](#)

[Thread is Hung Messages on Shared Storage](#)

## 5.4.5.1. Solutions

---

### - SAP/Oracle Patching -

#### Patching Oracle nodes (SAP/Oracle) with DataKeeper

##### Solution Details

Most of the Oracle updates / patches only require access to the system tables drives and not to the data drives (protected DataKeeper data drives).

The following procedure can be used in general:

Prior to the procedure, set block failover on primary and confirm failover on target.

1. On the standby / target server:
  - a. Apply the appropriate upgrades, patches, etc.
  - b. Reboot your server
2. Then, switchover your resources from the source / active server to the standby / target server:
  - a. Verify that you can access the Oracle database
3. On the other node (which now is a standby after the switchover and the original source / active server):
  - a. Apply the appropriate upgrades, patches, etc.
  - b. Reboot your server
4. (Optional) Switchover again to verify connectivity and access to Oracle on the original / source /active server

In some cases where the patches require access to the DataKeeper volumes (e.g. running catsbp), you need to pause / unlock the mirrors to perform the upgrades on the standby / target system.

Verify that Oracle can be started on the backup, then stop Oracle on the backup node and continue the mirrors. Then repeat the upgrade on the source system.

At the end of the procedure remove the flags 'block failover' on primary and the 'confirm failover' on target.

### - Storage Quorum -

#### Heartbeat recommendations for using Storage Quorum

##### Solution Details

##### ISSUE:

[What are SIOS Heartbeat recommendations for using Storage Quorum?](#)

##### SOLUTION:

In order to use storage quorum, SIOS recommends that you increase the `LCMHEARTBEATS` to allow for a longer time before the path is marked as failed. This will change the timeout period from the default of 15 seconds to 45 seconds.

Edit the `/etc/default/LifeKeeper` file and change `LCMNUMHBEATS` to 9:

```
LCMNUMHBEATS=9
```

Since you will be making a change to a LifeKeeper core parameter, you will need to recycle LifeKeeper.

To minimize downtime, you can use `lkstop -f` that will leave the resources running. While LifeKeeper is stopped, failures of protected resources will not be detected or acted upon.

```
# lkstop -f
```

```
# lkstart
```

 These changes should be made on each node in the cluster.

## Storage quorum failed to prevent failover when communication between nodes is lost

### Solution Details

#### ISSUE:

Incomplete storage quorum configuration caused failures during lost communication processing

#### SOLUTION:

All comm paths between cluster nodes must be created and **“ALIVE”** before running `qwk_storage_init` on each node in the cluster.

If this is not the case, execute the following commands to reinitialize the storage quorum configuration **once all comm paths are ALIVE**.

```
# /opt/LifeKeeper/bin/qwk_storage_exit
```

```
# /opt/LifeKeeper/bin/qwk_storage_init
```

 This will reinitialize storage quorum.

## How do you tune parameters for storage quorum when using Amazon S3 storage?

### Solution Details

**ISSUE:**

How do you tune parameters for storage quorum when using Amazon S3 storage?

**SOLUTION:**

Once you set up storage quorum using the documentation, there may be questions on how to tune the heart beat parameters.

[Click here for more information.](#)

Here are **several things** that you can do to help determine if the default values are sufficient in your environment:

There are 2 main parameters in `/etc/default/LifeKeeper` that affect the storage quorum timeout value::

- `QWK_STORAGE_HBEATTIME` (default is 6) – Specifies the **interval in seconds** between reading and writing the QWK objects.
- `QWK_STORAGE_NUMHBEATS` (default is 4) – Specifies the **number of consecutive heartbeat checks** that, when missed, indicates that the target node has failed. A **missed heartbeat** occurs when the QWK object has not been updated since the last check.

When using an Amazon S3 bucket to store the QWK objects (i.e., `QWK_STORAGE_TYPE=aws_s3`), **SIOS suggests** running the following commands to ensure **good connectivity** in your environment:

1. Execute **ping s3.amazonaws.com** and make sure the time is under a second. This **ensures good connectivity** from the EC2 node to the global AWS domain.

 **Note:** Even though S3 is a global service, the S3 buckets are located in a specific region.

2. Execute **ping <bucketname>.s3.amazonaws.com**, which will resolve to the IP address of the hosting S3 service. This should also be less than a second.

Another thing to consider is the **amount of data** being transferred for overall S3 activities on this node. It is possible that file transfers are taking place. You may measure the response time using ping, as in the examples above. (See above [ping format](#)).

 **Note:** In AWS, S3 is a global service but the EC2 nodes are regional. The S3 URI (Amazon S3 storage Uniform Resource Identifier or the web address) **can be in a specific region**.

By comparing network responsiveness in both high-traffic and low-traffic situations, you can tune the number of missed heartbeats (`QWK_STORAGE_NUMHBEATS`) and the heartbeat time

(QWK\_STORAGE\_HBEATTIME) mentioned above.

The [parameters](#) above must be specified in `/etc/default/LifeKeeper` before running `qwk_storage_init`.

In most cases where your ping to the Amazon S3 service resolves to less than a second, the default is sufficient. However, if the connection to the Amazon S3 service is slow or you see degradation, we **recommend** increasing the `QWK_STORAGE_HBEATTIME` (see parameter above) from 6 to 7. This will increase the loss timeout from 24 seconds ( $6 \times 4$ ) to 28 seconds ( $7 \times 4$ ). SIOS **does not** recommend increasing the timeout much larger than 30 seconds.

If you change the default settings, be sure to change them **on each system in the cluster** and reinitialize storage quorum by executing the following commands on each system in the cluster while all comm paths are **“ALIVE”**:

```
# /opt/LifeKeeper/bin/qwk_storage_exit

# /opt/LifeKeeper/bin/qwk_storage_init
```

## - Majority Mode -

### Shutdown and restart procedure for LifeKeeper with a witness node

1. **Stop the backup (target) server** using `lkstop`

Wait for `lkstop` command to complete. Verify `lkstop` is complete by log entry:

```
NOTIFY:shutdown:::010055:LifeKeeper stopped
```

Quorum message:

```
NOTIFY:event.comm_down:::010469:We do have quorum on comm_down, continuing
```

Typically this is less than 2 minutes

2. **Stop the primary (source) server** using `lkstop`.

Wait for `lkstop` command to complete. Verify `lkstop` is complete by log entry:

```
NOTIFY:shutdown:::010055:LifeKeeper stopped
```

Quorum message:

```
NOTIFY:event.comm_down:::010469:We do have quorum on comm_down, continuing
```

Typically this is less than 2 minutes

3. **Stop LifeKeeper on the witness node** using `lkstop`.

Wait for `lkstop` command to complete. Verify `lkstop` is complete by log entry:

```
NOTIFY:shutdown:::010055:LifeKeeper stopped
```

Typically this is less than 2 minutes

#### 4. Bring the witness server up using `lkstart`.

Wait for `lkstart` command to complete. Verify `lkstart` is complete by log entry:

```
INFO:event.lcm_avail:::010479:RESOURCE INITIALIZATION FINISHED
```

Typically this is less than 2 minutes

#### 5. Bring the backup (target) server up using `lkstart`.

Wait for `lkstart` command to complete. Verify `lkstart` is complete by log entry:

```
INFO:event.lcm_avail:::010479:RESOURCE INITIALIZATION FINISHED
```

Quorum message:

```
INFO:event.comm_up:::010490:We do have quorum on comm_up to , putting resources into service if needed
```

Typically this is less than 2 minutes

#### 6. Bring the primary (source) server up using `lkstart`.

Wait for `lkstart` command to complete. Verify `lkstart` is complete by log entry:

```
INFO:event.lcm_avail:::010479:RESOURCE INITIALIZATION FINISHED
```

Quorum message:

```
INFO:event.comm_up:::010490:We do have quorum on comm_up to <node>, putting resources into service if needed
```

Typically this is less than 2 minutes

**Note:** By taking the servers down in this order, the resources will stay in-service on the primary server (not prompt a failover). Since the mirrors were not in-service on the backup server when it went down they will not go into service when the server comes back up. They will stay in the out-of-service (OSU) state.

**Note:** When LifeKeeper starts on the primary server, LK will bring the mirrors into service because they were ISP when LifeKeeper went down and they are not in-service on the backup server. LifeKeeper will determine the backup server is ready and will start the resync from the primary

server to the backup server.

! If you do NOT bring up the systems in the following order, you could end up with data\_corrupt flags that require user intervention.

## - SAP HANA -

### Instructions for patching the HANA database

#### Solution Details

#### ISSUE:

**What are the step by step instructions for patching the HANA database?**

#### SOLUTION:

**Node 1** = source

**Node 2** = backup

1. **Disable** quickCheck (set LKCHECKINTERVAL=0, killall lkcheck) on node 1
2. **Stop** LifeKeeper using "lkstop -f" on node 2
  - a. **Verify** that HSR is running from node 1 to 2
3. **Stop** HANA on node 2, patch
4. **Start** HANA on node 2
  - a. **Verify** that HSR is still running
5. **Start** LifeKeeper using "lkstart" on node 2
6. **Switchover** resources to node 2
7. **Repeat** steps above to patch node 1
8. **Switch** to node 1 if necessary
9. **Re-enable** quickCheck (set LKCHECKINTERVAL to previous value)

## - Quickcheck for mirror is constantly failing and recovering -

### Quickcheck for mirror is constantly failing and recovering

#### Solution Details

#### ISSUE:

**Looking at the lifekeeper.log, you can see that the mirrors are constantly failing the quickcheck, but the recover always works.**

Sample messages:

```
NOTIFY:lcd.remain:recover:datarep-data:011115:BEGIN recover of "datarep-data" (class=netraid
event=recover)
```

```
INFO:dr:recover:datarep-data:104008:/dev/md0: merging bitmap from target "SV-GCS-LIVEB"*
```

```
*Oct 5 05:57:07 SP-GCS-LIVEA recover [11754]: INFO:dr:recover:datarep-data:104009:/dev/md0:
```

bitmap merged, resyncing 2.3

\*Oct 5 05:57:12 SP-GCS-LIVEA recover [11754]:

INFO:dr:recover:datarep-data:104095:Partial resynchronization of component “/dev/nbd1” has begun for mirror “/dev/md0”\*

This usually coincides nbd with errors in the message logs:

**Oct 5 05:56:59 SP-GCS-LIVEA kernel: [5499433.410998] nbd (pid 7278: nbd-client) got signal 9**

**Oct 5 05:56:59 SP-GCS-LIVEA kernel: [5499433.411003] nbd1: shutting down socket**

**Oct 5 05:56:59 SP-GCS-LIVEA kernel: [5499433.411015] nbd1: Receive control failed (result -4)**

**Oct 5 05:56:59 SP-GCS-LIVEA kernel: [5499433.411039] nbd1: queue cleared**

**Oct 5 05:57:11 SP-GCS-LIVEA nbd-client: Begin Negotiation**

**Oct 5 05:57:11 SP-GCS-LIVEA nbd-client: size = 268434407424**

**Oct 5 05:57:11 SP-GCS-LIVEA nbd... (truncated, see original email for full text)**

## **SOLUTION:**

These messages seem to indicate nbd issues that are causing the replication connections to drop. The resync recoveries in the LifeKeeper logs are the reaction to the connection being dropped and the need to re-establish them.

When a mirror is created, its state is monitored via quickCheck and via a mdadm process on the source. The quickCheck process checks a number of items and will issue a recovery event if it finds the mirror out of sync (from /proc/mdstat info), it can't ping the target or the target state is not alive, or if it finds that nbd-client/nbd-server processes are not running.

A recovery will also be initiated based on events from the md driver via the mdadm monitoring process. These include Fail, FailSpare and DegradedArray events (these are documented in the mdadm man page).

These events indicate issues that occurred that the md driver detected and LifeKeeper must react to so that it has the same state as the driver. Additionally, if the comm path over which the mirror is using goes down this will also lead to a recovery event.

There is one tuning option:

Increase the NBD\_XMIT\_TIMEOUT parameter. The default value for NBD\_XMIT\_TIMEOUT is 6 seconds.

Keep in mind that you do not want to raise this value by much to ensure a true hang condition on packet transmissions is detected and an abort is done to reset the connection and restart mirroring. Waiting too long could lead to hung writes on the source and eventually a hung system.

## 5.4.5.2. Common Causes of a LifeKeeper Initiated Failover

---

In the event of a failure, LifeKeeper has two methods of recovery: local recovery and inter-server recovery. If local recovery fails, a “failover” is implemented. A failover is defined as automatic switching to a backup server upon the failure or abnormal termination of the previously active application, server, system, hardware component or network. Failover and switchover are essentially the same operation, except that failover is automatic and [usually operates without warning](#), while switchover requires human intervention. This automatic failover can occur for a number of reasons. Below is a list of the most common examples of a LifeKeeper initiated failover.

### Server Level Causes

#### Server Failure

LifeKeeper has a built-in heartbeat signal that periodically notifies each server in the configuration that its paired server is operating. A failure is detected if a server fails to receive the heartbeat message.

- Primary server loses power or is turned off.
- CPU Usage caused by excessive load — Under very heavy I/O loads, delays and low memory conditions can cause system to become unresponsive such that LifeKeeper may detect a server as down and initiate a failover.
- Quorum/Witness – As part of the I/O fencing mechanism of quorum/witness, when a primary server loses quorum, a “[fastboot](#)”, “[fastkill](#)” or “[osu](#)” is performed (*based on settings*) and a failover is initiated. When determining when to fail over, the witness server allows resources to be brought in service on a backup server only in cases where it verifies the primary server has failed and is no longer part of the cluster. This will prevent failovers from happening due to simple communication failures between nodes when those failures don’t affect the overall access to, and performance of, the in-service node.

Relevant Topics
<a href="#">Supported Storage List</a>
<a href="#">Server Failure Recovery Scenario</a>
<a href="#">Tuning the LifeKeeper Heartbeat</a>
<a href="#">Quorum/Witness</a>

#### Communication Failures/Network Failures

LifeKeeper sends the heartbeat between servers every five seconds. If a communication problem causes the heartbeat to skip two beats but it resumes on the third heartbeat, LifeKeeper takes no action. However, if the communication path remains dead for three beats, LifeKeeper will label that

communication path as dead but will initiate a failover only if the redundant communication path is also dead.

- Network connection to the primary server is lost.
- Network latency.
- Heavy network traffic on a TCP comm path can result in unexpected behavior, including false failovers and LifeKeeper initialization problems.
- Using STONITH, when LifeKeeper detects a communication failure with a node, that node will be powered off and a failover will occur.
- Failed NIC.
- Failed network switch.
- Manually pulling/removing network connectivity.

Relevant Topics
<a href="#">Creating a Communication Path</a>
<a href="#">Tuning the LifeKeeper Heartbeat</a>
<a href="#">Network Configuration</a>
<a href="#">Verifying Network Configuration</a>
<a href="#">LifeKeeper Event Forwarding via SNMP</a>
<a href="#">Network-Related Troubleshooting</a>
<a href="#">Running LifeKeeper With a Firewall</a>
<a href="#">STONITH</a>

## Split-Brain

If a single comm path is used and the comm path fails, then LifeKeeper hierarchies may try to come into service on multiple systems simultaneously. This is known as a false failover or a “split-brain” scenario. In the “**split-brain**” scenario, each server believes it is in control of the application and thus may try to access and write data to the shared storage device. To resolve the split-brain scenario, LifeKeeper may cause servers to be powered off or rebooted or leave hierarchies out-of-service to assure data integrity on all shared data. Additionally, heavy network traffic on a TCP comm path can result in unexpected behavior, including false failovers and the failure of LifeKeeper to initialize properly.

The following are scenarios that can cause split-brain:

- Any of the comm failures listed above

- Improper shutdown of LifeKeeper
- Server resource starvation
- Losing all network paths
- DNS or other network glitch
- System lockup/thaw

## Resource Level Causes

LifeKeeper is designed to monitor individual applications and groups of related applications, periodically performing local recoveries or notifications when protected applications fail. Related applications, by example, are hierarchies where the primary application depends on lower-level storage or network resources. LifeKeeper monitors the status and health of these protected resources. If the resource is determined to be in a failed state, an attempt will be made to restore the resource or application on the current system (in-service node) without external intervention. If this local recovery fails, a resource failover will be initiated.

## Application Failure

- An application failure is detected, but the local recovery process fails.
- Remove Failure – During the resource failover process, certain resources need to be removed from service on the primary server and then brought into service on the selected backup server to provide full functionality of the critical applications. **If this remove process fails, a reboot of the primary server will be performed** resulting in a complete server failover.

Examples of remove failures:

- Unable to unmount file system
- Unable to shut down protected application (oracle, mysql, postgres, etc)

Relevant Topics
<a href="#">File System Health Monitoring</a>
<a href="#">Resource Error Recovery Scenario</a>

## File System

- Disk Full — LifeKeeper's File System Health Monitoring can detect disk full file system conditions which may result in failover of the file system resource.
- Unmounted or Improperly Mounted File System — User manually unmounts or changes options on an in-service and LK protected file system.
- Remount Failure — The following is a list of common causes for remount failure which would lead to a failover:

- corrupted file system (fsck failure)
- failure to create mount point directory
- mount point is busy
- mount failure
- LifeKeeper internal error

Relevant Topics
-----------------

<a href="#">File System Health Monitoring</a>
-----------------------------------------------

## IP Address Failure

When a failure of an IP address is detected by the IP Recovery Kit, the resulting failure triggers the execution of the IP local recovery script. LifeKeeper first attempts to bring the IP address back in service on the current network interface. If the local recovery attempt fails, LifeKeeper will perform a failover of the IP address and all dependent resources to a backup server. During failover, the remove process will un-configure the IP address on the current server so that it can be configured on the backup server.

**Failure of this remove process will cause the system to reboot.**

- IP conflict
- IP collision
- DNS resolution failure
- NIC or Switch Failures

Relevant Topics
-----------------

<a href="#">Creating Switchable IP Address</a>
------------------------------------------------

<a href="#">IP Local Recovery</a>
-----------------------------------

## Reservation Conflict

- A reservation to a protected device is lost or stolen
- Unable to regain reservation or control of a protected resource device (caused by manual user intervention, HBA or switch failure)

Relevant Topics
-----------------

<a href="#">SCSI Reservations</a>
-----------------------------------

<a href="#">Disabling Reservations</a>
----------------------------------------

## SCSI Device

- Protected SCSI device could not be opened. The device may be failing or may have been removed from the system.

## 5.4.5.3. Known Issues and Restrictions

---

Included below are the restrictions or known issues open against LifeKeeper for Linux, broken down by functional area.

[Installation Known Issues / Restrictions](#)

[LifeKeeper Core Known Issues / Restrictions](#)

[Internet/IP Licensing Known Issues / Restrictions](#)

[GUI Known Issues / Restrictions](#)

[Data Replication Known Issues / Restrictions](#)

[IPv6 Recovery Kit Known Issues / Restrictions](#)

[Apache Recovery Kit Known Issues / Restrictions](#)

[Oracle Recovery Kit Known Issues / Restrictions](#)

[MySQL Recovery Kit Known Issues / Restrictions](#)

[NFS Server Recovery Kit Known Issues / Restrictions](#)

[SAP Recovery Kit Known Issues / Restrictions](#)

[LVM Recovery Kit Known Issues / Restrictions](#)

[Multipath Recovery Kits \(DMMP / HDLM / PPATH / NECSPS\) Known Issues / Restrictions](#)

[DMMP Recovery Kit Known Issues / Restrictions](#)

[DB2 Recovery Kit Known Issues / Restrictions](#)

[Sybase ASE Recovery Kit Known Issues / Restrictions](#)

[WebSphere MQ Recovery Kit Known Issues / Restrictions](#)

[MaxDB Known Issues / Restrictions](#)

[EC2 Recovery Kit Known Issues / Restrictions](#)

[Known Issues/Restrictions when using LifeKeeper on Oracle Cloud Infrastructure](#)

## 5.4.5.3.1. Installation – Known Issues / Restrictions

Description
<p><b>A functional yum or zypper configuration is required for the successful installation of LfieKeeper.</b> A misconfigured or non-functional yum or zipper configuration can result in the failure of the LifeKeeper installation script. Output such as the following may be seen:</p> <pre>Install LifeKeeper and dependent packages done. Setup high availability data replication features.. done. Setup NFS high availability features... Configure LifeKeeper management group  Setup failed. Fix the problem and try again. sed: can't read /etc/default/LifeKeeper: No such file or directory /tmp/mnt/setuplibs/install.sh: line 118: /config/system_information: No such file or directory /tmp/mnt/setuplibs/install.sh: line 122: /config/system_information: No such file or directory /tmp/mnt/setuplibs/install.sh: line 122: /config/system_information: No such file or directory /tmp/mnt/setuplibs/install.sh: line 122: /config/system_information: No such file or directory</pre> <p>Additionally, the following may be seen in the /var/log/LK_install.log file:</p> <pre>Abort, retry, ignore? [a/r/i/? shows all options] (a): a done.</pre>
<p><b>In Release 7.4 and forward, relocation of the SIOS product RPM packages is no longer supported.</b></p>
<p><b>Linux Dependencies</b></p> <p>Installing LifeKeeper for Linux including the optional Recovery Kits requires several packages which have dependencies. If the package manager is properly configured, the required package is automatically installed by the package manager.</p> <p>If the installation can not be done automatically, the setup script will be interrupted. After manually installing dependent packages (see Linux dependencies for details), re-execute the setup script.</p> <p><b>Note:</b> If the installation of these dependent packages is not completed successfully, it could affect the ability to start LifeKeeper for Linux as well as the loading of the LifeKeeper for Linux GUI.</p>

**The multipathd daemon will log errors in the error log when the nbd driver is loaded as it tries to scan the new devices**

**Solution:** To avoid these errors in the log, add **devnode “^nbd”** to the blacklist in */etc/multipath.conf*.

### **mksh conflicts with LifeKeeper for Linux setup needing ksh**

If the `mksh` package is installed, the LifeKeeper for Linux setup will fail indicating a package conflict. The LifeKeeper for Linux requires the `ksh` package.

**Workaround:** On RHEL, CentOS or Oracle Linux, remove the `mksh` package and install the `ksh` package. After installing the `ksh` package, re-run the LifeKeeper for Linux setup.

#### **Example:**

1. Remove the `mksh` package

```
yum remove mksh
```

2. Install the `ksh` package

```
yum install ksh
```

3. Re-run setup

### **Unexpected termination of daemons**

Daemons using IPC terminate unexpectedly after update to Red Hat Enterprise Linux 7.2 and Red Hat 7.2 derivative systems. A new `systemd` feature was introduced in Red Hat Enterprise Linux 7.2 related to the cleanup of all allocated inter-process communication (IPC) resources when the last user session finishes. A session can be an administrative cron job or an interactive session. This behavior can cause daemons running under the same user, and using the same resources, to terminate unexpectedly.

To work around this problem, edit the file */etc/systemd/logind.conf* and add the following line:

```
RemoveIPC=no
```

Then, execute the following command, so that the change is put into effect:

```
systemctl restart systemd-logind.service
```

After performing these steps, daemons no longer crash in the described situation. Applications (such as MQ, Oracle, SAP, etc) using shared memory and semaphores may be affected by this issue and therefore require this change.

### **Re-execution of LifeKeeper’s “setup” script may be required after updating the kernel**

In Red Hat Enterprise Linux 7.x/CentOS 7.x/Oracle Linux 7.x environment, DataKeeper may not function properly when updating the kernel to 7.3 or later.

**Workaround:**

The problem will be solved by re-running the “setup” script that was executed when installing LifeKeeper on the updated system.

**Description:**

A loaded kernel module cannot be used after updating the kernel to 7.3 or later due to the compatibility of OS kernel modules.

DataKeeper uses a kernel module called nbd.ko, which accesses disks through the network. A correct nbd.ko module is installed when executing setup script for LifeKeeper installation.

nbd.ko for the new kernel will be installed by executing setup script again after updating the kernel.

**Unnecessary warning, displayed from the setup script**

Depending on the installation status of the LifeKeeper packages, the following warning is displayed when the setup script is executed.

Found changes in following files.

These files are overwritten in install process.

If you want to keep changes, please backup these files.

missing /opt/LifeKeeper/lkadm/subsys/scsi/DEVNAME

missing /opt/LifeKeeper/lkadm/subsys/scsi/DEVNAME/bin

missing /opt/LifeKeeper/subsys/scsi/resources/DEVNAME

This warning is caused by package management issues, but does not affect the setup and operation of LifeKeeper.

This issue will be fixed in a future version.

## 5.4.5.3.2. LifeKeeper Core – Known Issues / Restrictions

Description
<p><b>New or Deprecated Mount Options After a Kernel Upgrade</b></p> <p>When upgrading the Linux kernel, it is possible that some existing file system mount options may be deprecated in the new kernel or that the new kernel may add new default mount options to existing mounts. For example, the “nobarrier” mount option was deprecated in RedHat Enterprise Linux 8, and some kernel versions have added new default mount options such as “logbufs=8” and “logbsize=32k”.</p> <p>If a LifeKeeper-protected file system resource contains mount options which become deprecated after a kernel upgrade, the deprecated options should be removed from the list of mount options for the LifeKeeper resource on every server in the cluster. See the <a href="#">Modifying Mount Options for a LifeKeeper File System Resource</a> section for more details.</p> <p>If new default mount options are added by the kernel to an existing LifeKeeper-protected mount point after a kernel upgrade, then the new options should be added to the list of mount options for the LifeKeeper resource on every server in the cluster. See the <a href="#">Modifying Mount Options for a LifeKeeper File System Resource</a> section for more details.</p>
<p><b>If you set your shutdown strategy to “Do not Switchover Resources” (default), do not start LifeKeeper immediately after stopping it. If the time between stopping and starting LifeKeeper is too short, a split brain may occur. This is especially important for Quorum configurations in storage mode.</b></p> <p>Conflicts between LifeKeeper’s stop and start processes can cause a split brain. Allow a few seconds between stopping and starting LifeKeeper. Since Quorum configurations in storage mode take longer to stop, you need to allow more time than QWK_STORAGE_HBEATTIME * QWK_STORAGE_NUMHBEATS (24 seconds by default).</p>
<p><b>A split brain may occur if a resource fails while processing a restored comm path if quorum is configured</b></p> <p>If a cluster is configured with quorum and both quickCheck and local recovery fail during processing of a restored communication path, it can result in a race condition between the hierarchy failover and quorum processing resulting in a split brain between the nodes.</p>
<p><b>If quorum is configured and the active node fails or is rebooted with the Shutdown Strategy set to Switchover Resources, resources do switchover to the secondary node. However, when the original active node comes back up after the reboot, an attempt to switch the hierarchy back to the original active node is made, leaving the hierarchy in a failed / out of service state.</b></p>

With large hierarchies, there is a race condition that occurs during the LifeKeeper restart on the original active node that results in an attempt to restore the hierarchy on the original primary node. This results in removing the hierarchy on new active node and resulting in the hierarchy not being in service on any node. Please contact Support for a patch for this issue.

### **If there is a problem with a network connection, stop the service that automatically configures the network**

In an environment where IP addresses are protected using LifeKeeper, IP resources may conflict with daemons and services that automatically configure the network, such as `avahi-daemon`. If there is a problem when restoring communication paths or starting IP resources, stop the services that automatically configure the network.

### **Do not disconnect the network using the `ifconfig down` or the `ip link down` command**

When a network interface is disconnected using the `ifconfig down` or `ip link down` command, a communication path may not be restored after reconnecting, if a virtual IP resource is configured on the interface.

### **LifeKeeper does not start with `systemd` target set to multi-user**

In order for LifeKeeper to function properly, when running `systemctl set-default` or `systemctl isolate`, you must use the `lifekeeper-graphical.target` (for graphical mode) or `lifekeeper-multi-user.target` (for console mode). Do not use the normal `graphical.target` and `multi-user.target` `systemd` targets.

### **DataKeeper Disk UUID Restriction**

Starting in version 9.5.0, DataKeeper can no longer mirror disks that do not present a UUID to the operating system. The best way to mirror such a disk is to partition it with a GPT (GUID Partition Table). The “parted” tool can be used for this purpose. **Caution:** partitioning a disk will destroy any data that is already stored on the disk.

**Workaround:** See [DataKeeper for Linux Troubleshooting](#)

### **On SLES 15, LifeKeeper logging may not appear in the LifeKeeper log file following a log rotation**

If `logrotate` is run on the command line or if a background log rotation occurs due to the size of the log, LifeKeeper will stop logging.

**Workaround:** Run `systemctl reload rsyslog` to resume LifeKeeper logging.

### **File system labels should not be used in large configurations**

The use of file system labels can cause performance problems during boot-up with large clusters. The

problems are generally the result of the requirement that to use labels all devices connected to a system must be scanned. For systems connected to a SAN, especially those with LifeKeeper where accessing a device is blocked, this scanning can be very slow.

To avoid this performance problem on Red Hat systems, edit `/etc/fstab` and replace the labels with the path names.

### **lkscsid will halt the system when it should issue a sendevent when a disk fails in certain environments**

When `lkscsid` detects a disk failure, it should, by default, issue a `sendevent` to LifeKeeper to recover from the failure. The `sendevent` will first try to recover the failure locally and if that fails, will try to recover the failure by switching the hierarchy with the disk to another server. On some versions of Linux (RHEL 5 and SLES11), `lkscsid` will not be able to issue the `sendevent` but instead will immediately halt the system. This only affects hierarchies using the SCSI device nodes such as `/dev/sda` in a shared storage configuration.

### **DataKeeper Create Resource fails**

When using DataKeeper in certain environments (e.g., virtualized environments with IDE disk emulation, or servers with HP CCISS storage), an error may occur when a mirror is created:

```
ERROR 104052: Cannot get the hardware ID of the device "/dev/hda3"
```

This is because LifeKeeper does not recognize the disk in question and cannot get a unique ID to associate with the device.

**Workaround:** Use a GUID Partition so that LifeKeeper can recognize the disk in question.

### **Specifying hostnames for API access**

The key name used to store LifeKeeper server credentials must match the hostname of the other LifeKeeper server **exactly** (as displayed by the `hostname` command on that server). If the hostname is an FQDN, then the credential key must also be the FQDN. If the hostname is a short name, then the key must also be the short name.

**Workaround:** Make sure that the hostname(s) stored by [credstore](#) match the hostname exactly.

### **Restore of an lkbackup after a resource has been created may leave broken equivalencies**

The configuration files for created resources are saved during an `lkbackup`. If a resource is created for the first time after an `lkbackup` has been taken, that resource may not be properly accounted for when restoring from this previous backup.

**Solution:** Restore from `lkbackup` prior to adding a new resource for the first time. If a new resource has

been added after an `lkbakup`, it should either be deleted prior to performing the restore, or delete an instance of the resource hierarchy, then re-extend the hierarchy after the restore. **Note:** It is recommended that an `lkbakup` be run when a resource of a particular type is created for the first time.

#### Resources removed in the wrong order during failover

In cases where a hierarchy shares a common resource instance with another root hierarchy, resources are sometimes removed in the wrong order during a cascading failover or resource failover.

**Solution:** Creating a common root will ensure that resource removals in the hierarchy occur from the top down.

1. Create a `gen/app` that always succeeds on restore and remove.
2. Make all current roots children of this new `gen/app`.

**Note:** Using `/bin/true` for the restore and remove script would accomplish this.

#### Delete of nested file system hierarchy generates “Object does not exist” message

**Solution:** This message can be disregarded as it does not create any issues.

#### filesyshier returns the wrong tag on a nested mount create

When a database has nested file system resources, the file system kit will create the file system for both the parent and the nested child. However, `filesyshier` returns only the child tag. This causes the application to create a dependency on the child but not the parent.

**Solution:** When multiple file systems are nested within a single mount point, it may be necessary to manually create the additional dependencies to the parent application tag using `dep_create` or via the UI Create Dependency.

#### DataKeeper: Nested file system create will fail with DataKeeper

When creating a DataKeeper mirror for replicating an existing file system, if a file system is nested within this structure, you must unmount it first before creating the File System resource.

**Workaround:** Manually unmount the nested file systems and remount / create each nested mount.

#### Changing the mount point of the device protected by Filesystem resource may lead data corruption

The mount point of the device protected by LifeKeeper via the File System resource (`filesys`) must not be changed. Doing so may lead to the device being mounted on multiple nodes and if a switchover is done and

this could lead to data corruption.

### **XFS file system usage may cause quickCheck to fail.**

In the case CHECK\_FS\_QUOTAS setting is enabled for LifeKeeper installed on Red Hat Enterprise Linux 7 / Oracle Linux 7 / CentOS 7, quickCheck fails if uquota, gquota option is set to the XFS file system resource, which is to be protected.

**Solution:** Use usrquota, grpquota instead of uquota, gquota for mount options of XFS file system, or, disable CHECK\_FS\_QUOTAS setting.

### **Btrfs is not supported**

Btrfs (or any other LifeKeeper for Linux unsupported filesystem) cannot be used for LifeKeeper files (/opt/LifeKeeper), bitmap files if they are not in /opt/LifeKeeper, lbackupfiles, or any other LifeKeeper related files. In addition, LifeKeeper does not support protecting Btrfs (or any other LifeKeeper for Linux unsupported filesystem) within a resource hierarchy.

**Solution:** A simple work around for placing /opt/LifeKeeper on a Btrfs file system is to add a small disk to your instances and format that disk with ext4 or xfs, and mount this filesystem as /opt/LifeKeeper.

1. Create a small disk to be used for /opt/LifeKeeper
  - A minimum of 110MB is required for software installs
  - Note: In Azure, you can create a 1 GB data disk at a minimum.
  - Note: Additional ARKs and the number of mirrors may increase the total required space.
2. Once the disk is added to the node and visible, partition the disk (or use lvm).
  - Example: `gdisk /dev/sdb`
  - Note: How to add a disk to your system is outside the scope of this KBA (contact your sysadmin for your environment)
3. Format the partition with a supported filesystem (see <http://docs.us.sios.com/spslinux/9.4.1/en/topic/sios-protection-suite-for-linux-release-notes>).
  - Example: `mkfs.ext4 /dev/sdb1` (where sdb1 was created in step 2)
4. Add the newly created and formatted partition to /etc/fstab and set it to be automatically mounted on system boot.
5. Mount the new partition as /opt/LifeKeeper
  - Example: `mount /dev/sdb1 /opt/LifeKeeper`
  - verify filesystem is mounted
6. Install LifeKeeper for Linux
7. After the installation, edit /etc/fstab and add the entry, so the disk can be mounted on reboot.
  - Example: `/dev/sdb            /opt/LifeKeeper    ext4`

**SLES12 SP1 or later on AWS**

The following restrictions apply with SLES12 SP1 or later on AWS:

- Cannot set static routing configuration

Automatic IP address configuration via DHCP does not work if a static routing configuration is set in `/etc/sysconfig/network/routes`. This causes the network not to start correctly.

**Solution:** Update the routing information in the configuration file by modifying the “ROUTE” parameter in `/etc/sysconfig/network/ifroute-ethX`

- Hostname is changed even if the “Change Hostname via DHCP” setting is disabled. The LK service does not work properly if the hostname is rewritten. In SLES12 SP1 or later on AWS, the hostname is changed even after the “Change Hostname via DHCP” setting is disabled.

**Solution:**

- Update `/etc/cloud/cloud.cfg` to comment out the “update\_hostname” parameter
- Update `/etc/cloud/cloud.cfg` to set the `preserve_hostname` parameter to “true”
- Update `/etc/sysconfig/network/dhcp` to set the `DHCLIENT_SET_HOSTNAME` parameter to “no”

**Shutdown Strategy set to “Switchover Resources” may fail when using Quorum/Witness Kit in Witness mode**

Hierarchy switchover during LifeKeeper shutdown may fail to occur when using the Quorum/Witness Kit in

Witness mode.

**Workaround:** Manually switchover resource hierarchies before shutdown.

### Edit /etc/services

If the following entry in /etc/service is deleted, LifeKeeper cannot start up.

```
lcm_server 7365/tcp
```

Don't delete this entry when editing the file.

**Any storage unit which returns a string including a space for the SCSI ID cannot be protected by LifeKeeper.**

### Using bind mounts is not supported

Bind mounts (mount —bind) cannot be used for the file system protected by LifeKeeper.

On SLES running on AWS or Azure, change the network interface configuration file in order to prevent a cloud network plug-in from removing the virtual IP address.

Click [here](#) for more details.

**log ID 4739 – Switchover request failed when the core hierarchy reservation lock for the switchover was reset by a recover event for a mirror resync.**

When the switchover completed successfully and the core attempted to clear the lock and found it owned by another process it failed the switchover and marked the root resource as failed even though it had restored successfully.

**Solution/Workaround:** Raise the value of the RESRVRECTIMEOUT tunable. If RESRVRECTIMEOUT was raised above 600 then the value of RESRVTIMEOUT should also be raised to match it.

**Note:** When setting the timeout for the tunables RESRVRECTIMEOUT and RESRVTIMEOUT the value will very depending on the number of resources as well as the type of each resource in the hierarchy. The value must include the time to remove each resource on the current source node plus the amount of time to restore each resource on the new source node. Resource restore and remove times can vary from cluster to cluster and over time so to get an estimated value for these tunables follow the steps below.

Steps on how to determine the tunable setting:

- Perform multiple switchovers (we recommend at least 3).

- Document the times each switchover takes to complete and note the longest total switchover time.
  - Once the “longest total switchover” is recorded, set the value of the RESRVRECTIMEOUT to a time (in seconds) that is greater than the longest in-service time seen in any of the hierarchies containing a DataKeeper resource.

**Example:** The value set for RESRVRECTIMEOUT will be x and RESRVTIMEOUT will be y.

If the “x” value of RESRVRECTIMEOUT > 600, then set RESRVTIMEOUT to x. **[If x>600, then x=y]**

## 5.4.5.3.3. Internet/IP Licensing – Known Issues / Restrictions

Description
<p data-bbox="124 387 544 418"><b><i>/etc/hosts</i> settings dependency</b></p> <p data-bbox="124 472 379 504"><b><i>/etc/hosts</i> settings:</b></p> <p data-bbox="124 557 1436 629">When using internet-based licensing (IPv4 address), the configuration of <i>/etc/hosts</i> can negatively impact license validation. If LifeKeeper startup fails with:</p> <pre data-bbox="165 714 1356 875">Error in obtaining LifeKeeper license key: Invalid host. The hostid of this system does not match the hostid specified in the license file.</pre> <p data-bbox="124 972 1436 1086">and the listed internet hostid is correct, then the configuration of <i>/etc/hosts</i> may be the cause. To correctly match <i>/etc/hosts</i> entries, IPv4 entries must be listed before any IPv6 entries. To verify if the <i>/etc/hosts</i> configuration is the cause, run the following command:</p> <pre data-bbox="165 1171 1005 1202">/opt/LifeKeeper/bin/lmutil lmhostid -internet -n</pre> <p data-bbox="124 1296 1436 1368">If the IPv4 address listed does not match the IPv4 address in the installed license file, then <i>/etc/hosts</i> must be modified to place IPv4 entries before IPv6 entries to return the correct address.</p>

## 5.4.5.3.4. GUI – Known Issues / Restrictions

### Description

**GUI does not immediately update IP resource state after network is disconnected and then reconnected**

When the primary network between servers in a cluster is disconnected and then reconnected, the IP resource state on a remote GUI client may take as long as 1 minute and 25 seconds to be updated due to a problem in the RMI/TCP layer.

**Java Mixed Signed/Unsigned Code Warning – When loading the LifeKeeper Java GUI client applet from a remote system, the following security warning may be displayed:**



Enter “Run” and the following dialog will be displayed:



Block? Enter **“No”** and the LifeKeeper GUI will be allowed to operate.

**Recommended Actions:** To reduce the number of security warnings, you have two options:

1. Check the **“Always trust content from this publisher”** box and select **“Run”**. The next time the LifeKeeper GUI Java client is loaded, the warning message will not be displayed.

or

2. Add the following entry to your Java **“deployment.properties”** file to eliminate the second dialog about blocking. The security warning will still be displayed when you load the Java client, however, the applet will not be blocked and the Block **“Yes”** or **“No”** dialog will not be displayed. Please note this setting will apply to all of your Java applets.

```
deployment.security.mixcode=HIDE_RUN
```

To bypass both messages, implement 1 and 2.

### **steeleye-lighttpd process fails to start if Port 778 and 779 are in use**

If a process is using Port 778 and 779 when steeleye-lighttpd starts up, steeleye-lighttpd fails which can cause GUI connect failures and resource hierarchy extend issues.

**Solution:** Set the following tunables on all nodes in the cluster and then restart LifeKeeper on all the nodes:

Add the following lines to */etc/default/LifeKeeper*:

```
API_SSL_PORT=port_number  
LKAPI_WEB_PORT=port_number
```

where *port\_number* is the new port to use.

## 5.4.5.3.5. Data Replication – Known Issues / Restrictions

Description
<p><b>When using DataKeeper on Oracle Linux 8.5 (RHCK), it is necessary to replace the HADR package when installing the LifeKeeper.</b></p> <p>If you are using DataKeeper for Linux v9.6.1, perform the following steps on all nodes in the cluster. These steps are not needed if using UEK (Unbreakable Enterprise Kernel).</p> <p>Uninstall the existing HADR package.</p> <pre># rpm -e HADR-RHAS-4.18.0-all-9.4.0-6882.x86_64</pre> <p>Install the HADR package used for Oracle Linux 8.5. The following is an example where the LifeKeeper installation image (sps.img) is mounted on /mnt.</p> <pre># rpm -i /mnt/RHAS/HADR-RHAS-4.18.0-348.el8.x86_64-9.6.1-7412.x86_64.rpm</pre> <p>Re-read the nbd module.</p> <pre># modprobe -r nbd; modprobe nbd</pre> <p>Now a DataKeeper resource can be created and utilized.</p>
<p><b>A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</b></p> <p>In the DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode, the read/write process for mirrors may hang within the kernel.</p> <p>Run the following command on each node to determine if the DataKeeper resource is synchronous or asynchronous. 0 is synchronous mode and non zero is asynchronous mode. Resources with all synchronous or resources with all asynchronous on all nodes are acceptable. To avoid this issue do not mix synchronous and asynchronous modes.</p> <pre>perl -nle 'my @x = split(/\x01/, \$_); print "\$x[0]:\$x[3]";' /opt/LifeKeeper/subsys/scsi/resources/netraid/mirrorinfo_&lt;md num&gt;</pre> <p><b>Solution:</b> Currently no workaround is available. Recreate a DataKeeper resource and select synchronous mode at the time of creating and extending.</p>
<p><b>Partitions with an odd number of sectors are not supported when running kernel 4.12 or later</b></p> <p>The use of a partition with an odd number of sectors is not supported in a DataKeeper mirror in environments running kernel 4.12 or later. This is due to an issue where a resync may fail when attempting to write past the end of the disk.</p>

**Important reminder about DataKeeper for Linux asynchronous mode in an LVM over DataKeeper configuration**

Kernel panics may occur in configurations where LVM resources sit above multiple asynchronous mirrors. In these configurations data consistency may be an issue if a panic occurs. Therefore the required configurations are a single DataKeeper mirror or multiple synchronous DataKeeper mirrors.

**In symmetric active SDR configurations with significant I/O traffic on both servers, the filesystem mounted on the mirror stops responding and eventually the whole system hangs**

Due to the single threaded nature of the Linux buffer cache, the buffer cache flushing daemon can hang trying to flush out a buffer which needs to be committed remotely. While the flushing daemon is hung, all activities in the Linux system with dirty buffers will stop if the number of dirty buffers goes over the system accepted limit (set in `/proc/sys/kernel/vm/bdflush`).

Usually this is not a serious problem unless something happens to prevent the remote system from clearing remote buffers (e.g. a network failure). LifeKeeper will detect a network failure and stop replication in that event, thus clearing a hang condition. However, if the remote system is also replicating to the local system (i.e. they are both symmetrically replicating to each other), they can deadlock forever if they both get into this flushing daemon hang situation.

The deadlock can be released by manually killing the nbd-client daemons on both systems (which will break the mirrors). To avoid this potential deadlock entirely, however, symmetric active replication is not recommended.

**High CPU usage reported by top for md\_raid1 process with large mirror sizes**

With the `mdX_raid1` process (*with X representing the mirror number*), high CPU usage as reported by `top` can be seen on some OS distributions when working with very large mirrors (500GB or more).

Solution: To reduce the CPU usage percent, modify the chunk size to 1024 via the LifeKeeper tunable `LKDR_CHUNK_SIZE` then delete and recreate the mirror in order to use this new setting.

**The use of lkbakup with DataKeeper resources requires a full resync**

Although `lkbakup` will save the `instance` and `mirror_info` files, it is best practice to perform a full resync of DataKeeper mirrors after a restore from `lkbakup` as the status of source and target cannot be guaranteed while a resource does not exist.

**DataKeeper does not support using Network Compression on SLES12 SP1 or later**

DataKeeper does not support using Network Compression on SLES12 SP1 or later due to disk I/O performance problem.

**Certain kernel versions do not support DataKeeper asynchronous mode.**

It has been observed that kernel panic will occur with certain kernel versions when using DataKeeper resource asynchronous mode with LifeKeeper for Linux. Since this is a kernel dependent problem, there is no fundamental solution with LifeKeeper. In order to use DataKeeper asynchronous mode configuration, it is necessary to update or downgrade the kernel.

The kernel versions that do not support the DataKeeper asynchronous mode are as follows.

3.10.0-693. series for 3.10.0-693.24.1.el7.x86\_64 or later

3.10.0-862.el7.x86\_64 ~ 3.10.0-862.26.x.el7.x86\_64

3.10.0-957.el7.x86\_64 ~ 3.10.0-957.3.x.el7.x86\_64

If you use the kernel version listed above and use DataKeeper resources in asynchronous mode, please update (or downgrade) to the following kernel version.

3.10.0-693. series kernel for before 3.10.0-693.24.1.el7.x86\_64

3.10.0-862.29.1.el7.x86\_64 or later

3.10.0-957.4.1.el7.x86\_64 or later

If you cannot update (or downgrade) the kernel, do not use DataKeeper asynchronous mode.

### **Some kernel versions do not support the Secure Boot feature**

If Secure Boot is enabled on RHEL7 or later, CentOS7 or later, or Oracle Linux 7 or later, the nbd module fails to load. Also, in some kernel versions of SUSE Linux Enterprise Server and Oracle Linux UEK kernel, loading of the md / raid1 kernel module fails when Secure Boot is enabled.

**Solution:** Take one of the following actions:

1. Disable Secure Boot – Disable Secure Boot in the UEFI configuration.
2. Disable signature verification – Disable signature verification with the “`mokutil --disable-validation`” command. See mokutil documentations for details.

**Solution 1 is recommended. Both require a system reboot.**

## 5.4.5.3.6. IPv6 – Known Issues / Restrictions

### Description

SIOS has migrated to the use of the `ip` command and away from the `ifconfig` command. Because of this change, customers with external scripts are advised to make a similar change. Instead of issuing the `ifconfig` command and parsing the output looking for a specific interface, scripts should instead use “`ip -o addr show`” and parse the output looking for lines that contain the words “`inet`” and “`secondary`”.

```
# ip -o addr show
1: lo:  mtu 16436 qdisc noqueue state UNKNOWN
    \   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
1: lo    inet 127.0.0.1/8 scope host lo
1: lo    inet6 ::1/128 scope host
    \    valid_lft forever preferred_lft forever
2: eth0: mtu 1500 qdisc pfifo_fast state UP qlen 1000
    \   link/ether d2:05:de:4f:a2:e6 brd ff:ff:ff:ff:ff:ff
2: eth0  inet 172.17.100.77/22 brd 172.17.103.255 scope global eth0
2: eth0  inet 172.17.100.79/22 scope global secondary eth0
2: eth0  inet 172.17.100.80/22 scope global secondary eth0
2: eth0  inet6 2001:5c0:110e:3364::1:2/64 scope global
    \    valid_lft forever preferred_lft forever
2: eth0  inet6 2001:5c0:110e:3300:d005:deff:fe4f:a2e6/64 scope global dynam
ic
    \    valid_lft 86393sec preferred_lft 14393sec
2: eth0  inet6 fe80::d005:deff:fe4f:a2e6/64 scope link
    \    valid_lft forever preferred_lft forever
```

So for the above output from the `ip` command, the following lines contain virtual IP addresses for the `eth0` interface:

```
2: eth0    inet 172.17.100.79/22 scope global secondary eth0
2: eth0    inet 172.17.100.80/22 scope global secondary eth0
```

### **'IPV6\_AUTOCONF = No' for /etc/sysconfig/network-scripts/ifcfg-<nicName> is not being honored on reboot or boot**

On boot, a stateless, auto-configured IPv6 address is assigned to the network interface. If a comm path is created with a stateless IPv6 address of an interface that has `IPV6_AUTOCONF=No` set, the address will be removed if any system resources manage the interface, e.g. `ifdown <nicName>;ifup <nicName>`.

Comm path using auto-configured IPv6 addresses did not recover and remained dead after rebooting

primary server because IPV6\_AUTOCONF was set to No.

**Solution:** Use Static IPv6 addresses only. The use of auto-configured IPv6 addresses could cause a comm loss after a reboot, a NIC change, etc.

While IPv6 auto-configured addresses may be used for comm path creation, the system administrator should be aware of the following conditions:

- IPv6 auto-configured/stateless addresses are dependent on the network interface (NIC) MAC address. If a comm path was created and the associated NIC is later replaced, the auto-configured IPv6 address will be different and LifeKeeper will show the comm path is dead. The comm path will need to be recreated.
- With Red Hat Enterprise Linux, implementing the intended behavior for assuring consistent IPv6 auto-configuration during all phases of host operation requires specific domain knowledge for accurately and precisely setting the individual interface config files AS WELL AS the sysctl.conf, net.ipv6.\* directives (i.e. explicitly setting IPV6\_AUTOCONF in the ifcfg-<nic> which is referenced by the 'if/ip' utilities AND setting directives in /etc/sysctl.conf which impact NIC control when the system is booting and switching init levels).

#### **IP: Modify Source Address Setting for IPv6 doesn't set source address**

When attempting to set the source address for an IPv6 IP resource, it will report success when nothing was changed.

**Workaround:** Currently no workaround is available. This will be addressed in a future release.

#### **IP: Invalid IPv6 addressing allowed in IP resource creation**

Entering IPv6 addresses of the format 2001:5c0:110e:3368:000000:000000001:61:14 is accepted when the octets contain more than four characters.

**Workaround:** Enter correctly formatted IPv6 addresses.

#### **IPv6 resource reported as ISP when address assigned to bonded NIC but in 'tentative' state**

IPv6 protected resources in LifeKeeper will incorrectly be identified as 'In Service Protected' (ISP) on SLES systems where the IPv6 resource is on a bonded interface, a mode other than 'active-backup' (1) and Linux kernel 2.6.21 or lower. The IPv6 bonded link will remain in the 'tentative' state with the address unresolvable.

**Workaround:** Set the bonded interface mode to 'active-backup' (1) or operate with an updated kernel which will set the link state from 'tentative' to 'valid' for modes other than 'active-backup' (1).

## 5.4.5.3.7. Apache – Known Issues / Restrictions

---

Description
<p><b>Apache Kit does not support IPv6; doesn't indentify IPv6 in <i>httpd.conf</i></b></p> <p>Any IPv6 addresses assigned to the 'Listen' directive entry in the <i>httpd.conf</i> file will cause problems.</p> <p><b>Solution:</b> Until there is support for IPv6 in the Apache Recovery Kit, there can be no IPv6 address in the <i>httpd.conf</i> file after the resource has been created.</p>

## 5.4.5.3.8. Oracle – Known Issues / Restrictions

Description
<p><b>The Oracle Recovery Kit does not include support for Connection Manager and Oracle Names features</b></p> <p>The LifeKeeper Oracle Recovery Kit does not include support for the following Oracle Net features of Oracle: Oracle Connection Manager, a routing process that manages a large number of connections that need to access the same service; and Oracle Names, the Oracle-specific name service that maintains a central store of service addresses.</p> <p>The LifeKeeper Oracle Recovery Kit does protect the Oracle Net Listener process that listens for incoming client connection requests and manages traffic to the server. Refer to the <a href="#">LifeKeeper for Linux Oracle Recovery Kit Administration Guide</a> for LifeKeeper configuration specific information regarding the Oracle Listener.</p>
<p><b>The Oracle Recovery Kit does not support the ASM or grid component features</b></p> <p>The Oracle Automatic Storage Manager (ASM) feature provided in Oracle is not currently supported with LifeKeeper. In addition, the grid components are not protected by the LifeKeeper Oracle Recovery Kit. Support for raw devices, file systems, and logical volumes are included in the current LifeKeeper for Linux Oracle Recovery Kit. The support for the grid components can be added to LifeKeeper protection using the gen/app recovery kit.</p>
<p><b>The Oracle Recovery Kit does not support NFS Version 4</b></p> <p>The Oracle Recovery Kit supports NFS Version 3 for shared database storage. NFS Version 4 is not supported at this time due to NFSv4 file locking mechanisms.</p>
<p><b>Oracle listener stays in service on primary server after failover</b></p> <p>Network failures may result in the listener process remaining active on the primary server after an application failover to the backup server. Though connections to the correct database are unaffected, you may still want to kill that listener process.</p>
<p><b>DataKeeper: Nested file system create will fail with DataKeeper</b></p> <p>When creating a DataKeeper mirror for replicating an existing file system, if a file system is nested within this structure, you must unmount it first before creating the File System resource.</p> <p><b>Proper Procedure:</b> Create the /oracle/BPP file system before creating the sub-file systems under it, such as /oracle/BPP/mirrorlogA. (i.e. the /oracle/BPP file system should appear at the bottom of the GUI list)</p> <p><b>Workaround:</b> Manually unmount the nested file systems and remount / create each nested mount.</p>

## 5.4.5.3.9. MySQL – Known Issues / Restrictions

---

Description
<p><b>The “include” directive is not supported</b></p> <p>The “include” directive is not supported. All the setup configuration information must be described in a single my.cnf file.</p>
<p><b>Crash Recovery</b></p> <p>Restarting MySQL after an abnormal termination initiates a MySQL crash recovery. While in this recovery state MySQL client connections are denied. This will prevent LifeKeeper from checking the state of MySQL causing a possible failover to the standby node.</p>

## 5.4.5.3.10. NAS Recovery Kit – Known Issues / Restrictions

---

Description
<p><b>Autofs is not supported</b></p> <p>File system managed by autofs cannot be protected.</p>

## 5.4.5.3.11. NFS Server – Known Issues / Restrictions

Description
<p><b>Top level NFS resource hierarchy uses the switchback type of the hanfs resource</b></p> <p>The switchback type, which dictates whether the NFS resource hierarchy will automatically switch back to the primary server when it comes back into service after a failure, is defined by the hanfs resource.</p>
<p><b>IPv6 address cannot be used to specify the client</b></p> <p>Resource creation will fail when the client is specified using an IPv6 address in the <i>/etc/exports file</i>.</p> <p><b>Solution:</b> Use hostnames or wildcards to specify the client.</p>
<p><b>File Lock switchover fails</b></p> <p>Switchover file locks during resources switchover/failover does not work. The result is the same when mounting with both NFS v3 or v4. Do not use file lock from client applications.</p>
<p><b>Resource creation process will hang when gssproxy is not running</b></p> <p>If the gssproxy daemon is not running, LifeKeeper starts it when a resource is created. However, the resource creation process is not completed since the gssproxy does not work as LifeKeeper expected. This issue can occur in both lkGUIapp and lkcli. If the gssproxy is not installed and rpc.svcgssd is installed, this issue does not occur because the rpc.svcgssd is used.</p> <p><b>Solution:</b> Start gssproxy daemon. Enable the auto-start to keep it running all the time. The following are examples of RHEL 7.</p> <pre>systemctl start gssproxy.service systemctl enable gssproxy.service</pre>
<p><b>If rpc.idmapd and rpc.svcgssd are not running in the LifeKeeper Single Server Protection environment, the restore function will fail</b></p> <p>In LifeKeeper Single Server Protection, the directory to mount rpc_pipefs is not created. Mounting rpc_pipefs is tried before starting rpc. idmapd and rpc.svcgssd daemons, but it fails because the directory is missing.</p> <p><b>Solution:</b> Start rpc.idmapd and rpc.svcgssd daemons. Enable the auto-start to keep it running all the time. The following are examples of SLES15. If the gssproxy is installed, it is not necessary to have the rpc.svcgssd daemon running.</p> <pre>systemctl start nfs-idmapd.service systemctl enable nfs-idmapd.service systemctl start rpc-svcgssd.service</pre>

```
systemctl enable rpc-svcgssd.service
```

## 5.4.5.3.12. SAP Recovery Kit – Known Issues / Restrictions

Description
<p><b>A split brain may occur for ERS v2 resources while processing a restored comm path</b></p> <p>If an SAP cluster is configured with quorum and both the ASCS and ERS v2 resources reside on the same node when the communication path is restored to an eligible node for the ERS v2 resource, then a split brain may occur while attempting to switch the ERS v2 resource to the eligible node. Please contact Support for a patch for this issue.</p>
<p><b>SAP resources fail to come in-service due to csh bug in RHEL 8</b></p> <p>Due to a bug in the tcsh package available for RHEL 8, the SAP administrative user's .cshrc and .login files are not sourced correctly in certain situations. Due to this, important environment variables that the SAP Recovery Kit depends on are not properly exported, which may cause SAP resources to fail to come in-service on RHEL 8. See RedHat Bug 1714267 for more details and a workaround.</p>
<p><b>SAP Dual Stack Environment Restriction</b></p> <p>The <a href="#">redesigned ERS resource type</a> introduced in LifeKeeper for Linux 9.4.0 (which operates in a hierarchy separate from the corresponding central services instance) does not support an SAP dual stack (ABAP+Java) environment where there are two pairs of central services and enqueue replication server instances (e.g., ASCS00/ERS10 and SCS01/ERS11) installed under the same SID. Customers with an SAP dual stack (ABAP+Java) environment installed under the same SID should continue to use the pre-9.4.0 ERS resource design (which is located at the top of the SAP hierarchy with a dependency on the corresponding ASCS/SCS resource).</p>
<p><b>Failed delete or unextend of a SAP hierarchy</b></p> <p>Deleting or unextending a SAP hierarchy that contains the same IP resource in multiple locations within the hierarchy can sometimes cause a core dump that results in resources not being deleted.</p> <p>To correct the problem, after the failed unextend or delete operation, manually remove any remaining resources using the LifeKeeper GUI. You may also want to remove the core file from the server.</p>
<p><b>Handle Warnings gives a syntax error at -e line 1</b></p> <p>When changing the default behavior of <b>No</b> in <b>Handle Warnings</b> to <b>Yes</b>, an error is received.</p> <p><b>Solution:</b> Leave this option at the default setting of <b>No</b>. <b>Note:</b> It is highly recommended that this setting be left on the default selection of <b>No</b> as Yellow is a transient state that most often does not indicate a failure.</p>
<p><b>Choosing same setting causes missing button on Update Wizard</b></p> <p>If user attempts to update the <b>Handle Warning</b> without changing the current setting, the next screen, which indicates that they must go back, is missing the <b>Done</b> button.</p>

**When changes are made to res\_state, monitoring is disabled**

If **Protection Level** is set to **BASIC** and SAP is taken down manually (i.e. for maintenance), it will be marked as FAILED and monitoring will stop.

**Solution:** In order for monitoring to resume, LifeKeeper will need to start up the resource instead of starting it up manually.

**ERS in-service fails on remote host if ERS is not parent of Core/CI**

**Note:** This only applies to the pre-9.4.0 ERS resource design (which is located at the top of the SAP hierarchy with a dependency on the corresponding ASCS/SCS resource). For more details see [ERS Resource Types in LifeKeeper](#).

Creating an ERS resource without any additional SAP resource dependents will cause initial in-service to fail on switchover.

**Solution:** Create ERS as parent of CI/Core instance (SCS or ASCS), then retry in-service.

**SAP instance processes in an inconsistent state**

Issuing concurrent administrative commands while a migration of an SAP resource is in-progress may leave the SAP instance processes in an inconsistent state, which may require manual intervention to resolve.

## 5.4.5.3.13. LVM – Known Issues / Restrictions

---

Description
<p><b>Important reminder about DataKeeper for Linux asynchronous mode in an LVM over DataKeeper configuration</b></p> <p>Kernel panics may occur in configurations where LVM resources sit above multiple asynchronous mirrors. In these configurations data consistency may be an issue if a panic occurs. Therefore the required configurations are a single DataKeeper mirror or multiple synchronous DataKeeper mirrors.</p>
<p><b>Use of IkID incompatible with LVM overwritten on entire disk</b></p> <p>When IkID is used to generate unique disk IDs on disks that are configured as LVM physical volumes, there is a conflict in the locations in which the IkID and LVM information is stored on the disk. This causes either the IkID or LVM information to be overwritten depending on the order in which IkID and pvcreate are used.</p> <p><b>Workaround:</b> When it is necessary to use IkID in conjunction with LVM, partition the disk and use the disk partition(s) as the LVM physical volume(s) rather than the entire disk.</p>

## 5.4.5.3.14. Multipath Recovery Kits (DMMP / HDLM / PPATH / NECSPS) Known Issues / Restrictions

---

Description
<p data-bbox="124 459 1428 533"><b>Multipath Recovery Kits (DMMP / HDLM / PPATH / NECSPS): Registration conflict error occurs on lkstop when resource OSF</b></p> <p data-bbox="124 584 1452 658">The multipath recovery kits (DMMP, HDLM, PPATH, NECSPS) can have a system halt occur on the active (ISP) node when LifeKeeper is stopped on the standby (OSU) node if the Multipath resource state is OSF.</p> <p data-bbox="124 710 316 741"><b>Workarounds:</b></p> <p data-bbox="124 792 1034 824">a) Switch the hierarchy to the standby node before LifeKeeper is stopped</p> <p data-bbox="124 875 172 907"><b>OR</b></p> <p data-bbox="124 958 1326 1032">b) Run <code>ins_setstate</code> on the standby node and set the Multipath resource state to OSU before LifeKeeper is stopped</p>

## 5.4.5.3.15. DMMP – Known Issues / Restrictions

Description
<p><b>DMMP: Write issued on standby server can hang</b></p> <p>If a write is issued to a DMMP device that is reserved on another server, then the IO can hang indefinitely (or until the device is no longer reserved on the other server). If/when the device is released on the other server and the write is issued, this can cause data corruption.</p> <p>The problem is due to the way the path checking is done along with the IO retries in DMMP. When “no_path_retry” is set to 0 (fail), this hang will not occur. When the path_checker for a device fails when the path is reserved by another server (MSA1000), then this also will not occur.</p> <p><b>Workaround:</b> Set “no_path_retry” to 0 (fail). However, this can cause IO failures due to transient path failures.</p>
<p><b>DMMP: Multiple initiators are not registered properly for SAS arrays that support ATP_C</b></p> <p>LifeKeeper does not natively support configurations where there are multiple SAS initiators connected to a SAS array. In these configurations, LifeKeeper will not register each initiator correctly, so only one initiator will be able to issue IOs. Errors will occur if the multipath driver (DMMP for example) tries to issue IOs to an unregistered initiator.</p> <p><b>Solution:</b> Set the following tunable in <i>/etc/default/LifeKeeper</i> to allow path IDs to be set based on SAS storage information:</p> <pre>MULTIPATH_SAS=TRUE</pre>
<p><b>Two or more different storage can not be used concurrently in case of the parameter configuration of DMMP recovery kit is required for some storage model.</b></p>
<p><b>DMMP RK doesn't function correctly if the disk name ends with “p&lt;number&gt;”.</b></p> <p>The DMMP RK doesn't function correctly if the disk name ends with “p&lt;number&gt;”.</p> <p><b>Workaround:</b> Do not create disk names ending in “p&lt;number&gt;”.</p>

## 5.4.5.3.16. DB2 – Known Issues / Restrictions

---

Description
<p><b>DB2 Recovery Kit reports unnecessary error</b></p> <p>If DB2 is installed on a shared disk, the following message may be seen when extending a DB2 resource.</p> <pre data-bbox="165 629 1461 703">LifeKeeper was unable to add instance "%s" and/or its variables to the DB2 registry.</pre> <p>This message will not adversely affect the behavior of the DB2 resource extend.</p>

## 5.4.5.3.17. Sybase ASE – Known Issues / Restrictions

Description
<p><b>User Name/Password Issues:</b></p> <ul style="list-style-type: none"> <li>• <b>If the default user name is password-protected, the create UI does not detect this until after all validation is complete</b></li> </ul> <p>When creating the Sybase resource, you are prompted to enter the user name. The help to front displays a message that if no user is specified, the default of 'sa' will be used. However, no password validation is done for the default account at this time. When LifeKeeper attempts to create the Sybase resource, the resource creation fails because the password has not been validated or entered. The password validation occurs on the user/password dialog, but only when a valid user is actually entered on the user prompt. Even if using the default user name, it must be specified during the <b>create</b> action.</p> <ul style="list-style-type: none"> <li>• <b>Password prompt skipped if no user name specified</b></li> </ul> <p>User/password dialog skips the password prompt if you do not enter a user name. When updating the user/password via the UI option, if you do not enter the Sybase user name, the default of 'sa' will be used and no password validation is done for the account. This causes the monitoring of the database to fail with invalid credential errors. Even if using the default user name, it must be specified during the update action. To fix this failure, perform the following steps:</p> <ol style="list-style-type: none"> <li>1. Verify that the required Sybase data files are currently accessible from the intended server. In most instances, this will be the backup server due to the monitoring and local recovery failure on the primary.</li> <li>2. Start the Sybase database instance from the command line on this server (see the Sybase product documentation for information on starting the database manually).</li> <li>3. From the command line, change directory (cd) to the LKROOT/bin directory (/opt/LifeKeeper/bin on most installations).</li> <li>4. Once in the bin directory, execute the following: <pre>./ins_setstate -t &lt;SYBASE_TAG&gt; -S ISP</pre> <p>where &lt;SYBASE_TAG&gt; is the tag name of the Sybase resource</p> </li> <li>5. When the command completes, immediately execute the <b>Update User/Password Wizard</b> from the UI and enter a valid user name, even if planning to use the Sybase default of 'sa'. <b>Note:</b> The</li> </ol>

**Update User/Password Wizard** can be accessed by right-clicking on the Sybase resource instance and selecting **Change Username/Password**.

6. When the hierarchy has been updated on the local server, verify that the resource can be brought in service on all nodes.

#### 7. **Protecting backup server fails when Sybase local user name >= eight characters**

The Sybase user name must consist of less than eight characters. If the Sybase local user name is greater than eight characters, the process and user identification checks used for resource creation and monitoring will fail. This will also prevent the protection of a valid Sybase Backup Server instance from being selected for protection. This problem is caused by the operating system translation of user names that are >= eight characters from the name to the UID in various commands (for example, ps). You must use a user name that is less than eight characters long.

#### **Resource Create Issue:**

- Default Sybase install prompt is based on ASE 16.0 SP02 (/opt/sybase). During the LifeKeeper resource creation, the default prompt for the location of the Sybase installation shows up relative to Sybase Version 16.0 SP02 (/opt/sybase). You must manually enter or browse to the correct Sybase install location during the resource create prompt.

#### **Extend Issues:**

- **The Sybase tag prompt on extend is editable but should not be changed.** The Sybase tag can be changed during extend, but this is not recommended. Using different tags on each server can lead to issues with remote administration via the command line.

#### **Properties Page Issues:**

- **Image appears missing for the Properties pane update user/password.** Instead of the proper image, a small square appears on the toolbar. Selecting this square will launch the **User/Password Update Wizard**.

**Sybase Monitor server** is not supported in 15.7 or later with LifeKeeper. If the Sybase Monitor server process is configured in Sybase 15.7 or later, you must use a Generic Application (gen/app) resource to protect this server process.

#### **Cannot create a resource using the Sybase RK OR experiencing issues on a LifeKeeper upgrade when using SLES 15**

##### **Symptoms:**

When setting up SLES 15 (this includes SLES 15 SP1 or SP2), the Sybase RK may not work on an upgrade or on a fresh install of LifeKeeper.

During a fresh install of LifeKeeper you won't be able to create the Sybase resource, or during a LifeKeeper

upgrade there could be issues in controlling the resource from LifeKeeper. The underlying error is “An error occurred when attempting to allocate localization-related structures”.

**Resolution:**

There are 2 possible solutions:

1. Add a line in the [Linux] section of the file `/sybase/<SID>/locales/locales.dat`

```
locale = POSIX, us_english, utf8
```

OR

2. Add this line to the Sybase profile for the `syb<SID>` user profile.

This profile is typically located in `/sybase//SYBASE.sh`

```
export LANG=en_US.UTF-8
```

<SID> is the 3 letter system ID.

These changes force the correct language to be used and is appropriate for US English locale. If another locale is used for Japan or China, please consult the local country and the system locale.

You can also su to the `syb<SID>` user and run the command `echo $LANG`. The `cshell` parameter will point to the correct language for the appropriate locale.

Which change is appropriate to make? It depends on the comfort level of the Sybase (SAP ASE) database administrator. One is changing the shell script, the other will add the POSIX support to the locales.

**Unable to detect that Sybase ARK is running**

**Symptom:** Unable to detect that Sybase ARK is running.

**Cause:** The Sybase ARK uses the default sql interface tool (`isql`). On some 64 bit systems, the `isql` tool is installed as `isql64` and not `isql`.

**Solution:** The `isql64` tool can be copied, in the same path, to `isql`. Or a link can be created between the `isql64` executable and `isql`.

## 5.4.5.3.18. WebSphere MQ – Known Issues / Restrictions

Description
<p><b>Error when lksupport command is executed:</b></p> <p>The following error can be output when lksupport command is executed in the case MQ queue manager protected by MQ RK is set on the disk shared by NFS.</p> <pre data-bbox="165 674 1214 701">cat: &lt;PATH&gt;/mqm/qmgrs/tkqmgr/qm.ini: Operation not permitted</pre> <p>This happens because the root access to NFS area is prohibited. This error output doesn't cause any problem.</p>
<p><b>Quickcheck fails if queue has long messages if a message is in the test queue of size &gt; 101 characters the put/get fails and the queue fills up</b></p>
<p><b>Install fails if the only installed MQ is a relocated install (non-standard and not likely)</b></p>
<p><b>Package install fails if the MQ package does not have the default name</b></p>
<p><b>Compile samples fails if the software is not installed under /opt/mqm</b></p>
<p><b>If two listeners are defined for a single instance and one is set to manual and the other is automatic failures can occur in create and quickCheck</b></p>

## 5.4.5.3.19. SAP HANA – Known Issues / Restrictions

Description
<p><b>Local recovery may fail if the local database is stopped with a sapcontrol Stop/StopWait command, resulting in a failover of the SAP HANA resource hierarchy</b></p> <p><b>ISSUE:</b> When a user issues a <b>Stop/StopWait</b> request for the HDB instance using the sapcontrol utility (which is also what is used internally when a user issues an ‘HDB stop’ command), sapstartsrv begins an asynchronous process of gracefully stopping all of the HANA database processes, and does not stop this process until either the database is completely shut down or the process times out. Therefore any other action issued via sapcontrol while sapstartsrv is in the process of gracefully shutting down the database will compete with the already-in-progress stop action, and will ultimately fail and time out.</p> <p>In particular, the following sequence of events may lead to a failover of the SAP HANA resource hierarchy, even when local recovery is enabled for the protected database:</p> <ol style="list-style-type: none"> <li>1. A user initiates a graceful shutdown of the HANA database while it is running on the primary server by issuing a ‘sapcontrol Stop/StopWait’ or ‘HDB stop’ command.</li> <li>2. The ‘quickCheck’ script in the SAP HANA Recovery Kit detects that at least one database process is no longer running, which results in an attempt to locally restart the database.</li> <li>3. The ‘recover’ script in the SAP HANA Recovery Kit issues a ‘sapcontrol StartWait’ command to attempt to restart the protected HDB instance.</li> <li>4. Because the ‘sapcontrol Stop/StopWait’ command issued in step 1 is still actively stopping the HANA database processes, the ‘sapcontrol Start’ command issued by the SAP HANA Recovery Kit fails and times out.</li> <li>5. Since the SAP HANA Recovery Kit is unable to restart the database locally, local recovery fails and the SAP HANA resource hierarchy fails over to the standby server.</li> </ol> <p><b>WORKAROUND/SOLUTION:</b> If the database is being stopped manually as part of pre-production cluster testing to simulate local recovery after a failure of the primary database, consider forcefully killing the database processes (e.g., with ‘HDB kill-9’) to more accurately simulate a primary database crash. See <a href="#">Testing Your SAP HANA Resource Hierarchy</a> for sample test cases.</p>

## 5.4.5.3.20. EC2 Recovery Kit Known Issues / Restrictions

---

Description
<p data-bbox="124 383 855 416"><b>Set the UTC time of the instance to the AWS UTC time.</b></p> <p data-bbox="124 468 1316 501">If the time of the instance does not match the AWS time, restoring of the EC2 resources will fail.</p>

## 5.4.5.3.21. Known Issues/Restrictions when using LifeKeeper on Oracle Cloud Infrastructure (OCI)

---

### Oracle Cloud Infrastructure (OCI) Support Configuration

Oracle provides images for the following Operating Systems support by LifeKeeper.

- Oracle Linux 7 (UEK, RHCK)
- Oracle Linux 8 (UEK, RHCK)
- CentOS 7

The OS configured in BYOI is the same as the OS supported by LifeKeeper (e.g., RHEL). The following Operating Systems are not supported.

- SUSE Linux Enterprise Server (SLES)
- CentOS 8

The following Recovery Kits are not supported at the time of v9.6.1 release.

- WebSphere MQ Recovery Kit
- SAP Recovery Kit
- SAP HANA Recovery Kit
- SAP MaxDB Recovery Kit
- Sybase Recovery Kit

The following Recovery Kits are not available on OCI.

- DB2 Recovery Kit
- Recovery Kit for EC2
- Recovery Kit for Route53
- VMDK as Shared Storage Recovery Kit
- Multipath Recovery Kits

 Refer to the [Requirement](#) and [Restrictions](#) pages for Recovery Kit for Oracle Cloud Infrastructure restrictions.

## 5.4.5.4. Communication Paths Going Up and Down

---

If you find the communication paths failing then coming back up repeatedly (the LifeKeeper GUI showing them as Alive, then Dead, then Alive), the heartbeat tunables may not be set to the same values on all servers in the cluster.

This situation is also possible if the tunable name is misspelled in the LifeKeeper defaults file `/etc/default/LifeKeeper` on one of the servers.

### Suggested Action

1. Shut down LifeKeeper on all servers in the cluster.
2. On each server in the cluster, check the values and spelling of the `LCMHBEATTIME` and `LCMNUMHBEATS` tunables in `/etc/default/LifeKeeper`. Ensure that for each tunable, the values are the same on ALL servers in the cluster.
3. Restart LifeKeeper on all servers.

## 5.4.5.5. Incomplete Resource Created

---

If the resource setup process is interrupted leaving instances only partially created, you must perform manual cleanup before attempting to install the hierarchy again. Use the LifeKeeper GUI to delete any partially-created resources. See [Deleting a Hierarchy from All Servers](#) for instructions. If the hierarchy list does not contain these resources, you may need to use the `ins_remove` (see LCDI-instances(1M)) and `dep_remove` (LCDI-relationship(1M)) to clean up the partial hierarchies.

## 5.4.5.6. Incomplete Resource Priority Modification

---

A hierarchy in LifeKeeper is defined as all resources associated by parent/child relationships. For resources that have multiple parents, it is not always easy to discern from the GUI all of the root resources for a hierarchy. In order to maintain consistency in a hierarchy, LifeKeeper requires that priority changes be made to all resources in a hierarchy for each server. The GUI enforces this requirement by displaying all root resources for the hierarchy selected after the OK or Apply button is pressed. You have the opportunity at this point to accept all of these roots or cancel the operation. If you accept the list of roots, the new priority values will be applied to all resources in the hierarchy.

You should ensure that no other changes are being made to the hierarchy while the Resource Properties dialog for that hierarchy is displayed. Before you have edited a priority in the Resource Properties dialog, any changes being made to LifeKeeper are dynamically updated in the dialog. Once you have begun making changes, however, the values seen in the dialog are frozen even if underlying changes are being made in LifeKeeper. Only after selecting the Apply or OK button will you be informed that changes were made that will prevent the priority change operation from succeeding as requested.

In order to minimize the likelihood of unrecoverable errors during a priority change operation involving multiple priority changes, the program will execute a multiple priority change operation as a series of individual changes on one server at a time. Additionally, it will assign temporary values to priorities if necessary to prevent temporary priority conflicts during the operation. These temporary values are above the allowed maximum value of 999 and may be temporarily displayed in the GUI during the priority change. Once the operation is completed, these temporary priority values will all be replaced with the requested ones. If an error occurs and priority values cannot be rolled back, it is possible that some of these temporary priority values will remain. If this happens, follow the suggested procedure outlined below to repair the hierarchy.

### Restoring Your Hierarchy to a Consistent State

If an error occurs during a priority change operation that prevents the operation from completing, the priorities may be left in an inconsistent state. Errors can occur for a variety of reasons, including system and communications path failure. If an error occurs after the operation has begun, and before it finishes, and the program was not able to roll back to the previous priorities, you will see a message displayed that tells you there was an error during the operation and the previous priorities could not be restored. If this should happen, you should take the following actions to attempt to restore your hierarchy to a consistent state:

1. If possible, determine the source of the problem. Check for system or communications path failure. Verify that other simultaneous operations were not occurring during the same time that the priority administration program was executing.
2. If possible, correct the source of the problem before proceeding. For example, a failed system or communications path must be restored before the hierarchy can be repaired.
3. Re-try the operation from the Resource Properties dialog.

4. If making the change is not possible from the Resource Properties dialog, it may be easier to attempt to repair the hierarchy using the command line `hry_setpri`. This script allows priorities to be changed on one server at a time and does not work through the GUI.
5. After attempting the repair, verify that the LifeKeeper databases are consistent on all servers by executing the `eqv_list` command for all servers where the hierarchy exists and observing the priority values returned for all resources in the hierarchy.
6. As a last resort, if the hierarchy cannot be repaired, you may have to delete and re-create the hierarchy.

## 5.4.5.7. No Shared Storage Found When Configuring a Hierarchy

---

When you are configuring resource hierarchies there are a number of situations that might cause LifeKeeper to report a “No shared storage” message:

**Possible Cause:** Communications paths are not defined between the servers with the shared storage. When a hierarchy is configured on the shared storage device, LifeKeeper verifies that at least one other server in the cluster can also access the storage.

**Suggested Action:** Use the LifeKeeper GUI or `lcdstatus (1M)` to verify that communication paths are configured and that they are active.

**Possible Cause:** Communication paths are not operational between the servers with the shared storage.

**Suggested Action:** Use the LifeKeeper GUI or `lcdstatus (1M)` to verify that communication paths are configured and that they are active.

**Possible Cause:** Linux is not able to access the shared storage. This could be due to a driver not being loaded, the storage not being powered up when the driver was loaded, or the storage device is not configured properly.

**Suggested Action:** Verify that the device is properly defined in `/proc/scsi/scsi`

**Possible Cause:** The storage was not configured in Linux before LifeKeeper started. During the startup of LifeKeeper, all SCSI devices are scanned to determine the mappings for devices. If a device is configured (powered on, connected or driver loaded) after LifeKeeper is started, then LifeKeeper must be stopped and started again to be able to configure and use the device.

**Suggested Action:** Verify that the device is listed in `$(LKROOT)/subsys/scsi/resources/hostadp/device_info` where `$(LKROOT)` is by default `/opt/LifeKeeper`. If the device is not listed in this file, LifeKeeper will not try to use the device.

**Possible Cause:** The storage is not supported. The [Supported Storage List](#) shows specific SCSI devices that are supported and have been tested with LifeKeeper. However, note that this list includes known devices; there may be other devices that SIOS Technology Corp. has not tested which meet LifeKeeper requirements.

**Suggested Action:** Verify that the device is listed in `$(LKROOT)/subsys/scsi/resources/hostadp/device_info` where `$(LKROOT)` is by default `/opt/LifeKeeper`. If the device is listed in this file but the ID following the device name begins with “NU-” then LifeKeeper was unable to get a unique ID from the device. Without a unique ID LifeKeeper cannot determine if the device is shared.

**Possible Cause:** The storage may require a specific LifeKeeper software to be installed before the device can be used by LifeKeeper. Examples are the **steeleye-ikRAW** kit to enable Raw I/O support and the **steeleye-ikDR** software to enable data replication.

**Suggested Action:** Verify that the necessary LifeKeeper packages are installed on each server. See the [LifeKeeper for Linux Release Notes](#) for software requirements.

**Additional Tip:**

The `test_1k (1M)` tool can be used to help debug storage and communication problems.

## 5.4.5.8. Recovering from a LifeKeeper Server Failure

---

If any server in your LifeKeeper cluster experiences a failure that causes re-installation of the operating system (and thus LifeKeeper), you will have to re-extend the resource hierarchies from each server in the cluster. If any server in the cluster has a shared equivalency relationship with the re-installed server, however, LifeKeeper will not allow you to extend the existing resource hierarchy to the re-installed server. LifeKeeper will also not allow you to unextend the hierarchy from the re-installed server because the hierarchy does not really exist on the server that was re-installed.

### Suggested Action:

1. On each server where the resource hierarchies are configured, use the `eqv_list` command to obtain a list of all the shared equivalencies (see `LCDI-relationship` for details).

The example below shows the command and resulting output for the IP resource `iptag` on `server1` and `server2` where `server2` is the server that was re-installed and `server1` has the hierarchy configured:

```
eqv_list -f:
```

```
server1:iptag:server2:iptag:SHARED:1:10
```

2. On each server where the resource hierarchies are configured, use `eqv_remove` to manually remove the equivalency relationship for each resource in the hierarchy (see `LCDI-relationship` for details).

For example, execute the following command on `server1` using the example from step 1 above:

```
eqv_remove -t iptag -S server2 -e SHARED
```

3. In clusters with more than two servers, steps 1-2 should be repeated on each server in the cluster where equivalency relationships for these resource hierarchies are defined.
4. Finally, extend each resource hierarchy from the server where the resource hierarchy is in-service to the re-installed server using the GUI.

## 5.4.5.9. Recovering from a Non-Killable Process

---

If a process is not killable, LifeKeeper may not be able to unmount a shared disk partition. Therefore, the resource cannot be brought into service on the other system. The only way to recover from a non-killable process is to reboot the system.

## 5.4.5.10. Recovering from a Panic during a Manual Recovery

---

A PANIC during manual switchover may cause incomplete recovery. If a PANIC or other major system failure occurs during a manual switchover, complete automatic recovery to the back-up system cannot be assured. Check the backup system to make sure all resources required to be in-service are in-service. If they are not in-service, use the LifeKeeper GUI to manually bring the missing resources into service. See [Bringing a Resource In-Service](#) for instructions.

## 5.4.5.11. Recovering Out-of-Service Hierarchies

---

As a part of the recovery following the failure of a LifeKeeper server, resource hierarchies that are configured on the failed server, but are not in-service anywhere at the time of the server failure, are recovered on the highest priority alive server at the time of the failure. This is the case no matter where the out-of-service hierarchy was last in-service, including the failed server, the recovering server, or some other server in the hierarchy.

## 5.4.5.12. Resource Tag Name Restrictions

---

### Tag Name Length

All tags within LifeKeeper may not exceed the 256 character limit.

### Valid “Special” Characters

- \_ . /

The first character in a tag should not contain “.” or “/”.

### Invalid Characters

+ ; : ! @ # \$ \* = “space”

## 5.4.5.13. Serial (TTY) Console WARNING

---

If any part of the serial console data path is unreliable or goes out of service, users who have a serial (RS-232 TTY) console can experience severe problems with LifeKeeper service. During operation, LifeKeeper generates console messages. If your configuration has a serial console (instead of the standard VGA console), the entire data path from LifeKeeper to the end-user terminal must be operational in order to ensure the delivery of these console messages.

If there is any break in the data path—such as terminal powered off, modem disconnected, or cable loose—the Linux STREAMS facility queues the console message. If the STREAMS queue becomes full, the Unix kernel suspends LifeKeeper until the STREAMS buffer queue again has room for more messages. This scenario could cause LifeKeeper to HANG.

 **Note:** The use of serial consoles in a LifeKeeper environment is strongly discouraged. SIOS recommends using the VGA console. If you must use a serial console, be sure that your serial console is turned on, the cables and optional modems are connected properly, and that messages are being displayed.

## 5.4.5.14. Taking the System to init state S

### WARNING

---

When LifeKeeper is operational, the system must not be taken directly to init state S. Due to the operation of the Linux init system, such a transition causes all the LifeKeeper processes to be killed immediately and may precipitate a fastfail. Instead, you should either stop LifeKeeper manually (using `/etc/init.d/lifekeeper stop-nofailover`) or take the system first to init state 1 followed by init state S.

## 5.4.5.15. Thread is Hung Messages on Shared Storage

---

In situations where the device checking threads are not completing fast enough, this can cause messages to be placed in the LifeKeeper log stating that a thread is hung. This can cause resources to be moved from one server to another and in worse case, cause a server to be killed.

### Explanation

The FAILFASTTIMER (in `/etc/default/LifeKeeper`) defines the number of seconds that each device is checked to assure that it is functioning properly, and that all resources that are owned by a particular system are still accessible by that system and owned by it. The FAILFASTTIMER needs to be as small as possible to guarantee this ownership and to provide the highest data reliability. However if a device is busy, it may not be able to respond at peak loads in the specified time. When a device takes longer than the FAILFASTTIMER then LifeKeeper considers that device as possibly hung. If a device has not responded after 3 loops of the FAILFASTTIMER time period then LifeKeeper attempts to perform recovery as if the device has failed. The recovery process is defined by the tunable SCSIERROR. Depending on the setting of SCSIERROR the action can be a sendevent to perform local recovery and then a switchover if that fails or it can cause the system to halt.

### Suggested Action:

In cases where a device infrequently has a hung message printed to the error log followed by a message that it is no longer hung and the number in parenthesis is always 1, there should be no reason for alarm. However, if this message is frequently in the log, or the number is 2 or 3, then two actions may be necessary:

- Attempt to decrease the load on the storage. If the storage is taking longer than 3 times the FAILFASTTIMER (3 times 5 or 15 seconds by default) then one should consider the load that is being placed on the storage and re-balance the load to avoid these long I/O delays. This will not only allow LifeKeeper to check the devices frequently, but it should also help the performance of the application using that device.
- If the load can not be reduced, then the FAILFASTTIMER can be increased from the default 5 seconds. This value should be as low as possible so slowly increase the value until the messages no longer occur, or occur infrequently.

 **Note:** When the FAILFASTTIMER value is modified LifeKeeper must be stopped and restarted before the new value will take affect.

## 5.5. DataKeeper

---

SIOS DataKeeper for Linux provides an integrated data mirroring capability for LifeKeeper environments. This feature enables LifeKeeper resources to operate in shared and non-shared storage environments.

[Mirroring with SIOS DataKeeper for Linux](#)

[How SIOS DataKeeper Works](#)

## 5.5.1. Mirroring with SIOS DataKeeper for Linux

---

SIOS DataKeeper for Linux offers an alternative for customers who want to build a high availability cluster (using SIOS LifeKeeper) without shared storage or who simply want to replicate business-critical data in real-time between servers.

SIOS DataKeeper uses either synchronous or asynchronous volume-level mirroring to replicate data from the primary server (mirror source) to one or more backup servers (mirror targets).

### DataKeeper Features

SIOS DataKeeper includes the following features:

- Allows data to be reliably, efficiently and consistently mirrored to remote locations over any TCP/IP-based Local Area Network (LAN) or Wide Area Network (WAN).
- Supports synchronous or asynchronous mirroring.
- Transparent to the applications involved because replication is done at the block level below the file system.
- Supports multiple simultaneous mirror targets including cascading failover to those targets when used with LifeKeeper.
- Built-in network compression allows higher maximum throughput on Wide Area Networks.
- Supports all major file systems (see the [LifeKeeper for Linux Release Notes](#) product description for more information regarding journaling file system support).
- Provides failover protection for mirrored data.
- Integrates into the LifeKeeper Graphical User Interface.
- Fully supports other LifeKeeper Application Recovery Kits.
- Automatically resynchronizes data between the primary server and backup servers upon system recovery.
- Monitors the health of the underlying system components and performs a local recovery in the event of failure.
- Supports STONITH devices for I/O fencing. For details, refer to the [STONITH](#) topic.

# Synchronous vs. Asynchronous Mirroring

Understanding the differences between synchronous and asynchronous mirroring will help you choose the appropriate mirroring method for your application environment.

## Synchronous Mirroring

SIOS DataKeeper provides real-time mirroring by employing a synchronous mirroring technique in which data is written simultaneously on the primary and backup servers. For each write operation, DataKeeper forwards the write to the target device(s) and awaits remote confirmation before signaling I/O completion. The advantage of synchronous mirroring is a high level of data protection because it ensures that all copies of the data are always identical. However, the performance may suffer due to the wait for remote confirmation, particularly in a WAN environment.

## Asynchronous Mirroring

With asynchronous mirroring, each write is made to the source device and then a copy is queued to be transmitted to the target device(s). This means that at any given time, there may be numerous committed write transactions that are waiting to be sent from the source to the target device. The advantage of asynchronous mirroring is better performance because writes are acknowledged when they reach the primary disk, but it can be less reliable because if the primary system fails, any writes that are in the asynchronous write queue will not be transmitted to the target. To mitigate this issue, SIOS DataKeeper makes an entry to an intent log file for every write made to the primary device. If a large amount of data is written, the I/O performance may decrease temporarily because that data takes priority in the queue for transmission to the other nodes.

The intent log is a bitmap file indicating which data blocks are out of sync between the primary and target mirrors. In the event of a server failure, the intent log can be used to avoid a full resynchronization (or resync) of the data.

 A full resync synchronizes an entire disk partition.

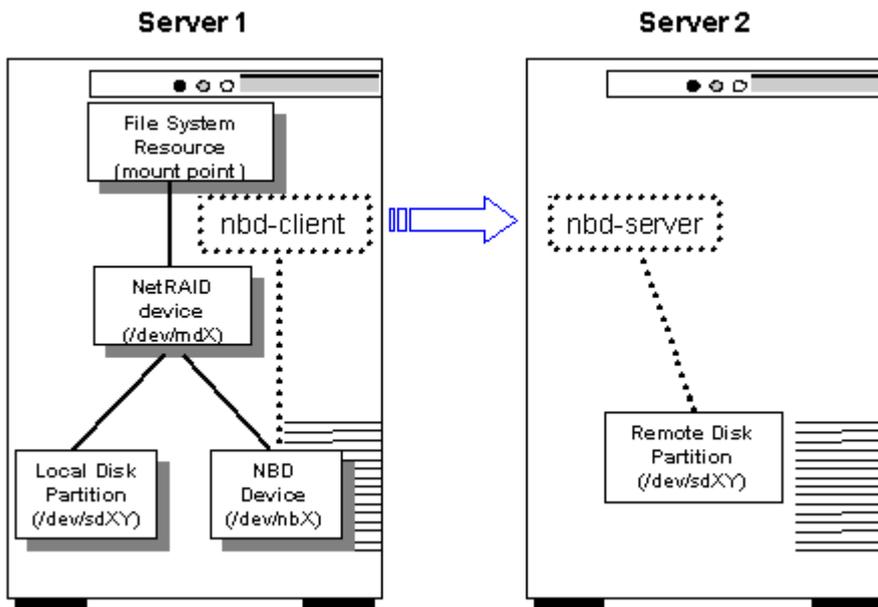
 **INFORMATION:** For async mirrors, DataKeeper allows up to 4096 outstanding target writes to be queued. Set the `LKDR_ASYNC_LIMIT` in `/etc/default/LifeKeeper` to allow more writes to be queued.

Once the number of outstanding writes to the target reaches the limit, the mirror will revert to synchronous mode until the number drops below the set value.

Assuming a 4K block size, if the value is left at the default (asynchronous limit of 4096 writes), you would have a maximum of 16MB of data in transit.

## 5.5.2. How SIOS DataKeeper Works

SIOS DataKeeper creates and protects NetRAID devices. A NetRAID device is a RAID1 device that consists of a local disk or partition and a Network Block Device (NBD) as shown in the diagram below.



A LifeKeeper supported file system can be mounted on a NetRAID device like any other storage device. In this case, the file system is called a replicated file system. LifeKeeper protects both the NetRAID device and the replicated file system.

The NetRAID device is created by building the DataKeeper resource hierarchy. Extending the NetRAID device to another server will create the NBD device and make the network connection between the two servers. SIOS DataKeeper starts replicating data as soon as the NBD connection is made.

The nbd-client process executes on the primary server and connects to the nbd-server process running on the backup server.

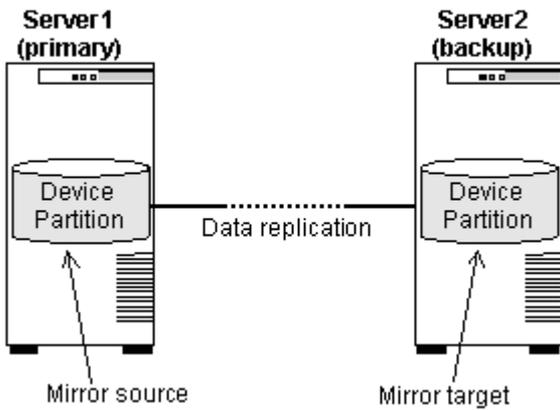
### Synchronization (and Resynchronization)

After the DataKeeper resource hierarchy is created and before it is extended, it is in a degraded mode; that is, data will be written to the local disk or partition only. Once the hierarchy is extended to the backup (target) system, SIOS DataKeeper synchronizes the data between the two systems and all subsequent writes are replicated to the target. If at any time the data gets “out-of-sync” (i.e., a system or network failure occurs) SIOS DataKeeper will automatically resynchronize the data on the source and target systems. If the mirror was configured to use an intent log (bitmap file), SIOS DataKeeper uses it to determine what data is out-of-sync so that a full resynchronization is not required. If the mirror was not configured to use a bitmap file, then a full resync is performed after any interruption of data replication.

### Standard Mirror Configuration

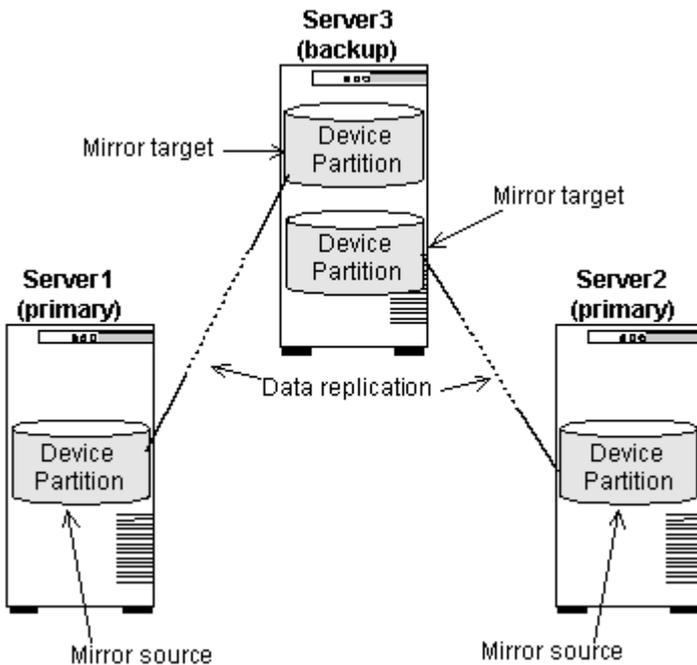
The most common mirror configuration involves two servers with a mirror established between local

disks or partitions on each server, as shown below. Server1 is considered the primary server containing the mirror source. Server2 is the backup server containing the mirror target.



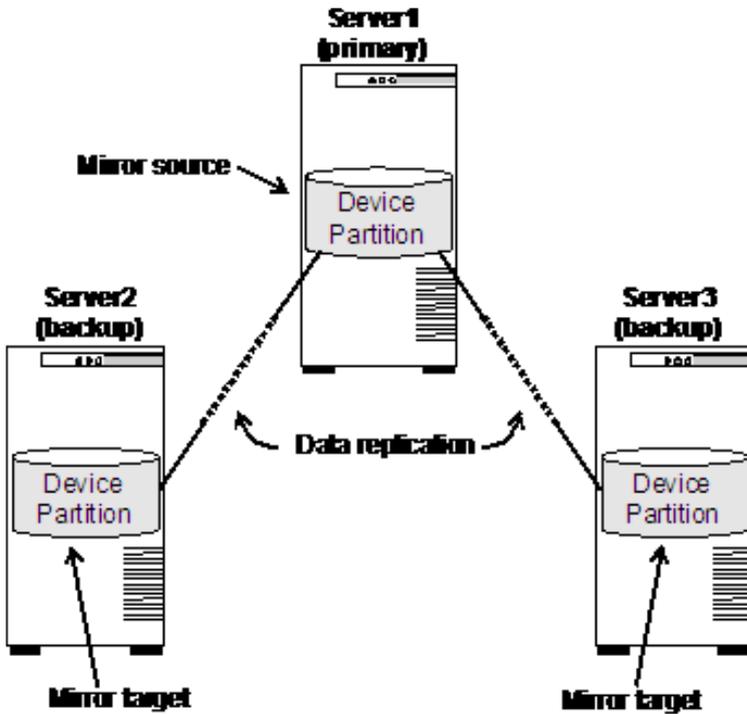
## N+1 Configuration

A commonly used variation of the standard mirror configuration above is a cluster in which two or more servers replicate data to a common backup server. In this case, each mirror source must replicate to a separate disk or partition on the backup server, as shown below.



## Multiple Target Configuration

When used with an appropriate Linux distribution and kernel version 2.6.7 or higher, SIOS DataKeeper can also replicate data from a single disk or partition on the primary server to multiple backup systems, as shown below.



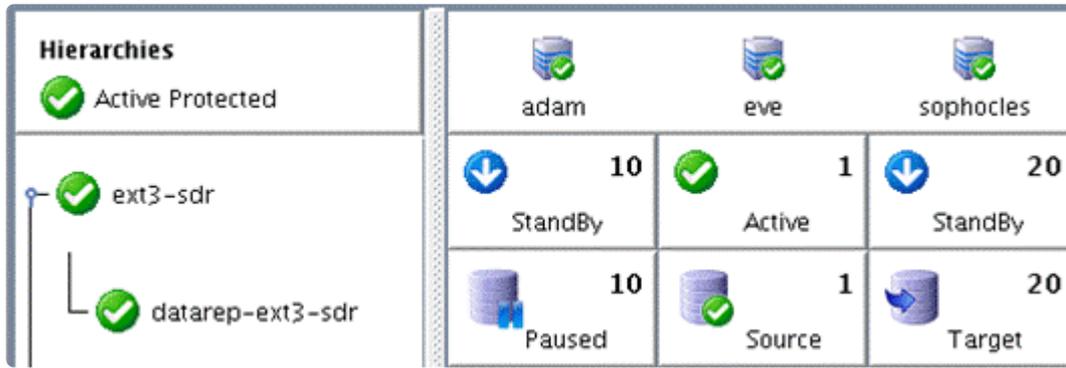
A given source disk or partition can be replicated to a maximum of 7 mirror targets, and each mirror target must be on a separate system (i.e. a source disk or partition cannot be mirrored to more than one disk or partition on the same target system).

This type of configuration allows the use of LifeKeeper’s cascading failover feature, providing multiple backup systems for a protected application and its associated data.

To avoid a full resync to all targets when a mirror is started, the bitmap from the previous source must first be merged before the remaining targets in the cluster can be reconnected. Prior to v9.3.2, if the previous source was not available when the mirror was started, a full resync was automatically done to each target. Starting with v9.3.2, when the mirror is started on a system it will wait for the previous source to join the cluster before connecting targets. When the previous source joins the cluster, its bitmap is merged so that all targets can join with a partial resync. When the mirror is stopped and targets are in-sync, no previous source is needed to start the mirror and replicate to targets. If the previous source is not available to rejoin the cluster, targets can manually be resynced with a full resync using the “mirror\_action fullresync” command. The variable `LKDR_WAIT_FOR_PREVIOUS_SOURCE_TIMEOUT` in `/etc/default/LifeKeeper` determines the resync behavior (refer to the [DataKeeper Parameters List](#) for more information).

## SIOS DataKeeper Resource Hierarchy

The following example shows a typical DataKeeper resource hierarchy as it appears in the LifeKeeper GUI:



The resource *datarep-ext3-sdr* is the NetRAID resource, and the parent resource *ext3-sdr* is the file system resource. Note that subsequent references to the DataKeeper resource in this documentation refer to both resources together. Because the file system resource is dependent on the NetRAID resource, performing an action on the NetRAID resource will also affect the file system resource above it.

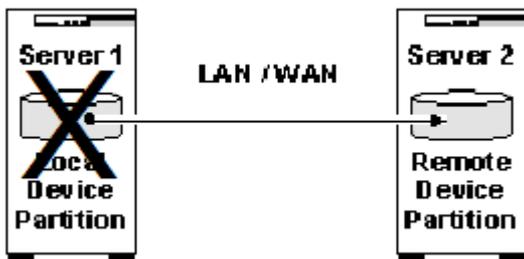
## Failover Scenarios

### Failover Scenarios – 2 nodes

The following four examples show what happens during a failover using SIOS DataKeeper. In these examples, the LifeKeeper for Linux cluster consists of two servers, Server 1 (primary server) and Server 2 (backup server).

#### Scenario 1

Server 1 has successfully completed its replication to Server 2 after which Server 1 becomes inoperable.



**Result:** Failover occurs. Server 2 now takes on the role of primary server and operates in a degraded mode (with no backup) until Server 1 is again operational. SIOS DataKeeper will then initiate a resynchronization from Server 2 to Server 1. This will be a full resynchronization on kernel 2.6.18 and lower. On kernels 2.6.19 and later or with Red Hat Enterprise Linux 5.4 kernels 2.6.18-164 or later (or a supported derivative of Red Hat 5.4 or later), the resynchronization will be partial, meaning only the changed blocks recorded in the bitmap files on the source and target will need to be synchronized.

**Note:** SIOS DataKeeper sets the following flag on the server that is currently acting as the mirror source:

```
$LKROOT/subsys/scsi/resources/netraid/$TAG_last_owner
```

When Server 1 fails over to Server 2, this flag is set on Server 2. Thus, when Server 1 comes back up; SIOS DataKeeper removes the last owner flag from Server 1. It then begins resynchronizing the data from Server 2 to Server 1.

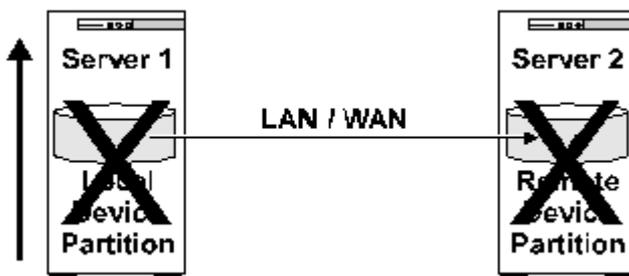
## Scenario 2

Considering scenario 1, Server 2 (still the primary server) becomes inoperable during the resynchronization with Server 1 (now the backup server).

**Result:** Because the resynchronization process did not complete successfully, there is potential for data corruption. As a result, LifeKeeper will not attempt to fail over the DataKeeper resource to Server 1. Only when Server 2 becomes operable will LifeKeeper attempt to bring the DataKeeper resource in-service (ISP) on Server 2.

## Scenario 3

Both Server 1 (primary) and Server 2 (target) become inoperable. Server 1 (primary) comes back up first.



**Result:** Server 1 will not bring the DataKeeper resource in-service. The reason is that if a source server goes down, and then it cannot communicate with the target after it comes back online, it sets the following flag:

```
$LKROOT/subsys/scsi/resources/netraid/$TAG_data_corrupt
```

This is a safeguard to avoid resynchronizing data in the wrong direction. In this case you will need to force the mirror online on Server 1, which will delete the data\_corrupt flag and bring the resource into service on Server 1. [See Force Mirror Online](#)

**Note:** The user must be certain that Server 1 was the last primary before removing the \$TAG\_data\_corrupt file. Otherwise data corruption might occur. You can verify this by checking for the presence of the last\_owner flag.

## Scenario 4

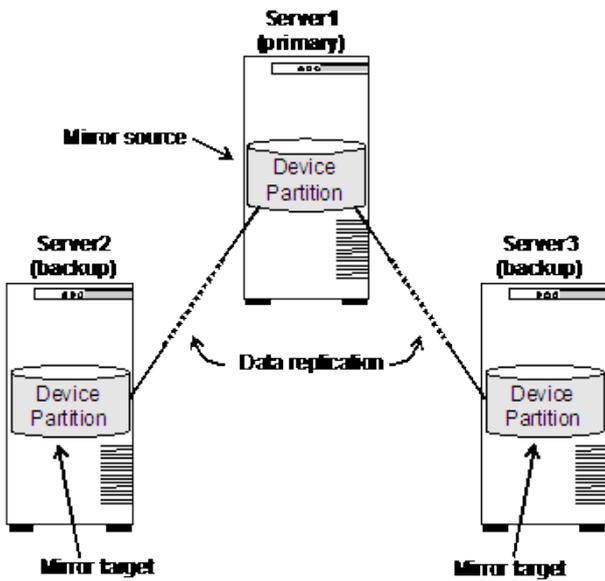
Both Server 1 (primary) and Server 2 (target) become inoperable. Server 2 (target) comes back up first.



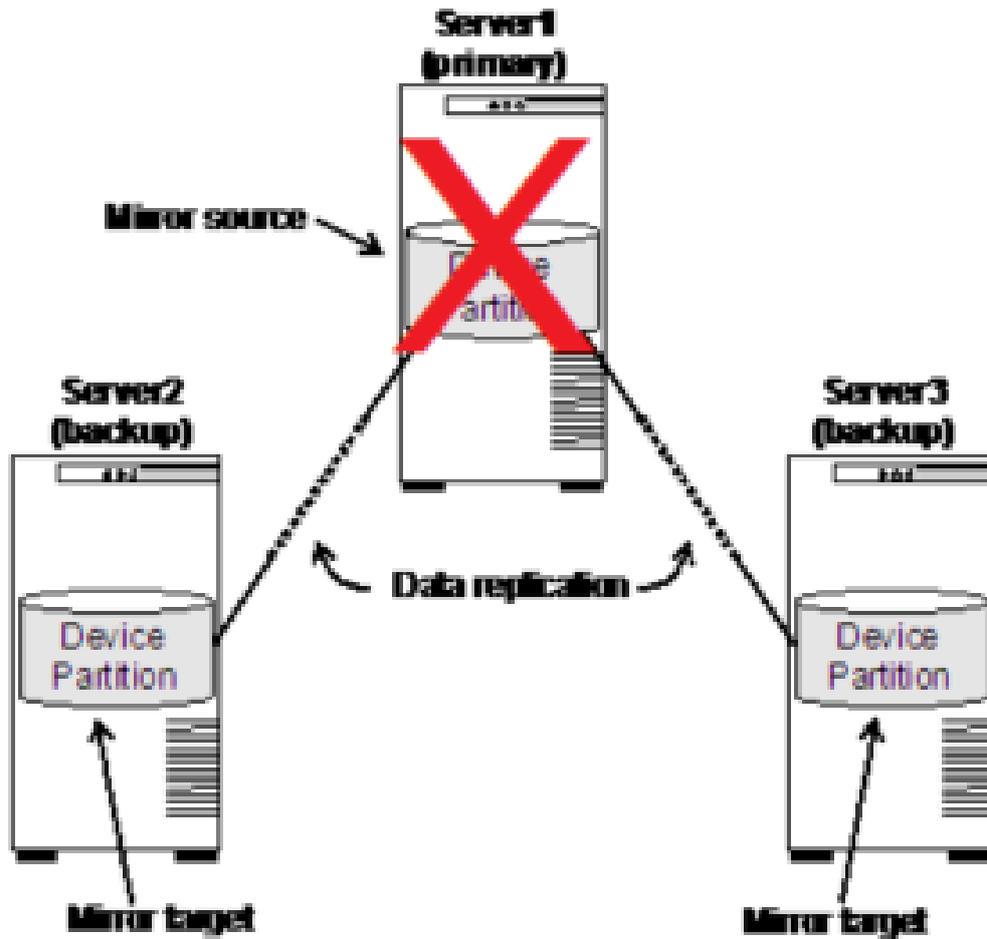
**Result:** LifeKeeper will not bring the DataKeeper resource ISP on Server 2. When Server 1 comes back up, LifeKeeper will automatically bring the DataKeeper resource ISP on Server 1.

### Failover Scenario – 3 nodes

The following example shows what happens during a failover using SIOS DataKeeper. In this example the LifeKeeper for Linux cluster consists of three servers, Server 1 (primary server), Server 2 (backup server) and Server 3 (backup server).



Server 1 (priority 1) has successfully completed its replication to Server 2 (priority 10) and Server 3 (priority 20) after which Server 1 becomes inoperable.



**Result:** Failover occurs to the next highest priority, Server 2. Server 2 now takes on the role of primary server. Prior to release v9.3.2 Server 3 will be added to the mirror with a full resynchronization. With v9.3.2 Server 2 **waits** for Server 1 (previous server) to return to the cluster before resuming replication to Server 3. This allows the bitmap from Server 1 to be merged with the bitmap from Server 2, allowing for a partial resync to both Server 1 and Server 3. While waiting for Server 1 to reconnect to the cluster, the LifeKeeper GUI will show the status of Server 3 as “Out of Sync (Wait for Previous Source)”. The status of Server 1 will be “Unknown” while the server is not connected. When it initially connects the GUI status will show “Out of Sync”. The properties page for the mirror will identify it as “Out of Sync (Previous Source)”.

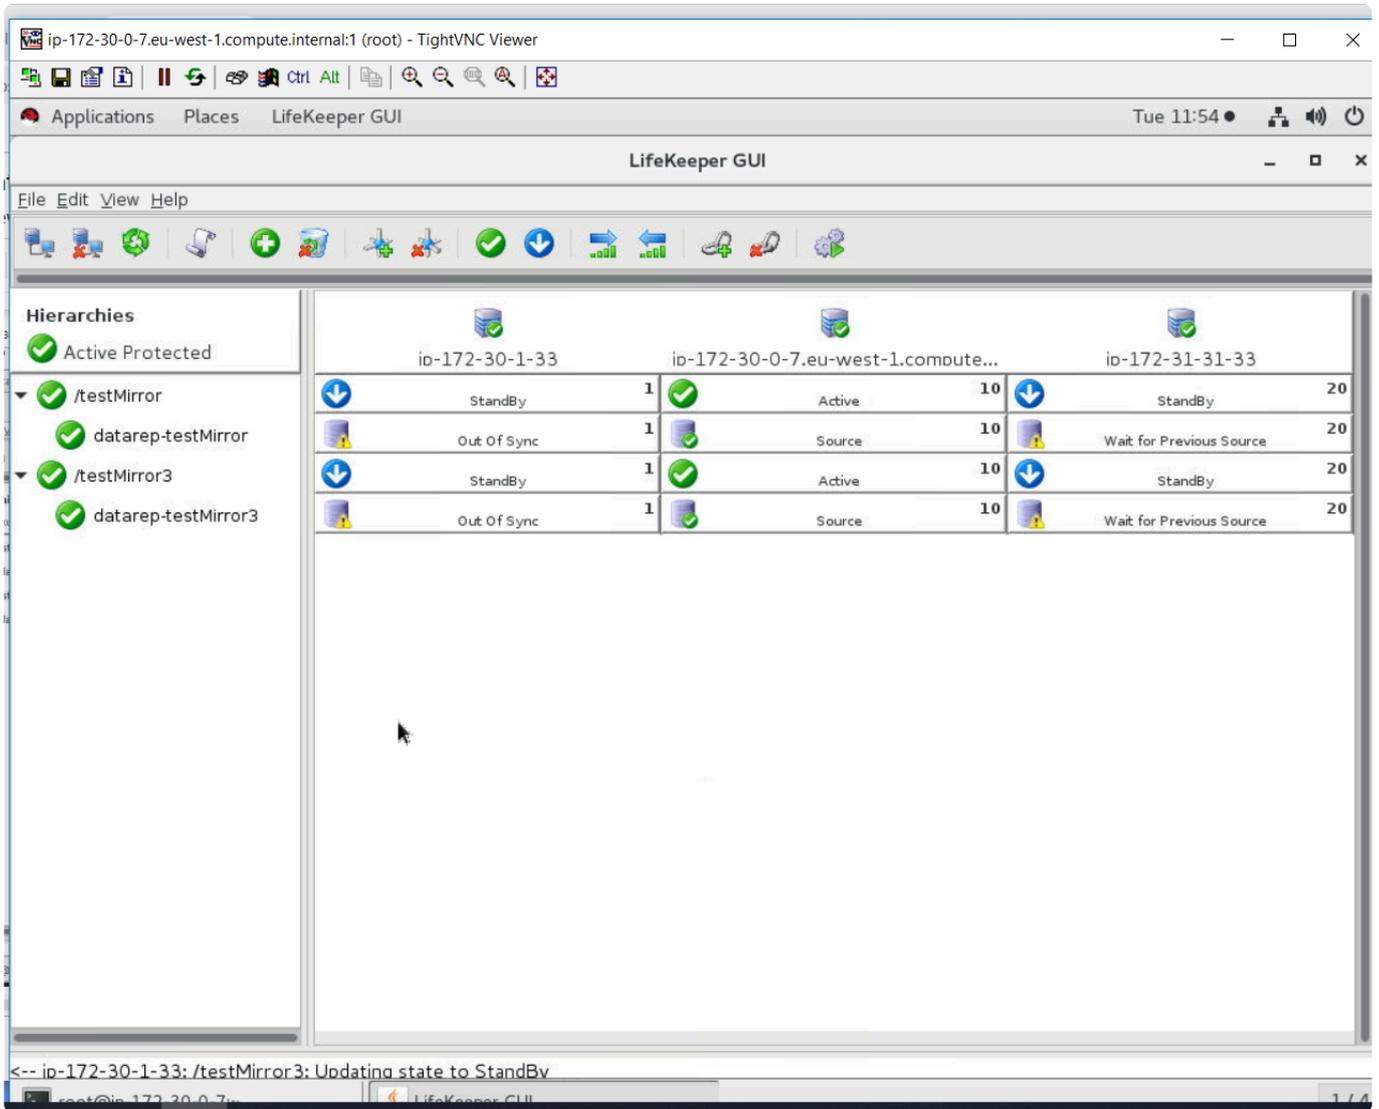
The screenshot shows the LifeKeeper GUI interface. A dialog box titled "Resource Properties for datarep-testMirror3" is open, displaying the following information:

- Select Resource: datarep-testMirror3
- Select Server for Resource: ip-172-30-0-7.eu-west-1.c...
- Replication Status: General Equivalencies Relations
- Mirror Configuration:**
  - ip-172-30-0-7.eu-west-1.compute.internal → ip-172-31-31-33 (172.31.31.33)  
Status: Out of Sync (Wait for Previous Source)  
Type: Synchronous
  - ip-172-30-0-7.eu-west-1.compute.internal → ip-172-30-1-33 (172.30.1.33)  
Status: Out of Sync (Previous Source)  
Type: Synchronous
- Bitmap: 4092 bits (chunks), 9 dirty (0.2%)

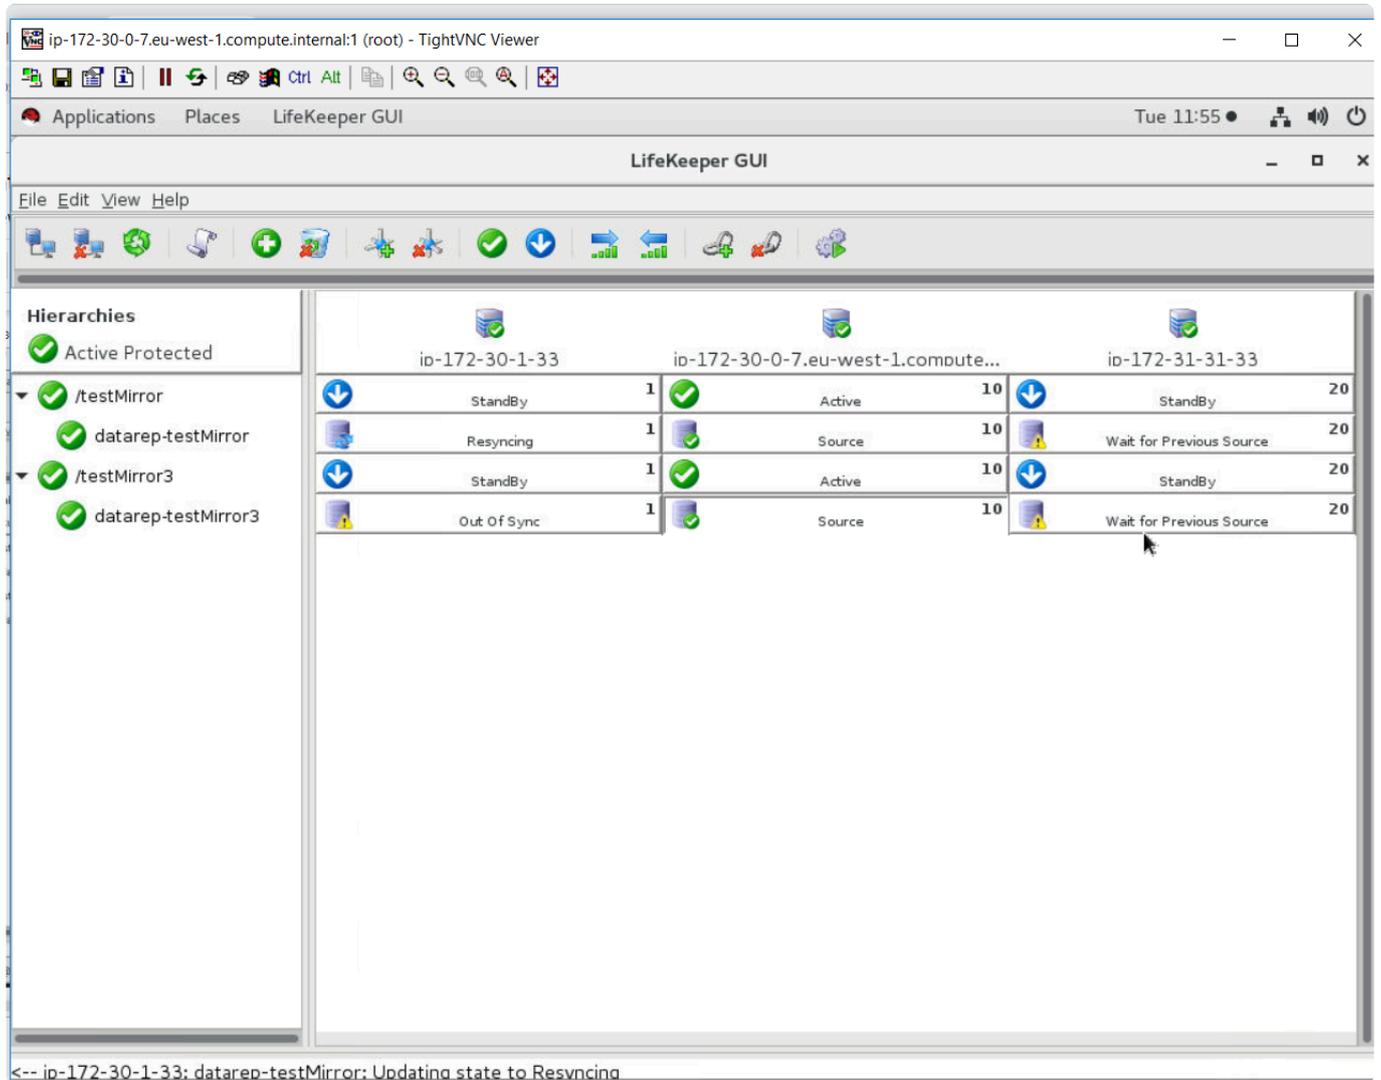
The background interface shows a "Hierarchies" panel on the left with a tree structure including "Backup Not Stan", "/testMirror", "datarep-testM", "/testMirror3", and "datarep-testM". On the right, a table lists operations for resource "id-172-31-31-33":

StandBy	20
Wait for Previous Source	20
StandBy	20
Wait for Previous Source	20

The bottom status bar shows "root@ip-172-30-0-7:~" and "LifeKeeper GUI" with a page indicator "1 / 4".



Once Server 1 reconnects, its bitmap is merged and a resynchronization begins, at which point its status is shown as “Resyncing”.



When resynchronization completes, its status will update to “Target” and Server 3 will begin resynchronization with its status set to “Resyncing”.

**Note:** SIOS DataKeeper sets the follow flags to track the mirror source:

`$LKROOT/subsys/scsi/resources/netraid/$TAG_last_owner`

`$LKROOT/subsys/scsi/resources/netraid/$TAG_source`

The `$TAG_last_owner` flag is on the system that is currently acting as the mirror source while the `$TAG_source` flag contains the name of the system that was source at the last point in time that the local node was part of the mirror.

When Server 1 fails over to Server 2, `$TAG_last_owner` flag is set on Server 2. The `$TAG_source` flag on Server 2 identifies Server 1 as the previous source (that has the bitmap needed to do a partial resync to Server 1 and Server 3). When Server 1 comes back up, SIOS DataKeeper removes the `$TAG_last_owner` flag from Server 1. Server 2 then merges the bitmap from Server 1 and begins resynchronizing the data from Server 2 to Server 1. When resynchronization is complete to Server 1 the `$TAG_source` flag on Server 1 is updated with the name of Server 2. After Server 1 is synchronized, Server 2 will perform the same resynchronization to Server 3. When that resynchronization is complete to Server 3 the `$TAG_source` flag on Server 3 is updated with the name of Server 2.

## 5.5.3. SIOS DataKeeper Installation and Configuration

---

### Installing and Configuring SIOS DataKeeper for Linux

[Hardware/Software Requirements](#)

### Before Configuring Your DataKeeper Resources

The following topics contain information for consideration before beginning to create and administer your DataKeeper resources. They also describe the three types of DataKeeper resources. Please refer to the [LifeKeeper Configuration](#) section for instructions on configuring LifeKeeper Core resource hierarchies.

---

[Requirements](#)

[General Configuration](#)

[Network Configuration](#)

[Changing the Data Replication Path](#)

[Network Bandwidth Requirements](#)

[Measuring Rate of Change on a Linux System](#)

[WAN Configuration](#)

[Resource Types](#)

[I/O Fencing with DataKeeper Configuration](#)

[Resource Configuration Tasks](#)

[Creating a Resource Hierarchy](#)

[Extending Your Hierarchy](#)

[Unextending Your Hierarchy](#)

[Deleting a Resource Hierarchy](#)

[Taking a Resource Out of Service](#)

[Bringing a Resource In Service](#)

[Testing Your Resource Hierarchy](#)

## 5.5.3.1. Hardware and Software Requirements

---

Your LifeKeeper configuration should meet the following requirements prior to the installation of SIOS DataKeeper.

### Hardware Requirements

- **Servers** – Two or more LifeKeeper for Linux supported servers.
- **IP Network Interface Cards** – Each server requires at least one network interface card. Remember, however, that a LifeKeeper cluster requires two communication paths; two separate LAN-based communication paths using dual independent sub-nets are recommended, and at least one of these should be configured as a private network. However using a combination of TCP and TTY is also supported.

✿ **Note:** Due to the nature of software mirroring, network traffic between servers can be heavy. Therefore, it is recommended that you implement a separate private network for your SIOS DataKeeper devices which may require additional network interface cards on each server.

- **Disks or Partitions** – Disks or partitions on the primary and backup servers that will act as the source and target disks or partitions. The target disks or partitions must be at least as large as the source disk or partition.

✿ **Note:** With the release of SIOS Data Replication 7.1.1, it became possible to replicate an entire disk, one that has not been partitioned (i.e. `/dev/sdd`). Previous versions of SIOS Data Replication required that a disk be partitioned (even if it was a single large partition; i.e. `/dev/sdd1_`) before it could be replicated. SIOS Data Replication 7.1.1 removed that restriction.

✿ With the release of LifeKeeper for Linux v9.5.0, all disks must be uniquely identifiable. DataKeeper had allowed the device name (i.e. `/dev/sdd`) to be used to identify a device but in some situations the device names can change that can lead to data corruption. The use of a GPT partition table can provide a unique identifier.

### Software Requirements

- **Operating System** – SIOS DataKeeper can be used with any major Linux distribution based on the 2.6 Linux kernel. See the [LifeKeeper for Linux Release Notes](#) for a list of supported distributions. Asynchronous mirroring and intent logs are supported only on distributions that use a 2.6.16 or later Linux kernel. Multiple target support (i.e., support for more than 1 mirror target) requires a 2.6.7 or later Linux kernel.

- **LifeKeeper Installation Script** – In most cases, you will need to install the following package (see the “Product Requirements” section in the [LifeKeeper for Linux Release Notes](#) for specific SIOS DataKeeper requirements):

#### **HADR-generic-2.6**

This package must be installed on each server in your LifeKeeper cluster prior to the installation of SIOS DataKeeper. The HADR package is located on the LifeKeeper Installation Image File, and the appropriate package is automatically installed by the Installation **setup** script.

- **LifeKeeper Software** – You must install the same version of the LifeKeeper Core on each of your servers. You must also install the same version of each recovery kit that you plan to use on each server. See the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.
- **SIOS DataKeeper software** – Each server in your LifeKeeper cluster requires SIOS DataKeeper software. Please see the [LifeKeeper for Linux Installation Guide](#) for specific instructions on the installation and removal of SIOS DataKeeper.

## 5.5.3.2. General Configuration

---

- The size of the target disks or partitions (on the backup servers) must be equal to or greater than the size of the source disk or partition (on the primary server).
- Once the DataKeeper resource is created and extended, the synchronization process will delete existing data on the target disks or partitions and replace it with data from the source partition.

## 5.5.3.3. DataKeeper for Linux Network Configuration

---

- The network path that is chosen for data replication between each pair of servers must also already be configured as a LifeKeeper communication path between those servers. To change the network path, see [Changing the Data Replication Path](#).
- Avoid a configuration to add virtual IP address with IP Recovery Kit to the network interface that DataKeeper users for data replication. Since the communication line is temporarily disconnected while IP Recovery Kit uses the network interface, data replication may stop at unexpected timing and unnecessary resynchronization may occur.
- This release of SIOS DataKeeper does not support Automatic Switchback for DataKeeper resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource sitting on top of a DataKeeper resource.
- If using Fusion-io, see the Network section of [Clustering with Fusion-io](#) for further network configuration information.

## 5.5.3.4. DataKeeper Events Table

---

The following table contains the list of DataKeeper specific events and associated trap numbers. The events will generate email notices when `LK_NOTIFY_ALIAS` is set.

LifeKeeper Event/Description	Trap #	Object ID
<b>DataKeeper Target Disk Mounted</b> Sent from the in-service node when the target disk of a DataKeeper mirror is directly mounted on the target.	142	.1.3.6.1.4.1.7359.1.0.142
<b>DataKeeper Resynchronization Complete</b> Sent from the in-service node when the DataKeeper mirror has resynchronized with a target.	143	.1.3.6.1.4.1.7359.1.0.143

## 5.5.3.5. Changing the Data Replication Path

Starting with LK 7.1, IP addresses for mirror endpoints can be modified using `lk_chg_value`. For example, to change a mirror endpoint from IP address 192.168.0.1 to 192.168.1.1:

```
# lkstop (lk_chg_value cannot be run while LifeKeeper is running)

# lk_chg_value -o 192.168.0.1 -n 192.168.1.1

# lkstart
```

Execute these commands on all servers involved in the mirror(s) that are using this IP address.



**Note:** This command will also modify communication paths that are using the address in question.

## 5.5.3.6. Network Bandwidth Requirements

---

Prior to installing SIOS DataKeeper, you should determine the network bandwidth requirements for replicating your current configuration whether you are employing virtual machines or using physical Linux servers. If you are employing virtual machines (VMs), use the method [Measuring Rate of Change on a Linux System \(Physical or Virtual\)](#) to measure the rate of change for the virtual machines that you plan to replicate. This value indicates the amount of network bandwidth that will be required to replicate the virtual machines.

After determining the network bandwidth requirements, ensure that your network is configured to perform optimally. If your network bandwidth requirements are above your current available network capacity, you may need to consider one or more of the following options:

- Enable compression in SIOS DataKeeper (or in the network hardware, if possible)
- Increase your network capacity
- Reduce the amount of data being replicated
- Create a local, non-replicated storage repository for temporary data and swap files
- Manually schedule replication to take place daily at off-peak hours

## 5.5.3.6.1. Measuring Rate of Change on a Linux System (Physical or Virtual)

---

Data can be replicated across any available network. In Wide Area Network (WAN) configurations, special consideration must be given to the question, “Is there sufficient bandwidth to successfully replicate the partition and keep the mirror in the mirroring state as the source partition is updated throughout the day?”

Keeping the mirror in the mirroring state is critical because a switchover of the partition is not allowed unless the mirror is in the mirroring state.

Short bursts of write activity are handled by adding the data to the async queue. However, make sure that over any extended period of time, the disk write activity for all replicated volumes combined remains, on average, below the amount of change that your network can transmit.

If the network capacity is not sufficient to keep up with the rate of change that occurs on your disks, and the async queue fills up, the mirror will revert to synchronous behavior, which can negatively affect performance of the source server.

### Measuring Basic Rate of Change

Use the following command to determine file(s) or partition(s) to be mirrored. For example `/dev/sda3`, and then measure the amount of data written in a day:

```
MB_START=`awk '/sda3 / { print $10 / 2 / 1024 }' /proc/diskstats`
```

... wait for a day ...

```
MB_END=`awk '/sda3 / { print $10 / 2 / 1024 }' /proc/diskstats`
```

The daily rate of change, in MB, is then `MB_END - MB_START`.

SIOS DataKeeper can mirror daily, approximately:

T1 (1.5Mbps) – 14,000 MB/day (14 GB)

T3 (45Mbps) – 410,000 MB/day (410 GB)

Gigabit (1Gbps) – 5,000,000 MB/day (5 TB)

### Measuring Detailed Rate of Change

The best way to collect Rate of Change data is to log disk write activity for some period of time (one day, for instance) to determine what the peak disk write periods are.

To track disk write activity, create a cron job which will log the timestamp of the system followed by a dump of `/proc/diskstats`. For example, to collect disk stats every two minutes, add the following link to `/etc/crontab`:

```
*/2 * * * * root ( date ; cat /proc/diskstats ) >> /path_to/  
filename.txt
```

... wait for a day, week, etc ... then disable the cron job and save the resulting data file in a safe location.

## Analyze Collected Detailed Rate of Change Data

The `roc-calc-diskstats` utility analyzes data collected in the previous step. This utility takes a `/proc/diskstats` output file that contains output, logged over time, and calculates the rate of change of the disks in the dataset.

[Click Here](#) to download `roc-calc-diskstats`

### Usage:

```
# ./roc-calc-diskstats <interval> <start_time> <diskstats-data-file> [dev-  
list]
```

### Usage Example (Summary only):

```
# ./roc-calc-diskstats 2m "Jul 22 16:04:01" /root/diskstats.txt  
sdb1,sdb2,sdc1 > results.txt
```

The above example dumps a summary (with per disk peak I/O information) to `results.txt`

### Usage Example (Summary + Graph Data):

```
# export OUTPUT_CSV=1  
  
# ./roc-calc-diskstats 2m "Jul 22 16:04:01" /root/diskstats.txt  
sdb1,sdb2,sdc1 2> results.csv > results.txt
```

The above example dumps graph data to `results.csv` and the summary (with per disk peak I/O information) to `results.txt`

### Example Results (from results.txt)

```
Sample start time: Tue Jul 12 23:44:01 2011
```

```
Sample end time: Wed Jul 13 23:58:01 2011
```

Sample interval: 120s #Samples: 727 Sample length: 87240s

(Raw times from file: Tue Jul 12 23:44:01 EST 2011, Wed Jul 13 23:58:01 EST 2011)

Rate of change for devices dm-31, dm-32, dm-33, dm-4, dm-5, total

dm-31 peak:0.0 B/s (0.0 b/s) (@ Tue Jul 12 23:44:01 2011) average:0.0 B/s (0.0 b/s)

dm-32 peak:398.7 KB/s (3.1 Mb/s) (@ Wed Jul 13 19:28:01 2011) average:19.5 KB/s (156.2 Kb/s)

dm-33 peak:814.9 KB/s (6.4 Mb/s) (@ Wed Jul 13 23:58:01 2011) average:11.6 KB/s (92.9 Kb/s)

dm-4 peak:185.6 KB/s (1.4 Mb/s) (@ Wed Jul 13 15:18:01 2011) average:25.7 KB/s (205.3 Kb/s)

dm-5 peak:2.7 MB/s (21.8 Mb/s) (@ Wed Jul 13 10:18:01 2011) average:293.0 KB/s (2.3 Mb/s)

**total peak:2.8 MB/s (22.5 Mb/s) (@ Wed Jul 13 10:18:01 2011) average:349.8 KB/s (2.7 Mb/s)**

## Graph Detailed Rate of Change Data

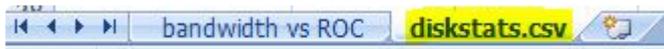
To help understand your specific bandwidth needs over time, SIOS has created a template spreadsheet called diskstats-template.xlsx. This spreadsheet contains sample data which can be overwritten with the data collected by roc-calc-diskstats.

[Click Here](#) to download diskstats-template.xlsx

1. Open results.csv, and select **all rows**, including the total column.

	A	B	C	D	E	F	G	H	I	J	K	L
1	dm-31	0	0	0	0	0	0	0	0	0	0	0
2	dm-32	3549.867	6826.667	3549.867	273.0667	7099.733	3549.867	6826.667	3686.4	341.3333	7099.733	3549.8
3	dm-33	3857.067	4505.6	3310.933	1911.467	4846.933	2935.467	4471.467	3310.933	1911.467	4710.4	2935.4
4	dm-4	2218.667	2389.333	2150.4	1570.133	2525.867	2116.267	2389.333	2013.867	4300.8	2833.067	1809.0
5	dm-5	25326.93	26683.73	23449.6	25164.8	26419.2	23048.53	28612.27	21367.47	35140.27	32145.07	40797.
6	total	34952.53	40405.33	32460.8	28919.47	40891.73	31650.13	42299.73	30378.67	41693.87	46788.27	49092.

2. Open **diskstats-template.xlsx**, select the **diskstats.csv** worksheet.



3. In cell 1-A, right-click and select **Insert Copied Cells**.

4. Adjust the **bandwidth** value in the cell towards the bottom left of the worksheet to reflect an amount of bandwidth you have allocated for replication.

Units: Megabits/second (Mb/sec)

 **Note:** The cells to the right will automatically be converted to bytes/sec to match the raw data collected.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
1	dm-31	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	dm-32	3545.86667	6826.66667	3549.867	273.0667	7099.733	3549.867	6826.667	3686.4	341.3333	7099.733	3549.867	6826.667	3276.8	273.0667	6826.667	3549.867	6826.667
3	dm-33	3827.06667	4205.0	3310.333	1911.467	4846.933	2935.467	4471.467	3310.333	1911.467	4710.4	2935.467	4710.4	2935.467	2798.333	4878.267	3837.067	4710.4
4	dm-4	2218.66667	2389.33333	2150.4	1570.133	2525.867	2116.267	2389.333	2013.867	4300.8	2833.067	1809.067	27955.2	1570.133	2286.933	2525.867	2116.267	2628.267
5	dm-5	25326.9333	26683.73333	23449.6	25164.8	26419.2	23048.53	28612.27	21367.47	35140.27	32145.07	40797.87	28492.8	23338.67	28561.07	27667.73	27784.53	29849.6
6	total	34952.5333	40405.33333	32468.8	28919.47	40891.73	31650.13	42299.73	30378.67	41893.87	46788.27	49092.27	67985.87	31121.07	33920	41096.53	37507.73	44014.93
7																		
8	bandwidth (Mb/s)																	
9		10	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720					

5. Make a note of the following row/column numbers:

- a. Total (row 6 in screenshot below)
- b. Bandwidth (row 9 in screenshot below)
- c. Last datapoint (column R in screenshot below)

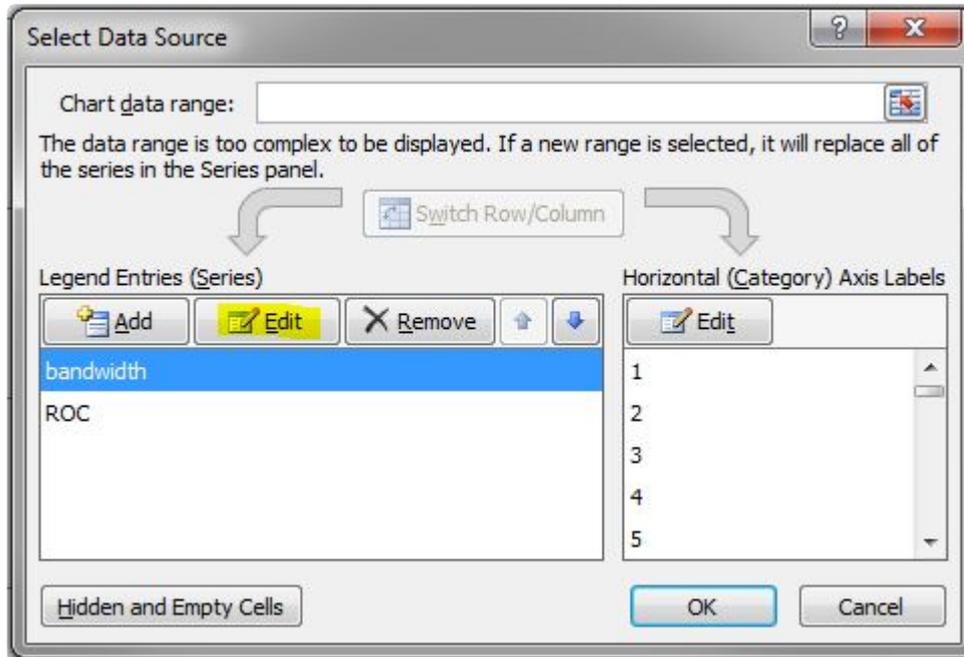
	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
1	dm-31	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	dm-32	3549.86667	6826.66667	3549.867	273.0667	7099.733	3549.867	6826.667	3686.4	341.3333	7099.733	3549.867	6826.667	3276.8	273.0667	6826.667	3549.867	6826.667
3	dm-33	3827.06667	4205.0	3310.333	1911.467	4846.933	2935.467	4471.467	3310.333	1911.467	4710.4	2935.467	4710.4	2935.467	2798.333	4878.267	3837.067	4710.4
4	dm-4	2218.66667	2389.33333	2150.4	1570.133	2525.867	2116.267	2389.333	2013.867	4300.8	2833.067	1809.067	27955.2	1570.133	2286.933	2525.867	2116.267	2628.267
5	dm-5	25326.9333	26683.73333	23449.6	25164.8	26419.2	23048.53	28612.27	21367.47	35140.27	32145.07	40797.87	28492.8	23338.67	28561.07	27667.73	27784.53	29849.6
6	total	34952.5333	40405.33333	32468.8	28919.47	40891.73	31650.13	42299.73	30378.67	41893.87	46788.27	49092.27	67985.87	31121.07	33920	41096.53	37507.73	44014.93
7																		
8	bandwidth (Mb/s)																	
9		10	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720	1310720

6. Select the **bandwidth vs ROC** worksheet.



7. Right-click on the graph and select **Select Data...**

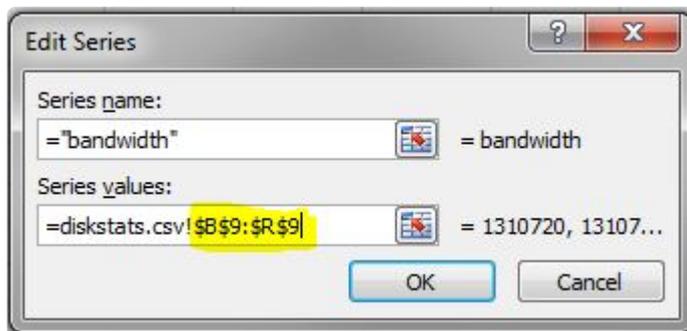
- a. Adjust **Bandwidth Series**
  - i. From the **Series** list on the left, select **bandwidth**
  - ii. Click **Edit**



iii. Adjust the **Series Values**: field with the following syntax:

```
"=diskstats.csv!$B$<row>:$<final_column>$<row>"
```

example: `"=diskstats.csv!$B$9:$R:$9"`

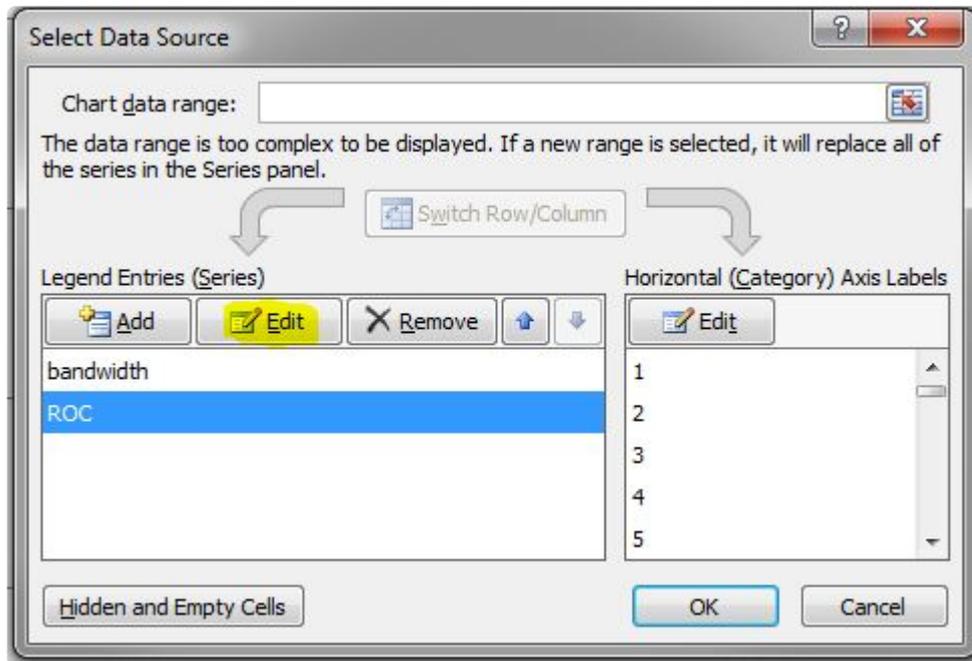


iv. Click **OK**

b. Adjust **ROC Series**

i. From the **Series** list on the left, select **ROC**

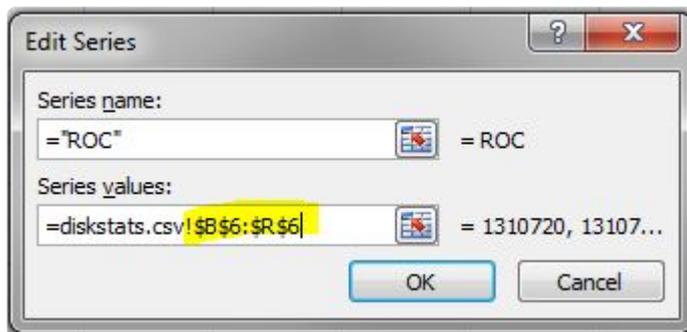
ii. Click **Edit**



iii. Adjust the **Series Values:** field with the following syntax:

```
"=diskstats.csv!$B$<row>:$<final_column>$<row>"
```

example: `"=diskstats.csv!$B$6:$R:$6"`



iv. Click **OK**

c. Click **OK** to exit the Wizard

8. The Bandwidth vs ROC graph will update. Please analyze your results to determine if you have sufficient bandwidth to support replication of your data.

# roc-calc-diskstats

```
#!/usr/bin/perl
# Copyright (c) 2011, SIOS Technology, Corp.
# Author: Paul Clements
use strict;
sub msg {
    printf STDERR _;
}
sub dbg {
    return if (! $ENV{'ROC_DEBUG'});
    msg _;
}
$0 =~ s@^\.*/@@; # basename
sub usage {
    msg "Usage: $0 <interval> <start-time> <iostat-data-file> [dev-list]\n";
    msg "\n";
    msg "This utility takes a /proc/diskstats output file that contains\n";
    msg "output, logged over time, and calculates the rate of change of\n";
    msg "the disks in the dataset\n";
    msg "OUTPUT_CSV=1 set in env. dumps the full stats to a CSV file on STDERR\n";
    msg "\n";
    msg "Example: $0 1hour \"jun 23 12pm\" steeleye-iostat.txt sdg,sdh\n";
    msg "\n";
    msg "interval - interval between samples\n";
    msg "start time - the time when the sampling starts\n";
    msg "iostat-data-file - collect this with a cron job like:\n";
    msg "\t0 * * * * (date ; cat /proc/diskstats) >> /root/diskstats.txt\n";
    msg "dev-list - list of disks you want ROC for (leave blank for all)\n";
    exit 1;
}
usage if (ARGV < 3);
my $interval = TimeHuman($ARGV[0]);
my $starttime = epoch($ARGV[1]);
my $file = $ARGV[2];
my $blksize = 512; # /proc/diskstats is in sectors
my %devs = map { $_ => 1 } split /,/, $ARGV[3];
my %stat;
my $firsttime;
my $lasttime;
```

```

# datestamp divides output
my %days = ( 'Sun' => 1, 'Mon' => 1, 'Tue' => 1, 'Wed' => 1,
              'Thu' => 1, 'Fri' => 1, 'Sat' => 1);
my %fields = ( 'major'   => 0,
              'minor'   => 1,
              'dev'     => 2,
              'reads'   => 3,
              'reads_merged' => 4,
              'sectors_read' => 5,
              'ms_time_reading' => 6,
              'writes'  => 7,
              'writes_merged' => 8,
              'sectors_written' => 9,
              'ms_time_writing' => 10,
              'ios_pending' => 11,
              'ms_time_total' => 12,
              'weighted_ms_time_total' => 13 );
my $devfield = $fields{'dev'};
my $scalffield = $ENV{'ROC_CALC_FIELD'} || $fields{'sectors_written'};
dbg "using field $scalffield\n";
open(FD, "$file") or die "Cannot open $file: $!\n";
foreach (<FD>) {
    chomp;
    _ = split;
    if (exists($days{$_[0]})) { # skip datestamp divider
        if ($firsttime eq '') {
            $firsttime = join ' ', _[0..5];
        }
        $lasttime = join ' ', _[0..5];
        next;
    }
    next if ($_[0] !~ /[0-9]/); # ignore
    if (!%devs || exists $devs{$_[$devfield]}) {
        push {$stat{$_[$devfield]}}, $_[$scalffield];
    }
}
{$stat{'total'}} = totals(\%stat);
printf "Sample start time: %s\n", scalar(localtime($starttime));
printf "Sample end time: %s\n", scalar(localtime($starttime + (($stat{'total'} - 1) * $interval)));
printf "Sample interval: %ss #Samples: %s Sample length: %ss\n", $interval,
(($stat{'total'} - 1), (($stat{'total'} - 1) * $interval);
print "(Raw times from file: $firsttime, $lasttime)\n";
print "Rate of change for devices " . (join ' ', sort keys %stat) . "\n";
foreach (sort keys %stat) {
    my vals = {$stat{$_}};
    my ($max, $maxindex, $roc) = roc($_, $blksize, $interval, vals);
}

```

```

        printf "$_ peak:%sB/s (%sb/s) ( %s) average:%sB/s (%sb/s)\n", HumanSize($max), HumanSize($max * 8), scalar localtime($starttime + ($maxindex * $interval)), HumanSize($roc), HumanSize($roc * 8);
    }
# functions
sub roc {
    my $dev = shift;
    my $blksize = shift;
    my $interval = shift;
    my ($max, $maxindex, $i, $first, $last, $total);
    my $prev = -1;
    my $first = $_[0];
    if ($ENV{'OUTPUT_CSV'}) { print STDERR "$dev," }
    foreach (__) {
        if ($prev != -1) {
            if ($_ < $prev) {
                dbg "wrap detected at $i ($_ < $prev)\n";
                $prev = 0;
            }
            my $this = ($_ - $prev) * $blksize / $interval;
            if ($this > $max) {
                $max = $this;
                $maxindex = $i;
            }
            if ($ENV{'OUTPUT_CSV'}) { print STDERR "$this," }
        }
        $prev = $_; # store current val for next time around
        $last = $_;
        $i++;
    }
    if ($ENV{'OUTPUT_CSV'}) { print STDERR "\n" }
    return ($max, $maxindex, ($last - $first) * $blksize / ($interval * ($i - 1)));
}
sub totals { # params: stat_hash
    my $stat = shift;
    my totalvals;
    foreach (keys %$stat) {
        next if (!defined($stat{$_}));
        my vals = {$stat{$_}};
        my $i;
        foreach (vals) {
            $totalvals[$i++] += $_;
        }
    }
    return totalvals;
}

```

```

# converts to KB, MB, etc. and outputs size in readable form
sub HumanSize { # params: bytes/bits
    my $bytes = shift;
    my suffixes = ( '', 'K', 'M', 'G', 'T', 'P' );
    my $i = 0;
    while ($bytes / 1024.0 >= 1) {
        $bytes /= 1024.0;
        $i++;
    }
    return sprintf("%.1f %s", $bytes, $suffixes[$i]);
}

# convert human-readable time interval to number of seconds
sub TimeHuman { # params: human_time
    my $time = shift;
    my %suffixes = ('s' => 1, 'm' => 60, 'h' => 60 * 60, 'd' => 60 * 60 *
24);

    $time =~ /^([0-9]*)(.*)$/;
    $time = $1;
    my $suffix = (split //, $2)[0]; # first letter from suffix
    if (exists $suffixes{$suffix}) {
        $time *= $suffixes{$suffix};
    }
    return $time;
}

sub epoch { # params: date
    my $date = shift;
    my $seconds = `date +%s' --date "$date" 2>&1`;
    if ($? != 0) {
        die "Failed to recognize time stamp: $date\n";
    }
    return $seconds;
}

```

## 5.5.3.7. WAN Configuration

Using SIOS DataKeeper in a WAN environment requires special configuration due to the nature of WAN networking. The following tips are recommended:

- To prevent false failover, you should enable manual failover confirmation. Because most WANs are somewhat less reliable than LANs and because typical WAN mirror configurations will have only one comm path, this is usually a good idea. With this option enabled, a LifeKeeper failover will proceed only if the user confirms the failover by using the `lk_confirmso` command. Refer to the `lk_confirmso` man page for more details.
- Determine the proper value for `LKDR_ASYNC_LIMIT`, based upon the latency and throughput of the WAN. The `LKDR_ASYNC_LIMIT` parameter (which is set in `/etc/default/LifeKeeper`) determines the number of outstanding asynchronous write operations (per mirror) that SIOS DataKeeper will allow. The default value for this parameter is 4096, but a larger number may increase write performance of the mirror. The disadvantage to increasing this value is that more data will be allowed to be out of sync between the primary and secondary at any given time. See the [Asynchronous Mirroring Information](#) in **Mirroring with SIOS DataKeeper for Linux** for further information on `LKDR_ASYNC_LIMIT`.
- If you are mirroring a large amount of data over a slow WAN link, it may be desirable to avoid the initial full data resynchronization and instead ship or otherwise transport a copy of the source disk or partition to the remote (disaster recovery) site. To avoid the initial resynchronization, follow the steps in [Avoiding Full Resynchronizations](#).

 **IMPORTANT:** This procedure is not necessary if you created your hierarchy using the “New Replicated Filesystem” option in the LifeKeeper GUI. The “New Replicated Filesystem” option has been optimized to avoid the full initial resync.

- If the WAN link experiences periods of downtime in excess of 15 seconds on a regular basis, it may also be wise to tune the LifeKeeper heartbeat parameters. See [Tuning the LifeKeeper Heartbeat](#) for details.

## 5.5.3.8. SIOS DataKeeper for Linux Resource Types

---

When creating your DataKeeper resource hierarchy, LifeKeeper will prompt you to select a resource type. There are several different DataKeeper resource types. The following information can help you determine which type is best for your environment.

### Replicate New File System

Choosing a [New Replicated File System](#) will create/extend the NetRAID device, mount the given mount point on the NetRAID device and put both the LifeKeeper supported file system and the NetRAID device under LifeKeeper protection. The local disk or partition will be formatted.

When this resource is extended to the second node, it will not do a full resync but only the file system metadata data. However when it is extended to the 3rd or more node, a full resync will be done to that node. To avoid a full resync, please follow the directions provided in [Avoiding Full Resynchronizations](#).

**!** **CAUTION: All data will be deleted.**

### Replicate Existing File System

Choosing [Replicate Existing File System](#) will use a currently mounted disk or partition and create a NetRAID device without deleting the data on the disk or partition. SIOS DataKeeper will unmount the local disk or partition, create the NetRAID device using the local disk or partition and mount the mount point on the NetRAID device. It will then put both the NetRAID device and the LifeKeeper supported file system under LifeKeeper protection.

### DataKeeper Resource

Choosing a [DataKeeper Resource](#) will create/extend the NetRAID device and put it under LifeKeeper protection without a file system. You might choose this replication type if using a database that can use a raw I/O device.

In order to allow the user continued data access, SIOS DataKeeper will not attempt to unmount and delete a NetRAID device if it is currently mounted. The user must manually unmount it before attempting a manual switchover and mount it on the other server after the manual switchover.

**Note:** After the DataKeeper resource has been created, should you decide to protect a manually mounted file system with LifeKeeper, you can do so as follows:

1. Format the NetRAID device with a LifeKeeper supported file system.
2. Mount the NetRAID device.
3. Create and extend a file system hierarchy using the NetRAID device as if it were a shared storage

disk or partition.

LifeKeeper's file system recovery kit will now be responsible for mounting/unmounting it during failover.

## 5.5.3.9. I/O Fencing with DataKeeper Configuration

---

In principle, I/O fencing using storage reservations is not available in DataKeeper configuration and split brain can occur. Therefore, you need to take steps to prevent a split brain from occurring via the following controls.

### Exclusive Control using IP Resources

IP resources have an exclusive control functionality using duplication checking to ensure the same IP resource is not activated on multiple nodes. This can be used to avoid a split brain with DataKeeper resources.

Adding an IP resources as a child resource to all DataKeeper resources in the hierarchy can prevent the DataKeeper resource from starting on multiple nodes at the same time.

This method can only be used in environments where all the nodes in the cluster reside in the same subnet. This is required to perform the duplicate IP address checking.

### Exclusive Control with Quorum/Witness Functionality

You can use the quorum/witness functionality in LifeKeeper to prevent multiple nodes from becoming active at the same time.

For details, please refer to the [Quorum/Witness](#) topic in the Technical Documentation.

## 5.5.3.10. Resource Configuration Tasks

---

You can perform all SIOS DataKeeper configuration tasks via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer and monitor SIOS DataKeeper resources.

### Overview

The following tasks are available for configuring SIOS DataKeeper:

- **Create a Resource Hierarchy** – Creates a DataKeeper resource hierarchy.
- **Delete a Resource Hierarchy** – Deletes a DataKeeper resource hierarchy.
- **Extend a Resource Hierarchy** – Extends a DataKeeper resource hierarchy from the primary server to a backup server.
- **Unextend a Resource Hierarchy** – Unextends (removes) a DataKeeper resource hierarchy from a single server in the LifeKeeper cluster.
- **Create Dependency** – Creates a child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- **Delete Dependency** – Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- **In Service** – Activates a resource hierarchy.
- **Out of Service** – Deactivates a resource hierarchy.
- **View/Edit Properties** – View or edit the properties of a resource hierarchy.

## 5.5.3.10.1. Creating a DataKeeper Resource Hierarchy

Perform the following on your primary server:

1. Select **Edit > Server > Create Resource Hierarchy**

The **Create Resource Wizard** dialog will appear.

2. Select the **Data Replication** option from the drop down list and click **Next** to continue.
3. You will be prompted for the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter any error requiring you to correct the previously entered information. You may click **Cancel** at any time to cancel the entire creation process.

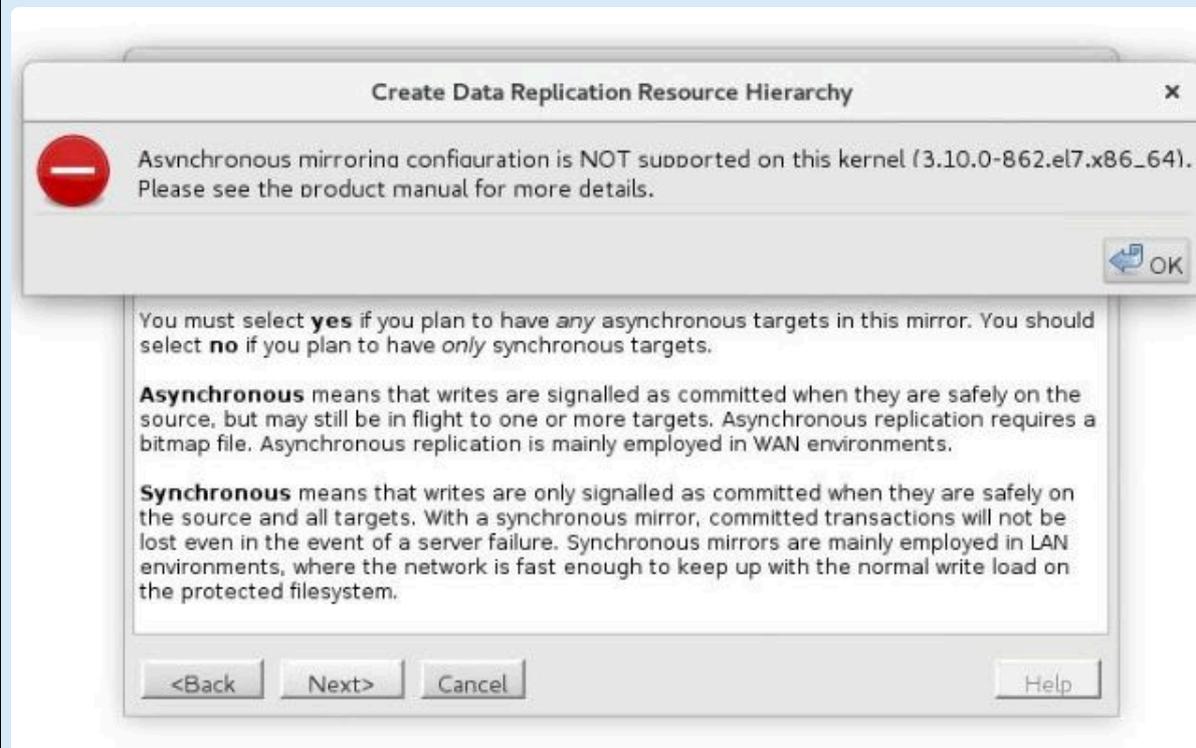
Field	Tips
Switchback Type	<p>You must select <b>intelligent switchback</b>. This means that after a failover to the backup server, an administrator must manually switch the DataKeeper resource back to the primary server.</p> <p><b>CAUTION:</b> This release of SIOS DataKeeper does not support Automatic Switchback for DataKeeper resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource sitting on top of a DataKeeper resource.</p>
Server	Select the name of the server where the NetRAID device will be created (typically this is your primary server). All servers in your cluster are included in the drop down list box.
Hierarchy Type	<p>Choose the DataKeeper Resource</p> <ul style="list-style-type: none"> <li>• <a href="#">Replicate New File System</a></li> <li>• <a href="#">Replicate Existing File System</a></li> <li>• <a href="#">DataKeeper Resource</a></li> </ul>
Bitmap File	<p>Select or edit the name of the bitmap file used for intent logging. If you choose <b>None</b>, then an intent log will not be used and every resynchronization will be a full resync instead of a partial resync.</p> <p><b>Important:</b> The bitmap file should not reside on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem). Placing data replication bitmap files on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem) will result in an “invalid argument” error when LifeKeeper tries to configure the mirror. The default location for the bitmap file is under /opt/LifeKeeper. This default location should be changed if /opt/LifeKeeper resides on a btrfs filesystem.</p> <p><b>Note:</b> btrfs is currently not supported by LifeKeeper for Linux.</p>

Enable Asynchronous Replication?

Select **Yes** to allow this replication resource to support asynchronous replication to target systems. Select **No** if you will use synchronous replication to all targets. You will be asked later to choose the actual type of replication, asynchronous or synchronous, when the replication resource is extended to each target server. (See [Mirroring with SIOS DataKeeper](#) for a discussion of both replication types.) If you want the replication to any of these targets to be performed asynchronously, you should choose **Yes** here, even if the replication to other targets will be done synchronously.

**Note:** If you select asynchronous mirroring in an environment where asynchronous mirroring is not supported, the following message is displayed.

**Asynchronous mirroring configuration is NOT supported on this kernel (3.10.0-862.el7.x86\_64).**



The next sequence of dialog boxes depends on which **Hierarchy Type** you have chosen. While some of the dialog boxes may be the same for each Hierarchy Type, their sequence and the required information may be slightly different. The next three topics take you through the remainder of the Hierarchy creation process.

- [DataKeeper Resource](#)
- [New Replicated File System](#)
- [Replicate Existing File System](#)

## 5.5.3.10.1.1. Replicate New File System

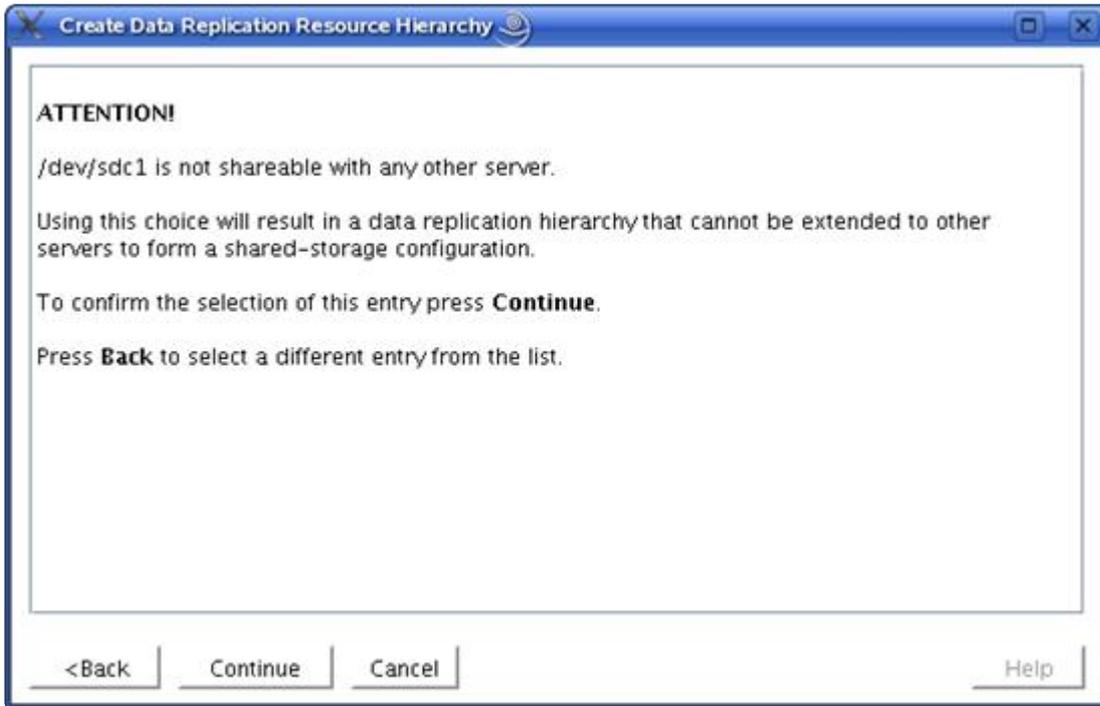
This option will create a NetRAID device, format it with a LifeKeeper supported file system type, mount the file system on the NetRAID device and place both the mounted file system and the NetRAID device under LifeKeeper protection. The NetRAID device and the local disk or partition will be formatted causing existing data to be deleted. You should select this option if you want to create a mirror on a new file system and place it under LifeKeeper protection. You will need one free disk or partition for this resource type.

**! CAUTION:** This option will cause your local disk or partition to be formatted and all existing data will be deleted.

1. Enter the following information when prompted:

Field	Tip
<a href="#">Source Disk or Partition</a>	<p>The list of Source Disks or Partitions in the drop-down list contains all the available disks or partitions that are <u>not</u>:</p> <ul style="list-style-type: none"> <li>• currently mounted</li> <li>• swap disks or partitions</li> <li>• LifeKeeper-protected disks or partitions</li> </ul> <p>The drop-down list will also filter out special disks or partitions, for example, root (/), boot (/boot),_ /proc_, floppy and cdrom.</p> <p><b>Note:</b> The source disk or partition must be uniquely identifiable. Starting in v9.5.0, LifeKeeper will no longer allow the device name to be used to identify a device.</p>

2. The following screen will display if you select a source disk or partition that is not shared.



3. Select **Back** to select different source disk or partition that is shared. Provide the remaining information to finish configuring the resource.

Field	Tips
New Mount Point	Enter the <b>New Mount Point</b> of the new file system. This should be the mount point where the replicated disk or partition will be located.
New File System Type	Select the <b>File System Type</b> . You may only choose from the LifeKeeper supported file system types.
DataKeeper Resource Tag	Select or enter a unique DataKeeper Resource Tag name for the DataKeeper resource instance.
File System Resource Tag	Select or enter the File System Resource Tag name for the file system resource instance.
Bitmap File	<p>Select the bitmap file entry from the pull down list.</p> <p>Displayed in the list are all available shared file systems that can be used to hold the bitmap file. The bitmap file must be placed on a shared device that can switch between the local nodes in the cluster.</p> <p><b>Important:</b> The bitmap file should not reside on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem). Placing data replication bitmap files on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem) will</p>

	result in an “invalid argument” error when LifeKeeper tries to configure the mirror. The default location for the bitmap file is under <i>/opt/LifeKeeper</i> . This default location should be changed if <i>/opt/LifeKeeper</i> resides on a btrfs filesystem.
--	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

4. Click **Next** to continue to the **Confirmation** Screen.
5. A confirmation screen noting the location where the new file system will be created and a warning indicating the pending reformat of the local disk or partition will display. Click **Create** to begin **Resource Creation**.
6. LifeKeeper will verify that you have provided valid data to create your resource on a new file system. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created. Note that the creation of the file system may take several minutes depending upon the disk or partition size.

Click **Next** to continue.

7. An information box appears announcing the successful creation of your new replicated file system resource hierarchy. You must **Extend** the hierarchy to another server in your cluster to begin data replication and in order to place it under LifeKeeper protection.

Click **Next** to extend the resource or click **Cancel** if you wish to extend your resource at another time.

If you click **Continue** LifeKeeper will launch the **Pre-extend Wizard**.

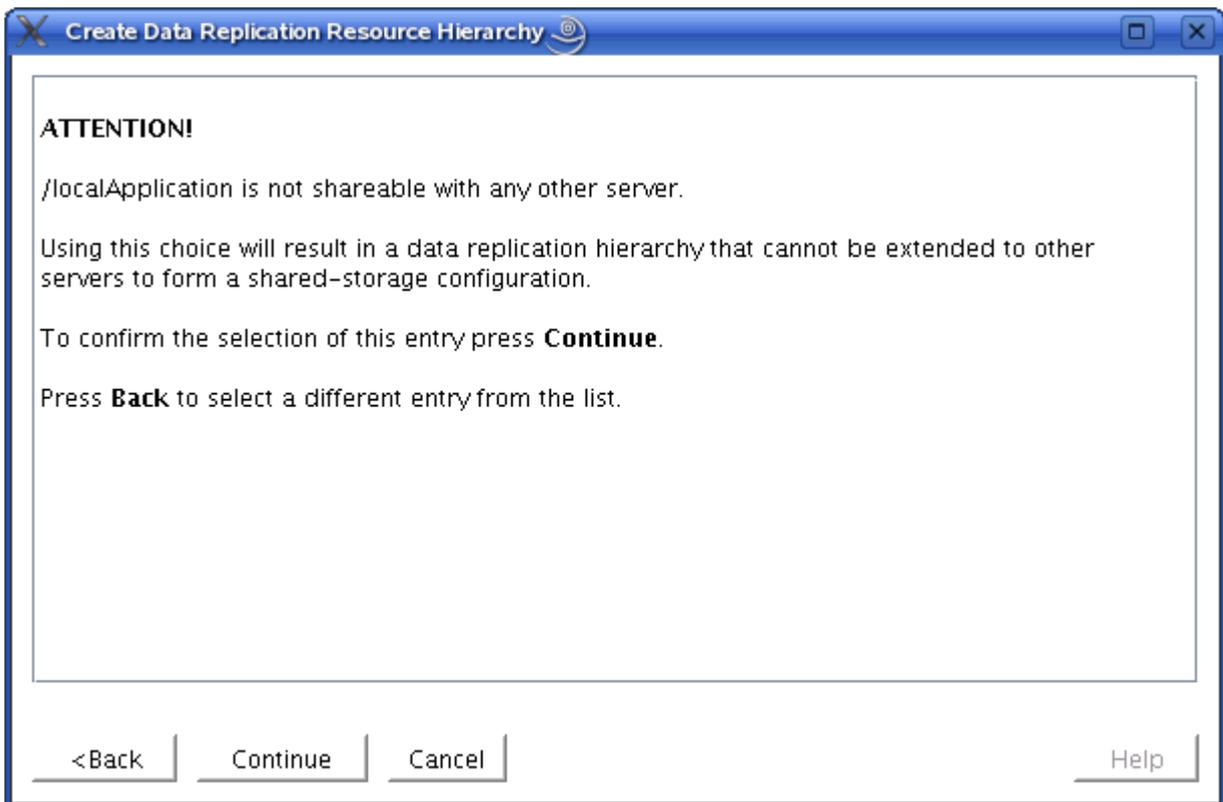
## 5.5.3.10.1.2. Replicate Existing File System

This option will unmount a currently mounted file system on a local disk or partition, create a NetRAID device, then re-mount the file system on the NetRAID device. Both the NetRAID device and the mounted file system are placed under LifeKeeper protection. You should select this option if you want to create a mirror on an existing file system and place it under LifeKeeper protection.

1. Enter the following information when prompted:

Field	Tip
Existing Mount Point	<p>This should be the mount point for the NetRAID device on the primary server. The local disk or partition should already be mounted at this location.</p> <p><b>Note:</b> The source disk or partition must be uniquely identifiable. Starting in v9.5.0, LifeKeeper will no longer allow the device name to be used to identify a device.</p>

2. The following screen will display if you select a mount point that is not shared.



3. Select **Back** to select a shared mount point. Provide the remaining information to finish configuring the resource.

Field	Tips
DataKeeper	Select or enter a unique <b>DataKeeper Resource Tag</b> name for the DataKeeper

Resource Tag	resource instance.
File System Resource Tag	Select or enter the <b>File System Resource Tag name</b> .
Bitmap File	<p>Select the bitmap file entry from the pull down list.</p> <p>Displayed in the list are all available shared file systems that can be used to hold the bitmap file. The bitmap file must be placed on a shared device that can switch between the local nodes in the cluster.</p> <p><b>Important:</b> The bitmap file should not reside on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem). Placing data replication bitmap files on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem) will result in an “invalid argument” error when LifeKeeper tries to configure the mirror. The default location for the bitmap file is under /opt/LifeKeeper. This default location should be changed if /opt/LifeKeeper resides on a btrfs file system.</p> <p><b>Important:</b> Do not select the shared disk area used for replication if displayed in the pull-down selection as the storage area for bitmaps. The shared file system allocated for replication cannot be used as the storage destination of bitmap files. You must use a shared file system location that has been allocated for just bitmap file.</p>

- Click **Next** to create your DataKeeper resource on the primary server.
- LifeKeeper will verify that you have provided valid data to create your DataKeeper resource. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created.

Click **Next**.

- An information box appears announcing that you have successfully created an existing replicated file system resource hierarchy. You must **Extend** the hierarchy to another server in your cluster to begin replication and to place it under LifeKeeper protection.

Click **Next** to extend the resource, or click **Cancel** if you wish to extend your resource at another time.

If you click **Continue** LifeKeeper will launch the *Pre-extend Wizard*. Refer to Step 2 under Extending Your Hierarchy for details on how to extend your resource hierarchy to another server.

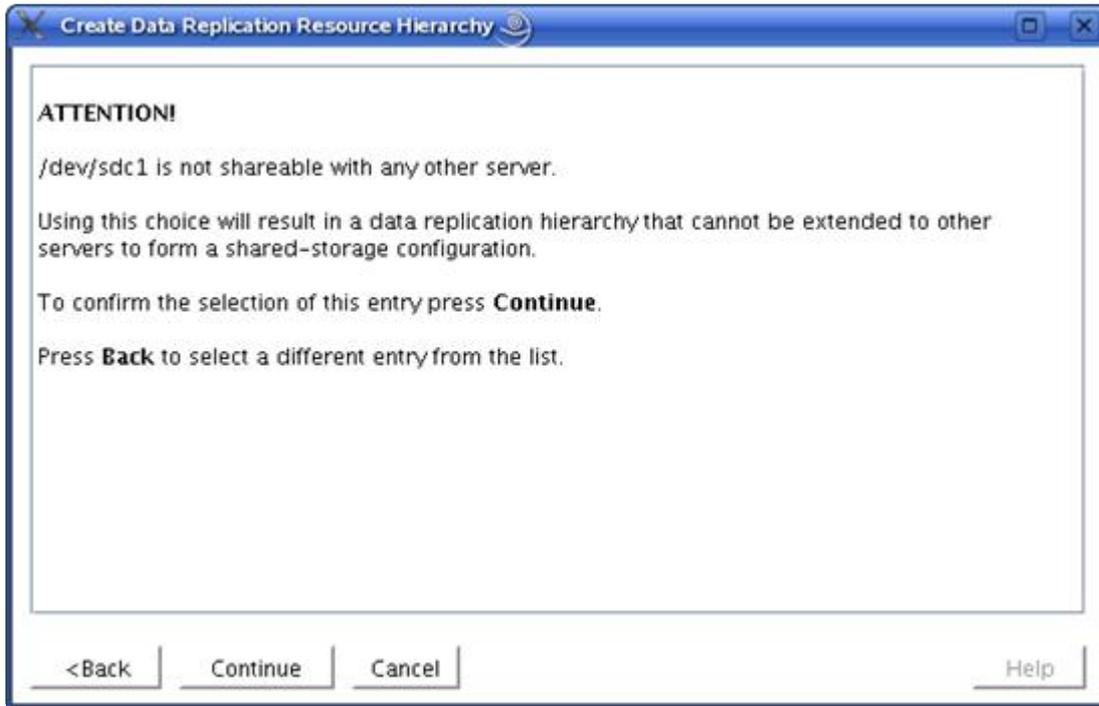
## 5.5.3.10.1.3. DataKeeper Resource

This option will create only the NetRAID device (not a file system) and place the device under LifeKeeper protection. You should select this option if you only want to create a DataKeeper device on a disk or partition and place the device under LifeKeeper protection. You will need to manually make and mount a file system on this device in order to create a readable mirror. You will need one free disk or partition for this resource type.

1. Enter the following information when prompted:

Field	Tip
Source Disk or Partition	<p>The list of Source Disks or Partitions in the drop-down list contains all the available disks or partitions that are not:</p> <ul style="list-style-type: none"> <li>° currently mounted</li> <li>° swap disks or partitions</li> <li>° LifeKeeper-protected disks or partitions</li> </ul> <p>The drop-down list will also filter out special disks or partitions, for example, root (/), boot (/boot), /proc, floppy and cdrom.</p> <p><b>Note:</b> The source disk or partition must be uniquely identifiable. Starting in v9.5.0, LifeKeeper will no longer allow the device name to be used to identify a device.</p>

2. The following screen will display if you select a source disk or partition that is not shared.



3. Select **Back** to select a different source disk or partition that is shared. Provide the remaining information to finish configuring the resource.

Field	Tips
DataKeeper Resource Tag	Select or enter a unique <b>DataKeeper Resource Tag name</b> for the DataKeeper resource instance.
Bitmap File	<p>Select the bitmap file entry from the pull down list.</p> <p>Displayed in the list are all available shared file systems that can be used to hold the bitmap file. The bitmap file must be placed on a shared device that can switch between the local nodes in the cluster.</p> <p><b>Important:</b> The bitmap file should not reside on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem). Placing data replication bitmap files on a btrfs filesystem (or any other LifeKeeper for Linux unsupported filesystem) will result in an “invalid argument” error when LifeKeeper tries to configure the mirror. The default location for the bitmap file is under /opt/LifeKeeper. This default location should be changed if /opt/LifeKeeper resides on a btrfs file system.</p>

4. Click **Next**.
5. An information window appears notifying you that you will have to manually make the file system and mount the NetRAID device (*/dev/mdX*) before being able to use it.

Click **Create** to create your DataKeeper device on the local disk or partition.

6. An information box appears and LifeKeeper will verify that you have provided valid data to create your DataKeeper resource. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created.

Click **Next** to continue.

7. An information box appears announcing the successful creation of your DataKeeper resource device. You must **Extend** the hierarchy to another server in your cluster to begin data replication and in order to place it on the backup/target server and under LifeKeeper protection.

Click **Continue** to extend the resource, or click **Cancel** if you wish to extend your resource at another time.

If you click **Continue** LifeKeeper will launch the ***Pre-extend Wizard***.

## 5.5.3.10.2. Extending Your DataKeeper Hierarchy

This operation should be started on the Primary Server to the Secondary Server from the **Edit** menu or initiated automatically upon completing the **Create Resource Hierarchy** option in which case you should refer to Step 2 below.

1. On the **Edit menu**, select **Resource** then **Extend Resource Hierarchy**. The **Pre-Extend Wizard** appears. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper **Extend Resource Hierarchy** defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information.

 **Note:** The first two fields appear only if you initiated the Extend from the **Edit** menu.

Field	Tips
Template Server	<p>Select the <b>Template Server</b> where your DataKeeper resource hierarchy is currently <i>in service</i>. It is important to remember that the <b>Template Server</b> you select now and the <b>Tag to Extend</b> that you select in the next dialog box represent an in-service (activated) resource hierarchy.</p> <p>An error message will appear if you select a resource tag that is not in service on the template server you have selected. The drop down box in this dialog provides the names of all the servers in your cluster.</p>
Tag to Extend	<p>This is the name of the DataKeeper instance you wish to extend from the template server to the target server. The drop down box will list all the resources that you have created on the template server.</p>
Target Server	<p>Enter or select the server you are extending to.</p>
Switchback Type	<p>You must select <b>intelligent switchback</b>. This means that after a failover to the backup server, an administrator must manually switch the DataKeeper resource back to the primary server.</p> <p><b>CAUTION:</b> This release of SIOS DataKeeper does not support Automatic Switchback for DataKeeper resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource sitting on top of a DataKeeper resource.</p>
Template Priority	<p>Select or enter a <b>Template Priority</b>. This is the priority for the DataKeeper hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.</p>

	<b>Note:</b> This selection will appear only for the initial extend of the hierarchy.
Target Priority	Select or enter the <b>Target Priority</b> . This is the priority for the new extended DataKeeper hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server’s priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number “1” to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.

After receiving the message that the pre-extend checks were successful, click **Next**.

Depending upon the hierarchy being extended, LifeKeeper will display a series of information boxes showing the Resource Tags to be extended, some of which cannot be edited.

3. Click **Next** to launch the **Extend Resource Hierarchy** configuration task.
4. The next section lists the steps required to complete the extension of a DataKeeper resource to another server.

## Extending a DataKeeper Resource

1. After you have been notified that your pre-extend script has executed successfully, you will be prompted for the following information:

Field	Tips
Mount Point	Enter the name of the file system mount point on the target server. (This dialog will not appear if there is no LifeKeeper-protected filesystem associated with the DataKeeper Resource.)
Root Tag	Select or enter the <b>Root Tag</b> . This is a unique name for the filesystem resource instance on the target server. (This dialog will not appear if there is no LifeKeeper-protected filesystem associated with the DataKeeper Resource.)
Target Disk or Partition	<p>Select the disk or partition where the replicated file system will be located on the target server.</p> <p>The list of disks or partitions in the drop down box contains all the available disks or partitions that are <u>not</u>:</p> <ul style="list-style-type: none"> <li>• already mounted</li> <li>• swap disks or partitions</li> <li>• LifeKeeper-protected disks or partitions</li> </ul> <p>The drop down list will also filter out special disks or partitions, for example, root (/), boot (/boot), /proc, floppy and cdrom.</p>

	<b>Note:</b> The size of the target disk or partition must be greater than or equal to that of the source disk or partition.
DataKeeper Resource Tag	Select or enter the <b>DataKeeper Resource Tag name</b> .
Bitmap File	Select the name of the bitmap file used for intent logging. If you choose <b>None</b> , then an intent log will not be used and every resynchronization will be a full resync instead of a partial resync.
Replication Path	<p>Select the pair of local and remote IP addresses to use for replication between the target server and the other indicated server in the cluster. The valid paths and their associated IP addresses are derived from the set of LifeKeeper communication paths that have been defined for this same pair of servers. Due to the nature of DataKeeper, it is strongly recommended that you use a private (dedicated) network.</p> <p>If the DataKeeper Resource has previously been extended to one or more target servers, the extension to an additional server will loop through each of the pairings of the new target server with existing servers, prompting for a Replication Path for each pair.</p>
Replication Type	<p>Choose “synchronous” or “asynchronous” to indicate the type of replication that should be used between the indicated pair of servers.</p> <p>As for the previous Replication Path field, if the DataKeeper Resource has previously been extended to one or more target servers, the extension to an additional server will loop through each of the pairings of the new target server with existing servers, prompting for a Replication Type for each pair.</p>

2. Click **Next** to continue. An information box will appear verifying that the extension is being performed.
3. Click **Finish** to confirm the successful extension of your DataKeeper resource instance.
4. Click **Done** to exit the **Extend Resources Hierarchy** menu selection.



**Note:** Be sure to test the functionality of the new instance on all servers by performing a manual switchover. See [Testing Your Resource Hierarchy](#) for details.

During resynchronization, the DataKeeper resource and any resource that depends on it will not be able to fail over. This is to avoid data corruption.

## 5.5.3.10.3. Unextending Your DataKeeper Hierarchy

---

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. On the **Edit** menu, select **Resource** then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the DataKeeper resource. It cannot be the server where the DataKeeper resource is currently in service (active).

**Note:** If you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance, this dialog box will not appear.

Click **Next**.

3. Select the **DataKeeper Hierarchy to Unextend** and click **Next**. (This dialog will not appear if you selected the **Unextend** task by right-clicking on a resource instance in either pane).
4. An information box appears confirming the target server and the DataKeeper resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the DataKeeper resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

**Note:** At this point, data is not being replicated to the backup server.

## 5.5.3.10.4. Deleting a DataKeeper Resource Hierarchy

---

To delete a DataKeeper resource from all servers in your LifeKeeper configuration, complete the following steps.

 **Note:** It is recommended that you take the DataKeeper resource out of service BEFORE deleting it. Otherwise, the **md** and **NetRAID** devices will not be removed, and you will have to unmount the file system manually. See [Taking a DataKeeper Resource Out of Service](#).

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the name of the **TargetServer** where you will be deleting your DataKeeper resource hierarchy.  
**Note:** If you selected the Delete Resource task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.
3. Select the **Hierarchy to Delete**. (This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a resource instance in the left or right pane.) Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete**.
5. Another information box appears confirming that the DataKeeper resource was deleted successfully. Click **Done** to exit.

 **Note:** If the NetRAID device was mounted prior to the resource deletion then it will remain mounted. Otherwise, the NetRAID device will also be deleted.

## 5.5.3.10.5. Taking a DataKeeper Resource Out of Service

---

Taking a DataKeeper resource out of service removes LifeKeeper protection for the resource. It breaks the mirror, unmounts the file system (if applicable), stops the **md** device and kills the **nbd** server and client.

**!** **WARNING:** Do not take your DataKeeper resource out of service unless you wish to stop mirroring your data and remove LifeKeeper protection. Use the **Pause** operation to temporarily stop mirroring.

1. In the right pane of the LifeKeeper GUI, right-click on the **DataKeeper resource** that is in service.
2. Click **Out of Service** from the resource popup menu.
3. A dialog box confirms the selected resource to be taken out of service. Any resource dependencies associated with the action are noted in the dialog. Click **Next**.
4. An information box appears showing the results of the resource being taken out of service. Click **Done**.

## 5.5.3.10.6. Bringing a DataKeeper Resource In Service

---

Bringing a DataKeeper resource in service is similar to creating the resource: LifeKeeper starts the **nbd** server and client, starts the **md** device which synchronizes the data between the source and target devices, and mounts the file system (if applicable).

1. Right-click on the **DataKeeper resource instance** from the right pane.
2. Click **In Service** from the popup menu. A dialog box appears confirming the server and resource that you have selected to bring into service. Click **In Service** to bring the resource into service.
3. An information box shows the results of the resource being brought into service. Any resource dependencies associated with the action are noted in the confirmation dialog. Click **Done**.

## 5.5.3.10.7. Testing Your DataKeeper Resource Hierarchy

---

You can test your DataKeeper resource hierarchy by initiating a manual switchover. This will simulate a failover of the resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource, and InService**. For example, an in-service request executed on a backup server causes the DataKeeper resource hierarchy to be taken out-of-service on the primary server and placed in-service on the backup server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

The state of the DataKeeper resource on the new primary server is set to “Source” in the LifeKeeper GUI. During the switchover, the state of the DataKeeper resource on each target is initially set to “Out of Sync” to show that data is not replicating to that target yet. For a multi-target configuration the previous source will show “Out of Sync” and other targets will show “Out of Sync (Wait for Previous Source)”. Resynchronization will automatically begin (the state will transition to “Resyncing”) on each target starting with the previous source. Once resynchronization is complete, the state will change to “Target”, which is the normal **Standby** condition. These state transitions will often occur quickly, so they may not be seen in the GUI.

 **Note:** Manual failover is prevented for DataKeeper resources during resynchronization.

If you execute the **Out of Service** request, the resource hierarchy is taken out of service without bringing it in service on the other server. The resource can only be brought in service on the same server if it was taken out of service during resynchronization.

## 5.5.4. Administering SIOS DataKeeper for Linux

---

The following topics provide information to help in understanding and managing SIOS DataKeeper for Linux operations and issues after DataKeeper resources are created.

---

[Viewing Mirror Status](#)

[GUI Mirror Administration](#)

[Force Mirror Online](#)

[Pause and Resume](#)

[Set Compression Level](#)

[Command Line Mirror Administration](#)

[Monitoring Mirror Status via Command Line](#)

[Server Failure](#)

[Resynchronization](#)

[Avoiding Full Resynchronizations](#)

## 5.5.4.1. Viewing Mirror Status

---

You can view the **Replication Status** dialog to see the following information about your mirror:

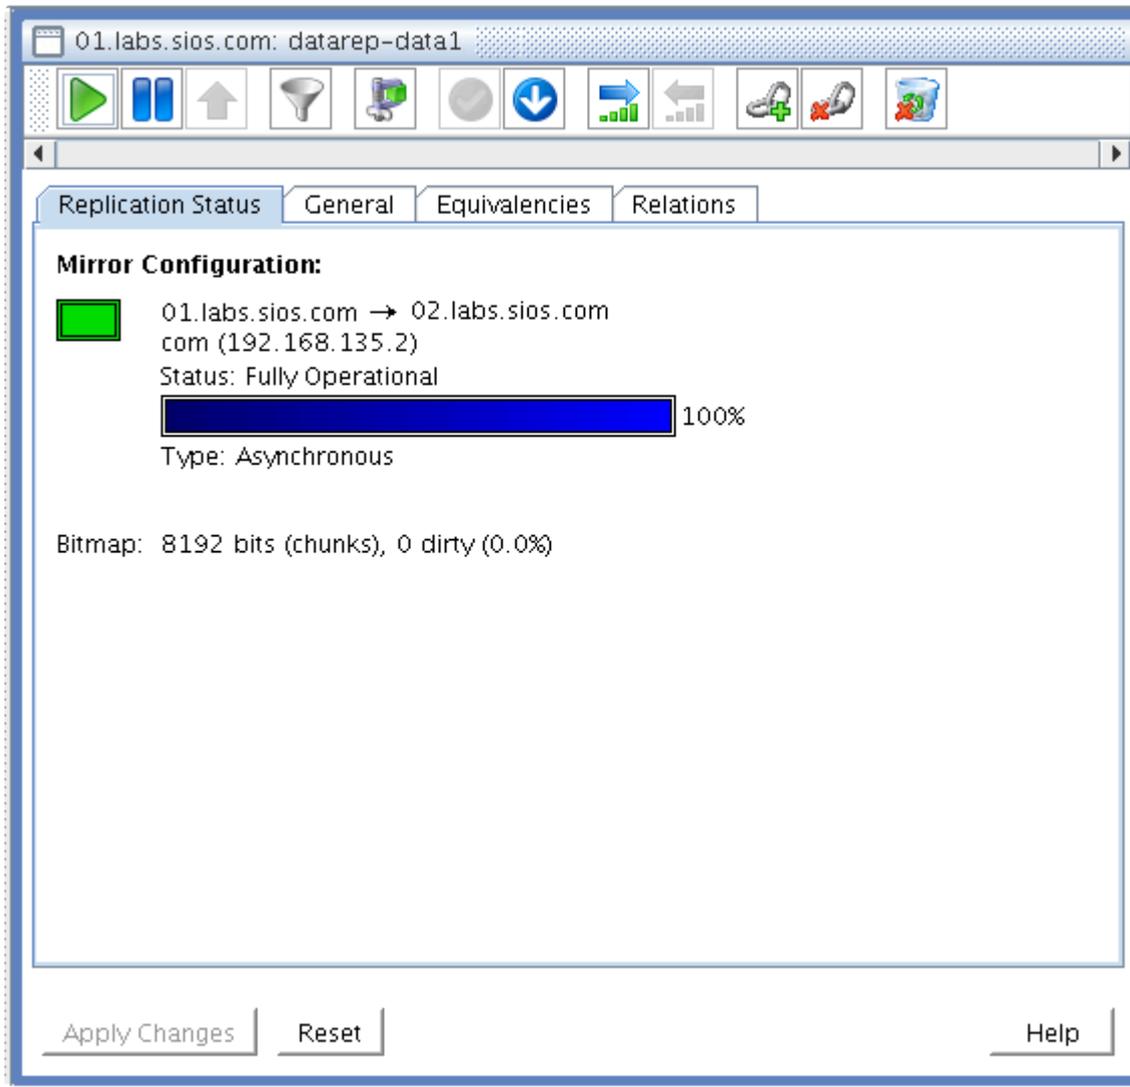
- **Mirror status:** Fully Operational, Paused, Resyncing, or Out Of Sync
- **Synchronization status:** percent complete
- **Replication type:** synchronous or asynchronous
- **Replication direction:** from source server to target server
- **Bitmap:** the state of the bitmap/intent log
- **Network Compression Level:** the compression level (if enabled)

To view the **Replication Status** dialog, do the following:

1. Click the **View** menu, and select **Properties Panel**.
2. Click the **DataKeeper resource** in the **LifeKeeper status** display.

OR

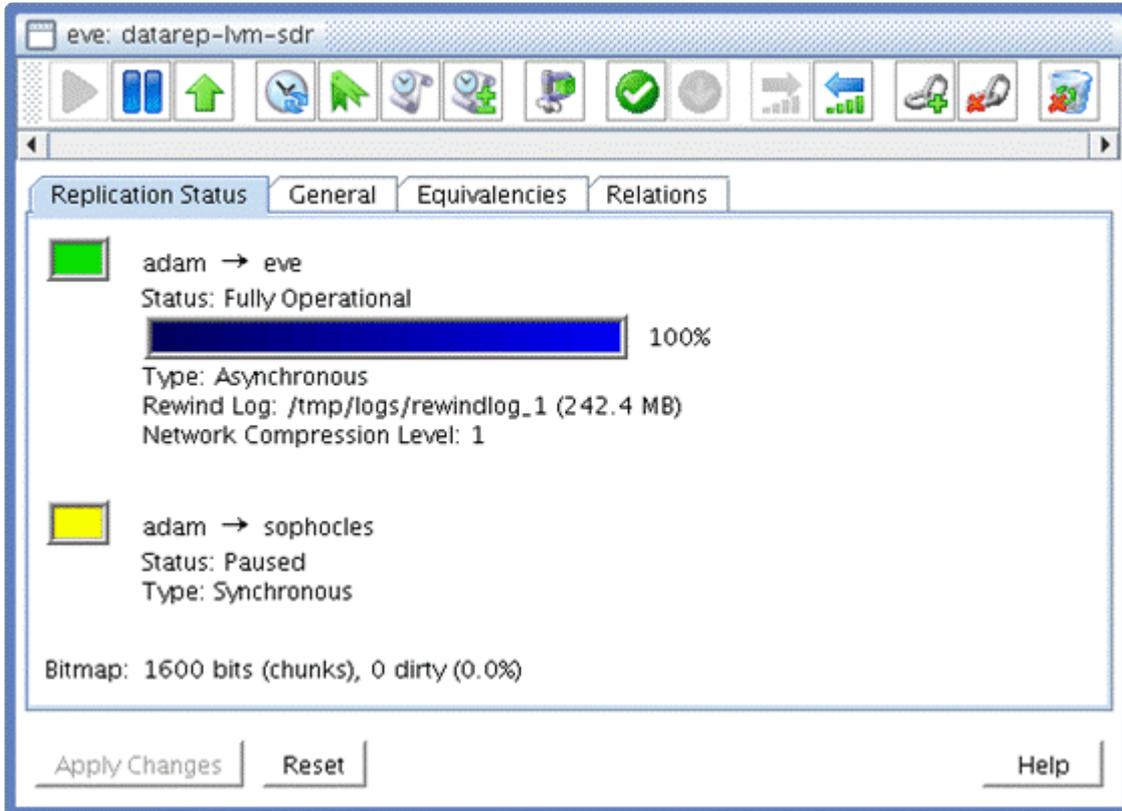
1. Right-click the **DataKeeper resource** in the LifeKeeper status display.
2. From the pop-up menu, select **Properties**.



## 5.5.4.2. GUI Mirror Administration

A SIOS DataKeeper mirror can be administered through the LifeKeeper GUI in two ways:

1. By enabling the **Properties Panel** and clicking the toolbar icons (shown in the screenshot).



Click on each icon below for a description.



OR

2. By right-clicking the **data replication resource** and selecting an action from the popup menu.

# Force Mirror Online

---



**Force Mirror Online** should be used only in the event that both servers have become inoperable and the primary server cannot bring the resource in service after rebooting.

Selecting **Force Mirror Online** removes the *data\_corrupt* flag and brings the DataKeeper resource in service. For more information, see Primary server cannot bring the resource ISP in the [Troubleshooting](#) section.

\* **Note:** `mirror_settings` should be run on the target system(s) (or on all systems, if you want the settings to take effect regardless of which system becomes the mirror source). The mirror must be **paused** and **restarted** before any settings changes will take effect.

## 5.5.4.2.1. Pause and Resume

---

### Pause Mirror



### Resume Mirror



You may pause a mirror to temporarily stop all writes from being replicated to the target disk. For example, you might pause the mirror to take a snapshot of the target disk or to increase I/O performance on the source system during peak traffic times.

When the mirror is paused, it will be mounted for read (or read/write with kernel 2.6.19 or higher) access at the normal filesystem mount point on the target system. Any data written to the target while the mirror is paused will be overwritten when the mirror is resumed.

## 5.5.4.2.2. Set Compression Level

---



The Network Compression Level may be set to a value from 0 to 9. A value of 0 disables compression entirely. Level 1 is the fastest but least aggressive compression level, while Level 9 is the slowest but best. Network compression is typically effective only on WANs.

## 5.5.4.3. Command Line Mirror Administration

---

In addition to performing actions through the LifeKeeper GUI, the mirror can also be administered using the command line. There are several commands (found in the `$LKROOT/bin` directory) that can be used to administer a DataKeeper resource.

### Mirror Actions

```
mirror_action <tag> <action> [source] [target(s)]
```

**<tag>** is the LifeKeeper resource tag of the DataKeeper resource

**<action>** is one of: `pause`, `resume`, `force`, `fullresync`

**[source]** (*optional*) is the current source system (if source is not specified, it will use the current system the command was run from)

**[target]** (*optional*) is the target system (or list of systems) that the action should affect (if target(s) is not specified, it will use all of the applicable target(s))

 **Note:** When using the `force` action, `source` argument is required to specify source node and `target(s)` argument is unnecessary. When using the `pause`, `resume` or `fullresync` action, if specifying `target(s)` argument, `source` argument is also required.

#### Examples:

To pause the mirror named `datarep-ext3`:

```
mirror_action datarep-ext3 pause
```

To resume replication from `adam` to both `eve` and `sophocles`:

```
mirror_action datarep-ext3 resume adam eve sophocles
```

To force the mirror online on system `eve` (force mirror online):

```
mirror_action datarep-ext3 force eve
```

To resume replication and force a full resynchronization from `adam` to `sophocles`:

```
mirror_action datarep-ext3 fullresync adam sophocles
```

## Mirror Resize

The `mirror_resize` command performs a mirror resize without having to recreate the resource. The underlying devices should be resized on both the source and target system before resizing the mirror. Then run the command on the source system. The size of the underlying devices will be auto-detected and used as the new mirror size. Optionally the mirror size can be specified.

With LifeKeeper v9.5.0 or later, the mirror can be resized even when the resource is in service. However, when reducing the mirror size, the resource must be out of service. Please note that some file systems do not support reducing the mirror size.

 **Note:** When the command is interrupted due to an error etc., add the “-f” option and execute again. If not completed successfully, data may become inconsistent.

```
mirror_resize [-f] [-s <size>] <tag>
```

<tag> is the tag of a mirror resource

-f forces the resize without user prompts (not recommended)

-s <size> specifies alternate mirror size (in KB) This parameter is required.

### Requirements for Mirror Resize

- Underlying devices should be a logical volume (LV)
- Only a configuration with a single target is supported

 **NOTE:** `mirror_resize` is NOT supported in multi-target configurations.

### Recommended Steps for Mirror Resize

1. Perform the resize of the underlying device. Perform this on both the source and the target. (Please note that the target size must be greater than or equal to the source size.)
2. Run `mirror_resize` on the source system. This will update the internal metadata and bitmap for the mirror to reflect the newly expanded disk size.

**Example:** `mirror_resize -s <size in KB> <tag>`

3. When the resource is out of service, bring only the mirror (i.e., *datarep*) resource in service. A resync of the newly expanded device will occur.
4. Perform the file system resize on the mirror device (e.g. `resize2fs /dev/mdX` where X is md device number for the mirror being resized such as `/dev/md0`).

**Note:** An `fsck` may be required before being able to resize the file system.

**Note:** Some file systems may be required to be mounted before being resized. Bring the resource in service if it is not mounted.

5. Bring the file system and application resources in service if they are out of service.

## Recommended Steps for Mirror Resize with an XFS file system:

1. Perform the resize of the underlying device. Perform this on both the source and the target. (Please note that the target size must be greater than or equal to the source size.)
2. Run `mirror_resize` on the source system. This will update the internal metadata and bitmap of the mirror to reflect the newly expanded device size.

**Example:** `mirror_resize -s <size in KB> <tag>`

3. When the resource is out of service, bring the mirror (i.e., `datarep`) and file system resource in service. A resync of the newly expanded device or partition will occur.
4. Perform the file system resize on the file system (e.g. `xfs_growfs -D size /path/to/file/system`).
5. When the resource is out of service, bring the application resource in service.

## Bitmap Administration

```
bitmap -a <num>|-c|-d|-x <size_kb>|-X <bitmap_file>
```

`-a <num>` adds the asynchronous write parameter to the bitmap file. It is needed if a synchronous mirror is upgraded to include an asynchronous target. The default value for `<num>` is 256. To calculate the correct value for this limit, see the [Asynchronous Mirroring Information](#) in **Mirroring with SIOS DataKeeper for Linux**.

`-c` cleans the bitmap file (*zeroes all the bits*). This can be used to avoid a full resync in case an exact replica of the source disk exists on the target. **Use this option with extreme caution.**

`-d` dirties the bitmap file (*sets all the bits to ones*). This option can be used to force a full resync, for example after a split-brain situation has occurred.

`-m` reads the bitmap and produces merge stream.

`-X <bitmap file>` examines the bitmap file and displays useful information about the bitmap and the mirror.

`-x <size_kb>` extends bitmap file to be valid with disk of `size_kb`.

**(Note:** This option is only used internally for mirror resizing.)

In addition, the `mdadm` command may also be used to administer a DataKeeper resource, as the DataKeeper resource is actually an md device. Refer to the `mdadm(8)` man page for details. **Note:** When using `mdadm`, be sure to use `$LKROOT/bin/mdadm`, not the one included with the operating system.

## 5.5.4.4. Monitoring Mirror Status via Command Line

---

Normally, the mirror status can be checked using the **Replication Status** tab in the **Resource Properties** dialog of the LifeKeeper GUI. However, you may also monitor the status of your mirror by executing:

```
$LKROOT/bin/mirror_status <tag>
```

### Example:

```
# mirror_status datarep-ext3-sdr
```

```
[-]     eve -> adam
```

```
      Status: Paused
```

```
      Type: Asynchronous
```

```
[-]     eve -> sophocles
```

```
      Status: Resynchronizing
```

```
      [=>           ] 11%
```

```
      Resync Speed: 1573K/sec
```

```
      Type: Synchronous
```

```
Bitmap: 4895 bits (chunks), 4895 dirty (100.0%)
```

**The following command may also be helpful:**

```
cat /proc/mdstat
```

**A sample *mdstat* file is shown below:**

```
eve:~ # cat /proc/mdstat
```

```
Personalities : [raid1]
```

```
md1 : active raid1 nbd10[1] nbd8[3](F) sdb1[0]
```

```
      313236 blocks super non-persistent [3/2] [UU_]
```

```
      bitmap: 3/3 pages [12KB], 64KB chunk, file: /opt/LifeKeeper/
```

bitmap\_ext3-sdr

unused devices: <none/></tag>

## 5.5.4.5. Server Failure

---

If both your primary and backup servers become inoperable, your DataKeeper resource will be brought into service/activated only when **both** servers are functional again. This is to avoid data corruption that could result from initiating the resynchronization in the wrong direction. If you are certain that the only operable server was the last server on which the resource was “**In Service Protected**” (**ISP**), then you can force it online by right-clicking the DataKeeper resource and then selecting **Force Mirror Online**.

## 5.5.4.6. Resynchronization

During the resynchronization of a DataKeeper resource, the state of this resource instance on the target server is “**Resyncing**”. However, the resource instance is “**Source**” (**ISP**) on the primary server. The LifeKeeper GUI reflects this status by representing the DataKeeper resource on the target server with the following icon:



and the DataKeeper resource on the primary server with this icon:



As soon as the resynchronization is complete, the resource state on the target becomes “**Target**” and the icon changes to the following:



✿ A full resync synchronizes an entire disk partition.

The following points should be noted about the resynchronization process:

- A SIOS DataKeeper resource and its parent resources cannot fail over to a target that was in the synchronization process when the primary failed.
- If your DataKeeper resource is taken out of service/deactivated during the synchronization of a target server, that resource can only be brought back into service/activated on the same system or on another target that is already in sync (if multiple targets exist), and the resynchronization will continue.
- If your primary server becomes inoperable during the synchronization process, any target server that is in the synchronization process will not be able to bring your DataKeeper resource into service. Once your primary server becomes functional again, a resynchronization of the mirror will continue.

## 5.5.4.7. Avoiding Full Resynchronizations

When replicating large amounts of data over a WAN link, it is desirable to avoid full resynchronizations which can consume large amounts of network bandwidth and time. With newer kernels, SIOS DataKeeper can avoid almost all full resyncs by using its bitmap technology. However, the initial full resync, which occurs when the mirror is first set up, cannot be avoided when existing data is being replicated. (For brand new data, SIOS DataKeeper does not perform a full resync, so the steps below are not necessary.)

✿ A full resync synchronizes an entire disk partition.

There are a couple of ways to avoid an initial full resync when replicating existing data. Two recommended methods are described below.

### Method 1 – Replicating to a 2nd node

The first method consists of taking a raw disk image and shipping it to the target site. This results in minimal downtime as the mirror can be active on the source system while the data is in transit to the target system.

#### Procedure

1. Create the mirror (selecting Replicate Existing Filesystem), but do not extend the mirror to the target system.
2. Take the mirror out of service.
3. Take an image of the source disk or partition. For this example, the chosen disk or partition is */dev/sda1*:

```
root@source# dd if=/dev/sda1 of=/tmp/sdr_disk.img bs=65536
```

(The block size argument of 65536 is merely for efficiency).

This will create a file containing the raw disk image of the disk or partition.

Note that instead of a file, a hard drive or other storage device could have been used.

4. Optional Step – Take a checksum of the source disk or partition:

```
root@source# md5sum /dev/sda1
```

5. Optional Step – Compress the disk image file:

```
root@source# gzip /tmp/sdr_disk.img
```

6. Clear the bitmap file (Replace the last argument with the path of the bitmap file which you specified when creating resources.):

```
root@source# /opt/LifeKeeper/bin/bitmap -c /opt/LifeKeeper/
bitmap__dr
```

7. Bring the mirror and dependent filesystem and applications (if any), into service. The bitmap file will track any changes made while the data is transferred to the target system.
8. Transfer the disk image to the target system using your preferred transfer method.
9. Optional Step – Uncompress the disk image file on the target system:

```
root@target# gunzip /tmp/sdr_disk.img.gz
```

10. Optional Step – Verify that the checksum of the image file matches the original checksum taken in Step 4:

```
root@target# md5sum /tmp/sdr_disk.img
```

11. Transfer the image to the target disk, for example, */dev/sda2*:

```
root@target# dd if=/tmp/sdr_disk.img of=/dev/sda2 bs=65536
```

12. Set `LKDR_NO_FULL_SYNC=1` in `/etc/default/LifeKeeper` on both systems:

```
root@source# echo 'LKDR_NO_FULL_SYNC=1' >>/etc/default/LifeKeeper
```

```
root@target# echo 'LKDR_NO_FULL_SYNC=1' >>/etc/default/LifeKeeper
```

13. Extend the mirror to the target. A partial resync will occur.
14. Edit `/etc/default/LifeKeeper` to remove the `LKDR_NO_FULL_SYNC` entry.

## Extending to a 3rd node or any additional nodes without doing a full resync

### Procedure for copying data from the Source:

These steps assume the mirror has already been created and extended to the 2nd node, aka target1.

1. Set `LKDR_NO_FULL_SYNC=1` in `/etc/default/LifeKeeper` on each system:

```
root@source# echo 'LKDR_NO_FULL_SYNC=1' >>/etc/default/LifeKeeper
```

```
root@target1# echo 'LKDR_NO_FULL_SYNC=1' >>/etc/default/LifeKeeper
```

```
root@target2# echo 'LKDR_NO_FULL_SYNC=1' >>/etc/default/LifeKeeper
```

2. Extend the mirror to the new target (target2). A partial resync will occur.
3. Pause the mirror to the new target (target2).
4. Take the mirror out of service.
5. Unmount the file system on the paused mirror on target2 by running 'umount <filesystem>' on target2.
6. Stop the md device running on target2 by running 'mdadm -stop /dev/md#' on target2, where # is the value reported in */proc/mdstat*.
7. Make a copy of the source disk or partition on the source node. This could be done using dd or tools from the disk vendor or cloud vendor. The copy **must be** a block-for-block identical copy. It cannot be a file level copy.
8. Optional step – Collect checksum data to verify disk image (md5sum, sha256sum, etc).
9. Optional step – Compress disk image.
10. Bring the mirror and dependent filesystem and applications (if any), into service. The bitmap file will track any changes made while the data is transferred to the target system.
11. Verify that the mirror to target2 is still paused. If it is not then restart at step 4.
12. Verify that the file system and md device are not running on target2. If they are then unmount the file system and stop the md device.
13. Transfer the disk image to the target disk on target2.
14. Verify that the disk image is correct. Perhaps use md5sum or sha256sum to validate the disk contents.
15. Resume the paused mirror to target2. The bitmap on the source was keeping track of any changes made since target2 was paused. When the mirror is resumed these changes will be sent to target2.
16. Edit */etc/default/LifeKeeper* to remove the LKDR\_NO\_FULL\_SYNC entry.

#### Procedure for copying data from paused target:

This will allow no downtime but while the targets are paused, there is no data redundancy.

These steps assume the mirror has already been created and extended to the 2nd node, aka target1.

1. Set LKDR\_NO\_FULL\_SYNC=1 in */etc/default/LifeKeeper* on each system:

```
root@source# echo 'LKDR_NO_FULL_SYNC=1' >>/etc/default/LifeKeeper
```

```
root@target1# echo `LKDR_NO_FULL_SYNC=1` >>/etc/default/LifeKeeper
```

```
root@target2# echo `LKDR_NO_FULL_SYNC=1` >>/etc/default/LifeKeeper
```

2. Extend the mirror to the new target (target2). A partial resync will occur.
3. Pause the mirror to the new target (target2).
4. Pause the mirror to target1.
5. Unmount the file system on the paused mirror on target2 by running 'umount <filesystem>' on target2.
6. Stop the md running on target2 by running 'mdadm -stop /dev/md#' on target2, where # is the value reported in */proc/mdstat*.
7. Unmount the file system on the paused mirror on target1 by running 'umount <filesystem>' on target1.
8. Stop the md running on target1 by running 'mdadm -stop /dev/md#' on target1, where # is the value reported in */proc/mdstat*.
9. Make a copy of the target disk or partition on target1. This could be done using dd or tools from the disk vendor or cloud vendor. The copy **must be** a block-for-block identical copy of the full disk or partition. It cannot be a file level copy.
10. Optional step – Collect checksum data to verify disk image (md5sum, sha256sum, etc).
11. Optional step – Compress disk image.
12. Resume replication to target1. The bitmap file will track any changes made while the data is transferred to target2.
13. Verify that the mirror to target2 is paused. If it is not then restart at step 4.
14. Verify that the file system and md device are not running on target2. If they are, then unmount the file system and stop the md device.
15. Transfer the disk image to the target disk on target2.
16. Optional step – Decompress disk image.
17. Optional step – Verify the disk image is correct (md5sum, sha256sum, etc.).
18. Resume the paused mirror to target2. The bitmap on the source was keeping track of any changes made since target2 was paused. When the mirror is resumed these changes will be sent to target2.

19. Edit `/etc/default/LifeKeeper` to remove the `LKDR_NO_FULL_SYNC` entry.

## Method 2

This method can be used if the target system can be easily transported to or will already be at the source site when the systems are configured. This method consists of temporarily modifying network routes to make the eventual WAN mirror into a LAN mirror so that the initial full resync can be performed over a faster local network. In the following example, assume the source site is on subnet 10.10.10.0/24 and the target site is on subnet 10.10.20.0/24. By temporarily setting up static routes on the source and target systems, the “WAN” traffic can be made to go directly from one server to another over a local ethernet connection or loopback cable.

### Procedure

1. Install and configure the systems at the source site.
2. Add static routes:

```
root@source# ip route add 10.10.20.0/24 dev eth0
```

```
root@target# ip route add 10.10.10.0/24 dev eth0
```

The systems should now be able to talk to each other over the LAN.

3. Configure the communication paths in LifeKeeper.
4. Create the mirror and extend to the target. A full resync will occur.
5. Pause the mirror. Changes will be tracked in the bitmap file until the mirror is resumed.
6. Delete the static routes:

```
root@source# ip route del 10.10.20.0/24
```

```
root@target# ip route del 10.10.10.0/24
```

7. Shut down the target system and ship it to its permanent location.
8. Boot the target system and ensure network connectivity with the source.
9. Resume Replication. A partial resync will occur.

## 5.5.4.8. Verify Data Before Resync (Wait to Resync)

---

To avoid replicating corrupt or inconsistent data to targets, LifeKeeper can wait for resources to be in-service before replicating data. Starting with 9.5.2, LifeKeeper will by default wait for parent resources of application type 'filesys' to be in-service before replicating data. This provides assurance that the file system is consistent before replicating data. If the file system fails to mount, the appropriate recovery action could be to repair the file system on the current system or check the data on another system before the data is replicated.

The "Wait to Resync" feature is configured in `/etc/default/LifeKeeper` with the setting "LKDR\_WAIT\_TO\_RESYNC". There are 3 options to configure the "Wait to Resync" feature:

1. **False**. This will disable the feature. Parent resources will not be checked before initializing resynchronization. Resynchronization will begin during the initial restore of the DataKeeper resource if there are no other issues blocking resynchronization.
2. **<resource type>**. Specify a specific resource type. By default the 'filesys' resource type is specified. This setting can be any installed resource type. The command `/opt/LifeKeeper/bin/typ_list -f` will display a list of installed resource types. The output of this command is a list of the form "application:type". For example,

```
# typ_list -f:
gen:app
gen:filesys
gen:qsp
gen:hanfs
gen:nfs
scsi:DEVNAME
scsi:device
```

The first field is the application and the second field following ":" is the type (see [Short Status Display](#) for more information). For example, the first entry in the list above is application "gen" and type "app". Any type on the right side of the ":" can be specified. LifeKeeper will wait for parents of the DataKeeper resource which have the specified type to be in-service before resynchronization will begin. For example, suppose that the resource type "LKDR\_WAIT\_TO\_RESYNC=app" is defined in `/etc/default/LifeKeeper` and that the resource hierarchy is as follows:

```
# lcdstatus -q
LOCAL      TAG          ID          STATE  PRIO PRIMARY
ip-10-0-2-128 app1        app1        OSU    1 ip-10-0-2-128
ip-10-0-2-128 SPSL1      SPSL1      OSU    1 ip-10-0-2-128
ip-10-0-2-128 /maxdb1    /maxdb1    OSU    1 ip-10-0-2-128
ip-10-0-2-128 datarep-maxdb1 536681bc-9259-4254 OSU    1 ip-10-0-2-128
ip-10-0-2-128 SPSL2      SPSL2      OSU    1 ip-10-0-2-128
ip-10-0-2-128 /maxdb2    /maxdb2    OSU    1 ip-10-0-2-128
ip-10-0-2-128 datarep-maxdb2 952e4ebd-507d-4b43 OSU    1 ip-10-0-2-128
ip-10-0-2-128 app3        app3        OSU    1 ip-10-0-2-128
ip-10-0-2-128 app2        app2        OSU    1 ip-10-0-2-128
ip-10-0-2-128 /maxdb1    /maxdb1    OSU    1 ip-10-0-2-128
ip-10-0-2-128 datarep-maxdb1 536681bc-9259-4254 OSU    1 ip-10-0-2-128
```

where “app1”, “app2”, and “app3” have resource type “app”. In this example, netraid resource “datarep-maxdb1” will not synchronize until all three gen:app parent resources are in-service. The netraid resource “datarep-maxdb2” will not synchronize until “app1” is in-service since “app1” is its only gen:app parent.

3. **Hierarchy.** All parent resources of a DataKeeper resource must be in-service before synchronization will begin.

**NOTE:** The ‘filesys’ type is the default since often file systems will fail to come in-service when there is data inconsistency, especially with the XFS file system. However, file systems can mount when there is data inconsistency or even corruption. Requiring the application to be in-service or the full hierarchy to be in-service may provide higher assurance the data is correct.

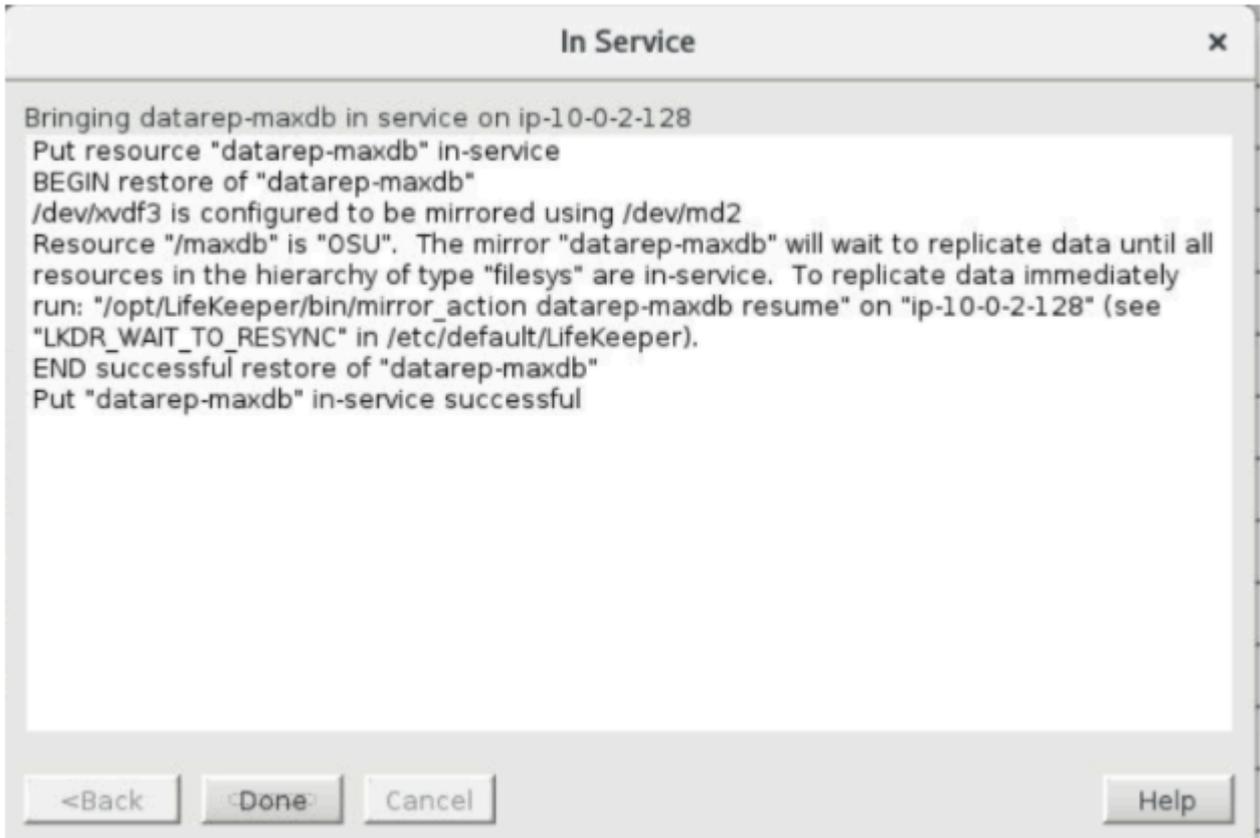
There are 3 ways a user will see information about this feature:

1. The GUI will show the target in the “Wait to Resync” state when a required parent resource is out of service:



2. A warning message is logged during the in-service operation when a required parent resource is out of service. The message will specify the type that is being checked, and will also specify that the full hierarchy is being checked in the case that “LKDR\_WAIT\_TO\_RESYNC=hierarchy”.

- a. Example log message shown during the in-service operation where LKDR\_WAIT\_TO\_RESYNC=filesys:



b. Example log message in the log file where LKDR\_WAIT\_TO\_RESYNC=hierarchy:

WARN:dr:recover:datarep-maxdb:104237:Mirror "datarep-maxdb" will wait to reconnect targets until parent file system "/maxdb" is in-service. To reconnect targets immediately run: "/opt/LifeKeeper/bin/mirror\_action datarep-maxdb resume" on "ip-10-0-2-128" (see "LKDR\_WAIT\_FOR\_FILE\_SYSTEM\_TO\_MOUNT" in /etc/default/LifeKeeper).

```

Nov  2 09:07:06 ip-10-0-2-128 restore[4037]: WARN:dr:recover:datarep-maxdb:104237:Resource "/maxdb" is "OSU". The mirror
"datarep-maxdb" will wait to replicate data until all resources in the hierarchy of type "filesystem" are in-service. To r
eplicate data immediately run: "/opt/LifeKeeper/bin/mirror_action datarep-maxdb resume" on "ip-10-0-2-128" (see "LKDR_WAI
T_TO_RESYNC" in /etc/default/LifeKeeper).

```

3. An emergency message is logged both to the LifeKeeper log file as well as all open terminals when a required parent resource is in the failed state (OSF).

a. In this case, the GUI will show the same "Wait to Resync" status.



b. The following message will be logged to all terminals and to the LifeKeeper log file:

EMERG:dr:recover:datarep-maxdb:104236:Resource "/maxdb" is "OSF". The mirror "datarep-maxdb" will wait to reconnect targets until parent file system "/maxdb" is in-service. This may indicate inconsistent data. Do not bring the resource "/maxdb" in-service until the data has

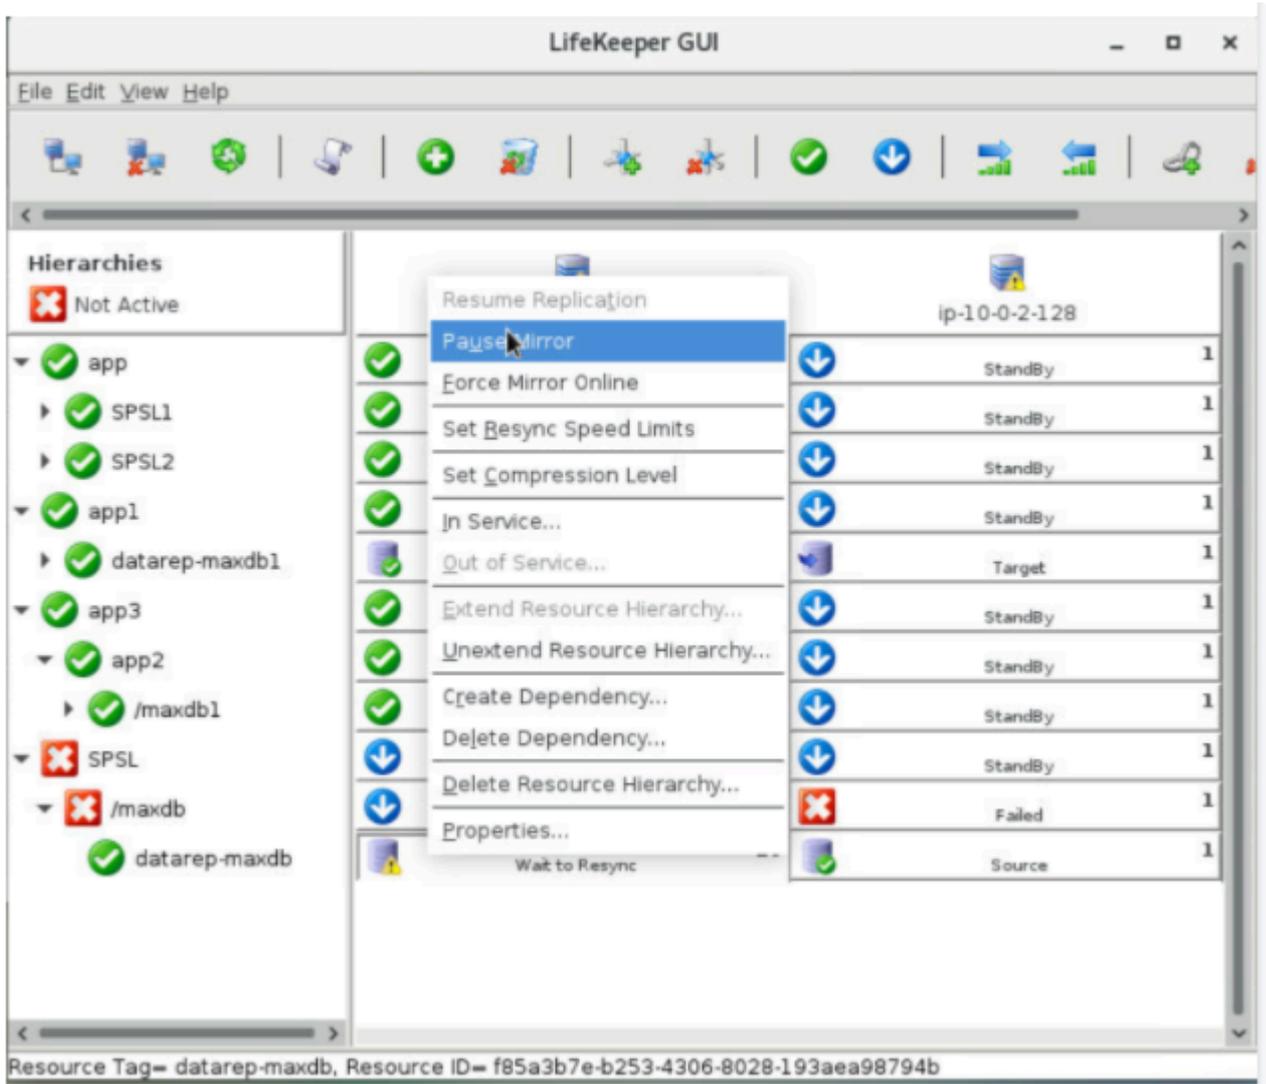
been verified; replication will continue when “/maxdb” is in-service. A full resync may be necessary (see “LKDR\_WAIT\_FOR\_FILE\_SYSTEM\_TO\_MOUNT” in /etc/default/LifeKeeper).

```
Oct 30 09:29:58 ip-10-0-2-128 recover[17040]: EMERG:dr:recover:datarep-maxdb:104236:Resource "/maxdb" is "OSF". The mirror "datarep-maxdb" will wait to replicate data until all resources in the hierarchy of type "filesys" are in-service. This may indicate inconsistent data. Verify the data is correct before replicating data; replication will continue when all resources in the hierarchy of type "filesys" are in-service. A full resync may be necessary (see "LKDR_WAIT_TO_RESYNC" in /etc/default/LifeKeeper).
```

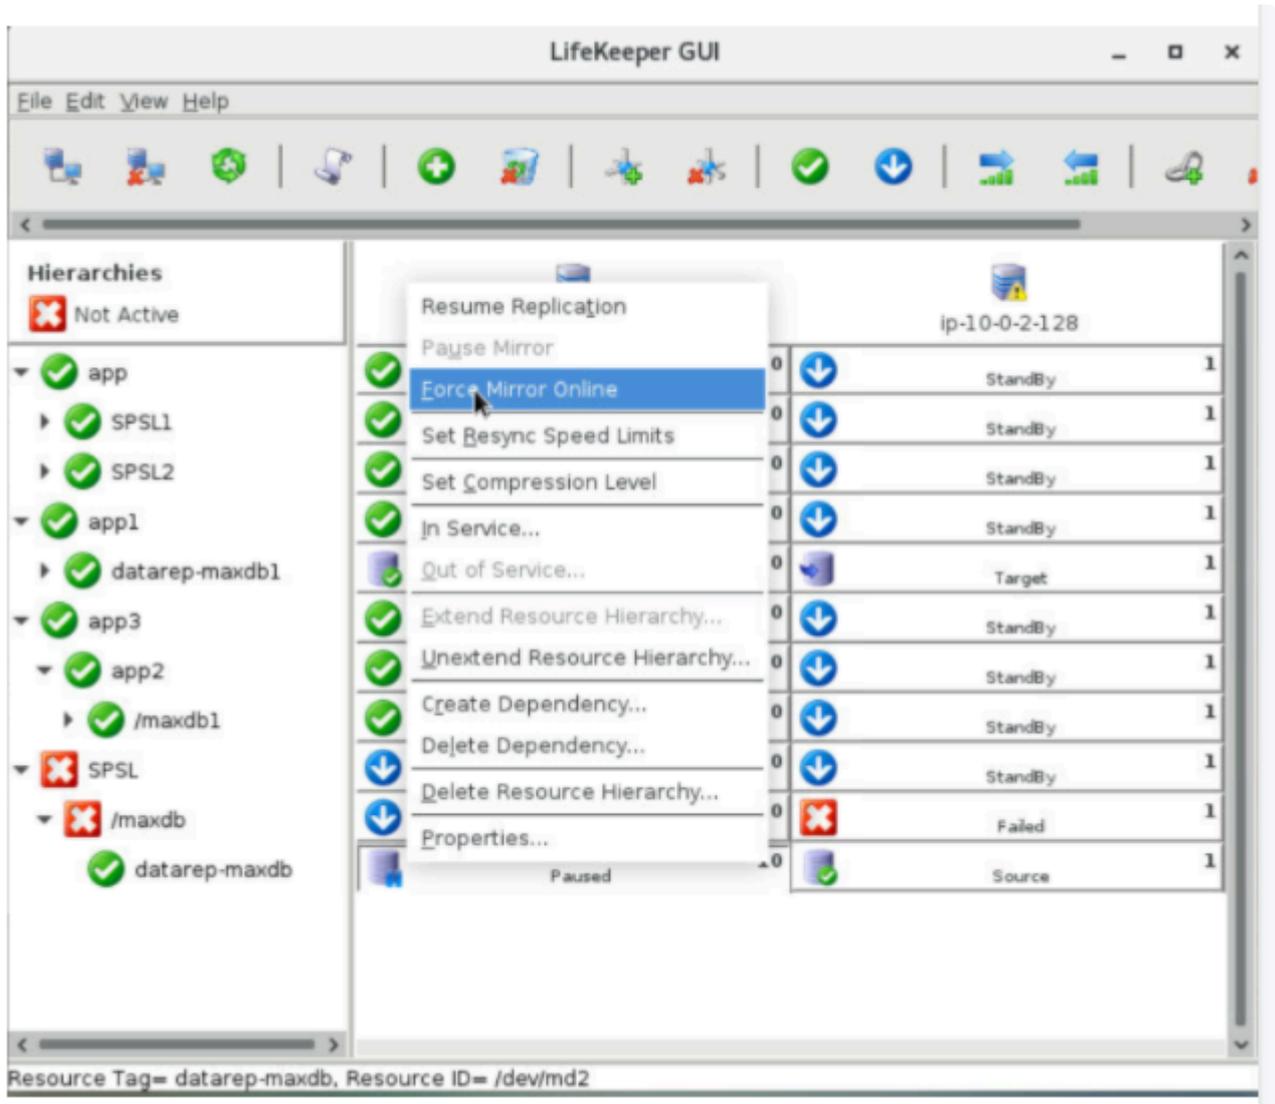
c. Do not bring the file system resource in-service until the data on the file system is verified. Once the file system resource is brought in-service replication will resume during the next quickCheck cycle for the DataKeeper resource.

d. The failed file system resource indicates that mount, log replay, and fsck are unable to repair the file system. This may indicate inconsistent data on the disk if the file system is able to come in-service on another node. Once the file system is repaired either by recovery on the node that failed or by switching to another node, a full resync may be necessary to ensure that all nodes have consistent data.

e. The mirror can be mounted on a target to check if the file system can be mounted and the data verified. This can be done using the “Pause Mirror” feature.



When the mirror is paused LifeKeeper will automatically mount the file system on the target. If this is successful, verify the data is correct. The mirror can be “forced online” by choosing the “Force Mirror Online” option on the target (“paused” server). **WARNING:** This operation will not resume replication until the appropriate parent resources are in-service as defined by the LK\_WAIT\_TO\_RESYNC on the server where the “force online” is being performed.



**DO NOT** choose the “resume Replication” option until there is high confidence that the data on the source is correct. Choosing “Resume Replication” after pausing the mirror will resume replication on the source though parent resources are not in-service (even if they are failed).

f. Once the data has been verified and the faulty parent resource has been repaired and brought in-service, the next quickCheck cycle for the DataKeeper resource will detect that it should no longer wait to resynchronize the data. There will be a time period (up to 2 minutes) before this quickCheck cycle where the “Wait to Resync” state is displayed even though the parent resource is in-service.

<ul style="list-style-type: none"> <li>✖ SPSL</li> <li>✓ /maxdb</li> <li>✓ datarep-maxdb</li> </ul>	<table border="1"> <tr> <td>↓</td> <td>StandBy</td> <td>10</td> <td>↓</td> <td>StandBy</td> <td>1</td> </tr> <tr> <td>↓</td> <td>StandBy</td> <td>10</td> <td>✓</td> <td>Active</td> <td>1</td> </tr> <tr> <td>⚠</td> <td>Wait to Resync</td> <td>10</td> <td>✓</td> <td>Source</td> <td>1</td> </tr> </table>	↓	StandBy	10	↓	StandBy	1	↓	StandBy	10	✓	Active	1	⚠	Wait to Resync	10	✓	Source	1
↓	StandBy	10	↓	StandBy	1														
↓	StandBy	10	✓	Active	1														
⚠	Wait to Resync	10	✓	Source	1														

g. If corruption is found (for example, the file system will mount on one node but will not on another) or suspected, then a full resync is advised. The mirror must first be paused to force a full resync. Once the mirror is paused run the following command on the source system:

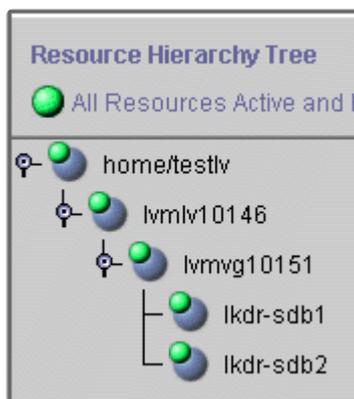
```
/opt/LifeKeeper/bin/mirror_action <tag> fullresync
```

## 5.5.5. Using LVM with DataKeeper

LifeKeeper for Linux currently supports both the use of DataKeeper “above” LVM and LVM “above” DataKeeper. In a standard DataKeeper configuration, using DataKeeper above LVM is supported and DO NOT install the LifeKeeper LVM Recovery Kit. DataKeeper is the only recovery kit necessary. However, using the LVM above DataKeeper configuration, the LVM Recovery Kit is required.

SIOS recommends using DataKeeper above LVM; however, if the LVM above DataKeeper configuration is being used, a two-phase hierarchy creation process must be used. The DataKeeper devices (i.e. hierarchies) must be configured using the DataKeeper “Data Replication Resource” option prior to the creation of the LVM volume groups and logical volumes on the primary server. Once the desired volume groups and logical volumes have been created, the remainder of the hierarchy is created according to the configuration instructions for the recovery kit associated with the application to be protected. The resulting hierarchy will look something like the one shown in Figure 3 below.

**Note:** For data consistency reasons, in an LVM over DataKeeper configuration, there must either be only one DataKeeper mirror or multiple **synchronous** mirrors.

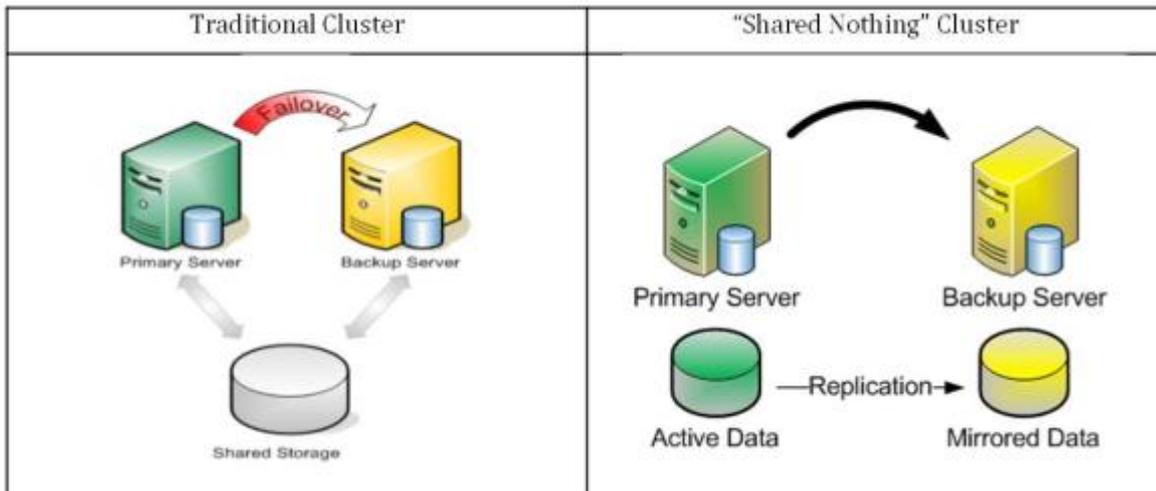


**Figure 3: Hierarchy with LVM above DataKeeper**

## 5.5.6. Clustering with Fusion-io

### Fusion-io Best Practices for Maximizing DataKeeper Performance

LifeKeeper for Linux includes integrated, block level data replication functionality that makes it very easy to set up a cluster when there is no shared storage involved. Using Fusion-io, LifeKeeper for Linux allows you to form “shared nothing” clusters for failover protection.



When leveraging data replication as part of a cluster configuration, it is critical that you have enough bandwidth so that data can be replicated across the network just as fast as it is being written to disk. The following best practices will allow you to get the most out of your “shared nothing” LifeKeeper cluster configuration when high-speed storage is involved:

#### Network

- **Use a 10 Gbps NIC:** Flash-based storage devices from Fusion-io (or other similar products from OCZ, LSI, etc.) are capable of writing data at speeds of HUNDREDS (750+) MB/sec or more. A 1 Gbps NIC can only push a theoretical maximum of approximately 125 MB/sec, so anyone taking advantage of an ioDrive's potential can easily write data much faster than 1 Gbps network connection could replicate it. To ensure that you have sufficient bandwidth between servers to facilitate real-time data replication, a 10 Gbps NIC should always be used to carry replication traffic.
- **Enable Jumbo Frames:** Assuming that your network cards and switches support it, enabling jumbo frames can greatly increase your network's throughput while at the same time reducing CPU cycles. To enable jumbo frames, perform the following configuration (example on a Red Hat/CentOS/OEL Linux distribution):

° Run the following command:

```
ip link set <interface_name> mtu 9000
```

° To ensure change persists across reboots, add “MTU=9000” to the following file:

```
/etc/sysconfig/network-scripts/ifcfg-<interface_name>
```

° To verify end-to-end jumbo frame operation, run the following command:

```
ping -s 8900 -M do <IP-of-other-server>
```

- **Change the NIC’s transmit queue length:**

° Run the following command:

```
ip link set <interface_name> txqueuelen 10000
```

° To preserve the setting across reboots, add to */etc/rc.local*.

- **Change the NIC’s netdev\_max\_backlog:**

° Set the following in */etc/sysctl.conf*:

```
net.core.netdev_max_backlog = 100000
```

## TCP/IP Tuning

- **TCP/IP tuning** that has shown to increase replication performance:

° Edit */etc/sysctl.conf* and add the following parameters (**Note:** These are examples and may vary according to your environment):

```
net.core.rmem_default = 16777216
```

```
net.core.wmem_default = 16777216
```

```
net.core.rmem_max = 16777216
```

```
net.core.wmem_max = 16777216
```

```
net.ipv4.tcp_rmem = 4096 87380 16777216
```

```
net.ipv4.tcp_wmem = 4096 65536 16777216
```

```
net.ipv4.tcp_timestamps = 0
```

```
net.ipv4.tcp_sack = 0
```

```
net.core.optmem_max = 16777216
```

```
net.ipv4.tcp_congestion_control=htcp
```

## Configuration Recommendations

- Allocate a small (~100 MB) disk partition, located on the Fusion-io drive to place the bitmap file. Create a filesystem on this partition and mount it, for example, at */bitmap*:

```
# mount | grep /bitmap

/dev/fioa1 on /bitmap type ext3 (rw)
```

- Prior to creating your mirror, adjust the following parameters in */etc/default/LifeKeeper*:

```
LKDR_CHUNK_SIZE=4096 (Default value is 256)
```

- Create your mirrors and configure the cluster as you normally would.
- The Bitmap file must be set up to be created in the partition, which is created as above.
- Set up for faster resynchronization. Select “Set Resync Speed Limits” from right menu of DataKeeper Resource and set up the following figure to the wizard. **Note:** Setting the resync speeds via the UI will cause the changes to take place immediately. If they are added to the default file the resource may need to be taken in and out of service.

```
Minimum Resync Speed Limit: 200000
```

- At the same time, set up Resync speed to be allowed during other I/Os operating. This figure must be set up under the half of the maximum write speed throughput of the drive as the empirical rule not to disturb the normal I/O functions during the Resync operation.

```
Maximum Resync Speed Limit: 1500000
```

- Set up the maximum bandwidth to use during Resync. This figure must be set up with enough high figure to execute Resync with the available maximum speed.

## 5.5.7. Using External Snapshot Functions for Disks and Devices Protected by DataKeeper

---

Full synchronization is required when using a snapshot process to restore data to a disk or device that is actively protected by DataKeeper.

Snapshot capability referred to in this document includes:

1. Snapshots provided by cloud environment services such as AWS
2. Snapshots provided by virtualization software such as vSphere
3. Snapshots provided by shared storage in physical environment

The process of restoring snapshots typically takes place without involvement of the operating system. When the snapshot is restored without involving the OS DataKeeper cannot properly synchronize these changes to the target system. In addition, the changes to the underlying disk or device are not written in the bitmap, so consistency is not maintained with differential synchronization.

The recommended procedure for using snapshots is as follows:

1. Stop the mirror with “pause mirror” command
2. Restore the snapshot
3. Perform a complete re-synchronization with restored data as the source

The source and target data are in an inconsistent state until the mirror is fully synchronized. Therefore, it is essential to execute full synchronization after the snapshot has been restored. Please note a full resync can take a long time depending on the mirror size, available bandwidth, and system resources. During the resync period, the DataKeeper resource and any dependent resources cannot be switched over.

Full synchronization may not be required when restoring the same snapshot to both the source and the target, but this operation is not supported.

## 5.5.8. DataKeeper for Linux Troubleshooting

This section provides information regarding issues that may be encountered with the use of DataKeeper for Linux. Where appropriate, additional explanation of the cause of an error is provided along with necessary action to resolve the error condition.

Messages specific to DataKeeper for Linux can be found in the [DataKeeper Message Catalog](#). Messages from other LifeKeeper components are also possible. In these cases, please refer to the [Combined Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

The following table lists possible problems and suggestions.

Symptom	Suggested Action
Wait to Resync	Resynchronization is waiting for a parent resource to come in-service. The LifeKeeper log will contain a message indicating the resource that is not in-service that is blocking replication as well as what needs to be in-service for resynchronization to begin. If the resources are simply OSU then bring the resources in-service or run the <code>mirror_action</code> command listed in the log message. If a resource is OSF then resolve the problem before bringing resources in-service that will resynchronize the data. See <a href="#">Verify Data Before Resync</a> for more details.
Import Failure	When creating a mirror using the GUI you are asked if you want to allow asynchronous replication. If you allow asynchronous during create then when you extend you are asked if you want to make the connection to that target synchronous or asynchronous. If you create the mirror asynchronous but extend synchronous then export the configuration and do an import, the import will fail.  <b>Note:</b> This configuration is not supported with the lkcli export/import. The lkcli export will not return an error but the lkcli import will fail.
Warning message that netraid mirror does not have a unique identifier.	The configuration should be modified as soon as possible to use a unique identifier.  The recommended steps to repair are: 1. Start with LifeKeeper running on all nodes. 2. Identify unsafe netraid resources and the underlying disks for each resource on each node. It is important to do this on each node as the mapping may be DIFFERENT on each node.  <pre># ins_list -r netraid -f:   grep DEVNAME   grep -v mapper   cut -f4,5 -d: datarep-test1:/dev/xvdb</pre>

```
datarep-test2:/dev/xvdc
datarep-filesys3:/dev/xvdd
```

**NOTE:** This is the list of netraid “tag:ID”. In the following instructions the device name mapping matching above is assumed.

3. Check if device is configured with GPT. For each device run the GPT getId:

```
#!/opt/LifeKeeper/lkadm/subsys/scsi/gpt/bin/getId -i /dev/
xvdb
```

```
4757cd62-e065-4013-8514-1031b446aa24
```

4. If getId returns a unique ID then update the instance:

```
#ins_setid -t datarep-test1 -i “4757cd62-e065-4013-8514-1031b446aa24”
```

<If there are any devices that are not GPT then continue with Step 5>

5. Identify resources that depend on the unsafe netraid resources.

```
# ins_list -r netraid -f: | grep DEVNAME | grep -v mapper | cut -f4 -d: | while
read entry; do dep_list -p $entry -f: 2>/dev/null | cut -f1 -d:; done
/test1
/test2
filesys3
```

**NOTE:** this is the **list of tags** for the file systems associated with the netraid devices.

6. Identify the mount points for the file system. Typically the tag for the file system resource is the same as the mount point. If that is not the case you can match the file system tag with the file system mount point using:

```
# ins_list -t filesys3 -f: | cut -d: -f5
/test3
```

7. Stop all activity leaving only the file system resources in-service on netraid devices.

	<p>a. Take all resources out-of-service.  b. Bring in-service only the file systems on unsafe netraid devices.</p> <p>Follow the steps below only for the devices that are unsafe.</p> <p>8. Backup all data on affected file systems where the resources are in-service.</p> <p>9. Take all resources out-of-service on all cluster nodes.</p> <p>10. Backup the LifeKeeper configuration (lkbackup -c —cluster).</p> <p>11. Delete the hierarchy with each unsafe netraid resource.</p> <p>a. Take note of the hierarchy, what application is affected is being deleted.  b. For complex configurations this may require deleting multiple hierarchies where the hierarchy is made up of multiple branches.</p> <p>12. At this point, the only things left are resources that do not have dependencies with unsafe netraid resources. In most cases that should be the IP resources, EC2 resources, etc.</p> <p>13. Run lkstop on all nodes.</p> <p>14. Verify all affected file systems are unmounted.</p> <p>15. Reconfigure devices on each node with a GPT partition table (using gdisk, parted, etc) or use LVM.</p> <p>16. Start LifeKeeper on all nodes.</p> <p>17. Create new Replicated file systems for each file system, extending each resource to all nodes.</p> <ul style="list-style-type: none"> <li>• /dev/xvdb1 -&gt; /test1</li> <li>• /dev/xvdc1 -&gt; /test2</li> <li>• /dev/xvdd1 -&gt; /test3</li> </ul> <p>18. Restore data from the backup to each mount point. The data will automatically resync to the target(s).</p> <p>19. Recreate application hierarchies deleted in step 11.</p> <p>Please refer to the <a href="#">SIOS Product Documentation</a> for details on DataKeeper storage configuration options.</p>
<p>After primary server panics,</p>	<p>Check the “switchback type” selected when creating your DataKeeper resource hierarchy. Automatic switchback is not supported for DataKeeper resources in this release. You can</p>

<p>DataKeeper resource goes ISP on the secondary server. When primary server reboots, the DataKeeper resource becomes OSF on both servers.</p>	<p>change the Switchback type to “Intelligent” from the resource properties window.</p>
<p>DataKeeper GUI wizard does not list a newly created partition.</p>	<p>The Linux OS may not recognize a newly created partition until the next reboot of the system. View the <i>/proc/partitions</i> file for an entry of your newly created partition. If your new partition does not appear in the file, you will need to reboot your system.</p>
<p>Errors during failover</p>	<p>Check the status of your device. If resynchronization is in progress you cannot perform a failover.</p>
<p>Error creating a DataKeeper hierarchy on currently mounted NFS file system</p>	<p>You are attempting to create a DataKeeper hierarchy on a file system that is currently exported by NFS. You will need to replicate this file system before you export it.</p>
<p>Extending to a target does not prompt for “Replication Type” to allow setting asynchronous or synchronous.</p>	<p>When the mirror was created, “no” was selected for “Enable Asynchronous Replication.” Delete the mirror and recreate selecting “yes” to “Enable Asynchronous Replication” when prompted.</p>
<p>NetRAID device not deleted after DataKeeper resource deletion.</p>	<p>Deleting a DataKeeper resource will not delete the NetRAID device if the NetRAID device is mounted. You can manually unmount the device and delete it by executing: <i>mdadm -S &lt;md_device&gt;</i> (<i>cat /proc/mdstat</i> to determine the <i>&lt;md_device&gt;</i>).</p>
<p>Primary server cannot bring the resource ISP when it reboots after both servers became inoperable.</p>	<p>If the primary server becomes operable before the secondary server, you can force the DataKeeper resource online by opening the resource properties dialog, clicking the <b>Replication Status</b> tab, clicking the <b>Actions</b> button, and then selecting <b>Force Mirror Online</b>. Click <b>Continue</b> to confirm, then <b>Finish</b>.</p>

<p>Replication Type is asynchronous instead of synchronous. Replication between two systems was initially configured for asynchronous replication, but synchronous replication is required instead.</p>	<p>Unextend the mirror and extend again, selecting “synchronous” when prompted for the connection.</p>
<p>Replication Type is synchronous instead of asynchronous. Replication between two systems was initially configured for synchronous replication, but asynchronous replication is required instead.</p>	<p>Unextend the mirror and extend again, selecting “asynchronous” when prompted for the connection.</p>
<p>Resources appear green (ISP) on both primary and backup servers.</p>	<p>This is a “split-brain” scenario that can be caused by a temporary communications failure. After communications are resumed, both systems assume they are primary.</p> <p>DataKeeper will not resync the data because it does not know which system was the last primary system. Manual intervention is required.</p> <p><b>If not using a bitmap:</b></p> <p>You must determine which server was the last backup, then take the resource out of service on that server. DataKeeper will then perform a FULL resync.</p> <p><b>If using a bitmap:</b></p> <p>You must determine which server was the last backup, then take the resource out of</p>

	<p>service on that server. DataKeeper will then perform a partial resync.</p>
<p>Target(s) are out of sync waiting for the previous source.</p>	<p>Connect the previous source to the cluster. If the previous source can not rejoin the cluster in a timely manner, then targets can be reconnected with a full resync by running the command “\$LKROOT/bin/mirror_action fullresync &lt;source&gt; &lt;target&gt;” on the current mirror source.</p>
<p>Core – Language Environment Effects</p>	<p>Some LifeKeeper scripts parse the output of Linux system utilities and rely on certain patterns in order to extract information. When some of these commands run under non-English locales, the expected patterns are altered and LifeKeeper scripts fail to retrieve the needed information. For this reason, the language environment variable LC_MESSAGES has been set to the POSIX “C” locale (LC_MESSAGES=C) in <i>/etc/default/LifeKeeper</i>. It is not necessary to install Linux with the language set to English (any language variant available with your installation media may be chosen); the setting of LC_MESSAGES in <i>/etc/default/LifeKeeper</i> will only influence LifeKeeper. If you change the value of LC_MESSAGES in <i>/etc/default/LifeKeeper</i>, be aware that it may adversely affect the way LifeKeeper operates. The side effects depend on whether or not message catalogs are installed for various languages and utilities and if they produce text output that LifeKeeper does not expect.</p>
<p>GUI – GUI login prompt may not re-appear when reconnecting via a web browser after exiting the GUI</p>	<p>When you exit or disconnect from the GUI applet and then try to reconnect from the same web browser session, the login prompt may not appear.</p> <p><b>Workaround:</b> Close the web browser, re-open the browser and then connect to the server. When using the Firefox browser, close all Firefox windows and re-open.</p>
<p>DataKeeper Create (and Extend) Resource fails</p>	<p>When using DataKeeper in certain environments (e.g., virtualized environments with IDE disk emulation, servers with HP CCISS storage, solid state devices (SSD), or Amazon EBS storage), an error may occur when a mirror is created:</p>  <p>This is because LifeKeeper does not recognize the disk in question and cannot get a unique ID to associate with the device.</p>

	<p><b>Workaround:</b> Create a GUID partition and assign a unique ID to the partition or use LVM.</p>
<p>The status of the mirror target becomes “Out of Sync” after upgrading</p>	<p>The use of an NU device is not recommended for LifeKeeper 9.2.2 or later. A “mirror out of sync” problem occurs in environments where DataKeeper resources are configured with NU devices.</p> <p>When upgrading to LifeKeeper 9.2.2 or later, add the following settings to <code>/etc/default/LifeKeeper</code> if NU devices are used:</p> <pre style="text-align: center;">LKDR_ALLOW_NU=TRUE</pre> <p><b>How to check whether NU devices are used:</b> Run the <code>lkdstatus</code> command. If a resource instance ID field contains a character string beginning with NU-, then NU devices are used.</p>

## 5.6. Command Line Interface

---

LifeKeeper's Command Line Interface can be used as an alternative to the Graphical User Interface. Administrator tasks may be automated by incorporating calls to the CLI in shell scripting.

### Document Contents

This guide contains the following topics:

- [Commands](#). Describes CLI commands.
- [Shell Script Examples](#). Provides some examples of CLI used in scripting.

### LifeKeeper Documentation

The following is a list of LifeKeeper related information available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper Installation Guide](#)
- [Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with other LifeKeeper Recovery Kits, is available online at:

<http://docs.us.sios.com/>

# Shell Script Examples

---

## Examples Pulled Out of Shell Scripts to Create and Extend Hierarchies.

Also how to create a dependency between two hierarchies.

### # Needed System Parameters

```
LKROOT=/opt/LifeKeeper  
  
OBJ_DIR=/opt/LifeKeeper/lkadm  
  
LKBIN=/opt/LifeKeeper/bin  
  
ExtendPath=/opt/LifeKeeper/lkadm/bin  
  
PATH=$PATH:$LKBIN
```

### # Generic ARGS

```
LocalServer=unix121.ha.uk.sbphrd.com  
  
TargetServer=unix122.ha.uk.sbphrd.com  
  
Node2Priority=10
```

### # The above variables are used for the commands that follow

## 5.6.1. Commands

### Groupings and Basic Descriptions of LifeKeeper bin Commands

The commands will be in: `/opt/LifeKeeper/bin`

To place this in your path, execute: ``. /etc/default/LifeKeeper``

#### Starting and Stopping LifeKeeper, the GUI, etc.

! When LifeKeeper starts and establishes communication with the other servers in a cluster, it **will not** allow data replication (*DataKeeper*) resources to come in-service (*manual or automatic*) until communication is established with all servers in the cluster. The data replication resources will be marked OSF if an attempt is made to *in-service* them before communication is established with all servers. (*Please refer to the system's log files for additional information.*)

`lkstart` – Start LifeKeeper core

#### Options:

None

`lkstop` – Stop LifeKeeper core. The behavior differs depending on the command line options.

! When '`lkstop -f`' is used to stop LifeKeeper, in-service resources are left configured and running, including data replication resources. **It is important that the hierarchy/resources are allowed to be brought back in-service on the correct server after an '`lkstop -f`'.** LifeKeeper will automatically bring all resources in-service once communication is established with all servers. Data replication resources may temporarily be marked OSF while LifeKeeper is waiting for all servers to restart and rejoin the cluster. (*Please refer to the system's log files for additional information.*)

\* **Note:** The `lkstop` process will continue even when a resource remove fails and will not result in a failed `lkstop`.

\* **Note:** Avoid running `lkstop` and/or shutting down multiple cluster nodes in parallel.

\* **Note:** We recommend waiting a minimum of two minutes between issuing each `lkstop`

command to allow processing to be completed within the cluster.

**Options:**

None – Stop LifeKeeper Core and the protected services. No switchover to the standby system is performed regardless of the shutdown strategy setting.

-r – Leave auto-start on system startup enabled.

-n – Stop LifeKeeper Core and the protected services. If the shutdown strategy is set to “Switchover Resources”, the system will be switched over to the standby system. Also, if LKSTOP\_FAILOVER\_N=1 is set in /etc/default/LifeKeeper, failover will occur regardless of the shutdown strategy setting.

-f – Stop LifeKeeper core but do not stop the protected resources.

-i – Stop LifeKeeper core but do not stop the protected resources. The user is prompted to confirm (yes/no) that they want to continue.

`lkGUIserver` – Start and stop the LifeKeeper GUI daemon processes

**Options:**

`start`

`stop`

`restart`

`lkGUIapp` – Starts the LifeKeeper Java application

**Options:**

None

## Monitoring LK and Other Misc. LifeKeeper Commands

`lcdstatus` – Display status of LifeKeeper resources, comm paths, etc.

**Options:**

-d – <node to run command on>

-q – short reports

`lcdsync` – Writes LifeKeeper configuration information from memory to disk

**Options:**

```
-d - <other node to run it on>
```

lcdrcp – Transfer files from one LifeKeeper node to another via the comm. path

**Options:**

```
lcdrcp <file names> {dest:ofile | dest:odir}
```

lcdremexec – Execute the given command on the given LifeKeeper node

**Options:**

```
-d <node to run command on> <command>
```

lcdrecover – Checks and sets the resource hierarchy instance status.

**Options:**

```
See man -M /opt/LifeKeeper/man lcdrecover.
```

## Bringing a Hierarchy into and out of Service

perform\_action – Performs a given action on a given resource and can be used to switch a given hierarchy to another node.

**Options:**

```
-a <action name>
```

```
-t <tag name>
```

**Examples:**

```
perform_action -a restore -t $LKTag – bring tier into service
```

```
perform_action -a remove -t $LKTag – take tier out of service
```

## Checking the LifeKeeper Configuration

lkchkconf : Performs the following checks to verify /etc/default/LifeKeeper settings.

- Checks that the running system is actually using the current settings found in /etc/default/LifeKeeper.

If the current setting of LifeKeeper is different from the /etc/default/LifeKeeper setting, an error message is output to inform the user.

- Checks for any inconsistencies between the resource health check time interval(LKCHECKINTERVAL) and the timeout value of each ARK.  
An error message will be logged if the timeout value of each ARK is longer than the resource health check time interval (LKCHECKINTERVAL).

**Options:** – None

## 5.6.1.1. lklogmsg

---

This command provides a command line interface to the LifeKeeper logging interface. It can be used in scripts that are part of the implementation of a Recovery Kit.

### NAME

`lklogmsg` — command line interface to the LifeKeeper logging interface

### SYNOPSIS

```
LKROOT/bin/lklogmsg [-o] [-e] [-h] [-p progame] -l level -s source -a action  
-t tag -i messageid [-c command - cmd args ]| [-- message]
```

### SYNTAX

```
lklogmsg [-o] [-e] [-h] [-p progame] -l level -s source -a action -t tag -i  
messageid [-c command - cmd args ]| [-- message]
```

Messages that are to be logged may either be passed in as the final arguments with the special argument “--” in front of it, or if no message parameters are present, this command will read input from `stdin`, and convert each line into a separate log message.

An additional mode of operation can be used where a command and its arguments are passed in via the `-c` option. In this mode, the `stdout` and `stderr` of the command are redirected through a pipe, and captured by `face="Courier New">lklogmsg`. The `stdin` of the command comes from `/dev/null`.

The `-o` option will cause the message to be echoed to `stdout`.

The `-e` option will cause the message to be echoed to `stderr`.

The `-h` option will cause `lklogmsg` to perform the same actions as `nohup`.

The `-p` should be the name of the script which is calling `lklogmsg`. It will be used as the application name that gets recorded in the log.

The `-l` option must be one of `LK_NOTIFY`, `LK_FATAL`, `LK_ERROR`, `LK_WARN`, `LK_TRACE`, `LK_INFO` or `LK_DEBUG`.

The `-s` option should be the Recovery Kit name which is calling `lklogmsg` (e.g., `mysql`, `lvm`, `oracle`).

The `-a` option should be the Recovery Kit action which is calling `lklogmsg` (e.g., `restore`, `remove`, `recover`, `create`, `extend`).

The `-i` option should be the (6-digit) message ID from the assigned range for the Recovery Kit.

## EXAMPLES

```
/opt/LifeKeeper/bin/lklogmsg -e -l LK_ERROR -s gopher -a recover -t  
sample_tagname -i 222016 Unable to restart gopher server sends "Unable to restart gopher server"  
to both the LifeKeeper log, and to stderr.
```

```
/opt/LifeKeeper/bin/lklogmsg -l LK_INFO -s webserver -a create -t  
sample_sdr_resource -i 100315
```

## 5.6.1.2. SYS – LifeKeeper Commands Related to the Systems in the LifeKeeper Cluster

---

`sys_list` – Lists out the systems known to a particular LifeKeeper node

**Options:**

`-d <other node to run it on>`

`sys_create` – Creates knowledge of another system on LifeKeeper node

**Options:**

`-s <remote system name>`

`-d <node to run command on>`

`sys_remove` – Removes knowledge of another system on a LifeKeeper node

**Options:**

`-d <dest>`

`-s <remote system name to be removed from list>`

`sys_getstate` – Lists the state of a given LifeKeeper node on the given LifeKeeper node

**Options:**

`-d <node to run command on>`

`-s <system concerning the state being checked>`

`sys_setstate` – Sets the state of a given LifeKeeper node on a given LifeKeeper node

**Options:**

`-d <node to run command on>`

`-s <system concerning the state being set>`

`-S <actual state> {DEAD|ALIVE|UNKNOWN}`

`-R <reason for state setting>`

`sys_getdescr` – Prints some information of why the system went to its current state

**Options:**

`-d <node to run command on>`

`-s <system to get data on>`

## 5.6.1.3. NET – Communication Paths Related Commands

---

`net_create` - Creates a communication path between two LifeKeeper nodes

**Options:**

- d <node to run command on>
- s <other system>
- D <device path>
- n <TTY or TCP>
- b <baud rate>
- r <remote IP address>
- l <local IP address>
- p <priority>

`net_remove` - Removes a communication path between two LifeKeeper nodes

**Options:**

- d <node to run command on>
- s <remote server name to be removed from>
- D <device path>
- r <remote IP address>

`net_list` - Lists communication path information on a given LifeKeeper node

**Option:**

- d <node to run command on>
- f: <field separator of ':'>
- s <system name>

`net_change` - Modify specific information about a given communication path

**Options:**

-d <node to run command on>

-s <server name for data to be modified>

-D <device>

**Cre1cm - Create a communication path**

```
/opt/LifeKeeper/lkadm/bin/cre1cm <node 1> <node 2> <net type> <baud  
rate> <IP address 1> <IP address 2> <prio>
```

**portio - Tests the serial connection between two LifeKeeper nodes**

## 5.6.1.4. FLAG – Commands Related to Internal LifeKeeper Flags

---

`flg_create` - Set a given LifeKeeper flag on a given node

**Options:**

`-d <node to run command on>`

`-f <flag name>`

`flg_remove` - Remove a given LifeKeeper flag on a given node

**Options:**

`-d <node to run command on>`

`-f <flag name>`

`flg_list` - List all LifeKeeper flags that are set on a given node

**Options:**

`-d <node to run command on>`

## 5.6.1.5. TYP – LifeKeeper Commands Related to Resource Hierarchy Types

---

`typ_create` - Creates a given resource type on a given node

### Options:

- d <node to run command on>
- a <app type> (need an app first)
- r <resource type>

`typ_remove` - Removes the given resource type from the configuration database set of known resource types on the specified system (or local system if no additional system is specified with the `-d dest` option)

### Options:

- d <node to run command on>
- a <application type>
- t <resource type>

`typ_list` - Lists all resource types on a given node

### Options:

- d <node to run command on>
- f: <field separator of ':'>
- a <app type>

## 5.6.1.6. APP – LifeKeeper Commands Related to Resource Applications (Group of Related Types)

---

`app_create` - Creates a given resource application on a given node

**Options:**

`-d <node to run on>`

`-a <application name>`

`app_remove` - Removes the given application from the configuration database set of known applications on the specified system (or local system if no additional system is specified with the `-d dest` option)

**Options:**

`-d <dest>`

`-a <application type>`

`app_list` - Lists all resource applications on a given node

**Options:**

`-d <node to run on>`

## 5.6.1.7. DEP – LifeKeeper Commands Related to How Resource Applications Relate to Each Other

### Running the Commands:

- **Start** from `/opt/LifeKeeper/bin`
- **Run** [one of the below commands] accompanied by [one of the “options”] (`./dep_list^-P`)

## Commands

\* [More information on these commands is available via the man pages installed with LifeKeeper.](#)

**dep\_create** – Creates a dependency between two resource instances

### Options/Descriptions

#### Options:

- p <parent tag>
- c <child tag>
- d <dest> (OPTIONAL)

**Description:** This function creates a dependency relationship between the resource instances with tags <parent tag> and <child tag> on one system.

#### Additional Information:

- Both resources must be on the system on which the command is run on the system specified by -d argument.
- If the destination is not specified, the current system is assumed. This implies the parent resource now requires the child for proper operation.
- **Both resource instances must already exist.**

\* **Note:** The dependency is only created on the system on which the command is run or on the system specified by the **-d arg**. To ensure complete setup of the dependency, the command should be run for all nodes. **Failure to do so can result in unexpected failures.**

**dep\_remove** – Removes a dependency between two resource instances

**Options/Description****Options:**

```
-p <parent tag> (OPTIONAL)
-c <child tag> (OPTIONAL)
-d <dest> (OPTIONAL)
```

**Description:** Removes the dependency relationship between the resource instances with tags <parent tag> and <child tag> on one system.

**Additional Information:**

- If dest is not specified, the current system is assumed.
- If child tag is not specified, all dependencies with parent tag are removed.
- If parent tag is not specified, all dependents with child tag are removed.
- The parent tag, the child tag, or both tags must be specified. If **neither are specified**, a usage error will appear.



**Note:** The dependency is only removed on the system that you run the command or on the system specified by the **-d arg**. Therefore, to fully remove dependencies the command should be run for all nodes. **Failure to do so can result in unexpected failures.**

**dep\_list - Lists the dependency relationship between two instances****Options/Descriptions****Options:**

```
-p <ofchildtag> (OPTIONAL)
-c <ofparenttag> (OPTIONAL)
-C <allchild> (OPTIONAL)
-P <allparent> (OPTIONAL)
-r <typ> (OPTIONAL)
-a <app> (OPTIONAL)
-d <dest> (OPTIONAL)
-f <field separator of `:`> (OPTIONAL)
```

**Description:** This function prints strings to stdout. It describes the dependency relationships between resource instances. **Note:** If dest (destination) is not specified, the current system is assumed.

When this command is ran each string will be displayed in the following form:

Parent = part10 Child = ha5

**varfs :** part10

part10:ha5

```
usrfs:part20
```

```
part20:ha5
```

There are two fields in each string that are separated by a **delimiter character**. Use the `-f <any character>` after the `dep_list` command to set a delimiter character to separate the two resource instances in the string. The example above shows a colon (`:`) as a delimiter. The first field indicates the parent tag name of the relationship and the field on the right is the child tag name.

**Example:**

```
[root@sios ~]# ./dep_list -f:
SAP-PDX_ERS10:ers-ip-12.1.2.0
ers-ip-12.1.2.0:ec2-12.1.2.0
SAP-PDX_ASCS00:ascs-ip-12.1.1.0
ascs-ip-12.1.1.0:ec2-12.1.1.0
```

**Note:** If you do not use the `-f` option the output still has a delimiter character of control->A (*This is an unprintable character that will not show up in the string*)

**Additional information:**

- If the `-p` option is specified, this command will print out only the direct parent dependents of the resource specified in **ofchildtag**.
- If the `-c` option is specified, this command will print out only the direct child dependencies of the resource specified in **ofparenttag**.
- If the `-C` option is specified, this command will print out all direct and indirect child dependencies of the resource specified in **allchild**.
- If the `-P` option is specified, this command will print out all direct and indirect parent dependencies of the resource specified in **allparent**.
- If no `-p`, `-c`, `-P`, or `-C` option is specified, all dependencies are printed.
- Specifying the `-r` option lists all the dependencies of child typ.
- Specifying the `-a` option lists all the dependencies of application app.

## eqv\_create – Creates an equivalency between two nodes

### Options/Descriptions

**Options:**

```
-t <tag>
-o <tag on equivalent system>
-S <equivalent system>
-e <{COMMON|SHARED|COMPOSITE}>
-p <Priority> (OPTIONAL)
-d <dest> (OPTIONAL)
-r <Priority on equivalent system (OPTIONAL)
```

**Description:** Creates an equivalency in the configuration database between the resource specified by `<tag>` on the local system, unless a `<dest>` system is specified via the `-d` arg, and the resource on the

<equivalent system> specified by <tag on equivalent system>.

**Additional Information:**

- If the <dest> argument is specified the command runs on the <dest> system, otherwise it runs on local system. See examples below on how using the -d arg impacts the values specified.
- The <Priority> and <Priority on equivalent system> arguments represent the resource priorities on the equivalent systems.
- The <Priority> and <Priority on equivalent system> options will default to 1 and 2 respectively if not provided.
- LifeKeeper will automatically add a equivalency on a remote system specified by the -S <equivalent system>
- For the equivalency type, -e arg, LifeKeeper currently supports **only SHARED**.



For purposes of the examples assume the following:  
 NodeA with resource **TagOnNodeA**  
 NodeB with resource **TagOnNodeB**

**Example:**

To create an equivalency between **TagOnNodeA** and **TagOnNodeB** run the following (both are run on NodeA):

- `eqv_create -t TagOnNodeA -p 1 -s NodeB -o TagOnNodeB -r 10 -e SHARED`
- OR**
- `eqv_create -d NodeB -t TagOnNodeB -p 10 -s NodeA -o TagOnNodeA -r 1 -e SHARED`

## **eqv\_remove** – Removes an equivalency of a given resource between two nodes

### Options/Descriptions

**Options:**

```
-S <equivalent system>
-t <tag>
-e <{COMMON|SHARED|COMPOSITE}>
-d <dest> (OPTIONAL)
-o <tag on equivalent system> (OPTIONAL)
```

**Description:** Removes an equivalency in the configuration database between the resource specified by <tag> on the local system, unless a <dest> system is specified via the -d arg, and the resource on the <equivalent system> specified by <tag on equivalent system>.

**Additional Information:**

- If the <dest> argument is specified the command runs on the <dest> system, otherwise it

runs on local system. See examples below on how using the -d arg impacts the values specified.

- The <tag> represents the resource tag on the local node, unless <dest> is specified in which case the <tag> is the resource tag on <dest>.
- For the equivalency type, -e arg, LifeKeeper currently supports **only SHARED**.
- LifeKeeper will automatically remove the equivalency on the remote system specified by the -S <equivalent system>.



For purposes of the examples assume the following:  
 NodeA with resource **TagOnNodeA**  
 NodeB with resource **TagonNodeB**

#### **Example:**

To remove an equivalency between **TagOnNodeA** and **TagOnNodeB** run the following (both are run on NodeA):

- `eqv_remove -t TagOnNodeA -S NodeB -e SHARED`

#### **OR**

- `eqv_remove -d NodeB -t TagOnNodeB -S NodeA -e SHARED`

## **eqv\_list - Lists equivalency relationships between resource instances**

### **Options/Descriptions**

#### **Options:**

-d <dest> (OPTIONAL)  
 -s <system> (OPTIONAL)  
 -t <tag name> (OPTIONAL)  
 -f: <field separator of `:'> (OPTIONAL)

#### **Description:**

This function prints strings to stdout describing equivalency relationships between resource instances.

#### **Additional Information:**

- If <dest> is specified, then the equivalency listing will be from that system, otherwise it will be from the local system.
- If <system> is specified, then the equivalency listing will be for any equivalencies that exist between the local system.

#### **Example:**

- `iwstp:varfs:remote:varfs_backup:SHARED:1:2`
- `iwstp:usrfs:remote:usrfs_backup:SHARED:1:2`

Each line in the example above contains fields in each string that are separated by a

**delimiter character.** Use the `-f <any character>` after the `eqv_list` command to set a delimiter character to separate the fields of the printed string(s). The above example shows a colon (:) as a delimiter. The fields are as follows (fields 1, 2, and 6 are information for the local system; fields 3, 4, and 7 pertain to the remote system, and field 5 is the equivalency type):

1. Local system name where resource tag 1 of equivalency resides: ***iwstp***
2. Tag name of resource tag 1: ***varfs***
3. Remote system name where resource tag 2 of equivalency resides: ***remote***
4. Tag name of resource tag 2: ***varfs backup***
5. Equivalency type: ***\* \_SHARED.\****
6. Priority value for local system/resource: ***1***
7. Priority value for the remote system/resource: ***2***

The remaining arguments to this function limit the information output as specified below:

8. `-e SHARED` This option prints all SHARED equivalency information.
9. `-t tag` This option limits the output to include only the equivalencies relating to the tag specified by the tag argument.

## **hry\_setpri** - Sets the priority of a given node or hierarchy on the node

### Options/Descriptions

#### Options:

```
-t <tag(s)>
-p <priority>
-d <dest> (OPTIONAL)
-q <Details in Description> (OPTIONAL)
-l: <Details in Description> (OPTIONAL)
```

**Description:** This function sets the resources instances in the specified <tag(s)> resource hierarchy to <priority>. By default the change will be performed on the local node unless `-d <dest>` is specified.

#### Additional Information:

- All associated equivalences are updated.
- The <priority> must not be in use by any existing equivalency.
- The hierarchy is identified by the root resource <tag(s)> specified by the `-t` option. If multiple root hierarchies exist with common resources all the roots must be specified via a comma separated list of tags such as `"-t H1,H2"`. **No imbedded spaces are allowed.**
- The new priority is specified using the `-p` option.
- If the `-q` option is specified, all normal output to `stdout` is suppressed.
- If the `-l` option is specified, a list of systems this hierarchy is resident on, and their associated priorities is displayed. For example:

```
teak 10  
ash 20  
plum 30
```

 **Note:** It is not possible to set the priority of a resource hierarchy that exists on only one system. The `hry_setpri` command will issue a warning message for each resource it encounters that has not been extended beyond its primary system.

## 5.6.1.8. INS – Commands Related to Individual LifeKeeper Hierarchy Instances

---

`ins_list` - Lists the current information of the given resource hierarchy instance

**Options:**

`-d <dest>`

`-f: <field separator of ':'>`

`-a / -r / -t / -i` specify optional app, type, tag, and id info

`ins_setas` - Sets the automatic switchback strategy for a given hierarchy

**Options:**

`-d <dest>`

`-t <tag name>`

`-s <switchback typ> {INTELLIGENT|AUTOMATIC}`

`ins_setinit` - Defines how a given resource should initialize when LifeKeeper starts

**Options:**

`-d <dest>`

`-t <tag name>`

`-I <init state> {AUTORES_ISP|INIT_ISP| INIT_OSU}`

`ins_setinfo` - Defines an information string for a given resource hierarchy

**Options:**

`-d <dest>`

`-t <tag name>`

`-v <string of information>`

`ins_setstate` - Sets the state of a given resource hierarchy on a given node

**Options:**

-d <dest>

-t <tag name>

-S <state to set instance> {ISP|ISU|OSU}

-R <reason for state setting>

-A <recursively set all resources that depend on this one>

ins\_gettag - Lists the tag name of the associated ID

**Options:**

-i <id>

## 5.6.1.8.1. Unextend a Hierarchy

---

```
/opt/LifeKeeper/lkadm/bin/unextmgr <Node_Name> <Tag_Name>
```

## 5.6.1.9. Accessing man (Manual) pages

Manual pages are included with the installation of LifeKeeper for Linux. They are brought in with the **IkMAN package** (see below for the full package name).

```
steel-eye-1kMAN-9.5.1-7154.noarch
```

### Follow the steps below to view the manual pages for the LifeKeeper commands:

1. **Confirm** LifeKeeper for Linux is installed on your OS. ([For help installing LifeKeeper for Linux, click here](#))
2. From a command prompt, **Run**:

```
rpm -qa | grep steel
```

**Note:** This will show you all of the LK rpm packages that are installed (as seen below).

3. **Once you have verified that the IkMAN package is installed**, make sure the manpage rpms are installed on the OS in use.
4. **Run** the following commands to set the MANPATH:

```
[ec2-user@uselpdxcs01 ~]$ MANPATH=/opt/Lifekeeper/man  
[ec2-user@uselpdxcs01 ~]$ export MANPATH  
[ec2-user@uselpdxcs01 ~]$ echo $MANPATH  
/opt/Lifekeeper/man
```

 **Note:** The `MANPATH` may need to be set after rebooting each time.

You should now be able to access the man pages.

### Which commands have manuals?

Manuals for certain commands can be found in `/opt/LifeKeeper/bin` as seen below:

```
[root@use1pdxcs01 bin]# ls
a2p                lcdmachfail       openssl
app_create         lcdrcp            pcregrep
app_list          lcdrecover        pcretest
app_remove         lcdremexec        perform_action
backupadm          lcdreserve        perl
c2ph              lcdrmipc          perl5.8.8
can_talk           lcdstatus         perlbug
certtool          lcdsync           perlcc
checkrklic        lcdunname         perldoc
cpan              lcdwait           perlivp
create_terminal_leaf lcm                piconv
credstore          libnetcfg         pl2pm
c_rehash          lkbackup          pod2html
cre_mbox           lkcheck           pod2latex
curl              lk_chg_value      pod2man
delallsys         lkchkconf         pod2text
dep_create        lkcli             pod2usage
dep_list          lk_confignotifyalias podchecker
dep_remove        lk_configsntp     podselect
dprofpp           lk_confirmso      portio
empty_mbox        lkcore_parameter prove
enc2xs            lkdefault         psed
eqv_create        lkdiskmon         psktool
eqv_list          lkdisktest        pstruct
eqv_remove        lkexterrlog       qwk_storage_exit
eventslcm         lkGUIapp          qwk_storage_init
find2perl         lkGUIserver       rcv
flg_create        lkID              rcv_msg
```

## Command for accessing man pages

To view a man page for the above commands run the following:

- `man <command>`

**Example:**

```
# man dep_list
```

## Troubleshooting

- ! If you are getting the error below and have verified the previous steps 1-3, set the `MANPATH` before access to man pages can be granted (continue to step 4).

```
[ec2-user@use1pdxcs01 ~]$ man dep_list
No manual entry for dep_list
```

## 5.6.2. LKCLI (LifeKeeper Command Line Interface)

---

LKCLI provides functions that can be performed with the LifeKeeper GUI through the command line interface. LifeKeeper provides export/import of communication paths and resource information which the GUI does not provide. Export/import functionality enables duplicating a created system and easy deployment from a testing environment to a production environment.

### Supported Environments

LKCLI is supported in the following environments:

- Number of nodes – Only a two-node cluster environment is supported.
- OS – All operating systems supported by LifeKeeper are supported.
- Communication path – Supported only in environments where nodes are connected via TCP/IP.
- Application Recovery Kits – Only environments configured with the Application Recovery Kits in the “[Supported ARK List](#)” are supported. If unsupported ARK resources have been created on a node, the environment is not supported. If you want to perform command line operations in an unsupported environment, please consider using the [Command Line Interface](#).

### How to Run as User Other Than Root

To run LKCLI, users must be a member of one of the following three groups: lkadmin, lkoper or lkguest. Please refer to [Configuring GUI Users](#) for more information. When running LKCLI as a user other than root, run using a sudo command. The LKCLI command must be specified with an absolute path. When running the command, authentication of the user who runs the sudo command is performed. The following is an example of LifeKeeper startup.

```
$ sudo /opt/LifeKeeper/bin/lkcli start
```

LifeKeeper sudo settings can be performed in */etc/sudoers.d/lifekeeper*. If LifeKeeper settings are not included in */etc/sudoers*, include “`#includedir /etc/sudoers.d`” or “`#include /etc/sudoers.d/lifekeeper`” in */etc/sudoers*. You can also specify paths by configuring the `secure_path` option in */etc/sudoers* and modify the authentication settings. Refer to `man sudoers(5)` for more information.

### Restriction

- When the hostname is not configured properly, the GUI cannot be started and the GUI login fails. At this point, some of the lkcli commands can be used but many of the processes of the lkcli commands will stop. Ensure that the hostname is configured properly.

### Commands

<a href="#">lkcli license</a>	Install a license key
-------------------------------	-----------------------

<a href="#">lkcli start</a>	Starts LifeKeeper
<a href="#">lkcli stop</a>	Stops LifeKeeper
<a href="#">lkcli import</a>	Creates communication paths, resources
<a href="#">lkcli export</a>	Exports communication paths, resources information
<a href="#">lkcli clean</a>	Deletes communication paths, resources
<a href="#">lkcli commpath</a>	Operates on communication paths
<a href="#">lkcli dependency</a>	Operates on dependencies
<a href="#">lkcli resource</a>	Operates resources
<a href="#">lkcli status</a>	Lists resource status
<a href="#">lkcli log</a>	Displays a LifeKeeper log
<a href="#">lkcli server</a>	Configures/operates servers
<a href="#">lkcli mirror</a>	Operates DataKeeper mirroring
<a href="#">lkcli esxi</a>	Setting up the ESXi host used by the VMDK Recovery Kit

### Common Options Available for All Commands

Option	Default	Description
[-- remote <str>]		Hostname of the machine where you want to run a command. If the option is not specified, the command is executed on the local machine. Before you can execute commands remotely, the machine on which you want to execute the command and the communication path must be bidirectionally connected. <b>Note:</b> The <code>start</code> , <code>stop</code> , <code>commpath</code> and <code>clean</code> commands are not supported with this option.

 **Note:** The brackets “[ ]” in the options table indicate that you don’t have to use this option.

### lkcli license

Registers your LifeKeeper license.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--file <str>		A path to the license file

## lkcli start

<b>Permission</b>	lkadmin
-------------------	---------

Starts LifeKeeper

## lkcli stop

Stops LifeKeeper

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
[-f]		Stops only the LifeKeeper daemon and does not stop protected services.
[-r]		Stops LifeKeeper without changing the settings for automatic startup of the LifeKeeper daemon.
[-n]		Performs failover when LifeKeeper stops. This option cannot be used with -f, -i, or -r.
[-i]		Stops only the LifeKeeper daemon and does not stop protected services. The user is prompted to confirm (yes/no) that they want to continue.

## lkcli import

Reads the LifeKeeper settings from a file and creates communication paths and resources.

### lkcli import commpath

Reads LifeKeeper settings from a file and creates communication paths. In order to connect a communication path bidirectionally, execute this command on both the local machine and the remote machine.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--file <str>		A LifeKeeper configuration file path. Create a communication path from a file (YAML format) where the output was saved with <code>lkcli export</code> . The communication path protocol that can be created is TCP/IP (socket).

### lkcli import resource

Reads LifeKeeper settings from a file and creates resources.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--file <str>		A LifeKeeper configuration file path. Create a communication path from a file (YAML format) where the output was saved with <code>lkcli export</code> . Refer to the <a href="#">ARKs list</a> for files that can be created.

**Notes:**

- This command fails if the environment including the hostname, IP address or application is not prepared.
- Rollback is not performed even if it fails. Only some resources may be created.
- Execution of the command where resources already exist is not supported.

### lkcli export

Exports LifeKeeper settings.

The current LifeKeeper settings are exported. The node where the command is executed and all nodes to which the communication path is connected are targeted.

Save the output in a file in YAML format.

```
# lkcli export > lk_export.yml
```

**Notes:**

- The communication path protocol that can be export is TCP/IP (socket).
- Refer to the [ARKs list](#) for resources that can be exported.
- The resource status (In Service, Out of Service, etc.) is not exported.
- The server property value is not exported.
- Only the configured values on LifeKeeper are exported. The settings of the protected applications are not exported.
- For the exported configuration file, edit only the hostname and IP address manually. Manual changes to other items are not supported.

<b>Permission</b>	lkadmin
-------------------	---------

### lkcli clean

Deletes LifeKeeper settings.

Delete the communication path and resource settings of the node where the command was executed. If you want to delete LifeKeeper configurations on all nodes, execute this command on all nodes. Please note that if the communication path is deleted at this time, commands cannot be executed remotely.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--mode <str>		Specify "all" or "resource". <ul style="list-style-type: none"> <li>• <b>all</b> – Delete all communication paths and all resources.</li> <li>• <b>resource</b> – Delete all resources.</li> </ul>

## Ikcli commpath

Operates communication paths.

### Ikcli commpath create

Creates a communication path for the node where the command is executed. The created communication path is the path from the local machine to the remote machine. Execute this command on both the local machine and the remote machine to connect a communication path in both directions. The communication path protocol that can be created is TCP/IP (socket).

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--laddr <str>		IP address on the local machine to be set for the communication path.
--raddr <str>		IP address on the remote machine to be set for the communication path.
--dest <str>		Hostname of the remote machine to be set for the communication path.
[--priority <str>]	Maximum value of existing paths +1	Priority of the communication path.  <b>Note:</b> Specify the same value on both the local machine and the remote machine.

### Ikcli commpath delete

Deletes the communication path of the node where the command was executed.

The communication path to be deleted is the path from the local machine to the remote machine. Execute this command on both the local machine and the remote machine to delete communication paths in both directions,

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--------	---------	-------------

--laddr <str>		IP address on the local machine to be set for the communication path.
--raddr <str>		IP address on the remote machine to be set for the communication path.
--dest <str>		Hostname of the remote machine to be set for the communication path.

## Ikcli dependency

Creates/deletes dependencies for LifeKeeper resources.

### Ikcli dependency create

Creates a new dependency between two resources.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--parent <str>		The tag name of the parent resource.
--child <str>		The tag name of the child resource.

### Ikcli dependency delete

Deletes the dependency between the two resources.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--parent <str>		The tag name of the resource that is a parent of the dependency you want to delete.
--child <str>		The tag name of the resource that is a child of the dependency you want to delete.

## Ikcli resource

Operates on the LifeKeeper resources.

### Ikcli resource create

Creates resources.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--tag <str>		Tag name of the resource to create.

		Existing tag names cannot be created. See <a href="#">Resource Tag Name Restrictions</a> for more details.
<code>--switchback &lt;str&gt;</code>	INTELLIGENT	Switchback settings for the resource. Specify "INTELLIGENT" or "AUTOMATIC".
OPTIONS FOR EACH ARK		Options vary by ARK. See <a href="#">Subcommands for Each ARK</a> for options for each ARK.

**ikcli resource extend**

Extends resources.

 **Note:** Only the target resource is extended even if there are dependencies. An Extend must be done for each resource.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
<code>--tag &lt;str&gt;</code>		Tag name of the resource to extend. Tag names that exist on the extension target cannot be extended. See <a href="#">Resource Tag Name Restrictions</a> for more details.
<code>--dest &lt;str&gt;</code>		The hostname of the target server where the resource hierarchy is extended.
<code>[--switchback &lt;str&gt;]</code>	INTELLIGENT	Switchback settings for the resource. Specify "INTELLIGENT" or "AUTOMATIC".
<code>[--template_priority &lt;num&gt;</code>	1	The priority of the resource hierarchy from which to extend. For the priority, unused values from 1 to 999 are valid, with lower numbers having higher priority.
<code>[--target_priority &lt;num&gt;</code>	10	The priority of the extension target resource hierarchy. For the priority, unused values from 1 to 999 are valid, with lower numbers having higher priority.
OPTIONS FOR EACH ARK		Options vary by ARK See <a href="#">Subcommands for Each ARK</a> for options for each ARK.

**ikcli resource config**

Changes resource settings.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--------	---------	-------------

--tag <str>		The tag name of the resource to change.
OPTIONS FOR EACH ARK		Options vary by ARK See <a href="#">Subcommands for Each ARK</a> for options for each ARK.

### ikcli resource unextend

Unextends the resource.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--tag <str>		Tag name of the resource to unextend.
--dest <str>		The hostname of the target server where you want to unextend the resource hierarchy.

### ikcli resource delete

Deletes a resource.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--tag <str>		Tag name of the resource to delete.

### ikcli resource restore

Brings the resource hierarchy on the active node in service.

<b>Permission</b>	lkadmin, lkoper
-------------------	-----------------

Option	Default	Description
--tag <str>		The tag name of the resource to bring in service.

### ikcli resource remove

Takes the resource hierarchy on the active node out of service.

<b>Permission</b>	lkadmin, lkoper
-------------------	-----------------

Option	Default	Description
--tag <str>		The tag name of the resource to take out of service.

**Ikcli resource info**

Outputs the resource property information. The output differs for each ARK.

<b>Permission</b>	lkadmin, lkoper, lkguest
-------------------	--------------------------

Option	Default	Description
--tag <str>		Tag name of the resource for which the property information is output.

**Ikcli resource eqv**

Outputs the equivalency information of the resource.

<b>Permission</b>	lkadmin, lkoper, lkguest
-------------------	--------------------------

Option	Default	Description
--tag <str>		Tag name of the resource that for which the equivalency information is output.

**Ikcli resource reorder-priority**

Changes the priority of a resource on an active node.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--tag <str>		The tag name of the resource where the priority is to be changed.
--priority <num>		Priority after the change. For the priority, unused values from 1 to 999 are valid, with lower numbers having higher priority.

**Ikcli resource switchback**

Change the switchback settings of the resource on the active node.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--tag <str>		The tag name of the resource to change.
--switchback <str>	INTELLIGENT	Switchback settings for the resource. Specify "INTELLIGENT" or "AUTOMATIC".

## Ikcli status

Displays LifeKeeper status.

See [Detailed Status Display](#) for the output status information.

<b>Permission</b>	lkadmin, lkoper, lkguest
-------------------	--------------------------

Option	Default	Description
[-q]		Outputs the information of each resource in the local system in a short report.
[-e]		Outputs the information of each resource in the local system in a short report and lists the backup system (with the next highest priority).
[-u]		Suppresses duplicate resource entries in the output of the command with -q or -e options.
[-r <str>]		Specifies the resource root tag. Restrict the report to specific resource root tags.

## Ikcli log

Displays LifeKeeper logs.

<b>Permission</b>	lkadmin, lkoper, lkguest
-------------------	--------------------------

Option	Default	Description
--lines <num>	10	Number of log lines to display.

## Ikcli server

Performs operations related to the LifeKeeper server.

### Ikcli server info

Checks the server shutdown method and failover functionality.

<b>Permission</b>	lkadmin, lkoper, lkguest
-------------------	--------------------------

### Ikcli server shutdown-strategy

Configures the server shutdown method.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--mode <str>		The value of <b>switchover</b> or <b>do_not_switchover</b> . <ul style="list-style-type: none"> <li>• <b>switchover</b> – LifeKeeper starts the backup server resources with a graceful shutdown.</li> <li>• <b>do_not_switchover</b> – LifeKeeper does not start backup server resources upon a graceful shutdown.</li> </ul>

### lkcli server confirmso

Configures whether to confirm the user for switching to the backup node when a failover occurs due to a node failure on the LifeKeeper cluster.

See Confirm Failover section in [Confirm Failover and Block Resource Failover](#) for details.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--server <str>		Name of the failover target server.
--value <str>		<b>enable</b> or <b>disable</b> <ul style="list-style-type: none"> <li>• <b>enable</b> – Enable user confirmation at failover.</li> <li>• <b>disable</b> – Do not confirm the user at failover.</li> </ul>

### lkcli server block-failover

Configures settings to block failover caused by a resource failure in the specified system.

See the Block Resource Failover section in [Confirm Failover and Block Resource Failover](#) for details.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--server <str>		Name of the failover target server.
--value <str>		<b>enable</b> or <b>disable</b> <ul style="list-style-type: none"> <li>• <b>enable</b> – Block failover to the specified server.</li> <li>• <b>disable</b> – Do not block failover to the specified server.</li> </ul>

### lkcli server block-all-failovers

Configures settings to block failover due to either a node failure or a resource failure on all servers. When you deactivate the block setting, review the following settings.

- Failover block setting due to a node failure.
  - [lkcli server confirmso](#)
  - [CONFIRMSODEF](#) parameter in the /etc/default/LifeKeeper

- [CONFIRMSOTO](#) parameter in the /etc/default/LifeKeeper
- Failover block setting due to a resource failure.
  - [lkcli server block-failover](#)

## lkcli mirror

Performs mirroring with DataKeeper.

See [Mirroring with SIOS DataKeeper for Linux](#) for more information.

 **Note:** `lkcli mirror` command requires DataKeeper to be installed.

Option	Default	Description
<code>--tag &lt;str&gt;</code>		Tag name of the DataKeeper resource.

### lkcli mirror status

Displays the status of the mirror.

<b>Permission</b>	lkadmin, lkoper, lkguest
-------------------	--------------------------

### lkcli mirror resume

Resumes the mirror.

<b>Permission</b>	lkadmin
-------------------	---------

### lkcli mirror pause

Pauses the mirror.

<b>Permission</b>	lkadmin
-------------------	---------

### lkcli mirror fullresync

Resyncs the entire mirror with a full disk resynchronization.

<b>Permission</b>	lkadmin
-------------------	---------

### lkcli mirror force

Forces the mirror to come online, even if LifeKeeper has marked the mirror disk as possibly out of sync.

 **Note:** Forcing a mirror online should be done with great caution, since this can cause data loss.

<b>Permission</b>	lkadmin
-------------------	---------

## Ikcli esxi

Refer to the [VMDK Shared Storage Recovery Kit Administration Guide](#) for more information about the VMDK Recovery Kit. Refer to [Register ESXi Host](#) and [VMDK Maintenance](#) for more information on ESXi host settings.

### Ikcli esxi add

Adds the ESXi host information. The username and password should be entered interactively if they are not specified as options.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--host <host>		ESXi host name to add.
[--user <user>]		Username to log in to the ESXi host.
[--password <password>]		Password to log in to the ESXi host.

### Ikcli esxi list

Displays the list of ESXi hosts that have been registered.

<b>Permission</b>	lkadmin, lkoper, lkguest
-------------------	--------------------------

### Ikcli esxi update

Updates the ESXi host information that has been registered. The username and password should be entered interactively if they are not specified as options.

<b>Permission</b>	lkadmin
-------------------	---------

Option	Default	Description
--host <host>		ESXi host name to update.
[--user <user>]		Username to log in to the ESXi host.
[--password <password>]		Password to log in to the ESXi host.

### Ikcli esxi delete

Deletes the ESXi host information that has been registered.

<b>Permission</b>	lkadmin
-------------------	---------

<b>Option</b>	<b>Default</b>	<b>Description</b>
--host <host>		ESXi host name to delete.

## 5.6.2.1. LKCLI Subcommands for Each ARK

### ARK List

- [Apache](#)
- [DataKeeper](#)
- [DB2](#)
- [EC2](#)
- [FileSystem](#) ( Multipath Recovery Kits – DMMP / HDLM )
- [Generic Application](#)
- [IP](#)
- [MySQL](#)
- [NFS](#)
- [Recovery Kit for Oracle Cloud Infrastructure](#)
- [Oracle](#)
- [OracleListener](#)
- [OraclePDB](#)
- [PostgreSQL](#)
- [Postfix](#)
- [QSP](#)
- [Raw](#)
- [Route53](#)
- [Samba](#)
- [SAP](#)
- [SAP HANA](#)
- [SAP MaxDB](#)
- [Sybase](#)
- [WebSphere MQ](#)

### Apache

See the Configuring LifeKeeper section of the [Apache Recovery Kit Administration Guide](#).

#### create apache

Option	Default	Description
--root <str>		Full path of Apache Web Server root directory. Relative paths and symbolic links cannot be used.
--path <str>		Full path name (including the file name) of Apache Web Server daemon.

#### extend apache

No options.

## config apache

No options.

## DataKeeper

See the Configuring Resources section of [SIOS DataKeeper for Linux](#) for more information.

### create dk

Option	Default	Description
--mode <str>		Replication type “synchronous” or “asynchronous”. Specify “synchronous” to limit replication to all targets to be synchronous.
[--bitmap <str>]	/opt/ LifeKeeper/ bitmap_<tag name>	Path of bitmap file used for an intent log. The bitmap file is required for asynchronous replication and recommended for synchronous replication. Refer to <a href="#">Creating a DataKeeper Resource Hierarchy</a> for more information.
-- hierarchy <str>		The type of data replication to create.  Options vary depending on the type.  < new   existing   dronly >

### --hierarchy new

Option	Default	Description
--device <str>		Source disk or partition.
--fstype <str>		File system type. Only file system types supported by LifeKeeper can be specified.
--mount_point <str>		New mount point for new file system.
--fstag <str>		Tag name of the file system resource.

### --hierarchy existing

Option	Default	Description
--mount_point <str>		Mount point to mount on the primary server's NetRAID device.
--fstag <str>		Tag name of the file system resource.

### --hierarchy dronly

Option	Default	Description
--------	---------	-------------

--device <str>		Source disk or partition.
----------------	--	---------------------------

### extend dk

Option	Default	Description
--mode <str>		Replication type. Specify “synchronous” or “asynchronous”.
--laddr		Local IP address.
--raddr		Remote IP address.
[--bitmap <str>]	/opt/LifeKeeper/bitmap_< Tag name >	Path of bitmap file used for an intent log.
[--device <str>]	Same as the extension source.	Target disk or partition.
[--fstag <str>]		Tag name of the file system resource.

### config dk

Option	Default	Description
--resync_speed_min <num>		Set the minimum resync speed limit (KB/s).
--resync_speed_max <num>		Set the maximum resynch speed limit (KB/s.)
--compression_level <num>		Set the network compression level (0-9).

## DB2

See [LifeKeeper for Linux DB2 Recovery Kit Configuration Tasks](#) for more details.

### create db2

Option	Default	Description
--instance <str>		DB2 instance Specify the DB2 instance to protect.

### extend db2

No options.

### config db2

Nothing to configure.

## EC2

See the Configuration section of the [Recovery Kit for EC2 Administration Guide](#) for more details.

### create ec2

Option	Default	Description
--type <str>		Specify the type of EC2 resource to be create. Specify "RouteTable" to select a route table scenario, "Elastic IP" to select an Elastic IP scenario.

#### --type RouteTable

Option	Default	Description
--ip_resource <str>		Specify the tag name of the IP resource created in advance.

#### --type ElasticIP

Option	Default	Description
--eip <str>		The IP address of the Elastic IP you want to protect.
--dev <str>		Network interface name to which EIP is attached.

### extend ec2

No options.

### config ec2

No options.

## FileSystem

See [Creating a File System Resource Hierarchy](#) or [Extending a File System Resource Hierarchy](#) for more details.

### create fs

Option	Default	Description
--mountpoint <str>		Specify the mount point of the file system.

### extend fs

Option	Default	Description
--------	---------	-------------

[--mountpoint <str>]	Source system mount point.	Specify the mount point of the file system.
----------------------	----------------------------	---------------------------------------------

## config fs

Option	Default	Description
--mountopt <str>		Specify the mount options for the file system.

## Generic Application

See Creating a [Generic Application Resource Hierarchy](#) or [Extending a Generic Application Resource Hierarchy](#) for more details.

## create gen

Option	Default	Description
--restore <str>		Specify the path of the restore script.
--remove <str>		Specify the path of the remove script.
[--quickCheck <str>]		Specify the path of the quickCheck script.
[--recover <str>]		Specify the path of the recover script.
[--appinfo <str>]		Specify optional information about the application.

## extend gen

Option	Default	Description
[--appinfo <str>]	Source system appinfo.	Specify optional information about the application.

## config gen

Option	Default	Description
[--restore <str>]		Specify the path of the restore script to be updated.
[--remove <str>]		Specify the path of the remove script to be updated.
[--quickCheck <str>]		Specify the path of the quickCheck script to be updated.
[--recover <str>]		Specify the path of the recover script to be updated.
[--all <str>]	No	Specify <b>Yes</b> or <b>No</b> <ul style="list-style-type: none"> <li><b>Yes</b> – Update scripts on all of the cluster nodes.</li> <li><b>No</b> – Update the script on the node where the command is executed.</li> </ul>

## IP

See the Configuration section of the [IP Recovery Kit Administration Guide](#) for more details.

### create ip

Option	Default	Description
--ipaddr <str>		Virtual IP address.  If the actual IP address is to be protected, specify '0.0.0.0'.
[--netmask <str>]	An appropriate value determined from ipaddr.	Virtual IP netmask.  If '0.0.0.0' is specified for ipaddr, the network interface netmask specified for the device will be used.
[--device <str>]	An appropriate value determined from ipaddr and netmask.  <b>Note:</b> If '0.0.0.0' is specified for ipaddr, this device must be specified.	Network interface name associated with the virtual IP or the actual IP.

### extend ip

Option	Default	Description
[--ipaddr <str>]	Source system ipaddr.	Virtual IP address on the extension destination node.  If the actual IP address is to be protected, specify '0.0.0.0'.
[--netmask <str>]	An appropriate value determined from ipaddr.	Virtual IP netmask on the extension destination node.  If '0.0.0.0' is specified for ipaddr, the network interface netmask specified for the device will be used.
[--device <str>]	An appropriate value determined from ipaddr and netmask.  <b>Note:</b> If '0.0.0.0' is specified for ipaddr, this device must be specified.	Network interface name associated with the virtual IP or the actual IP on the extension destination node.

### config ip

Option	Default	Description
[--pinglist <str>]		Ping the destination list for options (multiple designations are specified separated

		by comma).
[--srcaddr <str>]		Specify <b>0</b> or <b>1</b> Specify whether to use the virtual IP address as the source address for external communication IP traffic to the same subnet. <ul style="list-style-type: none"> <li>• <b>0</b> – Do not use</li> <li>• <b>1</b> – Use</li> </ul>
[--restoremode <str>]		Specify <b>Enabled</b> or <b>Disabled</b> . Enable/disable restoration and recovery for IP resources. <ul style="list-style-type: none"> <li>• <b>Enabled</b> – Enable restoration and recovery.</li> <li>• <b>Disabled</b> – Disable restoration and recovery.</li> </ul>

## MySQL

See the Installation section of the [MySQL Recovery Kit Administration Guide](#) for more details.

### create mysql

Option	Default	Description
--cnf <str>		Absolute path of the MySQL configuration file.
--bin <str>		Absolute path of the directory where the MySQL executable binary is located.
[--instance <str>]	None	MySQL instance number you want to protect. If you are using MySQL on a single instance, do not specify this number.

### extend mysql

Option	Default	Description
[--bin <str>]		Absolute path of the directory where the MySQL executable binary is located on the node to which the node is extended.

### config mysql

No options.

## NFS

See the Configuration section of the [NFS Recovery Kit Administration Guide](#) for more details.

### create nfs

Option	Default	Description
--export <str>		Export point for the NFS file system.

<code>--ip &lt;str&gt;</code>		Tag name of the IP resource corresponding to the virtual IP address used by the client to access the NFS file system.
-------------------------------	--	-----------------------------------------------------------------------------------------------------------------------

## extend nfs

No options.

## config nfs

No options.

# Recovery Kit for Oracle Cloud Infrastructure

The resource types are: `app: comm`, `typ: ocivip`.

See the Configuring LifeKeeper section of the [Recovery Kit for Oracle Cloud Infrastructure Administration Guide](#).

## create ocivip

Option	Default	Description
<code>--ipaddr &lt;str&gt;</code>	None	Specify the secondary private IP address. The IP address specified here is assigned to the VNIC.
<code>--device &lt;str&gt;</code>	None	Specify the name of the network interface to which the IP address is assigned.

## extend ocivip

<code>--device &lt;str&gt;</code>	Network interface name specified on the source node	Specifies the network interface name to which the IP address is assigned on the extended node.
-----------------------------------	-----------------------------------------------------	------------------------------------------------------------------------------------------------

## config ocivip

No ARK-specific options.

# Oracle

See the Configuring LifeKeeper section of the [Oracle Recovery Kit](#) for more details.

## create oracle

Option	Default	Description
<code>--sid &lt;str&gt;</code>		ORACLE_SID of the database.

[--listener <str>]	None	The tag name of the Oracle Listener resource that is included depending on the Oracle resource.
[--user <str>]	None	Oracle database username.
[--password <str>]	None	Oracle database user password.

## extend oracle

No options.

## config oracle

Option	Default	Description
--user <str>		Oracle database username.
--password <str>		Oracle database user password.
--role <str>		User role. Specify sysdba or sysoper.

## OracleListener

See the Creating a Shared Oracle Listener for Multiple Resources section of the [Oracle Recovery Kit Administration Guide](#) for more details.

## create listener

Option	Default	Description
--exe <str>		Execution path of the Listener.
--config <str>		Path of the execution setting file of the Listener.
-- protection <str>		Protection level of the Listener: <ul style="list-style-type: none"> <li>• <b>Full</b> – Start, stop, monitor and recover</li> <li>• <b>Intermediate</b> – Start, monitor and recover</li> <li>• <b>Minimal</b> – Only start and monitor</li> </ul>
-- recovery <str>		Recovery level of the Listener: <ul style="list-style-type: none"> <li>• <b>Standard</b> – Enable the standard LifeKeeper recovery. When all listeners fail locally, perform failover to a valid backup server if necessary.</li> <li>• <b>Optional</b> – Enable option LifeKeeper recovery. Even when all listeners fail locally, failover to a valid backup server will not be performed.</li> </ul>
[--user <str>]	None	System username. Specify a system user that has permission to start, stop, monitor and recover the

		Listener.
[--listener <str>]	LISTENER	The name of the Oracle Listener to protect.
[--iptag <str>]	None	The tag name of the IP resource that is protected as a dependency on this resource hierarchy.

## extend listener

Option	Default	Description
[--exe <str>]	Source system exe value.	Execution path of the Listener.
[--config <str>]	Source system config value.	Path of the execution setting file of the Listener.

## config listener

Option	Default	Description
--type <str>		The item name to change. Options vary depending on the item. <ul style="list-style-type: none"> <li>• <b>ProtectionLevel</b> – Protection level of the Listener.</li> <li>• <b>RecoveryLevel</b> – Recovery level of the Listener.</li> <li>• <b>Listener</b> – The name of the Oracle Listener to protect.</li> </ul>

### --type ProtectionLevel

Option	Default	Description
--value <str>		Same as the protection option of create.

### --type RecoveryLevel

Option	Default	Description
--value <str>		Same as the recovery option of create.

### --type Listener

Option	Default	Description
--value <str>		Name of the Oracle Listener to protect.
[--iptag <str>]	None	The tag name of the IP resource that is protected as a dependency.

## OraclePDB

See the Configuring a Pluggable Database with Oracle Multitenant section of the [Oracle Recovery Kit Administration Guide](#) for more details.

## create pdb

Option	Default	Description
--sid <str>		ORACLE_SID of the database
--pdb <str>		PDB list to protect (multiple designations are specified separated by comma)

## extend pdb

No options.

## config pdb

Option	Default	Description
--pdb <str>		PDB list to protect (multiple designations are specified separated by comma)

## PostgreSQL

See the Installation section of the [PostgreSQL Recovery Kit](#) for more details.

## create pgsq

Option	Default	Description
--datadir <str>		Absolute path of the directory where the database data is located.
--port <num>		Port number used by PostgreSQL.
--socket <str>		The path of the socket used by PostgreSQL.
--dbuser <str>		Username used by PostgreSQL.
--logfile <str>		Absolute path where the logs are output.
[--exepath <str>]	/usr/bin	Absolute path of the directory where the executable is located.
[--clientexe <str>]	<exepath>/psql	Absolute path of the executable "psql".
[--adminexe <str>]	<exepath>/pg_ctl	Absolute path of the executable "pg_ctl".

## extend pgsq

Option	Default	Description
[--exepath <str>]	Source system exepath.	Absolute path of the directory where the executable file is located on the node to which the resource is extended. If not specified, the setting of the extension origin is inherited.

## config pgsq

Option	Default	Description
[--dbuser <str>]	None	Username used by PostgreSQL.

## Postfix

See [Creating a Postfix Resource Hierarchy](#) and [Extending a Postfix Resource Hierarchy](#) for more information.

## create postfix

Option	Default	Description
[--binary <binary location>]	/usr/sbin	The directory where the postfix command of postfix is located.
[--config <config file location>]	/etc/postfix	The directory where the setting file (main.cf) of postfix is located.

## extend postfix

No options.

## QSP

See the [\(blank\)Quick Service Protection \(QSP\) Recovery Kit](#) for more details.

## create qsp

Option	Default	Description
--service <str>		Name of the service to protect.
[--quickCheck <str>]	enable	Enable/disable the monitoring function. Specify "enable" to enable and "disable" to disable the monitoring.

## extend qsp

No options.

## config qsp

Option	Default	Description
[--quickCheck <str>]	None	Enable/disable monitoring. Specify "enable" or "disable".
[--timeout_restore	None	Number of seconds for restore timeout. When 0 is specified, timeout

<num>]		does not occur.
[--timeout_remove <num>]	None	Number of seconds for remove timeout. When 0 is specified, timeout does not occur.
[--timeout_quickcheck <num>]	None	Number of seconds for quickCheck timeout. When 0 is specified, timeout does not occur.
[--timeout_recover <num>]	None	Number of seconds for recover timeout.

## Raw

See [Creating a Raw Device Resource Hierarchy](#) and [Extending a Raw Device Resource Hierarchy](#) for more details.

### create raw

Option	Default	Description
--partition <str>		Raw device partition.

### extend raw

No options.

## Route53

See the Configuration section of the [Route53 Recovery Kit](#) for more details.

### create route53

Option	Default	Description
--domain <str>		Domain name that exists in the Route53 to protect.
--hostname <str>		Name of the host to protect.
--ip_resource <str>		Tag name of the IP resource created in advance.

### extend route53

No options.

### config route53

No options.

## Samba

See [Creating a Samba Resource Hierarchy](#) or [Extending Your Samba Resource Hierarchy](#) for more details.

### create samba

Option	Default	Description
<code>--config &lt;config file name&gt;</code>	<code>/etc/samba/smb.conf</code>	Absolute path of the Samba configuration file.

### extend samba

No options.

## SAP

See the Configuration section of the [SAP Recovery Kit Administration Guide](#) for more details.

### create sap

Option	Default	Description
<code>--sid &lt;str&gt;</code>		SID for the SAP installation.
<code>--instance &lt;str&gt;</code>		SAP instance (e.g., ASCS10).
<code>[--protection_level &lt;FULL STANDARD BASIC MINIMUM&gt;]</code>	FULL	<a href="#">Protection Level</a> for the protected SAP resource.
<code>[--recovery_level &lt;FULL REMOTE LOCAL OFF&gt;]</code>	FULL	<a href="#">Recovery Level</a> for the protected SAP resource.
<code>[--ip &lt;str&gt;]</code>		LifeKeeper IP resource tag protecting the virtual IP for the given SAP instance. The IP resource must be in-service on the server where the SAP resource is being created. A comma-separated list of IP resource tags is allowed for this parameter. The IP resources specified by this argument will become child resources of the SAP resource.
<code>[--dependent_fs &lt;str&gt;]</code>		A comma-separated list of LifeKeeper resource tag(s) for the file system (gen:filesystem) or NFS (gen:nfs) resource(s) protecting the file system(s) associated to the given SAP instance. (e.g., /sapmnt,/usr/sap/SID/ASCS10). All given resources must be in-service on the server where the SAP resource is being created. The file system and/or NFS resources specified by this argument will become child resources of the SAP

		resource. Note: This parameter cannot be used in conjunction with the <code>--automate_fs_creation</code> parameter.
<code>[--dependent_cs &lt;str&gt;]</code>		LifeKeeper SAP resource tag for a central services resource (e.g., ASCS10) to be added as a child of the SAP resource being created. The given central services resource must be in-service on the server where the SAP resource is being created. Note: This parameter may be applicable when creating a resource to protect a PAS instance. See the “Create the Primary Application Server Resource” section of <a href="#">Creating an SAP Resource Hierarchy</a> for more information.
<code>[--automate_fs_creation]</code>		During SAP resource creation, LifeKeeper will scan the existing SAP directory structures and attempt to automatically create appropriate LifeKeeper resources to protect any shared file system mounts that it finds. Any resources that are created during this process will become children of the SAP resource. Inspect the resulting hierarchy for accuracy once the process completes. Notes: Replicated file system hierarchies cannot be created automatically by this process. This parameter cannot be used in conjunction with the <code>--dependent_fs</code> parameter.

### extend sap

No options.

### config sap

Option	Default	Description
<code>[--protection_level &lt;FULL STANDARD BASIC MINIMUM&gt;]</code>		Update the <a href="#">Protection Level</a> for the protected SAP resource. This command must be run on the server where the SAP resource is currently in-service.
<code>[--recovery_level &lt;FULL REMOTE LOCAL OFF&gt;]</code>		Update the <a href="#">Recovery Level</a> for the protected SAP resource. This command must be run on the server where the SAP resource is currently in-service.
<code>[--sshcc_action &lt;start stop migrate[:&lt;node&gt;]&gt;] maintenance:&lt;enable disable check&gt;&gt;]</code>		Perform an <a href="#">SAP SIOS HA Cluster Connector (SSHCC)</a> action. This command must be run on the server where the SAP resource is currently

	<p>in-service. The possible values of this parameter are:</p> <p>start – Bring the SAP resource in-service via the HA cluster connector.</p> <p>stop – Take the SAP resource out-of-service via the HA cluster connector.</p> <p>migrate[:&lt;node&gt;] – Migrate the SAP resource to a different server. Optionally, a target migration server may be provided by adding “:&lt;hostname&gt;” to the migrate parameter (e.g., <code>lkcli resource config sap --tag SAP-SID_ASCS10 --sshcc_action migrate:node2</code>). If no target migration server is explicitly provided, the HA cluster connector will attempt to determine the best target server based on current cluster conditions.</p> <p>maintenance:&lt;enable disable check&gt; – Enable, disable, or check the status of <a href="#">SAP Maintenance Mode</a> for the hierarchy containing the given SAP resource.</p>
--	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## SAP HANA

See the Configuration section of the [SAP HANA Recovery Kit Administration Guide](#) for more details.

### create hana

Option	Default	Description
<code>--sid &lt;str&gt;</code>		SID for the SAP HANA installation.
<code>--instance &lt;str&gt;</code>		SAP HANA instance (e.g., HDB00).
<code>[--ip &lt;str&gt;]</code>		LifeKeeper IP resource tag protecting the virtual IP for the SAP HANA database. The IP resource must be in-service on the server where the SAP HANA resource is being created.

[--set_local_recovery_policy <enable disable>]	enable	This parameter may be used to enable or disable local recovery for the SAP HANA resource on the server where it is being created.
------------------------------------------------	--------	-----------------------------------------------------------------------------------------------------------------------------------

## extend hana

Option	Default	Description
[--set_local_recovery_policy <enable disable>]	enable	This parameter may be used to enable or disable local recovery for the SAP HANA resource on the server where it is being extended.

## config hana

Option	Default	Description
[--stop_all_dbs]		Stop the SAP HANA database on all cluster nodes (e.g., for maintenance). This command must be run on the cluster node where the HANA resource is currently in-service.
[--takeover_with_handshake <target server>]	Local server	Perform a “takeover with handshake” of the SAP HANA database on the given target server. This is a feature of SAP HANA 2.0 SPS04 and later which reduces downtime of the primary database during switchover by suspending (rather than completely stopping) the primary database before performing a takeover of SAP HANA System Replication on the new database host.
[--set_local_recovery_policy <enable disable>]		This parameter may be used to enable or disable local recovery for the SAP HANA resource on the local server.

## SAP MaxDB

See the Configuration section of the [SAP MaxDB Recovery Kit Administration Guide](#) for more details.

## create sapdb

Option	Default	Description
[--prog <str>]	Path found in /etc/opt/sdb	The SAP MaxDB Program Path. You may type in another directory path. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
[--sid <str>]	The first instance found that is not already	The SAP MaxDB instance ID.

	configured.	
[--user <str>]	User found in /etc/opt/sdb	This is the system user that owns or has permission to execute SAP MaxDB commands. This user must already exist on the template server.
[--xuser <str>]	The first XUSER key configured.	The XUSER key is used to store database user data for use with SAP MaxDB Tools. The XUSER key must already be configured on the corresponding server for the system user and database instance combination.

### extend sapdb

Option	Default	Description
[--prog <str>]	Path found in /etc/opt/sdb on the target.	The SAP MaxDB Program Path. You may type in another directory path. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
[--user <str>]	User defined for this instance on the template server.	This is the system user that owns or has permission to execute SAP MaxDB commands on the target server. This user must already exist on the target server.
[--xuser <str>]	XUSER key for this instance on the template server.	The XUSER key is used to store database user data for use with SAP MaxDB Tools. The XUSER key must already be configured on the target server for the system user and database instance combination.

### config sapdb

No options.

## Sybase

See the Install section of the [Sybase ASE Recovery Kit Administration Guide](#) for more details.

### create sybase

Option	Default	Description
--basedir <str>		Sybase installation directory.
--datadir <str>		Sybase instance data directory.
--instance <str>		Database instance.
--user <str>		Database user. *If not specified, try to access and operate the database as the sa user.
--password <str>		Database user password.
--backup <str>	none	Backup server.

## extend sybase

No options.

## config sybase

Option	Default	Description
--user <str>		Database user.
--password <str>		Database user password.
--backup <str>	Current value	Backup server.
--monitor <str>	Current value	Monitor server.

## WebSphere MQ

See the Configuration section of the [WebSphere MQ Kit Administration Guide](#) for more details.

### create mq

Option	Default	Description
--qmgr <str>		Name of the queue manager.
[--protect_listener <YES NO>]	YES	Whether to protect the MQ listener associated with the queue manager.
[--channel <str>]	SYSTEM.DEF.SVRCONN	Connection channel for the queue manager.
[--ip <str>]		LifeKeeper IP resource tag protecting the listener IP address. The IP resource must be in-service on the server where the MQ resource is being created. The IP resource specified by this argument will become a child resource of the MQ resource.

### extend mq

No options.

### config mq

See [Editing MQ Resource Configuration Properties](#) for more details.

Option	Default	Description
[--protect_listener <YES NO>]		Whether to protect the MQ listener associated with the queue manager. This parameter may

		only be changed on the server where the MQ resource is in-service (ISP).
<code>--test_queue &lt;str&gt;</code>		The name of the test queue for the LifeKeeper MQ Recovery Kit to use when performing PUT/GET health checks. The value may be set to an empty string (i.e., <code>--test_queue ""</code> ) to disable PUT/GET tests.
<code>--log_level &lt;ERROR INFORMATIONAL DEBUG FINE&gt;</code>		The logging level for the MQ resource.
<code>--immediate_stop_timeout &lt;str&gt;</code>		The timeout value (in seconds) to be used when attempting an immediate stop of the queue manager.
<code>--preemptive_stop_timeout &lt;str&gt;</code>		The timeout value (in seconds) to be used when attempting a preemptive stop of the queue manager.
<code>--channel &lt;str&gt;</code>		The connection channel for the queue manager. This parameter may only be changed on the server where the MQ resource is in-service (ISP).
<code>--cmd_server_protection &lt;Full Minimal&gt;</code>		The protection level for the MQ command server.

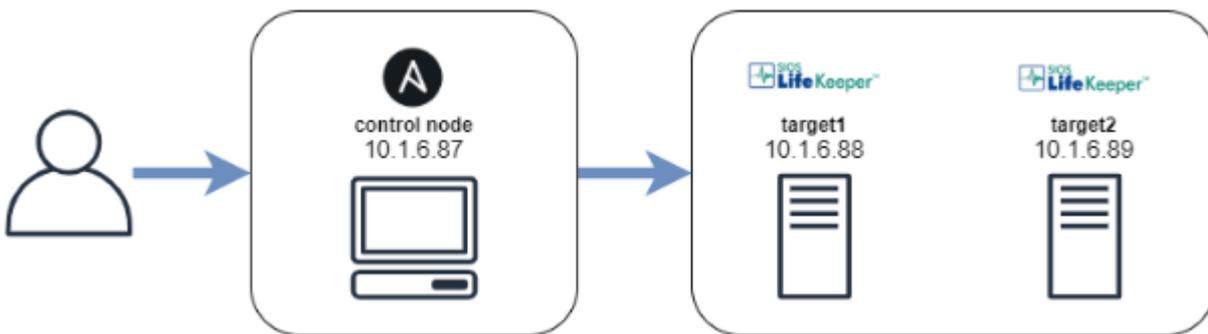
## 5.6.3. Setting up LifeKeeper with Ansible

### Overview

This guide describes the steps to automate the installation of LifeKeeper on multiple nodes with Ansible based on a response file. It is primarily intended to be used as a reference by users who are already using Ansible when adding LifeKeeper into their environment. This guide provides a sample of the required Ansible tasks. The settings should be changed according to your environment.

### Environment

In this document, we are creating a 2-node configuration.



OS: CentOS 8.0

Python: 3.6.8

Ansible: 2.9.9

LifeKeeper: 9.5.1

## Setting up Targets to Install LifeKeeper

### Configuring a Host Name and Name Resolution Service

LifeKeeper must be able to resolve target names in the cluster. Configure the DNS service or `/etc/hosts` so that name resolution can be performed between cluster nodes. In this guide we configure `/etc/hosts` on each node. The settings are changed as shown below.

#### hosts.yml

```

---
# Register all the hosts on which ansible is executed in /etc/hosts
- name: update /etc/hosts
  lineinfile:
    dest: /etc/hosts
    state: present
    insertafter: EOF
  
```

```

  regexp: "^{{ item.address }}"
  line: "{{ item.address }}\t{{ item.hostname }}"
with_items:
  - address: "10.1.6.88"
    hostname: "target1"
  - address: "10.1.6.89"
    hostname: "target2"
  - address: "10.1.6.87"
    hostname: "control-node"

```

## Configuring a Firewall

If a firewall is enabled it is necessary to configure the firewall on each target. LifeKeeper communicates between targets using TCP port 7365. In this guide the port is set to the default: public zone.

### firewall.yml

```

---
- name: enable lifekeeper port
  firewallld:
    zone: public
    port: 7365/tcp
    permanent: yes
    state: enabled
    register: firewallld

- name: reload
  command: firewall-cmd --reload
  when: firewallld.changed

```

## Configuring SELinux

In order to use LifeKeeper, SELinux must be disabled. The settings are changed as shown below.

### selinux.yml

```

---
- name: disable SELinux
  become: yes
  selinux: state=disabled
  register: selinux

- name: reboot a node
  become: yes
  reboot:
  when: selinux.reboot_required

```

## Preparing for LifeKeeper Installation and Licensing

In order to install LifeKeeper non-interactively, a response file is required. If you already have a response file you can skip this step.

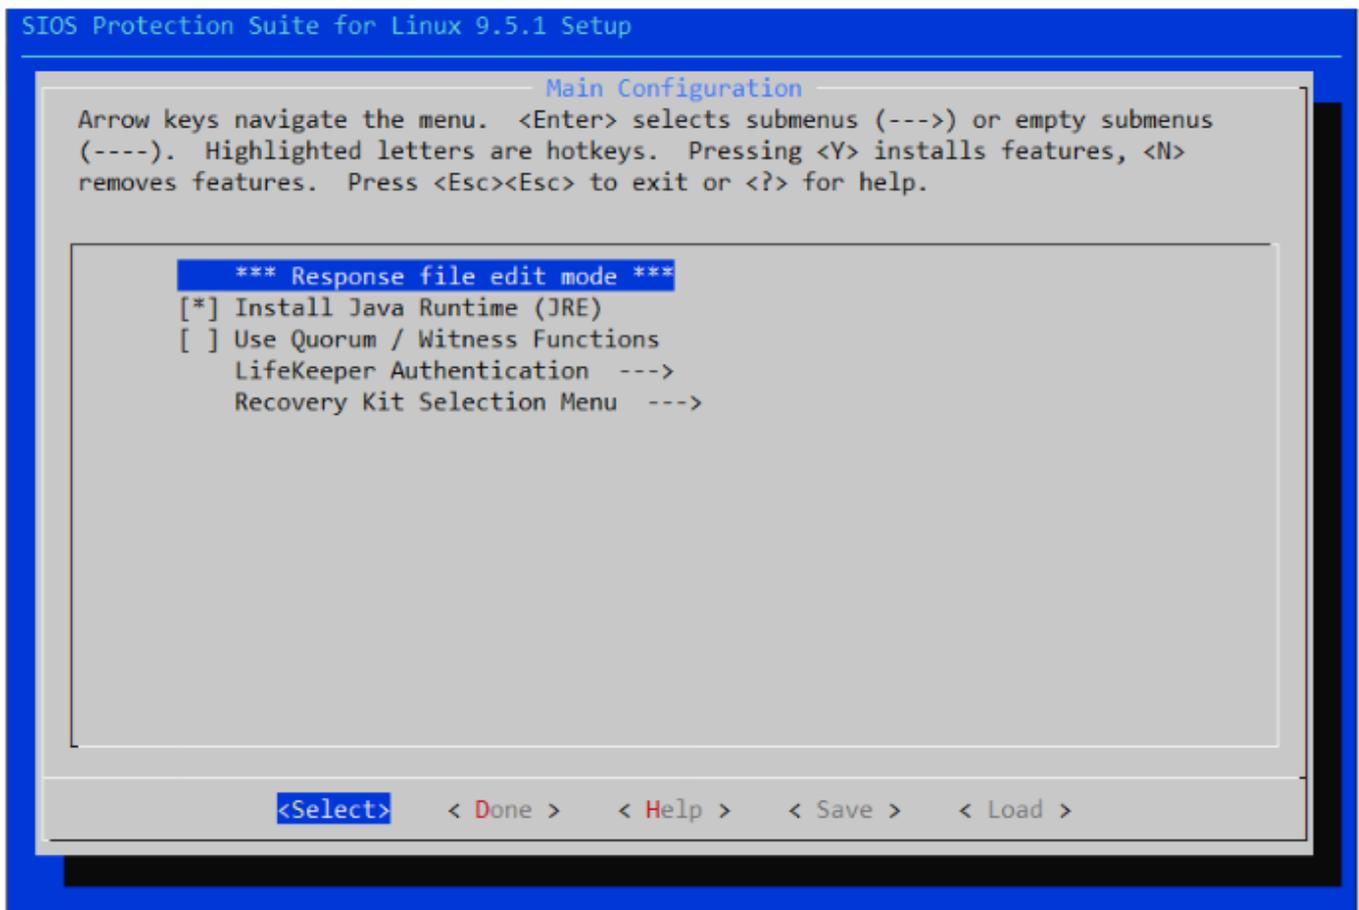
1. Download the create\_response\_file script.

The create\_response\_file script can be obtained from the same FTP directory that contains the LifeKeeper for Linux product image.

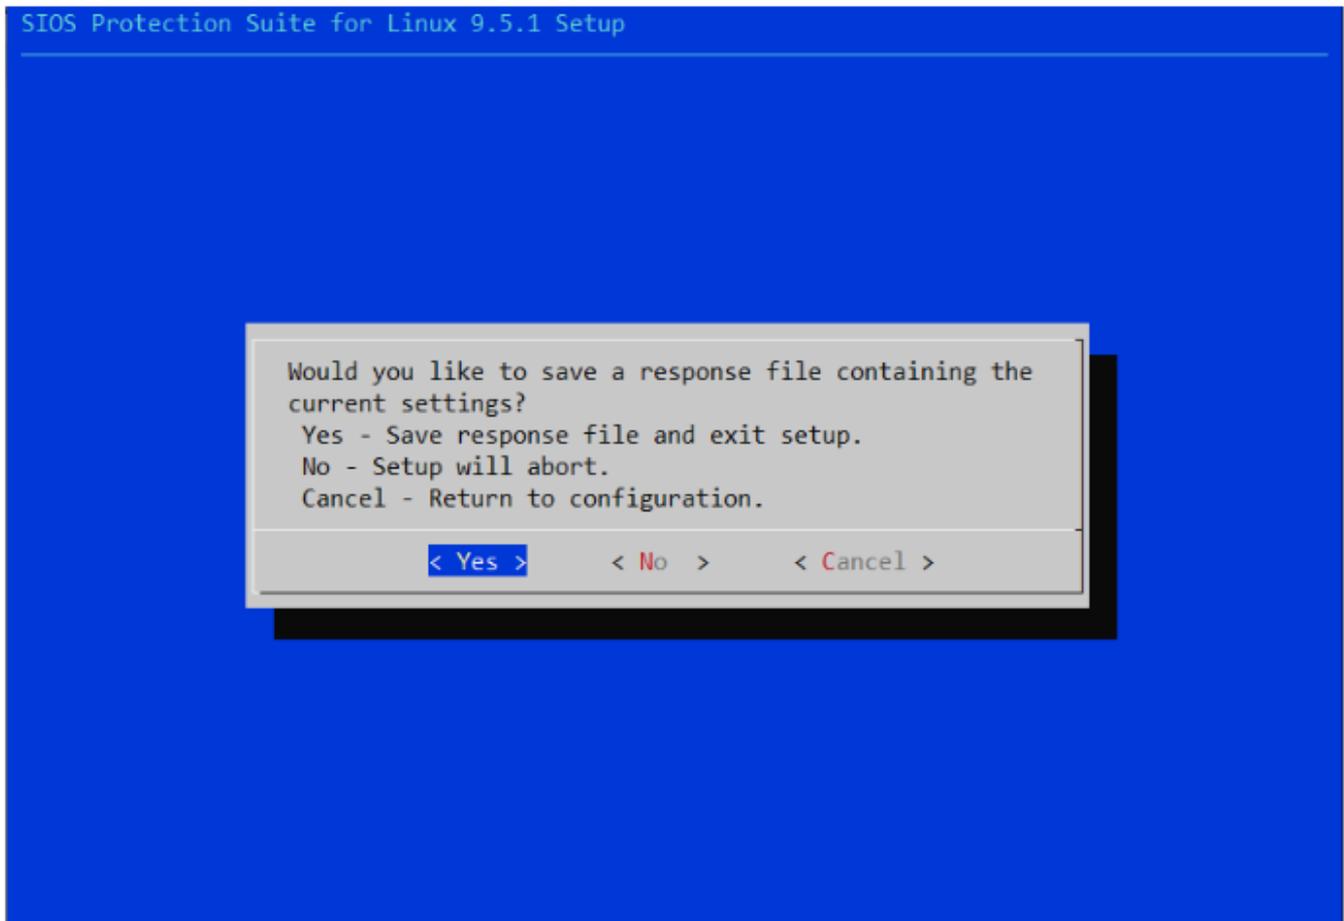
2. Specify the file name and run the create\_response\_file script to create the configuration file.

```
# ./create_response_file /root/LifeKeeper/LKCONF
```

3. When starting the create\_response\_file script the following menu screen is displayed.



4. Select the Recovery Kits you want to install from the Recovery Kit Selection Menu.
5. After setting up, select **Done**. Select **Yes** to exit and save the response file.



6. Inspect the generated LKCONF file to confirm that your setup is correct. The following example shows a sample response file in which the user has selected the Oracle, PostgreSQL, MySQL, DataKeeper for Linux, and Apache Recovery Kits.

```
# LifeKeeper setup response file
# DO NOT EDIT MANUALLY
LKCONF_INSTALL_JRE="y"
LKCONF_SELONLY="y"
LKCONF_AUTH="y"
LKCONF_LKUSER_lkadmin="root"
LKCONF_steeleye_lkORA="y"
LKCONF_steeleye_lkPGSQL="y"
LKCONF_steeleye_lkSQL="y"
LKCONF_steeleye_lkDR="y"
LKCONF_steeleye_lkAPA="y"
```

## Distribute each Obtained File and Specify the Directory

1. Distribute the obtained installation image, licenses for all cluster nodes, and the previously created response file to the control-node machine. The following example shows the directory structure used in this guide.

```

/root/lifekeeper/
|- -sps.img: # LifeKeeper installation image
|- -LKCONF: # Setup file for non-interactive LifeKeeper installation
|- -licenses/
    |- - target1 # Set of license files for target1
    |- - target2 # Set of license files for target2

```

## inventory

```

[target]
target1
target2

```

2. Specify the directories to use.

## defaults/main.yml

```

---
# Directory to save files required for installation on the control-node
control_node_dir: /root/lifekeeper
# Directory to save files that are temporarily required for installation on the
# target
target_dir: /tmp/lifekeeper
# Mount directory of the installation image on the target
target_mnt: /mnt/LifeKeeper

```

The following steps must be performed on the targets.

## Distribute the LifeKeeper Installation Image, License and Configuration File

### deploy.yml

```

---
- name: deploy to lifekeeper install files
  copy:
    src: "{{ control_node_dir }}"
    dest: "{{ target_dir }}"

```

## Mount the LifeKeeper Installation Image

### mount.yml

```

---
- name: mount setup image
  mount:

```

```

path: "{{ target_mnt }}"
src: "{{ target_dir }}/sps.img"
fstype: iso9660
state: mounted

```

## Install LifeKeeper

### install.yml

```

---
- name: install lifekeeper
  shell:
    # If you run setup directly, it will fail.
    cmd: script -q -e -c "{{ target_mnt | quote }}/setup -f {{ target_dir | quote }}/LKCONF -q y" {{ target_dir | quote }}/lifekeeper_install.log

```

## Add LifeKeeper Environment Changes

### bash\_profile.yml

```

---
- name: update lifekeeper PATH in bash_profile
  blockinfile:
    path: /root/.bash_profile
    marker: "# {mark} ANSIBLE MANAGED BLOCK: LifeKeeper"
    block: |
      PATH=$PATH:/opt/LifeKeeper/bin
      MANPATH=$MANPATH:/opt/LifeKeeper/man
      export PATH MANPATH

```

## Installing the LifeKeeper License

### license.yml

```

---
- name: install lifekeeper license
  shell:
    cmd: |
      for file in `ls`; do
        /opt/LifeKeeper/bin/lkcli license --file $file
      done
    chdir: "{{ target_dir }}/licenses/{{ inventory_hostname }}"

```

## Start LifeKeeper

### lkstart.yml

```
---
- name: start lifekeeper
  command: /opt/LifeKeeper/bin/lkcli start

- name: lifekeeper status
  command: /opt/LifeKeeper/bin/lkcli status
```

## Delete Unnecessary Files after Setting up LifeKeeper

### clean.yml

```
---
- name: unmount a mounted lifekeeper image
  mount:
    path: "{{ target_mnt }}"
    state: absent
- name: delete lifekeeper files
  file:
    path: "{{ target_dir }}"
    state: absent
```

## Sample Ansible Playbook

Once all of the task files in the previous sections have been created, they can be combined to create a play similar to the one given in the example playbook shown below.

### deploy\_lifekeeper.yml

```
---
- name: deploy and install lifekeeper on all targets
  hosts: targets
  tasks:
    - include_vars: main.yml
    - include_tasks: hosts.yml
    - include_tasks: firewall.yml
    - include_tasks: selinux.yml
    - include_tasks: deploy.yml
    - include_tasks: mount.yml
    - include_tasks: install.yml
    - include_tasks: bash_profile.yml
    - include_tasks: license.yml
```

```
- include_tasks: lkstart.yml  
- include_tasks: clean.yml
```

This sample playbook can be executed as root with the command:

```
# ansible-playbook deploy_lifekeeper.yml -i inventory
```

or can be incorporated into an existing Ansible playbook.

After completing all tasks described in this guide, Ansible has been successfully configured to automatically install LifeKeeper on all cluster nodes specified in the inventory file.

## 5.6.4. LKCLI Guide

---

### Overview

This guide is designed to assist you with configuring your LifeKeeper environment via the command line using LKCLI (LifeKeeper Command Line Interface). Please follow the order below when creating your environment via LKCLI:

1. [Communication Path Creation and Deletion](#)
2. [Resource Creation](#)
3. [Checking Cluster Status](#)
4. [Verifying Switchover Behavior](#)
5. [Maintenance Tasks](#)
6. [Replicate the Existing Cluster Settings](#)

Please refer to [LKCLI \(LifeKeeper Command Line Interface\)](#) for more details for each LKCLI command.

## 5.6.4.1. LKCLI – Communication Path Creation and Deletion

What is a communication path?

LifeKeeper configures clusters by connecting multiple nodes to each other. Availability can be maintained by switching to resources on another node in the cluster in the event that a node fails. To achieve this, it is necessary to configure paths for communication between nodes in advance. These paths are called “communication paths” in LifeKeeper.

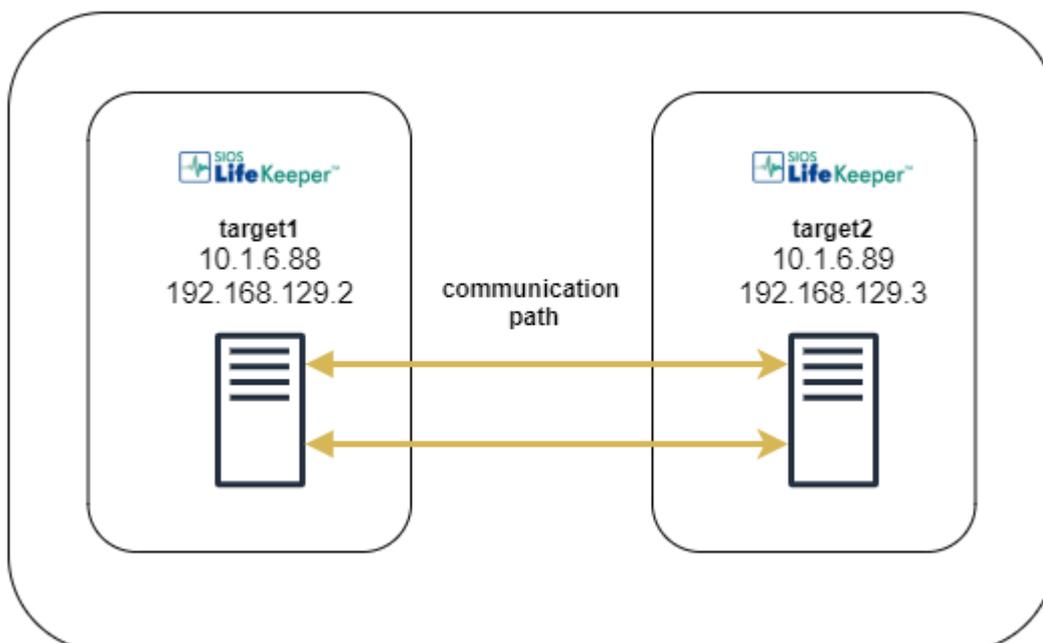
When communication paths are configured, each node sends a signal called a “heartbeat”. The heartbeat indicates that LifeKeeper is working properly. If the heartbeat is not received, the node is considered to have failed. SIOS recommends connecting cluster nodes with two or more communication paths so that in the case of a network failure, the nodes will still be able to receive heartbeats over an alternate path.

Communication paths are used for LifeKeeper internal communication, in addition to heartbeats. A remote LifeKeeper node can be managed over the network as long as communication paths are connected between the nodes.

This guide describes the steps to create and delete a single communication path between two nodes.

### Configuration

The commands and other information in this guide are based on the following diagram.



Before you begin, please check that the following conditions are met in your environment:

- TCP Port 7365 is available for the communication path

- The firewall allows communication or the firewall is disabled
- There are two machines running LifeKeeper

Next, record information about your systems that will be required in order to execute the CLI commands.

#### 1. Available IP address on each system

```
# ip address | grep 'inet'
inet 127.0.0.1/8 scope host lo
inet6 ::1/128 scope host
inet 10.1.6.88/16 brd 10.1.255.255 scope global noprefixroute ens192
inet6 fe80::a633:2758:4976:b42/64 scope link noprefixroute
inet 192.168.129.2/24 brd 192.168.129.255 scope global noprefixroute ens22
4
inet6 fe80::34a4:417f:b58b:dc15/64 scope link noprefixroute
inet 192.168.122.1/24 brd 192.168.122.255 scope global virbr0
```

#### 2. Host name of each system

```
# hostname
target1
```

## Creating a LifeKeeper Communication Path

**Note:** The command for creating a communication path creates the path *in only one direction*. Therefore, in order for the communication path to be able to communicate in both directions, you need to run the command on both systems.

### Perform the following steps on target1

1. Make sure that a communication path is not connected.

```
[target1]# lkcli status -q
LOCAL    TAG    ID    STATE    PRIO    PRIMARY
```

2. Create a communication path from target1 to target2.

```
[target1]# lkcli commpath create --laddr 10.1.6.88 --raddr 10.1.6.89 --dest ta
rget2
Performing commpath 'target2:10.1.6.88/10.1.6.89' create...
Commpath 'target2:10.1.6.88/10.1.6.89' created successful.
```

### Command Argument

Item	Input Value
------	-------------

--laddr	IP address on the local node to connect from
--raddr	IP address on the remote node to connect to
--dest	Host name of the remote node to connect to

3. Verify that a communication path to target2 has been created. (**Note:** At this point, the STATE is DEAD because only a one-way connection is established.)

```
[target1]# lkcli status -q
LOCAL    TAG    ID    STATE    PRIO  PRIMARY

MACHINE  NETWORK ADDRESSES/DEVICE    STATE    PRIO
target2  TCP    10.1.6.88/10.1.6.89  DEAD    1
```

### Perform the following steps on target2

4. Make sure a communication path is not connected.

```
[target2]# lkcli status -q
LOCAL    TAG    ID    STATE    PRIO  PRIMARY
```

5. Create a communication path from target2 to target1.

```
[target2]# lkcli commpath create --laddr 10.1.6.89 --raddr 10.1.6.88 --dest ta
rget1
Performing commpath 'target1:10.1.6.89/10.1.6.88' create...
Commpath 'target1:10.1.6.89/10.1.6.88' created successful.
```

### Command Arguments

Item	Input Value
--laddr	IP address on the local node to connect from
--raddr	IP address on the remote node to connect to
--dest	Host name of the remote node to connect to

6. Verify that the communication path is established. Confirm that the STATE is ALIVE.

\*It may take a few seconds for the communication path to become ALIVE.

```
[target2]# lkcli status -q
LOCAL    TAG    ID    STATE    PRIO  PRIMARY

MACHINE  NETWORK ADDRESSES/DEVICE    STATE    PRIO
```

```
target1 TCP 10.1.6.89/10.1.6.88 ALIVE 1
```

## Deleting a LifeKeeper Communication Path

**Note:** The command for deleting a communication path deletes the path *in only one direction*. Therefore, in order for the communication path to be deleted in both directions, you need to run the command on both systems.

### Perform the following steps on target2

1. Make sure that there is a connected communication path.

```
[target2]# lkcli status -q
LOCAL TAG ID STATE PRIO PRIMARY
MACHINE NETWORK ADDRESSES/DEVICE STATE PRIO
target1 TCP 10.1.6.89/10.1.6.88 ALIVE 1
```

2. Delete the communication path from target2 to target1.

```
[target2]# lkcli comppath delete --laddr 10.1.6.89 --raddr 10.1.6.88 --dest target1
```

### Command Arguments

Item	Input Value
--laddr	IP address on the local node to connect from
--raddr	IP address on the remote node to connect to
--dest	Host name of the remote node to connect to

3. Confirm that the communication path to target1 has been deleted.

```
[target2]# lkcli status -q
LOCAL TAG ID STATE PRIO PRIMARY
```

### Perform the following steps on target1

4. Make sure that a communication path to target2 still exists.

```
[target1]# lkcli status -q
LOCAL TAG ID STATE PRIO PRIMARY
```

MACHINE	NETWORK	ADDRESSES/DEVICE	STATE	PRIO
target2	TCP	10.1.6.88/10.1.6.89	DEAD	1

**5. Delete the communication path to target2.**

```
[target1]# lkcli commpath delete --laddr 10.1.6.88 --raddr 10.1.6.89 --dest target2
```

**6. Confirm that the communication path to target2 has been deleted.**

```
[target1]# lkcli status -q
```

LOCAL	TAG	ID	STATE	PRIO	PRIMARY
-------	-----	----	-------	------	---------

## 5.6.4.2. LKCLI – Resource Creation

---

The following topics describe how to create LifeKeeper resources for protected services and applications.

\* **Note:** Be sure to create a communication path first (refer to [LKCLI Communication Path Creation and Deletion](#)).

- [Creating a File System Resource](#)
- [Creating an IP Resource](#)
- [Creating a PostgreSQL Resource](#)
- [Creating a DataKeeper Resource](#)

\* Refer to [LKCLI \(LifeKeeper Command Line Interface\)](#) for a complete list of LKCLI commands and [Subcommands for Each ARK](#) for a complete list of LKCLI subcommands

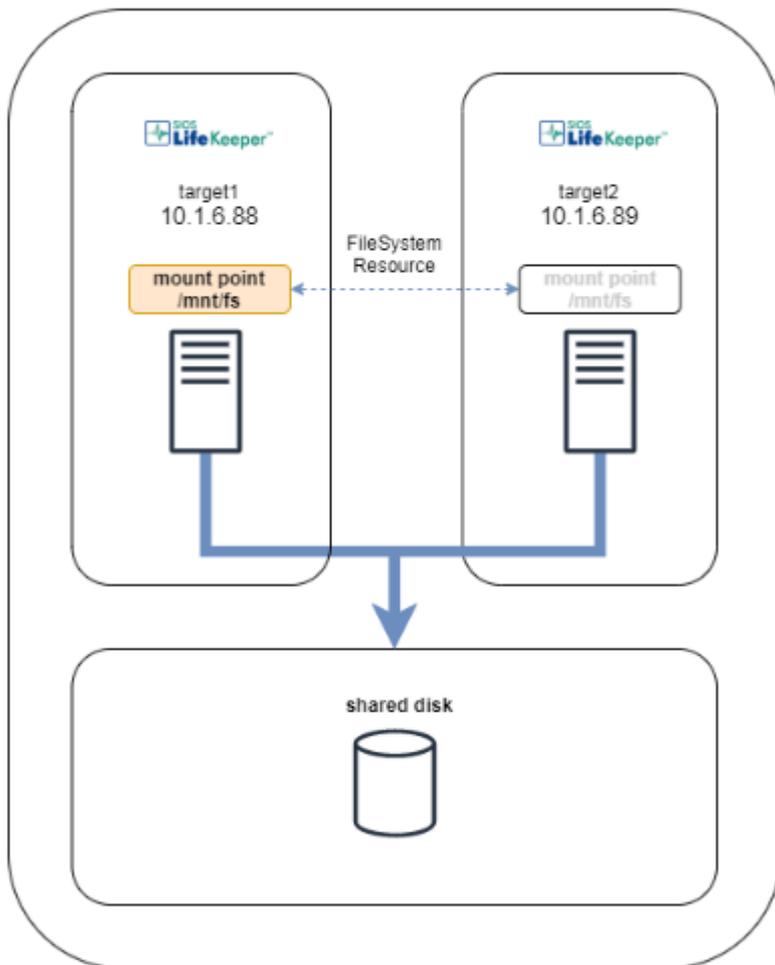
\* Refer to [Application Recovery Kits](#) for the Administration Guides for each Recovery Kit.

## 5.6.4.2.1. Creating a File System Resource

The steps for creating file system resources are described below. File system resources provide the capability to switch file systems on shared storage between cluster nodes. To create a replicated file system hierarchy using SIOS DataKeeper, see [Creating a DataKeeper Resource Hierarchy](#).

### System Configuration

The environment created in this guide is a two-node configuration as shown below.



To create file system resources, the following conditions must be satisfied:

- Shared storage (e.g., iSCSI, Fibre Channel) is physically connected to each node
- A file system has been created using a utility such as mkfs
- The file system can be mounted/unmounted on each node

### Perform the following steps on target1

#### 1. Mounting a file system

Mount a file system for which you want to create a file system resource.

In the example, `/dev/sdb1` is mounted on `/mnt/fs`. The mount point is `/mnt/fs`.

```
[target1]# df
Filesystem                1K-blocks    Used Available Use% Mounted on
/dev/mapper/centos-root  14034944 6904924   7130020   50% /
devtmpfs                  929204      0    929204    0% /dev
tmpfs                     941312      0    941312    0% /dev/shm
tmpfs                     941312    25948   915364    3% /run
tmpfs                     941312      0    941312    0% /sys/fs/cgroup
/dev/sda1                 1038336 148528   889808   15% /boot
tmpfs                    188264      0    188264    0% /run/user/0
/dev/sdb1                 1044132  32992   1011140    4% /mnt/fs
```

2. Creating a resource

Run the following command.

```
[target1]# lkcli resource create fs --tag fs-tag --mountpoint /mnt/fs
```

Resource Settings

Item	Input Value
--tag	Tag name
--mountpoint	Mount point

3. Extending a resource

Run the following command.

```
[target1]# lkcli resource extend fs --tag fs-tag --dest target2
```

Resource Settings

*Item	Input Value
--tag	Tag name of the created resource
--dest	Backup node name

4. Checking the resource

After creating and extending the resource, run the following command.

The resource information is displayed.

```
[target1]# lkcli status -q

LOCAL  TAG          ID              STATE  PRIO  PRI
MARY
target1 fs-tag    /mnt/fs        ISP    1    tar
get1
```

```
target1 device28856 36000c292eb0c693b2efb44ed56556636-1 ISP 1 tar
get1
target1 disk28786 36000c292eb0c693b2efb44ed56556636 ISP 1 tar
get1
```

When you create a file system resource, multiple resources are automatically created with dependencies as shown above.

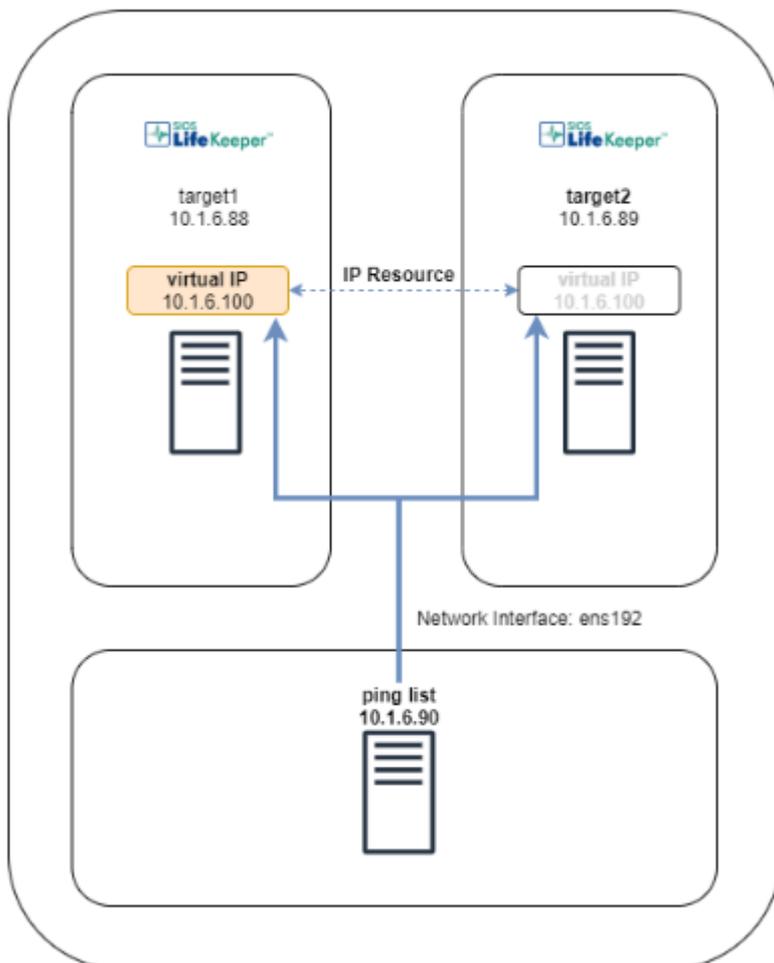
---

## 5.6.4.2.2. Creating an IP Resource

This section describes how to create an IP resource. The IP resource creates and protects a virtual IP address that can be switched between cluster nodes.

### System Configuration

The environment created in this guide is a two-node configuration as shown below.



Prepare a virtual IP address that can be pinged (10.1.6.100 in the above figure).

Also prepare a system that can be pinged (10.1.6.90 in the above).

Check for a ping response using the following command:

```
# ip -4 addr add 10.1.6.100/24 dev ens192  
  
# ping -c3 -I 10.1.6.100 10.1.6.90
```

There should be a ping response.

Once the ping response is verified, remove the IP address from the interface.

```
# ip -4 addr delete 10.1.6.100/24 dev ens192
```

### Restrictions:

- Make sure that the virtual IP address you are trying to create is unique.
- Make sure that there is a system (other than the cluster nodes) that can respond to pings on the same network as the virtual IP address.

**Note:** IP resources use ping to validate the health of the network. Therefore, you need a system outside the cluster that can respond to pings.

### Perform the following steps on target1

1. Creating a resource

Execute the following command:

```
[target1]# lkcli resource create ip --tag ip-tag --ipaddr 10.1.6.100
```

### Resource Settings

Item	Input Value
--tag	Tag name
--ipaddr	Virtual IP address

2. Configuring a ping list

Run the following command to configure a ping list.

```
[target1]# lkcli resource config ip --tag ip-tag --pinglist 10.1.6.90
```

Then bring the resource in service.

```
[target1]# lkcli resource restore --tag ip-tag
```

3. Extending a resource

Execute the following command:

```
[target1]# lkcli resource extend ip --tag ip-tag --dest target2
```

### Resource Settings

Item	Input Value
--tag	Tag name of the created resource

--dest	Backup node name
--------	------------------

#### 4. Setting up a ping list for the extended resource

Run the following command to set up a ping list for the extended resource as well.

```
[target1]# lkcli resource config ip --tag ip-tag --pinglist 10.1.6.90 --remote target2
```

#### 5. Checking the resource

After creating and extending the resource, run the following command.

The resource information is displayed.

```
[target1]# lkcli status -q
LOCAL    TAG      ID                STATE    PRIO  PRIMARY
target1  ip-tag   IP-10.1.6.100    ISP      1     target1
```

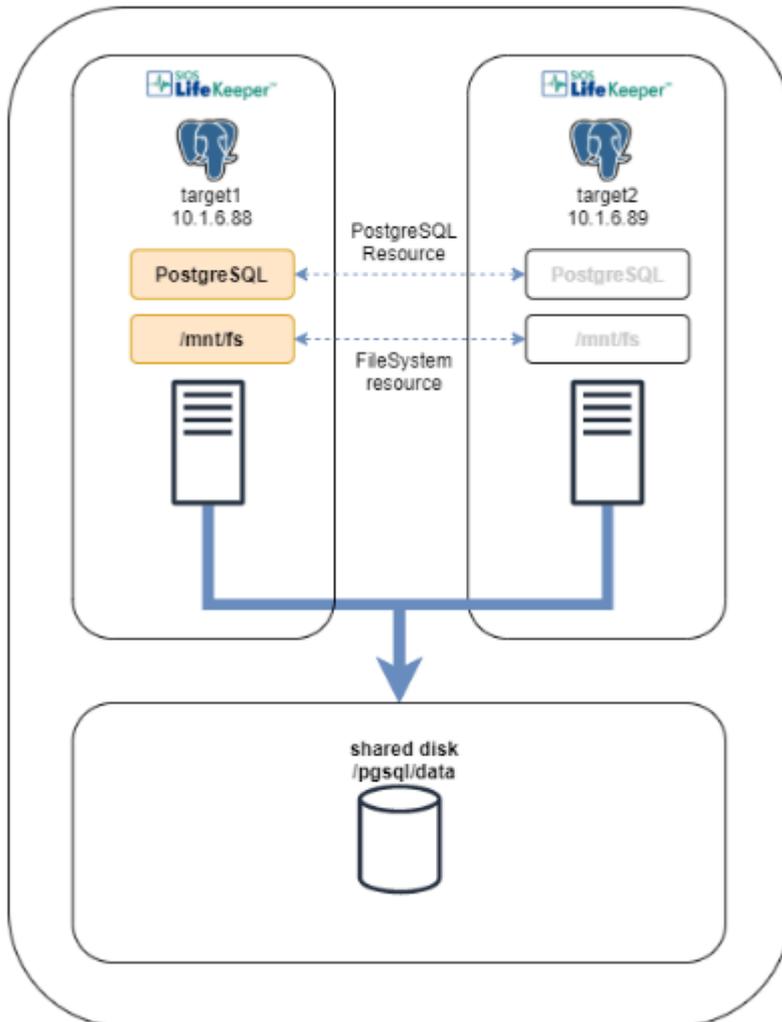
## 5.6.4.2.3. Creating a PostgreSQL Resource

This section describes how to create a PostgreSQL resource. The PostgreSQL resource provides the functionality to protect PostgreSQL database instances on LifeKeeper nodes.

### System Configuration

The environment created in this guide is a two-node configuration as shown below.

Install PostgreSQL ARK in advance.



The environment in this guide is PostgreSQL

- PostgreSQL 9.2.24 is used in this guide.
- Create a PostgreSQL data directory in the file system on the shared storage.
- The PostgreSQL database admin user should use the “postgres” created when the database is initialized.
- Protect a PostgreSQL resource instance with an active/standby configuration.

\*Make sure that your system does not fall under any of file system restrictions. See [\(LifeKeeper Core – Known Issues / Restrictions\)](#) for details.

## Perform the following steps on target1 and target2

### 1. Installing PostgreSQL

After installing, disable the autostart of the PostgreSQL service as follows:

```
# systemctl disable postgresql.service
```

## Perform the following steps on target1

### 2. Mounting the file system

Mount the file system in which the data directory will be created, refer to [Step 1 in Creating file system resources](#).

### 3. Creating a data directory

Create a PostgreSQL data directory on the shared disk.

```
[target1]# mkdir -p /mnt/fs/pgsql/data
[target1]# chown -R postgres:postgres /mnt/fs/pgsql
```

Change the following path described in the `/usr/lib/systemd/system/postgresql.service` file to the data directory on the shared disk.

```
Environment=PGDATA=/mnt/fs/pgsql/data
```

### 4. Database initialization

Execute the following command to initialize the database.

The database is created under the data directory.

```
[target1]# postgresql-setup initdb
```

### 5. Starting the PostgreSQL service

Execute the following command to start the PostgreSQL service.

```
[target1]# systemctl start postgresql.service
```

### 6. Creating a resource

Execute the following command:

```
[target1]# lkcli resource create pgsql --tag pgsql-tag --datadir /mnt/fs/pgsq
l/data --port 5432 --socket /tmp/.s.PGSQL.5432 --dbuser postgres --logfile /tm
p/pgsql-5432.lk.log
```

## Resource Settings

Item	Input Value
------	-------------

--tag	Tag name
--datadir	Absolute path of the directory that contains the PostgreSQL database data
--port	Port number used by PostgreSQL
--socket	Path of the socket used by PostgreSQL
--dbuser	PostgreSQL database administrator user name
--logfile	Absolute path to the pg_ctl log file used to start and stop PostgreSQL

## 7. Extending the resource

Execute the following command:

```
[target1]# lkcli resource extend pgsq1 --tag pgsq1-tag --dest target2
```

## Resource Settings

Item	Input Value
--tag	Tag name of the created resource
--dest	Backup node name

## 8. Checking the resource

After creating and extending the resource, run the following command.

The resource information is displayed.

```
[target1]# lkcli status -q
LOCAL    TAG          ID                STATE    PRIO  P
PRIMARY
target1  pgsq1-tag    target1.pgsq1-5432  ISP      1    t
arget1
target1  /mnt/fs      /mnt/fs           ISP      1    t
arget1
target1  device17885  36000c292eb0c693b2efb44ed56556636-1  ISP      1    t
arget1
target1  disk17816    36000c292eb0c693b2efb44ed56556636    ISP      1    t
arget1
```

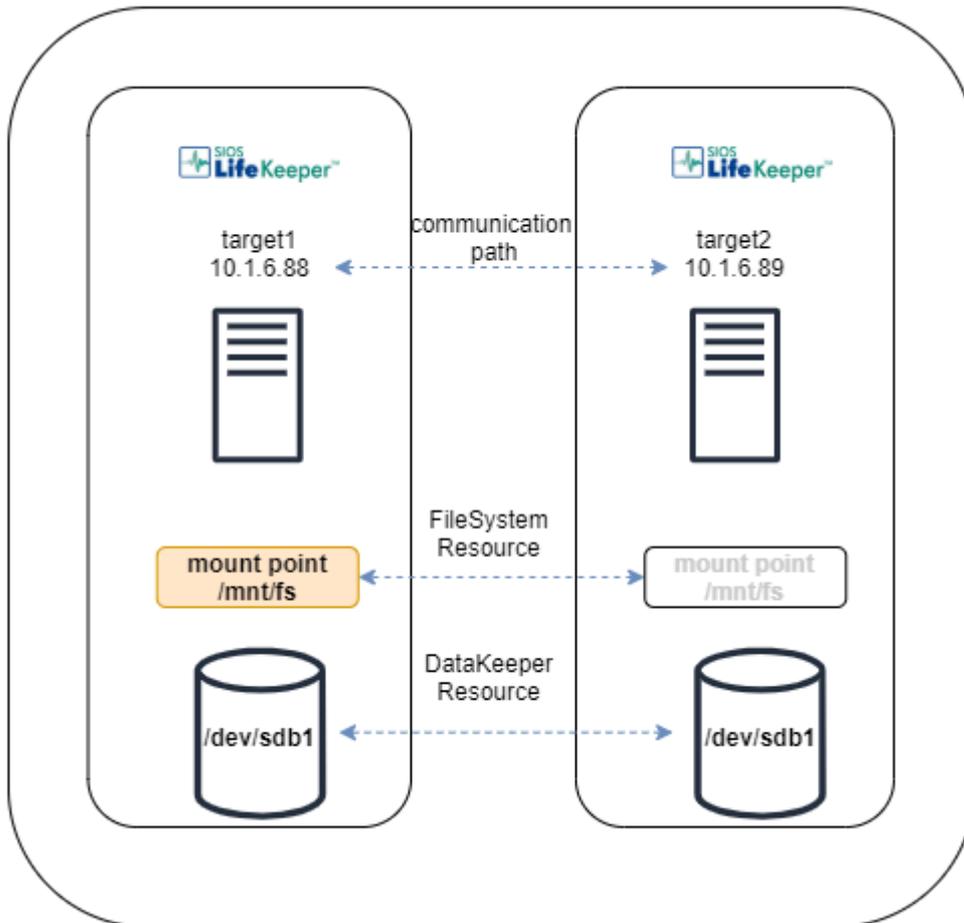
When a PostgreSQL resource is created, a file system resource is automatically created as shown above.

## 5.6.4.2.4. Creating a DataKeeper Resource

This section describes how to create a DataKeeper resource. DataKeeper resources provide the ability to build highly available clusters without the use of shared storage.

### System Configuration

The environment created in this guide has the following two-node configuration. Create a mirror on your new file system and protect it with LifeKeeper.



Make sure that your configuration satisfies the following DataKeeper requirements.

[Hardware and Software Requirements](#)

### Perform the following steps on “target1” and “target2”

1. Checking the device

```
# parted /dev/sdb print
Model: VMware Virtual disk (scsi)
Disk /dev/sdb: 1074MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

```
Number Start End Size File system Name Flags
```

2. Creating a partition

Create a partition in /dev/sdb of the device.

```
# parted -s /dev/sdb mklabel gpt
# parted /dev/sdb mkpart primary 0% 100%
# parted /dev/sdb print
Model: VMware Virtual disk (scsi)
Disk /dev/sdb: 1074MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 1049kB 1073MB 1072MB ext4 primary
```

Perform the following steps on the “target1”

1. Creating a resource

Execute the following command:

```
[target1]# lkcli resource create dk --tag dk-tag --mode synchronous --hierarch
y new --device /dev/sdb1 --fstype xfs --mount_point /mnt/dk --fstag fs-tag
```

Resource settings

Item	Input Value
--tag	DataKeeper resource tag name
--mode	Replication type
--hierarchy	Type of data replication to create
--device	Source disk or partition
--fstype	File system type
--mount_point	Mount point
--fstag	File system resource tag name

2. Extending a DataKeeper resource

Execute the following command:

```
[target1]# lkcli resource extend dk --tag dk-tag --dest target2 --mode synchro
nous --laddr 10.1.6.88 --raddr 10.1.6.89
```

## Resource settings

Item	Input Value
--tag	Tag name of the created DataKeeper resource
--dest	The hostname of the target server where the resource hierarchy is extended
--mode	Replication type
--laddr	IP address on the local machine to be set for the communication path
--raddr	IP address on the remote machine to be set for the communication path

### 3. Extending a file system resource

Execute the following command:

```
[target1]# lkcli resource extend fs --tag fs-tag --dest target2
```

## Resource settings

Item	Input Value
--tag	Tag name of the created file system resource
--dest	The hostname of the target server where the resource hierarchy is extended

### 4. Checking the resource

After creating and expanding the resource, run the following command.

The resource information is returned.

```
[target1]# lkcli status -q

LOCAL   TAG          ID                               STATE   PRIO  PRI
MARY
target1 fs-tag      /mnt/dk                         ISP     1    tar
get1
target1 dk-tag     36000c292eb0c693b2efb44ed56556636-1  ISP     1    tar
get1
```

A DataKeeper resource is automatically created with the file system resource dependent on the upper level as shown above.

### 5. Check the mirroring

Once the DataKeeper resource has been created, a full resync will be performed.

When the status is Fully Operational, the full resync is completed.

```
[target1]# lkcli mirror status --tag dk-tag
```

## Resource settings

Item	Input Value
--tag	Tag name of the created DataKeeper resource

## Export/Import a DataKeeper Resource

Refer to [Replicate the Existing Cluster Settings](#) to replicate the DataKeeper resource and the file system resource. Execute import with the file system unmounted on each node.

## 5.6.4.3. LKCLI – Checking Cluster Status

### Checking the Status of LiKeeper using `lkcli status`

The `lkcli status -q` command provides current resource information and communication path information.

```
# lkcli status -q
LOCAL      TAG          ID          STATE      PRIO  PRIMARY
target1    ip-10.1.6.100 ip-10.1.6.100  ISP        1     target1

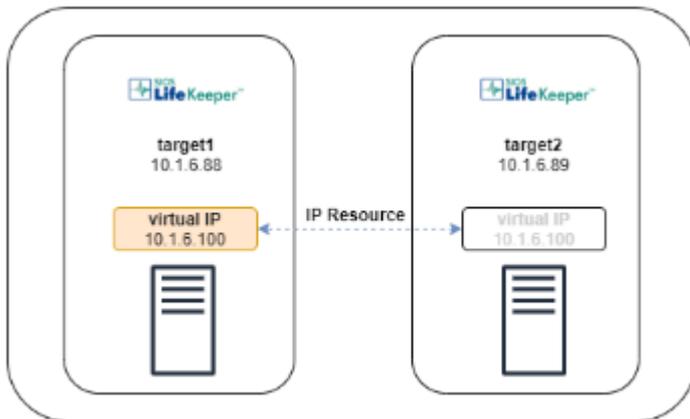
MACHINE    NETWORK ADDRESSES/DEVICE  STATE      PRIO
target2    TCP       10.1.6.88/10.1.6.89  ALIVE      1
```

Check the resource status on other cluster nodes by using the `--remote` option. Refer to [LKCLI \(LifeKeeper Command Line Interface\)](#) for more information.

## 5.6.4.4. LKCLI – Verifying Switchover Behavior

This section explains how to perform a switchover from target1 to target2.

### Configuration



### Execute the command on target2.

1. Make sure the status of the resource that you will switch over is **ISP on target1**.

```
[target2]# lkcli status -q --remote target1
LOCAL   TAG           ID           STATE        PRIO  PRIMARY
target1 ip-10.1.6.100 ip-10.1.6.100 ISP          1     target1

MACHINE NETWORK ADDRESSES/DEVICE  STATE    PRIO
target2 TCP      10.1.6.88/10.1.6.89 ALIVE    1
```

2. Make sure the status of the resource that you want to switch over is **OSU on target2**. Make a note of the resource tag name.

```
[target2]# lkcli status -q
LOCAL   TAG           ID           STATE        PRIO  PRIMARY
target2 ip-10.1.6.100 ip-10.1.6.100 OSU          10    target1

MACHINE NETWORK ADDRESSES/DEVICE  STATE    PRIO
target1 TCP      10.1.6.89/10.1.6.88 ALIVE    1
```

3. Switch the resource to target2.

```
[target2]# lkcli resource restore --tag ip-10.1.6.100
BEGIN restore of "ip-10.1.6.100"
END successful restore of "ip-10.1.6.100"
```

#### 4. Make sure the resource is ISP on target2.

```
[target2]# lkcli status -q
```

LOCAL	TAG	ID	STATE	PRIO	PRIMARY
target2	ip-10.1.6.100	ip-10.1.6.100	ISP	10	target1

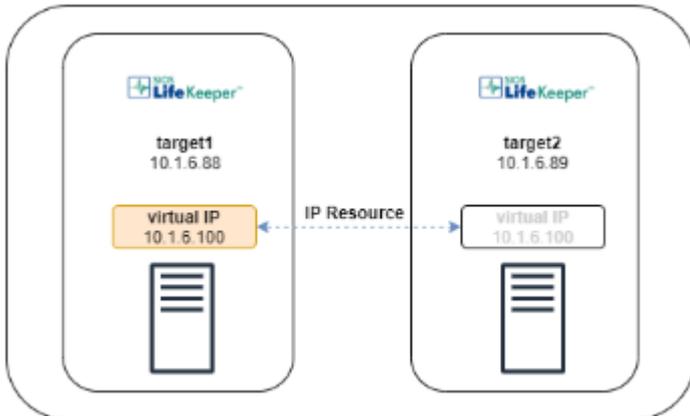
MACHINE	NETWORK	ADDRESSES/DEVICE	STATE	PRIO
target1	TCP	10.1.6.89/10.1.6.88	ALIVE	1

## 5.6.4.5. LKCLI – Maintenance Tasks

This section explains how to maintain the LifeKeeper-protected systems and resources.

### Configuration

These instructions assume that the configuration is 2 nodes.



### Maintaining a LifeKeeper Protected Machine

The maintenance tasks performed on target1 such as shutdown of the LifeKeeper protected machine, will have an impact on LifeKeeper and the resources.

#### Execute the command on target1.

1. Switch the active resource on target1 to target2
  - i. Check the status of all the resources on target1. Make a note of the resource tag name if there is a resource with an ISP status.

```
[target1]# lkcli status -q
LOCAL   TAG           ID           STATE      PRIO  PRIMARY
target1 ip-10.1.6.100 ip-10.1.6.100 ISP        1    target1

MACHINE NETWORK ADDRESSES/DEVICE  STATE      PRIO
target2 TCP      10.1.6.88/10.1.6.89 ALIVE      1
```

- ii. Switch the resources that are ISP on target1 to target2 one by one. You can execute the command from target1 by using "--remote" option.

```
[target1]# lkcli resource restore --tag ip-10.1.6.100 --remote target2
BEGIN restore of "ip-10.1.6.100"
END successful restore of "ip-10.1.6.100"
```

- iii. Make sure all the resources are OSU on target1.

```
[target1]# lkcli status -q
LOCAL      TAG                ID                STATE      PRIO  PRIMARY
target1    ip-10.1.6.100     ip-10.1.6.100   OSU        1     target1

MACHINE    NETWORK  ADDRESSES/DEVICE  STATE      PRIO
target2    TCP      10.1.6.88/10.1.6.89  ALIVE      1
```

2. Stop LifeKeeper on target1. It does not stop the resource by running “-f” option.

```
[target1]# lkcli stop -f
Removed /etc/systemd/system/lifekeeper-graphical.target.requires/lifekeeper.service.
Removed /etc/systemd/system/lifekeeper-multi-user.target.requires/lifekeeper.service.
```

3. Perform the necessary maintenance on target1.

4. Start LifeKeeper on target1.

```
[target1]# lkcli start
Created symlink /etc/systemd/system/lifekeeper-graphical.target.requires/lifekeeper.service → /usr/lib/systemd/system/lifekeeper.service.
Created symlink /etc/systemd/system/lifekeeper-multi-user.target.requires/lifekeeper.service → /usr/lib/systemd/system/lifekeeper.service.
```

5. Bring the resources in-service on target1, if desired.

```
[target1]# lkcli resource restore --tag ip-10.1.6.100
BEGIN restore of "ip-10.1.6.100"
END successful restore of "ip-10.1.6.100"
```

Refer to [Maintaining a LifeKeeper Protected System](#).

## Maintaining LifeKeeper Protected Resources

This section explains how to perform maintenance for specific resources only.

### Execute the command on target1.

1. Switch the active resources on target1 to target2.
  - i. Check the status of all of the resources on target1. If there is a resource that has an ISP status, perform a switchover following the steps below.

```
[target1]# lkcli status -q
```

LOCAL	TAG	ID	STATE	PRIO	PRIMARY
target1	ip-10.1.6.100	ip-10.1.6.100	ISP	1	target1

MACHINE	NETWORK	ADDRESSES/DEVICE	STATE	PRIO
target2	TCP	10.1.6.88/10.1.6.89	ALIVE	1

ii. Check the resource tag names on target2 which are ISP on target1.

```
[target1]# lkcli status -q --remote target2
```

LOCAL	TAG	ID	STATE	PRIO	PRIMARY
target2	ip-10.1.6.100	ip-10.1.6.100	OSU	10	target1

MACHINE	NETWORK	ADDRESSES/DEVICE	STATE	PRIO
target1	TCP	10.1.6.89/10.1.6.88	ALIVE	1

iii. Switch the resources that are ISP on target1 to target2 one by one.

```
[target1]# lkcli resource restore --tag ip-10.1.6.100 --remote target2
BEGIN restore of "ip-10.1.6.100"
END successful restore of "ip-10.1.6.100"
```

iv. Make sure all the resources have an OSU status on target1.

```
[target1]# lkcli status -q
```

LOCAL	TAG	ID	STATE	PRIO	PRIMARY
target1	ip-10.1.6.100	ip-10.1.6.100	OSU	1	target1

MACHINE	NETWORK	ADDRESSES/DEVICE	STATE	PRIO
target2	TCP	10.1.6.88/10.1.6.89	ALIVE	1

2. Perform maintenance for the resources that are OSU.

3. Bring the resources in-service on target1, if desired.

```
[target1]# lkcli resource restore --tag ip-10.1.6.100
BEGIN restore of "ip-10.1.6.100"
END successful restore of "ip-10.1.6.100"
```

Refer to [Maintaining a Resource Hierarchy](#).

## Changing the Resource Settings

This section explains how to change the IP resource settings. For other commands refer to [LKCLI Subcommands for Each ARK](#).

These instructions assume you are using LifeKeeper v9.5.0.

1. Check the resource tag name that you want to change the setting for.

```
# lkcli status -q
LOCAL      TAG                ID                STATE            PRIO  PRIMARY
target1    ip-10.1.6.100      ip-10.1.6.100    ISP              1     target1

MACHINE    NETWORK  ADDRESSES/DEVICE  STATE            PRIO
target2    TCP      10.1.6.88/10.1.6.89  ALIVE            1
```

2. Based on the tag name, check the resource type and current value. Refer to [LKCLI Subcommands for Each ARK](#) for more information.

```
# lkcli resource info --tag ip-10.1.6.100
---
app: comm
priority: 1
properties:
  device: ens192
  ipaddr: 10.1.6.100
  netmask: 255.255.255.0
  pinglist: ''
  realip: 0
  restoremode: Enabled
  srcaddr: 0
switchback: INTELLIGENT
tag: ip-10.1.6.100
typ: ip
```

3. Change the resource settings (example: `restoremode` will be Disabled at this time).

```
# lkcli resource config ip --tag ip-10.1.6.100 --restoremode Disabled
Performing restoremode change ...

restoremode change successful.
```

4. The resource settings have changed.

```
# lkcli resource info --tag ip-10.1.6.100
---
app: comm
priority: 1
properties:
  device: ens192
```

```
ipaddr: 10.1.6.100
netmask: 255.255.255.0
pinglist: ''
realip: 0
restoremode: Disabled
srcaddr: 0
switchback: INTELLIGENT
tag: ip-10.1.6.100
typ: ip
```

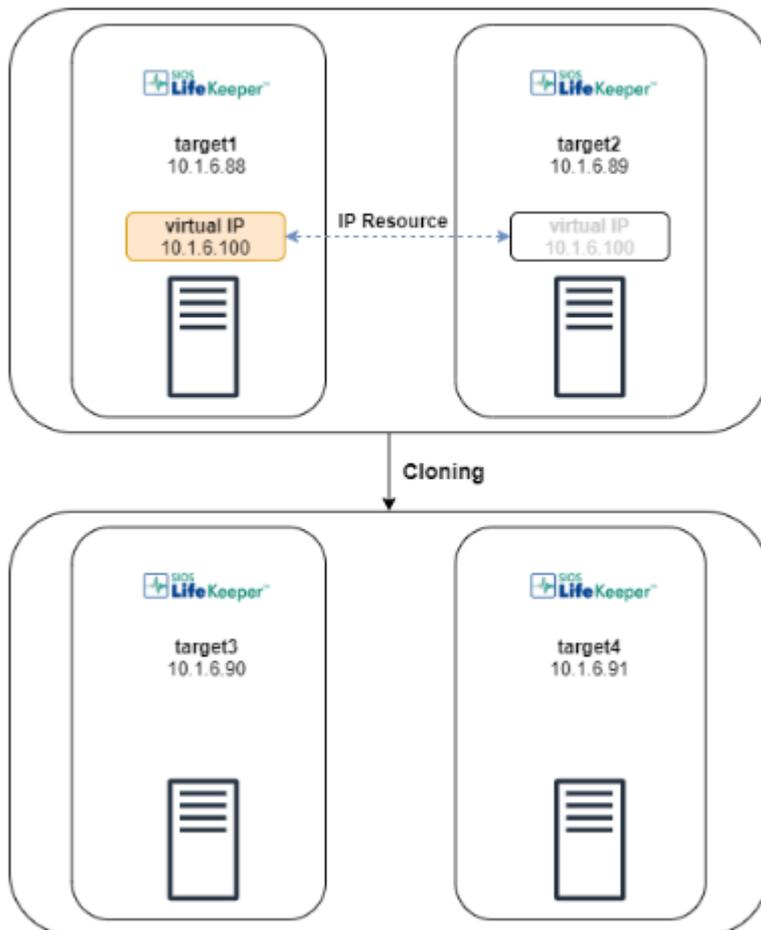
## 5.6.4.6. LKCLI – Replicate the Existing Cluster Settings

### Inherit and Duplicate the Cluster Settings

This section describes the procedure for replicating a cluster with the same settings based on the cluster with the communication path and resource set.

#### Configuration

The steps describe preparing a 2-node LifeKeeper cluster and another LifeKeeper cluster to be replicated. The other cluster is not configured with communication paths and resources.



#### Caution

The following are resource and communication path restrictions for the resource and communication path output when using the export command. Before performing the steps, check to see if the environment is supported.

- Make sure the communication path protocol that can be exported is TCP/IP.

```
# lkcli status -q
LOCAL      TAG                ID                STATE      PRIO  PRIMARY
target1    ip-10.1.6.100      ip-10.1.6.100    ISP        1     target1

MACHINE    NETWORK  ADDRESSES/DEVICE  STATE      PRIO
target2     TCP      10.1.6.88/10.1.6.89  ALIVE      1
```

- Make sure all resources in the cluster to be exported are supported ARKs with LKCLI. Refer to [LKCLI Subcommands for Each ARK](#) for a list of supported ARKs and resource types.

```
# lkcli resource info --tag ip-10.1.6.100
---
app: comm
priority: 1
properties:
  device: ens192
  ipaddr: 10.1.6.100
  netmask: 255.255.255.0
  pinglist: ''
  realip: 0
  restoremode: Disabled
  srcaddr: 0
switchback: INTELLIGENT
tag: ip-10.1.6.100
typ: ip
```

## Steps

### Execute the command on target1

1. Save the current settings to the file in the cluster where LifeKeeper is configured.

```
[target1]# lkcli export > src_settings.yml
```

2. Copy the exported file to the cluster to be replicated.

### Execute the command on target3.

3. For the exported configuration file, edit only the IP address and hostname manually.

```
[target3]# sed -e "s/10\.1\.6\.88/10\.1\.6\.90/" -e "s/target1/target3/" -e
"s/10\.1\.6\.89/10\.1\.6\.91/" -e "s/target2/target4/" src_settings.yml &gt; d
est_settings.yml
```

4. Make sure there are no communication path or resources in the cluster to be replicated (check all nodes in the cluster).

If any exist, run the following command to delete: `lkcli clean --mode all`. Note that the setting won't be restored after executing the clean command.

```
[target3]# lkcli status -q
LOCAL    TAG    ID    STATE    PRIO    PRIMARY
```

### Execute the command on target3 and target4

5. Import the resources in the replicated system and check that the communication path can be created. Execute this command on all of the nodes in the cluster.

```
# lkcli import commpath --file dest_settings.yml
Performing commpath 'target3:10.1.6.90/10.1.6.91' create...
Commpath 'target3:10.1.6.90/10.1.6.91' created successful.
```

```
# lkcli status -q
LOCAL    TAG    ID    STATE    PRIO    PRIMARY

MACHINE  NETWORK ADDRESSES/DEVICE    STATE    PRIO
target4  TCP    10.1.6.90/10.1.6.91  ALIVE    1
```

\*The communication path status will be ALIVE when it is created in both directions between 2 nodes.

6. Import the resources in the replicated system and check that the resource can be created. Execute this command only once in the cluster.

```
# lkcli import resource --file dest_settngs.yml
BEGIN create of "ip-10.1.6.100"
LifeKeeper application=comm on target1.
LifeKeeper communications resource type= ip on target1.
Creating resource instance with id ip-10.1.6.100 on machine target1
Resource successfully created on target1
BEGIN restore of "ip-10.1.6.100"
END successful restore of "ip-10.1.6.100"
END successful create of "ip-10.1.6.100".
Removing ping list for subnet 172.31.0.0...
Performing restoremode change ...

restoremode change successful.
Building independent resource list
Checking existence of extend and canextend scripts
Checking extendability for ip-10.1.6.100
Pre Extend checks were successful
```

```
Extending resource instances for ip-10.1.6.100
Creating dependencies
Setting switchback type for hierarchy
Creating equivalencies
LifeKeeper Admin Lock (ip-10.1.6.100) Released
Hierarchy successfully extended
Removing ping list for subnet 172.31.0.0...
```

## 7. The communication path and resource are created.

```
# lkcli status -q
LOCAL      TAG              ID              STATE          PRIO  PRIMARY
target3    ip-10.1.6.100   ip-10.1.6.100  ISP            1     target3

MACHINE    NETWORK ADDRESSES/DEVICE  STATE          PRIO
target4    TCP                10.1.6.90/10.1.6.91  ALIVE          1
```

## 6. Application Recovery Kits

---

LifeKeeper for Linux Application Recovery Kits (ARKs) include tools and utilities that allow LifeKeeper to manage and control a specific application. The following optional recovery kits are available with this release of LifeKeeper.

[Apache Recovery Kit Administration Guide](#)

[DB2 Recovery Kit Administration Guide](#)

[Recovery Kit for EC2 Administration Guide](#)

[Generic Application Kit for Load Balancer Health Checks](#)

[LVM Recovery Kit Administration Guide](#)

[IP Recovery Kit Administration Guide](#)

[MySQL Recovery Kit Administration Guide](#)

[WebSphere MQ Recovery Kit Administration Guide](#)

[NAS Recovery Kit Administration Guide](#)

[NFS Recovery Kit Administration Guide](#)

[Recovery Kit for Oracle Cloud Infrastructure Administration Guide](#)

[Oracle Recovery Kit Administration Guide](#)

[PostgreSQL Recovery Kit Administration Guide](#)

[Postfix Recovery Kit Administration Guide](#)

[Recovery Kit for Route 53™ Administration Guide](#)

[Samba Recovery Kit Administration Guide](#)

[SAP Recovery Kit Administration Guide](#)

[SAP HANA Recovery Kit Administration Guide](#)

[SAP MaxDB Recovery Kit Administration Guide](#)

[Sybase Recovery Kit Administration Guide](#)

[VMDK Shared Storage Recovery Kit Administration Guide](#)

# 6.1. Apache Recovery Kit Administration Guide

---

The LifeKeeper for Linux Apache Web Server Recovery Kit provides fault resilience for Apache Web Server software in a LifeKeeper environment.

This guide explains the following topics:

- [LifeKeeper for Linux Documentation](#). A list of all the LifeKeeper for Linux documentation and where the information is available.
- [Requirements](#). Before you can install and set up the recovery software, your server must meet certain hardware and software requirements. You should refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper for Linux Apache Recovery Kit.
- [Configuring Your Recovery Kit](#). To ensure that your LifeKeeper configuration provides the protection and flexibility you require, you need to be aware of the configuration rules. To appropriately plan your configuration, you must understand your network configuration, interface selection, user system setup, hierarchy options and the Apache configuration tasks. In addition to planning your configuration, this section also includes configuration examples and the specific tasks required to configure your Recovery Kit.
- [Troubleshooting](#). This section provides a list of informational and error messages with recommended solutions.

## 6.1.1. LifeKeeper Documentation and Apache References

---

The following is a list of LifeKeeper related information available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#) (available from the Help menu within the LifeKeeper GUI)
- [LifeKeeper for Linux Installation Guide](#)

This documentation, along with documentation associated with other LifeKeeper Recovery Kits, is provided online at:

<http://docs.us.sios.com>

### Reference Documents

The following is a list of reference documents associated with the Apache Web Server application and the LifeKeeper Apache Recovery Kit:

- Apache Online documentation
- Apache: The Definitive Guide, 2nd Edition, Ben Laurie and Peter Laurie, O'Reilly & Associates, Inc. 1999

## 6.1.2. Apache Recovery Kit Requirements

Before attempting to install or remove the Apache Recovery Kit, you must understand the hardware and software requirements for the package and the installation and removal procedures.

### Kit Hardware and Software Requirements

Before installing and configuring the LifeKeeper Apache Recovery Kit, be sure that your configuration meets the following requirements:

- **Servers.** The Recovery Kit requires two or more supported computers configured in accordance with LifeKeeper requirements described in the LifeKeeper for Linux Technical Documentation and the LifeKeeper for Linux Release Notes, which are located on our SIOS Technical Documentation site at [docs.us.sios.com](https://docs.us.sios.com).
- **LifeKeeper software.** You must install the same version of LifeKeeper software and any patches on each server. Please refer to the LifeKeeper for Linux Release Notes and LifeKeeper for Linux Technical Documentation for specific LifeKeeper requirements.
- **LifeKeeper IP Recovery Kit.** You must have the same version of this Recovery Kit on each server.
- **IP network interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

 **Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation, and so forth.

- **TCP/IP software.** Each server also requires the TCP/IP software.
- **Apache software.** Each server must have the Apache Web Server software installed and configured prior to configuring LifeKeeper and the LifeKeeper Apache Web Server Recovery Kit, including any DSO (Dynamic Shared Object) modules that will be used. The same versions of all web server software packages should be installed on each server. Consult the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information.

Refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper Apache Recovery Kit.

## 6.1.3. Configuring Apache Web Server with LifeKeeper

---

This section contains definitions and examples of typical LifeKeeper Apache Web Server configurations and information you should consider before you start to configure Apache Web Server.

Please refer to the [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring LifeKeeper Core resource hierarchies.

## 6.1.3.1. Configuration Definitions and Examples

---

Apache Web Server supports multiple instances of the httpd daemon running at the same time. Each LifeKeeper Apache Web Server hierarchy corresponds to a separate Apache instance with its own “server root” directory. Each instance may support one or more web sites, depending on whether or not it has been configured to use “virtual hosts.”

Primarily, the server root directory defines an Apache Web Server instance, since this directory will contain the conf/httpd.conf configuration file that specifies how the web instance is configured. The Apache configuration directives within this file will determine where the log files, web documents, other configuration files, etc. are located for the instance, as well as which IP and/or domain name addresses will be used.

It is useful to characterize Apache Web Server configurations with LifeKeeper based on whether or not a LifeKeeper file system (which uses shared storage) will be used. A single shared file system may be used for the server root directory (along with the configuration file conf/httpd.conf) and/or the document root directories (and optionally the httpd executable itself). Whether you choose to use a local or a shared configuration for a particular Apache instance will depend on two main factors: the difficulty of maintaining separate, identical copies of the configuration files and/or web site documents, and the availability and accessibility of storage which can be shared (or mirrored) between two or more servers. Note, however, that you may choose to configure both local and shared Apache instances on the same servers.

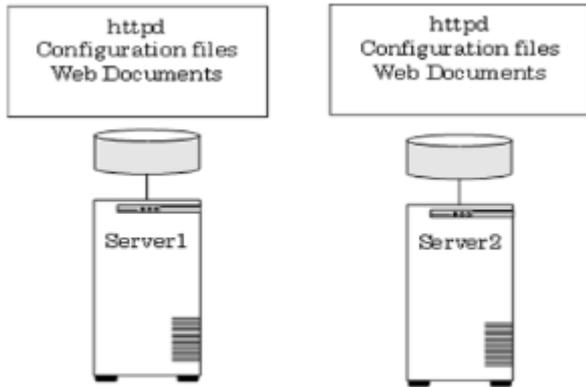
The following sections provide examples of [Local](#) and [Shared](#) Apache Web Server configurations in a LifeKeeper environment and summarize the main characteristics of each.

### Local Configuration

In a typical local configuration, nothing is shared between the servers. Identical copies of the Apache Web Server configuration file, web documents, DSO modules (and their configuration files, if any), and the httpd executable reside in exactly the same locations on each server. It is the responsibility of the Apache administrator to maintain identical copies of the Apache components on the different servers.

Each web site is assigned an IP address – or a domain address that maps to a particular IP address – through the configuration file, and a LifeKeeper IP address is created for each and added to the Apache resource hierarchy. When the Apache hierarchy is switched over from one server to another, this particular httpd instance is stopped and the IP addresses are deactivated on the first server, then the IP addresses are reactivated and the instance started on the other server. Clients will then be automatically connected via TCP/IP to the identical web site on the other server.

Figure 1. Local Configuration



Configuration Notes:

- Figure 1 is an example of a local configuration where nothing is on a shared file system.
- Each server has the same version of the Apache Web Server executable at the same location (typically /usr/bin/httpd).
- Each server has the same server root directory where identical copies of the configuration file for each instance are placed.
- Each server has the same document root directory(s) where identical copies of the web document for each instance are placed.
- If DSO modules are being used, each server has identical copies at the same location.

Creating an Apache Web Server resource hierarchy on Server 1:

Server:	Server 1
Web Server Binary Location:	/usr/sbin/httpd
Web Server Root Directory:	/home/www/examples/instance1/
Root Tag	apache-www.examples.instance1

Extending an Apache Web Server resource hierarchy to Server 2:

Template Server:	Server 1
Tag to Extend	apache-www.examples.instance1
Target Server	Server2
Target Priority:	10

Note that when an Apache resource hierarchy is extended to one or more additional servers, the same Web Server Binary Location and Web Server Root Directory must be used on all servers, regardless of

whether this is a local or a shared configuration. See the discussion above and the section on [Specific Configuration Considerations for Apache Web Server](#) for additional information. Also during hierarchy extension, LifeKeeper extends all the dependent resources which are part of the Apache resource hierarchy.

## Shared Configuration

In a typical shared configuration, the server root directory and the document root directories are all on the same shared file system. The same configuration file and web documents are shared between the servers, so there is no need to maintain identical copies on each server. If DSO modules are being used, they also can be located on the same shared file system, along with any configuration files or resources they may need.

Note that you may choose to place only the web documents on a shared file system. This will still appear much like a typical local configuration, since the server root directories will be local, but the hierarchy will also include a shared file system.

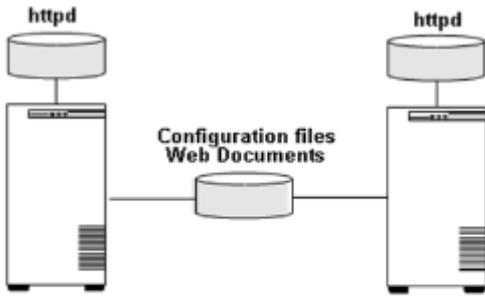
If you wish to use a particular version or a separate copy of the Apache executable for this Apache resource hierarchy, you may place this executable on the shared file system as well and it will be available only to this instance. To do this, simply enter the full path of the httpd executable on the shared file system when prompted for the Web Server Binary Location.

Note that only one shared file system may be used, since this assures that all required components which are on shared storage will be available at the same time. If you choose to use a Web Server Binary Location on a shared file system, you must also choose a Web Server Root Directory on the same shared file system, and all DocumentRoot directories configured for this server root must be on the same shared file system. Likewise, when you choose a Web Server Root Directory on a shared file system, all DocumentRoot directories must be on the same shared file system. If neither the binary nor the server root is placed on a shared file system, but any of the DocumentRoot directories are shared, all DocumentRoot directories must be shared on the same file system.

These rules can be summarized as follows:

- If the Apache executable is shared, then the server root directory must be shared.
- If the server root directory is shared, then all DocumentRoot directories must be shared.
- If any DocumentRoot directory is shared, they must all be shared.
- Only one shared file system is allowed for each Apache resource hierarchy.

Figure 2. Shared Configuration



Configuration Notes:

- Figure 2 is an example of a shared configuration with shared configuration files and web documents.
- You may choose to place only the web documents on a shared file system. This will still appear much like a typical local configuration, except that the hierarchy will also include a shared file system.
- If DSO modules are used, they may reside on the shared file system, along with any configuration files or resources they need.

Creating an Apache Web Server resource hierarchy on Server 1:

Server:	Server1
Web Server Binary Location:	/usr/sbin/httpd .....OR.... /shared/example/instance2/bin/httpd
Web Server Root Directory:	/shared/example/instance2
Root Tag	apache-shared.example.instance2

Extending an Apache Web Server resource hierarchy to Server 2:

Template Server:	Server1
Tag to Expand	apache-shared.example.instance2
Target Server	Server2
Target Priority:	10

## 6.1.3.1.1. Active/Standby and Active/Active Configurations

---

Apache Web Server is called an Active/Active application with LifeKeeper. This means that more than one instance of Apache can be running on a server at any time. For example, if two servers are running an instance of Apache and one server fails, the Apache instance on this server can fail over to the other server and it can continue to run its own instance as well. Some applications simply don't support this, so you would have to keep a server available for each instance of the application. These are called Active/Standby applications. Some applications can be configured either way.

There may be circumstances when you might want to operate Apache in an "active/standby" mode, particularly if only one of your servers is used primarily for running Apache. In this particular case, you should disable the automatic startup of the standard Apache default installation so that nothing is running on the backup server(s).

By manually bringing the Apache instances In Service on one or more particular servers, you can distribute the workload as you like. And by adjusting the server priorities for each of your instances, you can configure the Apache instances to fail over to a particular server only as a last resort, or to fail over to different servers to distribute the workload when a failure occurs.

If you disable automatic startup of Apache on all servers in the cluster, it is possible to use the default server root directory "/etc/http" for a single LifeKeeper Apache resource hierarchy by simply configuring this instance to use LifeKeeper IP addresses – and possibly using a shared file system for the document root directories. Note, however, that this would be an Active/Standby configuration (as described above), so you could no longer start up the default instance in the usual way. Of course, the default server root directory cannot be used for more than one hierarchy, since the server root must be unique.

## 6.1.3.2. Configuration Considerations for Apache Web Server

---

Before you create Apache resource hierarchies, you will need to make sure you have completed the following configuration tasks for the Apache Web Server application:

1. In the case Apache package attached to a distribution is installed, it is normally set to autostart at system startup, so it conflicts with LifeKeeper's protection. To avoid the conflict, refer to the manual of each distribution and disable the automatic start up.

 **Note: For Apache on SuSE:** The default installation of Apache on SuSE does not place the `httpd.conf` configuration file in a subdirectory of `ServerRoot` called `conf`. If you are using the default installation of Apache on SuSE, you must relocate the configuration file to the directory `/etc/httpd/conf`.

2. You must create a separate, distinct root directory for each LifeKeeper Apache Web Server hierarchy. This "server root" directory corresponds to the Apache "ServerRoot" configuration and command line parameters. Each LifeKeeper Apache resource hierarchy will correspond to a unique Apache instance with its associated server root directory. Note that the server root directory must be identical on all servers that are configured for a particular Apache hierarchy. You must place all configuration file information for the web site in the standard location relative to the server root (`conf/httpd.conf`) so that it can be found and accessed by the LifeKeeper software.
3. You must configure all web sites (virtual hosts) to listen on specific LifeKeeper IP addresses using `BindAddress` or `Listen` directives. These LifeKeeper-protected IP addresses must already be created and available to be brought in-service where the Apache hierarchy is to be created. They will automatically be added to the Apache resource hierarchy.

If you will be using a LifeKeeper shared file system, you must make all necessary preparations for the file system creation prior to creating the Apache hierarchy. In particular, the file system must be mounted on the server where the Apache hierarchy is to be created. If the LifeKeeper file system hierarchy has not already been created, it will automatically be created along with the Apache hierarchy, then joined to the Apache resource hierarchy.

Consult the Apache Web Server documentation for detailed information on configuring virtual hosts. As noted above, you must configure all Apache instances to listen on specific LifeKeeper-protected addresses. For example, the configuration file for an instance that combines IP-based and name-based virtual hosts would include directives like the following:

```
User webuser
Group webgroup
ServerName localhost
```

```
Listen 172.17.100.55:8000
NameVirtualHost 172.17.100.55:8000
```

```
Listen 172.17.100.56:80
```

```
<Virtualhost site.name_one:8000>  
ServerName site.name_one  
DocumentRoot /shared/site/name_one  
</VirtualHost>
```

```
<VirtualHost site.name_two:8000>  
ServerName site.name_two  
DocumentRoot /shared/site/name_two  
</VirtualHost>
```

```
<VirtualHost 172.17.100.56:80>  
ServerName site.ip  
DocumentRoot /shared/site/ip  
</VirtualHost>
```

4. If SSL support is enabled for your Apache instance, you must configure the SSL Listen directive, often found in a separate `ssl.conf` file, to use the appropriate LifeKeeper-protected IP address. Otherwise, the creation of your Apache hierarchy will fail with an error indicating that the IP address 0.0.0.0 is not LifeKeeper protected. Note that SSL support is enabled by default in the Apache configuration files of some Linux distributions.

For example, change the following entry in the default SSL configuration file at `/etc/httpd/conf.d/ssl.conf` from

```
Listen 0.0.0.0:443
```

to

```
Listen 172.17.100.55:443
```

5. For a Local configuration, you must install and configure Apache in the same location on both the primary and all backup servers and set up identical (or equivalent) configuration files in the same server root directory on all servers. Also, all document root directories must exist on all servers and should contain identical files. (See the section on [Local Configuration](#) in Configuration Definitions and Examples.)
6. For a Shared configuration, you will typically configure the server root directory on a LifeKeeper shared file system. Note that only one shared file system may be used, since this assures that all required components which are on shared storage will be available at the same time. Therefore, all document root directories must be subdirectories of the same shared file system, but they need not be subdirectories of the server root directory itself. You may place an Apache executable on the same shared file system as well, but this executable will only be available for use by this particular Apache resource hierarchy.

 **Note:** You don't necessarily need to place the server root directory on a shared file system in order to make use of shared storage. You may choose a local server root directory for configuration files, etc., and place only the document root directories on a shared file system. However, you must configure identical server root directories and identical (or equivalent) configuration files on all servers (as for a Local configuration as described above), and all document root directories must be on the same shared file system. (See the section on [Shared Configuration](#) in Configuration Definitions and Examples.)

7. Some web site implementations make use of DSO (Dynamic Shared Object) modules to extend Apache support for certain features. For example, there are modules available that implement functionality for PHP and Perl. These modules can be loaded and accessed at runtime by the Apache core. If you are using modules, they must be identically configured on every server in the cluster. Consult the documentation for the module package, and the vendor-supplied documentation for configuring Apache to use modules on your Linux platform. Depending on the module and the resources it uses, some objects may be required to reside on shared storage to facilitate proper failover. In some cases, a module may even need to be protected separately using the Generic Application Recovery Kit, or a custom recovery kit.
8. If you are using the SSL (Secure Sockets Layer) module with Apache, it is important that the server not be password protected. When the web server is password protected, the administrator must interactively type in the password at a prompt each time the daemon starts. Since this manual step is not consistent with a High Availability environment where recovery time is critical, LifeKeeper does not support password protected instances. Use the following command to remove the password:

```
openssl rsa -in server.key -out unprotected_server.key
```

Enter the server key password when prompted. To preserve the security of your site, make sure that the file is readable only by root!

```
chmod 400 unprotected_server.key
```

During the hierarchy creation of an Apache instance, the Recovery Kit checks that the resource is not password protected. If it is password protected, hierarchy creation will fail with an error message. However, when the instance is extended to another server, the Recovery Kit does not check for password protection on the backup server. You need to make sure that the hierarchy you are extending is not password protected.

The server key file(s) (specified by the SSLCertificateKeyFile directive(s) in the Apache configuration file) must have the same name and be at the same location on all servers in the cluster.

 **Note:** The PID file name of the httpd process that LifeKeeper uses has the following format:

```
"/var/run/httpd.<TAG name>.pid"
```

This PID file name is different from the default PID file name used by the OS. If you need to reference this PID file (ex. log rotate), please note the LifeKeeper PID file name and format.

9. In the case where "APACHE\_SERVER\_FLAGS" is defined in "/etc/sysconfig/apache2" in SuSE environments, add a "-D" in front of the flag name.

**Example:** APACHE\_SERVER\_FLAGS="-D SSL"

If there is no "-D", the resource creation and start up can fail.

## 6.1.4. LifeKeeper Configuration Tasks for Apache

---

You can perform the following configuration tasks from the LifeKeeper GUI. The following four tasks are described in this guide, as they are unique to an Apache resource instance and different for each Recovery Kit.

- [Create a Resource Hierarchy](#). Creates an application resource hierarchy in your LifeKeeper cluster.
- [Extend a Resource Hierarchy](#). Extends a resource hierarchy from the primary server to a backup server.
- [Unextend a Resource Hierarchy](#). Unextends (removes) a resource hierarchy from a single server in the LifeKeeper cluster.
- [Delete a Resource Hierarchy](#). Deletes a resource hierarchy from all servers in your LifeKeeper cluster.

The following tasks are described in the GUI Administration section within the LifeKeeper Technical Documentation, because they are common tasks with steps that are identical across all Recovery Kits.

- **Create a Resource Dependency**. Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- **Delete a Resource Dependency**. Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- **In Service**. Brings a resource hierarchy into service on a specific server.
- **Out of Service**. Takes a resource hierarchy out of service on a specific server.
- **View/Edit Properties**. View or edit the properties of a resource hierarchy on a specific server.

 **Note:** Throughout the rest of this section, we explain how to configure your Recovery Kit by selecting certain tasks from the Edit menu of the LifeKeeper GUI. You can also select each configuration task from the toolbar. You may also right click on a global resource in the Resource Hierarchy Tree (left-hand pane) of the status display window to display the same drop down menu choices as the Edit menu. This, of course, is only an option when a hierarchy already exists.

You can also right click on a resource instance in the Resource Hierarchy Table (right-hand pane) of the status display window to perform all the configuration tasks, except Creating a Resource Hierarchy, depending on the state of the server and the particular resource.

# 6.1.4.1. Creating an Apache Web Server Resource Hierarchy

**\* IMPORTANT:**

Before you create your Web Server resource hierarchy, you must make sure that your Apache configuration file has included an existing LifeKeeper-protected IP resource.

In a shared environment where the web documents and/or configuration files are on a shared disk, you must make sure that the shared file system is mounted. It is also important to remember that you require a working communication path (i.e. heartbeat) before you can extend your resource to a backup server.

To create a replicated file system hierarchy using SIOS DataKeeper, see [Creating a DataKeeper Resource Hierarchy](#).

To create a resource instance from the primary server, you should complete the following steps:

1. From the LifeKeeper GUI menu, select Edit, then Server. From the menu, select Create Resource Hierarchy.

The Apache Web Server should not be running when you create the resource. However, if you set up the listen variable in the system configuration file, the default daemon can be allowed to run.

The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized recovery kits installed within the cluster.

2. Select Apache Web Server and click **Next**.
3. You will be prompted to enter the following information. When the Back button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter an error requiring you to correct previously entered information. You may click Cancel at any time to cancel the entire creation process.

Field	Tips
Switchback Type	Choose either intelligent or automatic This dictates how the Apache instance will be switched back to this server when the server comes back up after a failover. The switchback type can be changed later from the General tab of the Resource Properties dialog box.
Server	Select the Server where you want to place the Apache Web Server (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop down list box.
Web Server Binary	Select or enter the full path name (including the file name) of the httpd Apache Web Server daemon. The default is /usr/sbin/httpd.

Location	
Web Server Root Directory	<p>You must provide the full path of the Web Server Root directory; a relative path or symbolic link may not be used. The Apache Web Server configuration file is located in conf/httpd.conf relative to the Server Root.</p> <p><b>Note:</b> At this point, LifeKeeper will check that there is a protected IP resource available. It will also validate that you have provided valid data to create your Apache Web Server resource hierarchy. If LifeKeeper detects a problem with either of these validations, an ERROR box will appear on the screen. If the Web Server Root Directory path is valid, but there are errors with the Apache configuration itself, you may pause to correct these errors and continue with the hierarchy creation. You may even pause to create any LifeKeeper IP resources that are required.</p>
Root Tag	<p>Select or enter the tag name given to the Web Server hierarchy. You can select the default, which is apache &lt;root directory&gt;, or enter your own tag name.</p>

4. Click **Create**. The Create Resource Wizard will then create your Apache resource hierarchy. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.
5. An information box will appear indicating that you have successfully created a Apache resource hierarchy, and you must Extend that hierarchy to another server in your cluster in order to achieve failover protection. Click **Next**.

 **Note:** You may encounter error messages indicating that the new Apache instance has failed to start correctly. Note that the new Apache hierarchy must be started (In Service) before it can be extended to another system. You may pause at this point and correct the problem based on the error message displayed, then bring the new hierarchy In Service before proceeding with extending the hierarchy.

6. Click **Continue**. LifeKeeper will then launch the Pre-ExtendWizard. Refer to Step 2 under Extending an Apache Resource Hierarchy (below) for details on how to extend your resource hierarchy to another server.

If you click Cancel, a dialog box will appear warning you that you will need to come back and extend your Apache resource hierarchy to another server at some other time to put it under LifeKeeper protection.

## 6.1.4.2. Extending an Apache Web Server Resource Hierarchy

This operation can be started from the Edit menu, or initiated automatically upon completing the Create Resource Hierarchy option, in which case you should refer to Step 2 below.

1. On the **Edit** menu, select **Resource**, then **Extend Resource Hierarchy**. The Pre-Extend Wizard appears. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The Pre-Extend Wizard will prompt you to enter the following information.

**Note:** The first two fields appear only if you initiated the Extend from the Edit menu.

Field	Tips
Template Server	Enter the server where your Apache resource is currently in service. It is important to remember that the Template Server you select now and the Tag to Extend that you select in the next dialog box represent an in service resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you selected. The drop down box in this dialog provides the names of all the servers in your cluster.
Tag to Extend	Select the name of the Web Server instance you wish to extend from the template server to the target server. The wizard will list in the drop down list box all the resources that you have created on the template server, which you selected in the previous dialog box.
Target Server	Select the Target Server where you are extending your Web Server resource hierarchy. The drop down box provides the names of the servers in your cluster that are not already in the selected hierarchy.
Switchback Type	Select either intelligent or automatic. This dictates how the Web Server instance will be switched back to this server when it comes back into service after a failover to the backup server. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.
Template Priority	Select or enter a priority for the template hierarchy. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. <b>Note:</b> This selection will appear only for the initial extend of the hierarchy.
Target Priority	Select or enter the Target Priority of your extended Web Server resource. The priority is a number between 1 and 999 indicating a server's priority in the cascading failover sequence for the resource. The hierarchy priorities are sorted numerically, where a lower number means a higher priority (the number 1 indicates the highest priority). Note that LifeKeeper automatically assigns the number "1" to the server that the hierarchy is created on. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.

	<p>After receiving the message that the pre-extend checks were successful, click Next.</p> <p>Depending upon the hierarchy being extended, LifeKeeper will display a series of information boxes showing the Resource Tags to be extended, which cannot be edited. Click Extend</p>
Network Interface	Select or enter the Network Interface. This is the name of the network interface (i.e. Ethernet card) the IP resource will use on the target server.
Backup Interface	Select a Backup Interface if you want to engage the IP Local Recovery feature on the server to which you are extending the IP resource. The default value is none; however, if you have another network interface card configured on this server, it should be listed in the drop down list.
IP Resource Tag	Select or enter the IP Resource Tag. This is the resource tag name to be used by the IP resource being extended to the target server.
Root Tag	Select or enter the Root Tag. This is the tag name given to the Web Server hierarchy. By default, the Root Tag name should be the same on both the template and target server.
Mount Point	<p>This selection appears only when the Web Server Root Directory is on a shared file system.</p> <p>Select or enter the Mount Point of the shared file system where the Web Server Root Directory is located. The Template Server and Target Server should have the same mount point for the shared Web Server Root Directory. The default mount point provided in the dialog box should be selected in most cases.</p>
Root Tag	<p>This selection appears only when the Web Server Root Directory is on a shared file system.</p> <p>Select or enter the Root Tag. This is the tag name of the shared file system.</p>

- An information box will appear verifying that the extension is being performed. Click Next Server if you want to extend the same Apache resource instance to another server in your cluster. This will repeat the Extend Resource Hierarchy operation.

If you click **Finish**, LifeKeeper will verify that the extension of the Web Server resource was completed successfully.

- Click **Done** in the last dialog box to exit from the Extend Resource Hierarchy menu selection.

**Note:** Be sure to test the functionality of the new instance on both servers.

## 6.1.4.3. Unextending an Apache Web Server Resource Hierarchy

---

1. On the Edit menu, select **Resource**, then **Unextend Resource Hierarchy**
2. Select the **Target Server** where you want to unextend the Web Server resource. It cannot be the server where the Web Server is currently in service. (This dialog box will not appear if you selected the Unextend task by right clicking on a resource instance in the right pane).

Click **Next**.

3. Select the Web Server hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the Unextend task by right clicking on a resource instance in either pane).
4. An information box appears confirming the target server and the Web Server resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the Web Server resource was unextended successfully.
6. Click **Done** to exit.

## 6.1.4.4. Deleting an Apache Web Server Resource Hierarchy

---

It is important to remember that if you delete a hierarchy before you take it out-of-service, the resource hierarchy will be removed from LifeKeeper protection, but the Apache instance will continue to run on the currently active server unless it is manually stopped or the system is rebooted. Attempting to recreate the same Apache hierarchy with a different IP address(s) or to create a new Apache hierarchy using the previously used IP address(s) (but using a different Server Root), will result in conflicts with the Apache instance that was left running with that same address.

To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Delete Resource Hierarchy**.
2. Select the name of the **TargetServer** where you will be deleting your Web Server resource hierarchy. Click **Next** to proceed to the next dialog box.

**Note:** If you selected the Delete Resource task by right clicking from the right pane on an individual resource instance, or from the left pane on a global resource where the resource is on only one server, the Target Server dialog will not appear.

3. Select the **Hierarchy to Delete**. (This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in the left or right pane.) Remember that the list box displays every hierarchy on the target server, both in service and out of service. If you want to stop the Apache instance and remove the resource hierarchy from LifeKeeper protection, you must make sure that the hierarchy you choose is out-of-service before deleting it.

Click **Next**.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Next**.
5. Another information box appears confirming that the Web Server resource was deleted successfully.
6. Click **Done** to exit.

## 6.1.4.5. Testing an Apache Web Server Resource Hierarchy

---

You can test your Apache resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting Edit, then Resource, then finally InService from the drop down menu. For example, an in service request executed on a backup server causes the application hierarchy to be placed in service on the backup server and taken out of service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the Out of Service request, the application is taken out of service without bringing it in service on the other server.

### Recovery Operations

When the primary server fails, the Apache Recovery Kit software performs the following tasks:

- Brings Apache into service on the backup server by bringing in service the IP address(s) on one/more of that server's physical network interfaces
- Mounts the file system—if one is being used—on the shared disk on that server
- Starts the daemon processes related to Apache

After recovery, Apache Web Server users may reconnect by clicking on the Reload/Refresh button of their browsers.

## 6.1.5. Apache Web Server Troubleshooting

---

This section provides a list of messages that you may encounter during the process of creating and extending a LifeKeeper Apache Web Server resource hierarchy, removing and restoring a resource, and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. Messages from other LifeKeeper components are also possible. In these cases, please refer to the Message Catalog(located on our Technical Documentation site under “Search for an Error Code”) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

Messages in this section fall under these topics:

- [Hierarchy Creation](#)
- [Extend Hierarchy](#)
- [Hierarchy Remove, Restore and Recovery](#)

## 6.1.5.1. Apache Hierarchy Creation Errors

---

The error messages that might be displayed during the Apache hierarchy creation are listed below, along with a suggested explanation for each. Error messages displayed by the LifeKeeper core and by other recovery kits are not listed in this guide. Note that you may stop to correct any of the problem(s) described here, and then continue with hierarchy creation from the point where you left off – including creating any new LifeKeeper resources you might need for your Apache configuration.

### During Validation of Web Server Binary Location

```
"Error: valid_httpd_path: Must specify absolute path to httpd executable."
```

Enter the full, absolute path name to a valid Apache httpd executable.

```
"Error: valid_httpd_path: File does not exist at path specified."
```

A valid Apache httpd executable does not exist at the location specified.

```
"Error: valid_httpd_path: Httpd failed to display Server version."
```

The httpd executable at the location specified does not display the standard Apache "Server version."

```
"Error: valid_httpd_path: Incorrect version $MAJOR.$MINOR.$POINT of Apache at $HTTPD_PATH."
```

The Apache httpd executable at the location specified displays the incorrect "Server version."

### During Validation of Web Server Root Directory

```
"Error: valid_http_root: Cannot find Apache configuration file at $CONF_FILE."
```

Must have valid Apache configuration file at conf/httpd.conf relative to the Server Root directory specified. Note that the default installation of Apache on SuSE does not place the httpd.conf configuration file in a subdirectory of ServerRoot called conf. If you are using the

default installation of Apache on SuSE, you must relocate the configuration file to the directory `/etc/httpd/conf`.

```
"Error: valid_http_root: Must specify absolute path to Apache server root directory."
```

Enter full, absolute path name to a Server Root directory.

```
"Error: valid_http_root: Apache instance at $HTTP_ROOT is already under LifeKeeper protection."
```

Each instance must have its own, unique Server Root directory, with configuration file located at `conf/httpd.conf`. The Server Root directory specified is already being used by another Apache instance.

```
"Syntax error on line <line number> of <configuration file path>, etc..."
```

Syntax error(s) were found in the Apache configuration file. These error messages were displayed by the `httpd -T` command when used to check the syntax of `$CONF_FILE`. See the error messages displayed for details.

```
"Error: valid_http_root: Since $HTTPD_PATH is shareable on $HTTPD_PATH_SHARED, $HTTP_ROOT must be also."
```

If the `httpd` executable is on shared/shareable storage, the Server Root and all DocumentRoot directories must be also.

```
"Error: valid_http_root: Since $HTTP_ROOT is shareable on $HTTP_ROOT_SHARED, all document root directories must shareable on this same filesystem."
```

If the Server Root is on shared/shareable storage, all DocumentRoot directories must be also.

```
"Error: http_docs_shared: Since one/more Apache document root directories are shareable on $docs_shared, $curr_root must be also."
```

If any DocumentRoot directories are on a shared/shareable file system, all DocumentRoot directories must be located on the same file system.

```
"Error: valid_http_root: Must include BindAddress or Listen directives
for each Apache instance. Check the Apache configuration file at $CONF_FILE."
```

In order to run multiple instances of Apache, each configuration file must contain BindAddress or Listen directives. Please refer to the [Configuration Considerations for Apache Web Server](#) section earlier in this guide for further detail.

```
"Error: valid_http_root: Default IP address * not allowed for LifeKeeper
protection. Check the Apache configuration file at $CONF_FILE."
```

You must specify at least one specific LifeKeeper protected IP address for each Apache instance.

```
"Error: valid_http_root: A Listen directive is being used which specifies
an IP address but no port. Check the Apache configuration file at $CONF_FILE."
```

The correct syntax for the Listen directive is Listen [IPAddress:] port number. This is not caught as a syntax error by Apache, but is interpreted incorrectly (as though the first number in the IP address was a port number specification).

```
"Error: valid_http_root: IP address $ip is not LifeKeeper protected."
```

The Apache configuration file refers to an IP address or domain name not configured under LifeKeeper protection. You must create these LifeKeeper IP address resources in advance.

## During Apache Resource Hierarchy Creation

```
"Error: Could not find IP resource for $IP_ADD on machine $MACH."
```

You must create this resource before the Apache resource creation will succeed.

```
"Error: Create Apache file system hierarchy failure for filesystem"
```

\$FSNAME used by server root \$HTTP\_ROOT.”

“Error: Failure bringing Apache Resource \$TAG into service on machine \$MACH.”

Check the Apache error logs for messages (default location is /var/log/httpd/error\_log, but other logs may be listed).

The most likely cause of this problem is an error in the Apache configuration file. You may be able to bring this resource into service manually after correcting the problem.

“LifeKeeper: RESTORE: \*ERROR\* Apache: The instance is Password Protected.”

The LifeKeeper Apache Web Server Recovery Kit cannot support password protected Private Key files for SSL-enabled web servers, since this would require manual interaction each time Apache starts up, and would prevent automatic restart and failover. The section Specific Configuration Considerations for Apache Web Server in this document explains how to remove password protection from the Private Key file (specified by the SSLCertificateKeyFile directive). This message applies only in an environment where the SSL module is used with Apache.

## 6.1.5.2. Apache Extend Hierarchy Errors

---

The error messages that might be displayed during Apache hierarchy extension are listed below, along with a suggested explanation for each. Note that these error messages appear when the GUI indicates it is “Executing the pre-extend script...” to validate the hierarchy prior to extending it to the new system.

Each will be preceded by an error message like:

*“Error – canextend(template\_server, tag, app\_type/resource\_type, target\_server) -”.*

Each will be followed by an error message like:

*“Error – extmgr(template\_server, tag, target\_tag, target\_server) -”.*

### During Validation of Web Server Binary Location

See errors listed for validation of Web Server Binary Location under [Hierarchy Creation Errors](#).

### During Validation of the Apache Configuration File on the Target System

*“Cannot find Apache configuration file at \$CONF\_FILE on \$TARGET\_SYS.”*

Must have a valid Apache configuration file at conf/httpd.conf relative to Server Root directory specified.

*“DocumentRoot directory “\$doc” in \$CONF\_FILE on \$TARGET\_SYS was not found in the configuration file on \$TEMPLATE\_SYS.”*

or

*“DocumentRoot directory “\$doc” in \$CONF\_FILE on \$TEMPLATE\_SYS was not found in the configuration file on \$TARGET\_SYS.”*

While comparing the configuration files on target and template servers, one or more DocumentRoot directories were found which do not match between the two. Check the details of the error messages displayed to determine the differences between the two. Note that if a DocumentRoot directory path is typed incorrectly, you will generally see both of these error messages, since each configuration file will appear to have an entry not in the other file.

*“IP:port combination “\$ipp” in \$CONF\_FILE on \$TARGET\_SYS was not found in the configuration file on \$TEMPLATE\_SYS.”*

or

*“IP:port combination “\$ipp” in \$CONF\_FILE on \$TEMPLATE\_SYS was not found in the configuration file*

on `$TARGET_SYS`.”

While comparing the configuration files on target and template servers, one or more IP/port combinations were configured for use on one server but not on the other. Note that the IP/port combinations used may be specified in terms of IP addresses, ports, and domain names using a variety of Apache configuration directives. It is the actual IP/port combinations used which are compared, not the directives used to specify them. Check the details of the error messages displayed to determine the differences between the two.

*“SSLCertificateKeyFile “\$file” in \$CONF\_FILE on \$SYS1 was not found in the configuration file on \$SYS2.”*

The filename specified for the SSLCertificateKeyFile in the Apache configuration file on the target system does not match the one specified on the template system. These configurations must be identical. This message applies only in an environment where the SSL module is used with Apache.

*“Apache SSLCertificateKeyFile exists on \$SYS1 but not on \$SYS2.”*

The SSLCertificateKeyFile specified in the Apache configuration files exists on one system, but not on the other. The file must be present on both nodes. This message applies only in an environment where the SSL module is used with Apache.

*“WARNING: PHP configuration file \$PHP\_CONFIG appears to be different on \$SYS1 and \$SYS2.”*

The configuration file for the PHP module on the target system is not identical to the one on the template system. Inspect the configuration on both servers to ensure that they are the same. This message applies only in an environment where the PHP module is used with Apache.

## **During Apache Resource Hierarchy Creation on Target Server**

See errors listed for Apache resource hierarchy creation under [Hierarchy Creation Errors](#).

## 6.1.5.3. Apache Hierarchy Restore, Remove, and Recover Messages and Errors

---

The following information and error messages are printed to the LifeKeeper error log.

They may be viewed by typing “lk\_log log”.

### Bringing an Apache Resource In Service (Restore)

*“LifeKeeper: RESTORE: APACHE: RESTORING \$TAG TO SERVICE START AT: <date>”*

**Informational message.** Records when the restore begins. Logged at the start of every restore.

*“LifeKeeper: RESTORE APACHE RESOURCE \$TAG END err=\$err AT: <date>”*

**Informational message.** Records when the restore completes. Logged at the end of every restore. If any errors occur during the restore, additional messages will be logged between these two messages and the value displayed for err=\$err will be non-zero.

*“Apache: No instance information found for Tag=\$TAG.”*

**Error:** Indicates no instance is defined with the tag value passed to the “restore” script. Unlikely to occur with the GUI, since only tags known to LifeKeeper are available as choices for the In Service and Out of Service actions.

*“LifeKeeper: RESTORE: Apache: Tag=\$TAG already running.”*

**Informational message.** Indicates that the instance appeared to already be up and running.

*“LifeKeeper: RESTORE: Apache: Existing processes terminated for ID=\$ID.”*

**Informational message.** Existing httpd processes were found running for this instance ID, but the PidFile is either missing or invalid. Therefore, the running processes were terminated.

*“LifeKeeper: RESTORE: Apache: Invalid PidFile=\$PIDFILE has been deleted.”*

**Informational message.** An existing PidFile was found for this instance, but its contents were invalid. Therefore, the PidFile was deleted.

*“LifeKeeper: RESTORE: Apache: Tag=\$TAG is being restarted.”*

**Informational message.** Indicates that the instance appeared to partially running, but needed to be restarted. If a PidFile still exists (which contains the process ID of the parent httpd process), the instance is restarted with a HUP signal. If the PidFile is missing, the instance is completely stopped and restarted.

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Error in web server configuration file \$CONF\_FILE for instance \$ID.”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Execute the following command to check the syntax of this file:”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: \$HTTPD\_PATH -t -d \$SERVER\_ROOT -f \$CONF\_FILE.”*

Prior to instance startup, the syntax of the configuration file is checked using the httpd -t option. The -d option checks the ServerRoot directory. Additional options related to modules may also be displayed if you have configured Apache to use modules. Any syntax errors caught during hierarchy creation are displayed in the LifeKeeper GUI, but syntax errors introduced later will not be displayed in the GUI or the LifeKeeper logs. You must manually run the following command to determine what is wrong with your configuration (add additional options for modules, if applicable):

```
$HTTPD_PATH -t -d $SERVER_ROOT -f $CONF_FILE
```

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Error starting web server instance \$INSTANCE.”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Restore of tag \$TAG failed.”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Examine the Apache error log at \$ERROR\_LOG”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: to determine the cause of the problem.”*

An error occurred executing the httpd daemon with the parameters specified. Check the httpd executable being used, configuration file, and general configuration for possible problems.

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Web server instance \$ID did not start correctly.”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Restore of tag \$TAG failed.”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: Examine the Apache error log at \$ERROR\_LOG”*

*“LifeKeeper: RESTORE: \*ERROR\* Apache: to determine the cause of the problem.”*

Note that in many cases the httpd daemon will appear to start even if its web sites don't respond as expected. The restore script checks all IP/port combinations used to make sure all sites configured are fully functional. If they are not, this message is printed and the restore fails.

Although the site is left in the Out of Service state, one/more httpd processes may still be left running. (This is intentional, since one/more web sites may be operational and we don't want to kill them off). You should resolve the problem as soon as possible and bring the instance In Service. If you don't, LifeKeeper will eventually attempt to recover the instance and restore it to service automatically. If it can't, it will fail over the hierarchy to another server.

## Taking an Apache Resource Out of Service (Remove)

*“LifeKeeper: REMOVE: APACHE: REMOVE \$TAG FROM SERVICE START AT: <date>”*

**Informational message.** Records when the remove begins. Logged at the start of every remove.

*“LifeKeeper: REMOVE APACHE RESOURCE \$TAG END err=\$err AT: <date>”*

**Informational message.** Records when the remove completes. Logged at the end of every remove.

If any errors occur during the remove, additional messages will be logged between these two messages and the value displayed for err=\$err will be non-zero.

*“LifeKeeper: REMOVE: \*WARNING\* APACHE: Error attempting to kill parent process for INSTANCE=\$INSTANCE.”*

There was an error attempting to kill the parent httpd process (whose process ID is stored in the Pidfile).

*“LifeKeeper: REMOVE: \*ERROR\* APACHE: Error attempting to kill all processes for INSTANCE=\$INSTANCE.”*

Although the parent httpd process appeared to be killed successfully, one/more processes for this instance are still running. Normally the remove will be able to terminate any/all processes for this

instance. When it cannot, this message is printed and the remove fails.

## Bringing an Apache Resource Back In Service (Recover)

The LifeKeeper core periodically checks the health of every Apache instance In Service on the local server by running an Apache “quickCheck” script, which checks the web sites using the same scripts used to check the state of the instance during restore and remove. If the instance is not fully functional, a “recover” script is invoked to attempt to restart the instance. This simply logs the first message shown below, invokes “restore,” prints the final error or success message shown below—depending on error or success of the “restore” script—and returns the same result as “restore.” If restore/recover fails, this instance is failed over to another server.

*“LifeKeeper: RECOVER: APACHE: Invoking restore for Apache instance “\$ID” at: <date>”*

*“LifeKeeper: RECOVER: APACHE: Restore for Apache instance \$ID returned error \$RET at: <date>”*

*“LifeKeeper: RECOVER: APACHE: Restore for Apache instance \$ID successful at: <date>”*

## 6.2. DB2 Recovery Kit Administration Guide

---

The LifeKeeper for Linux DB2 Recovery Kit provides fault resilient protection for DB2 database instances. LifeKeeper, together with the DB2 Universal Database product family afford increased availability to DB2 operating environments by effectively recovering database server failures without significant down-time or human intervention.

### Document Contents

This guide contains the following topics:

- [Documentation and References](#). A list of LifeKeeper for Linux documentation and where to find them.
- [Requirements](#). A description of the hardware and software necessary to properly setup, install, and operate the DB2 Recovery Kit. Refer to [LifeKeeper Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper for Linux software.
- [Overview](#). A description of the DB2 Recovery Kit's features and functionality.
- [Configuring the LifeKeeper for Linux DB2 Recovery Kit](#). A description of the procedures required to properly configure the DB2 Recovery Kit.
- [LifeKeeper Configuration Tasks](#). A description of the tasks for creating and managing your DB2 resource hierarchies using the LifeKeeper GUI.
- [Troubleshooting](#). A list of LifeKeeper for Linux error messages including a description for each.
- [Appendix](#). Steps for setting up DB2 to use raw I/O.

## 6.2.1. DB2 Documentation and References

---

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper Installation Guide](#)
- [Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is available on the SIOS Technology Corp. website at:

<http://docs.us.sios.com>

## 6.2.2. DB2 Recovery Kit Hardware and Software Requirements

---

Your LifeKeeper configuration must meet the following requirements **prior** to the installation of the LifeKeeper for Linux DB2 Recovery Kit. Please see [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the configuration of your LifeKeeper hardware and software.

### Hardware Requirements

- **Servers** – LifeKeeper for Linux supported servers configured in accordance with the requirements described in [LifeKeeper for Linux Installation Guide](#) and [LifeKeeper for Linux Release Notes](#).
- **IP Network Interface Cards** – Each server requires at least one Ethernet TCP/IP-supported network interface card. Remember, however, that a LifeKeeper cluster requires two communications paths; two separate LAN-based communication paths using dual independent sub-nets are recommended for heartbeats, and at least one of these should be configured as a private network. Using a combination of TCP and TTY heartbeats is also supported.

### Software Requirements

- **TCP/IP software** – Each server in your LifeKeeper configuration requires TCP/IP software.
- **IBM software** – Please refer to [LifeKeeper for Linux Release Notes](#) for specific DB2 version requirements on certain Linux distributions and hardware architectures.
- **LifeKeeper Software** – It is imperative that you install the same version of the LifeKeeper software and apply the same versions of the LifeKeeper software patches to each server in your cluster.
- **LifeKeeper for Linux DB2 Recovery Kit** – The DB2 Recovery Kit is provided on a CD. It is packaged, installed and removed via the Red Hat Package Manager, rpm. The following rpm file is supplied on the LifeKeeper for Linux DB2 Recovery Kit CD:

#### **steeleye-1kDB2**

Please see [LifeKeeper for Linux Installation Guide](#) for specific instructions on the installation and removal of the LifeKeeper for Linux software.

- **LifeKeeper for Linux NFS Recovery Kit-required for use of DB2 EEE and multiple partition ESE deployments.** This recovery kit is provided on CD in the **steeleye-1kNFS** package.

 **Important:** See [Issues Regarding DB2 EEE or multiple partition ESE and NFS](#) for important configuration information.

## 6.2.3. DB2 Recovery Kit Overview

---

### LifeKeeper for Linux DB2 Recovery Kit

In versions 8 and greater, DB2 UDB Enterprise Edition (EE) and Enterprise-Extended Edition (EEE) have been combined into a single product named DB2 UDB Enterprise Server Edition (ESE). Previous versions included two separate enterprise level database servers, the Enterprise Edition (EE) as a standard relational database management system and the Enterprise-Extended Edition (EEE) as an extension of the EE database server to support multi-partition databases.

The LifeKeeper for Linux DB2 Recovery Kit provides protection for the database manager in the EE, WE, and WSE environments, and for the database partition servers in an EEE environment. In a combined ESE environment, the recovery kit provides protection for both the database manager and the database partition servers.

Users may elect to define the DB2 Administration Server for each machine within the LifeKeeper cluster. When the DB2 Administration server is defined, LifeKeeper will attempt to start the DB2 Administration Server as a function of the DB2 hierarchy create and the DB2 hierarchy restore operations.

## 6.2.4. Configuring the LifeKeeper for Linux DB2 Recovery Kit

---

This section describes the LifeKeeper for Linux DB2 Recovery Kit configuration details. It also contains information you should consider before you start to configure and administer the DB2 Recovery Kit. Please refer to [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring LifeKeeper Core resource hierarchies.

## 6.2.4.1. Using DB2 with Raw I/O

---

If you plan to use DB2 with Raw I/O devices, you must install the LifeKeeper Raw I/O Recovery Kit from the LifeKeeper Core CD. You must also properly set up the Raw I/O devices prior to use. See the [Appendix](#) for instructions.

## 6.2.4.2. Running DB2

---

### Reducing the DB2 Process Startup Times

In some instances the startup times of the DB2 processes can be excessive when using DB2 8.x under LifeKeeper protection. Making the following change to the kernel network parameters can improve this situation. Add the following line to the `_ /etc/sysctl.conf_` file on each LifeKeeper clustered system that will be running DB2 8.x:

```
net.ipv4.tcp_syn_retries=1
```

Then running **sysctl -p** will cause this change to take effect.

### Preventing Frequent DB2 Instance Crashes (Panic)

If a LifeKeeper protected DB2 instance is encountering frequent crashes in a systemd environment (RHEL7, CentOS7, OEL7) then altering the automatic IPC cleanup configuration parameter may correct this issue. On each node in the LifeKeeper cluster, set the following configuration parameter in the `/etc/systemd/logind.conf` file.

```
RemoveIPC=no
```

Then, execute **systemctl restart systemd-logind** to make this change effective. Click [here](#) for more details.

## 6.2.4.3. Configuration Considerations for DB2 Single Partition

---

The following should be considered before operating the LifeKeeper for Linux DB2 Recovery Kit in the single partition or workgroup environment:

1. LifeKeeper requires the location of the DB2 instance home directory as well as associated databases, tablespaces, and resources be stored on shared drives. Replicated (SIOS DataKeeper) file system resources must be created before creating the DB2 resource. The shared drives are automatically protected at the time the hierarchy is created. During creation of the DB2 resource hierarchy, the DB2 database manager is created as the parent resource while the shared file systems containing instance home directories and actual databases are created as dependent resources. Consequently, if **after** the creation of your DB2 hierarchy you decide to create a database on a shared file system that is not protected by LifeKeeper, you will need to create a resource hierarchy for that file system and make it a dependency of your DB2 resource hierarchy.
2. When the database manager becomes inoperable on the primary system, the service fails over to a previously defined backup system. The database service on the backup system becomes available immediately after the dependent resources fail over and the database manager is brought into service. Previously connected DB2 clients are disconnected and must reconnect to the functioning server. Any uncommitted SQL statements are rolled back and should be re-entered.

## 6.2.4.4. Configuration Considerations for DB2 Multiple Partition

---

**DB2 Multiple Partition RESTRICTIONS:** All DB2 multiple database partition servers will be protected on a particular machine when the LifeKeeper DB2 resource hierarchy is created on that machine. The nodes to protect are determined by examining the following file:

*<instance home>/sqllib/db2nodes.cfg*

Future plans for this recovery kit include added functionality to allow for N-way failover.

## 6.2.4.4.1. Issues Regarding DB2 EEE or multiple partition ESE and NFS

---

If the NFS export point for the DB2 instance home directory becomes unavailable while the DB2 instances are running, the system will hang while waiting for the export point to become available again. Many system operations will not work correctly, including a system reboot. You should be aware that the NFS server for the DB2 multiple partitions cluster should be protected by LifeKeeper and should not be manually taken out of service unless all the partitions in the DB2 cluster are also taken out of service before shutting down the NFS resource. Additionally, the DB2 partitions cannot be brought into service unless the NFS resource is in service.

To avoid accidentally causing your cluster to hang by inadvertently stopping the NFS server, we make the following recommendations:

### NFS Recommendations

**Use additional servers:** It is highly recommended that you have a separate cluster for the NFS export point from which the DB2 instance home is mounted. The NFS export point on this cluster should be protected with the LifeKeeper NFS Server Recovery Kit.

If you do not have at least two additional servers available, you can reduce the chances of experiencing the problem described above by adding one additional server to the DB2 cluster. This additional server would export the NFS hierarchy. One of the other nodes in the cluster would serve as a backup. In this configuration the symptoms could occur if the NFS hierarchy were to failover to the backup node. The NFS export point on this cluster should be protected with the LifeKeeper NFS Server Recovery Kit.

**If you cannot use additional servers:** This is the least desirable option. However, if you decide to run your NFS server in the same cluster as your DB2 multiple partitions, the NFS export point should be protected with the LifeKeeper NFS Server Recovery Kit. You should note that LifeKeeper currently is not aware of the relationship between the DB2 partitions and the NFS server managing the DB2 partitions. Therefore, you must follow these manual procedures before stopping or starting LifeKeeper on any node in the cluster.

1. If you wish to stop LifeKeeper on a single server, you must make sure that the NFS server is active on another server in the cluster. Failure to do this may cause the LifeKeeper shutdown to hang trying to take the DB2 partitions out of service. Generally, you should make sure that all DB2 partitions are either switched to another server or manually taken out of service before you stop LifeKeeper to ensure you don't have problems trying to restart LifeKeeper.
2. To shut down the entire cluster, you should manually take all DB2 partition resources out of service. Next, take all the DB2 NFS server resources out of service, and finally shut down LifeKeeper.
3. If you remembered to take the DB2 resource out of service before shutting down LifeKeeper, you should be able to restart LifeKeeper normally. Then bring the NFS server resources into service, followed by any DB2 partitions you wish to restart.

4. If you forgot to take the DB2 partition out of service before shutting down LifeKeeper, you must make sure that the NFS server resources for that partition are active elsewhere in the cluster before you restart LifeKeeper.

## 6.2.4.4.2. DB2 Configuration Requirements

To ensure proper operation of the DB2 Recovery Kit in a multiple partition environment, LifeKeeper requires the following:

1. If you cannot use an additional cluster for your NFS hierarchy, be aware that the LifeKeeper for Linux DB2 Recovery Kit restricts the occurrence of active inodes on an underlying NFS-protected file system. Therefore, to prevent this condition, we recommend that users protect the top-level directory and export the instance home directory using the fully qualified directory name. The top-level directory is protected in order to prohibit users from changing directories directly into it (i.e. `cd<top level dir>`).
2. Verify the installation of IBM's latest Fix Pack (for EEE deployments) as described in the Software Requirements section of this document.
3. Ensure that the hostname value in your `db2nodes.cfg` file is the same as the value returned from issuing the **hostname** command.

Example:

`db2nodes.cfg` file:

```
0 server1.sc.steeleye.com 0
```

Additionally, the hostname value in your server's `/etc/hosts` file must be the same as the hostname value in your `db2nodes.cfg` file.

You must also verify that your server's `/etc/hosts` file contains both the local hostname and the fully qualified hostname for each server entry included in the file.

Example:

`/etc/hosts` file

```
127.0.0.1 localhost localhost.localdomain
```

```
9.21.55.53 server1.sc.steeleye.com server1
```

4. During execution of the `db2setup` script, **do not** opt to create the DB2 Warehouse Control Database (DWCNTRL) or the DB2 Sample Database at this time. The databases need to be created on a shared file system to ensure successful creation of the DB2 resource hierarchy. Electing to create either of these databases during execution of the `db2setup` script will cause the database to be created in the home directory and not on a shared file system. Users wishing to create these databases should do so external to the `db2setup` script in order to specify a shared file system.

In versions later than 8.1, the DB2 Tools Catalog should not be created during the setup script.

This database must be placed on a shared file system and should be created after setup has completed and prior to hierarchy creation, if necessary.

5. Active/Active or multiple partition server environments, each server in the configuration must be capable of running all database instances in a failover scenario. Please see the *IBM Getting Started Guide* for help determining the maximum number of DB2 instances or partition servers feasible for a given set of system resources.
6. Select or create a shared file system, then export this file system. (*i.e /export/db2home*). The file system will be used as the DB2 instance home. (Note: one exception from this requirement is for partitions that will all run on the same server at all times. In that case, no NFS export is necessary, and the instance homes can simply be located on shared storage.)
7. Protect your exported file system by creating a LifeKeeper NFS resource hierarchy. The file system should be included as a dependent resource in your NFS hierarchy.
8. NFS mount the shared file system on each server in the cluster including the server where it is being exported. See the *DB2 Quickstart Guide* for mount options. When creating the DB2 instance, the home directory of the instance must be located on the NFS mounted file system. Make certain that the file system is mounted using the LifeKeeper protected switchable IP address used when creating the NFS hierarchy. Additionally, the mount point of the home directory must be specified in the */etc/fstab* file on all servers in your LifeKeeper cluster. Each server in your configuration must have the file system mounted on identical mount points (*i.e. /db2/home*).

**Note:** We recommend that you create and test your NFS hierarchy prior to creating your DB2 resource hierarchy. Please see the [NFS Recovery Kit Administration Guide](#) for complete instructions on creating and testing a NFS hierarchy.

9. For all servers in your configuration, set the following DB2 environment variable to equal the total number of partitions in the instance. To set this variable, log on to the server as the instance owner and issue a **db2set** command. Adjusting this variable will accommodate all conceivable failover scenarios.

**db2setDB2\_NUM\_FAILOVER\_NODES=<partitions in the instance>**

10. Update your existing DB2 instances and your DB2 Administration servers using the following DB2 utilities:

*db2iupdt and dasiupdt*

11. A LifeKeeper DB2 hierarchy must be created on each server in the cluster that has a database partition server managing data for the instance. The databases and tablespaces must be on a shared file system. A separate LUN is required for each database partition server and for the NFS exported home directory. Dependent resources include the file systems where actual databases and tablespaces are located.
12. If you create a database on a non-protected LifeKeeper file system after the creation of your DB2 hierarchy, you will need to create a resource hierarchy for that file system and make it a

dependency of your DB2 resource hierarchy. The hierarchy will protect all of the partition servers that the *db2node.cfg* file indicates should run on the server.

13. To ensure proper execution of a failover, it is imperative that the file system of each database partition server is uniquely numbered.

Example:

The mount point for your database partition server *node0* should be:

**`/<FSROOT>/<db2instancename>/NODE0000`**

The mount point for your database partition server *node1* should be:

**`/<FSROOT>/<db2instancename>/NODE0001`**

**Note:** In this example there are two partition servers, and the file system for each is mounted on a separate LUN.

14. All database partition servers for a given machine must be running in order to assure the successful creation of your DB2 hierarchy.
15. When the database partition server becomes inoperable on the primary system, the service fails over to a previously defined backup system. The database service on the backup system becomes available immediately after the dependent resources fail over and the database partition server(s) is brought into service. Previously connected DB2 clients are disconnected and must reconnect to the functioning server. Any uncommitted SQL statements are rolled back and should be re-entered.

## 6.2.4.5. Configuration Considerations for All DB2 Configurations

---

1. DB2 instance names should contain alphanumeric characters only.
2. DB2 clients should be configured to connect to the database via a LifeKeeper protected IP address. Users can define:

**DB2SYSTEM=<Floating IP>” in \$instancehome/sqllib/profile.env**

and catalog the floating IP address on the clients.

3. The `/etc/services` file for each server in your configuration protecting a DB2 resource hierarchy must have identical service entries for the protected instance. Additionally, the User ID, Group ID and instance home directory for the protected DB2 instance must be the same on all servers where the resource will be protected.

DB2 adds the following entries as default in `/etc/services`:

*DB2\_db2inst1 60000/tcp*

*DB2\_db2inst1\_1 60001/tcp*

*DB2\_db2inst1\_2 60002/tcp*

*DB2\_db2inst1\_END 60003/tcp*

*db2c\_db2inst1 50001/tcp*

4. A recovery is what takes place after DB2 is terminated abruptly, as with a system crash. Following are tips that will significantly reduce the amount of time it takes for DB2 to recover from a failure.
  - Limit the log records that DB2 will process. You can accomplish this by properly configuring the **SOFTMAX** and **LOGFILSIZ** configuration parameters. You should use log files with a size of 4MB (1000 4KB pages) and keep the amount of active log space at 25% of the size of one log file (1MB):

**db2 UPDATE DB CFG FOR <dbname> USING SOFTMAX 25**

**db2 UPDATE DB CFG FOR <dbname> USING LOGFILSIZ 1000**

- Ensure that there is a sufficient number of page cleaners to accommodate your work load:

**db2 UPDATE DB CFG FOR <dbname> USING NUM\_IOCLEANERS <num>**

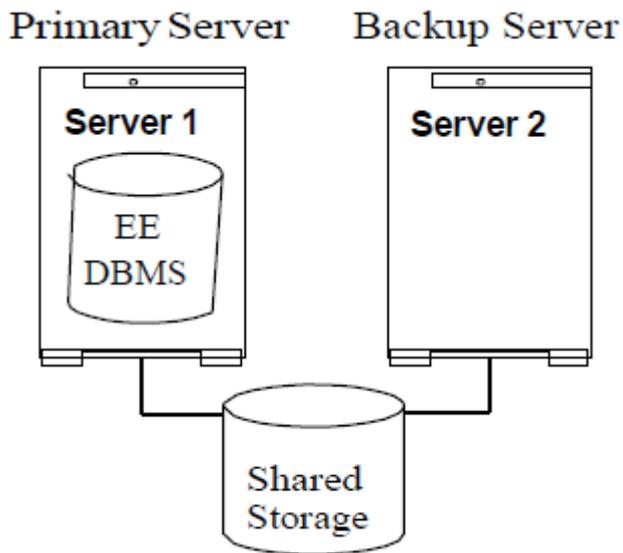
5. DB2 Fault Monitor should be disabled in all servers.

6. DB2 should be installed in all servers.

## 6.2.4.6. DB2 Configuration Examples

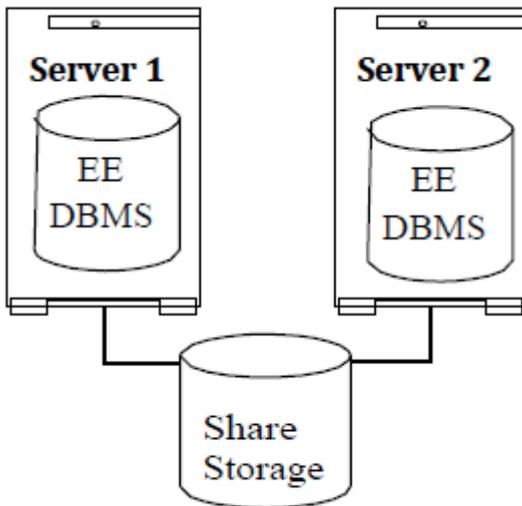
A few examples of what happens during a failover using LifeKeeper for Linux DB2 Recovery Kit are provided below. In the following pictures, EE and EEE are used to denote database configurations; ESE may be substituted wherever appropriate.

### Configuration 1: DB2 Single Partition Active/Standby Configuration



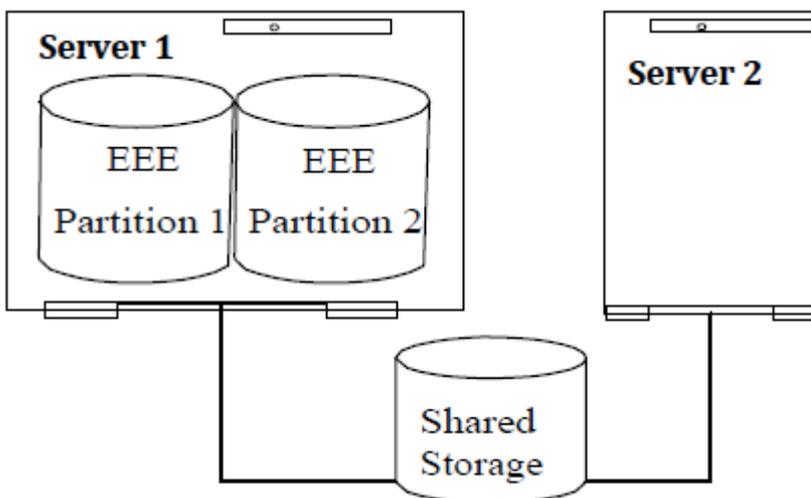
The DB2 instance is protected on Server 1. Server 2 will assume the DB2 resources when a failure occurs.

## Configuration 2: DB2 Single Partition Active/Active Configuration



One DB2 instance is protected on Server 1 and another DB2 instance is protected on Server 2. Each server will assume the other’s resources when a failure occurs.

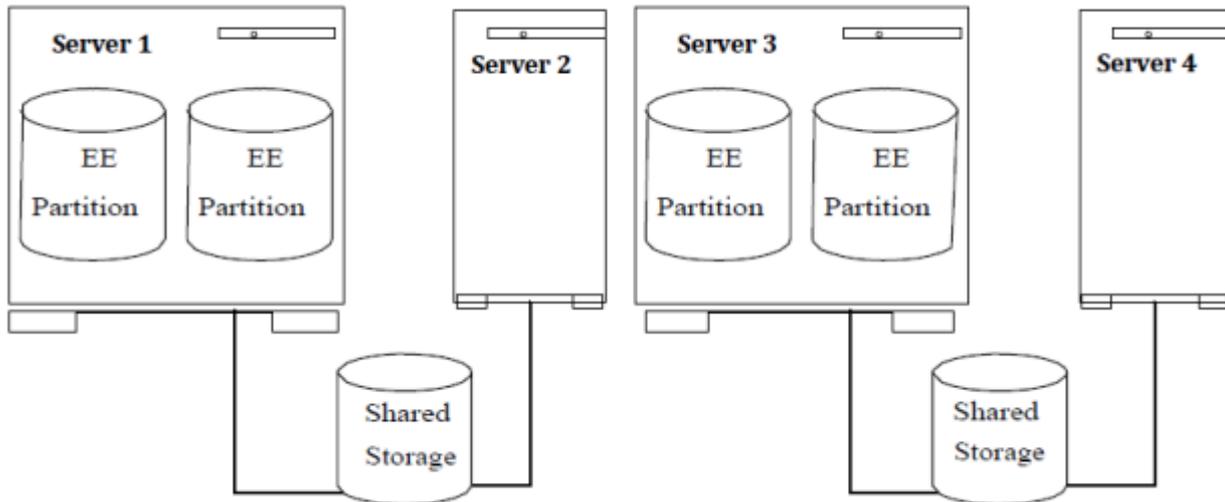
## Configuration 3: DB2 Multiple Partition Active/Standby (1 Cluster)



One DB2 instance with two database partition servers is protected on Server 1 with one LifeKeeper DB2 resource hierarchy. Server 2 will assume ownership of the DB2 resource hierarchy when a failure occurs.

**Note:** For all cluster of cluster configurations listed in the following section, users should be aware that the cluster of cluster configuration is protecting only one DB2 instance with multiple partitions on multiple physical nodes.

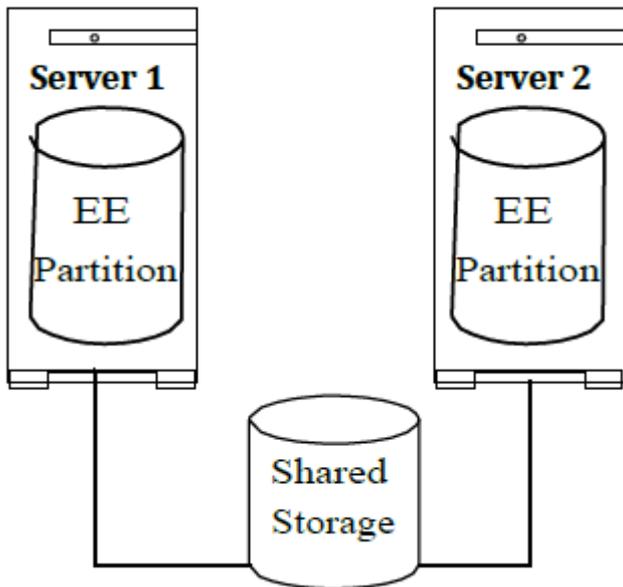
## Configuration 4: DB2 Multiple Partition Active/Standby (Cluster of Clusters)



One DB2 instance with two database partition servers is protected on Server 1 and two database partition servers protected on Server 3. There is one LifeKeeper DB2 resource hierarchy on Server 1, extended to Server 2, and another DB2 resource hierarchy on Server 3 extended to Server 4. When a failure occurs on Server 1, Server 2 will assume its resource. When a failure occurs on Server 3, Server 4 will assume its resource.

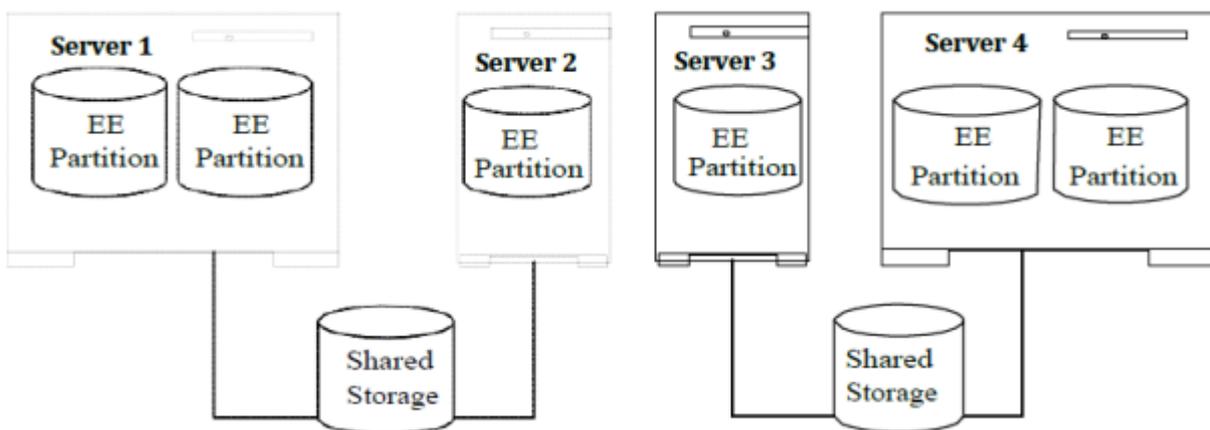
If the server that is exporting the DB2 instance home directory and its backup server become inoperable at once, the DB2 database is inaccessible. In addition, if the NFS hierarchy for the exported DB2 instance directory (primary and all backups) become inoperable at the same time, the DB2 database will be inaccessible until the NFS hierarchy can be restored.

## Configuration 5: DB2 Multiple Partitions Active/Active (1 Cluster)



One DB2 instance with one database partition server is protected on Server 1 and one database partition server protected on Server 2. There is one LifeKeeper DB2 resource hierarchy on Server 1 and another DB2 resource hierarchy on Server 2. When a failure occurs each server will assume the other's resources.

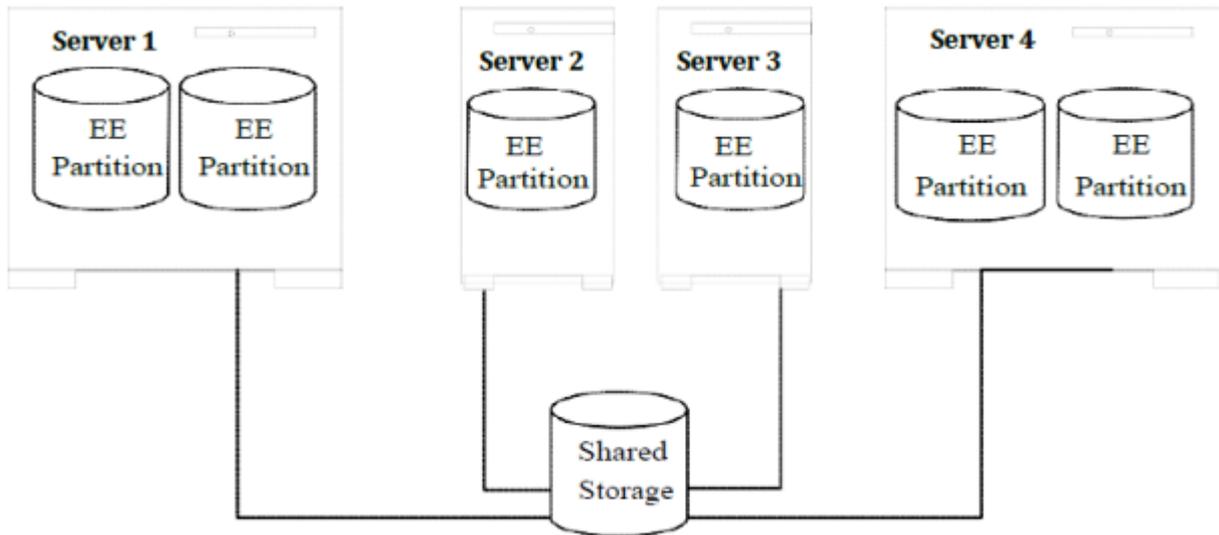
## Configuration 6: DB2 Multiple Partitions Active/Active (Cluster of Clusters)



One DB2 instance with two database partition servers is protected on Server 1, one database partition server protected on Server 2, one database partition server protected on Server 3 and two database partition servers protected on Server 4. There is one LifeKeeper DB2 resource hierarchy on each server in the cluster. Upon failure, Server 1 and Server 2 assume each other's resources and Server 3 and Server 4 assume each other's resources.

If the server that is exporting the DB2 instance home directory and its backup server become inoperable at once, the DB2 database is inaccessible. In addition, if the NFS hierarchy for the exported DB2 instance directory (primary and all backups), become inoperable at the same time, the DB2 database will be inaccessible until the NFS hierarchy can be restored.

## Configuration 7: DB2 Multiple Partition (4 Node Fibre Channel Cluster)



One DB2 instance with two database partition servers is protected on Server 1, one database partition server protected on Server 2, one database partition server protected on Server 3 and two database partition servers protected on Server 4. There is one LifeKeeper DB2 resource hierarchy on each server in the cluster. Each server in the cluster provides backup protection for the other in the event of failure.

If the server that is exporting the DB2 instance home directory and its backup server become inoperable at once, the DB2 database is inaccessible. In addition, if the NFS hierarchy for the exported DB2 instance directory (primary and all backups), become inoperable at the same time, the DB2 database will be inaccessible until the NFS hierarchy can be restored.

## 6.2.5. LifeKeeper for Linux DB2 Recovery Kit Configuration Tasks

---

You can perform all LifeKeeper for Linux DB2 Recovery Kit administrative tasks via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer, and monitor DB2 resources.

The following tasks are available for configuring the LifeKeeper for Linux DB2 Recovery Kit:

- [Create a Resource Hierarchy](#) – Creates a DB2 resource hierarchy. Replicated (SIOS DataKeeper) file system resources must be created before creating the DB2 resource.
- [Delete a Resource Hierarchy](#) – Deletes a DB2 resource hierarchy.
- [Extend a Resource Hierarchy](#) – Extends a DB2 resource hierarchy from the primary server to the backup server.
- [Unextend a Resource Hierarchy](#) – Unextends (removes) a DB2 resource hierarchy from a single server in the LifeKeeper cluster.
- [Create Dependency](#) – Creates a child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete Dependency](#) – Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#) – Activates a resource hierarchy.
- [Out of Service](#) – Deactivates a resource hierarchy.
- [View](#) / [Edit](#) Properties – View or edit the properties of a resource hierarchy.

## 6.2.5.1. Creating a DB2 Resource Hierarchy

Perform the following on your primary server:

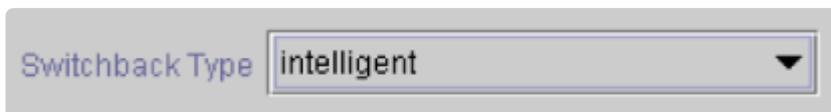
1. Select **Edit > Server > Create Resource Hierarchy**.
2. The “**Select Recovery Kit**” dialog appears. Select the **DB2 Database** option from the drop down list.



Click **Next** to continue.

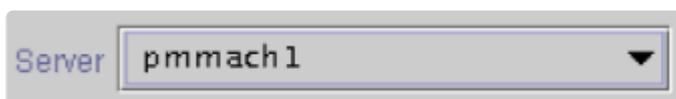
**CAUTION:** If you click the **Cancel** button at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

3. The “**Switchback Type**” dialog appears. The switchback type determines how the DB2 resource will be switched back to the primary server when it becomes in-service (active) on the backup server following a failover. Switchback types are either *intelligent* or *automatic*. Intelligent switchback requires administrative intervention to switch the resource back to the primary server while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.



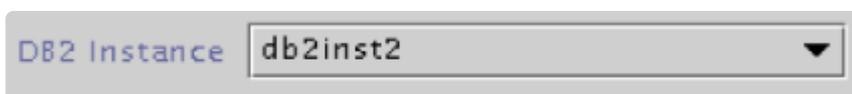
Click **Next** to continue.

4. The “**Server**” dialog appears. Select the name of the server where the DB2 resource will be created (typically this is your primary server). All servers in your cluster are included in the drop down list box.



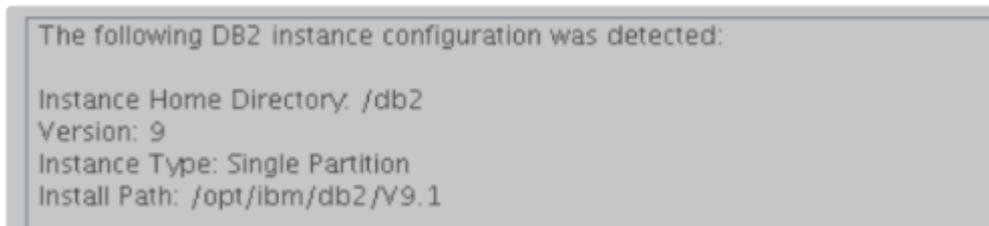
Click **Next** to continue.

5. The “**DB2 Instance**” dialog appears. Select or enter the name of the **DB2** instance that is being protected.



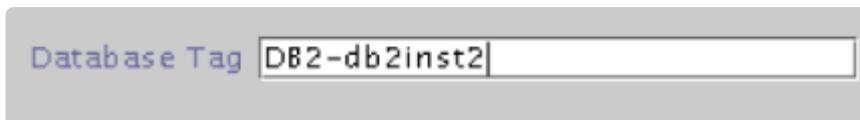
Click **Next** to continue.

6. An information box appears displaying information regarding the instance detected.



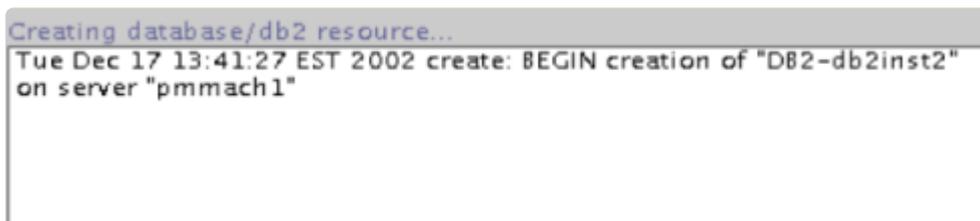
Click **Continue**.

7. The “**Database Tag**” dialog appears. This dialog is populated automatically with a unique tag name for the new DB2 database resource instance.



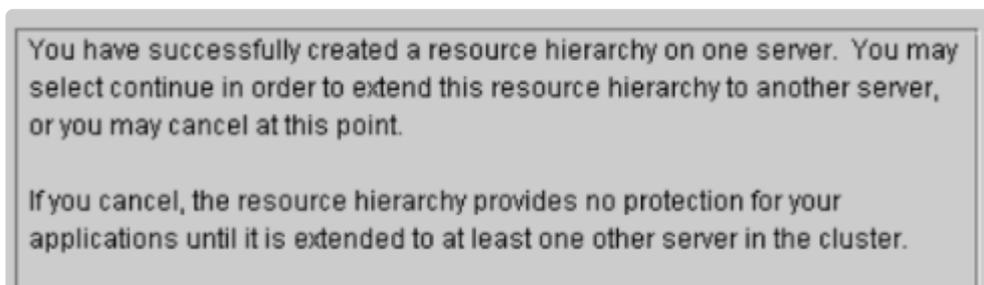
Click **Create** to continue.

8. An information box appears indicating the start of the hierarchy creation.



Click **Next** to continue.

9. An information box appears announcing the successful creation of your DB2 resource hierarchy. You must **Extend** the hierarchy to another server in your cluster in order to place it under LifeKeeper protection.



Click **Continue** to extend the resource.

Click **Cancel** if you wish to extend your resource at another time.

**Verifying Integrity of Extended Hierarchy...****Hierarchy Verification Finished**

**WARNING:** Your hierarchy exists on only one server. Your  
**WARNING:** application has no protection until you extend it  
**WARNING:** to at least one other server.

10. Click **Done** to exit the Create Resource Hierarchy menu selection.

## 6.2.5.2. Deleting a DB2 Resource Hierarchy

To delete a DB2 resource from all servers in your LifeKeeper configuration, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your DB2 resource hierarchy.

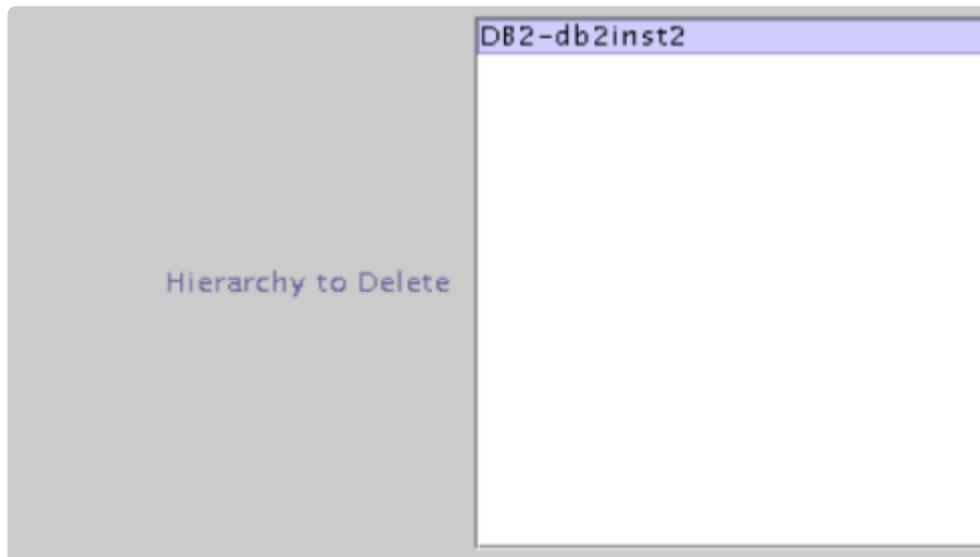
**Note:** If you selected the Delete Resource task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

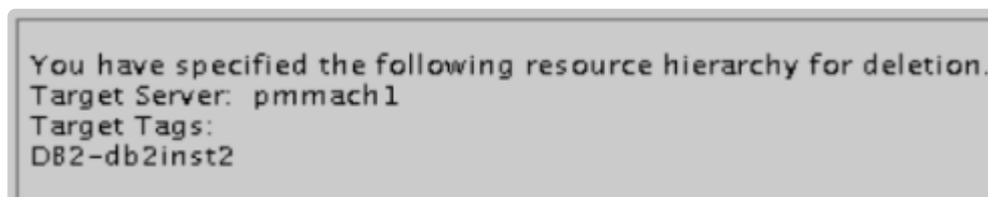
3. Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete, and highlight it.

**Note:** If you selected the Delete Resource task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete.



Click **Delete** to continue.

5. An information box appears confirming that the DB2 resource instance was deleted successfully.

```
Deleting resource hierarchy...  
Successfully removed  
ins_remove[701,Iraci.C]Thu Jun 1 07:06:54 EDT 2000:  
    fletch,priv_globact(1,delete): Running Post Global delete  
    Machine cornfed  
ins_remove[714,Iraci.C]Thu Jun 1 07:06:56 EDT 2000:  
    fletch,priv_globact(1,delete): Post Global delete Scripts F  
    Exiting 0 On Machine cornfed With Output Following:  
lcdrecovery[701,Iraci.C]Thu Jun 1 07:12:15 EDT 2000:
```

6. Click **Done** to exit the Delete Resource Hierarchy menu selection.

## 6.2.5.3. Extending Your DB2 Resource Hierarchy

---

After you have created a hierarchy, you should extend that hierarchy to another server in the cluster. There are three possible ways to extend your resource instance:

1. When you successfully create your DB2 resource hierarchy you will have an opportunity to select **Continue** which will allow you to proceed with extending your resource hierarchy to your backup server.
2. Right-click on an unextended hierarchy in either the left or right pane on the LifeKeeper GUI.
3. Select the “**Extend Resource Hierarchy**” task from the LifeKeeper GUI by selecting **Edit, Resource, Extend Resource Hierarchy** from the drop down menu. This sequence of selections will launch the Extend Resource Hierarchy wizard. The **Accept Defaults** button that is available for the **Extend Resource Hierarchy** option is intended for the user who is familiar with the LifeKeeper Extend Resource Hierarchy defaults and wants to quickly extend a LifeKeeper resource hierarchy without being prompted for input or confirmation. Users who prefer to extend a LifeKeeper resource hierarchy using the interactive, step-by- step interface of the GUI dialogs should use the **Next** button.

a. The first dialog box to appear will ask you to select the **Template Server** where your DB2 resource hierarchy is currently in service. Remember that the **Template Server** you select now and the **Tag to Extend** that you select in the next dialog box represent an in- service (activated) resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you have selected. The drop down box in this dialog provides the names of all the servers in your cluster.

**Note:** If you are entering the Extend Resource Hierarchy task by continuing from the creation of a DB2 resource hierarchy, this dialog box will not appear because the wizard has already identified the template server in the create stage. This is also the case when you right-click on either the DB2 resource icon in the left pane or right-click on the DB2 resource box in the right pane of the GUI window and choose Extend Resource Hierarchy.



**CAUTION:** If you click the **Cancel** button at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extend hierarchy process. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Click **Next** to continue.

b. Select the **Tag to Extend**. This is the name of the DB2 instance you wish to extend from the template server to the target server. The wizard will list in the drop down box all of the

resources that you have created on the template server.

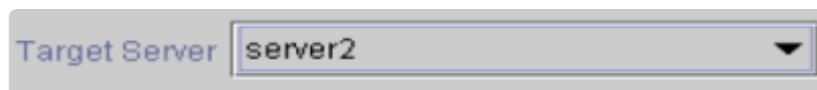
**Note:** Once again, if you are entering the Extend Resource Hierarchy task immediately following the creation of a DB2 hierarchy, this dialog box will not appear because the wizard has already identified the tag name of your resource in the create stage. This is also the case when you right-click on either the DB2 resource icon in the left pane or on the DB2 resource box in the right pane of the GUI window and choose *Extend Resource Hierarchy*.



A screenshot of a GUI dropdown menu. The label 'Tag to Extend' is on the left. The dropdown box contains the text 'DB2-db2inst2' and a downward-pointing arrow on the right side.

Click **Next** to continue.

c. Select the **Target Server** where you will extend your DB2 resource hierarchy.



A screenshot of a GUI dropdown menu. The label 'Target Server' is on the left. The dropdown box contains the text 'server2' and a downward-pointing arrow on the right side.

Click **Next** to continue.

d. The **Switchback Type** dialog appears. The switchback type determines how the DB2 resource will be switched back to the primary server when it becomes in service (active) on the backup server following a failover. Switchback types are either intelligent or automatic. Intelligent switchback requires administrative intervention to switch the resource back to the primary server while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.



A screenshot of a GUI dropdown menu. The label 'Switchback Type' is on the left. The dropdown box contains the text 'intelligent' and a downward-pointing arrow on the right side.

Click **Next** to continue.

e. Select or enter a **Template Priority**. This is the priority for the DB2 hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.

**Note:** This selection will appear only for the initial extend of the hierarchy.

Click **Next** to continue.

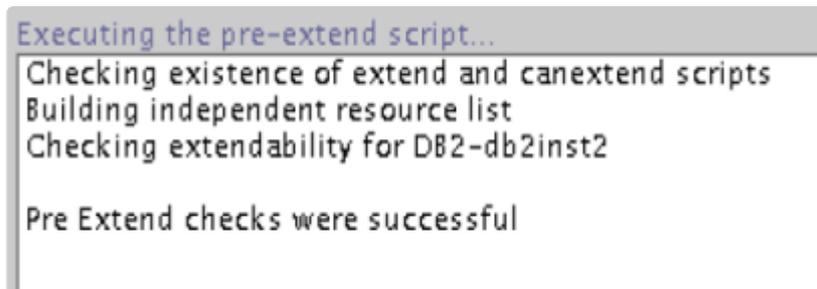
f. Select or enter the **Target Priority**. This is the priority for the new extended DB2 hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities

need not be consecutive, but no two servers can have the same priority for a given resource.

A screenshot of a user interface element. It consists of a light gray rounded rectangle containing the text "Target Priority" in blue, followed by a white text input field with the number "10" inside. To the right of the input field is a small gray square with a white downward-pointing arrow.

Click **Next** to continue.

g. An information box appears explaining that LifeKeeper has successfully checked your environment and that all requirements for extending this resource have been met. If there are requirements that have not been met, LifeKeeper will disable the **Next** button, and enable the **Back** button.

A screenshot of a terminal window. The title bar reads "Executing the pre-extend script...". The terminal text shows the following steps: "Checking existence of extend and canextend scripts", "Building independent resource list", and "Checking extendability for DB2-db2inst2". At the bottom, it states "Pre Extend checks were successful".

Click on the **Back** button to make changes to your resource extension.

Click **Cancel** to extend your resource another time.

Click **Next** to launch the Extend Resource Hierarchy configuration task.

Click **Finish** to confirm the successful extension of your DB2 resource instance.

4. Click **Done** to exit the Extend Resources Hierarchy menu selection.

 **Note:** Be sure to test the functionality of the new instance on *both* servers.

## 6.2.5.4. Unextending Your DB2 Resource Hierarchy

---

1. From the LifeKeeper GUI menu, select **Edit, Resource, and Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the DB2 resource. It cannot be the server where the resource is currently in service (active).

**Note:** If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

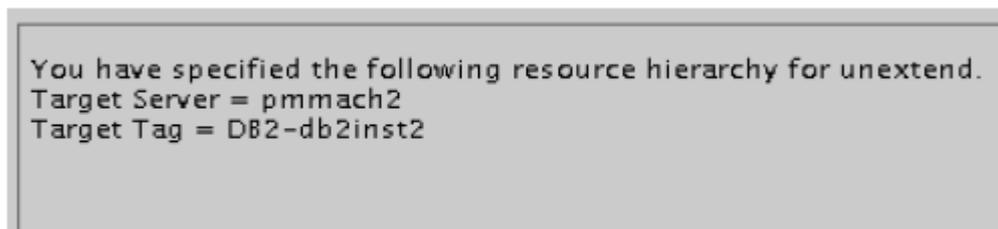
3. Select the DB2 **Hierarchy to Unextend**.

**Note:** If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

4. An information box appears confirming the target server and the DB2 resource hierarchy you have chosen to unextend.



Click **Unextend**.

5. Another information box appears confirming that the DB2 resource was unextended successfully.
6. Click **Done** to exit the Unextend Resource Hierarchy menu selection.

## 6.2.5.5. Testing Your DB2 Resource Hierarchy

---

Test your DB2 resource hierarchy by initiating a manual switchover that will simulate a failover of the resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource, and In Service**. For example, an in-service request executed on a backup server causes the DB2 resource hierarchy to be placed in-service on the backup server and taken out-of-service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the **Out of Service** request, the resource hierarchy is taken out-of-service without bringing it in-service on the other server. The resource can only be brought in-service on the same server, if it was taken out-of-service during resynchronization.

 **Important:** After bringing your resource hierarchy in service on the backup server, you should attempt to connect to the databases, especially when using raw devices as tablespace containers. This is necessary to ensure that all disk partitions are visible on the backup servers and the raw bindings are being established correctly.

If the raw bindings have not been established on the backup servers, it is most likely caused by the fact that new partitions were created on the primary server and added to the configuration, but the partition tables have not yet been updated on the backup servers.

The solution is to reboot the backup servers so that the partition tables are updated correctly.

## 6.2.6. DB2 Troubleshooting

Symptom	Possible Cause
One or more of your DB2 EEE partition servers fail to start	The <b>db2nodes.cfg</b> file's port number may have erroneously outgrown the range set in the <b>/etc/services</b> file. View the number of ports set in the <b>db2nodes.cfg</b> file and ensure that the ports range value in the <b>/etc/services</b> file is large enough to accommodate.
LifeKeeper "In-Service" or "Out-of-Service" operation hangs	The DB2 environment variable:  <b>DB2_NUM_FAILOVER_NODES</b> may not have been properly set. Ensure that for all servers in your configuration, this environment variable is set to equal the total number of partitions in the instance.  EXAMPLE:  <b>db2set DB2_NUM_FAILOVER_NODES =&lt;partitions in instance&gt;</b>
LifeKeeper "In-Service" operation hangs	The <b>dasupdt</b> command may not have been executed on the DB2 Administration server. Ensure that the <b>dasupdt</b> command was successfully executed on the DB2 Administration server.
LifeKeeper First Switch over operation fails	The DB2 Fenced User may not have been created on the backup server. Verify the DB2 Fenced User for the specified instances exists with the same user and group id for the primary. Ensure that the protected instance is also a member of the Administration Server group.
You need to add a new node to your existing DB2 resource hierarchy	Please see the <b>nodes</b> utility man page for complete instructions on adding a new node to your currently existing LifeKeeper DB2 resource hierarchy.
Administration Server fails to start	Verify another Administration Server is not already running on specified port.
Creating a DB2 Resource Hierarchy takes long time	Creating a resource may take long time to protect DB2 instance that has large DB. Activate before creating a resource.

### Error Messages

Refer to the [DB2 Recovery Kit Message Catalog](#) for a list of all messages that may be encountered while utilizing the DB2 kit.

## 6.2.7. Setting Up DB2 to use Raw I/O

There are several requirements for configuring RAW I/O devices for DB2 so that the DB2 instance can be protected by LifeKeeper.

### Requirements

- The Linux OS must support Raw I/O devices. For most distributions this support was included in the 2.4 kernel, but there are some distributions that support Raw I/O on a 2.2. kernel.
- All Raw I/O devices must be bound to a shared disk partition. A number of shared SCSI disk partitions is required. The exact number is determined by the number of tablespaces that will be located on Raw I/O devices. (Please see to DB2 documentation for guidelines for writing tablespaces on raw devices).
- DB2 Version 7.1 Fix Pack 3 or later OR DB2 Version 8 or higher is required.

### Raw I/O Setup Steps

The following steps 1-4 were taken from Section 7.3.1.1 (“Using Raw I/O on Linux”) of the *IBM Db2 Universal Database Release Notes Version 7.2/Version 7.1 Fix Pack 3*. In this example, the raw partition to be used is `/dev/sda5`. It should not contain any valuable data.

Note that step 4 or 5 will vary depending upon whether you are using Multiple Logical Nodes.

1. Calculate the number of 4 096-byte pages in this partition, rounding down if necessary.

Example:

```
# fdisk /dev/sda
```

```
Command (m for help):p
```

```
Disk /dev/sda:255 heads, 63 sectors, 1106 cylinders
```

```
Units = cylinders of 16065 * 512 bytes
```

Device Boot	Start	End	Blocks	System	ID
dev/sda1	1	23	4200997	83	Linux
/dev/sda2	524	1106	4682947+	5	Extended
/dev/sda5	524	1106	4682947	83	Linux

```
Command (m for help):q
```

```
#
```

The number of pages in `/dev/sda5` is:

```
num_pages = floor( ((1106-524+1)*16065*512)/4096 )
```

```
num_pages = 11170736
```

2. Bind an unused raw device node to this partition. Since this needs to be done every time the machine is rebooted, and requires root access, you may want to add the raw bindings to a system initialization file (i.e. `rc.local` or `boot.local`.) **These bindings must be removed once the hierarchy is under LifeKeeper protection.** LifeKeeper will re-establish the raw bindings for Raw I/O devices that are under LifeKeeper protection.

Use **raw -qa** to see which raw device nodes are already in use:

```
raw /dev/raw/raw1 /dev/sda5
```

```
/dev/raw/raw1:bound to major 8, minor 5
```

3. Set global read permissions on the raw device controller and the disk partition. Set global read and write permissions on the raw device:

```
# chmod a+r /dev/rawctl
```

```
# chmod a+r /dev/sdb1
```

```
# chmod a+rw /dev/raw/raw1
```

4. **Important:** This step only applies if you are using DB2 EE OR your DB2 EEE configuration will never run Multiple Logical Nodes (MLNs) even after failover. If the configuration may run MLNs at some point, proceed to step 5.

Create the tablespace in DB2, specifying the raw device, not the disk partition.

For example:

```
CREATE TABLESPACE dms1
```

```
MANAGED BY DATABASE
```

```
USING (DEVICE '/dev/raw/raw1' 11170736)
```

Tablespaces on raw devices are also supported for all other page sizes supported by DB2.

5. **Note:** This step must be followed if the configuration is running MLNs or will run MLNs at some point after failover.

Create the table space in DB2, specifying the raw device, not the disk partition, and specify a different raw I/O device node for each DB2 instance partition.

For example:

```
CREATE TABLESPACE dms1  
  
MANAGED BY DATABASE  
  
USING (DEVICE '/dev/raw/raw1' 11170736) on NODE (NODENUM)  
  
USING (DEVICE '/dev/raw/<different raw device node>' ##### ) on NODE  
(NODENUM)
```

**Note:** ON NODE must be used because each DB2 node (database partition server) must use a different raw I/O device. This must be specified even if the node is running on a different machine so that the failover will work correctly.

## 6.3. Recovery Kit for EC2™ Administration Guide

---

The Recovery Kit for EC2™ provides a mechanism to recover an Elastic IP from a failed primary server to a backup server. It also provides a mechanism to enable the IP Recovery Kit to work in multiple availability zones.

Please see the [Principles of Operation](#) for a comparison and additional information about the definition, scenarios, and operation of the Recovery Kit for EC2™.

### LifeKeeper Documentation

The following is a list of LifeKeeper for Linux related information available from SIOS Technology Corp.

- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Release Notes](#)
- [SIOS Technology Corp. Documentation](#)

Please refer to [Amazon Elastic Compute Cloud \(EC2\) Documentation](#) for more information.

 **Note:** “Amazon Web Services, the “Powered by Amazon Web Services” logo, “AWS”, “Amazon EC2”, “EC2”, “Amazon Elastic Compute Cloud”, “Amazon Virtual Private Cloud”, and “Amazon VPC” are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. Trademark symbols such as ® and ™ may be omitted from system names and product names in this document.

## 6.3.1. Recovery Kit for EC2™ Principles of Operation

---

Recovery Kit for EC2™ provides two functions.

1. The Route Table scenario (Backend Cluster) manages Route Table for LifeKeeper-protected IP resources to be reached from clients within the Amazon VPC™.
2. The Elastic IP scenario (Frontend Cluster) manages Elastic IP available from the Internet.

### Route Table scenario (Backend Cluster):

To clarify the administration and operation of Route Table, consider the scenario shown in Figure 1.

This example configuration contains one Amazon VPC™, two Availability Zones (AZ).

There are two Subnets in each AZ.

- A first Subnet (hereinafter referred to as “Public Subnet”) connects to the Internet via Internet Gateway by Route Table – see Route Table of 10.0.1.0/24 and 10.0.3.0/24.
- A second Subnet (hereinafter referred to as “Private Subnet”) connects to the Internet via NAT Instance by Route Table – see Route Table of 10.0.2.0/24 and Route Table of 10.0.4.0/24.

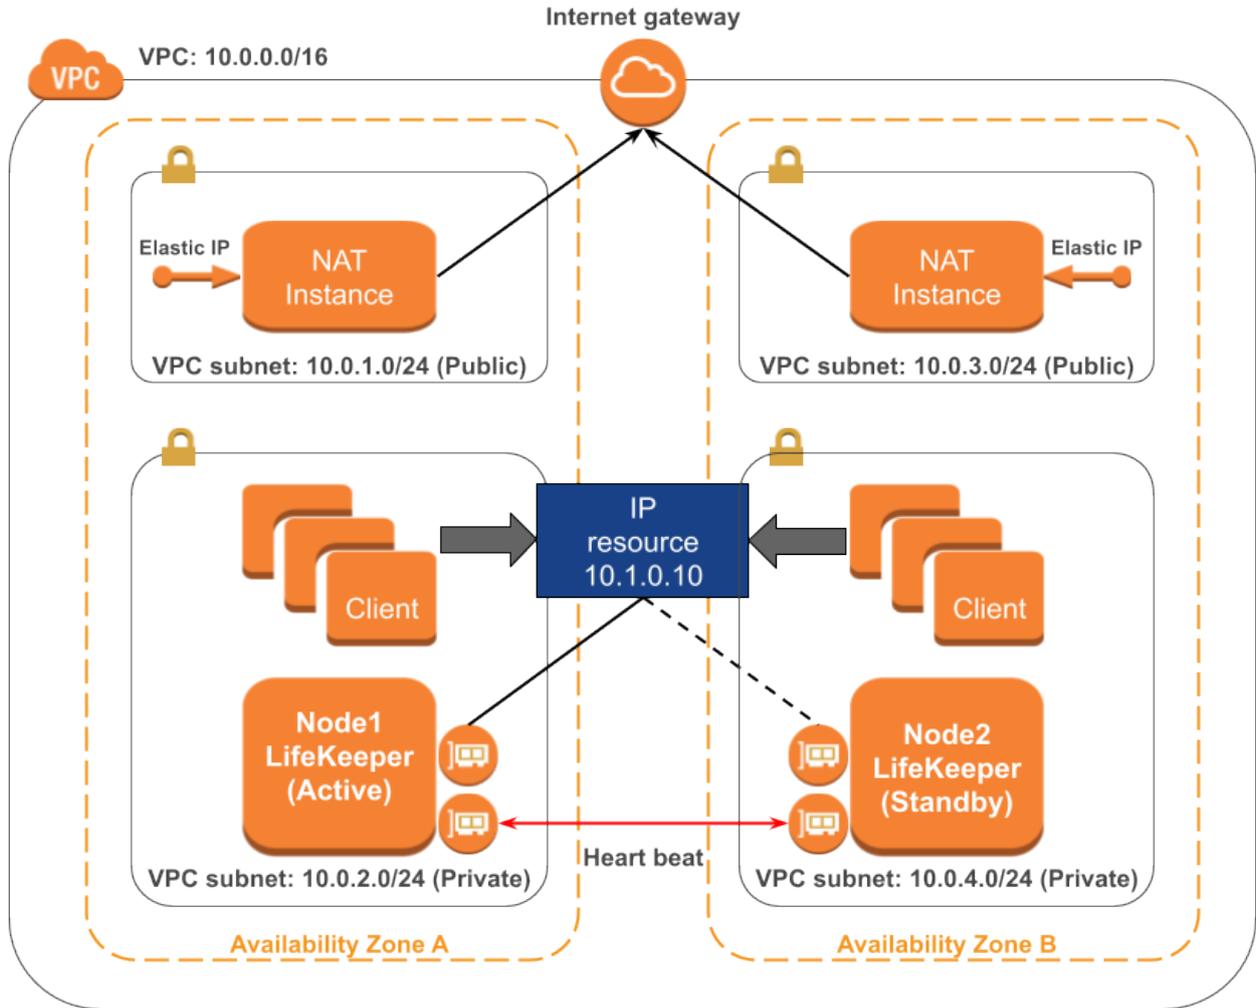
In each Public Subnet, there is an EC2™ instance to which you assigned an Elastic IP for NAT (hereinafter referred to as “NAT Instance”).

In each Private Subnet, there is an EC2™ instance for LifeKeeper Active/Standby (hereinafter referred to as “Node1” and “Node2”), and there are clients that will use the applications protected by Node1/Node2.

Each Node1/Node2 has two Elastic Network Interfaces (ENIs).

Configure the Network ACLs and Security Groups to be able to communicate between each Instance and each Node.

Figure 1. Route Table scenario



Route Table of 10.0.1.0/24 and 10.0.3.0/24

Destination	Target	Note
10.0.0.0/16	Local	Default
0.0.0.0/0	Internet Gateway	In order to connect to the Internet, requires the allocation of an Elastic IP.

Route Table of 10.0.2.0/24

Destination	Target	Note
10.0.0.0/16	Local	Default
10.1.0.10/32 (IP resource)	Elastic Network Interface (ENI) on LifeKeeper Active Node	This Target is updated by Recovery Kit for EC2™ during a switchover.
0.0.0.0/0	NAT instance (10.0.1.0)	Connect to the Internet via NAT

## Route Table of 10.0.4.0/24

Destination	Target	Note
10.0.0.0/16	Local	Default
10.1.0.10/32 (IP resource)	Elastic Network Interface (ENI) on LifeKeeper Active Node	This Target is updated by Recovery Kit for EC2™ during a switchover.
0.0.0.0/0	NAT instance (10.0.3.0)	Connect to the Internet via NAT

When a resource switchover is performed, LifeKeeper will take the IP resource out of service on Node 1. The Target entry of 10.1.0.10/32 in each Private Subnet will be updated to reflect the ENI of Node2. The IP resource will be brought in-service on Node2. Therefore IP address traffic to 10.1.0.10 is effectively redirected to Node2 by the new Route Table configuration changes in the Private Subnet.

If you need to access the IP address 10.1.0.10 from another subnet containing the public subnet, please add the destination route 10.1.0.10/32 to the route table entry for each subnet. LifeKeeper controls all entries for which the destination is set as “10.1.0.10/32” in the route table within the VPC.

### Elastic IP scenario (Frontend cluster):

To clarify the administration and operation of Elastic IP, consider the scenario shown in Figure 2.

This example configuration contains one Amazon VPC™, two Availability Zones (AZ).

There is one Subnet in each AZ.

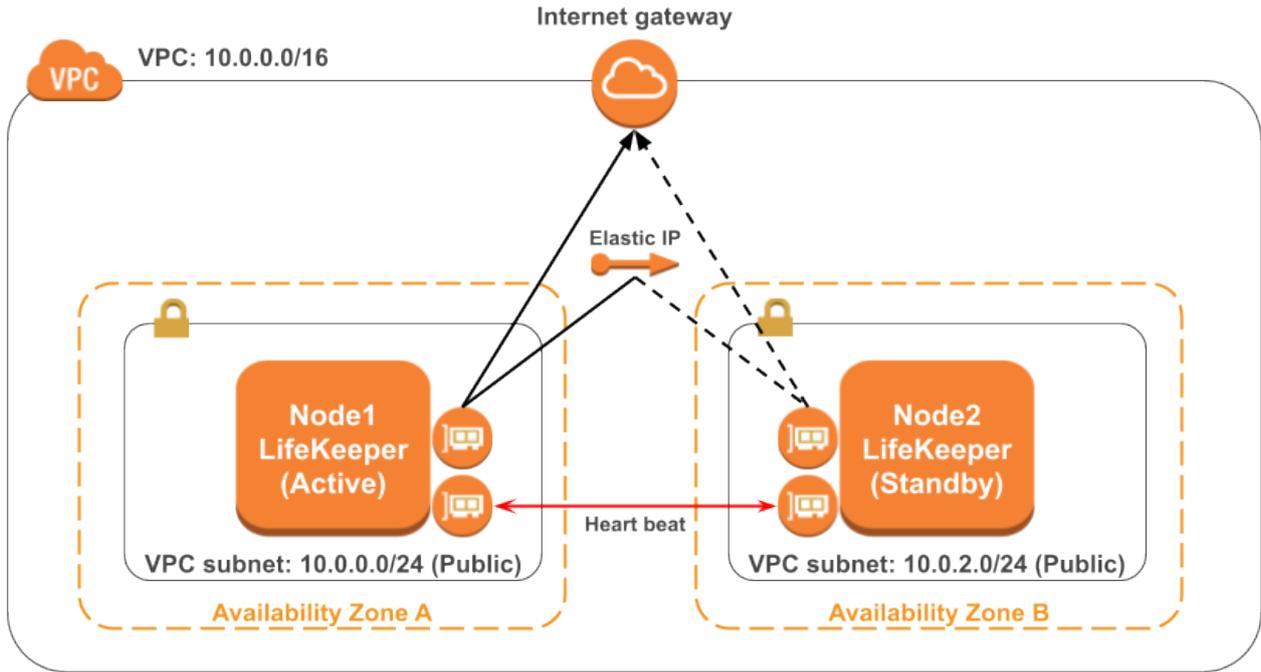
Each Subnet connects to the Internet via Internet Gateway by Route Table.

In Subnet, there is an EC2™ instance for LifeKeeper Active/Standby (hereinafter referred to as “Node1” and “Node2”).

Each Node1/Node2 has two Elastic Network Interfaces (ENIs).

Configure the Network ACLs and Security Groups to be able to communicate between each Node.

**Figure 2. Elastic IP scenario**



The system administrator allocates an Elastic IP address of frontend cluster to the ENI.

Assuming that Node1 is the primary server for the resource, the administrator creates the EC2™ resource hierarchy on Node1 using the wizard described in the section entitled [Creating a Resource Hierarchy](#).

When resource switchover is performed, Recovery Kit for EC2™ disassociates the Elastic IP from the ENI on Node 1. After that Recovery Kit for EC2™ determines if the elastic IP is associated with the ENI on Node 2, if not, associates the Elastic IP to the ENI. Therefore client on the Internet can reach Node 2 via the Elastic IP after switchover.

**Note:** Standby nodes need to have an access to the end point in order to control the EC2™ instance: that is, it is necessary to connect to the outside VPC. Please refer to “[Requirements](#)” for details. A public IP address is not necessary to access the endpoint when using PrivateLink. For details, please refer to “[VPC Endpoints](#).”

## 6.3.2. Recovery Kit for EC2™ Requirements

Before attempting to install or remove the Recovery Kit for EC2™ you must understand Amazon Web Service software requirements, as well as the installation and removal procedures for the Recovery Kit for EC2™ package.

### Amazon Web Service and Software Requirements

Before installing and configuring the Recovery Kit for EC2™, be sure that your configuration meets the following requirements:

#### Amazon Virtual Private Cloud (VPC):

- The recovery kit requires a VPC be configured within AWS
- Two or more Subnets created on different Availability Zones (AZ)
- Each Subnet contains associated Route Tables
- If you are configuring a Public (Frontend) Cluster, then one or more Elastic IPs must be allocated.

#### Amazon Elastic Compute Cloud (EC2):

- The recovery kit requires two or more EC2™ instances.
- The instances are associated on each Subnet.
- The instances are attached to an Elastic Network Interface (ENI).
- If creating a Route Table (backend cluster) resource, the network interface of each instance should have its network source/destination checks disabled.
- AWS Command Line Interface (AWS CLI) needs to be installed in each of EC2™ the instances. For the details, please refer to [AWS Command Line Interface Installation](#).
- All the EC2™ instances must be able to access Amazon EC2™ services endpoints ([AWS Regions and Endpoints](#)) using the protocols HTTP and HTTPS. Please configure EC2™ and the OS properly.
- In order to obtain metadata of Amazon EC2™ instances, it is necessary to have an access to IP address 169.254.169.254 using the HTTP protocol.
- Since the AWS CLI is used, outbound connections on TCP port 443 must be enabled.
- Since the Auto Recovery function may conflict with the recovery function of LifeKeeper, it is not recommended to use these functions together.

**Note:** If the path name of AWS CLI executable files is not specified on the “PATH” parameter in the LifeKeeper defaults file `/etc/default/LifeKeeper`, you must append the path name of AWS CLI executable files to the “PATH” parameter.

## AWS Identity and Access Management (IAM):

In order for LifeKeeper to operate AWS, an IAM user or IAM role with the following access privilege is required. Please configure an [EC2™ IAM role](#) or [configure AWS CLI](#) appropriately so that it can be accessed from root user of the EC2™ instance.

### Route Table (backend) configuration:

- ec2:DescribeRouteTables
- ec2:ReplaceRoute
- ec2:DescribeNetworkInterfaceAttribute
- ec2:ModifyNetworkInterfaceAttribute

### Elastic IP (frontend) configuration:

- ec2:DescribeAddresses
- ec2:AssociateAddress
- ec2:DisassociateAddress

## LifeKeeper Software:

You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.

### LifeKeeper Recovery Kit for EC2™:

You must install the same version of Recovery Kit for EC2™ software and any patches on each server.

### LifeKeeper IP Recovery Kit:

If you are using the Recovery Kit for EC2™ to provide protection for the Route Table (Backend Cluster), you must install the same version of LifeKeeper for Linux IP Recovery Kit software and any patches on each server.

**Note:** Please refer to the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information. You should refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper Recovery Kit for EC2™.

SIOS recommends using Quorum/Witness when using the Recovery Kit for EC2™. Please refer to [Quorum/Witness](#) for more information.

## 6.3.3. Recovery Kit for EC2™ Configuration

To ensure that your LifeKeeper configuration provides the protection and flexibility you require you'll need to be aware of the configuration requirements. To appropriately plan your configuration you must understand Amazon, Amazon Virtual Private Cloud, (VPC), Amazon Elastic Compute Cloud (EC2™), and the user system setup hierarchy options. In addition to planning your configuration, this section also includes the specific tasks required to configure your recovery kit.

### Specific Configuration Considerations for Amazon EC2™

In order to properly configure your Recovery Kit for EC2™ you should review the following topics to ensure that you have the information necessary to complete the configuration tasks:

- [User System Setup](#)

See the following topics for further configuration considerations:

- [EC2™ Resource Monitoring and Configuration Considerations](#)
- [EC2™ Local Recovery and Configuration Considerations](#)

### Specific Configuration Considerations for Amazon EC2™

The following configuration tasks for EC2™ resources are described in this section. They are unique to an EC2™ resource instance and different for each recovery kit.

- [Creating a Resource Hierarchy](#). Creates an application resource hierarchy in your LifeKeeper cluster.
- [Deleting a Resource Hierarchy](#). Deletes a resource hierarchy from all servers in your LifeKeeper cluster.
- [Extending Your Hierarchy](#). Extends a resource hierarchy from the primary server to a backup server.
- [Unextending Your Hierarchy](#). Unextends (removes) a resource hierarchy from a single server in your LifeKeeper cluster.
- [Viewing and Editing EC2™ Configuration Properties](#). Displays configuration details for an EC2™ resource and allows some of them to be modified.
- [Adjusting Recovery Kit for EC2™ Tunable Values](#). Tunes characteristics of the overall behavior of the Recovery Kit for EC2™.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#). They are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View Properties](#) / [Edit Properties](#) View or edit the properties of a resource hierarchy on a specific server.

The rest of this section explains how to configure your recovery kit by selecting certain tasks from the Edit menu of the LifeKeeper GUI. You may also select each configuration task from the toolbar.

- Right-click on a global resource in the Resource Hierarchy Tree (left-hand pane) of the status display window to display the same drop down menu choices as the Edit menu. This is only an option when a hierarchy already exists.
- Right-click on a resource instance in the Resource Hierarchy Table (right-hand pane) of the status display window to perform all the configuration tasks, except Creating a Resource Hierarchy, depending on the state of the server and the particular resource.

## 6.3.3.1. EC2™ Event Table

---

The following table contains the Recovery Kit for EC2™ specific event and associated trap number. The event will generate email notices when `LK_NOTIFY_ALIAS` is set.

LifeKeeper Event/Description	Trap #	Object ID
AWS API failure	180	.1.3.6.1.4.1.7359.1.0.180

## 6.3.3.2. Adjusting Recovery Kit for EC2™ Tunable Values

---

For the parameters that can be configured in the Recovery Kit for EC2™, refer to the [EC2™ Parameters List](#).

## 6.3.3.3. Creating an EC2™ Resource Hierarchy

To create a resource instance from the primary server, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop down menu, select **Create Resource Hierarchy**.
2. A dialog box will appear with a drop down list showing all of the recognized recovery kits installed within the cluster. Select “**Amazon EC2**” from the drop down list and click **Next**.
3. You will be prompted to enter the following information. (When the Back button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful in the event that you need to correct previously entered information.)

 **Note:** If you click the Cancel button at any time when creating your hierarchy, LifeKeeper will cancel the entire creation process.

Field	Tips
Switchback Type	<p>This dictates how the EC2™ resource will be switched back to this server when the server comes back up after a failover. You can choose either intelligent or automatic.</p> <ul style="list-style-type: none"> <li>• Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server.</li> <li>• Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.</li> </ul> <p><b>Note:</b> The switchback type can be changed later from the General tab of the Resource Properties dialog box.</p>
Server	<p>Select the Server for the EC2™ resource (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop down list.</p>
EC2™ Resource type	<p>The EC2 Recovery Kit™ provides protection for two AWS recovery scenarios. The Route Table and Elastic IP scenario.</p> <p>The Route Table scenario is used in conjunction with a local virtual IP address and is typically used for Backend Clusters.</p> <p>The Elastic IP scenario is used for protection of an Elastic IP and is typically used for Frontend Clusters.</p> <p>Select the EC2™ type to be used.</p>

IP resource	This field will only appear and be set in the Route Table scenario. Select the IP resource. This is the virtual IP resource that is protected by LifeKeeper and configured in the Route Table address in the VPC. <b>Note:</b> The list will only show IP resources that are ISP and IPv4 based.
Network Interface	This field will only appear and be set in the Elastic IP scenario. Select the Network Interface to associate with Elastic IPs.
Elastic IP	This field will only appear and be set in the Elastic IP scenario. Select the Elastic IP to be related to the network interface.
EC2™ Resource Tag	Select or enter a unique EC2™ Resource Tag name for the EC2™ resource instance you are creating. This field is populated automatically with a default tag name, ec2-<resource>, where <resource> is the resource name. This tag can be changed.

1. Click **Create**. The Create Resource Wizard will then create your EC2™ resource.
2. At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your EC2™ resource hierarchy. If LifeKeeper detects a problem an ERROR will appear in the information box. If the validation is successful your resource will be created. Click **Next**.

Another information box will appear confirming that you have successfully created an EC2™ resource hierarchy. You must extend that hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

When you click **Continue**, LifeKeeper will launch the Pre-Extend configuration task. Refer to [Extending Your Hierarchy](#) for details on how to extend your resource hierarchy to another server.

If you click **Cancel** now, another dialog box will appear alerting you that you will need to manually extend your EC2™ resource hierarchy to another server at some other time to put it under LifeKeeper protection.

## 6.3.3.4. Deleting an EC2™ Resource Hierarchy

---

To delete a resource hierarchy from all of the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the dropdown menu, select **Delete Resource Hierarchy**.
2. Select the name of the Target Server that you are deleting from your EC2™ resource hierarchy and click **Next**.

**Note:** This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in either pane.

3. Select the Hierarchy to Delete. Identify the resource hierarchy you wish to delete, highlight it then click **Next**.

**Note:** This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in the left or right pane.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete** to proceed.
5. An information box appears confirming that the EC2™ resource was deleted successfully.
6. Click **Done** to exit.

## 6.3.3.5. Extending Your EC2™ Hierarchy

After you have created a hierarchy, you must extend that hierarchy to another server in the cluster. There are three possible scenarios to extend your resource instance from the template server to a target server.

- Continue from creating the resource into extending that resource to another server.
- Enter the Extend Resource Hierarchy task from the edit menu as shown below.
- Right click on an unextended hierarchy in either the left or right hand pane.

Each scenario takes you through the same dialog boxes (with a few exceptions, detailed below).

1. If you are entering the Extend wizard from the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Extend Resource Hierarchy**. This will launch the Extend Resource Hierarchy wizard. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The Pre-Extend Wizard will prompt you to enter the following information.

 **Note:** The first two fields appear only if you initiated the Extend from the Edit menu. It should be noted that if you click Cancel at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extension process to that particular server. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Field	Tips
Switchback Type	<p>Select the Switchback Type. This dictates how the EC2™ instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either intelligent or automatic.</p> <ul style="list-style-type: none"> <li>• Intelligent switchback requires administrative intervention to switchback the instance to the primary/original server.</li> <li>• Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.</li> </ul> <p>The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.</p>
Template Priority	<p>Select or enter a Template Priority. This is the priority for the EC2™ hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will</p>

	<p>reject any priority for this hierarchy that is already in use by another system. The default value is recommended.</p> <p><b>Note:</b> This selection will appear only for the initial extend of the hierarchy.</p>
Target Priority	<p>Select or enter the Target Priority. This is the priority for the new extended EC2™ hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server’s priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest).</p> <p><b>Note:</b> LifeKeeper assigns the number “1” to the server on which the hierarchy is created by default. The priorities do not need to be consecutive and no two servers can have the same priority for a given resource.</p>

- An information box will appear explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this EC2™ resource have been met. If there were some requirements that have not been met, LifeKeeper will not allow you to select the **Next** button, and the **Back** button will be enabled. If you click **Back**, you can make changes to your resource extension according to any error messages that may appear in the information box. If you click **Cancel** now, you will need to manually extend your EC2™ resource hierarchy to another server to put it under LifeKeeper protection. When you click **Next**, LifeKeeper will launch you into the Extend Resource Hierarchy configuration task.
- The Extend Resource Hierarchy configuration task will prompt you to enter the following information:

Field	Tip
EC2™ Resource Tag	<p>Select or enter the EC2™ Resource Tag. This is the resource tag name to be used by the EC2™ resource being extended to the target server.</p> <p><b>Note:</b> The field is not editable.</p>

- An information box will appear verifying that the extension is being performed. Click **Next Server** if you want to extend the same EC2™ resource instance to another server in your cluster. This will repeat the Extend Resource Hierarchy operation. If you click **Finish**, LifeKeeper will verify that the extension of the EC2™ resource was completed successfully.
- Click **Done** to exit from the Extend Resources Hierarchy menu selection.

 **Note:** Be sure to test the functionality of the new instance on all servers.

## 6.3.3.6. EC2™ Local Recovery and Configuration

### Local Recovery scenario (Backend Cluster):

When a failure of the protected Route Table is detected by Recovery Kit for EC2™, the resulting failure triggers the execution of the EC2™ local recovery script. The local recovery gathers specified IP resource entries in all Route Tables and changes the entries' Target to the ENI on the active server. It also disables source/destination checks for the network interface. If the local recovery attempt fails, LifeKeeper will perform a failover of the EC2™ resource and all dependent resources to a standby server. See the [Principles of Operation](#) section for the configuration of this scenario.

 **Note:** Since the recovery kit will protect the configuration of the route table once the corresponding EC2™ resource gets created, the route table should not be modified manually.

The following example shows a typical scenario of the local recovery: When the recovery kit detects a wrong target setting of IP routing in the route table, the local recovery replaces the target to the ENI on the active server. During this process nothing will be changed regarding the entry of 10.1.0.20/32 on the Route Table B.

IP resource	10.1.0.10
ENI on Active Node	eni-01234567

#### Route Table A – Before

Destination	Target
10.1.0.10/32	eni-89abcdef
10.0.0.0/16	local

#### Route Table A – After

Destination	Target
10.1.0.10/32	eni-01234567
10.0.0.0/16	local

## Route Table B – Before

Destination	Target
10.1.0.10/32	eni-89abcdef
10.1.0.20/32	eni-89abcdef
10.0.0.0/16	local

## Route Table B – After

Destination	Target
10.1.0.10/32	eni-01234567
10.1.0.20/32	eni-89abcdef
10.0.0.0/16	local

## Elastic IP scenario (Frontend Cluster):

When a failure of the protected Elastic IP is detected by Recovery Kit for EC2™, the resulting failure triggers the execution of the EC2™ local recovery script. The local recovery allocates the Elastic IP to the ENI on the active node. If the local recovery attempt fails, LifeKeeper will perform a failover of the EC2™ resource and all dependent resources to a standby server. See the [Principles of Operation](#) section for the configuration of this scenario.

## 6.3.3.7. EC2™ Resource Monitoring and Configuration

---

### Route Table scenario (Backend Cluster):

The recovery kit uses AWS CLI to perform the monitoring of the Route Table settings to enable access from clients within the VPC to the protected IP resources. The recovery kit ensures that the target of the IP resources for all the Route Tables in the VPC is correctly set to the ENI on the active server. It also ensures that source/destination checks for the network interface are disabled. Otherwise, the recovery kit performs the EC2™ local recovery process.

### Elastic IP scenario (Frontend Cluster):

The recovery kit uses AWS CLI to monitor the association of the Elastic IP with the ENI on the active server. The recovery kit ensures that the Elastic IP is correctly associated with the ENI attached on the active server. Otherwise, the recovery kit performs the EC2™ local recovery process.

 **Note:** In both scenarios, when a timeout occurs at AWS CLI, no failover will be performed and the resource will remain in ISP state. Only a timeout related message will be logged in the LifeKeeper log. The recovery kit will execute the monitoring once again after a check interval. See the [EC2™ Parameters List](#) for more information about how to configure the value for timeout.

## 6.3.3.8. Unextending Your EC2™ Hierarchy

---

To unextend a hierarchy complete the following steps:

1. From the **LifeKeeper GUI menu**, select **Edit**, then **Resource**. From the dropdown menu, select **Unextend Resource Hierarchy**.
2. Select the Target Server that you are unextending from the EC2™ resource. It cannot be the server that the EC2™ resource is currently in service on. Click **Next**.

**Note:** If you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance, the dialog box will not appear.

3. Select the EC2™ hierarchy to unextend. Click **Next**.

**Note:** If you selected the **Unextend** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, the dialog will not appear.

4. An information box will appear confirming the target server and the EC2™ resource hierarchy you have chosen to unextend. Click **Unextend**.
5. An information box will appear confirming the EC2™ resource was unextended successfully.
6. Click **Done** to exit.

## 6.3.3.9. EC2™ User System Setup

---

### Route Table scenario (Backend Cluster):

The Route Table protection option in the Recovery Kit for EC2™ provides the ability to automatically update the routing in the VPC. During a failover the recovery kit will update the route table to reflect the new Elastic Network Interface (ENI) location of the virtual IP address on the target server. In order for LifeKeeper to protect, monitor and update the Route Table in the VPC, the following configuration steps must be performed (this also applies to the active/active configuration):

- The virtual IP address to be protected by the LifeKeeper for Linux IP Recovery Kit must be out of range of the allocated CIDR in the VPC.
- The virtual IP address must be protected by LifeKeeper prior to creating the Recovery Kit for EC2™ resource.
- The Source/Dest Checking of the ENI must be disabled. This is required in order for the instance to accept network packets for the virtual IP address.
- Broadcast PING checking of the LifeKeeper IP resources must be disabled. LifeKeeper monitors IP resources by executing the Broadcast PING test of the IP address on the local subnet. In multiple availability zone environments this feature would not be useable because of the different subnets that exist between multiple availability zones. To disable this feature you must set the NOBCASTPING entry in the */etc/default/LifeKeeper* configuration file as follows:

```
NOBCASTPING=1
```

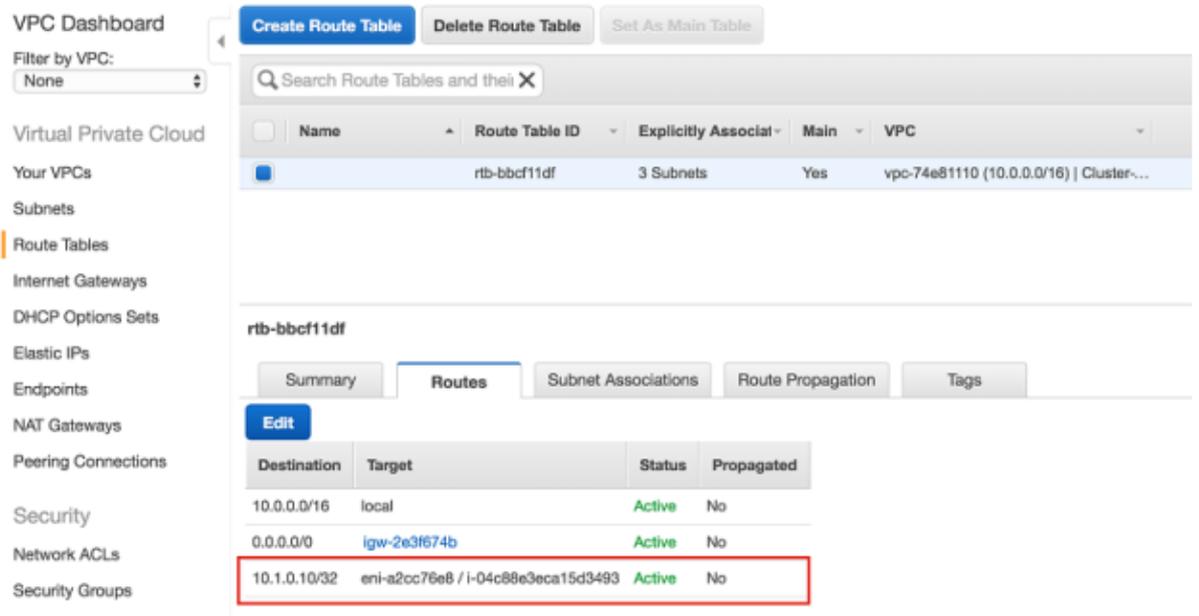
- The Route Table should have a route entry for the virtual IP address and the ENI of the active server.

**Note:** Since the EC2™ recovery kit will protect the configuration of the Route Table once the corresponding EC2™ resource has been created, the Route Table should not be modified manually after hierarchy creation.

#### Example:

Destination: VIP 10.1.0.10/32

Target: eni-a2cc76e8



### Elastic IP scenario (Frontend Cluster):

The Elastic IP (EIP) protection option in the Recovery Kit for EC2™ provides the ability to automatically re-associate an EIP with a specific ENI (the ENI used by the EC2™ resource on the active or backup server).

In order for LifeKeeper to protect, monitor and update the association of an EIP with the ENI on the active or backup server, the following configuration steps must be performed:

- One ENI can be associated with only one Elastic IP. No other EIPs (any EIPs other than the one used by EC2™ resource) should be associated with the specific ENIs. Otherwise the recovery kit will disassociate any other EIPs that are already associated with the specific ENIs.

#### Notes:

- Since an Elastic Block Store (EBS) of AWS can only be attached to one EC2™ instance, DataKeeper for Linux is recommended when creating an HA cluster configuration using EBS.
- We recommend increasing RESRVRECTIMEOUT in /etc/default/LifeKeeper to 300 from 150 as the default. RESRVRECTIMEOUT is the number of seconds that a LifeKeeper process will sleep when waiting to reserve a resource for “recovery”, while another process already has the resource reserved.

## 6.3.4. Recovery Kit for EC2™ Troubleshooting

---

The [Message Catalog](#) provides a list of all messages that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received, or you may go directly to the [Recovery Kit for EC2™ Message Catalog](#) which contains a list of all messages that may be encountered while utilizing the Recovery Kit for EC2™.

## 6.4. Generic Application Kit for Load Balancer Health Checks

---

### Overview

The Generic Application Kit for Load Balancer Health Checks (Gen LB) provides a mechanism to receive and respond to a health check probe for load balancer target instances of load balancer instances in Microsoft Azure (Azure) and Google Cloud™ environments.

This document explains how the Gen LB scripts work. Please refer to the [LifeKeeper for Linux Technical Documentation](#) for definitions of LifeKeeper terms and information on how to use LifeKeeper.

It is important to note that this recovery kit may only be used in the supported cloud environments (Microsoft Azure and Google Cloud), and there is no guarantee that the recovery kit will function correctly in unsupported environments.

### Prerequisites

- The scripts included with the Gen LB ARK may only be used with LifeKeeper for Linux v9.5.1 and later.
- The OS and configurations supported by LifeKeeper for Linux v9.5.1 and later are supported.  
**Note:** In v9.5.1 SAP is not supported on Google Cloud. For the supported OS and configurations, refer to the [Release Notes](#) and the [LifeKeeper for Linux Support Matrix](#). Please see the installation steps of the online documentation.
- Microsoft Azure and Google Cloud are the only supported cloud platforms for the Gen LB ARK.
- The Gen LB ARK only supports use of the TCP protocol for health check probes.

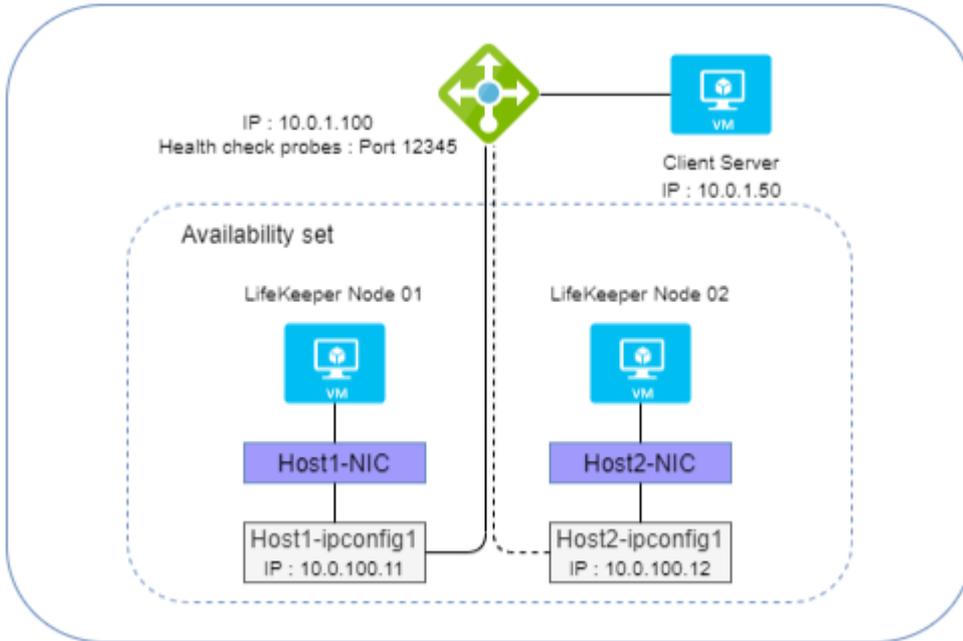


**Note:** Google Cloud, BigQuery and Google Compute Engine are trademarks of Google LLC.

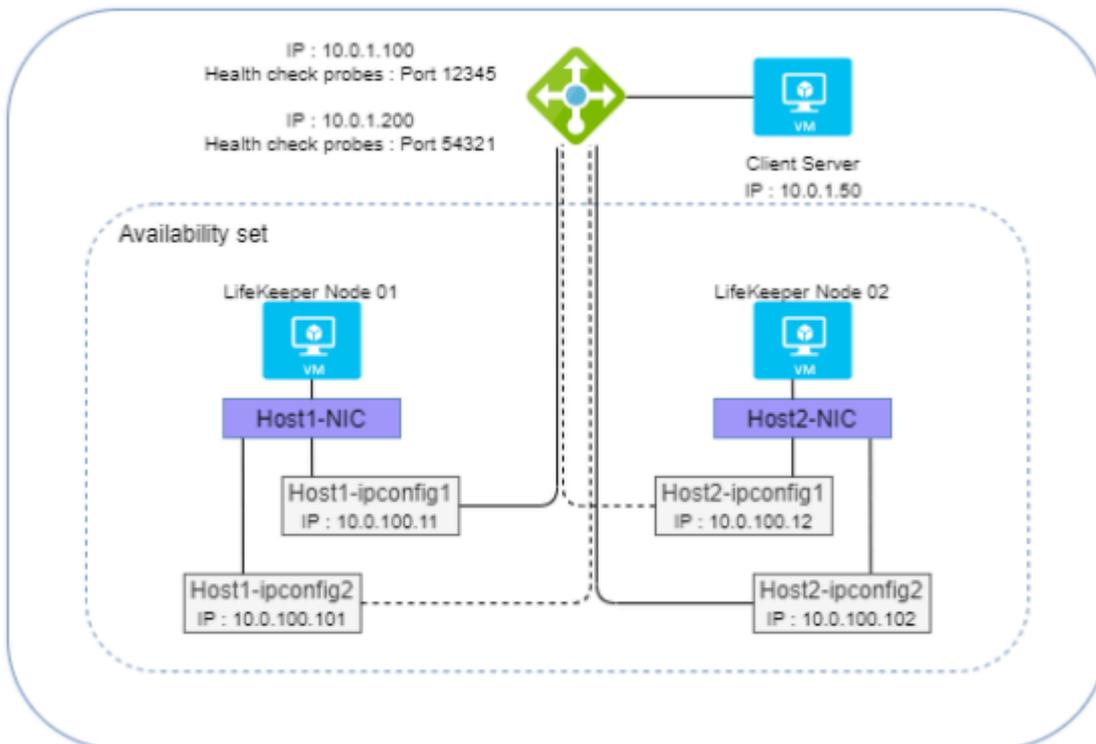
# 6.4.1. Configuration Examples

## Configuration Examples for Azure

### Active / Standby Configuration

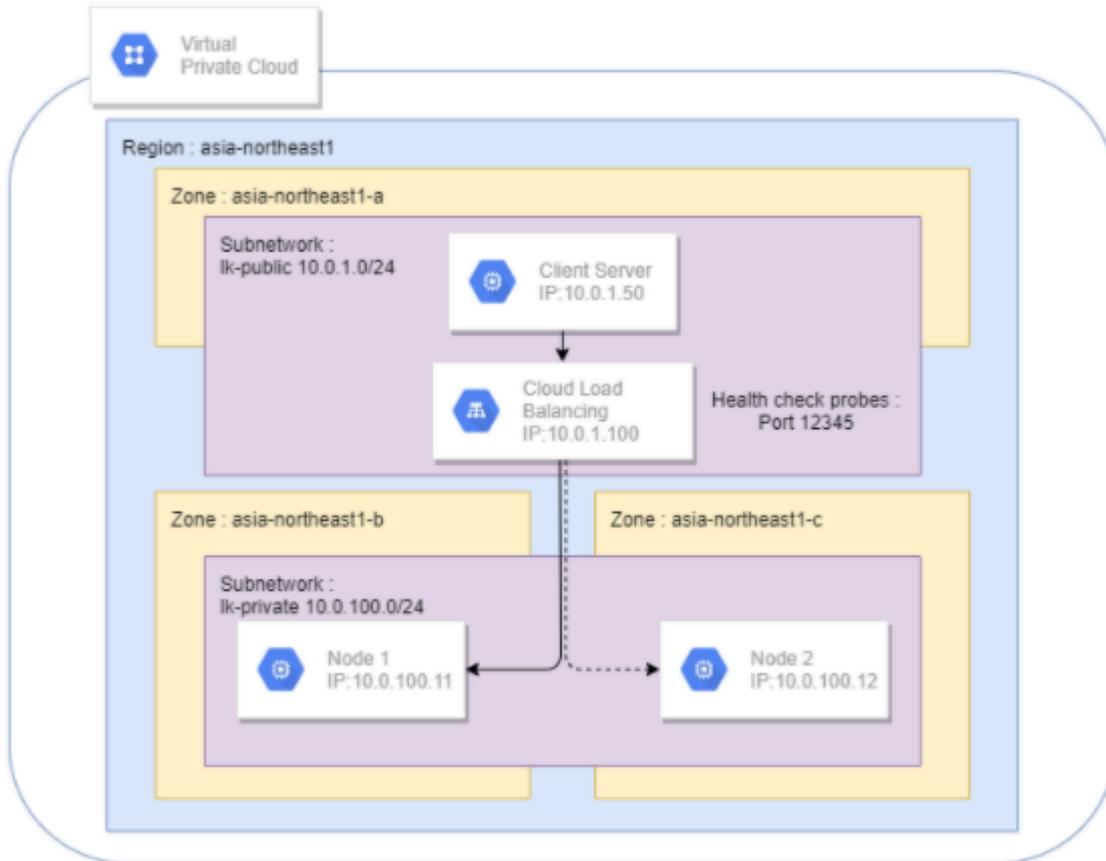


### Active / Active Configuration

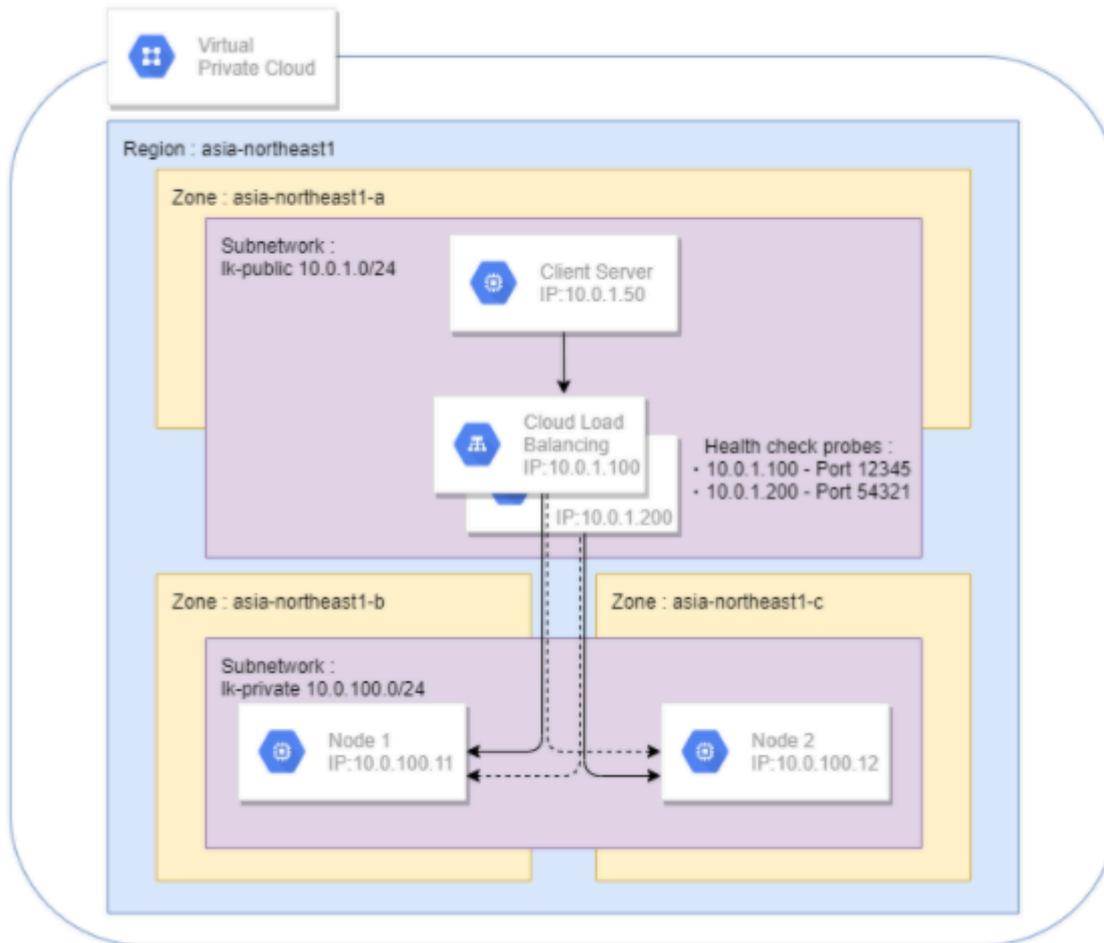


# Configuration Examples for Google Cloud

## Active / Standby Configuration



### Active / Active Configuration



## 6.4.2. Basic Behaviors

---

### Bringing a Resource in Service (restore)

When a Gen LB resource is brought in service via the LifeKeeper GUI client or LifeKeeper CLI, daemons are started which listen for a health check probe at a specified port. Once these processes are started successfully, the resource status changes to ISP (In-Service, Protected). The status of the Load Balancer Health Check resource will be set to OSF (Out of Service, Failed) when the processes cannot be started successfully. The same process is also performed during switching operations such as switchover and failover. Refer to [Script Specifications](#) for more information about this behavior.

### Taking a Resource Out of Service (remove)

When a Gen LB resource is taken out of service via the LifeKeeper GUI client or LifeKeeper CLI, the daemons are stopped. Refer to [Script Specifications](#) for more information about this behavior.

### Monitoring (quickCheck)

The Gen LB quickCheck script confirms that the appropriate daemons are running properly while the Load Balancer Health Check resource is in service (ISP). When the monitoring process detects that the daemon has not been started properly, a failover is performed. Refer to [Script Specifications](#) for more information about this behavior.

### Recovery (recover)

There is no recovery process.

## 6.4.3. Script Specifications

---

A Generic Application resource is created by specifying the appropriate restore, remove, quickCheck and recover scripts. Refer to the actual scripts for more information.

Examples of command line related file names and paths are described in the following explanation. Refer to the actual scripts for more details.

### Bringing a Resource In Service (restore)

The following steps are performed by the restore script to start the Load Balancer Health Check resource.

1. The tag name and resource ID of the Load Balancer Health Check resource are obtained from the `-t` and `-i` command line options, respectively.
2. The port number and reply message are obtained from the info field for the Load Balancer Health Check resource. If the resource information cannot be obtained, the restore script exits with exit code 1.
3. The restore script uses the given resource ID to determine whether the associated daemon process has already been started. If it has already been started, the restore script will exit with exit code 1.
4. The restore script generates the Socket object and waits for the connection from the Load Balancer. When no connection occurs within `$HC_TIMEOUT` sec (described in [Script Parameter List](#)), the restore script exits with exit code 1.
5. The health probe listening process described in step 4 runs as a daemon in the background. When no connection occurs for `$HC_TIMEOUT` sec (described in [Script Parameter List](#)), the `$id-hc-ng` flag is created. When the next connection attempt is detected, the flag will be deleted and the timeout will be reset. **Note:** The `$id-hc-ng` flag is used for failure detection by quickCheck.
6. The restore parent process exits with exit code 0.

### Taking a Resource Out of Service (remove)

The following steps are performed by the remove script to stop the Load Balancer Health Check resource.

1. The tag name and resource ID of the Load Balancer Health Check resource are obtained from the `-t` and `-i` command line options, respectively.
2. The PID of the running Load Balancer Health Check daemon process corresponding to the given resource ID is obtained.

3. The remove script verifies that the daemon process with the PID found in step 2 is still running. If it is already stopped, the remove script exits with exit code 0.
4. If the daemon process is still running, the remove script terminates it with a TERM signal.
5. The remove script exits with exit code 0.

## Monitoring (quickCheck)

The following steps are performed by the quickCheck script to monitor the Load Balancer Health Check resource.

1. The tag name and resource ID of the Load Balancer Health Check resource are obtained from the `-t` and `-i` command line options, respectively.
2. The resource status of the Load Balancer Health Check resource is obtained. When the status is not ISP (i.e., the resource has already been taken out of service), the quickCheck script exits with exit code 0.
3. The quickCheck script checks for the existence of the `$id-hc-ng` flag, which is generated by step 5 of the restore process if no connection attempt has been received within the past `$HC_TIMEOUT` seconds. If the flag exists, the quickCheck script exits with exit code 1.
4. Terminating the process with exit code 0.

## Execution Time of the Scripts

### restore

The process will timeout if there is no connection for `$HC_TIMEOUT` sec from Load Balancer when executing the restore. The absence of connection requests from a load balancer on the specified port for an extended period of time generally indicates a configuration error or a network issue that may be preventing the load balancer from communicating with the server.

### remove, quickCheck

No time required for the scripts.

## 6.4.4. Script Parameter List

---

The configurable parameters list for each script is below.

Parameter	Description
HC_TIMEOUT	<p>Set this parameter as the connection timeout value from the Load Balancer.</p> <p>If there is no connection for the specified number of seconds, it is treated as a failure and failover is performed. The default value is 60 seconds. This parameter should be set in the <i>/etc/default/LifeKeeper</i> file.</p> <p><b>Example:</b> HC_TIMEOUT=30</p>

Parameters other than the above cannot be modified. Note that modifying parameters other than the allowed parameters is not supported. Modification of the code in the action scripts for the Gen LB resource is also prohibited.

## 6.4.5. Creating/Extending a Resource

---

The Generic Application resource type is used in order to create/extend a Gen LB resource. Please refer to the [LifeKeeper for Linux Technical Documentation](#) for more information about Generic Application resources.

During Gen LB resource creation, you will be prompted to specify the port number and reply string in the Application Information (AppInfo) of the Generic Application resource creation wizard. The port number is mandatory (any value from 1024 to 65535 that is not currently in-use on the server is valid). The reply string is optional.

AppInfo : <Port number:1024-65535> [reply string]

Ex. AppInfo : 12345 "message"

\* There must be a single space between the port number and reply string.

\* The reply message string cannot contain spaces.

\* If the scripts were installed via the steeleye-1kHOTFIX-Gen-LB-PL-7172 rpm package, then they will be located in the /opt/LifeKeeper/SIOS\_Hotfixes/Gen-LB-PL-7172 directory.

## 6.4.6. Messages List

This section provides a list of messages from the scripts. Each message is logged in the event log. DEBUG messages can be output by setting the debug flag or debug\_gen-lb flag.

Code	Category	Message	Description
127001	ERROR	Daemon already started.	Restore process failed because the daemon is already started.
127002	ERROR	Socket open failed.	Opening Socket failed (e.g., the specified port is already used).
127003	ERROR	Health probe not received.	During the restore process, no probe from a Load Balancer was received within \$HC_TIMEOUT seconds.
127004	DEBUG	Health check received.	During the restore process, a probe from a Load Balancer was received.
127005	ERROR	Daemon stat failed.	Failed to start the daemon.
127006	ERROR	Health probe not received.	While the daemon is waiting for a probe, no probe is received from a Load Balancer for \$HC_TIMEOUT seconds.
127007	DEBUG	Health check received.	While the daemon was waiting for a probe, a probe from a Load Balancer was received.
127008	INFO	Daemon already stopped.	During the remove process, the daemon has already been stopped.
127009	ERROR	Resource out of service.	During quickCheck, the resource is not In Service (ISP).
127010	ERROR	Daemon not running.	During quickCheck, the daemon is not running.
127011	ERROR	Failed to accept the health probe:\$!	Failed to accept a probe for some reason.
127012	ERROR	Port number is not valid.	Restore failed because the specified port is not available.

## 6.5. LVM Recovery Kit Administration Guide

---

The LifeKeeper for Linux Logical Volume Manager (LVM) Recovery Kit provides logical volume support for other LifeKeeper Recovery Kits. Thus, LifeKeeper-protected applications can take advantage of the benefits offered by the Logical Volume Manager, including simplified storage management and the ability to dynamically re-size volumes as needs change.

The LVM Recovery Kit is different from most other LifeKeeper Recovery Kits in that it is never used alone but always as a dependency of another LifeKeeper resource. As such, many of the operations typically associated with a LifeKeeper Recovery Kit – for example, creating a hierarchy – are not directly applicable to the LVM Recovery Kit.

The logical volume subsystem within the Linux OS (LVM) has a mirror feature. This feature is not supported by the LVM Recovery Kit.

### Document Contents

This guide explains the following topics:

- [Documentation and References](#). Provides a list of related LifeKeeper for Linux documents and where to find them, along with references to a number of helpful documents about the LVM product.
- [Requirements](#). Describes the hardware and software necessary to properly set up, install and operate the LVM Recovery Kit. Refer to the LifeKeeper for Linux Installation Guide for specific instructions on how to install or remove the LifeKeeper for Linux software.
- [Overview](#). Provides a general description of the LVM Recovery Kit and corresponding resource types.
- [LifeKeeper LVM Hierarchy Creation and Administration](#). Includes a detailed description of LVM Recovery Kit administration tasks through LifeKeeper.
- [Troubleshooting](#). Provides a list of informational and error messages with recommended solutions.

## 6.5.1. LVM Documentation and References

---

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Installation Guide](#)
- [Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is available on the SIOS Technology Corp. website at:

<http://docs.us.sios.com>

For information on LVM, refer to:

- [LOGICAL VOLUME MANAGER ADMINISTRATION Red Hat Enterprise Linux 7](#)
- [CONFIGURING AND MANAGING LOGICAL VOLUMES Red Hat Enterprise Linux 8](#)
- [Logical Volumes \(LVM | Storage Administration Guide | SUSE Linux Enterprise Server 12 SP5](#)
- [Logical Volumes \(LVM | Storage Administration Guide | SUSE Linux Enterprise Server 15 SP2](#)
- [LVM HOWTO](#) (This document is out-of-date)

## 6.5.2. LVM Recovery Kit Requirements

---

Your LifeKeeper configuration must meet the following requirements prior to the installation of the LifeKeeper for Linux LVM Recovery Kit. Please see the [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the configuration of your LifeKeeper for Linux hardware and software.

## 6.5.2.1. LVM Hardware and Software Requirements

---

### Hardware Requirements

- **Servers.** This recovery kit requires two or more computers configured in accordance with the requirements described in the [LifeKeeper for Linux Release Notes](#) and the [LifeKeeper for Linux Installation Guide](#), which are shipped with the product media.
- **Data Storage.** The LVM Recovery Kit can be used in conjunction both with shared storage and with replicated storage provided by the DataKeeper for Linux product. It cannot be used with network attached storage (NAS). Otherwise, the kit has no specific requirements on storage configurations beyond the requirements of the recovery kit protecting the application sitting on top of the logical volume(s).

### Software Requirements

- **Operating System.** LVM is included in all major Linux distributions. See the [LifeKeeper for Linux Release Notes](#) for a list of supported distributions and LVM versions.
- **Logical Volume Manager.** The recovery kit installation requires that the `lvm orlvm2` rpm package be installed. This release of the LifeKeeper Logical Volume Manager Recovery Kit supports both LVM Version 1 and LVM Version 2 (LVM2). The specific versions of LVM supported are those delivered by the Linux distributions.
- **LifeKeeper Software.** You must install the same version of LifeKeeper core software and any recovery kits including the LVM Recovery Kit and any patches on each server. Please refer to the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.
- **LifeKeeper for Linux Logical Volume Manager Recovery Kit.** The Logical Volume Manager Recovery Kit is provided on the LifeKeeper Installation Image File (`sps.img`). It is packaged, installed and removed via the Red Hat Package Manager, rpm: `steeleye-lkLVM`.

During package installation, checks are made to ensure that supported versions of both the LifeKeeper Core package and the LVM package are present on the system where the LVM Recovery Kit is being installed. The [LifeKeeper for Linux Release Notes](#) contains information on the required versions of these packages.

Refer to the [LifeKeeper for Linux Installation Guide](#) for instructions on how to install or remove the LifeKeeper Core software and the LVM Recovery Kit.

The LVM Recovery Kit must be installed on each server in the cluster on which LVM is being used to manage disk resources that are to be protected by LifeKeeper.

The LVM Recovery Kit must be installed prior to the hierarchy creation and extension of applications that sit on top of an LVM volume.

## 6.5.3. LVM Recovery Kit Overview

### LVM Operation

LVM is currently the standard volume management product included with all of the major Linux distributions. LVM allows multiple physical disks and/or disk partitions to be grouped together into entities known as volume groups. Volume groups may then be divided or partitioned into logical volumes. Logical volumes are accessed as regular block devices and as such may be used by file systems or any application that can operate directly with a block device.

Logical volume managers are principally used to simplify storage management. Logical volumes can be resized dynamically as storage requirements change, and volume groups and logical volumes can be sensibly named with identifiers chosen by the administrator rather than physical disk or partition names such as `sda` or `sdc1`.

The following diagram shows the relationship of the LVM entities. File systems or applications use logical volumes. Logical volumes are created by partitioning volume groups. Volume groups consist of the aggregation of one or more physical disk partitions or disks.

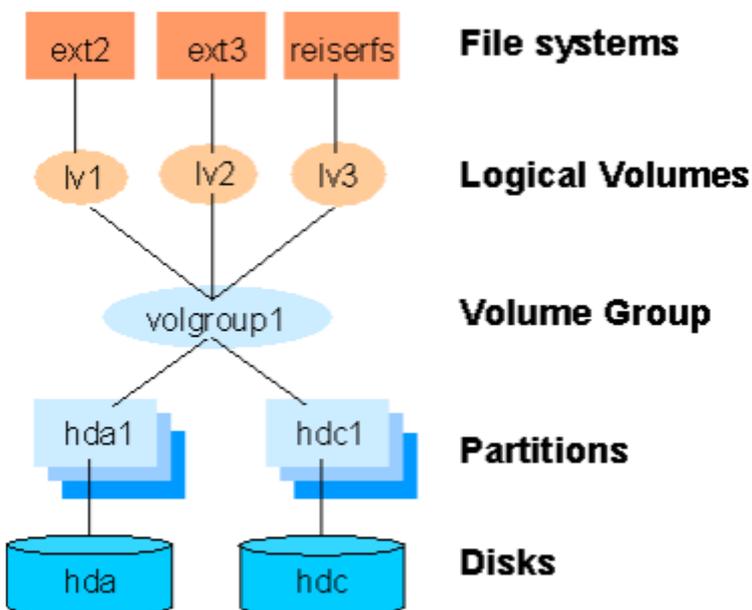


Figure 1: Logical Volume Manager Entity Relationships

### LifeKeeper for Linux LVM Recovery Kit

The LifeKeeper LVM Recovery Kit provides the support needed to allow other LifeKeeper recovery kits to operate properly on top of Linux logical volumes. To accomplish this support, the LVM Recovery Kit installs two new resource types: `lvmlv` and `lvmsg` which correspond to logical volumes and volume groups respectively. The `lvmlv` and `lvmsg` resources exist solely for internal use so that other LifeKeeper resources can operate.

As shown in Figure 1, each volume group has one or more logical volumes that depend on it.

Conversely, each logical volume must have a volume group on which it depends. A typical LifeKeeper hierarchy containing these two LVM resources looks much like the relationships shown in Figure 1. Refer to Figure 2 in the [LifeKeeper LVM Hierarchy Creation and Administration](#) section for an example of an actual LifeKeeper hierarchy.

The LVM Recovery Kit uses the commands provided by the lvm package to manage the volume group and logical volume resources in a LifeKeeper hierarchy. Volume groups and logical volumes are configured (or activated) when a hierarchy is being brought in service during a failover or switchover operation and are unconfigured when a hierarchy is being taken out of service.

## 6.5.3.1. LVM Recovery Kit Notes and Restrictions

---

The following notes and restrictions apply to this version of the LVM Recovery Kit.

### Support for Raw I/O and Entire Disks

While [Figure 1](#) shows logical volumes residing below various file systems and volume groups on top of disk partitions, it is important to note that the LVM Recovery Kit can support raw access to logical volumes when used in conjunction with the LifeKeeper Raw I/O Recovery Kit and can manage volume groups that are composed of one or more entire disks (e.g. `/dev/sdc`) rather than disk partitions (e.g. `/dev/sdc1`)

Also see the section [Using LVM with DataKeeper](#) for a further option in the use of LVM.

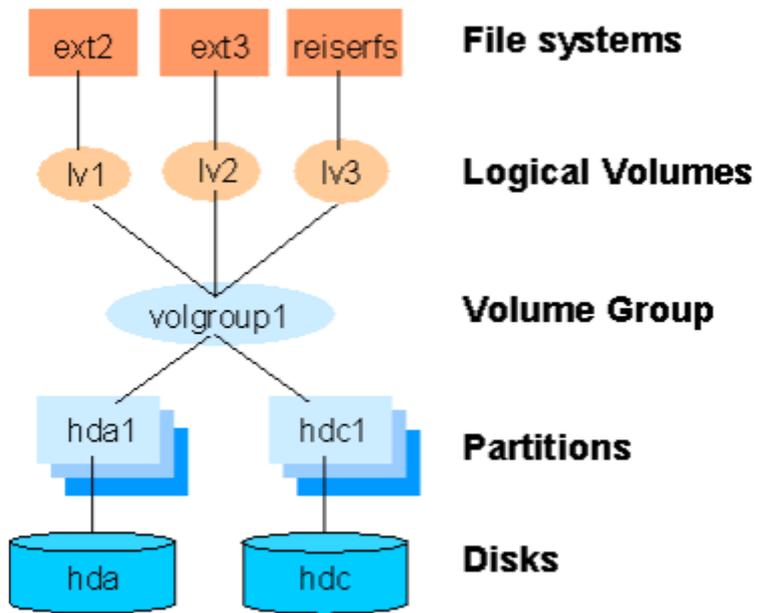
### Volume Group Activation

In the current LVM implementations, when a volume group is activated, all logical volumes associated with that volume group are also activated automatically. For LifeKeeper, this means that there will be times when a logical volume is active despite the fact that its associated resource instance is still marked as being Out-of-Service (OSU). In a typical failover or switchover operation, LifeKeeper will attempt to bring the logical volumes in service immediately after the volume groups anyway, and the resulting calls to the restore script will return immediately with a success indication. This unneeded attempt to bring the logical volumes in service has no usability impact.

### LVM mirroring functionality is not supported

The logical volume subsystem within the Linux OS (LVM) has a mirror feature. This feature is not supported by the LVM Recovery Kit.

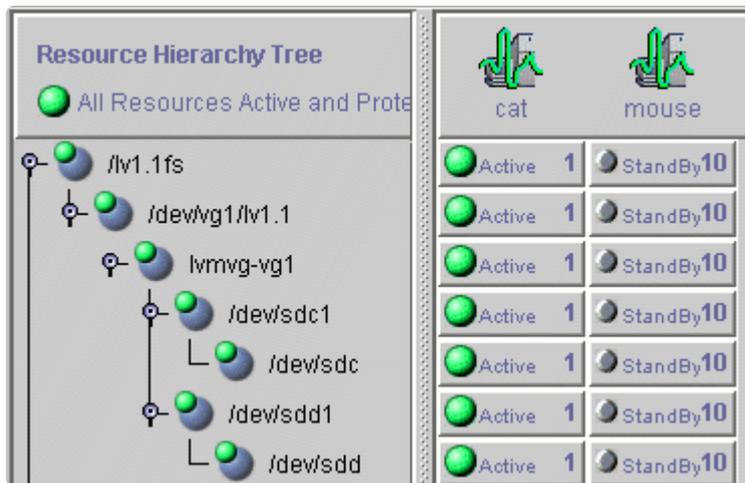
# LVM Figure 1



## 6.5.4. LifeKeeper LVM Hierarchy Creation and Administration

LifeKeeper LVM hierarchies are created automatically during the hierarchy creation process for resources that sit on top of logical volumes. The creation and extension of hierarchies containing the LVM resource types will always be driven by the create and extend processes of a higher-level resource type, likewise the delete and unextend.

The figure below is a LifeKeeper GUI screen shot showing a complete hierarchy containing LVM resources. Note that the resources in the hierarchy are displayed by their LifeKeeper IDs for clarity rather than the default display by tags.



**Figure 2: LifeKeeper Hierarchy Containing LVM Resources**

The hierarchy pictured in Figure 2 is a file system hierarchy created by selecting the File System Recovery Kit under the **Edit > Server > Create Resource Hierarchy** menu selection. It consists of a file system resource, `/lv1.1fs`, mounted on an LVM logical volume, `/dev/vg1/lv1.1`. That logical volume is a part of the `vg1` volume group represented with the LifeKeeper ID `lvmvg-vg1`. The volume group `vg1` is composed of two physical disk partitions, `/dev/sdc1` and `/dev/sdd1`. The hierarchy also includes the underlying disk devices, `/dev/sdc` and `/dev/sdd`, below each of the disk partitions.

## 6.5.4.1. LVM Hierarchy Creation Procedures

To create a hierarchy in which a file system or higher-level application uses an LVM logical volume, the following high-level procedure should be followed.

1. Determine the desired configuration of your LVM volume groups and logical volumes. In doing this, keep in mind the following points:
  - All of the disk resources associated with a given volume group must move together from one server to another in the LifeKeeper cluster.
  - All of the logical volumes associated with a given volume group (and any file systems or applications which use them) must move together from one server to another in the LifeKeeper cluster.
  - lvm2-lvmetad is disabled during the LVM RK installation. If you install lvm2 after installing the LVM RK, you need to disable it manually. Refer to the operating system documentation for more information.
2. On the system which is to be the primary server for your application, create and activate the desired volume groups and logical volumes using the tools provided by the LVM package and described in the *LVM HowTo* document referenced in the [Documentation and References](#) topic.

If you are using shared storage, you must ensure that all physical volumes assigned to a volume group are properly shared between the machines in the LifeKeeper cluster on which you intend to run the protected application. If you intend to use LVM with DataKeeper, see the [Using LVM with DataKeeper](#) topic.

3. Create file systems on each of the logical volumes. If instead you intend to use raw I/O, bind a raw device to each of the logical volume devices.
4. Configure the protected application on the file systems following the configuration instructions in the administration guide for the LifeKeeper recovery kit associated with the application.

Create and extend the application hierarchy following the instructions in the appropriate application recovery kit administration guide.

**!** **IMPORTANT:** Perform manual in-service operations to temporarily move the application hierarchy to each of the cluster nodes to which the hierarchy has been extended. This step must be done once prior to any node failover operations in order for the LVM subsystem on each cluster node to know about the configuration of the new volume groups and logical volumes. After you have performed these manual switchovers, move the application hierarchy back to the desired primary cluster node.

## 6.5.4.2. Using the LVM Recovery Kit with DataKeeper

LifeKeeper for Linux currently supports both the use of DataKeeper “above” LVM and LVM “above” DataKeeper. In a standard DataKeeper configuration, using DataKeeper above LVM is supported and DO NOT install the LifeKeeper LVM Recovery Kit. DataKeeper is the only recovery kit necessary. However, using the LVM above DataKeeper configuration, the LVM Recovery Kit is required.

SIOS recommends using DataKeeper above LVM; however, if the LVM above DataKeeper configuration is being used, a two-phase hierarchy creation process must be used. The DataKeeper devices (i.e. hierarchies) must be configured using the DataKeeper “Data Replication Resource” option prior to the creation of the LVM volume groups and logical volumes on the primary server. Once the desired volume groups and logical volumes have been created, the remainder of the hierarchy is created according to the configuration instructions for the recovery kit associated with the application to be protected. The resulting hierarchy will look something like the one shown in Figure 3 below.

**Note:** For data consistency reasons, in an LVM over DataKeeper configuration, there must either be only one DataKeeper mirror or multiple *synchronous* mirrors.

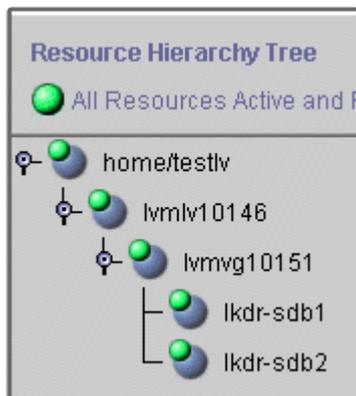


Figure 3: Hierarchy with LVM above DataKeeper

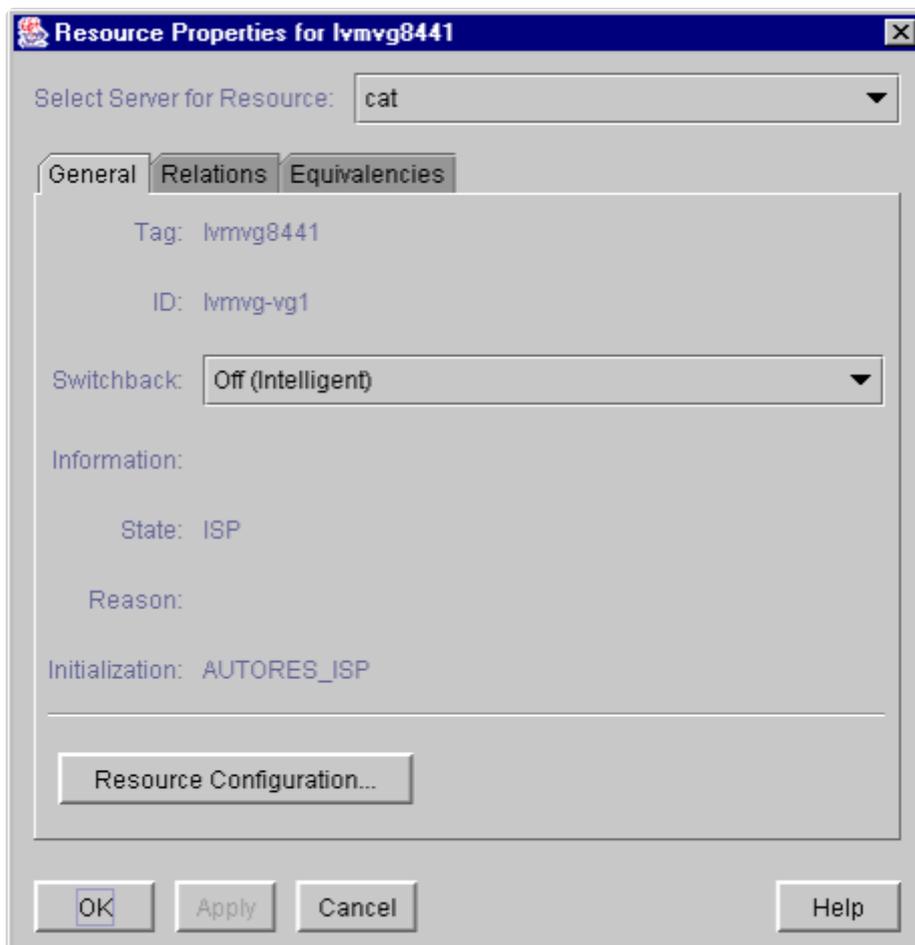
## 6.5.4.3. Volume Group Reconfiguration

One of the primary benefits of using a logical volume manager is the ability to dynamically resize logical volumes as storage requirements change. Because this may involve adding or deleting physical partitions or disks from an LVM volume group definition, the LVM Recovery Kit includes a mechanism for modifying an existing resource hierarchy to reflect such a change.

All volume group, logical volume and file system reconfiguration should be performed outside of LifeKeeper prior to modifying the LifeKeeper hierarchy to reflect the changes. Refer to the *LVM HowTo* document referenced in the [Documentation and References](#) section for information about how this is done. If any of the steps require you to **unmount** or **unconfigure** a resource that is being protected by LifeKeeper, be sure to use the LifeKeeper GUI to do so, using the **Out of Service** operation.

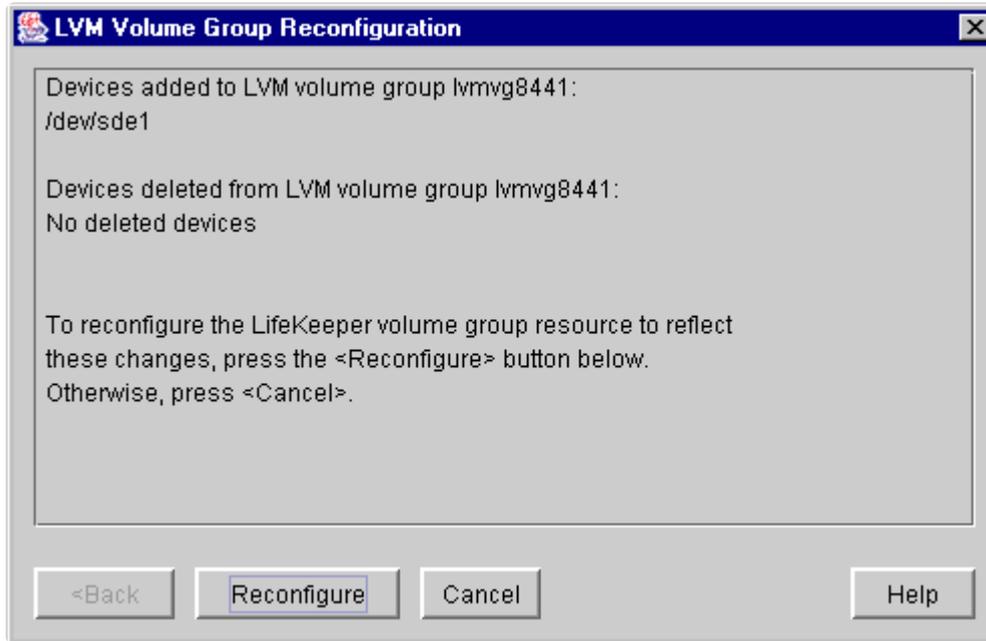
**! IMPORTANT:** The new device **MUST** be seen by both systems (shared) before LifeKeeper will allow the reconfiguration to take place.

To update a LifeKeeper hierarchy following these changes, first access the **Resource Properties** dialog for the modified volume group, either by right-clicking on the active volume group resource and selecting **Properties** or by using the **Edit > Resource > Properties** menu selection and selecting the appropriate volume group resource in the **Select Resource** field. The resulting **Resource Properties** dialog should look like the one pictured in Figure 4 below including the **Resource Configuration** button near the bottom.



#### Figure 4: LVM Volume Group Resource Properties Dialog

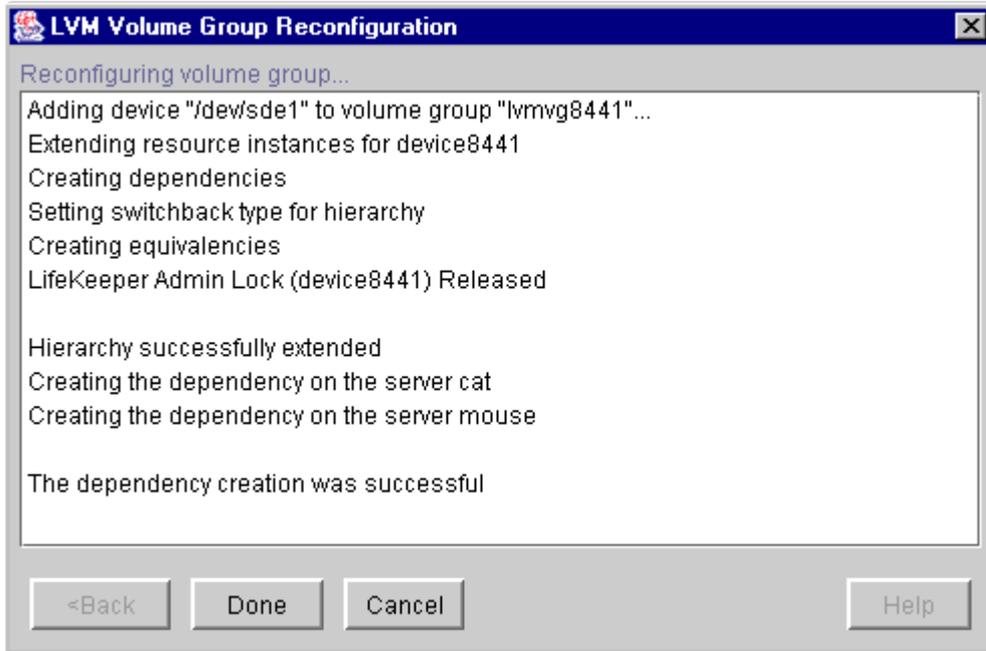
Clicking the **Resource Configuration** button initiates the mechanism for reconfiguring your hierarchy to reflect any modifications to the volume group resource. After a brief pause, an information box will display the volume group modifications that LifeKeeper has detected. Figure 5 below shows an example in which a single disk partition has been added to a volume group.



#### Figure 5: LVM Volume Group Reconfiguration for Added Device

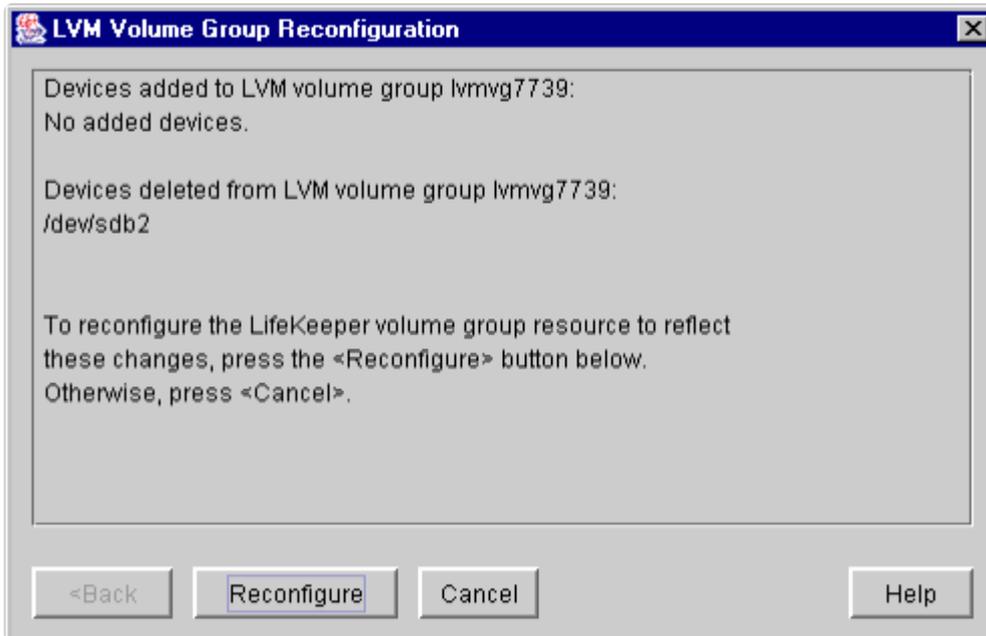
As stated in the information box, to reconfigure the LifeKeeper volume group to reflect the changes that have been detected, simply click the **Reconfigure** button. If you do not wish to proceed with the LifeKeeper hierarchy modification, click **Cancel**.

After clicking the **Reconfigure** button, an information box will appear showing the progress of the reconfiguration procedure as shown in **Figure 6** below. When the process has been completed successfully, the **Done** button will become enabled. Clicking **Done** will close the information box and return you to the display of the **Resource Properties** dialog.

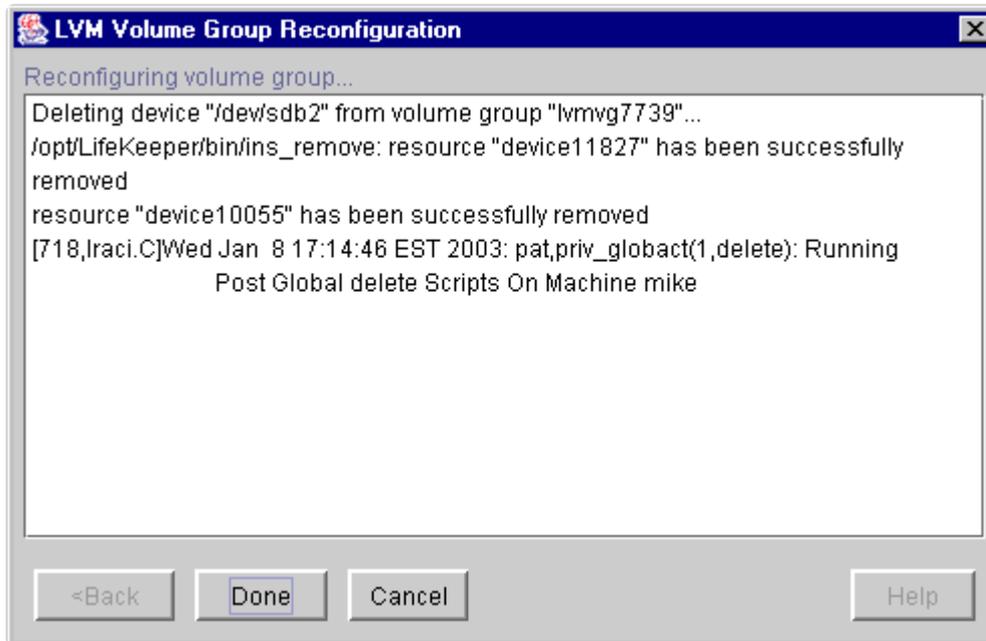


**Figure 6: LVM Volume Group Reconfiguration for Added Device**

The following two figures show examples of the information boxes that would be displayed during the reconfiguration process when a device partition has been removed from a volume group.



**Figure 7: LVM Volume Group Reconfiguration for Deleted Device**



**Figure 8: LVM Volume Group Reconfiguration for Deleted Device**

## 6.5.5. LVM Troubleshooting

### Error Messages

This section provides a list of messages that you may encounter with the use of the LifeKeeper LVM Recovery Kit. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Because the LVM Recovery Kit relies on other LifeKeeper components to drive the creation and extension of hierarchies, messages from these other components are also possible. In these cases, please refer to the [Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

### LVM Recovery Kit Error Messages

Error Number	Error Message
110000	<LVM resource type> resource type is not installed on <LifeKeeper server name>. <b>Action:</b> Install the LVM Recovery Kit on the identified system.
110001	This script must be executed on <LifeKeeper server name>.
110002	Failed to create <device name> hierarchy.
110003	Failed to create dependency <resource tag>-<resource tag> on machine <LifeKeeper server name>.
110004	LifeKeeper internal ID <resource ID> already in use.
110005	<LVM resource type> constructor requires a valid argument.
110006	Usage: adddelpv <VG tag> [addlist dellist]
110007	WARNING: Failure in updating list of LifeKeeper-controlled volume groups (/etc/lkvgs)
110008	WARNING: The device hierarchy for <device name>, with tag <device resource tag>, cannot be extended automatically. <b>Action:</b> If the device hierarchy is not already extended, extend it using the LifeKeeper GUI. Then create a dependency from the volume group resource to the device hierarchy.
110009	Failed to create a dependency between volume group resource <volume group tag> and device resource <device tag>. <b>Action:</b> Create the dependency using the LifeKeeper GUI.

110010	Failed to make the LVM logical volume <Logical Volume Path> active with error code <error code>.
110011	Failed to vgscan the LVM volume group <Volume Group Name> with error code <error code>.
110012	Failed to vgimport the LVM volume group <Volume Group Name> with error code <error code>.
110013	Failed to make the LVM volume group <Volume Group Name> active with error code <error code>.

## 6.6. IP Recovery Kit Administration Guide

---

The LifeKeeper for Linux Internet Protocol (IP) Recovery Kit provides a mechanism to recover an IP address from a failed primary server to a backup server in a LifeKeeper environment. The IP Recovery Kit can define an IP address that can be used to connect to a LifeKeeper-protected application. As with other LifeKeeper resources, IP resource switchovers can be initiated automatically as a result of a failure or manually by an administrative action.

The IP Recovery Kit supports the implementation of the TCP/IP protocol suite using secondary addresses on existing network interfaces, allowing it to provide switchover and failover of IP addresses without requiring extra standby network interface cards or *dummy* IP addresses. Starting with Release 7.4, the IP Recovery Kit supports both IPv4 and IPv6 addresses.

The following LifeKeeper product documentation is available from the SIOS Technology Corp. website:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [Optional Recovery Kit Documentation](#)

## 6.6.1. IP Recovery Kit Principles of Operation

---

### Virtual IP Resource

LifeKeeper brings an IP resource into service by creating an IP alias address on one of the physical network interfaces on the primary server. Users connect to the node using this alias address.

The IP Recovery Kit software performs checks to help ensure that the selected address, network mask and interface can function properly. The software verifies the following elements:

- **Unused resource.** The new IP address is not already assigned to any other IP resource in the LifeKeeper cluster.
- **Unique address.** The address cannot be currently active on the network. In addition to checking during creation, the software also performs the uniqueness check immediately before bringing the resource into service. If the software detects a duplicate address on the net, it does not bring the resource into service.

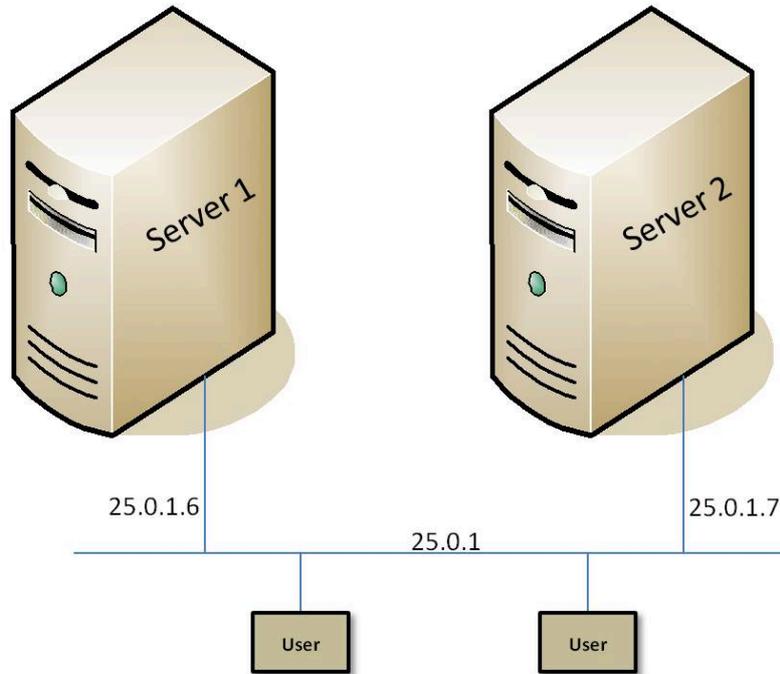
When the primary server fails, the IP Recovery Kit brings the IP resource into service on a backup server by configuring the IP alias on one of that server's physical network interfaces.

Since session context is lost following recovery, after the recovery, IP users must reconnect using exactly the same procedures they used to connect originally.

In a manual switchover, the IP Recovery Kit removes the alias address from service on the active server before adding it to the backup server.

To clarify the administration and operation of the IP Recovery Kit, consider the scenario shown in Figure 1. This example configuration contains two servers, Server 1 and Server 2. Each server has a single LAN interface, eth0, connected to subnet 25.0.1. The user systems are also on this subnet. The LAN interfaces on Server 1 and Server 2 have addresses 25.0.1.6 and 25.0.1.7, respectively.

**Figure 1. Administration and Operation Scenario**



The system administrator decides to use 25.0.1.10 as the alias address for an IP resource, to be called *ipname*. The administrator creates entries in the */etc/hosts* files (and in the DNS, if used), similar to the following:

25.0.1.6	server1
25.0.1.7	server2
25.0.1.10	ipname

Assuming that Server 1 is the primary server for the resource, the administrator creates the IP resource hierarchy for *ipname* on Server 1 using the wizard described in the section entitled [Creating an IP Resource Hierarchy](#). The software finds the address associated with *ipname* (25.0.1.10) from */etc/hosts*, verifies that it is available and brings it into service by configuring a secondary address on eth0 on Server 1. eth0 on Server 1 now responds to both *server1* and *ipname*.

With LifeKeeper 7.3 or earlier, the new alias address can be verified using the `ifconfig` or `ip addr show` command. Starting with LifeKeeper 7.4, the `ip addr show` command should be used (for more information, see the [IPv6 Known Issue](#)).

Users can then connect to Server 1 by entering, for example, `telnet ipname`. If Server 1 crashes, LifeKeeper automatically switches over the *ipname* address to eth0 on Server 2. The user sessions on Server 1 terminate. When users re-run `telnet ipname`, they are connected to Server 2.

Regardless of where *ipname* is actively in service, addresses *server1* and *server2* are active and usable, though not protected by LifeKeeper recovery. The addresses could be used for any cases that require connection to a specific server by name rather than to a switched application. Examples might include remote system management and the LifeKeeper communications path. (In this case, for example, 25.0.1.6 and 25.0.1.7 would be used for the LifeKeeper communications path.)

## Actual IP Resource

LifeKeeper can protect not only the virtual IP address but also an actual IP address (i.e Primary IP address which is configured for the network interface). This allows you to configure without the virtual IP address in the Amazon Web Services (AWS) environment by using Recovery Kit for Route 53. Refer to the [Recovery Kit for Route 53™ Administration Guide](#) for more information.

## IP Resource Monitoring

LifeKeeper monitors the health of the IP resources under its control on a periodic basis, using the following techniques, in this order.

1. Check the link status for the network interface on which the IP resource is configured to determine whether the interface is properly connected to the physical network.
2. Verify that the IP resource is still configured as an alias on the appropriate network interface.
3. Perform a broadcast ping test or ping a pre-configured list of addresses, using the protected IP address as the source address of the pings, to determine whether the IP resource can successfully send and receive data on the network.

The broadcast ping test is the default test mechanism. It operates by sending a broadcast ping packet to the broadcast address of the subnet associated with the IP resource, using the protected IP address as the source address. If a response is received from any address other than addresses on the local system, the test is considered successful.

For environments in which there are no systems on the network that can respond to the broadcast ping test (which is the default configuration of many systems), LifeKeeper also offers the ability to configure a list of addresses to be pinged as an alternative to the broadcast ping test. If such a list has been specified, the broadcast ping test is skipped, and all of the addresses in the list are pinged in parallel. The test is considered successful if a ping response is received from any one of the addresses in the Ping List. This technique is also useful to reduce broadcast storms on larger networks.

If any of these tests fail during the periodic health check of an IP resource, LifeKeeper is notified of the failure. LifeKeeper will first attempt a local recovery operation to try to restore the IP resource to a working state on the local node. See the section [IP Local Recovery and Configuration Considerations](#) for more information about the local recovery procedure. If local recovery is unsuccessful in restoring the IP resource to a working state, LifeKeeper will then attempt to migrate the application hierarchy containing the IP resource to another LifeKeeper system in the cluster.

 **Note:** When the actual IP resource is used, the local recovery attempt is skipped and forcibly determined as failed. LifeKeeper will then perform a failover of the IP resource and all dependent resources to a backup server.

LifeKeeper also uses these same health checks to verify the proper operation of an IP resource

immediately after it is brought in-service. A failure of any of the checks will cause the in-service operation to fail.

The IP health check mechanisms can be tuned and adjusted in many ways. See the sections [Viewing/Editing IP Configuration Properties](#) and [Adjusting IP Recovery Kit Tunable Values](#) for details.

## 6.6.2. IP Recovery Kit Requirements

---

Before attempting to install or remove the IP Recovery Kit, you must understand the hardware and software requirements for the package and the installation and removal procedures.

### Hardware and Software Requirements

Before installing and configuring the LifeKeeper IP Recovery Kit, be sure that your configuration meets the following requirements:

- **Servers.** The recovery kit requires two or more supported computers configured in accordance with LifeKeeper requirements described in the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#).
- **LifeKeeper software.** You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.
- **LifeKeeper IP Recovery Kit.** You must have the same version of this recovery kit on each server.
- **IP network interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications. This interface should be configured. If there are no *ifcfg\** files, IP switchover may fail when the interface is down.

**Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons; for example, heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation and local recovery support. Also, set an actual IP address for NIC used for the IP resource setting. By using the actual IP address, confirm the communication on IP network.

For the configuration of channel bonding and network teaming that have been tested, please click [here](#).

- **TCP/IP software.** Each server also requires the TCP/IP software.

Consult the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information.

You should refer to the [LifeKeeper Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper IP Recovery Kit.

## 6.6.3. IP Recovery Kit Configuration

---

To ensure that your LifeKeeper configuration provides the protection and flexibility you require, you need to be aware of the configuration rules. To appropriately plan your configuration, you must understand your network configuration, interface selection, user system setup, hierarchy options and the IP configuration tasks. In addition to planning your configuration, this section also includes configuration examples and the specific tasks required to configure your recovery kit.

### Configuring TCP/IP with LifeKeeper

This section contains information you should consider before you start to configure TCP/IP and examples of typical LifeKeeper IP configurations.

Please refer to the [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring your LifeKeeper Core resource hierarchies.

### Specific Configuration Considerations for TCP/IP

In order to properly configure your IP Recovery Kit, you should review the following topics to ensure that you have the information necessary to complete the configuration tasks:

- [Interface Selection](#)
- [User System Setup](#)
- [General IP Planning Considerations](#)

See the following topics for further configuration considerations and examples:

- [IP Resource Monitoring and Configuration Considerations](#)
- [IP Local Recovery and Configuration Considerations](#)
- [Configuration Examples](#)

### LifeKeeper Configuration Tasks

The following configuration tasks for IP resources are described in this section, as they are unique to an IP resource instance and different for each recovery kit.

- [Creating an IP Resource Hierarchy](#). Creates an application resource hierarchy in your LifeKeeper cluster.
- [Deleting a Resource Hierarchy](#). Deletes a resource hierarchy from all servers in your LifeKeeper cluster.

- [Extending Your Hierarchy](#). Extends a resource hierarchy from the primary server to a backup server.
- [Unextending Your Hierarchy](#). Unextends (removes) a resource hierarchy from a single server in your LifeKeeper cluster.
- [Testing Your Resource Hierarchy](#). Tests an IP resource hierarchy for proper configuration and operation.
- [Viewing/Editing IP Configuration Properties](#). Displays configuration details for an IP resource and allows some of them to be modified.
- [Adjusting IP Recovery Kit Tunable Values](#). Tunes characteristics of the overall behavior of the IP Recovery Kit.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View](#) / [Edit](#) Properties. View or edit the properties of a resource hierarchy on a specific server.

**Note:** Throughout the rest of this section, we explain how to configure your recovery kit by selecting certain tasks from the **Edit** menu of the LifeKeeper GUI. You can also select each configuration task from the toolbar. You may also right-click on a global resource in the Resource Hierarchy Tree (left-hand pane) of the status display window to display the same drop down menu choices as the **Edit** menu. This, of course, is only an option when a hierarchy already exists.

You can also right-click on a resource instance in the Resource Hierarchy Table (right-hand pane) of the status display window to perform all the configuration tasks, except Creating a Resource Hierarchy, depending on the state of the server and the particular resource.

## 6.6.3.1. IP Interface Selection

---

When creating an IP resource, select the IP resource address, the netmask to use with the address and the network interface. Not all combinations are allowed. The address/netmask pair provided and all the address/netmask pairs currently in-service determine choices. Also, see the section on [IP Local Recovery](#) for additional configuration considerations if planning on using this feature of the recovery kit.

The selected address/netmask determines the subnet for the resource. If another address on the same subnet (either a physical or logical interface address) is currently in service on any interface, then the IP resource must be configured on that interface. The software performs tests to determine the allowed choices based upon the current network configuration. Select from any of the choices provided.

Because the IP Recovery Kit software does not distinguish between physical media types, the physical network for the resource must be determined and the address selected appropriately. For example, assume that you have a server connected to an Ethernet backbone on subnet xx.yy.12 and Ethernet LANs on subnets xx.yy.20 and xx.yy.30. If you want to create a resource on the first Ethernet subnet, select an address on that subnet, such as xx.yy.20.120.

In general, even though the IP Recovery Kit software allows you to select almost any value for the netmask, you should avoid selecting multiple netmasks for the same physical interface because multiple masks can cause packet misrouting.

One further consideration is the need to be consistent in your selection of interfaces on all LifeKeeper servers. If you configure several IP resources on a single interface on Server A, they should also be configured on a single interface on Server B.

When creating an IP resource hierarchy, you may utilize any interface which is initially UP and has a corresponding and correct network interface configuration file on both the primary and backup hosts, i.e. if using *eth1*, *eth1* must be UP and *eth1* must have a corresponding and correct *ifcfg-eth1* file (test with `ifup/ifdown ifcfg-eth1`) even if the configuration is minimal without any address assignments or is DOWN on boot.

## 6.6.3.2. IP User System Setup

---

When the IP Recover Kit software switches an IP resource from one server to another, the MAC address associated with the switched IP address changes because the interface changes. Each router and user system on the LAN must reflect this change in its ARP table before it can contact the IP address at its new location. In certain operating systems, when a new IP address is added to a network interface, an ARP packet is automatically sent out by the operating system to update all clients' ARP tables on the subnet. This feature does not exist in Linux. LifeKeeper therefore must send out an ARP packet after adding a switchable IP address to an interface to force this client ARP cache update.

TCP/IP implementations differ in their ability to implement the required ARP updates in response to this ARP packet. The following list describes some important cases:

- **Full Linux TCP/IP implementation.** Fully functional TCP implementations in Linux and most other operating systems support ARP cache updates when the systems receive an ARP request packet. LifeKeeper uses this feature, as described above, to force ARP cache updates on such systems.
- **ARP cache.** User systems that do not support the ARP refinements but do support an ARP cache usually have a timer associated with the cache to maintain some level of currency. For some implementations, decreasing the timer value can minimize the time required for that particular user system to reflect the changed address mapping. If the number of users on the LAN is small, this option may be acceptable. For other systems, decreasing the timer value may not be necessary. For example, the TCP implementation shipped with Windows NT uses a ten second timer value, so no change in timer value would be needed.
- **Static address mapping.** For systems without a dynamic ARP cache or those where cache timing is not tunable, routers can be used to handle mapping changes. Such user systems would access the IP resource subnet by way of a router (gateway). In this configuration, cache update is needed only for the routers directly connected to the resource subnet and no changes are needed on the user systems themselves.

## 6.6.3.3. General IP Planning Considerations

After you have selected the addresses, netmasks and associated host/domain names you intend to use for IP resource hierarchies, add the appropriate entries to each server's */etc/hosts* file, and to the Domain Name Server (DNS), if used.

 **Note:** Even if you are using a DNS, it is strongly recommended that you place entries for the IP resources in the local */etc/hosts* files on all LifeKeeper servers. This will reduce recovery times. However, if the resource name that you enter when creating the IP instance is the IP address itself, then the host file entry is unnecessary.

Do not configure the protected IP addresses into your system as you would if you were creating a permanent logical interface to be activated at system boot time. The LifeKeeper software will manage them instead of the system software.

If any of the resource addresses are on new (logical) subnets, update routers to handle routing to these subnets.

## 6.6.3.4. IP Resource Monitoring and Configuration

---

By default, the LifeKeeper IP Recovery Kit monitors IP resources by executing a broadcast ping on the IP addresses logical subnet, then listening for replies. For this test to work properly, at least one additional non-LifeKeeper system capable of responding to broadcast pings must exist on the physical network, with an IP address on the same logical subnet as the IP resource. A router on the same logical subnet is usually sufficient to meet this need. Note that the default configuration of many devices is to not respond to broadcast pings, so it may be necessary to change the configuration of at least one device.

If this requirement cannot be met, you can choose to either disable the broadcast ping test completely, or you can configure a static list of IP addresses that should be pinged as an alternative to the broadcast ping test mechanism. See the [Adjusting IP Recovery Kit Tunable Values](#) and [Viewing/Editing IP Configuration Properties](#) topics for more information about how to configure these options.

## 6.6.3.5. IP Local Recovery and Configuration

The standard Linux NIC bonding mechanism is the recommended means of providing network interface redundancy in a high availability configuration. The LifeKeeper IP Recovery Kit fully supports the creation of virtual IP addresses on bonded interfaces.

### Local Recovery Scenario

When a failure of an IP address is detected by the IP Recovery Kit, the resulting failure triggers the execution of the IP local recovery script. LifeKeeper will first attempt to bring the IP address back in-service on the current network interface. If the local recovery attempt fails, LifeKeeper will perform a failover of the IP resource and all dependent resources to a backup server.

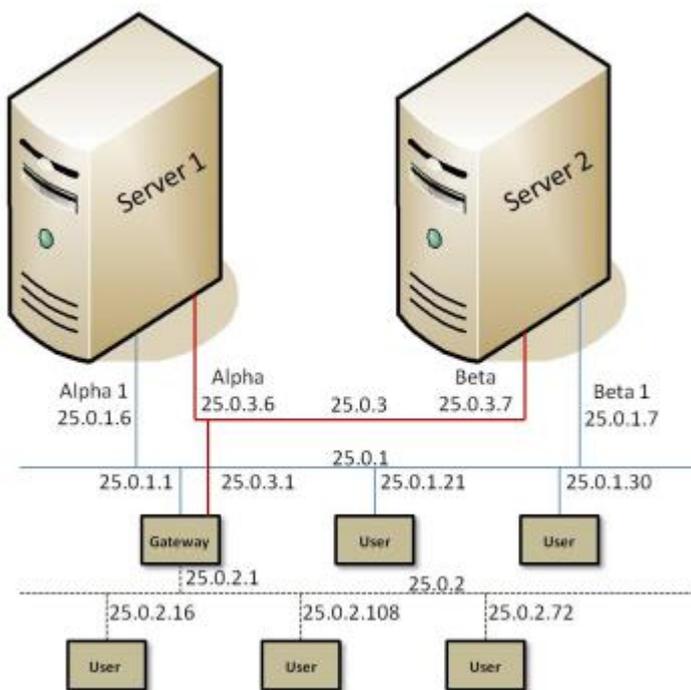
 **Note:** When the actual IP resource is used, the local recovery attempt is skipped and forcibly determined as failed. LifeKeeper will then perform a failover of the IP resource and all dependent resources to a backup server.

# 6.6.3.6. IP Recovery Kit Configuration Examples

This topic identifies example network configurations and then describes two sample IP configuration exercises. The first example illustrates a typical case of a database application dependent upon a single IP resource and configured on a pre-existing subnet. The second example illustrates an active/active scenario where multiple IP resources are configured.

## Network Configuration

The first two configuration examples assume the network configuration diagrammed in the following figure.



The network configuration has these components:

- **Servers.** The configuration has two servers, Server 1 and Server 2, each with the appropriate LifeKeeper and application software installed.
- **Interfaces.** Each server has two Ethernet interfaces, eth0 and eth1, configured as follows:

Interface	Server 1	Server 2
eth0	Server1 25.0.3.6	Server2 25.0.3.7
eth1	Server11	Server21

	25.9.1.8	25.0.1.7
--	----------	----------

- **Network.** The network consists of three subnetworks:
  - Low traffic backbone (25.0.3) primarily for servers
  - High traffic backbone (25.0.1) with both servers and clients
  - High traffic client network (25.0.2.)

A gateway provides interconnection routing between all LANs. A Domain Name Server (not shown) is used for address resolution.

- **Heartbeat.** TCP heartbeat communication paths would be configured using either or both of the server subnetworks.

## Typical Configuration Example

Server 1 and Server 2 have access to an application called mydatabase that resides on a shared disk. To ensure that the application mydatabase and the IP resources used to access it are switched together, the system administrator creates a mydatabase application resource and adds the IP resource to the application hierarchy as a dependency.

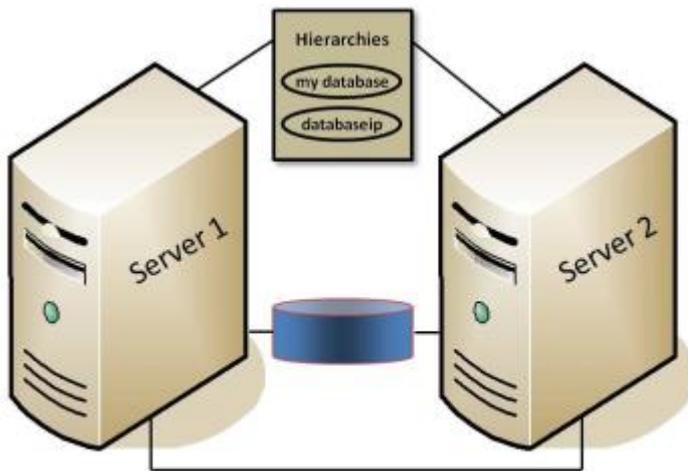
These are the configuration issues:

- **Application hierarchy.** The application hierarchy must exist before the administrator names it as a parent of the IP resource. For the purposes of this example, Server 1 is the primary server. The application resource tags are mydatabase-on-server1 and mydatabase-on-server2.
- **IP resource name.** The administrator adds the name and address of the IP resource to the `/etc/hosts` file on both Server 1 and Server 2 and to the DNS database. In this example, the IP resource name is databaseip and its network address is 25.0.1.2. If no name-to-IP address association is necessary, then this is not required.
- **Routers, gateways, and users.** Because databaseip is an address on an existing subnet, no additional configuration is necessary. The IP resource is on the 25.0.1 subnet. All users connect to databaseip via the route they currently use to get to the 25.0.1 subnet. For example, users on 25.0.2 go through the gateway and users on 25.0.1 connect directly.
- **IP instance definition.** When the administrator enters databaseip as the IP resource on the Resource Hierarchy Create screen, the software performs several tests. It verifies that Server 1 can determine the address that goes with databaseip (it is in the hosts file and/or can be retrieved from the DNS). It also verifies that the address retrieved, address 25.0.1.2, is not already in use. Since the IP resource is on the 25.0.1 subnet, the IP Recovery software will ensure that it is configured on the eth1 interface. If the IP resource is acceptable, the software fills in the remainder of the wizard dialog boxes with default values, as shown in the table below Figure 3. If you selected all the default values, an independent IP resource hierarchy called ip-databaseip

would be created.

**Note:** The tables associated with each configuration illustration provide examples of the appropriate information that would be entered in the Create Resource Hierarchy wizard for the primary server (Server 1) and Extend Resource Hierarchy wizard for the backup server (Server 2). For additional details on what information should be entered into the wizards, refer to the [LifeKeeper Configuration Tasks](#) section later in this section. These tables can be a helpful reference when configuring your recovery kit.

**Figure 3. Typical Configuration Example of IP Resource Creation**



**Configuration Notes:**

1. The application resource is mydatabase-on-server1.
2. The IP resource is databaseip with a tag name of ip-databaseip.
3. If mydatabase-on-server1 fails, LifeKeeper switches it to Server 2; (ip-databaseip is only switched if a dependency exists).
4. If Server 1 fails, both resources are brought in-service on Server 2.
5. During a switchover, databaseip users would be disconnected. When they log back in, they can access any applications on Server 2.
6. During a manual switchover, users connected to Server 1 via connections other than databaseip remain connected to Server 1.

**Creating an IP resource hierarchy on Server 1:**

Server:	Server1
IP Resource:	databaseip
Netmask:	255.255.252.0

Network Interface:	eth1
IP Resource Tag:	ip-databaseip

### Extending an IP resource hierarchy to Server 2:

Template Server:	Server1
Tag to Extend:	databaseip
Target Server:	Server2
Target Priority:	10
**IP Resource:	25.0.1.2
Netmask:	255.255.252.0
Network Interface:	eth1
IP Resource Tag	ip-databaseip

**Note:** The actual IP address associated with the DNS name is displayed in the *Extend Wizard* as the IP resource.

## Test Your IP Resource

To verify the successful creation of the IP resource, the administrator should perform the following tasks:

1. From the LifeKeeper GUI, observe whether ip-databaseip is in-service (ISP) on Server 1.
2. From a remote server, connect to address databaseip using ping or telnet.
3. Test manual switchover by selecting the in\_service option on Server 2 and selecting ip-databaseip. Verify that the IP address migrates to Server 2.

## Active/Active Configuration Example

The second example, using the same network configuration, describes two IP resources, one active on each server.

### Resource Addresses

For this example, the IP resources are server1ip (address 25.0.6.20) and server2ip (address 25.0.6.21). Entries for these resources must be in the /etc/hosts files on each server and in the DNS database.

### Router Configuration

Because the selected addresses are on a new (logical) subnet, they can be configured for either eth0 or eth1. However, both must go on the same interface.

For this example, choosing eth0 means that all users would have to go through the gateway. Choosing eth1 would allow the users on the 25.0.1 subnet to access the resources directly (assuming that the new subnet had been added to their internal routing tables). Users on subnet 25.0.2 would still require the gateway. For the purposes of this example, the selected interface is eth1.

Regardless of which physical network is chosen to support the new subnet, the network administrator would have to add routing information to the gateway system before creating the IP resources.

## First IP Resource Definition

The administrator creates the first IP resource on Server 1. eth0 is the first available interface on each server and would appear as the default. To define eth1 as the interface, the administrator selects it from the list of available interfaces.

### Creating an IP resource hierarchy on Server 1:

Server:	Server1
IP Resource:	server1ip
Netmask:	255.255.252.0
Network Interface:	eth1
IP Resource Tag:	ip-server1ip

### Extending an IP resource hierarchy to Server 2:

Template Server:	Server1
Tag to Extend:	server1ip
Target Server:	Server2
Target Priority:	10
**IP Resouce:	25.0.6.20
Netmask:	255.255.252.0
Network Interface:	eth1
IP Resource Tag:	ip-server1ip

**Note** The actual IP address associated with the DNS name is displayed in the *Extend Wizard* as the IP resource.

## Second IP resource definition

The administrator creates the second IP resource on Server 2. eth0 is the first available interface on each server and would appear as the default. To define eth1 as the interface, the administrator selects it from the list of available interfaces.

### Creating an IP resource hierarchy on Server 2:

Server:	Server2
IP Resource:	server2ip
Netmask:	255.255.252.0
Network Interface:	eth1
P Resource Tag:	ip-server2ip

### Extending an IP resource hierarchy to Server 1:

Template Server:	Server2
Tag to Extend:	server2ip
Target Server:	Server1
Target Priority:	10
**IP Resouce:	25.0.6.21
Netmask:	255.255.252.0
Network Interface:	eth1
IP Resource Tag:	ip-server2ip

**Note:** The actual IP address associated with the DNS name is displayed in the *Extend Wizard* as the IP resource.

 **Note:** Since subnet 25.0.6 is not active on Server 2, both eth0 and eth1 are available choices for the Primary network interface. On Server 1 (the backup server), the only choice is eth1 because the first IP resource, 25.0.6.20, is in service there. When the administrator saves the definition, LifeKeeper brings address 25.0.6.21 in-service on eth1 on Server 2.

## Testing IP Resources

The administrator should verify that the new resources are functioning on both servers by performing the following tests:

1. With each resource on its primary server, verify that each is accessible by using either ping or telnet. The administrator may also want to test connectivity from all user sites.
2. Test switchover by manually bringing ip-server1ip into service on Server 2. Verify both resources are functional on Server 2.
3. Bring both resources into service on Server 1. Verify both resources are functional on Server 1.
4. Bring ip-server2ip back into service on its primary server, Server 2.

## 6.6.3.7. Creating an IP Resource Hierarchy

To create a resource instance from the primary server, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop down menu, select **Create Resource Hierarchy**.
2. A dialog box will appear with a dropdown list box menu listing all recognized recovery kits installed within the cluster. Select **IP** from the dropdown list and click **Next**.
3. You will be prompted to enter the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful should you encounter an error that might require you to correct previously entered information.

If you click the **Cancel** button at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

Field	Tips
<b>Switchback Type</b>	<p>This dictates how the IP instance will be switched back to this server when the server comes back up after a failover. You can choose either <i>intelligent</i> or <i>automatic</i>. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths. The switchback type can be changed later from the General tab of the Resource Properties dialog box.</p>
<b>Server</b>	<p>Select the <b>Server</b> where you want to place the IP Address (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop down list.</p>
<b>IP Resource</b>	<p>Enter the virtual IP address to protect. This is the IP address or symbolic name that LifeKeeper will use for this resource.</p> <p>Not only a virtual IP address but also an actual IP address which is allocated for a network interface can be protected. If you protect an actual IP address, enter '0.0.0.0'. When you protect the actual IP, the IP address which is allocated for the network interface you will select later will be protected.</p> <p>A client connects to the IP address via the network interface which is configured later. If you use a symbolic name, it must exist in the local /etc/hosts file or be accessible via a Domain Name Service (DNS). Alias names and domain names are acceptable as long as they meet the criteria listed above. No defaults are provided for this information field.</p> <p><b>Note:</b> If you choose to use a symbolic name, be advised that when you extend this resource, the actual IP address will appear in one of the dialog boxes as the IP resource designation.</p>

	<p><b>Note:</b> When the actual IP address is used, you must allocate it to the physical network interface. Only IP v4 is supported for the actual IP resource.</p>
<b>Netmask</b>	<p>Select or enter the network mask, <b>Netmask</b>, which your IP resource will use on the target server. Any standard netmask for the class of the specific IP resource address is valid.</p> <p>If you specify the actual IP address protection ('0.0.0.0') in the IP resource field, this Netmask field will not appear.</p> <p><b>Note:</b> The netmask you choose, combined with the IP address, determines the subnet that will be used by the IP resource and should be consistent with the network configuration.</p>
<b>Network Interface</b>	<p>Select or enter the <b>Network Interface</b> where your IP resource will be placed under LifeKeeper protection. This is the physical Ethernet card that the IP address is interfacing with. Valid choices will depend on the existing network configuration and values chosen for the IP resource address and netmask. The default value is the interface within the set of valid choices which most closely matches the address and netmask values you have selected.</p>
<b>IP Resource Tag</b>	<p>Select or enter a unique IP Resource Tag name for the IP resource instance you are creating. This field is populated automatically with a default tag name as described below.</p> <p>Virtual IP address : ip-&lt;resource&gt; (&lt;resource&gt; is the resource name or IP address.)                  Actual IP address : realip-&lt;network interface&gt; (This is the value which is selected in the Network Interface field.)</p> <p>This tag can be changed if necessary.</p>

4. Click **Create**. The **Create Resource Wizard** will then create your IP resource.
5. At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your IP resource hierarchy. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created. Click **Next**.
6. Another information box will appear explaining that you have successfully created an IP resource hierarchy, and you must Extend that hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

When you click **Continue**, LifeKeeper will launch the **Pre-Extend configuration task**. Refer to the [Extending Your Hierarchy](#) topic for details on how to extend your resource hierarchy to another server.

If you click **Cancel** now, another dialog box will appear alerting you that you'll need to come back and extend your IP resource hierarchy to another server at some other time to put it under LifeKeeper protection.

## 6.6.3.8. Deleting an IP Resource Hierarchy

---

To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Delete Resource Hierarchy**.
2. Select the name of the Target Server where you will be deleting your IP resource hierarchy and then click **Next**. (This dialog will not appear if you selected the Delete Resource task by right-clicking on a resource instance in either pane.)
3. Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete, and highlight it and then click **Next**. (This dialog will not appear if you selected the Delete Resource task by right-clicking on a resource instance in the left or right pane.)
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete** to proceed with resource deletion.
5. Another information box appears confirming that the IP resource was deleted successfully.
6. Click **Done** to exit out of the Delete Resource Hierarchy menu selection.

## 6.6.3.9. Extending Your IP Hierarchy

After you have created a hierarchy, you will want to extend that hierarchy to another server in the cluster. There are three possible scenarios to extend your resource instance from the template server to a target server. The first scenario is when you “Continue” from creating the resource into extending that resource to another server. The second scenario is when you enter the Extend Resource Hierarchy task from the edit menu as shown below. The third scenario is when you right click on an unextended hierarchy in either the left or right hand pane. Each scenario takes you through the same dialog boxes (with a few exceptions, which are clearly detailed below).

1. If you are entering the **Extend wizard** from the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Extend Resource Hierarchy**. This will launch the **Extend Resource Hierarchy wizard**. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information. **Note:** The first two fields appear only if you initiated the **Extend** from the **Edit** menu. It should be noted that if you click **Cancel** at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extension process to that particular server. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Field	Tips
<b>Template Server</b>	Enter the server where your IP resource is currently in service.
<b>Tag to Extend</b>	Select the IP resource you wish to extend. This is the name of the IP instance you wish to extend from the template server to the target server. The wizard will list in the drop down list box all the resources that you have created on the template server that you selected in the previous dialog box.
<b>Target Server</b>	Select the <b>Target Server</b> where you are extending your IP resource hierarchy. The drop down box provides the names of the servers in your cluster that are not already in the selected hierarchy.
<b>Switchback Type</b>	Select the <b>Switchback Type</b> . This dictates how the IP instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either <i>intelligent</i> or <i>automatic</i> . Intelligent switchback requires administrative intervention to switchback the instance to the primary/original server. Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths. The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.
<b>Template Priority</b>	Select or enter a <b>Template Priority</b> . This is the priority for the IP hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will

	reject any priority for this hierarchy that is already in use by another system. The default value is recommended. <b>Note:</b> This selection will appear only for the initial extend of the hierarchy.
<b>Target Priority</b>	Select or enter the <b>Target Priority</b> . This is the priority for the new extended IP hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.

- An information box will appear explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this IP resource have been met. If there were some requirements that had not been met, LifeKeeper would not allow you to select the Next button, and the Back button would be enabled. If you click Back, you can make changes to your resource extension according to any error messages that may appear in the information box. If you click Cancel now, you will need to come back and extend your IP resource hierarchy to another server at some other time to put it under LifeKeeper protection. When you click Next, LifeKeeper will launch you into the Extend Resource Hierarchy configuration task.
- The Extend Resource Hierarchy configuration task will prompt you to enter the following information.

Field	Tips
<b>IP Resource</b>	This is the same <b>IP Resource</b> or address used in the Create Resource Wizard. If the actual IP address is to be protected, enter '0.0.0.0'.
<b>Netmask</b>	This is the same <b>Netmask</b> that was selected when the IP resource was created for the template server and will now be used by the IP resource for the target server. If you specify the actual IP address protection ('0.0.0.0') in the IP resource field, this Netmask field will not appear.
<b>Network Interface</b>	Select or enter the <b>Network Interface</b> . This is the name of the network interface (i.e. Ethernet card) the IP resource will use on the target server.
<b>IP Resource Tag</b>	Select or enter the <b>IP Resource Tag</b> . This is the resource tag name to be used by the IP resource being extended to the target server.

- An information box will appear verifying that the extension is being performed.

Click **Next Server** if you want to extend the same IP resource instance to another server in your cluster. This will repeat the Extend Resource Hierarchy operation.

If you click **Finish**, LifeKeeper will verify that the extension of the IP resource was completed

successfully.

6. Click **Done** to exit from the **Extend Resources Hierarchy** menu selection.

**Note:** Be sure to test the functionality of the new instance on all the servers.

## 6.6.3.10. Unextending Your IP Hierarchy

---

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the dropdown menu, select **Unextend Resource Hierarchy**.
2. Select the Target Server where you want to unextend the IP resource. It cannot be the server where the IP address is currently in service.

**Note:** If you selected the Unextend task by right-clicking from the right pane on an individual resource instance, this dialog box will not appear.

Click **Next** to proceed to the next dialog box.

3. Select the **IP Hierarchy to Unextend**.

**Note:** If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

Click **Next** to proceed to the next dialog box.

4. An information box appears confirming the target server and the IP resource hierarchy you have chosen to unextend.

Click **Unextend**.

5. Another information box appears confirming that the IP resource was unextended successfully.
6. Click **Done** to exit out of the **Unextend Resource Hierarchy menu**.

## 6.6.3.11. Testing Your IP Resource Hierarchy

You can test your IP resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit**, then **Resource**, then finally **In Service** from the dropdown menu. For example, an in-service request executed on a backup server causes the application hierarchy to be placed in service on the backup server and taken out of service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the **Out of Service** request, the application is taken out of service without bringing it in service on the other server.

In a manual switchover, the IP Recovery Kit removes the address from service on the active server before adding it to the backup server.

After switchover, the IP resource has a different hardware (MAC) address because it is associated with a different LAN interface. Before user systems can reconnect, the user systems' TCP/IP software must determine this new address mapping. The IP Recovery Kit automatically informs all connected servers that they must update their ARP (Address Resolution Protocol) tables to reflect the new mapping.

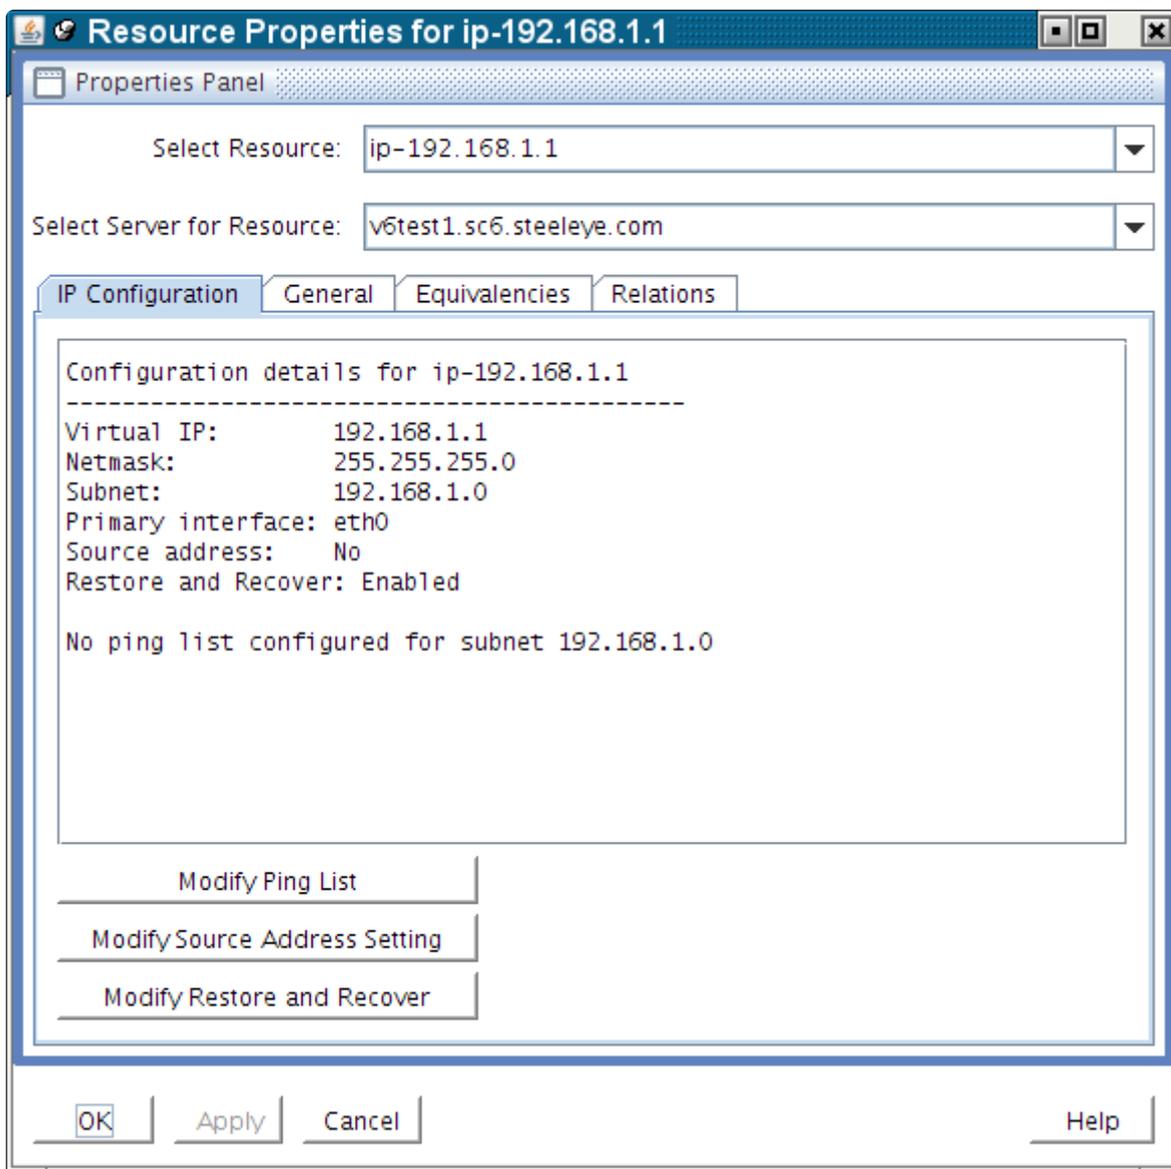
User systems running full TCP/IP implementations are updated immediately. User systems with less sophisticated implementations may have delayed update or may require routers as addressing intermediaries.

## 6.6.3.12. Viewing and Editing IP Configuration Properties

The **IP Configuration Properties** page allows you to view the configuration details for a specific IP resource, as well as to modify a number of selected configuration items.

To access the **IP Configuration Properties** page, from the LifeKeeper GUI menu select **Edit**, then **Resource**. From the dropdown menu, select **Properties**. Then select the resource for which you want to view properties from the **Resource list** and the server for which you want to view that resource from the **Server list**. You can also access the properties page using the context-sensitive menu that appears when you right-click on a specific IP resource instance.

Below is an example of the properties page that will appear for an IP resource.

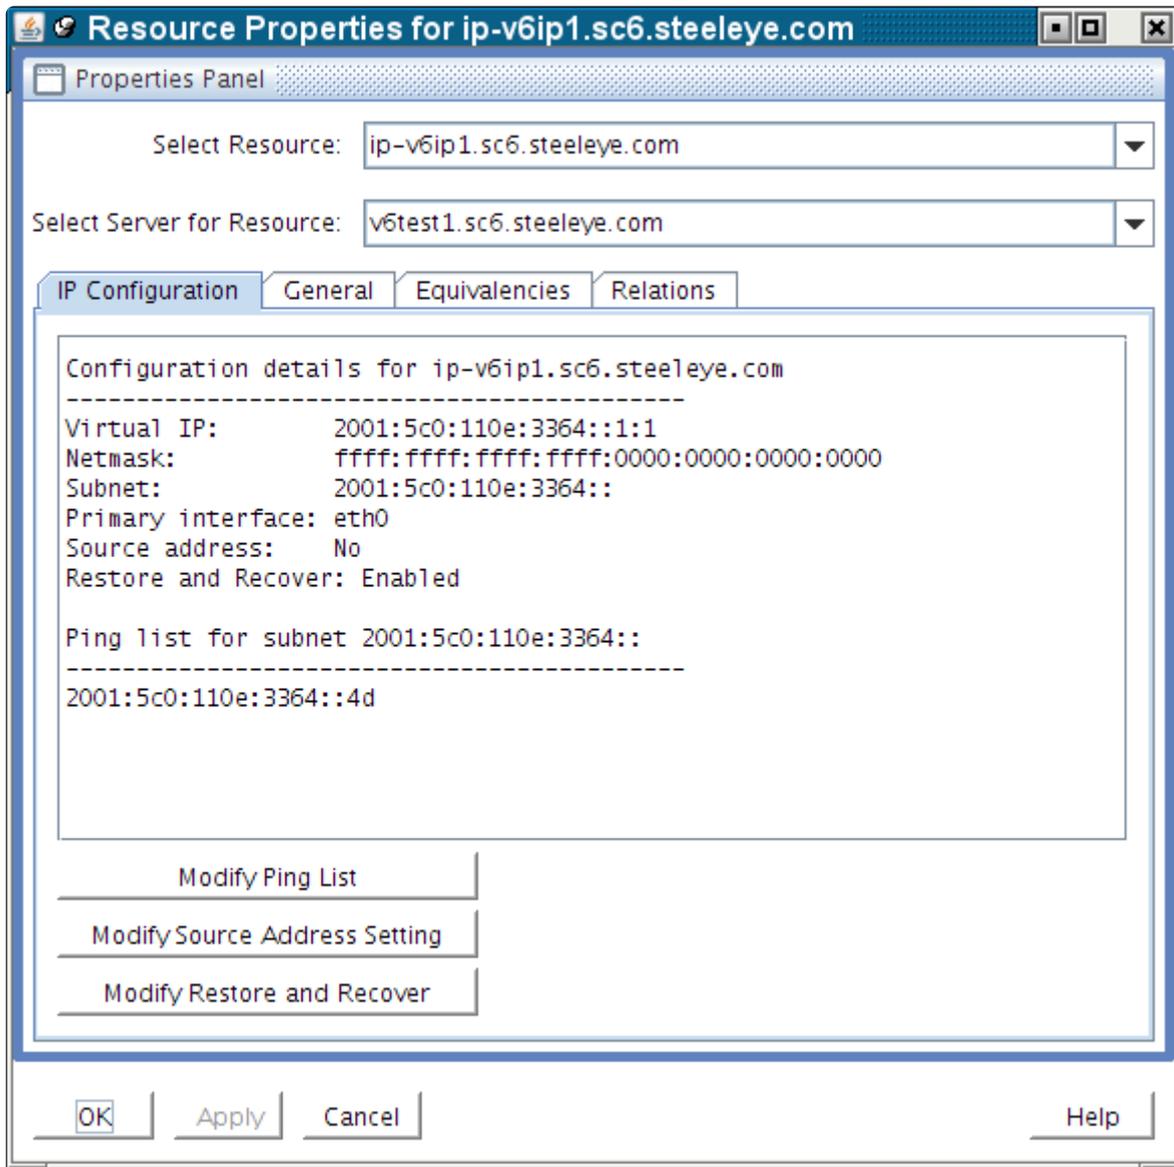


The resulting properties page contains four tabs. The first of those tabs, labeled **IP Configuration**, contains configuration information that is specific to IP resources. The remaining three tabs are available for all LifeKeeper resource types.

The IP Configuration tab displays the following information elements about the selected IP resource.

Virtual IP	The virtual IP address associated with this IP resource.
Netmask	The netmask for the virtual IP address. This value determines how much of the address makes up the subnet portion.
Subnet	The logical subnet address for the virtual IP address, including the number of bits included in the subnet portion of the address.
Primary interface	The network interface on which the virtual IP address should be configured when it is active.
Source address setting	Specifies whether the virtual IP address should be configured as the source address for outbound IP traffic onto its associated subnet.
Ping List	The optional list of IP addresses to be pinged during IP health checks for this IP resource (and others on the same subnet), as an alternative to the normal broadcast ping mechanism.

In the example above, there is no Ping List configured for this IP resource. When a Ping List is configured, the resulting properties page looks like the following example.

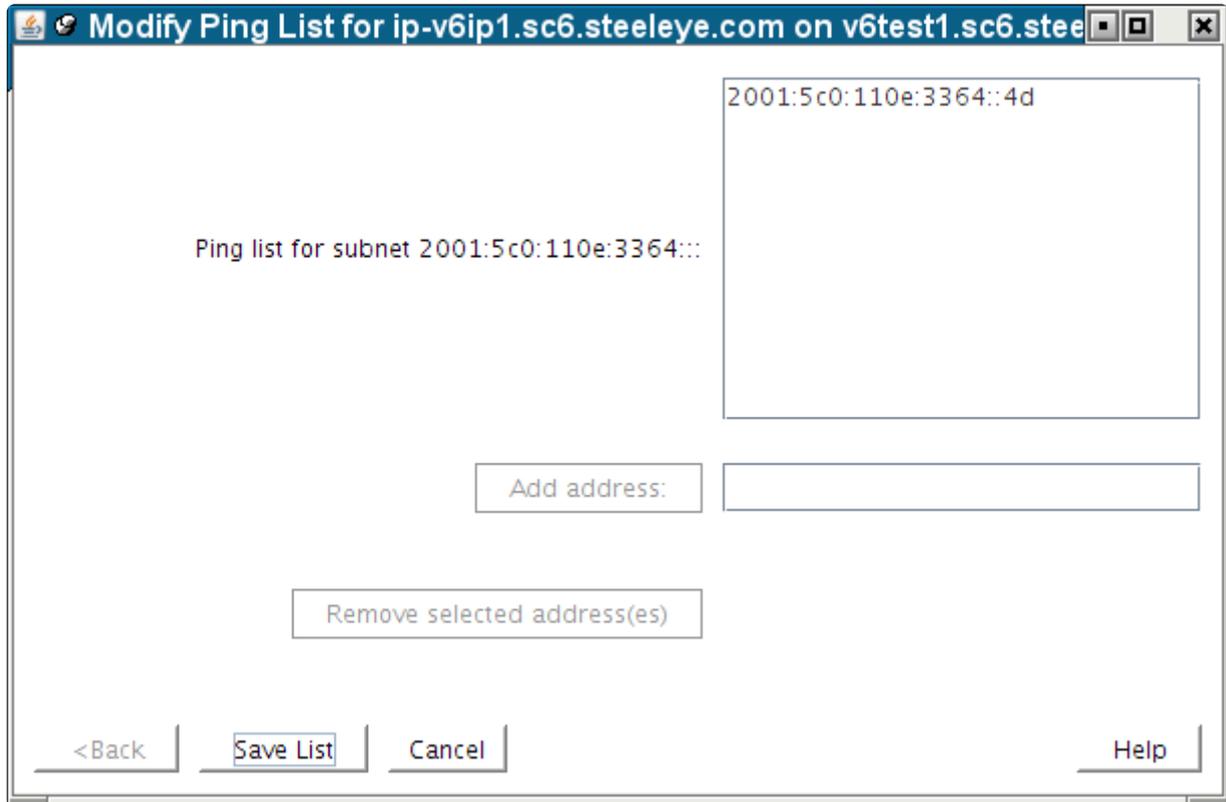


The **Modify Ping List** and **Modify Source Address Setting** buttons can be used to perform modifications to those configuration items, as described in the sections below.

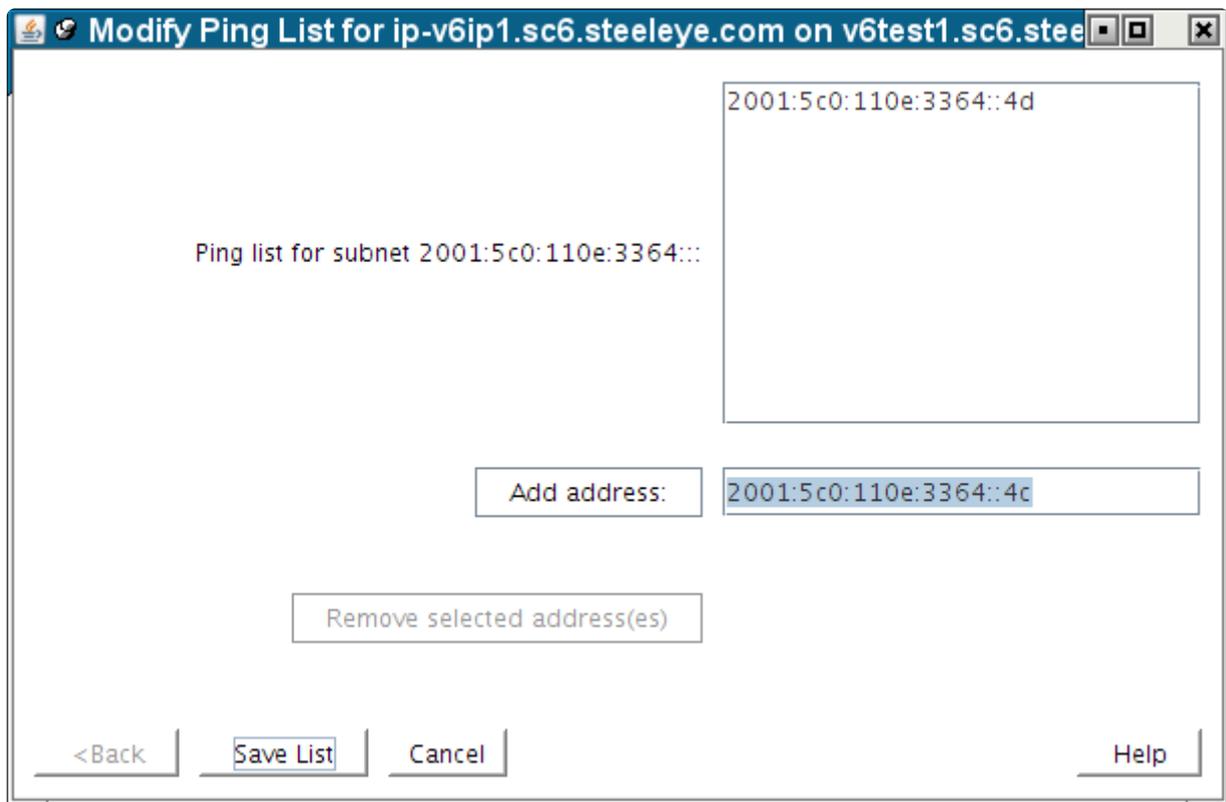
## Modifying the Ping List

For a description of the use and function of the Ping List for an IP resource, see the topic [IP Resource Monitoring](#).

To create a Ping List for an IP resource, or to modify an existing list, click the **Modify Ping List** button on the **IP Configuration properties page**. This brings up a dialog window similar to the following example.

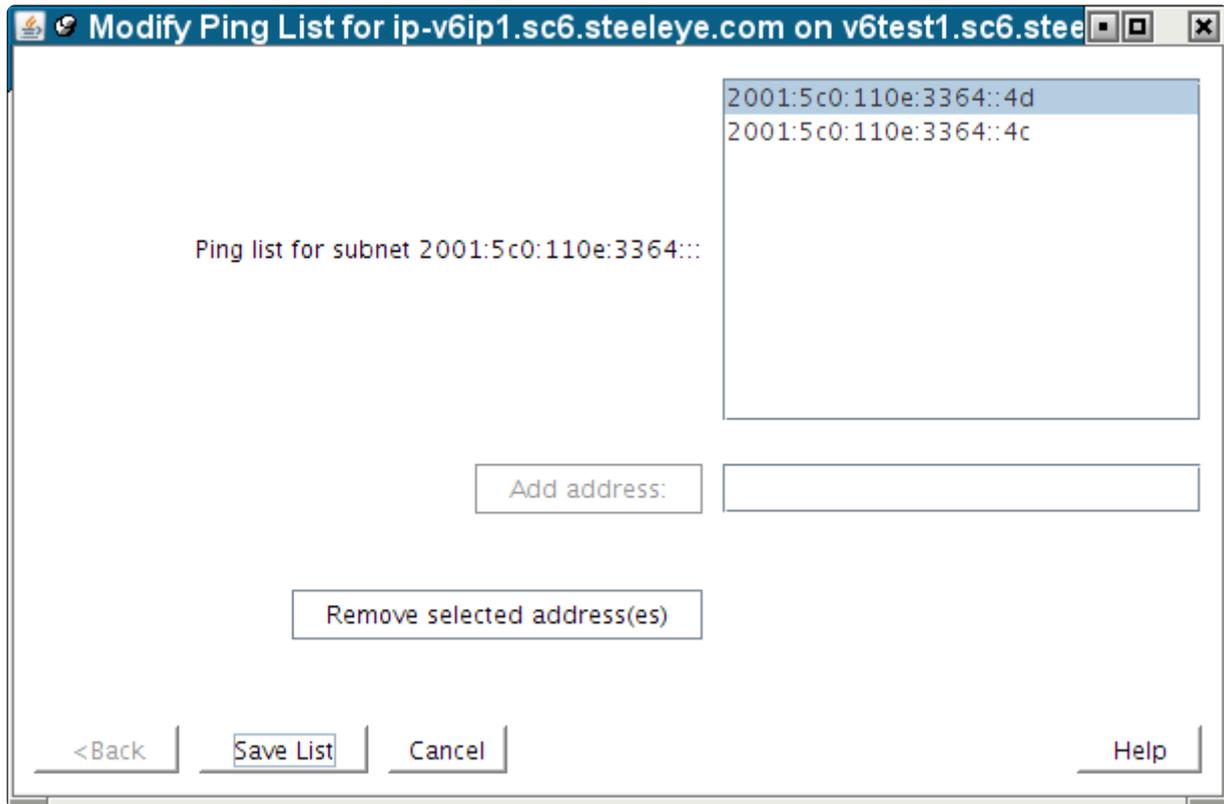


To add an address to the Ping List, type the address in the field next to the **Add address:** button, and push the button, as shown in the following two images. Note that the **Add address:** button is grayed out until you begin typing an address in the field.

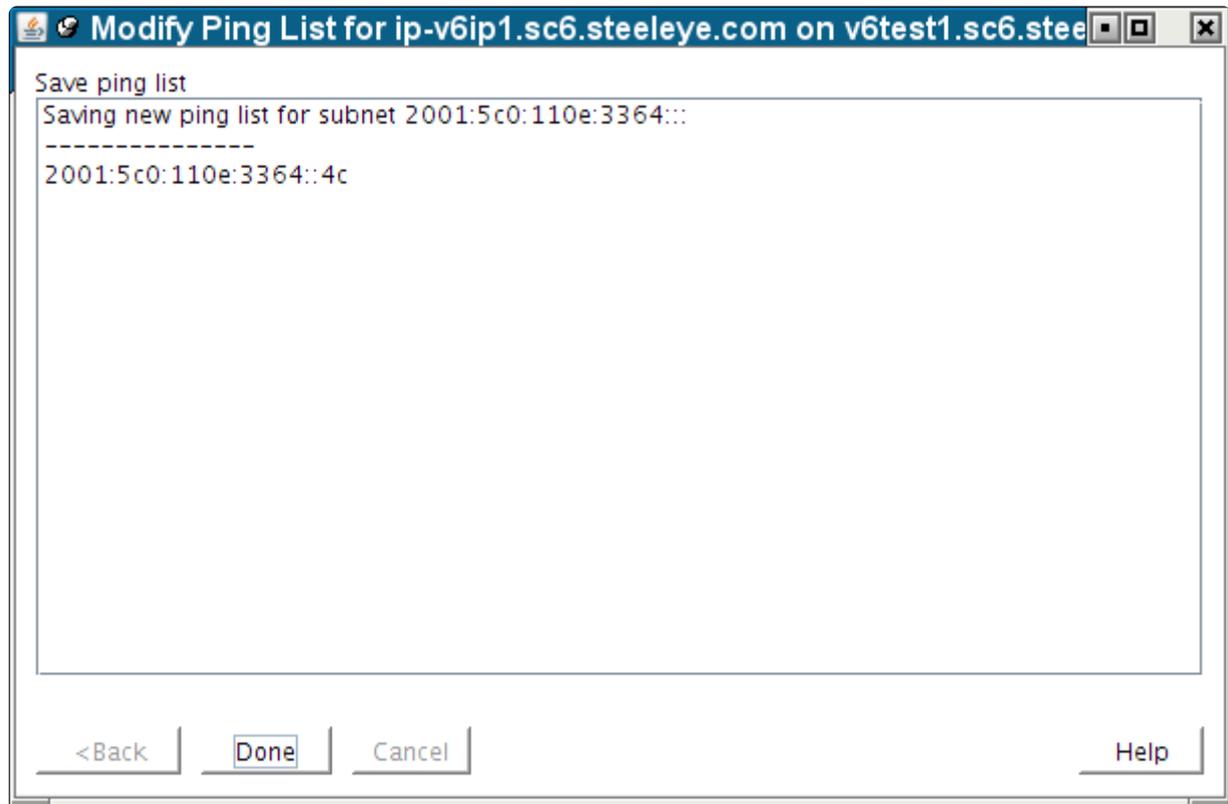


To remove one or more addresses from the Ping List, click to select the addresses to be removed and click the **Remove selected address(es)** button. The **Remove selected address(es)** button is also

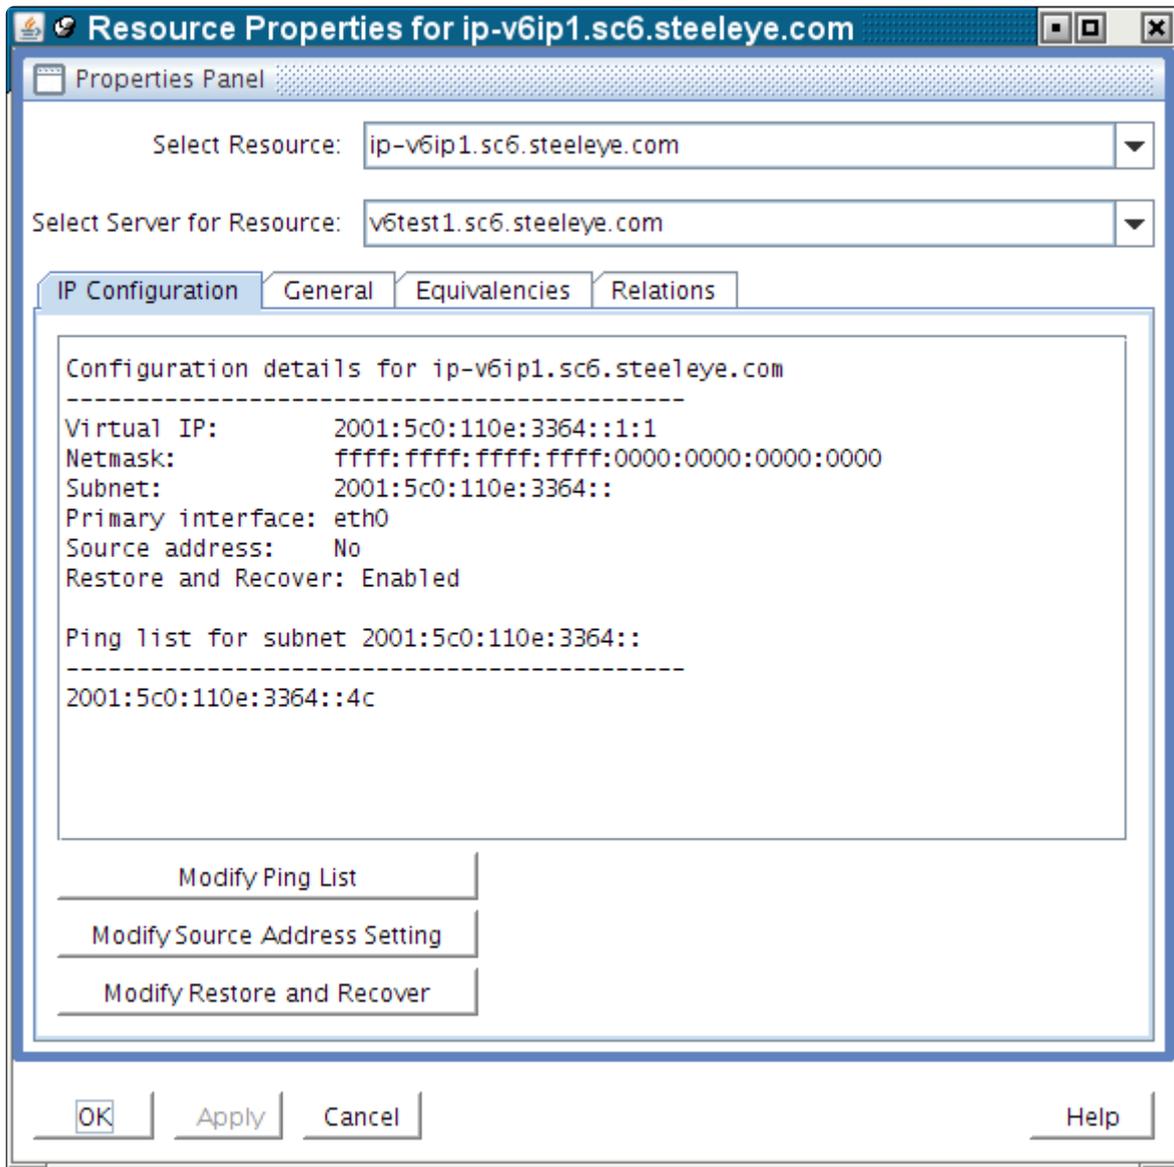
grayed out until at least one address in the list has been selected.



To save the modified list, click **Save List**. This produces the following confirmation window.



Click **Done** to close the window, bringing you back to the **IP Configuration properties page**, where you can see the modified Ping List.



## Important Notes About Using a Ping List

A *Ping List* for an IP resource is unique to the LifeKeeper system on which it is configured. If the IP resource is extended to another LifeKeeper system after the Ping List has been created, the Ping List will be copied to the other system as a part of the extension. However, if the IP resource has already been extended, the Ping List must be configured individually for each system on which the IP resource is defined. Ping List modifications can be made to an IP resource regardless of its state, so there is no need to perform switchovers of the IP resource in order to modify the Ping List on each system.

If there are multiple IP resources defined on the same logical subnet, all of those IP resources share a common Ping List. This is reflected in the IP Configuration properties page and the dialogs associated with modifying the Ping List, where the list is identified as being for the subnet associated with the IP resource.

Once a Ping List has been defined for an IP resource, all health checks for that resource will use the Ping List mechanism rather than the default broadcast ping mechanism. To revert back to the broadcast ping mechanism, you must delete the Ping List by removing all of the address entries in the list.

LifeKeeper performs no validation of the IP addresses entered into a Ping List, other than ensuring the validity of the formatting of the addresses. It is important that you ensure that the addresses you are entering actually exist on your network, can be pinged from the LifeKeeper systems, and are expected to be active at all times. You should not choose addresses that exist on the LifeKeeper systems themselves, because local pings to such addresses may be successful regardless of the actual status of the network interface on which the monitored IP resource is defined.

As mentioned above, the definition of a Ping List for an IP resource on a given system causes LifeKeeper to automatically use the Ping List mechanism rather than a broadcast ping for that resource and all other IP resources on the same subnet. It is not necessary to disable the broadcast ping mechanism using the NOBCASTPING setting described in the [Adjusting IP Recovery Kit Tunable Values](#) topic. However, if you have a configuration in which there are no systems available on the network to respond to a broadcast ping, you may have to use the NOBCASTPING=1 setting initially in order get the IP resource created, before you can then define a Ping List using the procedure described above. Once the Ping List has been created, you can revert back to the default NOBCASTPING=0 setting.

The contents of Ping List remains even after IP resource is deleted. Please note that the old Ping List setting will remain when IP resource with the old subnet address is created after deleting IP resource.

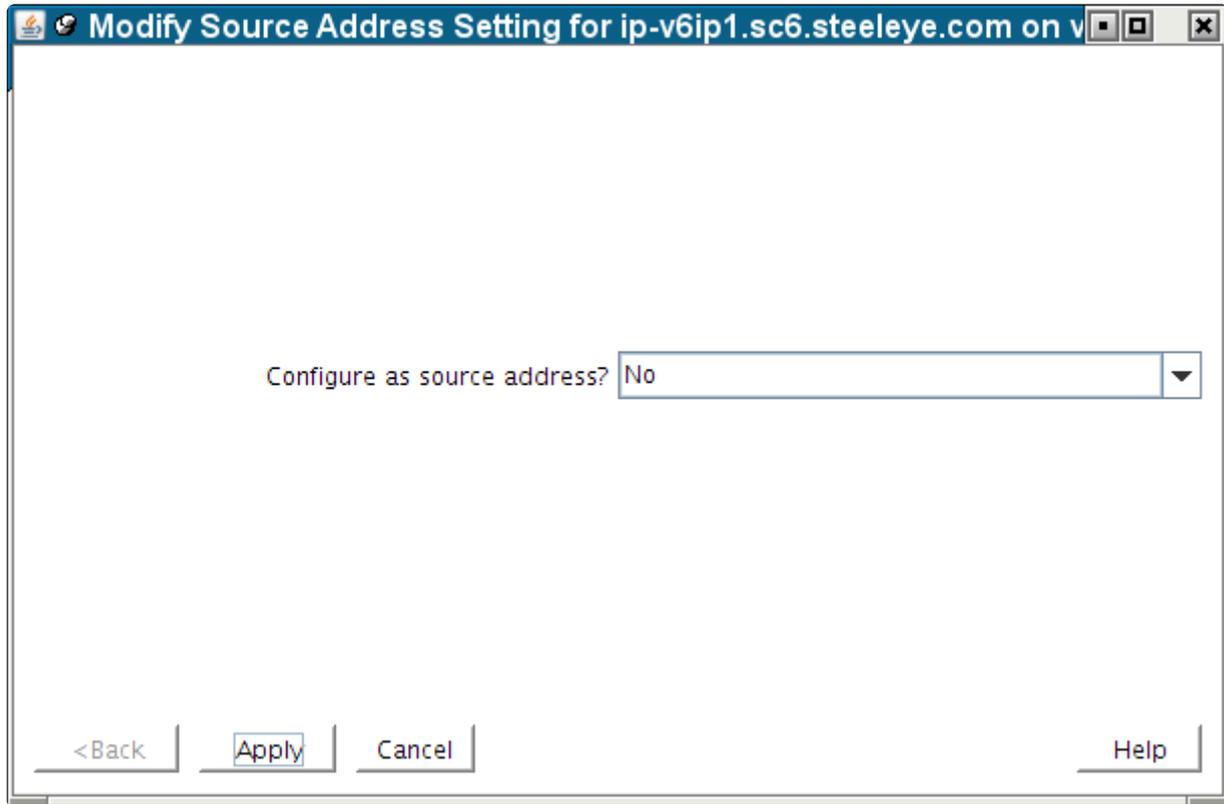
## Modifying the Source Address Setting

The Source Address Setting for an IP resource determines whether the virtual IP address should be used as the source address for outgoing traffic onto the subnet associated with that IP resource, when the IP resource is in-service. This value defaults to No, which means that if the virtual IP address is on the same subnet as the primary IP address for the network interface, outgoing traffic onto the subnet will normally appear to be coming from that primary IP address. This is usually appropriate for most configurations, because the virtual IP address is generally used as an incoming connection point for clients, meaning that all connections in which the virtual IP address is used are initiated as incoming traffic.

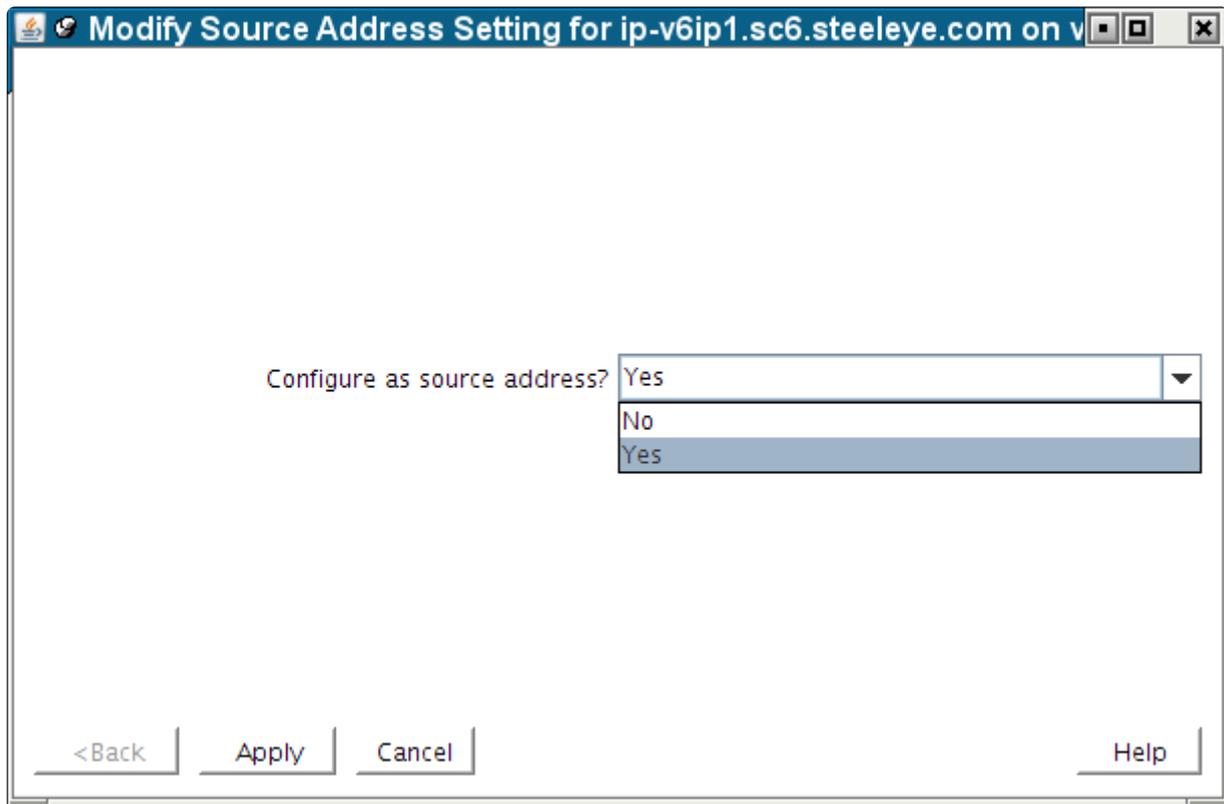
However, there may be situations or configurations in which it is important for connections initiated from the LifeKeeper system to appear to be coming from the virtual IP address. By changing the Source Address Setting for the IP resource to Yes, when the IP resource is brought in-service, the TCP/IP routes on the system are modified such that this will be the case.

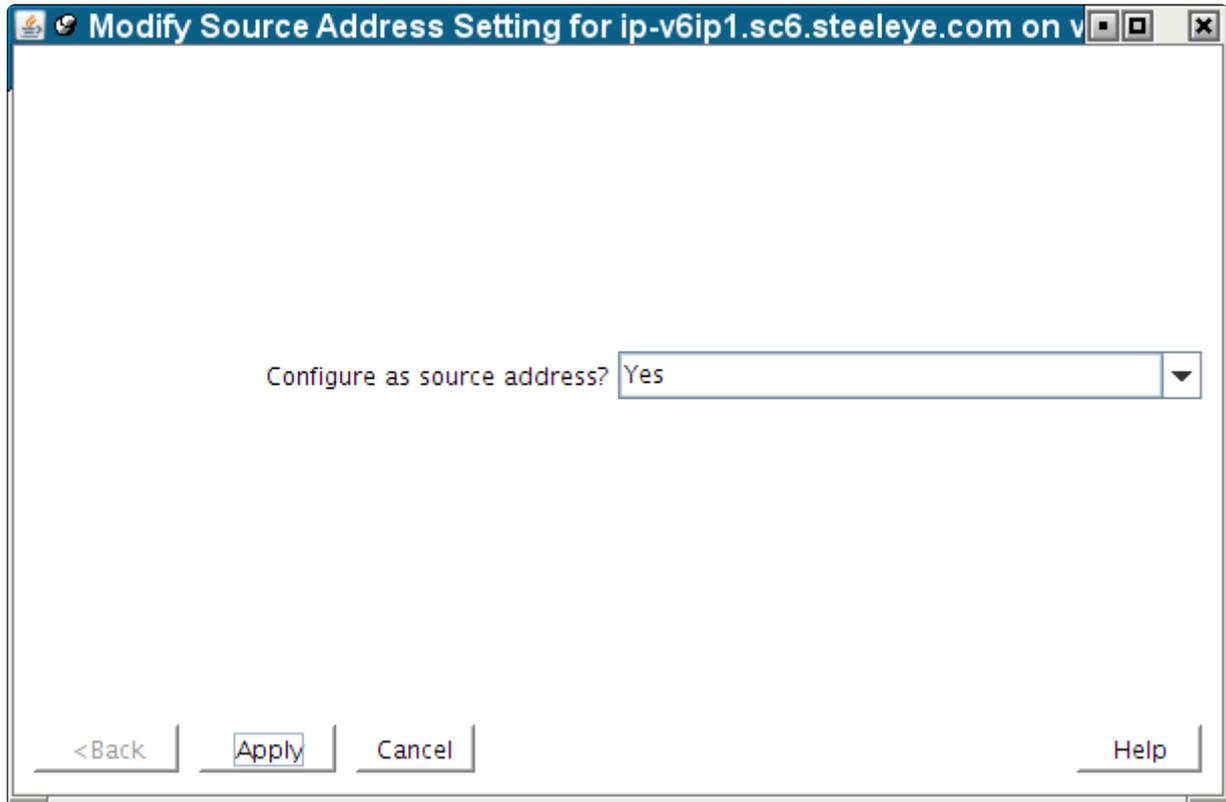
Note that if the virtual IP address is on its own distinct logical subnet from the permanent IP addresses on the system, all outgoing traffic onto that subnet will always come from the virtual IP address without any modifications to the Source Address Setting. Additionally, for EC2 route table configurations with a virtual IP outside the CIDR/subnet, the source address for traffic will be the virtual ip regardless of the source address setting since that's the only ip configured in the subnet.

To modify the Source Address Setting for an IP resource, click the Modify Source Address Setting button on the IP Configuration properties page. This brings up a dialog window similar to the following example.

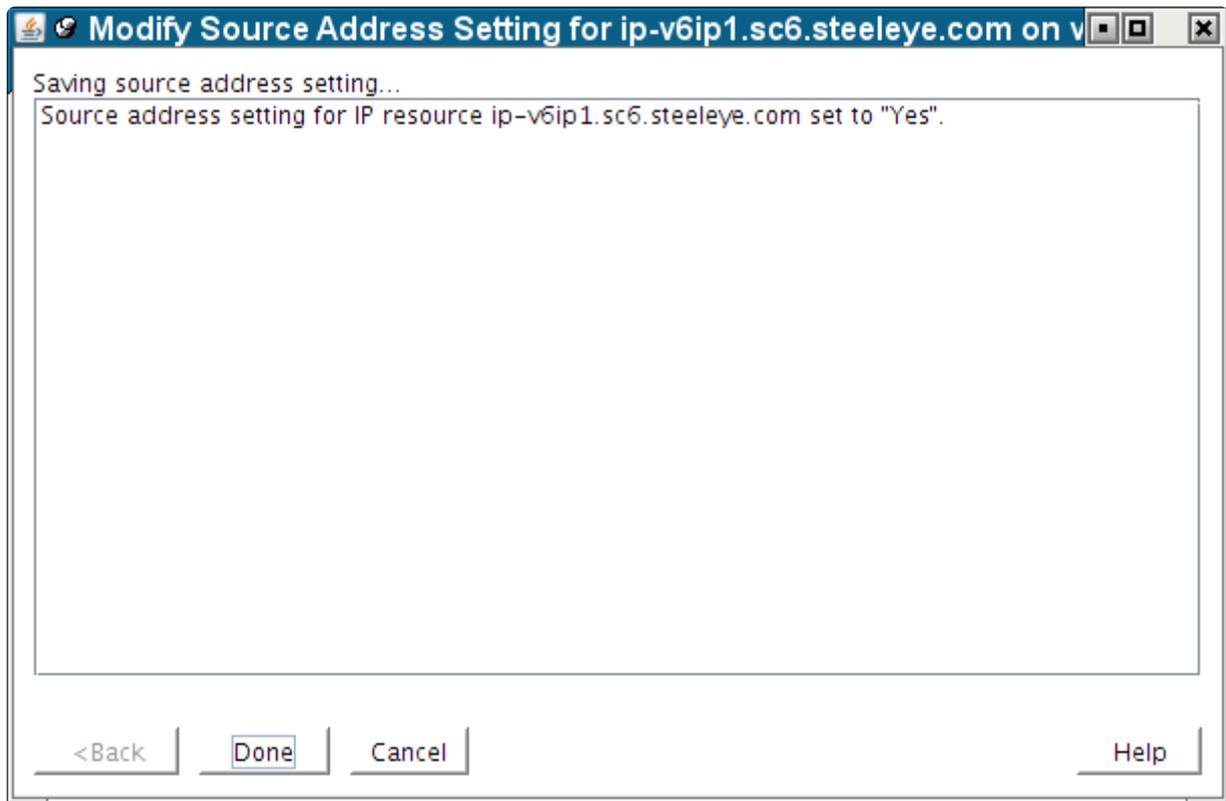


To change the setting, use the dropdown list to select the new value, either **Yes** or **No**.

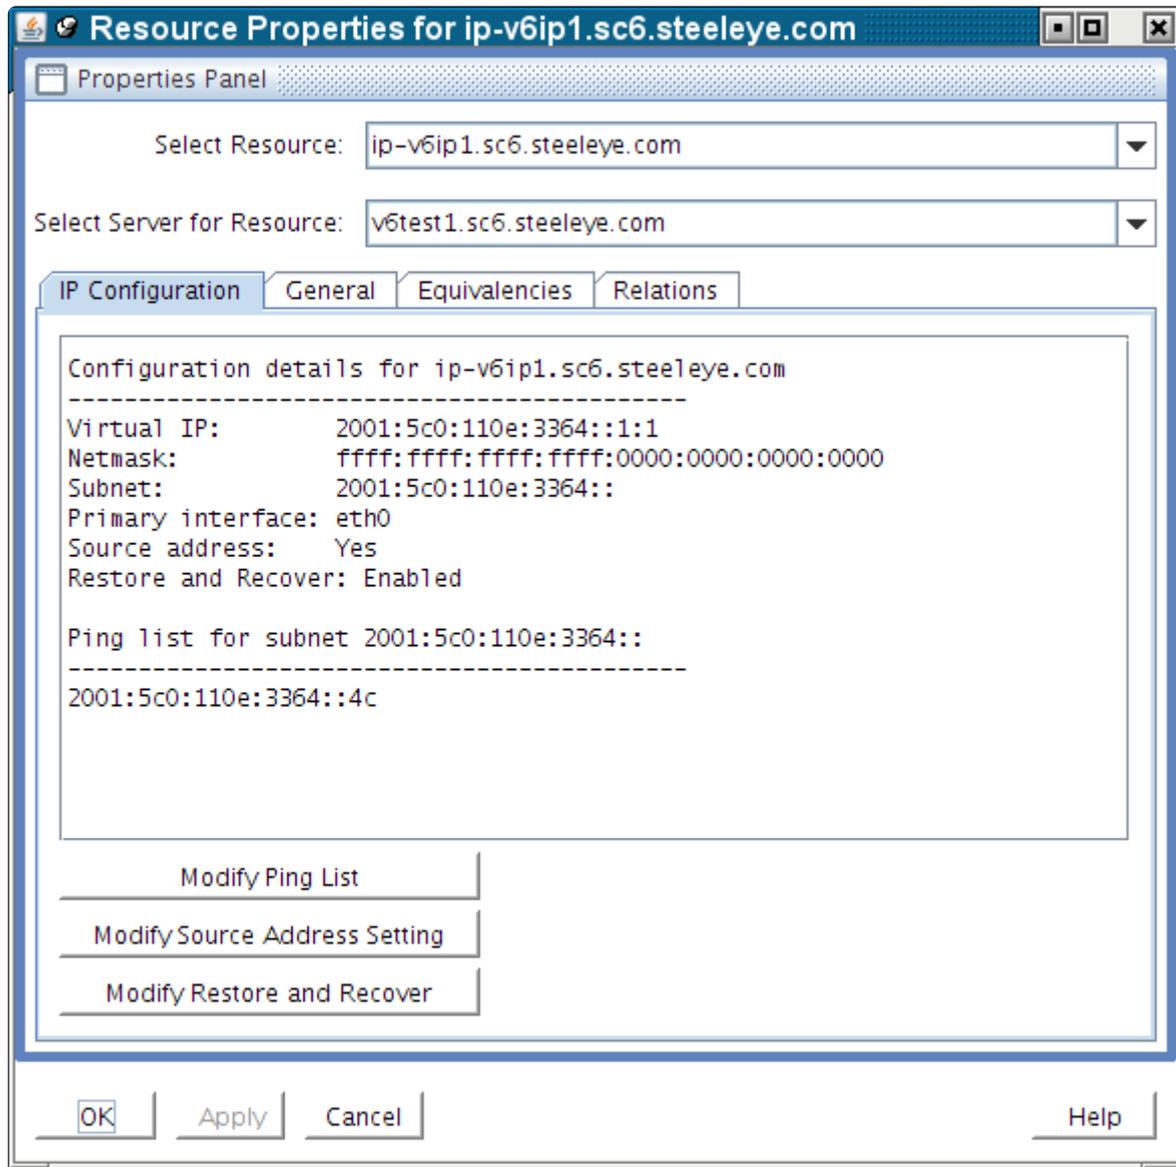




Click **Apply** to save the new setting. This produces the following confirmation window.



Clicking **Done** will close the window and take you back to the **IP Configuration properties page**, where you can see the modified setting.



## Important Notes About the Source Address Setting

The **Source Address Setting** for an IP resource is unique to the LifeKeeper system on which it is configured. If the IP resource is extended to another LifeKeeper system after the **Source Address Setting** has been modified, the setting will be copied to the other system as a part of the extension. However, if the IP resource has already been extended, the **Source Address Setting** modification must be made individually for each system on which the IP resource is defined.

It only makes sense for at most one IP resource on a given subnet to have its Source Address Setting set to Yes, because only a single IP address can actually be the source address for outgoing traffic onto the subnet. If there are multiple IP resources on the same subnet with a setting of Yes, the most recent resource to be brought in-service will override any others and become the source address for outgoing traffic onto the subnet.

The **Source Address Setting** only affects the local TCP/IP configuration when the IP resource is brought into service. So if the resource is already active when the setting is changed, the resource must be taken out-of-service and then back in-service before the change is reflected in the TCP/IP configuration.

The **Source Address Setting** only affects IPv4 addresses. This setting has no effect on an IPv6 address.

## Modifying Restore and Recover

This feature allows a user to choose to **Enable** or **Disable** the default restore and recovery behavior for an existing IP address resource. If configured with the **Enable** option, the IP address will be brought in-service as normal and the regular monitoring and recovery process will occur. The **Enable** option is the current default behavior for an IP address restore.

If the **Restore and Recover** option is set to **Disable**, the resource will come in-service, but the IP address will not be brought active on the network or network adapter. This setting allows hierarchies in a WAN environment that depend on an IP to be brought in-service on the Disaster Recovery (DR) system where it may be difficult to configure the IP on the DR system due to the WAN configuration.

This setting can be selected after the resource is created and extended.

**Important consideration for Active IP addresses (ISP):** Setting the action to **Disable** on an ISP and active IP address does not take the active IP out of service.

## 6.6.3.13. Adjusting IP Recovery Kit Tunable Values

---

For details about the tunable values that are available for modifying the behavior of the IP Recovery Kit, please click [here](#). These values are tuned by editing the `/etc/default/LifeKeeper` configuration file. Because none of the components of the IP Recovery Kit are memory resident, changes to these particular values become effective immediately after they are changed in `/etc/default/LifeKeeper`, without requiring a LifeKeeper restart.

## 6.7. MySQL Recovery Kit Administration Guide

---

The LifeKeeper for Linux MySQL Recovery Kit provides an easy way to add LifeKeeper fault-resilient protection for MySQL resources and databases. This enables a failure on the primary database server to be recovered on a designated backup server without significant lost time or human intervention.

### LifeKeeper Documentation

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#) (also available from the **Help** menu within the LifeKeeper GUI)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is available on the [SIOS Technical Documentation](#) website.

## 6.7.1. MySQL Recovery Kit Hardware and Software Requirements

---

Before you can install and set up the recovery software, your server must meet certain hardware and software requirements. You should refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper MySQL Recovery Kit.

Be sure that your configuration meets the following requirements:

- **Servers.** The Recovery Kit requires two or more LifeKeeper supported computers configured in accordance with the requirements described in [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#).
- **LifeKeeper software.** You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.
- **LifeKeeper IP Recovery Kit.** This kit is required if remote clients will be accessing the MySQL database. You must have the same version of this Recovery Kit on each server.
- **IP network interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

 **Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation and so forth.

- **TCP/IP software.** Each server also requires the TCP/IP software.
- **MySQL software.** Each server must have the MySQL software installed and configured prior to configuring LifeKeeper and the LifeKeeper MySQL Recovery Kit. The same version should be installed on each server. Consult the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information.

## 6.7.2. MySQL Recovery Kit Configuration

---

This section contains definitions and examples of typical LifeKeeper MySQL configurations and information you should consider before you start to configure MySQL.

Please refer to the [Resource Hierarchies](#) section of the [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring your LifeKeeper Core resource hierarchies.

---

[Configuration Considerations for MySQL](#)

[Client Configuration Considerations](#)

[Configuration Requirements](#)

[Configuration Examples](#)

[Active – Standby Configuration](#)

[Active – Active Configuration](#)

[Multiple Database Server Environment](#)

[Using mysqld Groups with LifeKeeper](#)

[Using Network Attached Storage](#)

[Considerations on MySQL use in Systemd environments](#)

## 6.7.2.1. Configuration Considerations for MySQL

---

Below are some specific considerations you need to think about concerning your LifeKeeper MySQL environment.

To operate MySQL database services on the primary and backup servers, file systems and disk partitions must be accessible from each server. Before you can begin configuring the MySQL Recovery Kit, be sure you have completed the following preliminary steps and have tested/run the databases on each server. In the instructions below, the user “mysql” refers to the operating system user that will start the MySQL server.

1. Install the MySQL server and client components on all servers. Be sure that all of the servers are running the same version of the MySQL client and server components. The MySQL executables can be located on a local or shared drive.

**Note:** If you use Red Hat Software Collections and need to export the X\_SCLS environment variable in order to run a specific version of MySQL with LifeKeeper, then set the X\_SCLS environment variable via `/etc/default/LifeKeeper` by adding the line `X_SCLS=VERSION` to the file (i.e. `X_SCLS=mysql55`). This is typically only the case if you want to enable MySQL 5.5 which is included in RHEL 5.10 (MySQL 5.0 is enabled as the default).

2. If `mysqld` is running on any of the servers on the socket and/or port where you wish to run the LifeKeeper protected MySQL database server, stop each MySQL server using the `mysqladmin` command.
3. Move the contents of the MySQL data directory to a shared location. By default, the MySQL data directory is installed on a local drive. This location depends on the distribution mechanism. The binary RPM installs the data directory at `/var/lib/mysql`. (Be sure that only the contents are moved and the directory remains intact. This allows the MySQL database server to write logs in this directory, if necessary. Make sure that the “mysql” user described in step 4 has permissions to write the logs to this location.)
4. If the installation process did not create the Linux user “mysql”, create this user. For security reasons, the MySQL server should not be run as “root.” (Refer to the [MySQL Administration Guide](#) for a full discussion of the security issues.) Make sure that “mysql” is the only user with read/write permissions in the database directories. The “mysql” user and group should be created on all servers. The user ID and group ID must be the same on all servers.
5. **IMPORTANT:** A server started by an automatic OS startup cannot be under LifeKeeper protection. In addition, the server can not use the same port number or socket as a server under LifeKeeper protection.
6. It is recommended that the socket be written to the data directory on the shared disk. If the socket will be written to a local disk, make sure the path exists on all LifeKeeper servers where your hierarchy will exist. Make sure that the user “mysql” has permissions to write the socket to this

location.

7. Start the MySQL server using the `mysql` daemon startup command appropriate for your configuration. For configurations defining a single instance in the *my.cnf* file, use the command:

```
<start command> -user=mysql -socket=<socket> -port=<port number>
  
-datadir=<path to the data directory> -log &
```

The `<start command>` for `mysql` versions 3.x is `safe_mysqld`, and the command for version 4.x is `mysqld_safe`.

For configurations using the `mysqld` Group feature in the *my.cnf* file, use the command:

```
mysqld_multi start <group number>
```

The `<group number>` represents the numerical instance defined in the *my.cnf* file for the `mysqld` Group. For more information on using `mysqld` groups with LifeKeeper, see: [Using `mysqld` Groups with LifeKeeper](#).

`systemctl` command must be utilized when MySQL (v5.7.6 or later) is set up to use Systemd in the distribution with Systemd. For the details, refer to “[Consideration about the use in Systemd environment](#)”.

8. Create a MySQL database user named “mysql”. Give this user a password and grant the user “shutdown” permissions. This only has to be done on one server. (Refer to the [MySQL Administration Guide](#) for details on creating users and granting permissions).
9. Copy the sample *my.cnf* configuration file to the desired location (*/etc* or */<datadir>*). This file contains options for the database server and for client programs.

The file can be located in either the *MySQL data* directory or the */etc* directory. The */etc/my.cnf* file contains global options. Place the *my.cnf* file in */etc* if only one database will run on the machine at any given time (i.e. an Active/Standby configuration) or if you are using the `mysqld` Group feature (see [Using `mysqld` Groups with LifeKeeper](#)). If the file is located in */etc*, you must copy it to each LifeKeeper backup server. The *my.cnf* file in the data directory should contain server-specific options. For multiple servers and Active/Active configurations, this file must be stored in the data directory for each resource instance unless you are using the `mysqld` Group feature (see [Using `mysqld` Groups with LifeKeeper](#)).

**Note:** The *my.cnf* file should not exist in both the */etc* and */* locations if both copies will contain server specific options. If a *my.cnf* file containing server specific options is located in */etc* along with a protected *my.cnf* file installed in the */* potential conflicts may result. Refer to the MySQL documentation on configuring global settings and server specific options.

Add or edit the following entries:

- a. In the “client” section of the file, specify the user and the password that should be used for

connections.

```
[client]
user =clientuser
password =password
.
.
.
```

b. In the “mysqld” section of the file, specify the socket and port that should be used for connections, as well as the pid-file location for the mysqld process. The user variable should specify the operating system user that will start the mysqld process.

```
[mysqld]
socket =/home1/test/mysql/mysql.sock
port =3307
pid-file =/home1/test/mysql/mysqld.pid
user =osuser
```

 **Note:** Make sure this file is properly protected and owned by the user “mysql.”

**Note:** Once the MySQL hierarchy is created, if you need to change any of the information in the `my.cnf` file, you must stop the mysql server instance by taking the hierarchy out-of-service (i.e. the OSU state) before making changes.

**Note:** The above example `my.cnf` configuration describes a single database instance `mysqld`. See [Using mysqld Groups with LifeKeeper](#) for configuration examples using `mysqld` groups.

**Note:** “include” directive is not supported. All the setups must be described in a single `my.cnf` file.

## 6.7.2.2. Client Configuration Considerations for MySQL

---

Following are some configuration considerations for MySQL database clients:

- If clients will connect from remote hosts, create an IP address under LifeKeeper to be used for client connections.
- Clients must be configured to connect to the database server through a LifeKeeper-protected IP address.
- If the clients will connect through a domain name instead, create an entry in each client's hosts file for the protected IP address, or configure the name in DNS. Test the protected IP address by pinging it from all clients and all LifeKeeper servers in the cluster.
- Although each user can have a *my.cnf* file in the home directory of their machine, LifeKeeper only uses the *my.cnf* file located in the */etc* directory or the data directory. The *my.cnf* file stores the client connection information (i.e. the port, socket identification, user and password).

## 6.7.2.3. MySQL Configuration Requirements

---

Each of the examples involves one or two database instances: databaseA and databaseB. The Database Tag names are arbitrary names that describe these database instances to LifeKeeper. The word on and the system identifier that follows provide clarification but are not required. The default tag name suggested by LifeKeeper is mysql or mysql for configurations using mysqld Groups (see [Using mysqld Groups with LifeKeeper](#)). To understand the configuration examples, keep these configuration requirements in mind:

- **LifeKeeper hierarchy.** When performing LifeKeeper administration, the primary hierarchy refers to the hierarchy being built on the server you are administering. For the configuration diagrams, the information entered in the first administration screen is from the perspective of Server 1. When a second screen is shown, it refers to the hierarchy being built while administering the second server. In the configuration examples, the second server is Server 2.
- **Shared disk locked by one server.** When you use LifeKeeper, one server reserves shared storage resources that are under LifeKeeper protection for use. This is done using SCSI reservations. If the shared device is a disk array, an entire LUN is reserved; if a shared device is a disk, then the entire disk is reserved. This prevents inadvertent corruption of the data by other servers in the cluster. When a server fails, the highest priority backup server breaks the old reservation and establishes its own reservation, locking out all other servers.
- **Data Directory on shared disk.** In order for the LifeKeeper MySQL Recovery Kit to function properly, the data directory (datadir) of the database instance must always be on a shared disk. The data directory must be on a file system. The file system must be mountable from both the primary and backup servers. The data directory (datadir) can also exist on replicated or [network attached storage](#).

## 6.7.2.4. MySQL Configuration Examples

---

The examples in this section show how MySQL database instances can be configured. Each diagram shows the relationship between the type of configuration and the MySQL parameters. Each configuration also adheres to the configuration rules and requirements described in this documentation that ensure compatibility between the MySQL configuration and the LifeKeeper software.

This section describes the configuration requirements and then provides these configuration examples:

- Active/Standby
- Active/Active

The examples in this section are only a sample of the configurations you can establish, but understanding these configurations and adhering to the configuration rules will help you define and set up workable solutions for your computing environment.

### [Configuration Requirements](#)

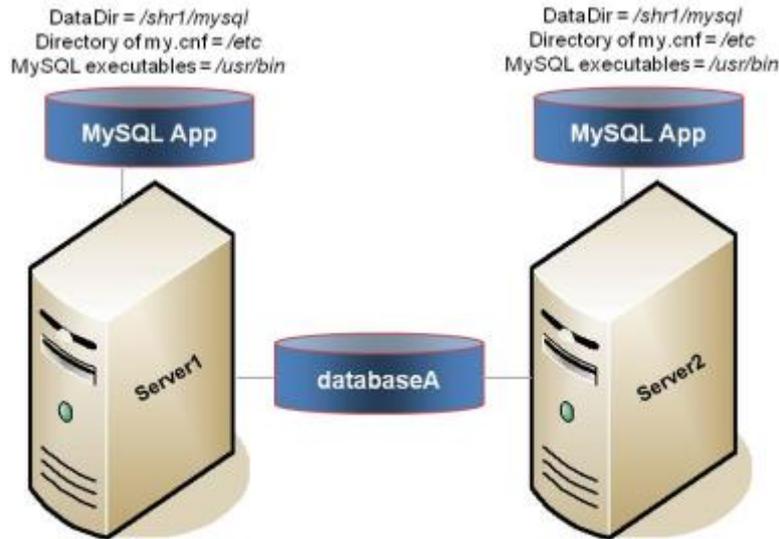
#### **Example 1** – [Active/Standby Configuration](#)

#### **Example 2** – [Active/Active Configuration](#)

## 6.7.2.5. Active/Standby MySQL Configuration

This section provides an example of an active/standby configuration. In this configuration, Server 1 is considered active because it has exclusive access to the database. Server 2 does other processing. If Server 1 fails, Server 2 gains access to the database, and LifeKeeper re-establishes the database operations.

**Figure 1. Active/Standby Configuration, Example 1**



### Configuration Notes:

- Both servers use the MySQL data directory (which includes the database (databaseA)) on a shared disk.
- The path to the MySQL data directory is the same on both servers.
- The *my.cnf* configuration file is located on a local disk in */etc*.
- The MySQL executables are located on a local drive on each server in */usr/bin*.
- Server 2 cannot access files and directories on the shared disk while Server 1 is active.

### Creating a resource hierarchy on Server 1:

Server:	Server 1
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>

Directory of MySQL Executables Location:	<i>/usr/bin</i>
Database Tag	mysql-on-server1

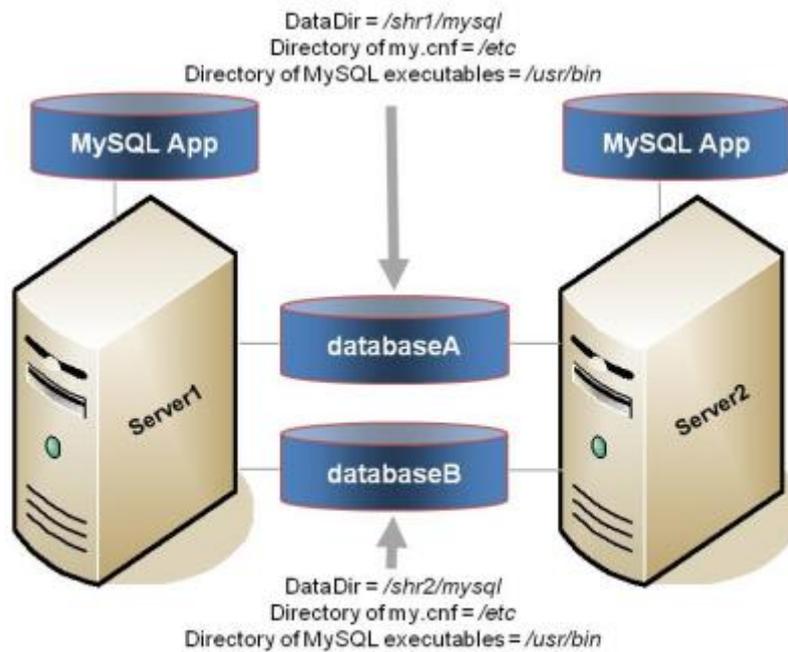
**Extending a resource hierarchy to Server 2:**

Template Server:	Server 1
Tag to Extend	mysql-on-server1
Target Server	Server2
Target Priority:	10
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>
Directory of MySQL Executables Location:	<i>/usr/bin</i>
Database Tag	mysql-on-server2

## 6.7.2.6. Active/Active MySQL Configuration

An active/active configuration consists of two or more servers actively running a different database instance with each serving as a backup for each other. The database instances must be on different shared physical disks. For LifeKeeper configurations supporting multiple MySQL database instances (of the same or different versions), SIOS recommends that the [mysqld Group](#) feature be used for versions of MySQL that support this feature. For these configurations, the *my.cnf* configuration file will reside in */etc*. For MySQL versions that do not support the *mysqld Group* feature, the *my.cnf* configuration file must reside in the MySQL data directory shared file system for each database instance (e.g. in Figure 2 below, */shr1/mysql* and */shr2/mysql*).

Figure 2. Active/Active Configuration, Example 2



### Configuration Notes:

- Each server uses a different MySQL data directory (which includes the database instances (database A and database B) on different shared disks
- The path to the MySQL data directory is different for each instance defined on the server.
- The *my.cnf* configuration file is located in */etc* and contains *mysqld* group sections for each database instance. Each section defines a unique MySQL data directory, port and socket for that database instance. The *my.cnf* configuration file must be kept in sync on all nodes in the cluster. For systems running versions of MySQL that do not support *mysqld Groups*, the *my.cnf* configuration file for each of the database instances is located on the shared drive in the data directory for the database instance. Each configuration file defines a unique MySQL data directory, port and socket definition for that database instance.

- The MySQL executables are located on a local drive on each server in */usr/bin*.
- Initially, Server 1 runs databaseA and Server 2 runs databaseB. In a switchover situation, one server can run both databases.

**Creating the first resource hierarchy on Server 1:**

Server:	Server 1
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>
Directory of my MySQL Executables Location:	<i>/usr/bin</i>
Database Tag:	mysql-shared.example.instance1

**Extending the first resource hierarchy to Server 2:**

Template Server:	Server 1
Tag to Extend:	mysql-shared.example.instance1
Target Server:	Server2
Target Priority:	10
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>
Directory of my MySQL Executables Location:	<i>/usr/bin</i>
Database Tag:	mysql-shared.example.instance1

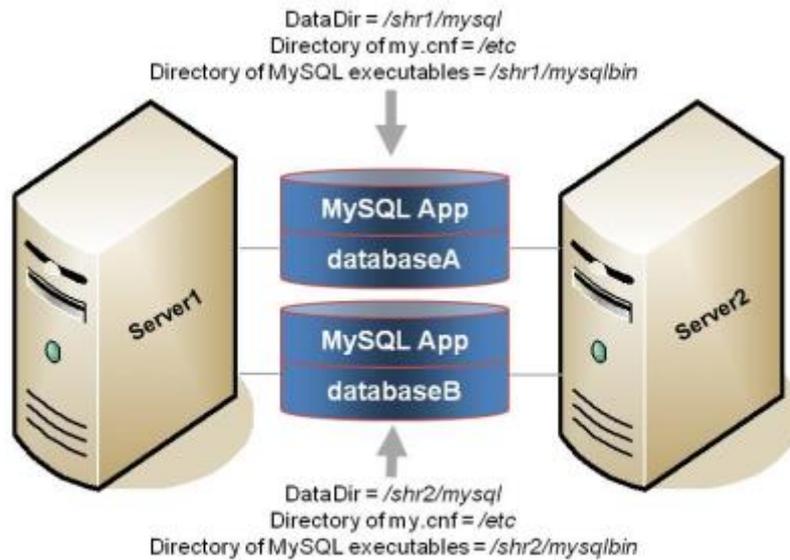
**Creating the second resource hierarchy on Server 2:**

Server:	Server 2
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>
Directory of my MySQL Executables Location:	<i>/usr/bin</i>
Database Tag:	mysql-shared.example.instance2

**Extending the second resource hierarchy to Server 1:**

Template Server:	Server 2
Tag to Extend:	mysql-shared.example.instance2
Target Server:	Server1
Target Priority:	10
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>

Directory of my MySQL Executables Location:	/usr/bin
Database Tag:	mysql-shared.example.instance2



Configuration Notes:

- Each server uses a different MySQL data directory (which includes the database instances (database A and database B) on different shared disks
- The path to the MySQL data directory is different for each database instance defined on the server.
- The *my.cnf* configuration file is located in /etc and contains mysqld group sections for each database instance. Each section defines a unique MySQL data directory, port and socket for that database instance. The *my.cnf* configuration file must be kept in sync on all nodes in the cluster. For systems running versions of MySQL that do not support mysqld Groups, the *my.cnf* configuration file for each of the database instances is located on the shared drive in the data directory for the database. Each configuration file defines a unique MySQL data directory, port and socket definition for that database instance.
- There is a copy of the MySQL executables on each of the shared disks that contains the data directories.
- Initially, Server 1 runs databaseA and Server 2 runs databaseB. In a switchover situation, one server can run both database instances.

Creating the first resource hierarchy on Server 1:

Server:	Server1
Directory of <i>my.cnf</i> File Location:	/etc

Directory of MySQL Executables Location:	<i>/shr1/mysqlbin</i>
Database Tag:	mysql-shared.example.instance1

### Extending the first resource hierarchy to Server 2:

Template Server:	Server1
Tag to Extend:	mysql-shared.example.instance1
Target Server:	Server2
Target Priority:	10
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>
Directory of MySQL Executables Location:	<i>/shr1/mysqlbin</i>
Database Tag:	mysql-shared.example.instance1

### Creating the second resource hierarchy on Server 2:

Server:	Server2
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>
Directory of MySQL Executables Location:	<i>/shr2/mysqlbin</i>
Database Tag:	mysql-shared.example.instance2

### Extending the second resource hierarchy to Server 1:

Template Server:	Server2
Tag to Extend:	mysql-shared.example.instance2
Target Server:	Server1
Target Priority:	10
Directory of <i>my.cnf</i> File Location:	<i>/etc</i>
Directory of MySQL Executables Location:	<i>/shr2/mysqlbin</i>
Database Tag:	mysql-shared.example.instance2

## 6.7.2.7. Multiple Database Server Environment

---

Following are some configuration considerations if you have multiple MySQL database servers and databases:

- If running active/active or multiple MySQL instances (of the same or different versions), please consider using the [mysqld Group](#) feature if possible. SIOS recommends using [mysqld Groups](#) (`mysqld_multi`) for multiple MySQL database server configurations.
- If running active/active or multiple instances of MySQL, do not mount a shared file system as `/var/lib/mysql`. This causes unexpected shutdown of MySQL servers by the `mysql` startup command (`safe_mysqld` or `mysqld_safe`).
- The `my.cnf` file must be stored in the data directory for each of the active/active or multiple servers if not using the `mysqld` group feature. For configurations using `mysqld Groups`, the `my.cnf` file should be stored in `/etc` and not in the data directory. For more information on LifeKeeper and the `mysqld` Group feature, see [Using mysqld Groups with LifeKeeper](#).
- Additional port numbers for MySQL must be specified in the `/etc/services` file.
- Each MySQL database server must be configured to run on a different port and access a different socket file. These configuration options are specified in the `my.cnf` file in the data directory.
- Each server must be configured to access data from a different shared location (i.e. each server must use a different data directory).

## 6.7.2.8. Using mysqld Groups with LifeKeeper

---

The MySQL Application Recovery Kit supports my.cnf files using the mysqld group feature managed via [mysqld\\_multi](#). This MySQL feature allows multiple MySQL instances to be easily configured via a single my.cnf file (typically stored in */etc*.) The kit now detects a my.cnf file using the mysqld group format and prompts the administrator to select the number of the mysqld group to be protected. The choice list provided to the administrator is determined by the group numbers defined in the my.cnf file minus any group numbers already being protected by the kit.

In general, it is easier to set up and control multiple MySQL instances using the mysqld group feature, and SIOS recommends that this approach be used when setting up active/active or multiple instance configurations.

### my.cnf File

When using the mysqld group feature, the following are imperative:

- a. A single my.cnf file should be used for defining mysqld groups for the database instances.
- b. The my.cnf file should NOT be placed on shared storage.
- c. An exact copy of the my.cnf file needs to exist on each cluster node (*/etc/my.cnf* is ideal).
- d. Any changes made to the my.cnf file must be propagated to every node in the LifeKeeper cluster.

The recovery kit uses `mysqld_multi` commands when it detects the my.cnf file is using mysqld groups. Based on this, you should be able to use `mysqld_multi` to test your MySQL instance before placing it under control of LifeKeeper.

The following is a relatively complex my.cnf file using mysqld groups that describes two database instances controlled by `mysqld_multi`. The `mysqld_multi` command (and the MySQL LifeKeeper recovery kit) gives the administrator a lot of options on how things get set up. In the example below, `[mysqld1]` defines a relatively simple MySQL instance that uses most of the default locations for various MySQL directives. The second example `[mysqld55]` moves things around more. The comments will help describe what each section is doing in terms of LifeKeeper's interaction with MySQL.

```
#The following client section defines which username/password combination will be used for
#LifeKeeper connections. The username/password combination needs to be defined in each MySQL
# Database instance that will be described in this my.cnf file.
[client]
user      = steeleye
password = password
```

```

# This next section describes the default version of mysqld and mysqldadmin th
at mysqld_multi
# will use when processing mysqld_multi commands. The username/password combo
defines the
# MySQL account that mysqld_multi will use when working with the database inst
ances. This
# username and password combo needs to be defined in each MySQL Database insta
nce that will be
# controlled by mysqld_multi. See how to set up the multi_admin account in th
e MySQL Reference
# Manual, by issuing "mysqld_multi --example".
[mysqld_multi]
mysqld      = /usr/bin/mysqld_safe
mysqldadmin = /usr/bin/mysqldadmin
user        = multi_admin
<>password  = password

```

```

# The next section defines the first of two MySQL Database instances in this m
y.cnf file. Note
# that each section starts with a [mysqldNN] where NN is the mysqld group num
ber (or instance).
# Each group name must have a number. There are a number of directives that t
he LifeKeeper MySQL
# Recovery Kit will be looking for in these sections.
[mysqld1]
datadir = /s11/mysql-data5077          #Defines where the data files for the ins
tance will live. For
                                         # LifeKeeper, this directory must be on L
ifeKeeper protected
                                         # (shared or replicated) storage.
mysqld = /usr/bin/mysqld_safe          # Defines specifically which mysqld comm
and will be used for
                                         # starting the instance. This one is u
sing the
                                         # default mysqld_safe that came with
the distribution.
socket=/s11/mysql-data5077/moe.socket # Defines the location of the socket fo
r this instance.
                                         # If the socket is not on LifeKeeper p
rotected storage, it
                                         # needs to be defined in exactly the s
ame place on each
                                         # node in the cluster and be owned by
the "user" defined
                                         # below.<
port = 3307                             # Each instance needs its own, unique TC

```

```

P/IP port.
pid-file = /var/run/mysqld/mysqld.pid # The pid-file can be on LifeKeeper protected or
# non-LifeKeeper protected storage.
log-error= /var/log/mysqld.log # Location of the MySQL error log for this instance. Can be
# on LifeKeeper protected or non-LifeKeeper protected
# storage.
user = mysql # The Linux user name that will run the MySQL processes.

```

```

#The next section defines the more complicated of the two MySQL instances. Instance "55" is not
#using the default MySQL that came with the Linux distribution as it is using the 5.5.12 version
#of MySQL that was installed from source. The binaries for this version were installed onto shared
#storage, and the binary directory is LifeKeeper protected.
[mysqld55]
datadir = /s11/mysql-data5512 # Same as above; this instance uses a different data
# directory, and this directory is
on LifeKeeper # protected storage.
mysqld =/s11/mysql5512/bin/mysqld_safe # For this instance, a different version of mysqld_safe
# is used; the one that is included with 5.5.12.
socket=/s11/mysql-misc5512/larry.socket # This instance has the socket on LifeKeeper protected
# storage, but not in the default location (datadir).
port = 3308 # This instance has a unique TCP/IP port as well.
pid-file = /var/run/mysqld/mysqld55.pid # This instance's pid-file is not on LifeKeeper protected
# storage.
log-error = /var/log/mysqld55.log # This instance's log-error (error log) is not on
# LifeKeeper protected storage.
log-bin = /s11/mysql-log5512/larry # The log-bin directive specifies where the binary
# transaction logs are located for this instance.
# These logs must be on LifeKeeper

```

```
protected storage
# (the recovery kit will enforce thi
s). By default,
# these logs are in the datadir.
user = mysql # The Linux user name that will run th
e MySQL processes.
```

When describing both sets up of [mysqld<N>] for multi instance and [mysqld] for single instance, the set up for single instance must be described at the last part.

### Example:

```
[mysqld1]
(set up for mysqld1)

[mysqld2]
(set up for mysqld2)

[mysqld55]
(set up for mysqld55)

[mysqld]
(set up for mysqld for single instance )
```

## mysqld\_multi Commands

For this example, issuing the mysql command:

```
# mysqld_multi start 1
```

would start the mysqld group 1 instance defined in my.cnf as [mysqld1], assuming all of the LifeKeeper protected resources that it depends on were in service on one of the LifeKeeper nodes.

Issuing the mysql command:

```
# mysqld_multi report 1
```

would report on the status of this instance (e.g. running or not running). Once this instance is running, creating a resource for it in LifeKeeper should be easy.

To get more information on setting up a mysqld\_multi style my.cnf file, issue the command:

```
# mysqld_multi -example
```

## 6.7.2.9. Using Network Attached Storage

---

There are a couple of special considerations to take into account when configuring LifeKeeper to use an NFS file server (Network Attached Storage) as cluster storage.

### Use the NAS Recovery Kit

The optional Network Attached Storage (NAS) recovery kit is required when using an NFS server as a shared storage array with LifeKeeper for Linux. Install the NAS recovery kit (and a license) on each cluster node. See the [NAS Recovery Kit](#) documentation for more details.

### Possible Error Message

When using Network Attached Storage (NAS) with MySQL, you may experience MySQL instances not restarting following a failover due to a system crash. The MySQL error log should indicate the cause of the error.

#### MySQL 5.0

```
110523 22:10:58 mysqld started
InnoDB: Unable to lock ./ibdata1, error: 11
InnoDB: Check that you do not already have another mysqld process
InnoDB: using the same InnoDB data or log files.
110523 22:10:58 InnoDB: Retrying to lock the first data file
InnoDB: Unable to lock ./ibdata1, error: 11
InnoDB: Check that you do not already have another mysqld process
InnoDB: using the same InnoDB data or log files.
```

#### MySQL 5.5

```
110524 10:52:20 InnoDB: The InnoDB memory heap is disabled
110524 10:52:20 InnoDB: Mutexes and rw_locks use GCC atomic builtins
110524 10:52:20 InnoDB: Compressed tables use zlib 1.2.3
110524 10:52:20 InnoDB: Initializing buffer pool, size = 128.0M
110524 10:52:20 InnoDB: Completed initialization of buffer pool
InnoDB: Unable to lock ./ibdata1, error: 11
InnoDB: Check that you do not already have another mysqld process
InnoDB: using the same InnoDB data or log files.
110524 10:52:20 InnoDB: Retrying to lock the first data file
InnoDB: Unable to lock ./ibdata1, error: 11
InnoDB: Check that you do not already have another mysqld process
InnoDB: using the same InnoDB data or log files.
```

This indicates that the MySQL mysqld process has set an NFS lock on the file "ibdata1" on the NFS file system that is being controlled by LifeKeeper. The lock was not cleared by the system crash, so

LifeKeeper is unable to bring the MySQL instance back into service. MySQL thinks that some other process is using the *ibdata1* file.

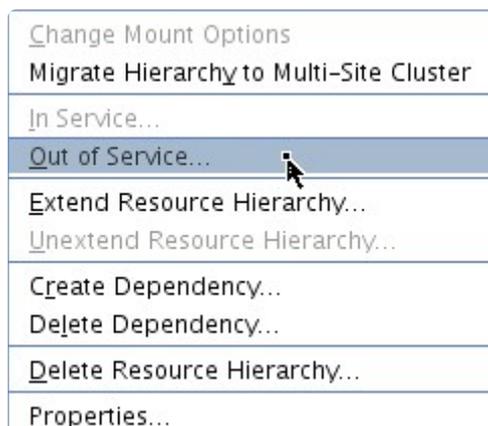
## Solution

To fix this, mount the NFS file system that will hold *ibdata1* with the “*no-lock*” NFS option before the File System resource is created. By default, NFS allows file locks to be set. If the “*no-lock*” option is used before resource creation, LifeKeeper will pick up this option and use it each time it brings the file system resource in service. Since LifeKeeper will be controlling access (from the cluster nodes) to the file system containing *ibdata1*, the lock is not typically critical. The NFS mount options used during testing were “*rw, sync, tcp, nfsvers=3, no-lock*”.

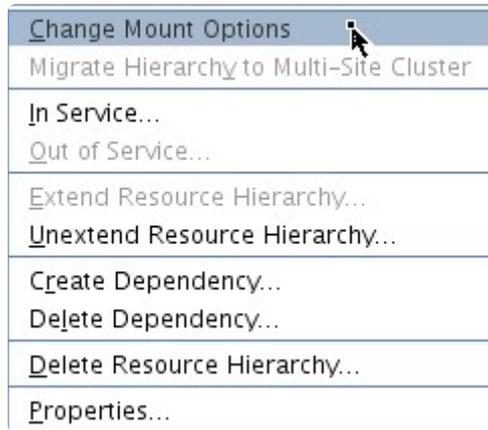
It is not necessary to use the “*no-lock*” on other file systems used by the MySQL resource hierarchy such as the file system where the MySQL binaries are located.

If the NAS File System resource has already been created without the “*no-lock*” option set, use the following procedure to change the mount option:

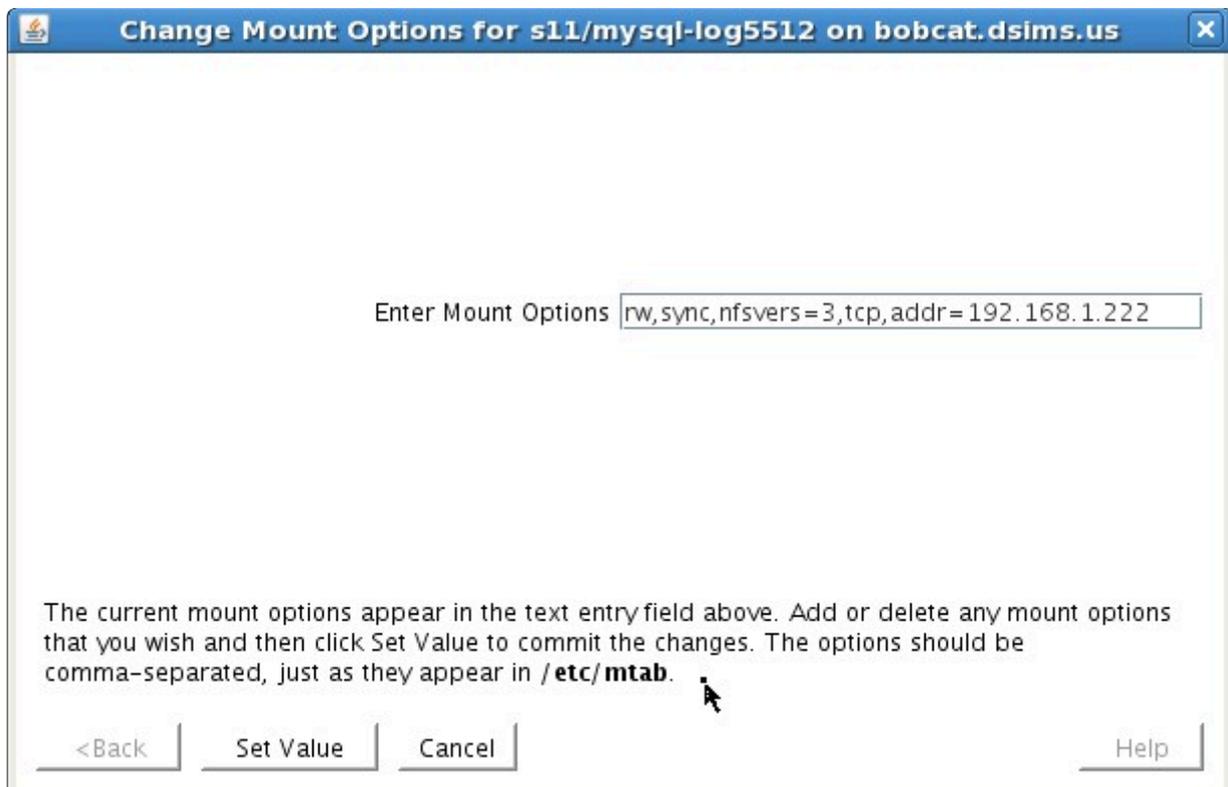
1. Using the LifeKeeper GUI, take the file system resource that needs to be changed out of service. This can be done from the LifeKeeper GUI putting the pointer on the file system resource and doing a right mouse click, and select **Out of Service** from the drop-down menu. This action may take parent resources out of service as well.



2. Confirm the **Out of Service** action and allow the process to complete.
3. Once the file system resource is out of service, you can put the pointer on the resource and do another right mouse click, and from the drop-down menu select **Change Mount Options**.



- In the popup window, add **noexec** to the line of options, and click **Set Value**. You will need to repeat steps 3 and 4 for each node in the cluster.



- Bring the NAS File System resource back in service by doing a right mouse click, and selecting **In Service**.
- The File System resource's property panel should now reflect that "noexec" is one of the current mount options.

bobcat.dsims.us: s11/mysql-log55 12

Filesystem Status | General | Equivalencies | Relations

Mount Point: /s11/mysql-log55 12  
Device: 192.168.1.222:/export/s11/mysql-log55 12  
Type: nfs  
Mount Options: rw, sync, nfsvers=3, tcp, addr=192.168.1.222, nolock

Size: 20G  
Used: 1.4G  
Free: 18G  
Usage: 8% 

Apply Changes | Reset | Help

## 6.7.2.10. Considerations on MySQL use in systemd Environments

If MySQL (version 5.7.6 or later) is installed on a OS distribution adopting systemd, the mysqld\_safe and mysqld\_multi commands are not installed and thus unavailable for LifeKeeper use. In these environments, LifeKeeper will use the systemctl command to start and stop the MySQL service.

Set up MySQL referring to the article [Managing MySQL Server with systemd](#). Specially set up the PID File as this is required for Systemd and LifeKeeper. Systemd MySQL set up and my.cnf set up must be the same on all nodes.

The following set up items must be defined for resource creation.

Set up item	Value to set
<i>Location of my.cnf</i>	The full path of the directory that contains the my.cnf file used when starting the MySQL service with the systemctl command
<i>Location of MySQL executables</i>	The full path of the directory where the mysqladmin command is installed

**\* Notes:** The setting of “Location of my.cnf” is used inside LifeKeeper to read the settings in my.cnf. The value set up here is not used by the systemctl startup command. In the case where the my.cnf file is in a different path from the default, set up MySQL for Systemd correctly and make sure the MySQL service starts and stops as expected. Also, the “include” directive is not supported. All the configuration information must be described in a single my.cnf file.

## 6.7.3. Installing/Configuring MySQL with LifeKeeper

---

[LifeKeeper Configuration Tasks](#)

[Creating a MySQL Resource Hierarchy](#)

[Deleting a Resource Hierarchy](#)

[Extending Your Hierarchy](#)

[Unextending Your Hierarchy](#)

## 6.7.3.1. LifeKeeper Configuration Tasks for MySQL

---

You can perform the following configuration tasks from the LifeKeeper GUI. The following four tasks are described in this section, as they are unique to a MySQL resource instance, and different for each Recovery Kit.

- [Create a Resource Hierarchy](#). Creates an application resource hierarchy in your LifeKeeper cluster.
- [Delete a Resource Hierarchy](#). Deletes a resource hierarchy from all servers in your LifeKeeper cluster.
- [Extend a Resource Hierarchy](#). Extends a resource hierarchy from the primary server to a backup server.
- [Unextend a Resource Hierarchy](#). Unextends (removes) a resource hierarchy from a single server in the LifeKeeper cluster.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View / Edit Properties](#). View or edit the properties of a resource hierarchy on a specific server.

 **Note:** Throughout the rest of this section, we explain how to configure your Recovery Kit by selecting certain tasks from the **Edit** menu of the LifeKeeper GUI. You can also select each configuration task from the toolbar. You can also right-click a global resource in the **Resource Hierarchy Tree** (left-hand pane) of the status display window to display the same drop-down menu choices as the **Edit** menu.

You can also right-click a resource instance in the **Resource Hierarchy Table** (right-hand pane) of the status display window to perform all the configuration tasks, except **Creating a Resource Hierarchy**, depending on the state of the server and the particular resource.

## 6.7.3.2. Creating a MySQL Resource Hierarchy

### ! IMPORTANT:

In a LifeKeeper cluster environment where the MySQL data directory (datadir) files are on a shared/replicated disk, you must make sure that the shared file system is mounted on the primary/template server. If the file system resource is created first, the shared file system **MUST** be mounted on the same mount point on each server. It is also important to remember that a working communication path (i.e. heartbeat) is required before you can create your resource. The MySQL data directory can exist on shared, replicated or network attached storage.

Replicated (SIOS DataKeeper) file system resources must be created before creating the MySQL resource. File systems on other storage types will automatically be created during MySQL hierarchy creation.

To create a resource instance from the primary server, you should complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

If you wish to change a selection you have already entered or encounter an error message during any step in the creation of your MySQL resource hierarchy, you will generally be able to back up and change your selection or make corrections (assuming the **Back** button is enabled).

\* **Important:** The MySQL database server daemon (mysqld) for the MySQL instance you want to protect must be running when you create the resource.

A dialog box will appear with a drop-down menu listing all recognized Recovery Kits installed within the cluster. Select **MySQL Database** from the drop-down menu.

Please Select Recovery Kit

Click **Next**.

If you click the **Cancel** button at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the MySQL instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either *intelligent* or *automatic*. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. Automatic switchback means the

switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.

Switchback Type

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box.

Click **Next**.

- 3. Select the **Server** where you want to place the MySQL database instance (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop-down menu.

Server

Click **Next** to proceed to the next dialog box.

- 4. Select or enter the **Location of my.cnf**. This is the full path name (excluding the file name) where the MySQL configuration file (*my.cnf*) is located.

Location of my.cnf

Click **Next** to proceed to the next dialog box.

- 5. Select the **Protection Instance Number** if you have a `mysqld_multi` style `my.cnf` file. If you are using a more traditional style `my.cnf` file, you will not see this screen.

Select protection instance number

- 6. Select or enter the\* **Location of MySQL executables location**\*. This is the full path name of the binaries used to start and monitor the MySQL database server daemon.

Location of MySQL executables

 **Note:** At this point, LifeKeeper will validate that you have provided valid data to create your MySQL resource hierarchy. If LifeKeeper detects a problem with either of this validation, an ERROR will appear on the screen. If the directory paths are valid, but there are errors with the MySQL configuration itself, you may pause to correct these errors and continue with the hierarchy creation.

Click **Next** to proceed to the next dialog box.

7. Select or enter the **Database Tag**. This is a tag name given to the MySQL hierarchy. You can select the default or enter your own tag name.

Database Tag

When you click **Create**, the **Create Resource Wizard** will create your MySQL resource.

Creating database/mysql resource...

```
Tue Jun 21 16:02:02 EDT 2011 create: BEGIN creation of "mysql-55" on server "bobcat.dsims.us"
Tue Jun 21 16:02:12 EDT 2011 create: 102045: Executable path "/s11/mysql5512/bin" is on a
shared file system.
Tue Jun 21 16:02:14 EDT 2011 create: 102045: socket path
"/s11/mysql-misc5512/larry.socket" is on a shared file system.
Tue Jun 21 16:02:28 EDT 2011 create: END successful creation of "mysql-55" on server
"bobcat.dsims.us"
***WARNING*** perform_action;Tue Jun 21 16:02:29 EDT 2011: License key (for Kit
database/mysql) will expire at midnight in 49 days
Tue Jun 21 16:02:29 EDT 2011 restore: BEGIN restore of "mysql-55" on server "bobcat.dsims.us"
Tue Jun 21 16:02:29 EDT 2011 restore: END successful restore of "mysql-55" on server
"bobcat.dsims.us"
```

 **Note:** The MySQL resource hierarchy should be created successfully at this point.

8. Another information box will appear explaining that you have successfully created an MySQL resource hierarchy, and you must **Extend** that hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

You have successfully created the resource hierarchy mysql-55 on bobcat.dsims.us. Select a target server to which the hierarchy will be extended.

If you cancel before extending mysql-55 to at least one other server, LifeKeeper will provide no protection for the applications in the hierarchy.

When you click **Continue**, LifeKeeper will launch the **Pre-Extend Wizard** that is explained in the next section.

If you click **Cancel** now, a dialog box will appear warning you that you will need to come back and extend your MySQL resource hierarchy to another server at some other time to put it under LifeKeeper protection.

Hierarchy Verification Finished

WARNING: Your hierarchy exists on only one server. Your  
WARNING: application has no protection until you extend it  
WARNING: to at least one other server.

9. Click **Done** to exit.

## 6.7.3.3. Deleting a MySQL Resource Hierarchy

To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, and then **Resource**. From the drop-down menu, select **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your MySQL resource hierarchy.

\* **Note:** If you selected the **Delete Resource** task by right-clicking from the right pane on an individual resource instance, or from the left pane on a global resource where the resource is on only one server, this dialog box will not appear.

Target Server bobcat.dsims.us

Click **Next**.

3. Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete, and highlight it.

\* **Note:** If you selected the **Delete Resource** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

Hierarchy to Delete

s11/mysql-log5077  
ip-moe-1.41  
s11/mysql-data5077  
ip-larry-1.42  
mysql-55-larry

Click **Next**.

- An information box appears confirming your selection of the target server and the hierarchy you have selected to delete.

```
You have specified the following resource hierarchy for deletion.  
Target Server: bobcat.dsims.us  
Target Tags:  
mysql-55-larry
```

Click **Delete**.

- Another information box appears confirming that the MySQL resource was deleted successfully.

```
Deleting resource hierarchy mysql-55-larry  
Removing root resource hierarchy starting at "mysql-55-larry":  
Mon Jun 27 17:28:42 EDT 2011 delete: BEGIN delete of "mysql-55-larry" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:42 EDT 2011 delete: END successful delete of "mysql-55-larry" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:43 EDT 2011 delete: BEGIN delete of "device-nfs31707" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:43 EDT 2011 delete: END successful delete of "device-nfs31707" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:43 EDT 2011 delete: BEGIN delete of "device-nfs30658" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:43 EDT 2011 delete: END successful delete of "device-nfs30658" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:43 EDT 2011 delete: BEGIN delete of "device-nfs30733" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:43 EDT 2011 delete: END successful delete of "device-nfs30733" on server  
"bobcat.dsims.us"  
Mon Jun 27 17:28:43 EDT 2011 delete: BEGIN delete of "device-nfs31659" on server  
"bobcat.dsims.us"  
Successfully removed  
Mon Jun 27 17:28:43 EDT 2011 delete: END successful delete of "device-nfs31659" on server  
"bobcat.dsims.us"
```

- Click **Done** to exit.

## 6.7.3.4. Extending Your MySQL Hierarchy

After you have created a hierarchy, you will want to extend that hierarchy to another server in the cluster. There are three possible scenarios to extend your resource instance from the template server to a target server. The first scenario is when you **Continue** from creating the resource into extending that resource to another server. The second scenario is when you enter the **Extend Resource Hierarchy** task from the edit menu as shown below. The third scenario is when you right-click on an unextended hierarchy in either the left or right pane. Each scenario takes you through the same dialog boxes (with a few exceptions, which are clearly detailed below).

1. If you are entering the **Extend** wizard from the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop-down menu, select **Extend Resource Hierarchy**. This will launch the **Extend Resource Hierarchy** wizard.
2. The first dialog box to appear will ask you select the **Template Server where your MySQL resource hierarchy is currently in service**. It is important to remember that the **Template Server** you select now and the **Tag to Extend** that you select in the next dialog box represent an in service resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you selected. The drop-down box in this dialog provides the names of all the servers in your cluster.

**Note:** If you are entering the **Extend Resource Hierarchy** task immediately following the creation of a MySQL resource hierarchy, this dialog box will not appear, since the wizard has already identified the template server in the create stage. This is also the case when you right-click either the MySQL resource icon in the left pane or right-click on the MySQL resource box in the right pane the of the GUI window and choose **Extend Resource Hierarchy**.



It should be noted that if you click **Cancel** at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extension process to that particular server. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

For example, let us say you have created your resource on Server 1 and extended that resource to Server 2. In the middle of extending the same resource to Server 3, you change your mind and click **Cancel** inside one of the dialog boxes. This will cancel only your action to extend the resource to Server 3, not the extension you created to Server 2. If you want to remove Server 2 from this hierarchy, you must unextend the resource from Server 2.

Click **Next** to proceed to the next dialog box.

3. Select the **Tag to Extend**. This is the name of the MySQL instance you wish to extend from the template server to the target server. The wizard will list in the drop-down menu all the resources that you have created on the template server, which you selected in the previous dialog box.

**Note:** Once again, if you are entering the Extend Resource Hierarchy task immediately following the creation of a MySQL resource hierarchy, this dialog box will not appear, since the wizard has already identified the tag name of your MySQL resource in the create stage. This is also the case when you right-click either the MySQL resource icon in the left hand pane or on the MySQL resource box in the right hand pane of the GUI window and choose **Extend Resource Hierarchy**.

Tag to Extend

Click **Next**.

4. Select the **Target Server** where you are extending your MySQL resource hierarchy. The drop-down box provides the names of the servers in your cluster that are not already in the selected hierarchy.

Target Server

Click **Next**.

5. Select the **Switchback Type**. This dictates how the MySQL instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either *intelligent* or *automatic*. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. Automatic switchback means the switchback will occur as soon as the primary server comes back online and reestablishes LifeKeeper communication paths.

Switchback Type

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box.

Click **Next**.

6. Select or enter a **Template Priority**. This is the priority for the MySQL hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. **Note:** This selection will appear only for the initial extend of the hierarchy.

Click **Next**.

7. Select or enter the **Target Priority**. This is the priority for the new extended MySQL hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid,

indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.



Target Priority

Click **Next**.

- An information box will appear explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this MySQL resource have been met. If there were some requirements that had not been met, LifeKeeper would not allow you to select the **Next** button, and the **Back** button would be enabled.

```
Executing the pre-extend script...
Building independent resource list
Checking existence of extend and canextend scripts
Checking extendability for mysql-55

Pre Extend checks were successful
```

If you click **Back**, you can make changes to your resource extension according to any error messages that may appear in the information box.

If you click **Cancel** now, you will need to come back and extend your MySQL resource hierarchy to another server at some other time to put it under LifeKeeper protection.

When you click **Next**, LifeKeeper will launch you into the **Extend Resource Hierarchy** configuration task.

- This dialog box is for information purposes only. You cannot change the **Location of my.cnf** that appears in the box. The MySQL instance acquired the location information from its configuration file.

Location of my.cnf

Click **Next**.

- Select or enter the **Location of MySQL executables**. This is the full path name of the binaries used to start and monitor the MySQL database server daemon.

Location of MySQL executables

Click **Next**.

11. Select or enter the **Database Tag**. This is a tag name given to the MySQL hierarchy. You can select the default or enter your own tag name.

Tag to Extend

Click **Extend**.

12. An information box will appear verifying that the extension is being performed.

```

Extending resource hierarchy mysql-55 to server lion.dsims.us
-----
Extending resource instances for mysql-55
Creating dependencies
Setting switchback type for hierarchy
Creating equivalencies
LifeKeeper Admin Lock (mysql-55) Released

Hierarchy successfully extended

```

Click **Next Server** if you want to extend the same MySQL resource instance to another server in your cluster. This will repeat the **Extend Resource Hierarchy** operation.

If you click **Finish**, LifeKeeper will verify that the extension of the MySQL resource was completed successfully.

13. If you clicked **Finish**, the following screen appears.

```

Verifying Integrity of Extended Hierarchy...
-----
Examining hierarchy on lion.dsims.us

Hierarchy Verification Finished

```

14. Click **Done** in the last dialog box to exit.

**Note:** Be sure to test the functionality of the new instance on both servers.

## 6.7.3.5. Unextending Your MySQL Hierarchy

1. From the LifeKeeper GUI menu, select **Edit**, and **Resource**. From the drop-down menu, select **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the MySQL resource. It cannot be the server where the MySQL resource is currently in service.

\* **Note:** If you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance, this dialog box will not appear.

Target Server

Click **Next**.

3. Select the MySQL **Hierarchy to Unextend**.

\* **Note:** If you selected the **Unextend** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

Hierarchy to Unextend

Click **Next**.

4. An information box appears confirming the target server and the MySQL resource hierarchy you have chosen to unextend.

```
You have specified the following resource hierarchy for unextend.  
Target Server = lion.dsims.us  
Target Tag = mysql-55-larry
```

Click **Unextend**.

5. Another information box appears confirming that the MySQL resource was unextended successfully.

```
Unextending resource hierarchy mysql-55-larry from lion.dsims.us
Hierarchy Unextend Manager Initializing
Checking Target Machine Communication Paths
LifeKeeper Admin Lock Flag (mysql-55-larry) Established
Removing Equivalencies
Removing Resources and Associated Dependencies
Mon Jun 27 11:56:24 EDT 2011 delete: BEGIN delete of "mysql-55-larry" on server
"lion.dsims.us"
Mon Jun 27 11:56:24 EDT 2011 delete: END successful delete of "mysql-55-larry" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: BEGIN delete of "device-nfs31707" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: END successful delete of "device-nfs31707" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: BEGIN delete of "device-nfs30658" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: END successful delete of "device-nfs30658" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: BEGIN delete of "device-nfs30733" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: END successful delete of "device-nfs30733" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: BEGIN delete of "device-nfs31659" on server
"lion.dsims.us"
Mon Jun 27 11:56:25 EDT 2011 delete: END successful delete of "device-nfs31659" on server
"lion.dsims.us"
LifeKeeper Admin Lock Flag (mysql-55-larry) Released
Synchronizing LifeKeeper Databases
Unextend completed successfully
```

6. Click **Done** to exit.

## 6.7.4. MySQL Administration

---

### Testing Your Resource Hierarchy

You can test your MySQL resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

[Performing a Manual Switchover from the GUI](#)

## 6.7.4.1. Performing a Manual Switchover from the GUI

---

You can test your MySQL resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource** and **In-Service** from the drop-down menu. For example, an **In-Service** request executed on a backup server causes the application hierarchy to be placed in service on the backup server and taken out of service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the **Out-of-Service** request, the application is taken out of service without bringing it in service on the other server.

LifeKeeper does not regulate or control internal operations such as rollbacks and backing up archives. Tape archiving and restoration are the responsibility of the application administrator.

### Recovery Operations

When the primary server fails, the MySQL Recovery Kit software performs the following tasks:

- Mounts the file system(s) – shared or replicated – on the backup server
- Starts the daemon processes related to MySQL

## 6.7.5. MySQL Troubleshooting

### Common Error Messages

This section provides a list of messages that you may encounter while creating and extending a LifeKeeper MySQL resource hierarchy or removing and restoring a resource. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Messages from other LifeKeeper components are also possible.

In these cases, please refer to the [Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

### MySQL Specific Error Messages

**Note:** In the Error Message column, a word in quotations and all capital letters refers to the name of a resource on the server (for example, "SERVER" might actually be a server named "Server1").

Error Number	Error Message
102001	Usage: "SCRIPT NAME" sysname dbvarname cnfpath exepath instance
102002	Usage: "SCRIPT NAME" cnfpath
102003	Usage: "SCRIPT NAME" exepath cnfpath
102004	Unable to obtain a valid value for the "socket" variable in "PATH"/my.cnf <b>Action:</b> There must be an entry for the "socket" in the 'mysqld' section of the my.cnf configuration file
102005	Unable to obtain a valid value for the "port" in "PATH"/my.cnf <b>Action:</b> There must be an entry for the "port" in the 'mysqld' section of the my.cnf configuration file
102006	Unable to obtain the data directory location "PATH" <b>Action:</b> Please make sure that the database is running using the socket and port specified.
102007	Must specify the absolute path to the my.cnf configuration file
102008	Must specify the absolute path to the MySQL executables
102009	The file my.cnf does not exist in the path specified

102010	The MySQL executables do not exist in the path specified
102011	LifeKeeper was unable to start the MySQL database server
102012	LifeKeeper successfully started the MySQL database server
102013	LifeKeeper was unable to stop the MySQL database server
102014	LifeKeeper successfully stopped the MySQL database server
102015	The port "PORT NUMBER" is in use on the target server "SERVER"
102016	The MySQL database server is not running on server "SERVER"
102017	Unable to open the configuration file "PATH"/my.cnf
102018	Unable to get the Data Directory information for resource "TAG" on server "SERVER"
102019	Unable to get the configuration file location information for resource "TAG" on server "SERVER"
102020	Unable to get the executable location information for resource "TAG" on server "SERVER"
102021	The argument for the configuration file path is empty
102022	The argument for the executable path is empty
102023	The path "PATH" for directive "DIRECTIVE" is not on a shared filesystem
102024	Unable to get the information for resource "TAG" on system "SYSTEM"
102025	The MySQL data directory "DATADIR" is already under LifeKeeper protection
102026	The port variables in the file /etc/my.cnf on "SERVER1" and "SERVER2" do not match
102027	The socket variables in the file /etc/my.cnf on "SERVER1" and "SERVER2" do not match
102028	The MySQL database server is not running on server "SERVER" <b>Action:</b> There must be a valid entry for the "user" variable in the 'client' section of the my.cnf configuration file
102029	Unable to obtain a valid value for the "password" variable in "PATH"/my.cnf <b>Action:</b> There must be a valid entry for the "password" variable in the 'client' section of the my.cnf configuration file
102030	The user variables in the file /etc/my.cnf on "SERVER1" and "SERVER2" do not match
102031	The password variables in the file /etc/my.cnf on "SERVER1" and

	"SERVER2" do not match
102032	Unable to obtain the pid file location <b>Action:</b> There must be an entry for the "pid-file" variable in the 'mysqld' section of the my.cnf configuration file
102033	Unable to obtain a valid value for the "user" variable in "PATH"/my.cnf <b>Action:</b> The OS user must be specified using the "user" variable in the 'mysqld' section of the my.cnf configuration file
102034	WARNING: A my.cnf file exists at "PATH", which may override the values specified in the file at "PATH"/my.cnf.
102035	The mysql system user "USER" does not exist on target server "SERVER"
102036	The mysql system user "USER" uids are different on target server "SERVER1" and template server "SERVER2"
102037	The mysql system user "USER" gids are different on target server "SERVER1" and template server "SERVER2"
102038	LifeKeeper was unable to stop the MySQL database server using a graceful shutdown. Issuing kill for pid(s): "PROCESS ID LIST".
102039	LifeKeeper will ignore failed connection as possible max connections error, due to existence of process pid "PROCESS ID".
102040	The mysql action for resource tag :TAG" returned: "COMMAND OUTPUT".
102041	LifeKeeper was unable to start the MySQL database server using the defaults-file option. Retrying with individual options.
102042	The LifeKeeper "ACTION" action detected the flag "FLAG", and will exit.
102043	END of "ACTION" action on due to a(n) "SIGNAL" signal.
102044	The file my.cnf does not exist in the stored path "PATH".
102045	"DIRECTIVE" path "PATH" is on a shared filesystem.
102046	Starting mysqld daemon with databases from "PATH".

# 6.8. WebSphere MQ Recovery Kit Administration Guide

---

The LifeKeeper for Linux WebSphere MQ Recovery Kit provides fault resilient protection for WebSphere MQ queue managers and queue manager storage locations. This kit enables a failure on a primary WebSphere MQ server or queue manager to be recovered on the primary server or a designated backup server without significant lost time or human intervention.

## Document Contents

This guide contains the following topics:

- [LifeKeeper for Linux Documentation](#). Provides a list of LifeKeeper for Linux documentation and where to find it.
- [Abbreviations](#). Contains a list of abbreviations that are used throughout this document along with their meaning.
- [Requirements](#). Describes the hardware and software necessary to properly set up, install and operate the WebSphere MQ Recovery Kit. Refer to the [LifeKeeper Installation Guide](#) for specific instructions on how to install or remove LifeKeeper for Linux software.
- [WebSphere MQ Recovery Kit Overview](#). Provides a brief description of the WebSphere MQ Recovery Kit's features and functionality as well as lists the versions of the WebSphere MQ software supported by this Recovery Kit.
- [WebSphere MQ Configuration Considerations](#). Provides a general description of configuration issues and shows file system layouts supported by the WebSphere MQ Recovery Kit.
- [Configuring WebSphere MQ for Use with LifeKeeper](#). Provides a step-by-step guide of how to install and configure WebSphere MQ for use with LifeKeeper.
- [Configuration Changes Post Resource Creation](#). Provides information on how WebSphere MQ configuration changes affect LifeKeeper WebSphere MQ resource hierarchies.
- [WebSphere MQ Configuration Examples](#). Provides examples of typical WebSphere MQ configurations and the steps to configure your WebSphere MQ resources.
- [LifeKeeper Configuration Tasks](#). Describes the tasks for creating and managing your WebSphere MQ resource hierarchies using the LifeKeeper GUI.
- [WebSphere MQ Troubleshooting](#). Provides a list of informational and error messages with recommended solutions.
- [Appendices](#). Provide sample configuration files for WebSphere MQ and a configuration sheet that can be used to plan your WebSphere MQ installation.

## Reference Documents

MQ documentation is located at the WebSphere MQ Library available at:

<http://www.ibm.com/software/integration/wmq/library/>

## LifeKeeper Documentation

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper Installation Guide](#)
- [Optional Recovery Kit Documentation](#)
- [LifeKeeper for Linux IP Recovery Kit Administration Guide](#)

This documentation, along with documentation associated with other LifeKeeper Recovery Kits, is available online at:

<http://docs.us.sios.com/>

## 6.8.1. MQ Recovery Kit Abbreviations

The following abbreviations are used throughout this document:

Abbreviation	Meaning
HA	Highly Available, High Availability
QMDIR	<p>WebSphere MQ queue manager directory. This directory holds the queue manager persistent queue data and is typically located in <code>/var/mqm/qmgrs</code> with the name of the queue manager as subdirectory name. The exact location of this directory is specified in the global <code>mqs.ini</code> configuration file.</p> <p>If the <code>DataPath</code> parameter is defined then the <code>DataPath</code> value along with queue manager name specifies the location of the queue manager persistent data, otherwise the default location as noted above is used.</p>
QMLOGDIR	WebSphere MQ queue manager log directory. This directory holds the queue manager log data and is typically located in <code>/var/mqm/log</code> with the queue manager name as subdirectory. The exact location of this directory is specified in the queue manager configuration file ( <code>QMDIR/qm.ini</code> ).
MQUSER	The operating system user running all WebSphere MQ commands. This user is the owner of the <code>QMDIR</code> . The user must be a member of the <code>MQGROUP</code> administrative group <code>mqm</code> (see below).
MQGROUP	The operating system user group that the <code>MQUSER</code> must be part of. This group must be named <code>mqm</code> .
UID	Numeric user id of an operating system user.
GID	Numeric group id of an operating system user group.

## 6.8.2. MQ Recovery Kit Requirements

---

Your LifeKeeper configuration must meet the following requirements prior to the installation of the WebSphere MQ Recovery Kit. Please see the [LifeKeeper Installation Guide](#) for specific instructions regarding the installation, removal and configuration of your LifeKeeper hardware and software.

## 6.8.2.1. MQ Hardware and Software Requirements

---

### Hardware Requirements

- **Servers.** The Recovery Kit requires two or more servers configured in accordance with the requirements described in the [LifeKeeper Installation Guide](#). See the [Linux Configuration Table](#) for supported Linux distributions.
- **Data Storage.** The WebSphere MQ Recovery Kit can be used in conjunction both with shared storage and with replicated storage provided by the DataKeeper product. It can also be used with network-attached storage (NAS).

### Software Requirements

- **LifeKeeper Software.** You must install the same version of **LifeKeeper** software and any patches on each server.
- **LifeKeeper WebSphere MQ Recovery Kit.** Version 7.5.1 or later of the WebSphere MQ Recovery Kit is required for systems running WebSphere MQ v7.1 or later.
- **LifeKeeper IP Recovery Kit.** You must have the same version of the LifeKeeper IP Recovery Kit on each server.
- **IP Network Interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

**Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation and so forth.

- **TCP/IP Software.** Each server also requires the TCP/IP software.
- **WebSphere MQ Software.** IBM WebSphere MQ must be ordered separately from IBM. See the [LifeKeeper Release Notes](#) for supported WebSphere MQ versions. The WebSphere MQ Software must be installed on each server of the cluster prior to installing the WebSphere MQ Recovery Kit. The following WebSphere MQ packages must be installed to successfully install the WebSphere MQ Recovery Kit:

MQSeriesServer, MQSeriesSamples, MQSeriesClient, MQSeriesRuntime, MQSeriesSDK

Beginning with IBM WebSphere MQ Version 7.0.1 Fix Pack 6, a new feature was introduced allowing multiple versions of WebSphere MQ to be installed and run on the same server (e.g. MQ Versions 7.0.1 Fix Pack 6 and 7.1). This feature, known as multi-instance support, is now supported starting with version 9.0.1 of the WebSphere MQ Recovery Kit. Protecting multiple

queue managers within a single IBM WebSphere MQ installation version, protection of queue managers from multiple IBM WebSphere MQ installation versions, as well as the use the DataPath parameter in the `mqc.ini` file introduced as part of the multi-instance feature are all now supported in this version of the recovery kit.

- **Optional C Compiler.** The WebSphere MQ Recovery Kit contains a modified `amqsget0.c` sample program from the WebSphere MQ samples package. This program has been modified to work with a timeout of 0 seconds instead of the default 15 seconds. It is used to perform `PUT/GET` tests for the queue manager. This program is compiled during RPM installation and therefore a C compiler must be installed and must be located in the `PATH` of the “root” user.
- **Syslog.pm.** If you want to use syslog logging for WebSphere MQ resources, the `Syslog.pm` PERL module must be installed. This module is part of the standard PERL distribution and is not required to be installed separately.

## 6.8.2.2. Upgrading an MQ LifeKeeper Cluster

---

1. Upgrade LifeKeeper on all nodes in the cluster including the WebSphere MQ Recovery Kit following the instructions documented in the Upgrading LifeKeeper section of the [LifeKeeper Installation Guide](#).
2. Upgrade IBM WebSphere MQ software on each node in the cluster using the following steps:
  - a. If one or more LifeKeeper IBM WebSphere MQ resource hierarchies are in service on the node being upgraded, they must be taken out of service before the upgrade of the IBM WebSphere MQ software can be done. This can be done by switching over to the standby node.
  - b. Follow the IBM WebSphere upgrade instructions.
3. Once the IBM WebSphere software has been installed on the node, bring the LifeKeeper IBM WebSphere MQ resource hierarchies in service (restore) and verify the operation of each Queue Manager.
4. Once the operation of each Queue Manager is confirmed, upgrade all the other nodes in the cluster.

## 6.8.3. WebSphere MQ Recovery Kit Overview

---

WebSphere MQ (formerly known as MQSeries) is an IBM software product that provides reliable and guaranteed one time only delivery of messages. The core element of WebSphere MQ is the queue manager which handles one or more queues that are used to send (put) and receive (get) messages. Once a message is put into a queue, it is guaranteed that this message is persistent and will be delivered only once.

The WebSphere MQ Recovery Kit enables LifeKeeper to protect WebSphere MQ queue managers including the command server, the listener and the persistent queue manager data. Protection of the queue manager listener can be optionally disabled on a per queue manager basis to support configurations that do not handle client connects or to enable the administrator to shut down the listener without causing a LifeKeeper recovery attempt.

The WebSphere MQ Recovery Kit provides a mechanism to recover protected WebSphere MQ queue managers from a failed primary server onto a backup server. LifeKeeper can detect failures either at the server level (via a heartbeat) or resource level (by monitoring the WebSphere MQ daemons) so that control of the protected WebSphere MQ services are transferred to a backup server.

The WebSphere MQ Recovery Kit also supports multiple installations of WebSphere MQ to be installed and run on the same system. With multi-version MQ support a Queue Manager from MQ software version 7.x and a Queue Manager from MQ software version 8.x can both be protected by the Recovery Kit. Prior to the addition of this support in version 9.0.2 of the WebSphere MQ Recovery Kit only 1 version of the MQ software could be installed and running on the system (NOTE: the installation required it be installed in the default location – /opt/mqm).

# 6.8.3.1. MQ Recovery Kit Resource Hierarchies

A typical WebSphere MQ hierarchy will be comprised of a WebSphere MQ queue manager resource. It also contains one or more file system resources, depending on the file system layout and zero or more IP resources. The exact makeup of the hierarchy depends on what is being protected. If the administrator chooses to include an IP resource in the WebSphere MQ resource hierarchy, that IP must be created prior to creating the WebSphere MQ queue manager resource and that IP resource must be active on the primary server. Replicated (SIOS DataKeeper) file system resources must be created before creating the MQ resource. The file system hierarchies are created automatically during the creation of the WebSphere MQ queue manager resource.



Figure 1 Typical WebSphere MQ hierarchy – symbolic links

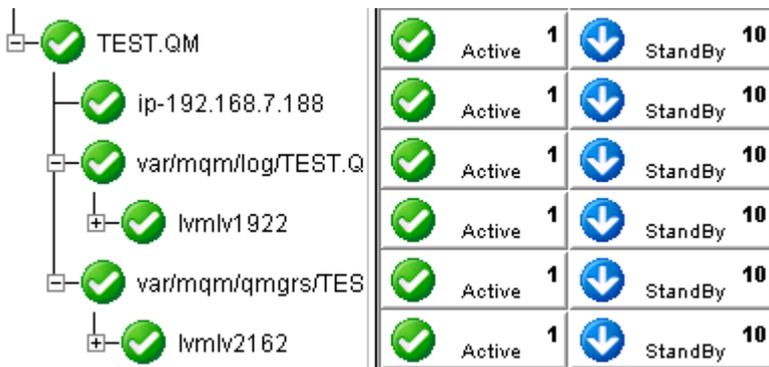


Figure 2 Typical WebSphere MQ hierarchy – LVM configuration

## 6.8.3.2. MQ Recovery Kit Features

---

The WebSphere MQ Recovery Kit provides the following features:

- Supports mult-instance Queue Managers (queue managers created with multiple versions of MQ software)
- Supports Active/Active configurations
- Supports LINEAR and CIRCULAR logging (detected automatically)
- Supports end to end application health check via server connect and client connect
- Supports optional PUT/GET tests (with definable test queue via GUI and command line)
- Supports customizable logging levels
- Supports all LifeKeeper supported storage types
- Supports optional listener protection (default: enabled)
- Supports additional syslog message logging (log facility local7)
- Supports multiple levels of Command Server protection (default: full)

## 6.8.4. WebSphere MQ Configuration Considerations

---

This section contains information that should be considered before beginning to configure WebSphere MQ. It also contains a step-by-step process for configuring and protecting a WebSphere MQ queue manager with LifeKeeper.

For instructions on installing WebSphere MQ on Linux distributions supported by LifeKeeper, please see [WebSphere MQ documentation](#).

## 6.8.4.1. MQ Configuration Requirements

The section [Configuring WebSphere MQ for Use with LifeKeeper](#) contains a process for protecting a queue manager with LifeKeeper. In general, the following requirements must be met to successfully configure a WebSphere MQ queue manager with LifeKeeper:

1. **Configure Kernel Parameters.** Please refer to the [WebSphere MQ documentation](#) for information on how Linux kernel parameters such as shared memory and other kernel resources should be configured.
2. **MQUSER and MQGROUP.** The MQGROUP and the MQUSER must exist on all servers in the LifeKeeper cluster. Websphere MQ software requires that the MQGROUP mqm exist and that it also have the MQUSER mqm defined that has its primary group membership set to the MQGROUP mqm. If the mqm user and mqm group do not exist at the time the Websphere MQ software is installed they will be automatically created. When installing the WebSphere MQ software most of the files and directories will have their user and group ownership set to the mqm user and mqm group. User and group ownership of the files and directories in the Queue Manager data and log directories will also be set to the mqm user and mqm group. Additionally, when a Queue Manager is started it will run as the mqm user. Therefore, the MQUSER user id (uid) and the MQGROUP group id (gid) must be the same on all servers in the cluster. The MQ Recovery Kit will verify this when attempting to extend the resource. If they do not match the resource extension will fail. Note: If you are using NIS, LDAP or another authentication tool besides the local password and group files you need to set up the MQUSER and MQGROUP prior to the installation of the Websphere MQ and LifeKeeper software. You may also need to create a home directory. If this is an upgrade from a prior release of the WebSphere MQ Recover Kit, then the MQUSER PATH environment variable setting may need to be modified. In prior releases of the Recovery Kit the MQUSER PATH environment variable needed to be modified to include the default install location of the WebSphere MQ software, /opt/mqm. If that change was made in a prior release it must be unset for this version of the Recovery Kit to function correctly.
3. **Alternate MQ user support.** Although the Websphere MQ software will always run as the mqm user an alternate user name can be specified for running all MQ commands provided the alternate user has primary or secondary membership in the mqm group. An alternate user name for starting WebSphere MQ may be required when integrating with other MQ Tools. To change to an alternate user see the MQS\_ALT\_USER\_NAME tunable in the “Changing LifeKeeper WebSphere MQ Recovery Kit Defaults” section of this document.
4. **Manual command server startup.** If you want to have LifeKeeper start the command server, disable the automatic command server startup using the following command on the primary server. Otherwise, the startup of the command server will be performed automatically when the Queue Manager is started:

```
runmqsc QUEUE.MANAGER.NAME
```

```
ALTER QMGR SCMDSERV(MANUAL)
```

5. **QMDIR and QMLOGDIR must be located on shared storage.** The queue manager directory

QM`DIR` and the queue manager log directory `QMLOGDIR` must be located on LifeKeeper-supported shared storage to let the WebSphere MQ on the backup server access the data. See [Supported File System Layouts](#) for further details.

6. **QM`DIR` and QMLOGDIR permissions.** The `QMDIR` and `QMLOGDIR` directories must be owned by `MQUSER` and the group `MQGROUP`. The ARK dynamically determines the `MQUSER` by looking at the owner of this directory. It also detects symbolic links and follows them to the final targets. Use the system command `chown` to change the owner of these directories if required.
7. **Disable Automatic Startup of Queue Manager.** Disable Automatic Startup for the queue manager(s) protected by LifeKeeper. You can disable using the `systemctl` command.
8. **Server Connection Channel Authorization.** Beginning with WebSphere MQ version 7.1 changes were made to channel authorization. By default the MQADMIN user (`mqm`) is unable to authenticate anonymously (no password) thus failing the resource hierarchy create (authorization for queue managers created with a WebSphere MQ release prior to 7.1 should continue to work). Starting with WebSphere MQ 7.1 one method to allow authorization for the MQADMIN user is to disable channel authorization. For WebSphere MQ 8.0 additional changes are required to the `authinfo` for `system.default.authinfo.idpwos` (in `runmqsc` run 'display `authinfo(system.default.authinfo.idpwos)`' to retrieve the current settings). The `chckclnt` setting of 'reqdamd' must be altered and set to 'optional'. Failure to allow the MQADMIN user anonymous authorization will result in the following error: 'MQCONNX ended with reason code 2035' during resource creation. See the [WebSphere MQ documentation](#) for details on how to create channels.
9. **MQSeriesSamples, MQSeriesSDK and MQSeriesClient Package.** LifeKeeper uses a client connection to WebSphere MQ to verify that the listener and the channel initiator are fully functional. This is a requirement for remote queue managers and clients to connect to the queue manager. Therefore, the `MQSeriesClient` package must be installed on all LifeKeeper cluster nodes running WebSphere MQ. Also, the `MQSeriesSDK` and `MQSeriesSamples` packages must be installed to perform client connect tests and `PUT/GET` tests.
10. **Optional C Compiler.** For the optional `PUT/GET` tests to take place, a C compiler must be installed on the machine. If not, a warning is issued during the installation.
11. **LifeKeeper Test Queue.** The WebSphere MQ Recovery Kit optionally performs a `PUT/GET` test to verify queue manager operation. A dedicated test queue has to be created because the recovery kit retrieves all messages from this queue and discards them. This queue should have set the default persistency setting to "yes" (`DEFPSIST=yes`) When you protect a queue manager in LifeKeeper, a test queue named "LIFEKEEPER.TESTQUEUE" will be automatically created. You can also use the following command to create the test queue manually before protecting the queue manager:

```
su - MQUSER
runmqsc QUEUE.MANAGER.NAME

define qlocal(LIFEKEEPER.TESTQUEUE) DEFPSIST(YES) DESCR('LifeKeeper
test queue')
```

**Note:** If you want to use a name for the LifeKeeper test queue other than the default “LIFEKEEPER.TESTQUEUE”, the name of this test queue must be configured. See [Editing Configuration Resource Properties](#) for details.

12. **TCP Port for Listener Object.** Alter the Listener object via runmqsc to reflect the TCP port in use. Use the following command to change the TCP port of the default Listener:

```
su – MQUSER
runmqsc QUEUE.MANAGER.NAME
```

```
alter LISTENER(SYSTEM.DEFAULT.LISTENER.TCP) TRPTYPE(TCP) PORT(1414)
IPADDR(192.168.1.100)
```

**Note:** The listener object must be altered even if using the default MQ listener TCP port 1414, but it is not necessary to set a specific IP address (`IPADDR`). If you skip the `IPADDR` setting, the listener will bind to all interfaces on the server. If you do set `IPADDR`, it is strongly recommended that a virtual IP resource be created in LifeKeeper using the `IPADDR` defined address. This ensures the IP address is available when the MQ listener is started.

13. **TCP Port Number.** Each WebSphere MQ listener must use a different port (default 1414) or bind to a different virtual IP with no listener binding to all interfaces. This includes protected and unprotected queue managers within the cluster.
14. **Queue Manager configured in `mqc.ini`.** In Active/Active configurations, each server holds its own copy of the global queue manager configuration file `mqc.ini`. In order to run the protected queue manager on all servers in the cluster, the queue manager must be configured in the `mqc.ini` configuration file of all servers in the cluster. Copy the appropriate QueueManager: stanza from the primary server and add it to the `mqc.ini` configuration files on all backup servers.

## 6.8.4.1.1. MQ Supported File System Layouts

---

Depending on your shared storage system and the file system layout, there are three different supported configurations. They differ in the file system layout. The following section describes the supported file system layouts.

- [Configuration 1 – /var/mqm on Shared Storage](#)
- [Configuration 2 – Direct Mounts](#)
- [Configuration 3 – Symbolic Links](#)

# 6.8.4.1.1.1. Configuration 1 – /var/mqm on Shared Storage

In this configuration, the whole /var/mqm directory is mounted on LifeKeeper supported shared storage (SCSI, SAN, NAS or replicated).

**Note:** This only works for Active/Passive configurations.

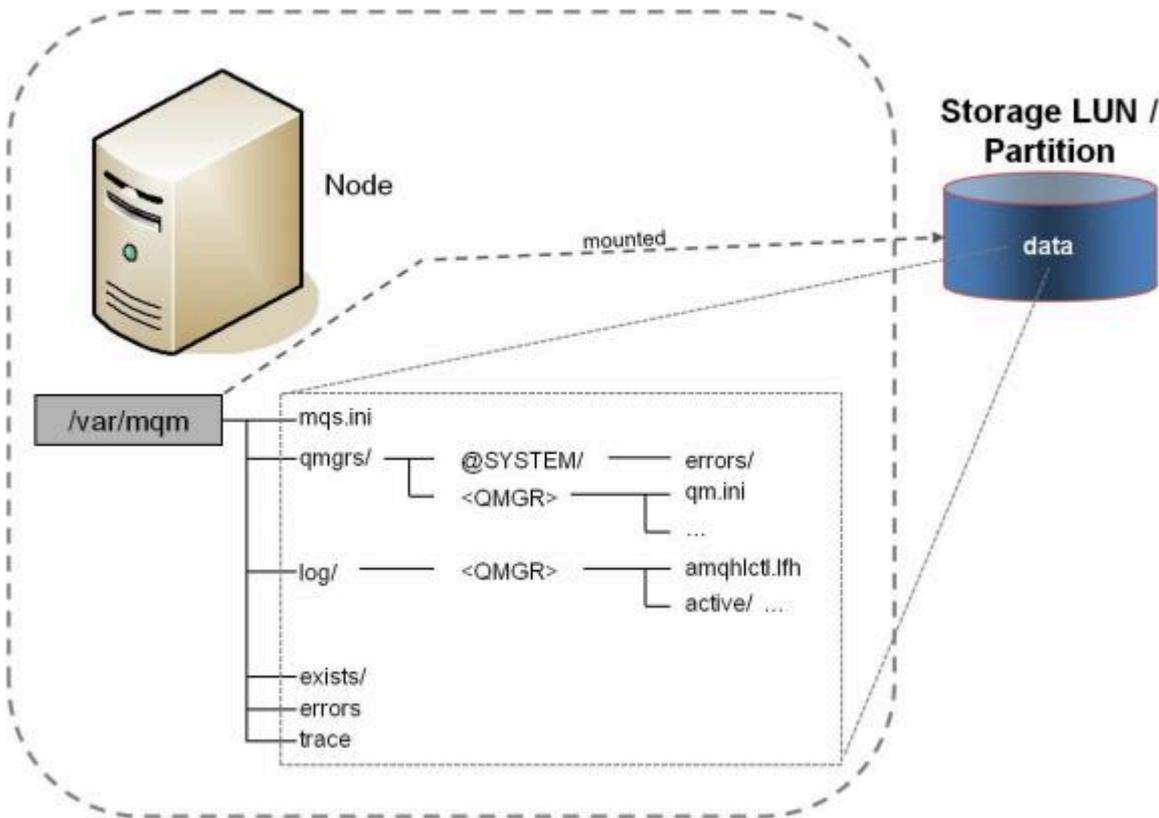


Figure 3 – File System Layout 1 – /var/mqm on Shared Storage

## 6.8.4.1.1.2. Configuration 2 – Direct Mounts

In this configuration, the *QMDIR* and the *QMLOGDIR* directories are located on shared storage. This requires two dedicated LUNS or partitions or the use of LVM for each queue manager. If LVM is used, two logical volumes from the same LUN can be created and separately mounted on the two directories.

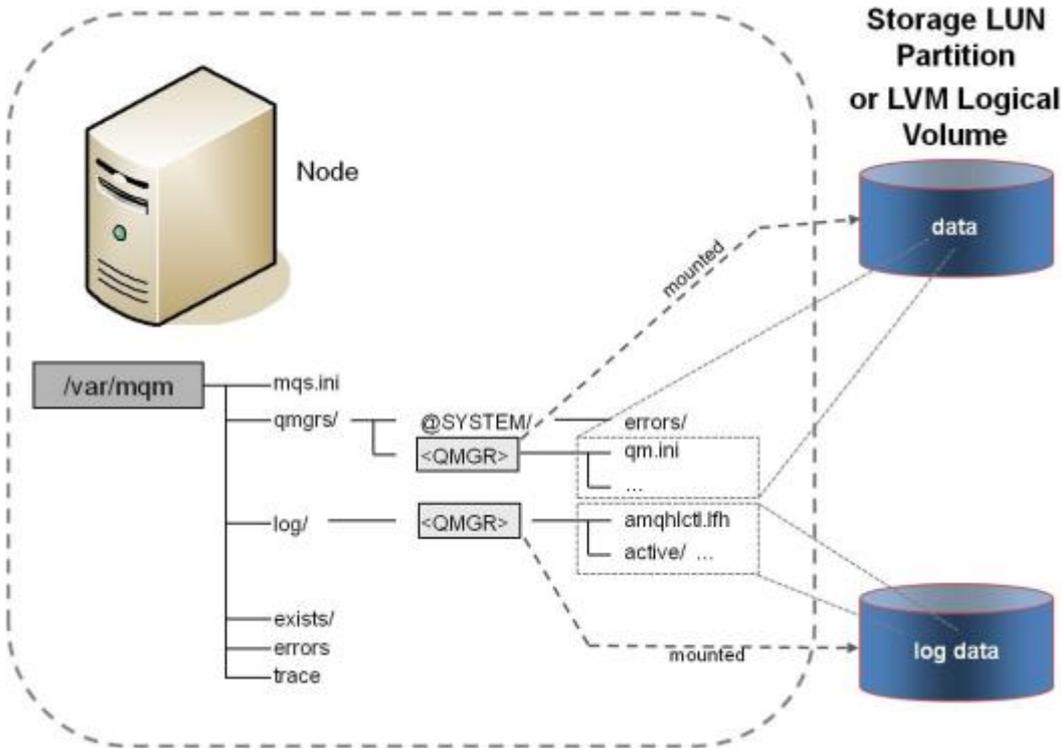


Figure 4 – File System Layout 2 – Direct Mounts

## 6.8.4.1.1.3. Configuration 3 – Symbolic Links

The recommended configuration for Active/Active configurations without LVM and with a large number of queue managers is the use of symbolic links. In this case, one or more dedicated mount points are created (e.g. /mq). A LifeKeeper protected file system is mounted there and subdirectories for each queue manager are created (e.g. /mq/QUEUE!MANAGER!NAME/log and /mq/QUEUE!MANAGER!NAME/qmgrs). The QMDIR and QMLOGDIR directories are then linked to this location.

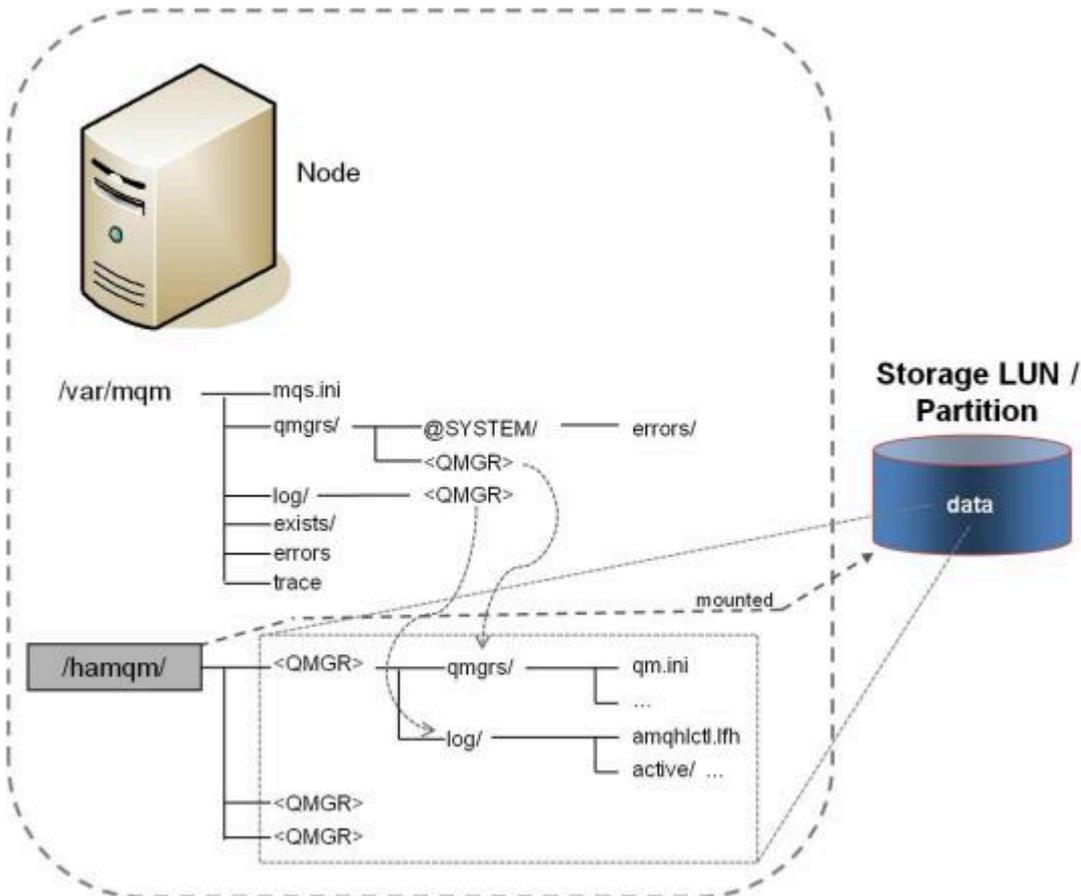


Figure 5 – File System Layout 3 – Symbolic Links

## 6.8.4.1.2. Configuring WebSphere MQ for use with LifeKeeper

---

There are a number of WebSphere MQ configuration considerations that need to be made before attempting to create LifeKeeper for Linux WebSphere MQ resource hierarchies. These changes are required to enable the Recovery Kit to perform **PUT/GET** tests and to make the path to WebSphere MQ persistent data highly available. If the WebSphere MQ queue manager handles remote client requests via TCP/IP, a virtual IP resource must be created prior to creating the WebSphere MQ resource hierarchy. Perform the following actions to enable LifeKeeper WebSphere MQ resource creation:

1. Plan your installation (see [Appendix C](#)).

Before installing WebSphere MQ, you must plan your installation. This includes choosing an **MQUSER**, **MQUSER UID** and **MQGROUP GID**. You must also decide which file system layout you want to use (see [Supported File System Layouts](#)). To ease this process, SIOS Technology Corp. provides a form that contains fields for all required information. See [Appendix C – WebSphere MQ Configuration Sheet](#). Fill out this form to be prepared for the installation process.

2. Configure Kernel Parameters on each server.

WebSphere MQ may require special Linux kernel parameter settings like shared memory. See the [WebSphere MQ documentation](#) for your release of WebSphere MQ for the minimum requirements to run WebSphere MQ. To make kernel parameter changes persistent across reboots, you can use the `/etc/sysctl.conf` configuration file. It may be necessary to add the command `sysctl -p` to your startup scripts (`boot.local`). On SuSE, you can run `insserv boot.sysctl` to enable the automatic setting of the parameters in the `sysctl.conf` file.

3. Create the **MQUSER** and **MQGROUP** on each server.

Use the operating system commands `groupadd` and `adduser` to create the **MQUSER** and **MQGROUP** with the UID and GID from the “WebSphere MQ Configuration Sheet” you used in Step 1.

If the **MQUSER** you have chosen is named `mqm` and has UID 1002 and the **MQGROUP** GID is 1000, you can run the following command on each server of the cluster (change the **MQUSER**, UID and GID values to reflect your settings):

```
groupadd -g 1000 mqm
useradd -m -u 1002 -g mqm mqm
```

**Note:** These settings must be same on all nodes in the cluster. If you are running **NIS** or **LDAP**, create the user and group only once. You may need to create home directories if you have no central home directory server.

4. Unconfigure the **PATH** environment variable (upgrade only).

If this is an upgrade from a prior release of the WebSphere MQ Recover Kit, then the MQUSER PATH environment variable setting may need to be modified. In prior releases of the Recovery Kit the MQUSER PATH environment variable needed to be modified to include the default install location of the WebSphere MQ software, /opt/mqm. If that change was made in a prior release it must be unset for this version of the Recovery Kit to function correctly.

5. Install required packages to install WebSphere MQ on each server.

MQSeries installation requires the installation of X11 libraries and Java for license activation (`mqlicense_lnx.sh`). Install the required software packages.

6. Install WebSphere MQ software and WebSphere MQ fix packs on each server.

Follow the steps described in the "[WebSphere MQ documentation](#)" for your release of WebSphere MQ.

7. **Server Connection Channel Authorization.** Beginning with WebSphere MQ version 7.1 changes were made to channel authorization. By default the MQADMIN user (mqm) is unable to authenticate anonymously (no password) thus failing the resource hierarchy create (authorization for queue managers created with a WebSphere MQ release prior to 7.1 should continue to work). Starting with WebSphere MQ 7.1 one method to allow authorization for the MQADMIN user is to disable channel authorization. For WebSphere MQ 8.0 additional changes are required to the `authinfo` for `system.default.authinfo.idpwos` (in `runmqsc` run 'display `authinfo(system.default.authinfo.idpwos)`' to retrieve the current settings). The `chckclnt` setting of 'reqdamd' must be altered and set to 'optional'. Failure to allow the MQADMIN user anonymous authorization will result in the following error: 'MQCONN ended with reason code 2035' during resource creation. See the [WebSphere MQ documentation](#) for details on how to authorize channels and set access permission.
8. If MQ Version 7.1 or later is being used, enable the MQADMIN user for the specified channel within MQ for the Queue Manager being used.

9. Install LifeKeeper and the WebSphere MQ Recovery Kit on each server.

See the [LifeKeeper Installation Guide](#) for details on how to install LifeKeeper.

10. Prepare the shared storage and mount the shared storage.

See section [Supported File System Layouts](#) for file system layouts supported. Depending on the file system layout and the storage type, this involves creating volume groups, logical volumes, creating file systems or mounting NFS shares.

Here is an example of file system layout 2 with NAS storage:

```
node1:/var/mqm/qmgrs # mkdir TEST\!QM
node1:/var/mqm/qmgrs # mkdir ../log/TEST\!QM
```

```
node1:/var/mqm/qmgrs # mount 192.168.1.30:/raid5/vmware/shared_NFS/
TEST.QM/qmgrs ./TEST\!QM/
```

```
node1:/var/mqm/qmgrs # mount 192.168.1.30:/raid5/vmware/shared_NFS/
TEST.QM/log ../log/TEST\!QM/
```

## 11. Set the owner and group of *QMDIR* and *QMLOGDIR* to *MQUSER* and **MQGROUP**.

The *QMDIR* and *QMLOGDIR* must be owned by *MQUSER* and *MQGROUP*. Use the following commands to set the file system rights accordingly:

```
chown MQUSER QMDIR
chgrp mqm QMDIR
chown MQUSER QMLOGDIR
chgrp mqm QMLOGDIR
```

The values of *MQUSER*, *QMDIR* and *QMLOGDIR* depend on your file system layout and the user name of your **MQUSER**. Use the sheet from Step 1 to determine the correct values for the fields.

Here is an example for *MQUSER mqm* and queue manager *TEST.QM* with default *QMDIR* and *QMLOGDIR* destinations:

```
node1:/var/mqm/qmgrs # chown mqm TEST\!QM/
node1:/var/mqm/qmgrs # chgrp mqm TEST\!QM/
node1:/var/mqm/qmgrs # chown mqm ../log/TEST\!QM/
node1:/var/mqm/qmgrs # chgrp mqm ../log/TEST\!QM/
```

## 12. Create the queue manager on the primary server.

Follow the steps described in the [WebSphere MQ documentation](#) for how to create a queue manager for the version(s) of the WebSphere MQ software being used..

Here is an example for *MQUSER mqm* and queue manager *TEST.QM*.

```
node1:/var/mqm/qmgrs # su - mqm
mqm@node1:~> crtmqm TEST.QM
WebSphere MQ queue manager created.
Creating or replacing default objects for TEST.QM.
Default objects statistics : 31 created. 0 replaced. 0 failed.
Completing setup.
Setup completed.
```

**Note:** If you want to protect an already existing queue manager, use the following steps to move the queue manager data to the shared storage:

- a. Stop the queue manager (`endmqm -i QUEUE.MGR.NAME`).

- b. Copy the content of the queue manager directory and the queue manager log directory to the shared storage created in Step 10.
- c. Change the global configuration file (*mqg.ini*) and queue manager configuration file (*qm.ini*) as required to reflect the new location of the QMDIR and the QMLOGDIR.
- d. Start the queue manager to verify its function (`strmqm QUEUE.MGR.NAME`).
- e. Stop the queue manager (`endmqm -i QUEUE.MGR.NAME`).

13. **Optional:** Configure a virtual IP resource in LifeKeeper on the primary server.

Follow the steps and guidelines described in the [LifeKeeper for Linux IP Recovery Kit Administration Guide](#) and the [LifeKeeper Installation Guide](#).

**Note:** If your queue manager is only accessed by server connects, you do not have to configure the LifeKeeper virtual IP.

14. Modify the listener object to reflect your TCP IP address and port:

```
su - MQUSER
runmqsc QUEUE.MANAGER.NAME

alter LISTENER(SYSTEM.DEFAULT.LISTENER.TCP) TRPTYPE(TCP) PORT(1414)
IPADDR(192.168.1.100)
```

**Note:** Use the same IP address used in the Step 13 to set the value for IPADDR. Do not set IPADDR to have WebSphere MQ bind to all addresses.

15. Start the queue manager on the primary server.

On the primary server, start the queue manager, the command server if it is configured to be started manually and the listener:

```
su - MQUSER
strmqm QUEUE.MANAGER.NAME
strmqcsv QUEUE.MANAGER.NAME
runmqslsr -m QUEUE.MANAGER.NAME -t TCP &
```

16. Verify that the queue manager has been started successfully:

```
su - MQUSER
echo `display qlocal()*` | runmqsc QUEUE.MANAGER.NAME
```

17. Add the queue manager stanza to the global queue manager configuration file *mqg.ini* on the backup server.

**Note:** This step is required for file system layouts 2 and 3.

18. **Optional:** Create the LifeKeeper test queue on the primary server.

```
runmqsc TEST.QM
```

```
5724-B41 © Copyright IBM Corp. 1994, 2002. ALL RIGHTS RESERVED.
```

```
Starting MQSC for queue manager TEST.QM.
```

```
define qlocal(LIFEKEEPER.TESTQUEUE) defpsist(yes) descr('LifeKeeper  
test queue')
```

```
1 : define qlocal(LIFEKEEPER.TESTQUEUE) defpsist(yes)  
    descr('LifeKeeper test queue')
```

```
AMQ8006: WebSphere MQ queue created.
```

19. If you want to have LifeKeeper start the command server, disable the automatic command server startup using the following command on the primary server. Otherwise, the startup of the command server will be performed automatically when the Queue Manager is started:

```
su - MQUSER  
runmqsc TEST.QM  
ALTER QMGR SCMDSERV(MANUAL)
```

20. Create queue manager resource hierarchy on the primary server.

See section [LifeKeeper Configuration Tasks](#) for details.

21. Extend queue manager resource hierarchy to the backup system.

See section [LifeKeeper Configuration Tasks](#) for details.

22. Test your configuration.

To test your HA WebSphere MQ installation, follow the steps described in [Testing a WebSphere MQ Resource Hierarchy](#).

## 6.8.4.1.3. MQ Configuration Changes After Resource Creation

---

The LifeKeeper WebSphere MQ Recovery Kit uses WebSphere MQ commands to start and stop the queue manager. Some exceptions to this rule follow.

---

[Relocating QMDIR and QMLOGDIR](#)

[Changing the Listener Port](#)

[Changing the IP for the Queue Manager](#)

## 6.8.4.1.3.1. Relocating QMDIR and QMLOGDIR

---

If the location of the `QMDIR` and `QMLOGDIR` are changed, the LifeKeeper configuration must be modified. You have the following options to do so:

1. Recreate the queue manager resource hierarchies.

This involves deletion of the queue manager hierarchy and creation of the queue manager hierarchy. See sections [Deleting a WebSphere MQ Hierarchy](#) and [Creating a WebSphere MQ Resource Hierarchy](#) for details.

2. Create the new file system hierarchies manually and add the new file system hierarchies to the WebSphere MQ hierarchy. Remove the old file system hierarchies from the WebSphere MQ hierarchy and remove the old file system hierarchies. See the [LifeKeeper Installation Guide](#) for details on how to create and remove file system hierarchies.

## 6.8.4.1.3.2. Changing the Listener Port

---

To change the listener port of a queue manager, follow these steps:

Alter the listener object in `runmqsc` then stop and start the listener:

```
su - MQUSER
runmqsc QUEUE.MANAGER.NAME

alter LISTENER(SYSTEM.DEFAULT.LISTENER.TCP) TRPTYPE(TCP) PORT(1415)
stop LISTENER(SYSTEM.DEFAULT.LISTENER.TCP)
start LISTENER(SYSTEM.DEFAULT.LISTENER.TCP)
```

See the section [Editing Configuration Resource Properties](#) for details.

## 6.8.4.1.3.3. Changing the IP for the Queue Manager

---

To change the LifeKeeper protected IP associated with the WebSphere MQ queue manager, follow these steps:

1. Create a new LifeKeeper virtual IP in the LifeKeeper GUI.
2. Add the new virtual IP to the WebSphere MQ hierarchy.
3. Remove the old virtual IP from the WebSphere MQ hierarchy.
4. Delete the old virtual IP resource.
5. If needed, modify your listener object in runmqsc and restart the listener:

```
su - MQUSER
runmqsc QUEUE.MANAGER.NAME

alter LISTENER(SYSTEM.DEFAULT.LISTENER.TCP) TRPTYPE(TCP) PORT(1414)
(IPADDR192.168.1.101)
stop LISTENER(SYSTEM.DEFAULT.LISTENER.TCP)
start LISTENER(SYSTEM.DEFAULT.LISTENER.TCP)
```

As an alternative, you can use the LifeKeeper `lk_chg_value` facility to change the IP. See the `lk_chg_value(8)` man page for details.

## 6.8.4.1.4. WebSphere MQ Configuration Examples

---

This section contains definitions and examples of typical WebSphere MQ configurations. Each example includes the configuration file entries that apply to LifeKeeper.

## 6.8.4.1.4.1. Active/Standby Configuration with /var/mqm on Shared Storage

In the Active/Standby configuration, Node1 is the primary LifeKeeper server. It protects the WebSphere MQ queue managers. All storage resides on a shared array between the cluster servers. While Node2 may be handling other applications/services, it acts only as a backup for the WebSphere MQ resources in LifeKeeper's context. The directory `/var/mqm` is located on shared storage. The primary server can run as many queue managers as it can handle.

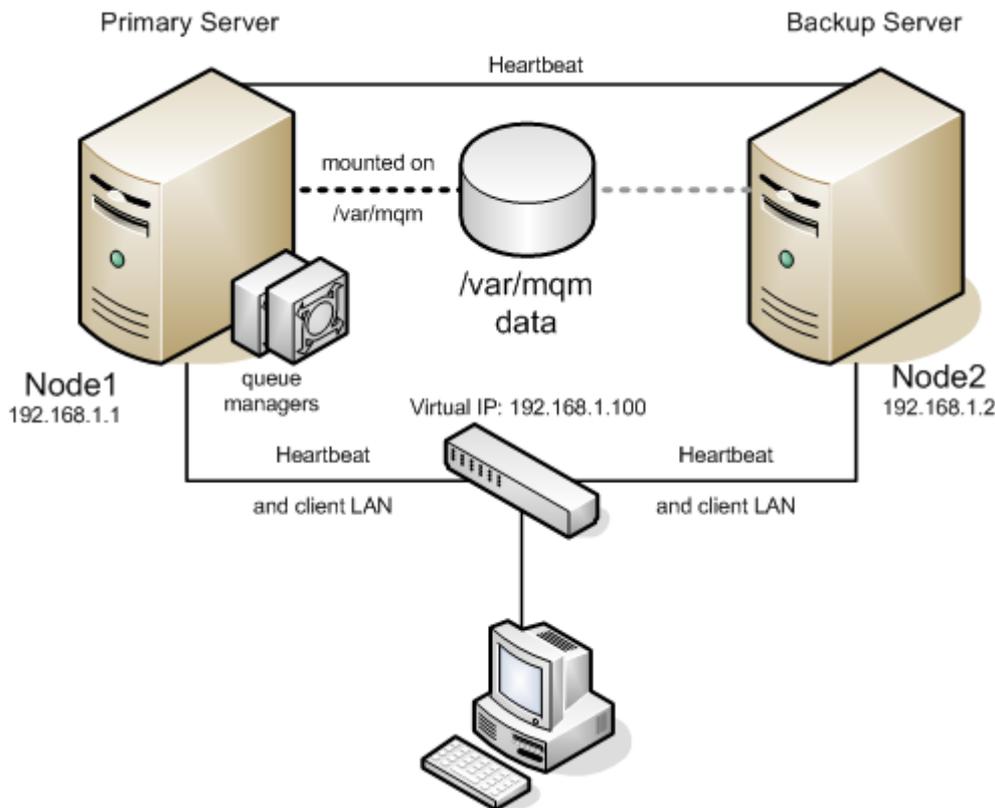


Figure 7 Active/Standby Configuration with Local Storage

### Configuration Notes

- The clients connect to the WebSphere MQ servers using the LifeKeeper protected IP 192.168.1.100 designated to float between the servers in the cluster.
- The directory `/var/mqm` is located on shared storage.
- Each queue manager has modified the listener object to contain a unique port number.

## 6.8.4.1.4.2. Active/Standby Configuration with NAS Storage

In the Active/Standby configuration, Node1 is the primary LifeKeeper server. It protects the WebSphere MQ queue managers. All storage resides on a NAS server with the IP 10.0.0.100. While Node2 may be handling other applications/services, it acts only as a backup for the WebSphere MQ resources in LifeKeeper's context. The directory `/var/mqm` is located from the NAS server's IP 10.0.0.100 and mounted on the active node only. The primary server can run as many queue managers as it can handle.

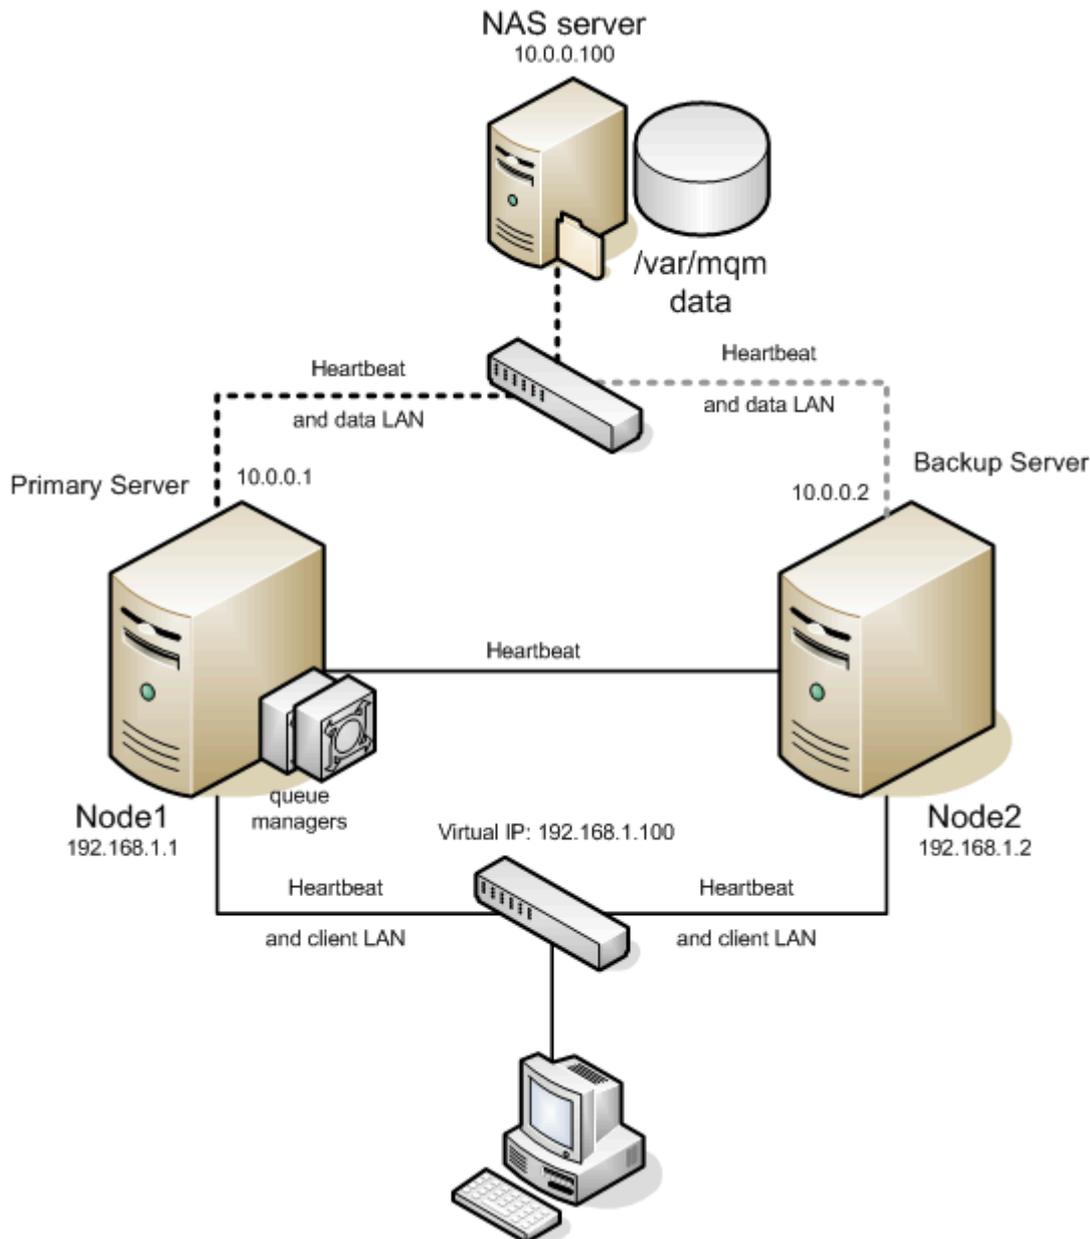


Figure 8 Active/Standby Configuration with NFS Storage

### Configuration Notes

- The clients connect to the WebSphere MQ servers using the LifeKeeper protected IP 192.168.1.100 designated to float between the servers in the cluster.

- The directory `/var/mqm` is located on the NAS server.
- The active server mounts the directory `/var/mqm` from the NAS server with IP 10.0.0.100 using a dedicated network interface.
- There are heartbeats configured on each network interface.
- Each queue manager has modified the listener object to contain a unique port number.

## 6.8.4.1.4.3. Active/Active Configuration with Shared Storage

In the Active/Active configuration below, both Node1 and Node2 are primary LifeKeeper servers for WebSphere MQ resources. Each server is also the backup server for the other. In this example, Node1 protects the shared storage array for queue manager `QMGR1`. Node2 protects the shared storage array for queue manager `QMGR2` as the primary server. Additionally, each server acts as the backup for the other, which in this example means that Node2 is the backup for the queue manager `QMGR1` on Node1, and Node1 is the backup for the queue manager `QMGR2` on Node2.

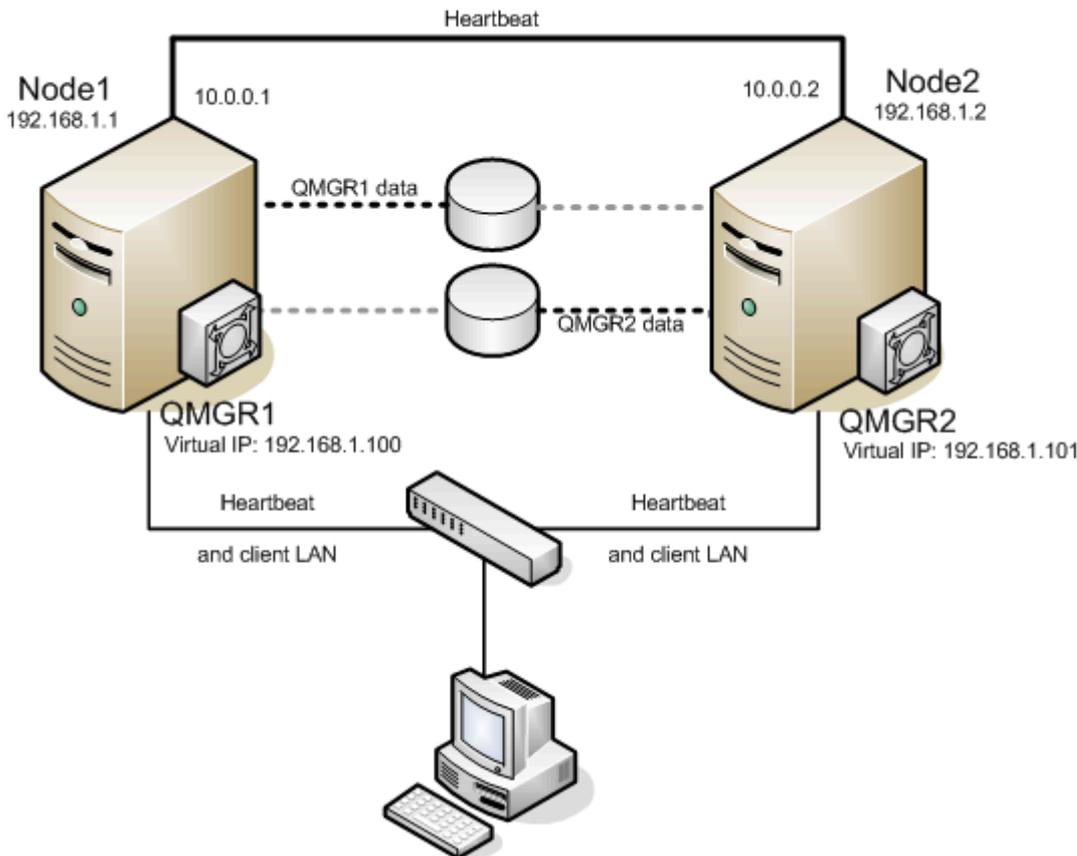


Figure 9 Active/Active Configuration with Shared Storage

### Configuration Notes

- The clients connect to the queue manager `QMGR1` using the LifeKeeper floating IP 192.168.1.100.
- The clients connect to the queue manager `QMGR2` using the LifeKeeper floating IP 192.168.1.101.
- There are heartbeats configured on each network interface.
- Each queue manager has modified the listener object to contain a unique port number.
- `QMGR1` data is located on a volume group on the shared storage with two logical volumes configured. Each logical volume contains a file system that is mounted on `QMDIR` or `QMLOGDIR`.

- **QMGR2 data is located on a secondary volume group on the shared storage with two logical volumes configured. Each logical volume contains a file system that is mounted on QMDIR or QMLOGDIR**

# 6.8.4.1.4.4. Active/Active Configuration with NAS Storage

In the Active/Active configuration below, both Node1 and Node2 are primary LifeKeeper servers for WebSphere MQ resources. Each server is also the backup server for the other. In this example, Node1 protects the NFS mount for queue manager QMGR1. Node2 protects the NFS mount for queue manager QMGR2 as the primary server. Additionally, each server acts as the backup for the other, which in this example means that Node2 is the backup for the queue manager QMGR1 on Node1, and Node1 is the backup for the queue manager QMGR2 on Node2.

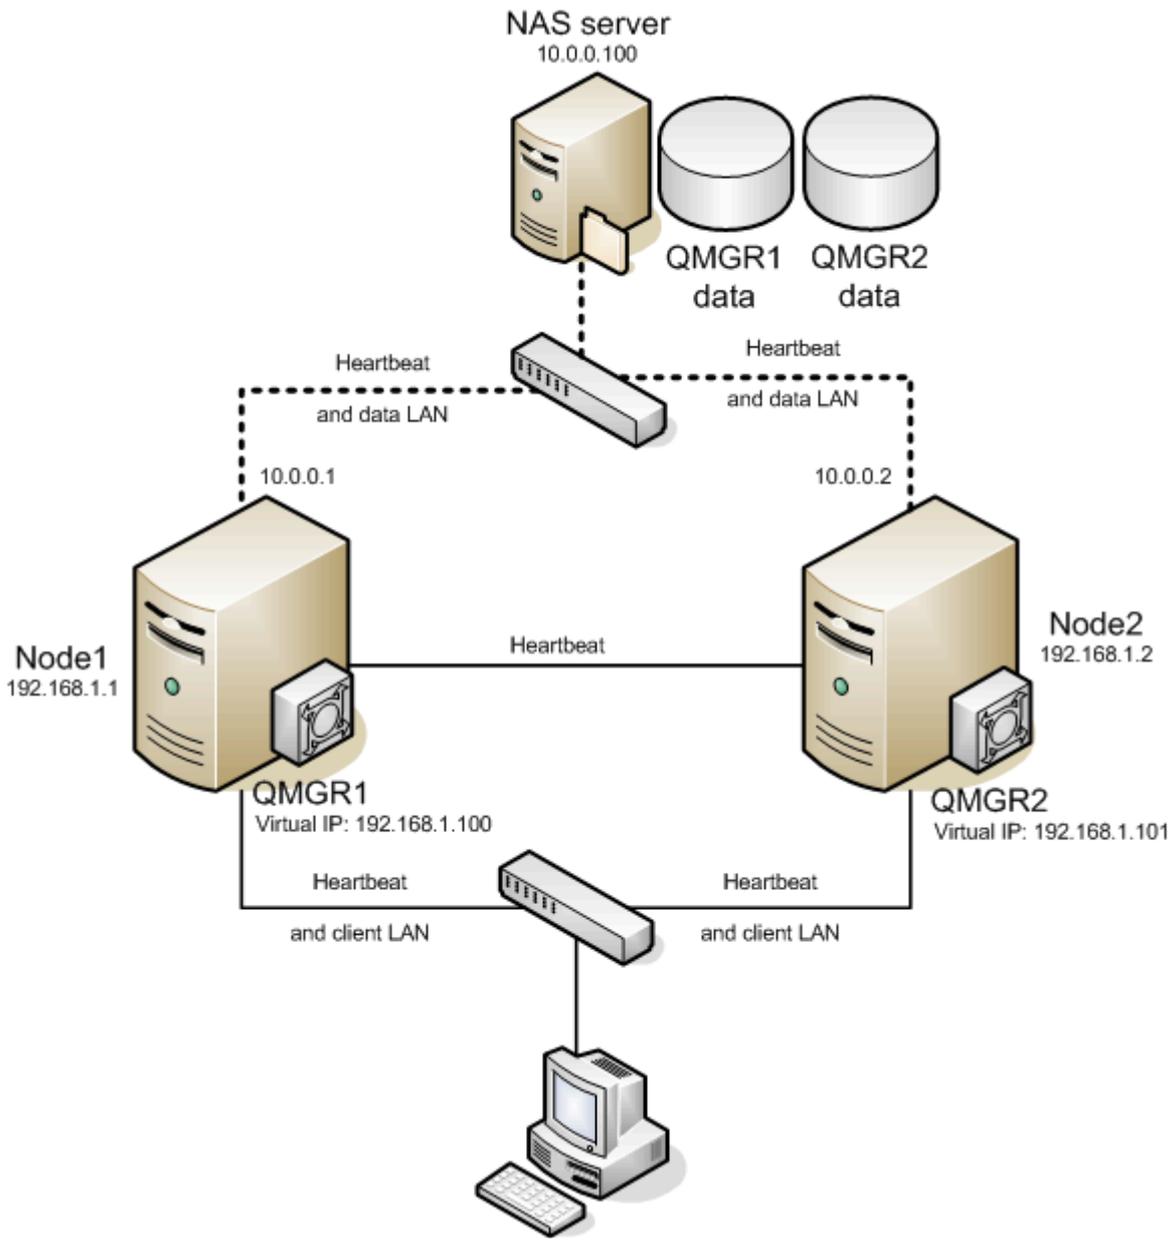


Figure 10 Active/Active Configuration with NFS Storage

## Configuration Notes

- The clients connect to the queue manager QMGR1 using the LifeKeeper floating IP 192.168.1.100.

- The clients connect to the queue manager `QMGR2` using the LifeKeeper floating IP 192.168.1.101.
- Each server has a dedicated network interface to access the NAS server.
- There are heartbeats configured on each network interface.
- Each queue manager has modified the listener object to contain a unique port number.
- `QMGR1` data is located on two NFS exports on the NAS server. The exports are mounted on `QMDIR` or `QMLOGDIR`. The NAS server IP is 10.0.0.100.
- `QMGR2` data is located on two NFS exports on the NAS server. The exports are mounted on `QMDIR` or `QMLOGDIR`. The NAS server IP is 10.0.0.100.

## 6.8.5. LifeKeeper Configuration Tasks for MQ

---

All LifeKeeper for Linux WebSphere MQ Recovery Kit administrative tasks can be performed via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer and monitor WebSphere resources.

### Overview

The following tasks are described in this guide, as they are unique to a WebSphere MQ resource instance and different for each Recovery Kit.

- [Create a Resource Hierarchy](#) – Creates a WebSphere MQ resource hierarchy.
- [Delete a Resource Hierarchy](#) – Deletes a WebSphere MQ resource hierarchy.
- [Extend a Resource Hierarchy](#) – Extends a WebSphere MQ resource hierarchy from the primary server to a backup server.
- [Unextend a Resource Hierarchy](#) – Unextends (removes) a WebSphere MQ resource hierarchy from a single server in the LifeKeeper cluster.
- [Editing Configuration Resource Properties](#) – Reconfigures WebSphere MQ resource parameters including LifeKeeper test queue, listener management and stop timeouts after creation of the WebSphere MQ resource hierarchy.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View](#) / [Edit](#) Properties. View or edit the properties of a resource hierarchy on a specific server.



**Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You can also perform most of these tasks:

- From the toolbar
- By right-clicking on a global resource in the left pane of the status display
- By right-clicking on a resource instance in the right pane of the status display

Using the right-click method allows you to avoid entering information that is required when using the **Edit** menu.

## 6.8.5.1. Creating a WebSphere MQ Resource Hierarchy

After completing the necessary setup tasks, use the following steps to define the WebSphere MQ resource hierarchy.

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From here, select **Create Resource Hierarchy**.

The **Create Resource Wizard** dialog box will appear with a drop-down list box displaying all recognized Recovery Kits installed within the cluster.

2. Select **IBM WebSphereMQ** and click **Next**.
3. You will be prompted to enter the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter an error requiring you to correct previously entered information. You may click **Cancel** at any time to cancel the entire creation process.

Field	Tips
Switchback Type	Choose either <b>Intelligent</b> or <b>Automatic</b> . This dictates how the WebSphere MQ instance will be switched back to this server when the server comes back up after a failover. The switchback type can be changed later from the <b>General</b> tab of the <b>Resource Properties</b> dialog box.  <b>Note:</b> The switchback strategy should match that of the IP or File System resource to be used by the WebSphere MQ resource. If they do not match the WebSphere MQ resource, creation will attempt to reset them to match the setting selected for the WebSphere MQ resource.
Server	Select the <b>Server</b> on which you want to create the hierarchy.
Queue Manager Name	Select the WebSphere MQ queue manager you want to protect. The queue manager must be created prior to creating the resource hierarchy. Queue managers already under LifeKeeper protection are excluded from this list. The queue managers are taken from the global <code>mqs.ini</code> configuration file.
Manage Listener	Select <b>"YES"</b> to protect and manage the WebSphere MQ queue manager listener. Select <b>"NO"</b> if LifeKeeper should not manage the WebSphere MQ listener.  <b>Note:</b> You can change this setting later. See <a href="#">Editing Configuration Resource Properties</a> for details.
Server Connection Channel	Select the server connection channel to use for connection tests. By default, the channel <code>SYSTEM.DEF.SVRCONN</code> will be used; however, beginning with MQ Version 7.1, changes in MQ's Channel Authentication require that a channel other than the default be used and that the <code>MQADMIN</code> user be enabled for the specified channel.

	<p><b>Note:</b> Make sure the Server Connection Channel has been created PRIOR to creating your resource. For more information, see <a href="#">Configuring WebSphere MQ for Use with LifeKeeper</a>.</p> <p><b>Note:</b> This setting can be changed later. See <a href="#">Editing Configuration Resource Properties</a> for details.</p>
Virtual IP	<p>Select the LifeKeeper virtual IP resource to include in the hierarchy. Select “None” if you do not want to include a LifeKeeper virtual IP in the WebSphere MQ hierarchy.</p> <p><b>Note:</b> The virtual IP must be ISP (active) on the primary node to appear in the selection list.</p>
IBM WebSphere MQ Resource Tag	<p>Either select the default root tag offered by LifeKeeper, or enter a unique name for the resource instance on this server. The default is the queue manager name. Letters, numbers and the following special characters may be used: – _ . /</p>

4. Click **Create**. The **Create Resource Wizard** will then create your WebSphere MQ resource hierarchy. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.
5. An information box will appear indicating that you have successfully created a WebSphere MQ resource hierarchy and that hierarchy must be extended to another server in your cluster in order to achieve failover protection. Click **Next**.
6. Click **Continue**. LifeKeeper will then launch the **Pre-Extend Wizard**. Refer to Step 2 under [Extending a WebSphere MQ Hierarchy](#) for details on how to extend your resource hierarchy to another server.

## 6.8.5.2. Extending a WebSphere MQ Hierarchy

This operation can be started from the **Edit** menu or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should refer to Step 2 below.

1. On the **Edit** menu, select **Resource**, then **Extend Resource Hierarchy**. The **Pre-Extend Wizard** appears. If you are unfamiliar with the **Extend** operation, click **Next**. If you are familiar with the **LifeKeeper Extend Resource Hierarchy** defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information.

**Note:** The first two fields appear only if you initiated the **Extend** from the **Edit** menu.

Field	Tips
Template Server	Enter the server where your WebSphere MQ resource is currently in service.
Tag to Extend	Select the WebSphere MQ resource you wish to extend.
Target Server	Enter or select the server you are extending to.
Switchback Type	Select either <b>Intelligent</b> or <b>Automatic</b> . The switchback type can be changed later, if desired, from the <b>General</b> tab of the <b>Resource Properties</b> dialog box. <b>Note:</b> Remember that the switchback strategy must match that of the dependent resources to be used by the WebSphere MQ resource.
Template Priority	Select or enter a priority for the template hierarchy. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (the number 1 indicates the highest priority). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. <b>Note:</b> This selection will appear only for the initial extend of the hierarchy.
Target Priority	Either select or enter the priority of the hierarchy for the target server.
Queue Manager Name	This informational field shows the queue manager name you are about to extend. You cannot change this value.
Root Tag	LifeKeeper will provide a default tag name for the new WebSphere MQ resource instance on the target server. The default tag name is the same as the tag name

	for this resource on the template server. If you enter a new name, be sure it is unique on the target server. Letters, numbers and the following special characters may be used: – _ . /
--	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

 **Note:** All configurable queue manager parameters like listener management, the name of the LifeKeeper test queue and the shutdown timeout values are taken from the template server.

3. After receiving the message that the pre-extend checks were successful, click **Next**.
4. Depending upon the hierarchy being extended, LifeKeeper will display a series of information boxes showing the Resource Tags to be extended which cannot be edited. Click **Extend**.
5. After receiving the message “**Hierarchy extend operations completed**”, click **Next Server** to extend the hierarchy to another server or click Finish if there are no other extend operations to perform.
6. After receiving the message “**Hierarchy Verification Finished**”, click **Done**.

## 6.8.5.3. Unextending a WebSphere MQ Hierarchy

---

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. On the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the WebSphere MQ resource. It cannot be the server where the WebSphere MQ resource is currently in service. (This dialog box will not appear if you selected the **Unextend** task by right-clicking on a resource instance in the right pane.) Click **Next**.
3. Select the WebSphere MQ hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the **Unextend** task by right-clicking on a resource instance in either pane.)
4. An information box appears confirming the target server and the WebSphere MQ resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the WebSphere MQ resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

## 6.8.5.4. Deleting a WebSphere MQ Hierarchy

It is important to understand what happens to dependencies and protected services when a WebSphere hierarchy is deleted.

- **Dependencies:** Before removing a resource hierarchy, you may wish to remove the dependencies. Dependent file systems will be removed. Dependent non-file system resources like IP or Generic Application will not be removed as long as the delete is done via the LifeKeeper GUI or the WebSphere MQ delete script. For LifeKeeper to not delete the dependent file systems of the WebSphere MQ queue manager, manually remove the dependencies prior to deleting the WebSphere MQ hierarchy.
- **Protected Services:** If the WebSphere resource hierarchy is taken out of service before being deleted, the WebSphere daemons for this queue manager will be stopped. If a hierarchy is deleted while it is in service, the WebSphere MQ daemons will continue running and offering services (without LifeKeeper protection) after the hierarchy is deleted.

To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the **Target Server** where you will be deleting your WebSphere MQ resource hierarchy and click **Next**. (This dialog will not appear if you selected the Delete Resource task by right-clicking on a resource instance in either pane.)
3. Select the **Hierarchy to Delete**. (This dialog will not appear if you selected the Delete Resource task by right-clicking on a resource instance in the left or right pane.) Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete**.
5. Another information box appears confirming that the WebSphere resource was deleted successfully.
6. Click **Done** to exit.

## 6.8.5.5. Testing a WebSphere MQ Resource Hierarchy

---

You can test your WebSphere MQ resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

On the **Edit** menu, select **Resource**, then **In Service**. For example, an In Service request executed on a backup server causes the application hierarchy to be taken out of service on the primary server and placed in service on the backup server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the **Out of Service** request, the application is taken out of service without bringing it in service on the other server.

### Testing Shared Storage Configuration

To test WebSphere MQ shared storage operations, perform the following steps:

1. Create a temporary test queue on the primary server with the default persistency of **“yes”**

```
mqm@node1:/opt/mqm/samp/bin> runmqsc TEST.QM
5724-B41 © Copyright IBM Corp. 1994, 2002. ALL RIGHTS RESERVED.
Starting MQSC for queue manager TEST.QM.
```

```
define qlocal(TEST) defpsist(yes)
  1 : define qlocal(TEST) defpsist(yes)
AMQ8006: WebSphere MQ queue created.
end
```

```
  2 : end
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
```

2. Put a message into the test queue created on the primary node:

```
mqm@node1:/opt/mqm/samp/bin> echo "HELLO WORLD on NODE1" | ./amqsput TEST
TEST.QM
Sample AMQSPUT0 start
target queue is TEST
Sample AMQSPUT0 end
```

3. Browse the test queue to see if the message has been stored:

```
mqm@node1:/opt/mqm/samp/bin> ./amqsbcbg TEST TEST.QM
```

You should see a message with the content “HELLO WORLD on NODE1” and some additional output. Look for the following line and verify that the persistency is 1:

```
[...]
    Priority : 0 Persistence : 1
[...]
```

4. Switch the resource hierarchy to the standby node.
5. On the standby server where the queue manager is now active, repeat Step 3. The message should be accessible on the standby server. If not, check your storage configuration.
6. On the standby server where the queue manager is now active, get the message from the test queue:

```
mqm@node1:/opt/mqm/samp/bin> ./amqsget TEST TEST.QM
Sample AMQSGETO start
message <HELLO WORLD on NODE1>
<now wait 15 seconds>
no more messages
Sample AMQSGETO end
```

7. Delete the test queue created in Step 1.

```
mqm@node1:/opt/mqm/samp/bin> runmqsc TEST.QM
5724-B41 © Copyright IBM Corp. 1994, 2002. ALL RIGHTS RESERVED.
Starting MQSC for queue manager TEST.QM.

delete qlocal(TEST)
    1 : delete qlocal(TEST)
MQ8007: WebSphere MQ queue deleted.
end
    2 : end
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
```

## 6.8.5.5.1. Testing MQ Client Connectivity

### Testing Client Connectivity

To test client connectivity, perform the following steps:

1. On the primary server, use the `amqsbcgc` command to connect to the queue manager:

```
export MQSERVER='SYSTEM.DEF.SVRCONN/TCP/192.168.1.90(1414)'
```

**Note:** Replace the IP 192.168.1.90 with the LifeKeeper protected virtual IP of the queue manager. If your queue manager uses a different port other than 1414, then replace the port number 1414 with the one being used. If the server connection channel being used is not the default `SYSTEM.DEF.SVRCONN` channel, then replace the server connection channel `SYSTEM.DEF.SVRCONN` with the one being used.

You should see the following output:

```
mqm@node1:/opt/mqm/samp/bin> ./amqsbcgc LIFEKEEPER.TESTQUEUE TEST.QM

AMQSBCG0 - starts here
*****

MQOPEN - 'LIFEKEEPER.TESTQUEUE'

No more messages
MQCLOSE
```

If you get a message like the following, then the test queue `LIFEKEEPER.TESTQUEUE` is not configured. Create the test queue as described in section [Configuring WebSphere MQ for Use with LifeKeeper](#) and repeat the test.

```
AMQSBCG0 - starts here
*****

MQOPEN - 'LIFEKEEPER.TESTQUEUE'
MQOPEN failed with CompCode:2, Reason:2085
```

2. Perform a switchover of the resource hierarchy.
3. Repeat Step 1 on the same server as before which is now the standby server after the switchover.

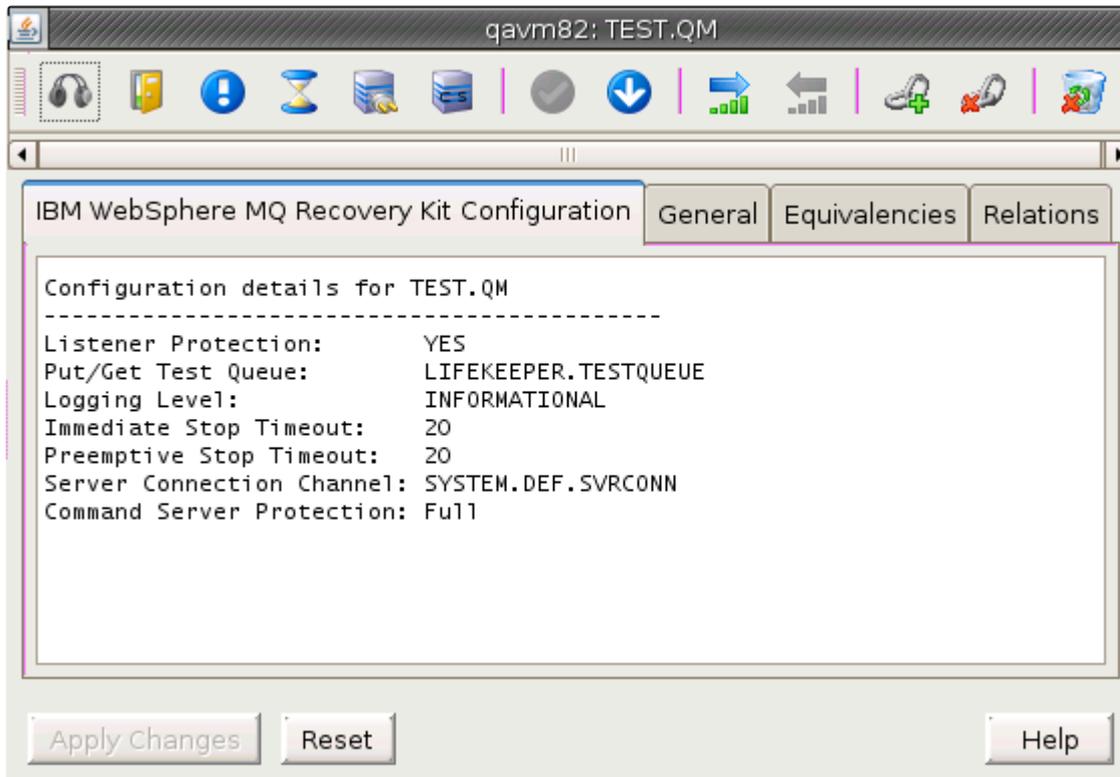
### Testing If PUT/GET Tests are Performed

To test if the WebSphere MQ Recovery kit performs all checks including the `PUT/GET` test, perform the following:

1. Make sure the queue manager is in service (ISP) on any server.
2. Increase the logging level of the queue manager as described in [Changing the Log Level](#) to **“FINE”**.
3. Open the log dialog on the machine where the queue manager is active and wait for the next check to happen (max. two minutes).
4. Analyze the log and verify that all checks are performed and none of the tests is skipped. The `PUT/GET` could be skipped for the following reasons:
  - a. No LifeKeeper test queue is configured (in this case, configure the test queue as described in [Changing the LifeKeeper Test Queue Name](#)).
  - b. LifeKeeper test queue does not exist (in this case, create the test queue as described in [Configuring WebSphere MQ for Use with LifeKeeper](#)).
  - c. The modified `amqsget©` executables are not available (in this case, install a C compiler and rerun the script `/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/compilesamples`).
5. Set the debug level to **“INFORMATIONAL”** again.

## 6.8.5.6. Viewing MQ Resource Properties

To view the IBM WebSphere MQ resource properties, right-click on the icon for the resource/server combination for which you want to view the properties. When the resource context menu appears, click **Properties**. The following dialog will appear.



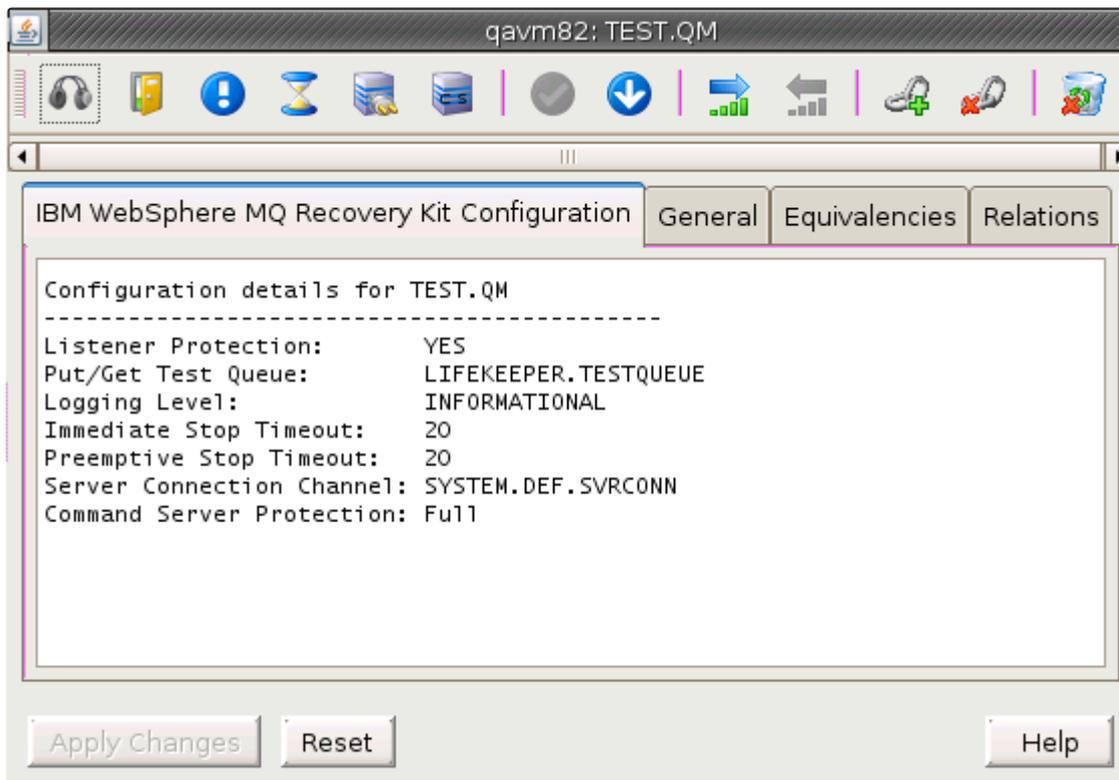
Resource properties will be displayed in the properties panel if it is enabled. You can also right-click on the icon for the global resource for which you want to view the properties. When the **Resource Context Menu** appears, click **Properties**. When the dialog comes up, select the server for which you want to view that resource from the Server list.

## 6.8.5.7. Editing MQ Configuration Resource Properties

The WebSphere MQ Properties page allows you to view and modify the configuration details for a specific WebSphere MQ resource via the properties panel if it is enabled. Specific WebSphere MQ resource configuration properties can also be modified via the **Resource Context Menu**.

To edit configuration details via the WebSphere MQ Configuration Properties page from the LifeKeeper GUI Properties Panel, you must first ensure the GUI Properties Panel is enabled. To enable the GUI Properties Panel, select **View**, then **Properties Panel** (must have a check mark to indicate it is enabled). Once enabled, left-click on the **WebSphere MQ** resource to display its configuration details in the LifeKeeper GUI Properties Panel.

Below is an example of the properties page that will appear in the LifeKeeper GUI Properties Panel for a WebSphere MQ resource.



The properties page contains four tabs. The first tab, labeled **IBM WebSphere MQ Recovery Kit Configuration**, contains configuration information that is specific to WebSphere MQ resources and allows modification via the resource specific icons. The remaining three tabs are available for all LifeKeeper resource types and their content is described in the topic [Resource Properties Dialog](#) in the [LifeKeeper for Linux Technical Documentation](#).

The following table displays the WebSphere MQ resource specific icons and the configuration component that can be modified when clicking on the icon.

	<b>Listener Protection Configuration</b>	Allows you to specify whether protection of the IBM WebSphere MQ listener is included with the other IBM WebSphere MQ queue manager components being protected.
	<b>PUT/GET Test Queue Configuration</b>	Allows you to change the name of the queue that the IBM WebSphere MQ Recovery Kit will use to perform PUT/GET tests for the queue manager being protected.
	<b>Logging Level Configuration</b>	Allows you to modify the log level that the IBM WebSphere MQ Recovery Kit will use for the queue manager being protected.
	<b>Shutdown Timeout Configuration</b>	Allows you to modify the timeout in seconds for the immediate shutdown and preemptive shutdown timers for the IBM WebSphere MQ queue manager being protected.
	<b>Server Connection Channel Configuration</b>	Allows you to modify the server connection channel that is used for client connection and the PUT/GET testing for the IBM WebSphere MQ queue manager being protected.
	<b>Command Server Protection Configuration</b>	Allows you to specify the protection/recovery level for command server component of the IBM WebSphere MQ queue manager being protected.

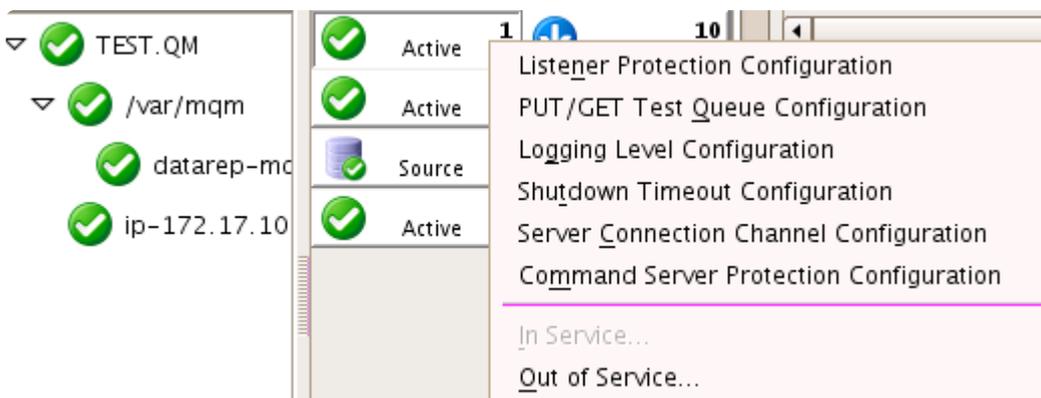
More details on each of these configuration options can be found below.

<b>Listener Management</b>	Specifies whether you want LifeKeeper to protect the listener for the queue manager or not. If listener management is disabled (value of NO), LifeKeeper will not monitor the listener and you can stop the listener without causing LifeKeeper recovery actions. If listener management is enabled (value of YES), LifeKeeper will monitor the listener and restart the listener if the listener is not running. If the recovery fails, a failover of the WebSphere MQ hierarchy to the backup server is initiated.
<b>LifeKeeper Test Queue</b>	LifeKeeper performs PUT/GET test to monitor queue manager operations. The WebSphere MQ Recovery Kit uses a <b>dedicated</b> test queue to put messages in and retrieve messages again. In case a failure is detected, no recovery or failover is performed. Instead, the Recovery Kit sends an event that you can register to receive. The events are called <code>putgetfail</code> and <code>putgetcfail</code> . You can add a notification script to the directories <code>/opt/LifeKeeper/events/mqseries/putgetfail</code> and <code>/opt/LifeKeeper/events/mqseries/putgetcfail</code> to react to those events.  <b>Note 1:</b> If the LifeKeeper test queue is not configured in the queue manager, the PUT/GET test is skipped. No recovery or failover takes place.  <b>Note 2:</b> If the listener is protected, a second client connect check will be done. If this check fails, a recovery or failover of the queue manager is attempted.
<b>Logging Level</b>	You can set the logging level of the WebSphere MQ Recovery Kit to four presets:

	<ul style="list-style-type: none"> <li>• ERROR  In this log level, only errors are logged. No informational messages are logged.</li> <li>• INFORMATIONAL (default)  In this log level, LifeKeeper informational messages about start, stop and recovery of resources are logged.</li> <li>• DEBUG  In this log level, the informational LifeKeeper messages and the command outputs from all WebSphere MQ commands in the restore, remove and recovery scripts are logged.</li> <li>• FINE  In this log level, all command outputs from WebSphere MQ commands issued in start, stop, recovery and quickCheck scripts are logged. Additional debug messages are also logged.</li> </ul> <p>It is recommended to set this debug level only for debugging purpose. As <code>quickCheck</code> actions are also logged, this fills up the log files each time a <code>quickCheck</code> for the WebSphere MQ queue manager runs.</p> <p>The default is <i>INFORMATIONAL</i>. This is equivalent to normal LifeKeeper logging of other recovery kits.</p> <p><b>Note:</b> Independent of the logging level setting, WebSphere MQ errors during start, stop, recovery or during the check routine are always logged with the complete command output of the last command run.</p>
<p><b>Stop Timeout Values</b></p>	<p>The WebSphere MQ Recovery Kit stops the queue manager in 3 steps:</p> <ol style="list-style-type: none"> <li>1. immediate stop</li> <li>2. preemptive stop</li> <li>3. kill all queue manager processes</li> </ol> <p>The timeout values specified determine the time the Recovery Kit waits in Steps 1 and 2 for a successful completion. If this timeout is reached, the next step in the shutdown process is issued. The default for the immediate and preemptive shutdown timeouts is <b>20 seconds</b>.</p>

<p><b>Server Connection Channel</b></p>	<p>The WebSphere MQ Recovery Kit allows the specification of the server connection channel. By default, the kit will use the channel <code>SYSTEM.DEF.SVRCONN</code>, but an alternate channel can be specified during resource creation or at any time after resource creation.</p>
<p><b>Command Server</b></p>	<p>The WebSphere MQ Recovery Kit allows two levels of protection and recovery for the command server component for the protected queue manager. The levels are <b>Full</b> and <b>Minimal</b>.</p> <p>With <b>Full</b> protection, the command server will be started, stopped, monitored and recovered or failed over if recovery is unsuccessful. The recovery steps with <b>Full</b> protection are:</p> <ul style="list-style-type: none"> <li>• Attempt to restart just the command server process.</li> <li>• If that fails, attempt a full restart of the queue manager including the command server process.</li> <li>• If both attempts are unsuccessful at restarting the command server, then initiate a failover to the standby node.</li> </ul> <p>With <b>Minimal</b> protection, the command server will only be started during restore or stopped during remove. No monitoring or recovery of the command server will be performed.</p> <p><b>NOTE:</b> Starting the command server will only be performed by the Recovery Kit during restore if the queue manager <code>SCMDSERV</code> parameter is set for manual startup. During a recovery, a failed command server restart will always be attempted regardless of the <code>SCMDSERV</code> setting unless the Command Server Protection Level is set to <b>Minimal</b>.</p>

As previously noted, these WebSphere MQ resource configuration components can be modified using the resource specific icons in the properties panel or via the **Resource Context Menu**.



The parameters above can be set for each queue manager separately either via the LifeKeeper GUI or via a command line utility.

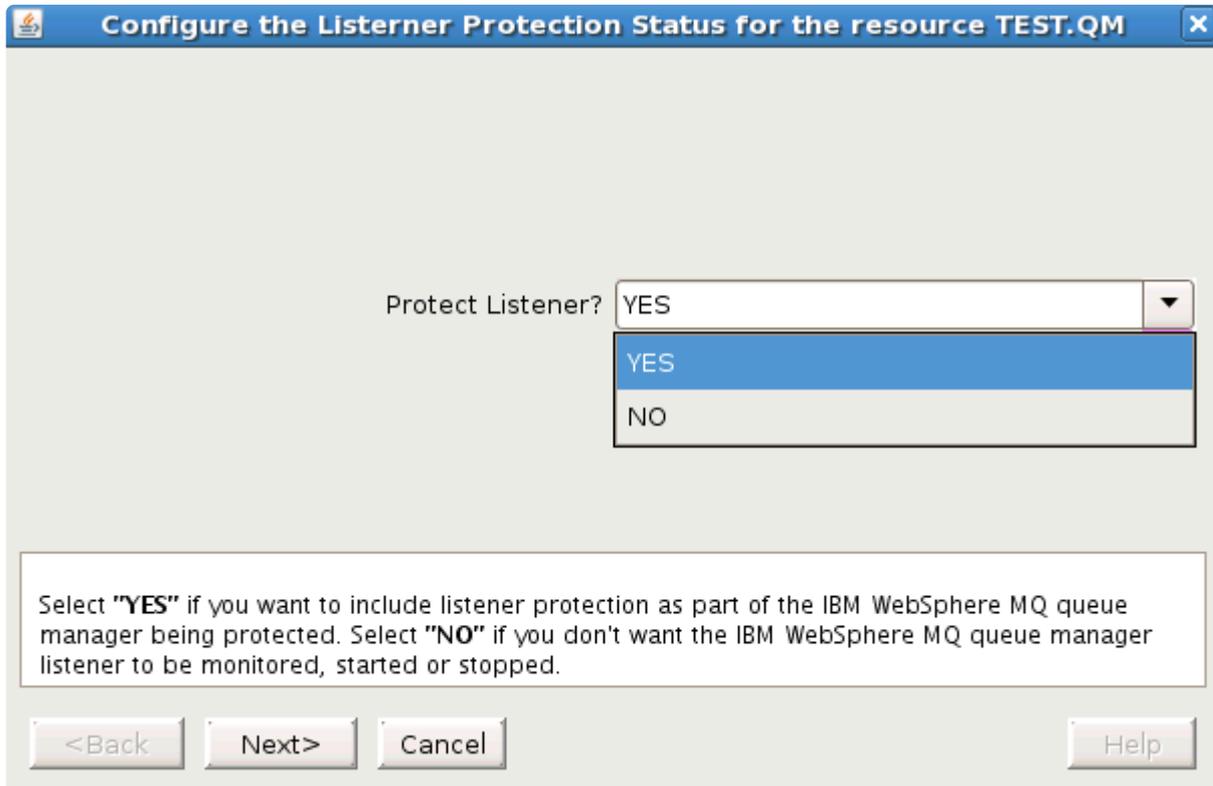
To set the parameters via the command line, use the script:

`$LKROOT/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam`

## 6.8.5.7.1. Enable/Disable Listener Protection

### GUI

First navigate to the **WebSphere MQ Resource Properties Panel** or the **Resource Context Menu** described above. The resource must be in service to modify the Listener Protection value. Then click on **Listener Protection Configuration** icon or menu item. The following dialog will appear:



Now select **YES** if you want LifeKeeper to start, stop and monitor the WebSphere MQ listener. Select **NO** if LifeKeeper should not start, stop and monitor the WebSphere MQ listener. Click **Next**. You will be asked if you want to enable or disable listener protection; click **Continue**. If you have chosen to enable listener management, the LifeKeeper GUI checks if the listener is already running. If it is not already running, it will try to start the listener. If the listener start was successful, the LifeKeeper GUI will enable listener management on each server in the cluster. If the listener is not running and could not be started, the LifeKeeper GUI will not enable listener management on the servers in the cluster.

### Command Line

To set the LifeKeeper listener management via command line, use the following command:

```
/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam -c -s
-i TEST.QM -p LISTENERPROTECTION -v YES
```

This will set (-s) the LifeKeeper listener management (-p) on each node of the cluster (-c) to **YES** (-v) (enable listener management) for queue manager TEST.QM (-i).

**Note:** You can either use the queue manager name (-i) or the LifeKeeper TAG (-t) name.

## 6.8.5.7.2. Changing the LifeKeeper Test Queue Name

### GUI

First navigate to the **WebSphere MQ Resource Properties Panel** or the Resource Context Menu described above. Then click on **PUT/GET TESTQUEUE Configuration** icon or menu item. The following dialog will appear:

LifeKeeper Test Queue:

Enter the name of the queue that the IBM WebSphere MQ Recovery Kit will use to perform PUT/GET tests for the queue manager being protected. Enter an empty queue name if you want to disable LifeKeeper PUT/GET tests.

<Back    Next>    Cancel    Help

Now enter the name of the LifeKeeper test queue and click **Next**. You will be asked if you want to set the new LifeKeeper test queue; click **Continue**. Next, the LifeKeeper GUI will set the LifeKeeper test queue on each server in the cluster. If you set the test queue to an empty value, no **PUT/GET** tests are performed.

### Command Line

To set the LifeKeeper test queue via command line, use the following command:

```
/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam -c -s
-i TEST.QM -p TESTQUEUE -v "LIFEKEEPER.TESTQUEUE"
```

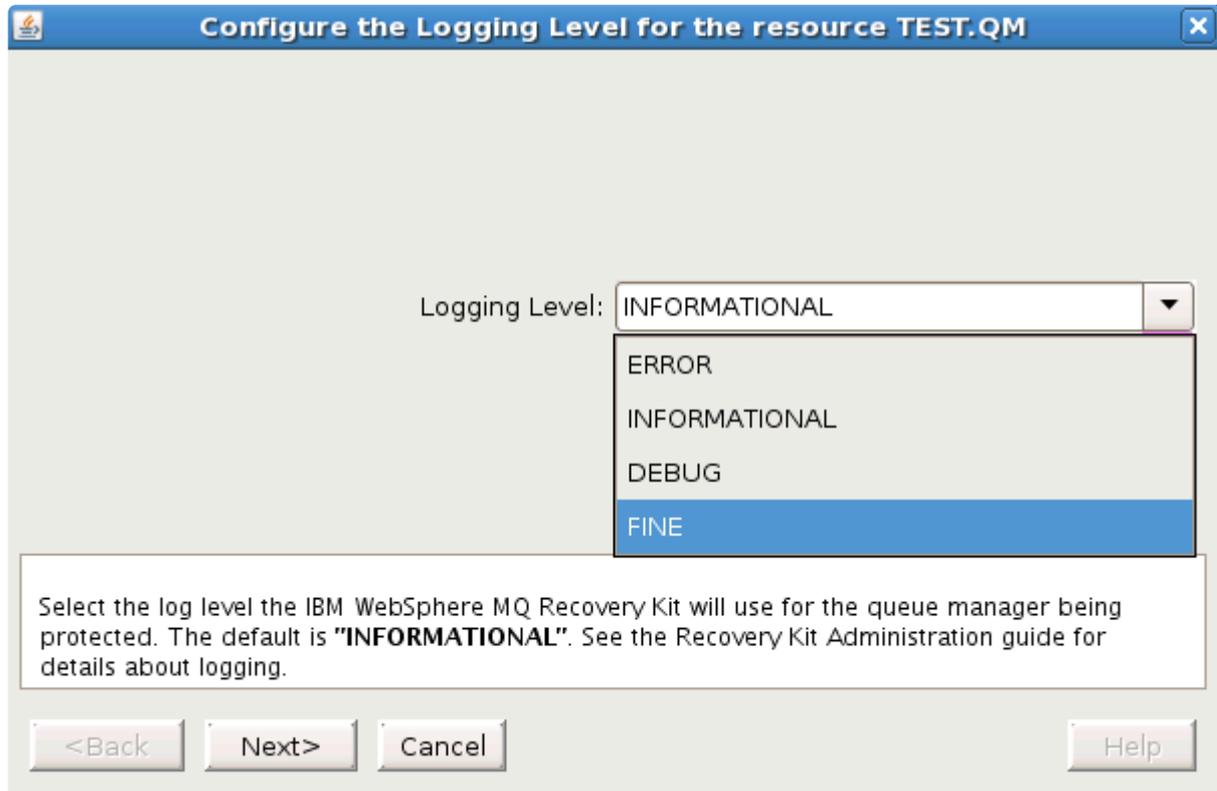
This will set (-s) the LifeKeeper test queue (-p) on each node of the cluster (-c) to LIFEKEEPER.TESTQUEUE (-v) for queue manager TEST.QM (-i).

**Note:** You can either use the queue manager name (-i) or the LifeKeeper TAG (-t) name.

## 6.8.5.7.3. Changing the Log Level

### GUI

First navigate to the **WebSphere MQ Resource Properties Panel** or the **Resource Context Menu** described above. Then click on **Logging Level Configuration** icon or menu item. The following dialog will appear:



Now select the **Logging Level** and click **Next**. You will be asked if you want to set the new LifeKeeper logging level; click **Continue**. Next, the LifeKeeper GUI will set the LifeKeeper logging level for the selected queue manager on each server in the cluster.

### Command Line

To set the LifeKeeper logging level via command line, use the following command:

```
/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam -c -s
-i TEST.QM -p DEBUG -v DEBUG
```

This will set (-s) the LifeKeeper logging level (-p) on each node of the cluster (-c) to `DEBUG` (-v) for queue manager `TEST.QM` (-i).

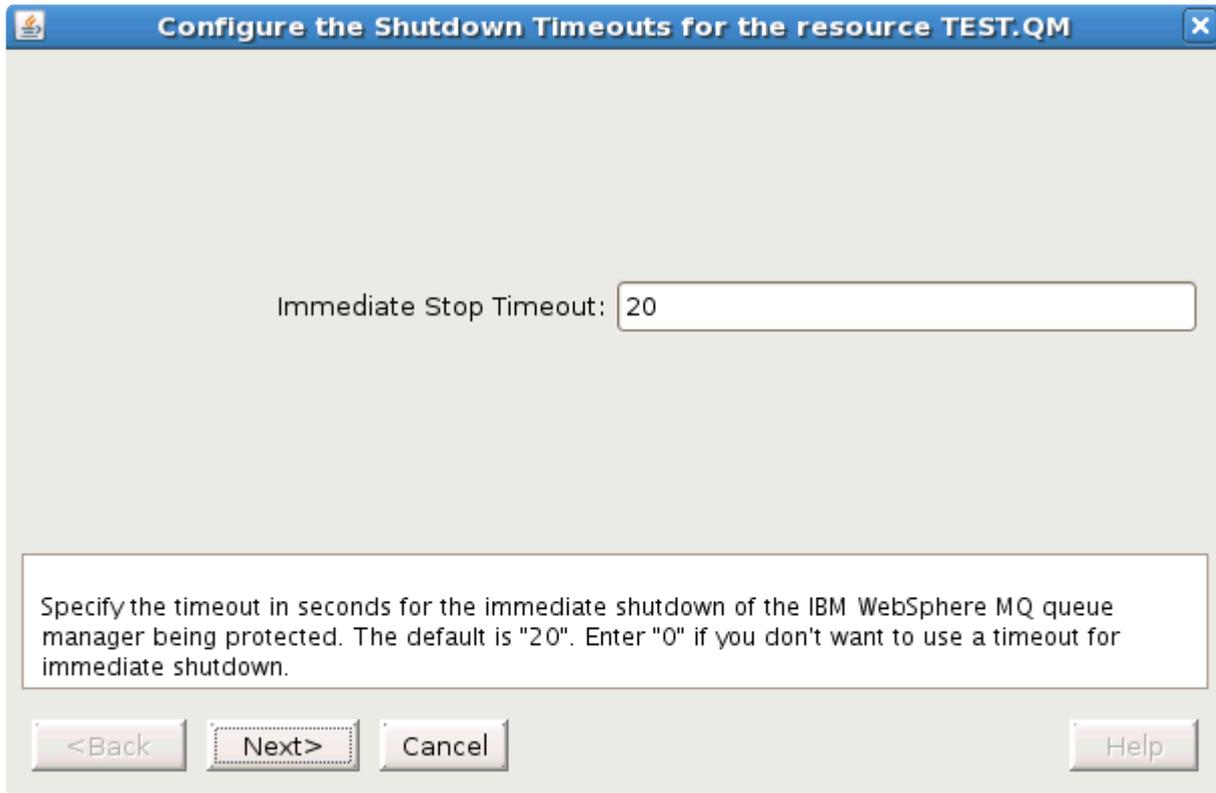
**Note:** You can either use the queue manager name (-i) or the LifeKeeper TAG (-t) name.

## 6.8.5.7.4. Changing Shutdown Timeout Values

---

### GUI

First, navigate to the WebSphere MQ resource properties panel or the resource context menu described above. Then click on **Shutdown Timeout Configuration** icon or menu item. The following dialog will appear:



Configure the Shutdown Timeouts for the resource TEST.QM

Immediate Stop Timeout:

Specify the timeout in seconds for the immediate shutdown of the IBM WebSphere MQ queue manager being protected. The default is "20". Enter "0" if you don't want to use a timeout for immediate shutdown.

<Back   Next>   Cancel   Help

Now enter the immediate shutdown timeout value in seconds and click **Next**. If you want to disable the immediate shutdown timeout, enter **0**. Now the following dialog will appear:

Configure the Shutdown Timeouts for the resource TEST.QM

Preemptive Stop Timeout:

Specify the timeout in seconds for the preemptive shutdown of the IBM WebSphere MQ queue manager being protected. The default is "20". Enter "0" if you don't want to use a timeout for preemptive shutdown.

<Back   Next>   Cancel   Help

Now enter the preemptive shutdown timeout value in seconds and click **Next**. If you want to disable the preemptive shutdown timeout enter 0. You will be asked if you want to set the new LifeKeeper timeout parameters, click **Continue**. Next, the LifeKeeper GUI will set the LifeKeeper immediate and preemptive timeout values on each server in the cluster.

## Command Line

To set the **preemptive shutdown** timeout values via command line, use the following command:

```
/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam -c -s
-i TEST.QM -p PREEMPTIVE_TIMEOUT -v 20
```

This will set (-s) the LifeKeeper preemptive shutdown timeout (-p) on each node of the cluster (-c) to *20 seconds* (-v) for queue manager TEST.QM (-i).

To set the **immediate shutdown** timeout values via command line, use the following command:

```
/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam -c -s
-i TEST.QM -p IMMEDIATE_TIMEOUT -v 20
```

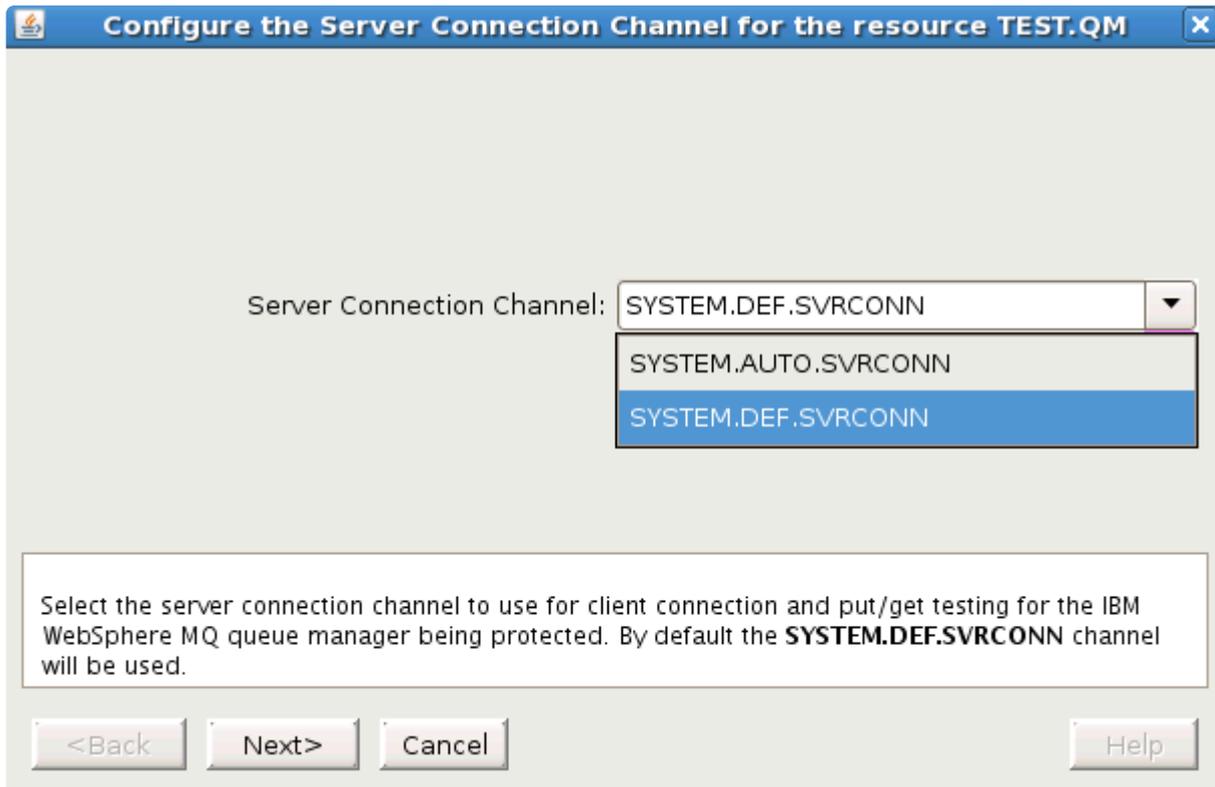
This will set (-s) the LifeKeeper immediate shutdown timeout (-p) on each node of the cluster (-c) to *20 seconds* (-v) for queue manager TEST.QM (-i).

**Note:** You can either use the queue manager name (-i) or the LifeKeeper TAG (-t) name.

## 6.8.5.7.5. Changing the Server Connection Channel

### GUI

First navigate to the WebSphere MQ resource properties panel or the resource context menu described above. The resource must be in service to modify the **Server Connection Channel** value. Then click on **Server Connection Channel Configuration** icon or menu item. The following dialog will appear:



Now select the **Server Connection Channel** to use and click **Next**. You will be asked if you want to change to the new **Server Connection Channel**, click **Continue**. Next, the LifeKeeper GUI will set the Server Connection Channel for the selected queue manager on each server in the cluster.

### Command Line

To set the Server Connection Channel via command line use the following command:

```
/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam -c -s
-i TEST.QM -p CHANNEL -v LK.TEST.SVRCONN
```

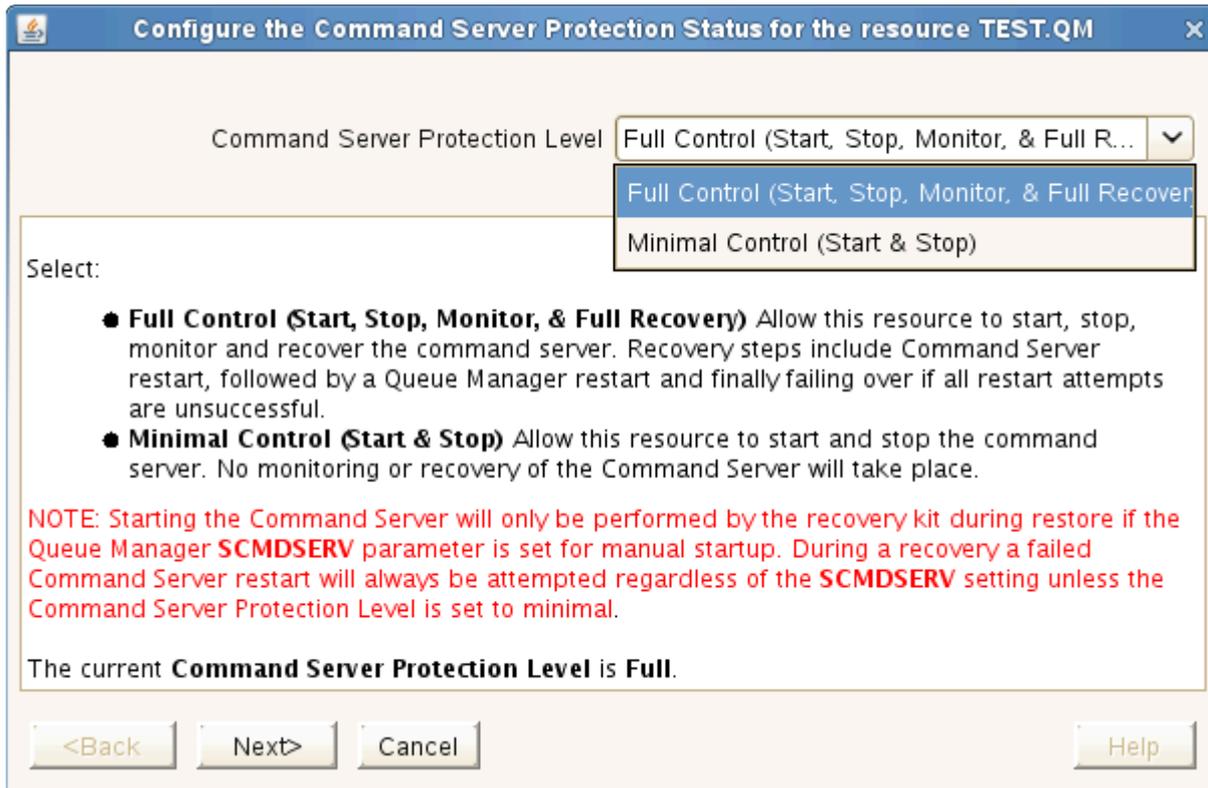
This will set (-s) the Server Connection Channel (-p) on each node of the cluster (-c) to *LK.TEST.SVRCONN* for queue manager *TEST.QM* (-i).

 **Note:** You can either use the queue manager name (-i) or the LifeKeeper TAG (-t) name.

## 6.8.5.7.6. Changing the Command Server Protection Configuration

### GUI

First navigate to the WebSphere MQ **Resource Properties Panel** or the **Resource Context Menu** described above. Then click on **Command Server Protection Configuration** icon or menu item. The following dialog will appear:



Select **Full Control** of the command server component of the WebSphere MQ queue manager to have LifeKeeper start, stop, monitor and attempt to recover and to then fail over if the recovery attempt is unsuccessful.

Select **Minimal Control** of the command server component of the WebSphere MQ queue manager to have LifeKeeper only start and stop but not monitor or attempt any recovery.

See [above table](#) for more details. Once the protection control is selected, click **Next**. You will be asked if you want to change the setting of the command server protection from its current setting to the new setting; click **Continue** to make the change on all nodes in the cluster.

### Command Line

To set the **LifeKeeper Command Server Protection Configuration** via the command line, use the following command:

```
/opt/LifeKeeper/lkadm/subsys/appsuite/mqseries/bin/mq_modqmgrparam -c -s
```

```
-i TEST.QM -p CMDSERVERPROTECTION -v LEVEL
```

where **LEVEL** is **Full** or **Minimal**.

This will set (-s) the LifeKeeper Command Server Protection Configuration (-p) on each node in the cluster (-c) to LEVEL (-v) for queue manager TEST.QM (-i).

**Note:** You can use either the queue manager name (-i) or the LifeKeeper TAG (-t) name.

## 6.8.5.7.7. Changing LifeKeeper WebSphere MQ Recovery Kit Defaults

The IBM WebSphere MQ Recovery Kit uses a number of default values which can be tuned and modified if you have problems with the default settings. The default settings should be sufficient for most environments. If you have problems with timeouts you can use the following table to identify tunable parameters. It is recommended that you do not change the parameters until you have problems with your WebSphere MQ resource hierarchies.

Variable Name in /etc/default/LifeKeeper	Default Value	Description
MQS_QUICKCHECK_TIMEOUT_SC	10 (seconds)	Timeout for the client connect check.
MQS_QUICKCHECK_TIMEOUT_CC	10 (seconds)	Timeout for the client connect check.
MQS_QUICKCHECK_TIMEOUT_PUTGET	10 (seconds)	Timeout for the PUT/GET check
MQS_QUICKCHECK_TIMEOUT_PS	5 (seconds)	Timeout for the check whether publish/subscribe is in use or not
MQS_QUICKCHECK_TIMEOUT_CLUSTER	5 (seconds)	Timeout for the check whether this queue manager is part of an WebSphere MQ cluster
MQS_QUICKCHECK_TIMEOUT	40 (seconds)	Timeout for the quickCheck script (must be at least 10 seconds).
MQS_QMGR_START_TIMEOUT	60 (seconds)	Timeout for the queue manager start command to complete.
MQS_CMDS_START_TIMEOUT	30 (seconds)	Timeout for the command server start command to complete.
MQS_LISTENER_START_TIMEOUT	30 (seconds)	Timeout for the listener start command to complete
MQS_LISTENER_LIST_TIMEOUT	10 (seconds)	Timeout for the listener list command to complete
MQS_CHECK_TIMEOUT_ACTION	ignore	The action in case a server connect check or client connect check times out. The default of "ignore" means that a message about the timeout is logged, but no recovery is initiated. If you set this variable to "sendevent" local recovery is initiated in case a server connect check timed out.

MQS_LISTENER_CHECK_DELAY	2 (seconds)	The time in seconds between the start of the listener and the check for the successful listener start. The default of 2 seconds should be sufficient to detect port in use conditions.
NO_AUTO_STORAGE_DEPS	0	If you set the variable to 1 the recovery kit does not check if the queue manager and log directory are located in shared storage. If set to 1 the recovery kit does not create file system hierarchies upon resource configuration too.
MQS_DSPMQVER_TIMEOUT	5 (seconds)	Timeout for the dspmqver command (needed to find out the version of WebSphere MQ), must be at least 2 seconds.
MQS_SKIP_CRT_MISSING_Q	0	Set to 1 to not automatically create a missing test queue.
MQS_ALT_USER_NAME	mqm if not set or the user does not have membership in the "mqm" group	The alternate user name to use for all WebSphere MQ commands. By default the user "mqm" is used. If set the alternate user must have its primary group set to the group "mqm" or must have secondary membership in that group.

To change the parameters add the appropriate variable in the table above to `/etc/default/LifeKeeper`. The line should have the following syntax:

```
[...]
MQS_CHECK_TIMEOUT_ACTION=sendevent
[...]
```

To disable a custom setting and fall back to the default value, just remove the line from `/etc/default/LifeKeeper` or comment out the corresponding line.

## 6.8.6. WebSphere MQ Troubleshooting

---

### WebSphere MQ Log Locations

If the queue manager name is known and the queue manager is available, WebSphere MQ error logs are located in the directory specified by the `LogPath` parameter defined in the queue manager configuration file `qm.ini`. If the queue manager is not available, error logs are located in: `/var/mqm/qmgrs/@SYSTEM/errors`. If an error has occurred with a client application, error logs are located on the client's root drive in: `/var/mqm/errors`.

If your application gets a return code indicating that a Message Queue Interface (MQI) call has failed, refer to the *WebSphere MQ Application Programming Reference Manual* for a description of that return code.

## 6.8.6.1. MQ Error Messages

This section provides a list of messages that you may encounter with the use of the LifeKeeper MQ Recovery Kit. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Because the MQ Recovery Kit relies on other LifeKeeper components to drive the creation and extension of hierarchies, messages from these other components are also possible. In these cases, please refer to the [Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

### Common Error Messages

Error Number	Error Message	Action
119001	Queue manager with TAG "TAG" failed to start on server "SERVER" with return code "Code"	<p>The <code>start</code> command was successful, but the check after the start failed.</p> <p>Check the IBM WebSphere MQ alert log on SERVER for possible errors and correct them.</p>
119002	Queue manager with TAG "TAG" start command failed on server "SERVER" with return code "Code".	<p>The <code>start</code> command for the queue manager TAG returned with non zero value.</p> <p>Check the IBM WebSphere MQ alert log on SERVER for possible errors and correct them.</p> <p>The return code Code is the return code of the <code>strmqm</code> command.</p>
119006	Command server start command for queue manager "TAG" failed on server "SERVER" with return code "Code".	<p>The start command for the command server returned with none zero value.</p> <p>Check the IBM WebSphere MQ alert log on SERVER for possible errors and correct them.</p> <p>The return code Code is the return code of the <code>runmqsc</code> command.</p> <p>For WebSphere MQ v6.0, verify that</p>

		the command server startup type is "MANUAL". See section <a href="#">Configuration Requirements</a> for details.
119007	Listener for queue manager "TAG" failed to start on server "SERVER".	Check the IBM WebSphere MQ alert log on SERVER for possible errors and correct them.
119008	Listener start command for queue manager with TAG "TAG" failed on server "SERVER" with return code "CODE".	Check the IBM WebSphere MQ alert log on SERVER for possible errors and correct them.
119013	Could not create queue manager object for queue manager "QMGR" with TAG "TAG".	Check the LifeKeeper and WebSphere MQ error logs.
119014	Could not create listener object for queue manager "QMGR" with TAG "TAG".	Check the LifeKeeper and WebSphere MQ error logs.
119015	No value for the "PARAMETER" specified.	Run the LifeKeeper MQ Recovery Kit script with the correct arguments.
119016	Instance with ID "ID" does not exist on server "SERVER".	Check the resource hierarchy.
119017	Instance with TAG "TAG" does not exist on server "SERVER".	Check the resource hierarchy.
119018	Invalid parameters specified.	Run the script with the correct options.
119019	Too few parameters specified	Run the script with the correct options.
119021	Failed to set "VALUE" for resource instance "TAG" on server "SERVER".	Check the LifeKeeper log for possible errors setting the value.
119025	Failed to update instance info for queue manager with TAG "TAG" on server "SERVER".	When the server is up and running again, retry the operation to synchronize the settings.
119026	The following program required does not exist or is not executable: "EXECUTABLE". Check failed.	The program EXECUTABLE cannot be found. Verify all installation requirements are met and install all required packages.  See section <a href="#">Configuration Requirements</a> for details.
119032	Script: usage error (error message)	Start the script Script with the correct arguments
119033	Script: error parsing config file "ConfigFile".	Make sure ConfigFile exists and is readable.
119034	CHECKTYPE check for queue manager with TAG "TAG" failed on server "SERVER" because the MQUSER could not be determined. This is probably because of a	The CHECKTYPE check for queue manager with tag TAG failed.

	removed configuration file – ignoring.	Make sure the global configuration file (mqs.ini) exists and is readable.  If it is removed, recreate the mqs.ini configuration file.
119035	CHECKTYPE check for queue manager with TAG “TAG” failed on server “SERVER” because no TCP PORT directive found in config file “CONFIGFILE” – ignoring.	Make sure the queue manager configuration file (qm.ini) exists and contains a TCP section as required during installation.  Add the TCP section to the queue manager configuration file.
119042	“CHECKTYPE” check for queue manager with TAG “TAG” failed on server “SERVER” because no TCP PORT information was found via runmqsc.	Verify that the port information for the listener objects has been defined and is accessible via runmqsc.
119043	TCP Listener configuration could not be read, reason: “REASON”.	Verify that MQ is running and the port information for the listener objects has been defined and is accessible via runmqsc.
119044	No TCP Listener configured, no TCP PORT information was found via runmqsc: “MESSAGE”.	Verify that the port information for the listener objects has been defined and is accessible via runmqsc.

## Create

Error Number	Error Message	Action
001022	END failed hierarchy “CREATE” of resource “TAG” on server “SERVER” with return value of “VALUE”.	Check the LifeKeeper log on server “SERVER” for possible errors creating the resource hierarchy. The failure is probably associated with the queue manager not starting.
119020	Create MQSeries queue manager resource with TAG “TAG” for queue manager “QMGR” failed.	Check the LifeKeeper log for possible errors creating the resource. The failure is probably associated with the queue manager not starting.
119022	Failed to create dependency between “PARENT” and “CHILD”.	Check the LifeKeeper log for possible errors creating the dependency.
119023	Creating the filesystem hierarchies for queue manager with TAG “TAG” failed. File systems: “Filesystems”.	Check the LifeKeeper log for possible errors creating the filesystem hierarchies.
119029	No TCP section configured in “CONFIGFILE” on server “SERVER”.	Add the TCP section to the queue manager configuration file on server SERVER. See section <a href="#">Configuration Requirements</a> for details.

119031	Queue manager "DIRTYTYPE" directory ("DIRECTORY") not on shared storage.	Move the directory DIRECTORY to shared storage and retry the operation.
119038	Creation of queue manager resource with TAG "TAG" failed on server "SERVER".	Check the LifeKeeper log on server SERVER for possible errors, correct them and retry the operation.
119039	TCP section in configuration file "FILE" on line "LINE1" is located before LOG section on line "LINE2" on server "SERVER".	It's recommended for the TCP section to be located after the LOG: section in the queue manager configuration file.  Move the TCP section to the end of the queue manager configuration file and retry the operation.
119040	Creation of MQSeries queue manager resource by create_ins was successful but no resource with TAG "TAG" exists on server "SERVER". Sanity check failed.	Check the LifeKeeper log for possible errors during resource creation.
119041	Creation of MQSeries queue manager resource was successful but no resource with TAG "TAG" exists on server "SERVER". Final sanity check failed.	Check the LifeKeeper log for possible errors during resource creation.

## Extend

Error Number	Error Message	Action
119024	Instance "TAG" can not be extended from "TEMPLATESYS" to "TARGETSYS". Reason:REASON	Correct the failure described in REASON and retry the operation.
119027	The user "USER" does not exist on server "SERVER".	Create the user USER on SERVER with the same UID as on the primary server and retry the operation.
119028	The user "USER" has a different numeric UID on server "SERVER1" (SERVER1UID) then it should be (SERVER2UID).	Change the UID so that USER has the same UID on all servers and reinstall WebSphere MQ on the server where you have changed the UID and retry the operation.
119029	No TCP section configured in "CONFIGFILE" on server "SERVER".	Add the TCP section to the queue manager configuration file on server SERVER.  See section <a href="#">Configuration Requirements</a> for details.

119030	Queue manager “QMGR” not configured in “CONFIGFILE” on server “SERVER”.	The queue manager QMGR you are trying to extend is not configured in the global configuration file on the target server SERVER. Add the queue manager stanza to the config file CONFIGFILE on server SERVER and retry the operation.
119036	Link “LINK” points to “LINKTARGET” but should point to “REALTARGET” on server “SERVER”.	For file system layout 3 symbolic links must point to the same location on the template and target server SERVER.  Correct the link LINK on server SERVER to point to REALTARGET and retry the operation.
119037	Link “LINK” that should point to “REALTARGET” does not exist on system “SERVER”.	For file system layout 3 symbolic links must also exist on the target server.  Create the required link LINK to REALTARGET on server SERVER and retry the operation.

## Remove

Error Number	Error Message	Action
119003	Failed to stop queue manager with TAG “TAG” on server “SERVER”.	The queue manager “TAG” on server “SERVER” could not be stopped through the Recovery Kit. For further information and investigation, change the logging level to DEBUG. Depending on the machine load the shutdown timeout values possibly have to be increased.
119004	Some orphans of queue manager with TAG “TAG” could not be stopped on server “SERVER”. Tried it “tries” times.	Try killing the orphans manually and restart the Queue Manager again. For further information change the logging level to DEBUG.
119010	Listener for queue manager with TAG “TAG” failed to stop on server “SERVER”.	This message will only appear if the monitoring for the listener is enabled. For further information change the logging level to DEBUG.

## Resource Monitoring

Error Number	Error Message	Action
119005	Queue manager with TAG “TAG” on server “SERVER” failed.	Check the IBM WebSphere MQ alert log on SERVER for possible errors. This message indicates a queue manager crash
119009	Listener for queue	This message will only appear if monitoring of the listener is enabled.

	manager with TAG "TAG" failed on server "SERVER".	For further information change the logging level to FINE.
119011	"CHECKTYPE" PUT/GET Test for queue manager with TAG "TAG" failed on server "SERVER" with return code "Code"	This message will only appear if the PUT/GET Test is enabled and the test queue exists. For further information change the logging level to FINE and check the IBM WebSphere queue manager error log (/var/mqm/errors) on SERVER for possible errors and correct them. Verify that the file systems are not full.
119012	Client connect test for queue manager with TAG "TAG" on server "SERVER" failed with return code "Code".	<p>This message will only appear if Listener management is enabled.</p> <p>This message indicates a problem with the listener or the queue manager.</p> <p>Check the log for possible errors and correct them.</p> <p>The return code Code is the return code of the amqscnxc command.</p>

### Warning Messages

Error Number	Error Message	Action
119201	Listener for queue manager with TAG "TAG" is NOT monitored on server "SERVER".	This is a warning that listener management is not enabled.
119202	Queue manager with TAG "TAG" is not running on server "SERVER" but some orphans are still active. This is attempt number "ATTEMPT" at stopping all orphans processes.	This is a warning that MQ was not stopped properly.
119203	Another instance of recover is running, exiting "EXITCODE".	Recovery was started, but another recovery process was already running, so this process will not continue.
119204	Queue manager server connect check for queue manager with TAG "TAG%" timed out after "SECONDS" seconds on server "SERVER".	<p>If you see this message regularly increase the value of MQS_QUICKCHECK_TIMEOUT_SC in /etc/default/LifeKeeper.</p> <p>See section <a href="#">Changing the Server Connection Channel</a> for details.</p>
119205	Queue manager client connect check for queue manager with TAG "TAG" timed out after "SECONDS" seconds on server "SERVER".	If you see this message regularly increase the value of MQS_QUICKCHECK_TIMEOUT_CC in /etc/default/LifeKeeper.

		See section <a href="#">Changing the Server Connection Channel</a> for details.
119206	Server "SERVER" is not available, skipping.	A server was not online while updating a queue manager configuration setting.  Wait for the server to be online again and repeat the configuration step.
119207	"CHECKTYPE" PUT/GET test for queue manager with TAG "TAG" failed because test queue "QUEUE" does not exist (reason code "REASONCODE") – ignoring.	Create the test queue QUEUE configured or reconfigure the test queue to an existing queue.  See section <a href="#">Configuration Requirements</a> for details on creating the test queue.
119208	Channel "CHANNEL" does not exist for queue manager with TAG "TAG" (reason code "REASONCODE") – ignoring.	Create the channel "CHANNEL" which does not appear to exist. By default the channel SYSTEM.DEF.SVRCONN is used.  See the WebSphere MQ documentation for details on how to create channels.
119209	PUT/GET test for queue manager with TAG "TAG" skipped because no test queue is defined.	Configure a LifeKeeper test queue for queue manager TAG,
19210	The following program required to perform the PUT/GET test does not exist or is not executable: "EXECUTBALE". Test skipped.	Install a C compiler on the system and make sure it is in the root users PATH environment variable. Run the script "LKROOT/lkadm/subsys/appsuite/mqseries/bin/compilesamples" to compile the modified sample amqsget and amqsgetc programs.
119211	Queue manager "CHECKTYPE" PUT/GET test for queue manager with TAG "TAG" timed out after "SECONDS" seconds on server "SERVER".	If you see this message regularly increase the value of MQS_QUICKCHECK_TIMEOUT_PUTGET in /etc/default/LifeKeeper.  See section <a href="#">Changing the Server Connection Channel</a> for details.
119212	QuickCheck for queue manager with TAG "TAG" timed out after SECONDS seconds on server "SERVER".	If you get this message regularly increase the value of MQS_QUICKCHECK_TIMEOUT in /etc/default/LifeKeeper
119213	mqseriesQueueManager::getMQVersion:: ERROR unexpected dspmqver output (OUTPUT) – using installation info	Reading the MQ version failed via runmqsc. If you get this message regularly, increase the value of MQS_DSPMQVER_TIMEOUT in /etc/default/

	instead (Queue QUEUE, Queuemanager QMGR).	LifeKeeper.
119214	mqseriesQueueManager::getMQVersion:: ERROR unexpected output retrieving MQ version information (Queue QUEUE, Queuemanager QMGR). Unexpected results *	Check if the following command yields some output when running as the mqm user: dspmqver -b -p1 -f2. Also, as the mqm user run the command dspmqinst and check what it returns.

## 6.8.7. Appendix A – Sample mqs.ini Configuration File

---

```

#*****#

#*

*#

#*

<START_COPYRIGHT>

*#

#* Licensed Materials - Property of
IBM                               *#

#*

63H9336

*#

#* © Copyright IBM Corporation 1994,
2000                               *#

#*

*#

#*

<END_COPYRIGHT>

```

```

#*****#
#* Module Name:
mqs.ini *#
#* Type : WebSphere MQ Machine-wide Configuration
File *#
#* Function : Define WebSphere MQ resources for an entire
machine *#
#*****#
#* Notes

```

```

*#
*# 1) This is the installation time default
configuration *#

```

```

*#
#*****#

```

AllQueueManagers:

```

#*****# #* The path to the qmgrs directory, below
which queue
manager data *# #* is
stored

```

```

*#
#*****# DefaultPrefix=/var/mqm

```

```

LogDefaults: LogPrimaryFiles=3 LogSecondaryFiles=2 LogFilePages=1024 LogType=CIRCULAR
LogBufferPages=0 LogDefaultPath=/var/mqm/log

```

QueueManager: Name=TEST.QM Prefix=/var/mqm Directory=TEST!QM

DefaultQueueManager: Name=TEST.QM

QueueManager: Name=TEST.QM.NEW Prefix=/var/mqm Directory=TEST!QM!NEW

QueueManager: Name=TEST.QM2 Prefix=/var/mqm Directory=TEST!QM2

```

QueueManager: Name=MULTIINS_1 Prefix=/var/mqm Directory=MULTIINS_1 DataPath=/opt/webmq/
MULTIINS_1/data

```

InstallationName=Installation1

```

QueueManager: Name=MULTIINS_2 Prefix=/var/mqm Directory=MULTIINS_2 DataPath=/opt/webmq/
MULTIINS_2/data InstallationName=Installation2

```



## 6.8.8. Appendix B – Sample qm.ini Configuration File

```

*****#
#* Module Name: qm.ini                                     *#
#* Type       : WebSphere MQ queue manager configuration file *#
# Function    : Define the configuration of a single queue manager *#
#*                                                  *#
*****#
#* Notes      :                                           *#
#* 1) This file defines the configuration of the queue manager
*#
*****#
ExitPath:
  ExitsDefaultPath=/var/mqm/exits/
  ExitsDefaultPath64=/var/mqm/exits64
#*                                                  *#
#*                                                  *#
Log:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=1024
  LogType=CIRCULAR
  LogBufferPages=0
  LogPath=/opt/MQ_log/MULTIINS_1
  LogWriteIntegrity=TripleWrite
Service:
  Name=AuthorizationService
  EntryPoints=14
ServiceComponent:
  Service=AuthorizationService
  Name=MQSeries.UNIX.auth.service
  Module=amqzfu
  ComponentDataSize=0

```

# 6.8.9. Appendix C – WebSphere MQ Configuration Sheet

Cluster name		
Contact information  (email or telephone number of person responsible for the cluster)		
LifeKeeper version		
Operating system		
Cluster nodes	name	public IP / netmask

Queue manager name		
Listener management by LifeKeeper	[ ] YES      [ ] NO	
WebSphere MQ operating system user	name	numeric (UID/GID)
user  (e.g. mqm/1002)		
group  (e.g. mqm/200)		
Virtual IP / netmask / network device  (eg. 192.168.1.1/24/eth0)		
Filesystem layout	__ Configuration 1 – /var/mqm on Shared Storage	
	__ Configuration 2 – Direct Mounts	
	__ Configuration 3 – Symbolic Links	
	__ Configuration 4 – Multi-Instance Queue Managers	

	__ other
Shared storage type	__ NAS (IP: _____)
	__ SCSI/FC (Type: _____)
	__ SDR
Queue manager <code>/var/mqm/qmgrs/QM.NAME</code> physical location (device, mount point or logical volume)  (e.g. LVM <code>/dev/mqm_test_qm/qmgrs</code> )	
Queue manager <code>/var/mqm/log/QM.NAME</code> physical location (device, mount point or logical volume) (e.g. LVM <code>/dev/mqm_test_qm/log</code> )	

## 6.9. NAS Recovery Kit Administration Guide

---

The LifeKeeper for Linux Network Attached Storage Recovery Kit (hereafter referred to as the NAS Recovery Kit) provides fault resilience for Network File System (NFS) software in a LifeKeeper environment. The NAS Recovery Kit affords LifeKeeper users the opportunity to employ an exported NFS file system as the storage basis for LifeKeeper hierarchies.

### Document Contents

This guide contain the following topics:

- [Documentation and References](#). Provides a list of LifeKeeper for Linux documentation and where to find them.
- [Requirements](#). A description of the hardware and software necessary to properly setup, install, and operate the NAS Recovery Kit. Refer to [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove LifeKeeper for Linux software.NAS Recovery Kit .
- [Overview](#). A description of the NAS Recovery Kit's features and functionality.
- [Configuring the LifeKeeper for Linux NAS Recovery Kit](#). A description of the procedures required to properly configure the NAS Recovery Kit.
- [LifeKeeper Configuration Tasks](#). A description of the tasks for creating and managing your NAS resource hierarchies using the LifeKeeper GUI.
- [Troubleshooting](#). A list of LifeKeeper for Linux error messages including a description for each.

## 6.9.1. NAS Documentation and References

---

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Installation Guide](#)
- [Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is available on the SIOS Technology Corp. website at:

<http://docs.us.sios.com>

## 6.9.2. NAS Recovery Kit Hardware and Software Requirements

---

### Requirements

Your LifeKeeper configuration must meet the following requirements **prior** to the installation of the LifeKeeper for Linux NAS Recovery Kit. Please see [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the configuration of your LifeKeeper hardware and software.

### Hardware Requirements

- **Servers** – LifeKeeper for Linux supported servers configured in accordance with the requirements described in [LifeKeeper for Linux Installation Guide](#) and [LifeKeeper for Linux Release Notes](#).
- **IP Network Interface Cards** – Each server requires at least one Ethernet TCP/IP-supported network interface card. Remember, however, that a LifeKeeper cluster requires two communications paths; two separate LAN-based communication paths using dual independent sub-nets are recommended for heartbeats, and at least one of these should be configured as a private network. Using a combination of TCP and TTY heartbeats is also supported.

### Software Requirements

- **TCP/IP software** – Each server in your LifeKeeper configuration requires TCP/IP software.
- **LifeKeeper Software** – It is imperative that you install the same version of the LifeKeeper for Linux software and apply the same versions of the LifeKeeper for Linux software patches to each server in your cluster.
- **LifeKeeper for Linux NAS Recovery Kit** – The NAS Recovery Kit is provided on the LifeKeeper Installation Image File (sps.img). It is packaged, installed and removed via the Red Hat Package Manager, rpm. The following rpm file is supplied on the LifeKeeper Installation Image File (sps.img):

#### **steeleye-lkNAS**

- **Linux software** – Each server in your cluster must have the **util-linux** package installed and configured prior to configuring LifeKeeper and the LifeKeeper NAS Recovery Kit. The NAS Recovery Kit requires version 2.9u or later of the **util-linux** package to assure proper functionality.

Please see the [LifeKeeper for Linux Installation Guide](#) for specific instructions on the installation and removal of the LifeKeeper for Linux software.

## 6.9.3. NAS Recovery Kit Overview

---

### LifeKeeper for Linux NAS Recovery Kit

The primary focus of the LifeKeeper for Linux NAS Recovery Kit is to offer LifeKeeper users an alternative storage method to shared storage and data replication.

The NAS Recovery Kit enables the creation of LifeKeeper resource hierarchies on LifeKeeper protected servers or clients that have imported (mounted) an exported Network File System (NFS) from either a Network Attached Storage device or an NFS server in the cluster. When a failure is detected on the node in the cluster where the exported file system is mounted, the NAS Recovery Kit initiates a fail over to the predetermined backup node.

Therefore, once the exported file system is mounted on a LifeKeeper server or client, it can be fully utilized as an additional storage basis for LifeKeeper hierarchies.

When you elect to use an exported file system as a storage medium, LifeKeeper does not require you to protect the server where the file system is exported. However, to achieve a greater degree of availability, users are encouraged to use the LifeKeeper for Linux NFS Server Recovery Kit to protect the server from failure where the file system is exported.

Resource hierarchies for the NAS Recovery Kit are created using the currently existing File System Recovery Kit available with the LifeKeeper Core product (**steeleye-ik** package).

While the NAS Recovery Kit delivers several advantages, the two most significant advantages are the elimination of the need for costly shared-storage devices and the capability to have multi-node cluster configurations.

### NAS Recovery Kit Restrictions

- This version of the NAS Recovery Kit does not include support for a local recovery when access to the NAS device fails. When a failure is detected, the default action is to initiate a transfer of the hierarchy to a backup server. Depending on the makeup of the resource hierarchy, this action can result in hung processes. To avoid hung processes, the default action can be changed to halt the server and force a failover to a backup server. To change the default switchover behavior, alter the setting of LKNASERROR in the LifeKeeper defaults file. See the section **Configuring the NAS Recovery Kit** later in this document for more discussion on LKNASERROR.
- The NAS Recovery Kit does not provide protection for your Network Attached Storage device. The objective of this kit is to expand LifeKeeper storage options into the Network Attached Storage arena.
- The NAS Recovery Kit does not permit the NFS file system to be mounted more than once on different mount points. Attempts to create hierarchies when the file system is found in the */etc/fstab* file multiple times will fail.

- The NAS Recovery Kit does not support the following format of NAS device or NFS.
  - “.” is included in the exported path (e.g. 192.168.1.10:/export/pa:th)
  - IPv6 is used (e.g. fc00::10:/export/data)

## 6.9.4. Configuring the LifeKeeper for Linux NAS Recovery Kit

---

This section describes the LifeKeeper for Linux NAS Recovery Kit configuration details. It also contains information you should consider before you start to configure and administer the NAS Recovery Kit. Please refer to [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring LifeKeeper Core resource hierarchies.

## 6.9.4.1. NAS Configuration Considerations

---

The following should be considered before operating the LifeKeeper for Linux NAS Recovery Kit:

1. Install the NAS Recovery Kit on the server(s) in your cluster configuration where you wish to mount your exported file systems and where you will extend your NAS resource hierarchy. You can export your file system from either a NFS server, which may be protected by LifeKeeper (this is the recommended configuration), or from a Network Attached Storage device.
2. To ensure proper execution of this kit, it is highly recommended that you mount your exported NFS file system using the server's IP address in place of the server name and that you perform your mount operation before you place your file system under LifeKeeper protection. Additionally, if you are mounting a file system that is currently protected by the LifeKeeper for Linux NFS Server Recovery Kit, we strongly suggest that the IP address used to create the NFS Server hierarchy be used to mount the file system on the LifeKeeper NAS server.
3. To eliminate the possibility of split-brain related problems (i.e. more than one node in the cluster has a hierarchy In Service Protected (ISP)), we highly recommend that you establish one of the communication paths between nodes in the cluster on the same network used to access the exported file system. Failure to comply with this recommendation can result in multiple nodes bringing the hierarchy ISP (split-brain) when a communication path failure occurs. To recover from a split-brain scenario, take all but one of the ISP hierarchies out of service. This will ensure that only one node has access to the exported file system.
4. The built-in file system recovery kit used to build NAS hierarchies cannot detect and remove processes not protected by LifeKeeper that are using the mounted file system in a fail over condition. Therefore, it is highly recommended that only LifeKeeper protected processes use the NAS protected file system.
5. The LKNFSTIMEOUT tunable represents the timeout in seconds the NAS Recovery Kit will use when attempting to determine the status of a NFS mounted file system. The default value for this tunable is set to 2 minutes. The LKNFSSYSCALLTO tunable represents the timeout in seconds the NAS Recovery Kit will use for alarms to interrupt system calls when attempting to determine the status of a mount point. Use the formula below to determine the value for this tunable:  
  
3 times your LKNFSSYSCALLTO value plus 5 should be less than the value of LKNFSTIMEOUT.
6. The LKNASERROR tunable controls the actions the NAS Recovery kit takes when access to the NAS device fails. The tunable has two values, **switch** and **halt**, with **switch** being the default. If the value is set to switch and access fails, the NAS Recovery Kit will initiate a transfer of the resource hierarchy to a backup server when the failure is detected. The attempt to transfer the resource hierarchy to the backup server can hang if any of the resources sitting above the NAS resource attempt to access anything on the NAS file system. To avoid this problem the tunable value can be set to **halt**, which will immediately halt the system when an access failure is detected. This action will force a failover of all resource hierarchies to the backup server.
7. STONITH devices or the Quorum/Witness package should be used so that a machine failure (all

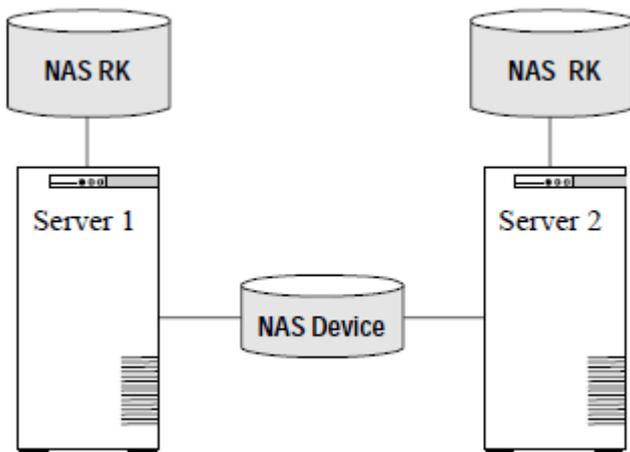
comm paths are down) does not result in a split brain where all the NAS resources are in service on all nodes in the cluster. This condition can lead to data corruption. More details on the Quorum/Witness package can be found in the LifeKeeper for Linux Technical Documentation.

## 6.9.4.2. NAS Configuration Examples

### Configuration Examples

A few examples of what happens during a fail over using LifeKeeper for Linux NAS Recovery Kit are provided below.

#### Configuration 1: Active/Standby Configuration Example



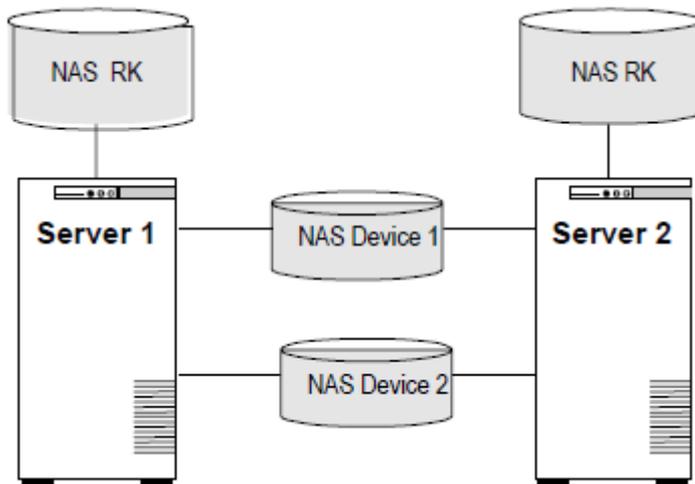
In this configuration, Server 1 is considered active because it is running the NAS Recovery Kit software and has imported (mounted) the file system from the NAS device. Server 2 does other processing. If Server 1 fails, Server 2 gains access to the file system and uses the LifeKeeper secondary hierarchy to make it available to clients.

#### Configuration Notes:

- The NAS software must be installed on both servers.
- The file system has been imported from a NAS device.
- Server 2 should not access files and directories on the NAS device while Server 1 is active.

\* **Note:** In an active/standby configuration, Server 2 might be running the NAS Recovery Kit, but does not have any other NAS resources under LifeKeeper protection.

## Configuration 2: Active/Active Configuration Example



An active/active configuration consists of two or more systems actively running the NAS Recovery Kit software and importing file systems from NAS device(s).

### Configuration Notes:

- The NAS software must be installed on both servers.
- Initially, Server 1 imports a file system and Server 2 imports a different file system. In a switchover situation, one system can import both file systems.

## 6.9.5. LifeKeeper Configuration Tasks for NAS

---

You can perform all LifeKeeper for Linux NAS Recovery Kit administrative tasks via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer, and monitor NAS resources.

The following tasks are available for configuring the LifeKeeper for Linux NAS Recovery Kit:

- [Create a Resource Hierarchy](#) – Creates a NAS resource hierarchy.
- [Delete a Resource Hierarchy](#) – Deletes a NAS resource hierarchy.
- [Extend a Resource Hierarchy](#) – Extends a NAS resource hierarchy from the primary server to the backup server.
- [Unextend a Resource Hierarchy](#) – Unextends (removes) a NAS resource hierarchy from a single server in the LifeKeeper cluster.
- [Create Dependency](#) – Creates a child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete Dependency](#) – Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#) – Activates a resource hierarchy.
- [Out of Service](#) – Deactivates a resource hierarchy.
- [View](#) / [Edit](#) Properties – View or edit the properties of a resource hierarchy.

 **Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You can also perform most of these tasks:

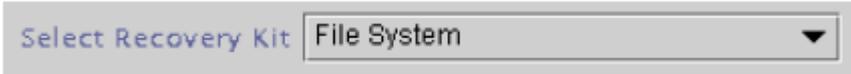
1. From the toolbar, right – click on a global resource in the left pane of the status display.
2. Right – click on a resource instance in the right pane of the status display.

 **Note:** Using the right-click method allows you to avoid entering information that is required when using the **Edit** menu.

## 6.9.5.1. Creating a NAS Resource Hierarchy

Perform the following on your primary server:

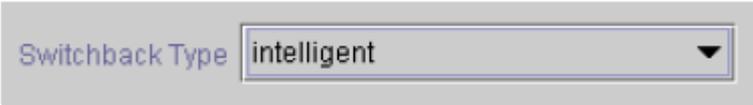
1. Select **Edit > Server > Create Resource Hierarchy**.
2. The **Select Recovery Kit** dialog appears. Select the **File System** option from the drop down list. Simply put, a NAS Resource Hierarchy is a File System Hierarchy created using a NFS mounted file system.



Click **Next** to continue.

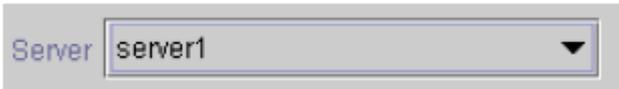
**! CAUTION:** If you click the **Cancel** button at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

3. The **Switchback Type** dialog appears. The switchback type determines how the NAS resource will be switched back to the primary server when it becomes in-service (active) on the backup server following a failover. Switchback types are either *intelligent* or *automatic*. Intelligent switchback requires administrative intervention to switch the resource back to the primary server while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.



Click **Next** to continue.

4. The **Server** dialog appears. Select the name of the server where the NAS resource will be created (typically this is your primary server). All servers in your cluster are included in the drop down list box.



Click **Next** to continue.

5. Select the **Mount Point** path to be protected by the NAS (File System) Resource Hierarchy. All "local" (i.e. file systems using shared storage) and NFS mounted file systems are listed. Select the NFS mounted file system from the drop down list box.

Mount Point

Click **Next** to continue.

- The **Root Tag** dialog is automatically populated with a unique name for the resource instance on the target server (i.e. the server selected above). You may accept the default or enter a unique tag consisting of letters, numbers and the following special characters: -, \_, ., or /.

Root Tag

Click **Create Instance**.

- An information box appears indicating the start of the hierarchy creation.

```

Creating gen/filesys resource...
07/27/2001 15:36:37 create: BEGIN creation of "device-nas20988" on
server "server1"
07/27/2001 15:36:37 create: END successful creation of
"device-nas20988" on server "server1"
07/27/2001 15:36:38 restore: BEGIN restore of "device-nas20988" on
server "server1"
07/27/2001 15:36:38 restore: END successful restore of
"device-nas20988" on server "server1"
Creating Resource Instance test-on-server1 with id /test on machine
"server1":
Resource test-on-server1 Successfully Created on machine "server1"
Creating Dependency test-on-server1-"device-nas20988" on machine
"server1":
Dependency test-on-server1-"device-nas20988" Successfully Created
on machine "server1"
Removing /etc/fstab entry

```

Click **Create** to continue.

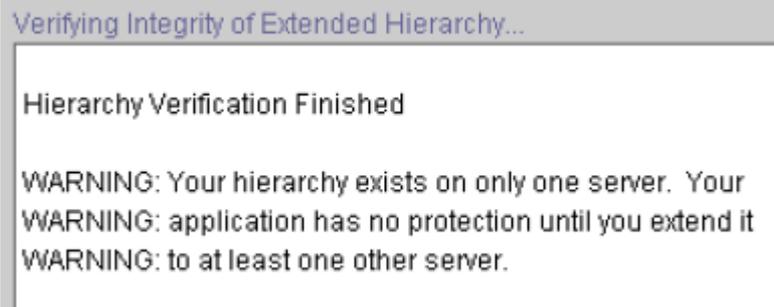
- An information box appears announcing the successful creation of your NAS resource hierarchy. You must **Extend** the hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

You have successfully created a resource hierarchy on one server. You may select continue in order to extend this resource hierarchy to another server, or you may cancel at this point.

If you cancel, the resource hierarchy provides no protection for your applications until it is extended to at least one other server in the cluster.

Click **Continue** to extend the resource.

Click **Cancel** if you wish to extend your resource at another time.



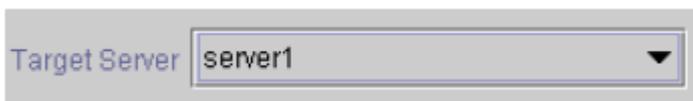
9. Click **Done** to exit the Create Resource Hierarchy menu selection.

## 6.9.5.2. Deleting a NAS Resource Hierarchy

To delete a NAS resource from all servers in your LifeKeeper configuration, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your NAS resource hierarchy.

\* **Note:** If you selected the Delete Resource task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

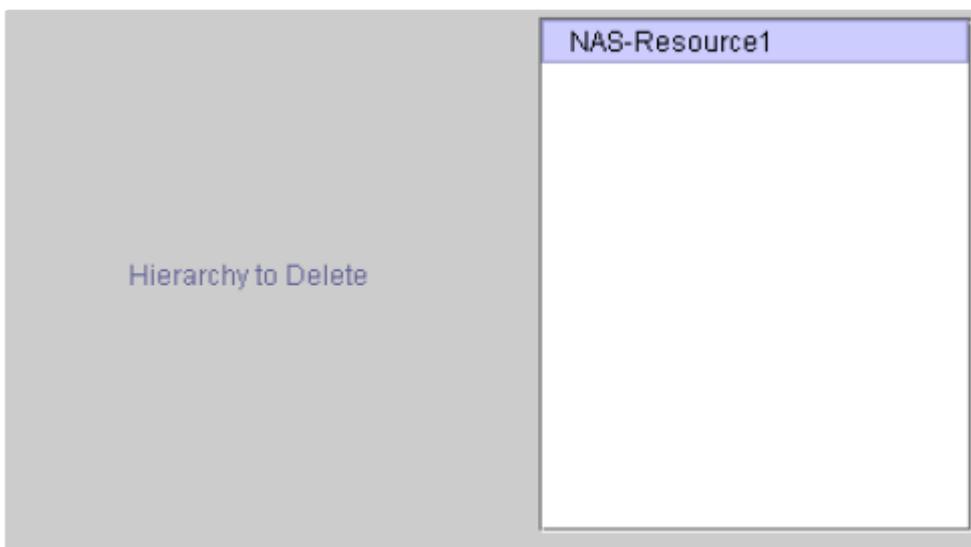


Target Server server1

Click **Next** to continue.

3. Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete, and highlight it.

\* **Note:** If you selected the Delete Resource task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Hierarchy to Delete

NAS-Resource1

Click **Next** to continue.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete.

You have specified the following resource hierarchy for deletion.  
Target Server: server1  
Target Tags: NAS-Resource1

Click **Delete** to continue.

5. An information box appears confirming that the NAS resource instance was deleted successfully.



```
Deleting resource hierarchy...
server1,priv_globact(1,delete): Running Post Global delete Scripts On
Machine server2
ins_remove[731,lraci.C]Mon Dec 11 12:53:44 EST 2000:
server1,priv_globact(1,delete): Post Global delete Scripts Finished
Exiting 0 On Machine server2 With Output Following:
lcdrecover[718,lraci.C]Mon Dec 11 13:46:23 EST 2000:
server2,priv_globact(1,delete): Running Post Global delete Scripts On
Machine server1
ins_remove[733,lraci.C]Mon Dec 11 12:53:44 EST 2000:
server1,priv_globact(1,delete): Local Post Global delete Scripts
Processing Continues...
```

6. Click **Done** to exit the Delete Resource Hierarchy menu selection.

## 6.9.5.3. Extending Your NAS Hierarchy

After you have created a hierarchy, you should extend that hierarchy to another server in the cluster. There are three possible ways to extend your resource instance:

1. When you successfully create your NAS resource hierarchy you will have an opportunity to select **Continue** which will allow you to proceed with extending your resource hierarchy to your backup server.
2. Right-click on an unextended hierarchy in either the left or right pane on the LifeKeeper GUI.
3. Select the **Extend Resource Hierarchy** task from the LifeKeeper GUI by selecting **Edit, Resource, Extend Resource Hierarchy** from the drop down menu. This sequence of selections will launch the Extend Resource Hierarchy wizard. The **Accept Defaults** button that is available for the **Extend Resource Hierarchy** option is intended for the user who is familiar with the LifeKeeper Extend Resource Hierarchy defaults and wants to quickly extend a LifeKeeper resource hierarchy without being prompted for input or confirmation. Users who prefer to extend a LifeKeeper resource hierarchy using the interactive, step-by- step interface of the GUI dialogs should use the **Next** button.
  - a. The first dialog box to appear will ask you to select the **Template Server** where your NAS resource hierarchy is currently in service. Remember that the **Template Server** you select now and the **Tag to Extend** that you select in the next dialog box represent an in- service (activated) resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you have selected. The drop down box in this dialog provides the names of all the servers in your cluster.

 **Note:** If you are entering the Extend Resource Hierarchy task by continuing from the creation of a NAS resource hierarchy, this dialog box will not appear because the wizard has already identified the template server in the create stage. This is also the case when you right-click on either the NAS resource icon in the left pane or right-click on the NAS resource box in the right pane of the GUI window and choose Extend Resource Hierarchy.

Template Server

 If you click the **Cancel** button at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extend hierarchy process. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Click **Next** to continue.

b. Select the **Tag to Extend**. This is the name of the NAS instance you wish to extend from the template server to the target server. The wizard will list in the drop down box all of the resources that you have created on the template server.

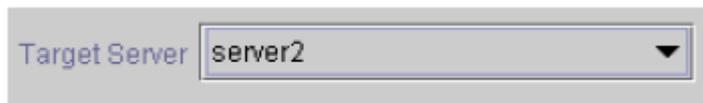
 **Note:** Once again, if you are entering the Extend Resource Hierarchy task immediately following the creation of a NAS hierarchy, this dialog box will not appear because the wizard has already identified the tag name of your resource in the create stage. This is also the case when you right-click on either the NAS resource icon in the left pane or on the NAS resource box in the right pane of the GUI window and choose *Extend Resource Hierarchy*.



A screenshot of a GUI element labeled "Tag to Extend" with a dropdown menu. The dropdown menu is open and shows the selected option "NAS-Resource1".

Click **Next** to continue.

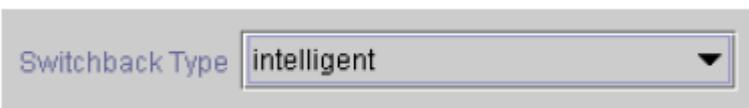
c. Select the **Target Server** where you will extend your NAS resource hierarchy.



A screenshot of a GUI element labeled "Target Server" with a dropdown menu. The dropdown menu is open and shows the selected option "server2".

Click **Next** to continue.

d. The **Switchback Type** dialog appears. The switchback type determines how the NAS resource will be switched back to the primary server when it becomes in service (active) on the backup server following a failover. Switchback types are either intelligent or automatic. Intelligent switchback requires administrative intervention to switch the resource back to the primary server while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.



A screenshot of a GUI element labeled "Switchback Type" with a dropdown menu. The dropdown menu is open and shows the selected option "intelligent".

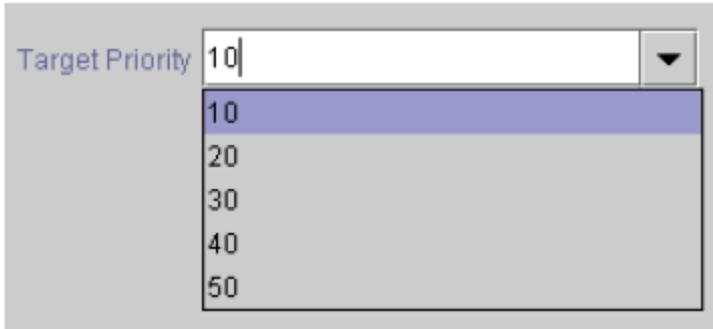
Click **Next** to continue.

e. Select or enter a **Template Priority**. This is the priority for the NAS hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.

 **Note:** This selection will appear only for the initial extend of the hierarchy.

Click **Next** to continue.

f. Select or enter the **Target Priority**. This is the priority for the new extended NAS hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.



Click **Next** to continue.

g. An information box appears explaining that LifeKeeper has successfully checked your environment and that all requirements for extending this resource have been met. If there are requirements that have not been met, LifeKeeper will disable the **Next** button, and enable the **Back** button.

```
Executing the pre-extend script..
Checking existence of extend and canextend scripts
Building independent resource list
Checking extendability for NAS-Resource1
Pre Extend checks were successful
```

Click on the **Back** button to make changes to your resource extension.

Click **Cancel** to extend your resource another time.

Click **Next** to launch the Extend Resource Hierarchy configuration task.

Click **Finish** to confirm the successful extension of your NAS resource instance.

4. Click **Done** to exit the Extend Resources Hierarchy menu selection.

 **Note:** Be sure to test the functionality of the new instance on both servers.

## 6.9.5.4. Unextending Your NAS Hierarchy

1. From the LifeKeeper GUI menu, select **Edit, Resource, and Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the NAS resource. It cannot be the server where the resource is currently in service (active).

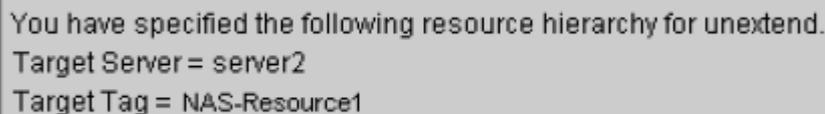
**Note:** If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

3. Select the NAS **Hierarchy to Unextend**.

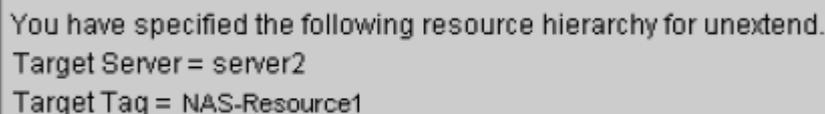
**Note:** If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

A screenshot of a grey dialog box with a white border. The text inside reads: 'You have specified the following resource hierarchy for unextend. Target Server = server2 Target Tag = NAS-Resource1'.

You have specified the following resource hierarchy for unextend.  
Target Server = server2  
Target Tag = NAS-Resource1

Click **Next** to continue.

4. An information box appears confirming the target server and the NAS resource hierarchy you have chosen to unextend.

A screenshot of a grey dialog box with a white border. The text inside reads: 'You have specified the following resource hierarchy for unextend. Target Server = server2 Target Tag = NAS-Resource1'.

You have specified the following resource hierarchy for unextend.  
Target Server = server2  
Target Tag = NAS-Resource1

Click **Unextend**.

5. Another information box appears confirming that the NAS resource was unextended successfully.
6. Click **Done** to exit the Unextend Resource Hierarchy menu selection.

## 6.9.5.5. Testing Your NAS Resource Hierarchy

---

You can test your NAS resource hierarchy by initiating a manual switchover that will simulate a fail over of the resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource, and In Service**. For example, an in-service request executed on a backup server causes the NAS resource hierarchy to be placed in-service on the backup server and taken out-of-service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the **Out of Service** request, the resource hierarchy is taken out-of-service without bringing it in-service on the other server.

## 6.9.6. NAS Troubleshooting

Symptom	Possible Cause
LifeKeeper fail over operation fails with umount busy error.	<p>The file system kit used to build NAS hierarchies cannot detect and remove processes not protected by LifeKeeper that are using the mounted file system in a fail over condition. Therefore, it is highly recommended that only LifeKeeper protected processes use the NAS protected file system. In the event of this failure, you must identify the processes using the file system and kill them. The <b>fuser -m</b> command can be used to determine the processes currently accessing the file system. Please see the <b>fuser</b> man pages for details on its use.</p>
LifeKeeper does local recovery of file system mounted via server name.	<p>If a file system protected by the NAS Recovery Kit was mounted via host name rather than IP address, then after creating the NAS resource, LifeKeeper logs a message similar to the following:</p> <pre data-bbox="288 757 959 786">. . . WARNING: Mon Aug 26 11:27:01 2002:</pre> <pre data-bbox="288 844 1422 918">LifeKeeper protected filesystem resource "tmp/mnt-on-tom.brown.com" (/tmp/mnt) is in service but not mounted</pre> <pre data-bbox="288 976 1011 1005">. . . Attempting Local Recovery of resource</pre> <p>LifeKeeper will re-mount the file system using the IP address at this point. However, if it encounters a problem, LifeKeeper will failover the NAS resource to the backup server (if it has been extended), or take the resource out of service (if the resource has not been extended).</p> <p><b>Suggested Action:</b> If the local recovery is successful, no further action is needed. However, if the local recovery fails, you should:</p> <ol data-bbox="323 1352 1193 1464" style="list-style-type: none"> <li>1. Delete the NAS resource in LifeKeeper.</li> <li>2. Re-mount the file system via IP address rather than host name.</li> <li>3. Re-create the NAS resource.</li> </ol>

## 6.9.6.1. NAS Error Messages

 **Note:** LifeKeeper takes resources out of service from the top of the hierarchy and works its way down to the other resources. When taking resources out of h2. Error Messages

This section provides a list of messages that you may encounter while creating and extending a LifeKeeper NAS resource hierarchy or removing and restoring a resource. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Messages from other LifeKeeper components are also possible. In these cases, please refer to the [Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

### NAS Recovery Kit Error Messages

Error Number	Error Message
107001	Creation of NAS device with tag id <tag id> on server <LifeKeeper server name> failed.
107002	Error getting list of IP addresses for NFS server device <NFS server name> on server <LifeKeeper server name>.
107003	Error attempting to find active address to NFS server <NFS server name> on server <LifeKeeper server name>.
107004	Error in format of device ID <resource device>.
107005	Cannot bring NAS resource <tag id> in service on server <LifeKeeper server name>. <b>Action:</b> After correcting the problem, try bringing the resource in service manually.
107006	create: Device not specified.
107007	Null Device returned by getlId on <LifeKeeper server name>.
107008	Cannot open /etc/mstab file on <LifeKeeper server name>.
107009	Illogical settings for NAS defaults on <LifeKeeper server name>. Using defaults of 120 for LKNFSTIMEOUT and 5 for LKNFSSYSCALLTO. <b>Action:</b> Reset NAS default values so that three times the value for LKNFSSYSCALLTO plus 5 is less than the value of LKNFSTIMEOUT.
107010	Error: detected conflict in expected tag name <tag id> on target machine <LifeKeeper server

	<p>name&gt;.</p> <p><b>Action:</b>Delete the conflicting resource and re-extend the hierarchy.</p>
107011	Error: mkdir of "/tmp/nas_mntpt.2915" on "mouse" failed: "permission denied".
107012	<p>Error: Exported file system &lt;NFS exported file system name&gt; cannot be accessed on &lt;server name&gt;.</p> <p>Possible causes:</p> <ul style="list-style-type: none"> <li>- The LifeKeeper node is not in the exported system list on the NFS server, or,</li> <li>- The exported system list has contradictory entries that are not displayed by the <b>showmount</b> command. (i.e. if exported system list exports a file system to both the world and to specific systems, <b>showmount</b> will report only the specific systems).</li> </ul> <p><b>Action:</b> Fix the exported file system access problem and re- extend the hierarchy.</p>
107013	<p>Error: Mount authorization check for "172.25.113.25:/ export" on "fred" appears to be hung. Exiting.</p> <p><b>Action:</b> Fix the access problem and re-extend the hierarchy.</p>
107015	Exported file system <tag> cannot be accessed on <LifeKeeper server name>.

## 6.9.6.2. LifeKeeper GUI Related Errors

---

Error Number	Error Message
104901	The mount point %s is mounted <b>Action:</b> Please specify a mount point that is not mounted.
104902	The mount point %s is not an absolute path <b>Action:</b> Please specify a mount point that begins with a slash.
104903	The mount point %s is not empty. <b>Action:</b> Please specify a mount point that does not exist or is empty.

## 6.10. NFS Server Recovery Kit Administration Guide

---

The LifeKeeper for Linux NFS Server Recovery Kit provides fault resilience for Network File System (NFS) software in a LifeKeeper environment. This enables a failure on the primary NFS server to be recovered on a designated backup server without significant lost time or human intervention.

[LifeKeeper for Linux NFS Server Recovery Kit Overview](#)

### LifeKeeper Documentation

The following is a list of LifeKeeper for Linux related information available from SIOS Technology Corp.:

- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Release Notes](#)
- [SIOS Technology Corp. Documentation](#)

### Reference Documents

The following is a list of reference documents associated with the LifeKeeper NFS Server Recovery Kit:

- [NFS Online documentation](#)
- *Managing NFS and NIS, Hal Stern, O'Reilly & Associates, Inc. 1991*

## 6.10.1. NFS Server Recovery Kit Overview

The NFS Server Recovery Kit provides a High Availability NFS service in hierarchical cooperation with the Filesystem Recovery Kit (provided as part of the steeleye-lk package) and the IP Recovery Kit (steeleye-lkIP).

The kit ensures that an IP resource and a file system resource containing the shared mount point are always in-service on the same server in the cluster. Clients who mount the file system using the LifeKeeper-protected IP resource can continue processing files on the volume virtually uninterrupted while the actual export service is switched between servers in the cluster (either manually or in response to a failure). Client recovery times will depend on the interaction between the client and the NFS server. For example, with NFSv3, the protocol timeouts for TCP are longer than that of UDP. In order to determine the best transport layer protocol to use with NFS, consider the recommendations of the OS vendor, the advantages and disadvantages of each transport protocol and your specific environment.

 **Note:** TCP transport is strongly recommended with NFSv4 by most OS vendors and the NFS Server Recovery Kit has been validated with only TCP transport and NFSv4.

All files on the file system become temporarily unavailable while a switchover or failover is in progress, but they become available again transparently when the resource transfer is complete. For a switchover, this can take between 5 and 30 seconds. For a failover, the recovery time depends on how long it takes to repair the file system. It is strongly recommended that you format the underlying disk volume with a Journaling File System (JFS) which is extremely robust to failure and can be repaired in a few seconds.

Beginning with LifeKeeper v9.6.0, NFS v4 pseudo file system is no longer supported. LifeKeeper cannot protect the fsid=0 export point and its sub directory's export point. If you are creating a fsid=0 resource using the version prior to LifeKeeper v9.6.0, it is necessary to delete the resource and then upgrade LifeKeeper.

Using the OS default setting, both NFS v3 and 4 are enabled. Specify the version using the mount option on clients.

## 6.10.2. NFS Server Recovery Kit Requirements

---

Before installing and configuring the LifeKeeper NFS Server Recovery Kit, be sure that your configuration meets the following requirements:

- **Servers.** The Recovery Kit requires two or more servers configured in accordance with the LifeKeeper requirements described in the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#). See the [Support Matrix](#) for supported Linux distributions.
- **LifeKeeper software.** You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.
- **LifeKeeper IP Recovery Kit.** You must have the same version of the LifeKeeper IP Recovery Kit on each server.
- **IP network interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP resource switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

 **Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation and so forth.

- **TCP/IP software.** Each server also requires the TCP/IP software.
- **NFS software.** The LifeKeeper Installation Support setup script will configure settings for use in an HA environment. The following software must be installed on each server prior to configuring LifeKeeper NFS Server Recovery Kit. The same version of the software should be installed on each server.
  - rpcbind
  - rpc.idmapd
  - gssproxy or rpc.svcgssd
  - rpc.nfsd
- Refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper NFS Server Recovery Kit.

## 6.10.3. NFS Server Recovery Kit Configuration Considerations

---

These following sections contain information to be considered before starting to configure and administer the NFS Server Recovery Kit as well as examples of typical LifeKeeper NFS configurations.

Please refer to [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring your LifeKeeper Core resource hierarchies.

- [Configuring NFS Server with LifeKeeper](#)
- [Specific Configuration Considerations](#)
- [Configuration Examples](#)

## 6.10.3.1. Configuring NFS Server with LifeKeeper

This section contains information to consider before starting to configure and administer the NFS Server Recovery Kit as well as examples of typical LifeKeeper NFS configurations.

Please refer to [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring your LifeKeeper Core resource hierarchies.

### NFS

The following table describes the NFS files, commands and daemons that are important to the NFS Server Recovery Kit:

NFS Component	Description
<i>exports(5) (/etc/exports)</i>	<p>Access control list for file systems exported to NFS clients. Each line of the file contains an export point, an optional list of clients that can mount the file system and an optional list of mount parameters.</p> <p><b>Note:</b> When you create a LifeKeeper-protected NFS resource, the export information for the file system is removed from the <i>exports</i> file and maintained under LifeKeeper. If you delete the NFS resource, the export information is restored to the <i>exports</i> file.</p>
<i>/var/lib/nfs</i>	Directory that contains NFS information on current exports, client mounts, locking status and more.
<i>/var/lib/nfs/etab</i>	<p>File that contains the current table of exported file systems for NFS. This file is maintained by the <b>exportfs</b> command; the user does not edit the file directly.</p> <p><b>Note:</b> When you bring an NFS resource into service on a backup server, the NFS file system is removed from the <i>etab</i> file on the primary server and inserted into the <i>etab</i> file on the backup server.</p>
<i>/var/lib/nfs/rpc_pipefs</i>	Used for kernel to userspace communication for NFS. This directory is relocated to <i>/var/lib</i> during installation of LifeKeeper.
<b>exportfs(8)</b> <i>(/usr/sbin/exportfs)</i>	Command used to maintain the table of exported file systems in <i>/var/lib/nfs/etab</i> .

<b>rpc.mountd(8)</b> <i>(/usr/sbin/ rpc.mountd)</i>	Daemon that authenticates a mount request and returns a filehandle if the client is permitted to mount the file system.
<b>rpc.nfsd(8)</b> <i>(/usr/sbin/ rpc.nfsd)</i>	Daemon that handles client file system requests.
<b>rpc.quotad(8)</b> <i>(/usr/sbin/ rpc.rquotad)</i>	The rpc server that returns quotas for a user of a local file system which is mounted remotely over NFS.
<b>rpc.lockd(8)</b> <i>(/sbin/rpc.lockd)</i>	Daemon that handles client file lock requests.
<b>rpc.statd(8)</b> <i>(/usr/sbin/ rpc.statd)</i>	Daemon that monitors the status of and makes status notifications for NFS clients and servers. This daemon must be running in order for NFS file locking to work properly.
<b>rpcbind</b>	Daemon process that converts RPC program numbers into port numbers and must be running for NFS. A failure of this process will force a switchover to a standby node. LifeKeeper also uses this for monitoring.
<b>rpc.idmapd</b>	NFS v4 ID to name mapper daemon process for translating user and group IDs to names and names to user and group IDs.

## Export Considerations

LifeKeeper protection for a given exported file system depends on the export options being exactly of the form as described in the `exports(5)` man page. In particular, pay attention to the host restriction format. There are only four legal host restrictions: (single host, netgroup, wildcard host `*name*` and netmask).

In particular, a wildcard IP address (like `172.13.4.*`) is not legal and will lead to potential stale filehandles on switchover or failover. Check very carefully by executing `exportfs -v` and manually comparing the returned export description against the format described in the man page (unfortunately, `exportfs`

doesn't check for you and will accept certain illegal export formats).

## Export option “fsid=0” is not supported

If “fsid=0” is specified as an export option, it is processed as a pseudo file system in NFS v4. LifeKeeper does not support this option. Do not specify “fsid=0” for the export point option to be protected by LifeKeeper. Also, LifeKeeper cannot protect the sub directory export point of the export point specifying fsid=0. Do not export by specifying “fsid=0” for the export which is not protected by LifeKeeper in order to avoid connection problems from a client.

## Bind mounts are not supported

Bind mounts are not supported by LifeKeeper and cannot be used for the export point.

## RPC.MOUNTD Restart

Under certain conditions with multiple NFS resource hierarchies, `rpc.mountd` fails to properly advertise the list of exports available. As such, the NFS Recovery Kit on a restore will stop and restart `rpc.mount` to ensure the proper list of exports is available to all clients. This action of stopping and restarting `rpc.mount` is controlled via the `RESTARTMOUNTD` entry in `/etc/default/LifeKeeper`. By default, this entry is set to true to cause the stop and restart of `_rpc.mount_` on all NFS restores:

```
RESTARTMOUNT=true
```

To turn off this action set:

```
RESTARTMOUNT=false
```

## NFS Resource Hierarchy

Create an IP address resource before creating an NFS resource.

When you create a LifeKeeper protected NFS resource, LifeKeeper creates the following hierarchy:

- NFS file system resource (parent or root)
  - IP resource
    - HA-NFS resource
      - File system resource (the underlying file system)

You have the option of creating the file system resource(s) before creating the NFS resource. If you do this, you can choose the name assigned to the file system resource(s). If not, the NFS Server Recovery Kit automatically creates the file system resource(s) when creating the NFS resource.

## Stopping and starting NFS subsystem adversely impacts LifeKeeper protected NFS exports

If the NFS subsystem is stopped while the NFS Server Recovery Kit is protecting NFS exports, then all protected exported directories will be impacted as the NFS stop action performs an un-export of all the directories. The quickCheck script will detect the stopped NFS processes and the un-exported directories, and run a local recovery to restart the processes and re-export the directories. However, you will need to run quickCheck for each protected export to recover everything. For example, if five exports are protected you will need to run quickCheck five times to recover all the exported directories the kit protects. Based on the default quickCheck time of two minutes, it could take up to ten minutes to recover all the exported directories. Do not stop the NFS subsystem while the NFS Server Recovery Kit is actively protecting exported directories on the system. If the NFS subsystem must be stopped, all NFS resources should be switched to the standby node before stopping the NFS subsystem. Use of the `exportfs` command should also be considered. This command line utility provides the ability to export and un-export a single directory thus bypassing the need to stop the entire NFS subsystem.

## 6.10.3.2. NFS Specific Configuration Considerations

---

The following should be considered before using the LifeKeeper NFS Server Recovery Kit:

- The NFS file system to be placed under LifeKeeper protection must be exported by the primary server (the server where the NFS resource is being created). This implies that NFS is running and the underlying file system is mounted.

 **Note:** If the */home* directory is shared via NFS, then */home* is the underlying file system.

- **When you export a read/write file system, use the *sync* option.** This option requests that all file system writes be committed to disk before the write request completes. **If the *sync* option is not used with an NFS file system under LifeKeeper protection, data may be lost during a failover.**
- The underlying file system must be on a shared device and mounted with write permission.
- If the underlying file system is already protected by LifeKeeper, it must be in service on the primary server and have the highest priority. If the underlying file system is not under LifeKeeper protection, then the Recovery Kit will place it under protection.
- The NFS Server Recovery Kit requires an IP resource that must be created and in service on the primary server. The IP resource must also have its highest priority on the primary server.
- Before creation of the NFS resource, clients must be able to mount the NFS file system using the LifeKeeper-protected IP address.
- When you extend an NFS file system resource, the file system must mount at the same mount point on each server.

 NFSv2 is not supported on RHEL 7/CentOS 7/OL 7 or later.

 NFS over UDP is not supported on RHEL 8 or later.

## 6.10.3.3. NFS Configuration Examples

---

The examples in this section show how NFS instances can be configured on shared (or replicated) disks. Each diagram shows the relationship between the type of configuration and the NFS parameters. Each configuration also adheres to the configuration rules and requirements described in this section that ensure compatibility between the NFS configuration and the LifeKeeper software.

The examples in this section are only a sample of the configurations than can be established, but understanding these configurations and adhering to the configuration rules will help define and set up workable solutions for your computing environment.

### Configuration Requirements

NFS Tag names are arbitrary names that describe protected file systems to LifeKeeper. The default tag name suggested by LifeKeeper is *"nfs-<export point>."*

To understand the configuration examples, keep in mind that the underlying file system must always be on shared or replicated disks. The file system(s) must be mountable from each of the systems.

### Examples

- [Active/Standby Configuration](#)
- [Active/Active Configuration](#)

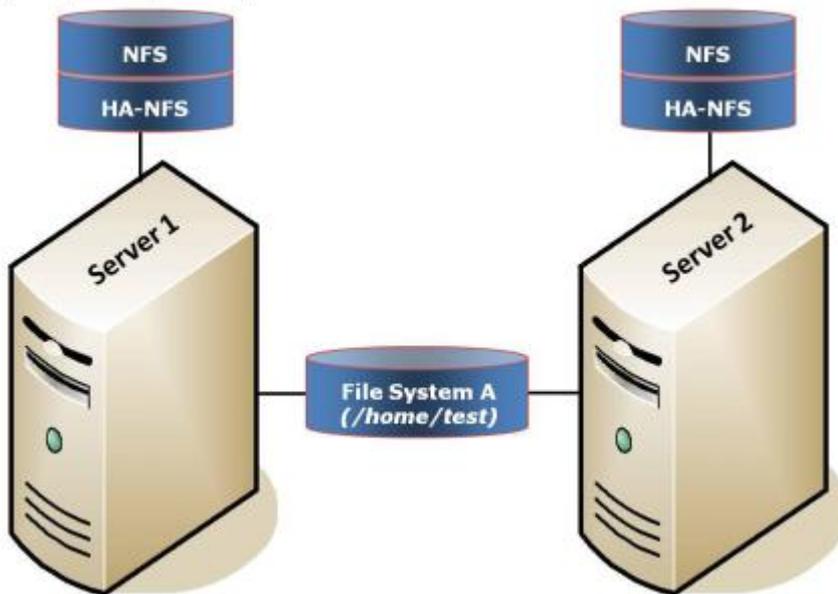
## 6.10.3.3.1. Active – Standby Configuration

This section provides an example of an **Active/Standby Configuration**. In this configuration, *Server 1* is considered active because it is running NFS and exporting the file system. If *Server 1* fails, *Server 2* mounts and exports the file system automatically.

**Note:** In an active/standby configuration, *Server 2* might be running NFS but does not have any other NFS resources under LifeKeeper protection.

### Active/Standby Configuration Example

Export point=*/home/test*



#### Configuration Notes:

- The NFS software must be installed on both servers.
- The underlying file system (*File System A*) must be on a shared (or replicated) disk.
- The NFS export point is */home/test*.
- The exported file system must have the same mount point on both *Server 1* and *Server 2*.
- *Server 2* cannot access files and directories on the shared disk while *Server 1* is active.

#### Creating a Resource Hierarchy to *Server 1*:

Server:	<i>Server1</i>
Export Point:	<i>/home/test</i>

IP Tag:	ip-172.17.100.202
NFS Tag:	<i>nfs-/home/test</i>

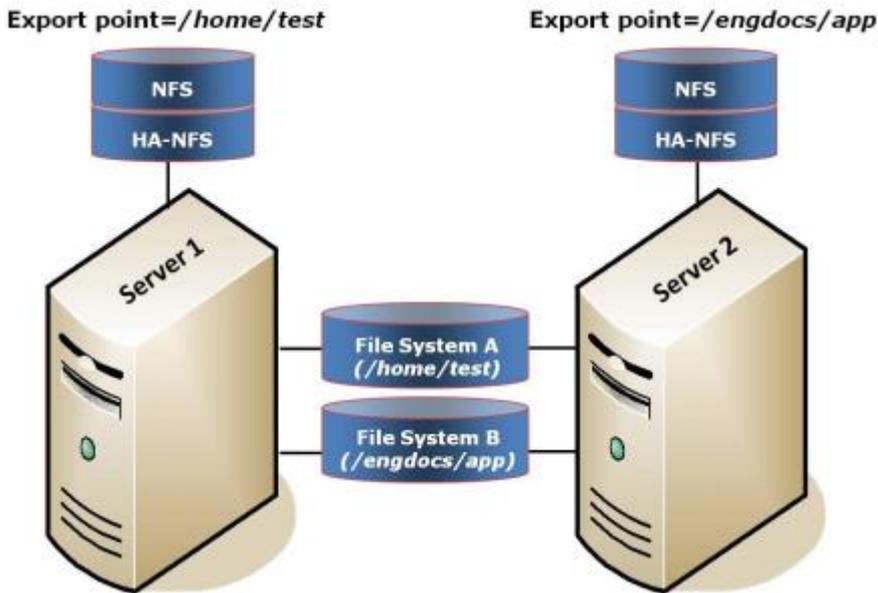
### Extending a Resource Hierarchy to *Server 2*:

Template Server:	<i>Server1</i>
Tag to Extend:	<i>nfs-/home/test</i>
Target Server:	<i>Server2</i>
Target Priority:	10

## 6.10.3.3.2. Active – Active Configuration

An example of **Active/Active** consists of two or more systems actively running NFS and exporting file systems.

### Active/Active Configuration Example



### Configuration Notes:

- The NFS software must be installed on both servers.
- Initially, *Server 1* exports */home/test* and *Server 2* exports */engdocs/app*. In a switchover situation, one system can export both file systems.
- File System A is the underlying file system for export point */home/test*. File System B is the underlying file system for export point */engdocs/app*.
- The underlying file systems are on different shared disks.

### Creating the First Resource Hierarchy on *Server 1*:

Server:	<i>Server1</i>
Export Point:	<i>/home/test</i>
IP tag:	<i>ip-172.17.100.202</i>
NFS Tag:	<i>nfs-/home/test</i>

**Extending the First Resource Hierarchy to Server 2:**

Template Server:	<i>Server1</i>
Tag to Extend:	<i>nfs-/home/test</i>
Target Server:	<i>Server2</i>
Target Priority:	10

**Creating the Second Resource Hierarchy on Server 2:**

Server:	<i>Server2</i>
Export Point:	<i>/engdocs/app</i>
IP Tag:	<i>ip-172.17.100.203</i>
NFS Tag:	<i>nfs-/engdocs/app</i>

**Extending the Second Resource Hierarchy to Server 1:**

Template Server:	<i>Server2</i>
Tag to Extend:	<i>nfs-/engdocs/ app</i>
Target Server:	<i>Server1</i>
Target Priority:	10

## 6.10.4. NFS Configuration Tasks

The following configuration tasks can be performed from the LifeKeeper GUI. These four tasks are described in this section as they are unique to an NFS Server resource instance and different for each Recovery Kit.

\* **Note:** Throughout this section, we explain how to configure your Recovery Kit by selecting certain tasks from the **Edit** menu of the LifeKeeper GUI. You can also select each configuration task from the toolbar. You may also right-click on a global resource in the **Resource Hierarchy Tree** (left-hand pane) of the status display window to display the same drop-down menu choices as the **Edit** menu. This, of course, is only an option when a hierarchy already exists. You can also right-click on a resource instance in the **Resource Hierarchy Table** (right-hand pane) of the status display window to perform all the configuration tasks, except **Creating a Resource Hierarchy**, depending on the state of the server and the particular resource.

- [Create a Resource Hierarchy](#). Creates an application resource hierarchy in your LifeKeeper cluster.
- [Delete a Resource Hierarchy](#). Deletes a resource hierarchy from all servers in your LifeKeeper cluster.
- [Extend a Resource Hierarchy](#). Extends a resource hierarchy from the primary server to a backup server.
- [Unextend a Resource Hierarchy](#). Unextends (removes) a resource hierarchy from a single server in the LifeKeeper cluster.
- [Test Your Resource Hierarchy](#). Tests your NFS resource hierarchy by initiating a manual switchover.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View](#) / [Edit](#) Properties. View or edit the properties of a resource hierarchy on a specific server.

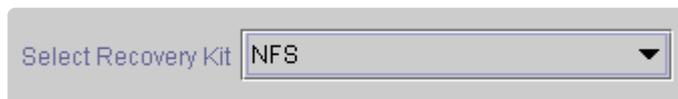
## 6.10.4.1. Creating an NFS Resource Hierarchy

To create a resource instance from the primary server, you should complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

To change a selection already entered or if an error message is encountered during any step in the creation of your NFS resource hierarchy, use the **Back** button to change your selection or make corrections (assuming the **Back** button is enabled).

A dialog box will appear with a drop-down menu listing all recognized Recovery Kits installed within the cluster. Select **NFS** from the drop-down menu.

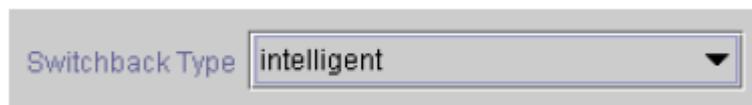


Select Recovery Kit

Click **Next** to proceed to the next dialog box.

**Note:** If you click the **Cancel** button at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the NFS instance will be switched back to the primary server when it comes back into service after a failover to the backup server. Choose either **Intelligent** or **Automatic**. **Intelligent switchback** requires administrative intervention to switch the instance back to the primary/original server. **Automatic switchback** means the switchback will occur as soon as the primary server comes back on line and is under LifeKeeper protection.

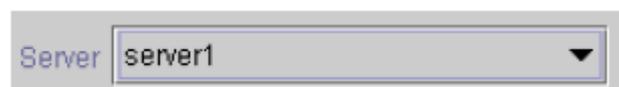


Switchback Type

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box.

Click **Next** to proceed to the next dialog box.

3. Select the **Server** where you want to create the NFS resource (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop-down menu.



Server

Click **Next** to proceed to the next dialog box.

4. The **Export Point** dialog displays a drop-down list of export points for NFS file systems that meet the following criteria:

- The export point has been exported by NFS.
- The export point is on a shared drive.
- If the underlying file system is LifeKeeper-protected, it must be in service and have the highest priority on the server selected on the **Server** dialog.
- NFSv4 criteria:
  - For NFS v4 root export with bind mounts, bind mounts must be on a shared drive just like the export, and if the file system is LifeKeeper-protected, it must be in service and have the highest priority on the server selected on the Server dialog.
  - If an NFS v4 root export is already being protected, no choices will be provided (there should only be one v4 and a mixture of V2/v3 with v4 cannot be protected).
  - If an NFS v2/v3 is already being protected, no NFS v4 will be listed in the choices.
  - If nothing is protected, then the list could contain both v2/v3 and v4.

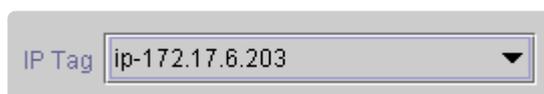
Select the NFS export point to be protected from the drop-down list.



Click **Next** to proceed to the next dialog box.

5. The **IP Tag** dialog displays a drop-down list of tags corresponding to virtual IP addresses currently under LifeKeeper protection and in service on the server where the NFS resource is being created.

Select the **tag** for the virtual IP address used by clients to access the protected NFS file system.



 **Note:** At this point, LifeKeeper will check to ensure that there is a protected IP resource available. It will also validate that you have provided valid data to create your NFS resource hierarchy. If LifeKeeper detects a problem with either of these validations, an ERROR box will appear on the screen. If the directory paths are valid but there are errors with the NFS configuration itself, you may pause to correct these errors and continue with the hierarchy creation. You may even pause to create any LifeKeeper IP resources that are required.

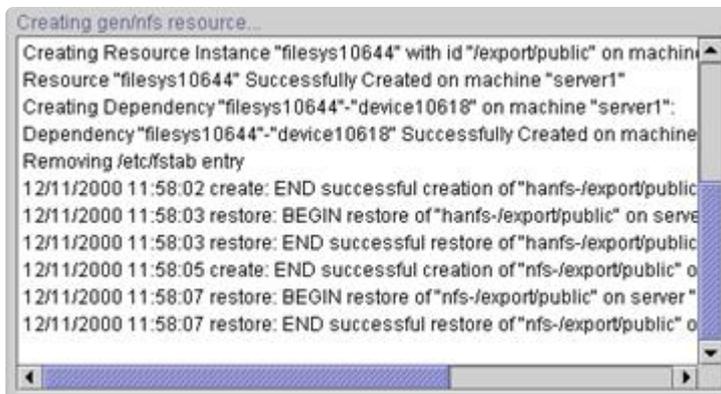
- \* **Note:** If you are using other LifeKeeper Recovery Kits that have virtual IP address dependencies, you might want to create a different virtual IP address for the NFS resource. Otherwise, if the virtual IP resource fails over to a backup server, all of the resources that depend on that IP resource will fail over at the same time.

Click **Next** to proceed to the next dialog box.

6. Select or enter the **NFS Tag**. This is a tag name given to the NFS hierarchy. You can select the default or enter your own tag name.

NFS Tag

When you click the **Create** button, the **Create Resource Wizard** will create your NFS resource.



When you click **Next**, LifeKeeper will launch the **Pre-Extend Wizard** that is discussed in [Extending Your Hierarchy](#)

- \* **Note:** The NFS resource hierarchy should be created successfully at this point. However, error messages may be encountered indicating that the new NFS instance has failed to start correctly. Note that the new NFS hierarchy must be started (In Service) before it can be extended to another system. A failure to start may remove the hierarchy, but if not, you may pause at this point and correct the problem based on the error message displayed. If the errors are not correctable, you will only be given the choice to cancel which cancels the resource create.

Bring the new hierarchy In Service before proceeding with [extending your hierarchy](#).

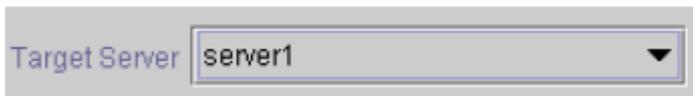
- \* **Note:** Please disable automatic startup of `nfs-server.service` after creating NFS resources on RHEL 7.1 or later and SLES12 SP1 or later. Since it is necessary for `rpcbind.service` to be running at the startup of NFS resources, please configure `rpcbind.service` to start automatically.

## 6.10.4.2. Deleting an NFS Resource Hierarchy

To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop-down menu, select **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your NFS resource hierarchy.

- \* **Note:** This dialog box does not appear if you select the **Delete Resource** task by right-clicking from either of the following:
- The right pane on an individual resource instance
  - The left pane on a global resource when the resource is on only one server



Click **Next** to proceed to the next dialog box.

3. Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete and highlight it.

- \* **Note:** If you selected the Delete Resource task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to proceed to the next dialog box.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete.

You have specified the following resource hierarchy for deletion.  
Target Server: server1  
Target Tags:  
nfs-/export/public

Click **Delete** to proceed to the next dialog box.

5. Another information box appears confirming that the NFS resource was deleted successfully.



6. Click **Done** to exit out of the **Delete Resource Hierarchy** menu selection.

## 6.10.4.3. Extending Your NFS Hierarchy

After you have created a hierarchy, you will want to extend that hierarchy to another server in the cluster. There are two possible scenarios to extend your resource instance from the template server to a target server. The first scenario is when you “Continue” from creating the resource into extending that resource to another server. The other scenario is when you enter the Extend Resource Hierarchy task from the edit menu as shown below. Both scenarios take you through the same dialog boxes (with a few exceptions, which are detailed below).

\* **Note:** If you click the **Cancel** button at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extension process to that particular server. LifeKeeper will also unextend any dependent resources in the hierarchy (IP address or file system) that are currently extended past the cancellation point. However, if you have already extended the NFS resource hierarchy to another server, that instance will continue to be in effect until you specifically unextend it. For example, you have created your resource on Server 1 and extended that resource to Server 2 and in the middle of extending the same resource to Server 3, you change your mind and click **Cancel** inside one of the dialog boxes. This will cancel only your action to extend the resource to Server 3, not the extension you created to Server 2. If you want to remove Server 2 from this hierarchy, you must unextend the resource from Server 2.

1. If you are entering the **Extend Wizard** from the **LifeKeeper GUI** menu, select **Edit**, then **Resource**. From the drop-down menu, select **Extend Resource Hierarchy**. This will launch the **Extend Resource Hierarchy** wizard.
2. The first dialog box to appear will ask you select the **Template Server** where your NFS resource hierarchy is currently in service. It is important to remember that the Template Server you select now and the **Tag to Extend** that you select in the next dialog box represent an in service resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you selected. The drop-down box in this dialog provides the names of all the servers in your cluster.

\* **Note:** If you are entering the **Extend Resource Hierarchy** task immediately following the creation of a NFS resource hierarchy, this dialog box will not appear, since the wizard has already identified the template server in the create stage. This is also the case when you right click on either the NFS resource icon in the left hand pane or right-click on the NFS resource box in the right hand pane on the of the GUI window and choose **Extend Resource Hierarchy**.



The image shows a graphical user interface element for a dialog box. It consists of a light gray rectangular container. On the left side, the text 'Template Server' is displayed in a dark gray font. To the right of this text is a white rectangular dropdown menu with a thin gray border. Inside the dropdown menu, the text 'server1' is visible. A small black downward-pointing triangle is located at the bottom right corner of the dropdown menu, indicating it is a list box.

Click **Next** to proceed to the next dialog box.

3. Select the **Tag to Extend**. This is the name of the NFS instance you wish to extend from the

template server to the target server. The wizard will list in the drop down menu all the resources that you have created on the template server, which you selected in the previous dialog box.

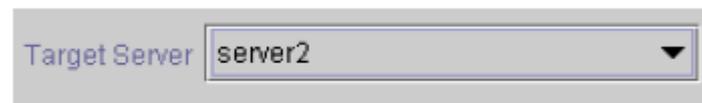
\* **Note:** Once again, if you are entering the **Extend Resource Hierarchy** task immediately following the creation of an NFS resource hierarchy, this dialog box will not appear, since the wizard has already identified the tag name of your NFS resource in the create stage. This is also the case when you right-click on either the NFS resource icon in the left hand pane or on the NFS resource box in the right hand pane of the GUI window and choose **Extend Resource Hierarchy**.



Tag to Extend

Click **Next** to proceed to the next dialog box.

4. Select the **Target Server** where you are extending your NFS resource hierarchy. The drop-down box provides the names of the servers in your cluster that are not already in the selected hierarchy.



Target Server

Click **Next** to proceed to the next dialog box.

5. Select the **Switchback Type**. This dictates how the NFS instance will be switched back to the primary server when it comes back into service after a failover to the backup server. You can choose either intelligent or automatic. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. Automatic switchback means the switchback will occur as soon as the primary server comes back on line and is under LifeKeeper protection.



Switchback Type

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box.

Click **Next** to proceed to the next dialog box.

6. Select or enter a **Template Priority**. This is the priority for the NFS hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. Note: This selection will appear only for the initial extend of the hierarchy.

Click **Next**.

7. Select or enter the **Target Priority**. This is the priority for the new extended NFS hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.



Target Priority

Click **Next**.

8. An information box will appear explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this NFS resource have been met. If there were some requirements that had not been met, LifeKeeper would not allow you to select Next, and the **Back** button would be enabled.

```
Executing the pre-extend script...
Checking existence of extend and canextend scripts
Building independent resource list
Checking extendability for nfs-/export/public
Checking extendability for ip-172.17.6.203
Pre Extend checks were successful
```

If you click **Back**, you can make changes to your resource extension according to any error messages that may appear in the information box.

If you click **Cancel** now, you will need to come back and extend your NFS resource hierarchy to another server at some other time to put it under LifeKeeper protection.

When you click **Next**, LifeKeeper will launch you into the **Extend Resource Hierarchy** configuration task and the **NFS Tag** dialog box will display.

9. This screen provides information about the **Template Server**, **Tag to Extend**, **Target Server** and the default **NFS Tag**. The **NFS Tag** is a tag name given to the NFS hierarchy extension. You can select the default or enter your own tag name.



NFS Tag

Click **Next** to proceed to the next dialog box.

10. An information box will appear verifying that the extension is being performed.

\* **Note:** If you have not already extended the IP resource to the target server, the NFS Server Recovery Kit extends it in the process of extending the NFS resource. Before displaying the extension verification information box, the Recovery Kit displays several additional dialog boxes related to the extension of the IP resource.

```
Extending resource hierarchy...
Extending resource instances for nfs-/export/public
Extending resource instances for ip-172.17.6.203
Creating dependencies
Setting switchback type for hierarchy
Creating equivalencies
LifeKeeper Admin Lock (nfs-/export/public) Released
Hierarchy successfully extended
```

Click **Next Server** if you want to extend the same NFS resource instance to another server in your cluster. This will repeat the Extend Resource Hierarchy operation.

If you click **Finish**, LifeKeeper will verify that the extension of the NFS resource was completed successfully.

11. If you click **Finish**, the following screen appears.

```
Verifying Integrity of Extended Hierarchy...
Examining hierarchy on server2
Hierarchy Verification Finished
```

12. Click **Done** to exit.

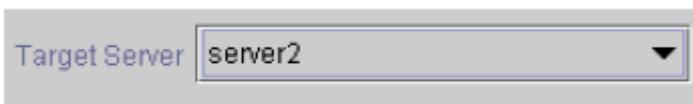
\* **Note:** Be sure to test the functionality of the new instance on both servers.

## 6.10.4.4. Unextending Your NFS Hierarchy

Perform the following steps to unextend a resource hierarchy:

1. From the **LifeKeeper GUI menu**, select **Edit** and **Resource**. From the drop-down menu, select **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the NFS resource. It cannot be the server where the NFS resource is currently in service.

\* **Note:** If you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance this dialog box will not appear.



Target Server server2 ▼

Click **Next** to proceed to the next dialog box.

3. Select the **NFS Hierarchy to Unextend**.

\* **Note:** If you selected the **Unextend** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Hierarchy to Unextend nfs-/export/public ▼

Click **Next** to proceed to the next dialog box.

4. An information box appears confirming the target server and the NFS resource hierarchy you have chosen to unextend.



You have specified the following resource hierarchy for unextend.  
Target Server = server2  
Target Tag = nfs-/export/public

Click **Unextend**.

5. Another information box appears confirming that the NFS resource was unextended successfully.



6. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection and return to the LifeKeeper GUI.

You will receive the warning **One or More Resources Unprotected** if the hierarchy is unextended down to one server.

## 6.10.4.5. Testing Your NFS Hierarchy

Before testing your NFS resource hierarchy, you should validate your client setup as described below. You can then test your NFS resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

### Validating the Client Setup

In general, clients must mount the file system using the LifeKeeper-protected IP address you selected during the Create NFS Resource Hierarchy task. There is no client-side checking to ensure that you select the correct IP address, so you must carefully follow the validation steps below to ensure the client is using the correct IP number for the file system.

To validate the client setup, do the following:

1. Verify that no NFS instances are in service on the backup server.
2. Mount the file system on the client using the correct LifeKeeper-protected IP address.
3. Perform a manual switchover to the backup server and ensure that the NFS instance you just switched over is the only NFS instance currently in service on the backup server.
4. When the switchover has completed, ensure that the client can still access the file system.

### Performing a Manual Switchover from the GUI

After you define the dependencies, LifeKeeper automatically controls the starting and stopping of the application whenever it detects faults, which initiate failover recovery. You can also manually initiate a switchover for administrative reasons, such as maintenance.

You can initiate a manual switchover from the LifeKeeper GUI by selecting Edit, Resource, and In Service from the drop-down menu. For example, an in service request executed on a backup server causes the application hierarchy to be placed in service on the backup server and taken out of service on the primary server.

If you execute the Out of Service request, the application is taken out of service without bringing it in service on the other server.

 **Note:** To take an NFS resource out of service, you must take both the NFS resource and the associated HA-NFS resource out of service.

For activities within the application, all actions are those defined in the application's documentation. LifeKeeper does not regulate or control internal operations such as rollbacks and backing-up archives. Tape archiving and restoration are the responsibility of the application administrator.

## Recovery Operations

When the primary server fails, the NFS Server Recovery Kit software performs the following tasks:

- Starts the NFS daemons if they are not running.
- Exports the NFS file system.

## 6.10.5. NFS Troubleshooting

---

This section provides a list of messages that you may encounter while creating and extending a LifeKeeper NFS resource hierarchy or removing and restoring a resource. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Messages from other LifeKeeper components are also possible. In these cases, please refer to the [Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

Click the following topics for Troubleshooting help.

- [HA nfs-utils Installation and Configuration](#)
- [NFS Hierarchy Creation Errors](#)
- [NFS Extend Hierarchy Errors](#)
- [NFS Hierarchy Restore, Remove and Recover Messages and Errors](#)
- [NFS Hierarchy Delete Messages and Errors](#)

## 6.10.5.1. HA nfs-utils Installation and Configuration

---

LifeKeeper NFS Server Recovery Kit requires the installation and configuration of a high availability enabled `nfs-utils` package (on some OS distribution versions, `nfs-utils` is provided via another package). The Recovery Kit will attempt to verify the presence of this HA enabled `nfs-utils` package. If it fails to detect a correctly configured `nfs-utils` package, the LifeKeeper Installation Support setup script may need to be rerun or the server may need to be rebooted.

The LifeKeeper Installation Support setup script will configure NFS for use in an HA environment. Restarting NFS service is required to reflect the configuration changes. Therefore, you need to select “Restart NFS Service” in the setup script or restart the NFS service manually after the installation is completed.

The configuration needed for **NFS v4** requires the movement of `rpc_pipefs` from `/var/lib/nfs` to `/var/lib`. To do this may require the unloading of kernel modules and the addition or modification of configuration and boot time scripts. A system reboot may be required if Installation Support is unable to unload and reload kernel modules after the change. If this should occur, the user will be notified of the need for a system reboot. Completing the `rpc_pipefs` setup including a system reboot is required for successful operation of LifeKeeper.

## 6.10.5.2. NFS Hierarchy Creation Errors

The error messages that might be displayed during the NFS hierarchy creation are listed below, along with a suggested explanation for each. The messages listed cover both the creation of the nfs and hanfs resources. Error messages displayed by the LifeKeeper core and by other recovery kits are not listed in this guide. Note that you may stop to correct any problem described here, and then continue with hierarchy creation from the point where you left off – including creating any new LifeKeeper resources you might need for your NFS configuration.

 **Note:** In the following error messages, *Command line* only indicates that you can only receive the message if you are entering commands on the command line; you cannot receive it if you are using the LifeKeeper GUI. Additionally, at the end of hierarchy create a resource restore is initiated. See *Hierarchy Restore, Remove and Recover Messages and Errors* for an explanation of messages and errors that can occur during that process.

Error Number	Error Message	Description
106000	Export point not specified	You must specify the export point for the NFS file system when you create the resource hierarchy. <i>Command line only.</i>
106001	The path "EXPORT POINT" is not exported by NFS	The export point you specified is not currently exported by NFS. Use <b>exportfs(8)</b> to export the path and verify the path is in the <i>/var/lib/nfs/etabfile</i> . <i>Command line only.</i>
106002	create: The export point "EXPORTPOINT" on server "SERVER" for client "CLIENT" either does not contain an FSID export option or the value is not unique. A unique FSID will be generated and "EXPORTPOINT" will be re-exported using the new FSID value.	All export points under LifeKeeper protection must use a unique fsid= export option for high availability NFS. The selected export did not meet this requirement so a unique value was generated followed by a re-export for the selected export point. Note: a client of "*" or "world" indicates the export point is available to all clients.
106003	Unable to create the export entry file for "EXPORT POINT" in LifeKeeper	<ul style="list-style-type: none"> <li>The file system is full on the target server.</li> <li>File system problems.</li> </ul>
106004	The export point "EXPORT POINT" is not on a shared file system on server "SERVER"	Make sure that the export point is for a shared file system. <i>Command line only.</i>
106005	Unable to create the HA-NFS hierarchy "TAG" with child resource ID "ID" on server "SERVER"	Review the other error messages to determine the action to take.
106006	Unable to remove entry for export point "EXPORT POINT" from <i>/etc/exports</i>	Verify that the <i>/etc/exports</i> file exists and is readable.

	server "SERVER"	
106014	Usage: USAGE STRING	Usage of command run with incorrect arguments. <i>Command line only.</i>
106016	"REQUIRED SOFTWARE" cannot be found or does not have the expected permissions on server "SERVER"	NFS must be installed on the primary server and all backup servers. Verify that the <b>nfs-utils</b> has been installed
106019	Executing command: "COMMAND"	This message is displayed when LifeKeeper restarts an NFS daemon or exports/unexports an export point. It provides additional information that can be useful if there is a problem.
106020	The generated id "ID" conflicts with an existing resource id	The internally generated resource ID for the nfs or hanfs resource has produced a duplicate.
106025	An unknown error has occurred while running "NEW TAG" on server "SERVER"	An unexpected error occurred while running the command newtag to generate a tag for the nfs for hanfs resource.
106026	Adding dependency between HA-NFS resource "HANFS TAG" and filesys resource "FILESYS TAG" on server "SERVER" failed.	Dependency creation between the selected hanfs resource and the filesys resource has failed for unknown reasons. See output for more information.
106027	Open of "ABC" on server "SERVER" failed: "File not found"	The attempted open failed for the reason listed.
106029	Unable to get exclusive lock on "FILE NAME" on server "SERVER"	Unable to obtain an exclusive lock for update after 20 seconds. This indicates a problem with the file.
106031	Re-export of "EXPORT POINT" to add FSID option failed on server "SERVER"	The export point did not contain a fsid argument and the re-export after one was generated failed. Manually add a fsid argument to the /etc/exports entry and re-export to determine the failure.
106032	Dependent IP resource tag name not specified	You must specify the resource tag for the protected IP address when you create the NFS resource hierarchy. Command line only.
106033	Selected IP resource "TAG" does not exist on server "SERVER"	You must create the IP resource on the specified server before you can create the NFS resource. Also, make sure that you typed the IP resource correctly when you entered the command. Command line only.
106034	Adding dependency between NFS resource "TAG" and IP resource "TAG" on server "SERVER" failed	Verify the IP resource is in-service on the server where the NFS resource is being created. Command line only.
106035	Creation of HA-NFS resource "TAG" on server "SERVER" failed	Review other error messages to determine the action to take.
106036	Adding dependency between IP resource	Verify that the IP resource is in-service on the server

	“TAG” and HA-NFS resource “TAG” on server “SERVER” failed	where the HA-NFS resource is being created. Command line only.
106037	Attempts to get exclusive lock on “FILE NAME” on server “SERVER” failed: “ERROR MSG”	Unable to obtain an exclusive lock for updating. See the error message for the cause.
106038	Unable to create directory “DIR NAME” on server “SERVER”: “ERROR MSG”	An attempt to create a directory on the exported file system has failed. See the error message for the cause.
106039	Open of “FILE” on server “SERVER” failed: “ERROR MSG” or Attempt to get exclusive lock on “FILE” on server “SERVER” failed: “ERROR MSG”	An attempt to open or obtain an exclusive lock on a file has failed. See the error message for the cause.
106041	The selected IP resource “IP TAG” is not ISP on server “SERVER”	The selected IP resource does exist on the server but is not currently in service. Bring the IP resource in service on the server and then re-attempt the creation. Command line only.
106048	Multiple NFS v4 root exports found on “SERVER”.	Multiple NFS v4 psuedo file systems found on SERVER where only one is supported.
106050	Unable to protect more than 1 NFS v4 export or a combination of NFS v4 and NFS v3 exports.	Attempting to protect a mix of NFS v2/v3 exports with NFS v4 which is not supported.

## 6.10.5.3. NFS Extend Hierarchy Errors

The error messages that might be displayed during NFS hierarchy extension are listed below, along with a suggested explanation for each. Note that these error messages appear when the GUI indicates it is “Executing the pre-extend script...” to validate the hierarchy prior to extending it to the new system.

### During NFS Resource Hierarchy Creation on Target Server

Error Number	Error Message	Description
106016	“REQUIRED SOFTWARE” cannot be found or does not have the expected permissions on server “SERVER”	NFS must be installed on the primary server and all backup servers. Verify that the <i>nfs-utils</i> has been installed.
106017	The file system “FILE SYSTEM” on template server “SERVER” has a different mount point “MOUNT POINT” on server “SERVER”	The resources must be created with the same mount point on each server. Either unextend the file system hierarchy from the target server or recreate it with the same mount point on the template and target servers.
106018	Unable to copy the file “FILENAME” from server “SERVER” to server “SERVER”	Possible causes: <ul style="list-style-type: none"> <li>• Communication path between the servers is down</li> <li>• File system is full of the target server</li> <li>• File system problems</li> </ul>
106020	The generated id “ID” conflicts with an existing resource id	The internally generated resource ID for the nfs or hanfs resource has produced a duplicate.
106022	The export point “EXPORT POINT” is in <i>/etc/exports</i> on the target server “SERVER”	Remove the export point from the <i>/etc/exports</i> file on the target server before trying to extend the resource.
106023	The export point “EXPORT POINT” is exported on the target server “SERVER”	Unexport the export point on the target server before trying to extend the resource.
106051	Unable to create active/active configurations with NFS v4 exports. Either “TEMPLATE SERVER” or “TARGET SERVER” currently protects an NFS v4 root export.	Unable to extend the NFS resource from the TEMPLATE SERVER to the TARGET SERVER as one or both of the servers already protects an NFS v4 export and active/active configurations with NFS v4 exports is not supported.

## 6.10.5.4. NFS Hierarchy Restore, Remove and Recover Messages and Errors

### Bringing an NFS Resource In-Service (Restore)

Error Number	Error Message	Description
106007	Cannot bring NFS or HANFS resource "TAG" in service on server "SERVER"	Review other error messages to determine the action to take. After correcting the problem, try bringing the resource in service manually.
106010	NFS is not running on server "SERVER". LifeKeeper will attempt to restart NFS.	This message is for information only. LifeKeeper will try to restart the NFS daemons automatically. If LifeKeeper encounters problems while restarting one of the daemons, you will receive a message that starting NFS failed.
106011	Starting NFS on server "SERVER" failed	There was a problem while restarting the NFS daemons. Try manually restarting NFS.
106012	The export point "EXPORT POINT" is not exported on server "SERVER". LifeKeeper will attempt to export the entry.	LifeKeeper has detected that the export point is no longer exported, and will try to export it.
106013	Unable to export "EXPORT POINT" on server "SERVER"	Try manually exporting the file system.
106014	Usage: USAGE STRING	Usage of command run with incorrect arguments. <i>Command line only.</i>
106019	Executing command: "COMMAND"	This message is displayed when LifeKeeper restarts an NFS daemon or exports/unexports an export point. It provides additional information that can be useful if there is a problem.
106024	Unable to stop and restart rpc.mountd on "SERVER"	During a hierarchy restore the rpc.mountd daemon process needed to be restarted and this process failed. Manually attempt to stop and restart the process to determine the error and the action to take.
106027	Open of "FILE NAME" on server "SERVER" failed: "ERROR MSG"	The attempted open failed for the reason listed.
106028	Mount of /proc/fs/nfsd failed on server "SERVER"	In 2.6 and later kernels /proc/fs/nfsd is used for client authentication and an attempt to mount it failed. Manually attempt to mount /proc/fs/nfsd to determine the failure.

106029	Unable to get exclusive lock on "FILE NAME" on server "SERVER"	Unable to obtain an exclusive lock for file update after 20 seconds. This indicates a problem with the file.
106030	Unable to restore client info for "CLIENT" on server "SERVER": "ERROR MSG"	Client lock failover failed. Correct the failure condition and attempt to restore the hierarchy again.
106037	Attempts to get exclusive lock on "FILE NAME" on server "SERVER" failed: "ERROR MSG"	Unable to obtain an exclusive lock for updating. See the error message for the cause.
106039	Open of "FILE" on server "SERVER" failed: "ERROR MSG" or Attempt to get exclusive lock on "FILE" on server "SERVER" failed: "ERROR MSG"	An attempt to open or obtain an exclusive lock on a file has failed. See the error message for the cause.
106040	Multiple virtual IP addresses detected. In this release NFS lock failover only supports one virtual IP address.	Recreate the NFS resource hierarchies to use only one virtual IP address or set FAILOVERNFSLOCKS to false in the LifeKeeper defaults file.
106052	Unable to mount <i>rpc_pipefs</i> on "SERVER". Reason: "REASON".	<i>rpc_pipefs</i> was not mounted on SERVER and the mount attempt failed for REASON.
106053	<i>rpc_pipefs</i> successfully mounted on "SERVER"	<i>rpc_pipefs</i> was successfully mounted on SERVER.
106064	Pseudo file system (fsid=0) protected by NFS RK is not supported in LifeKeeper v9.6.0 and later. Please delete the "TAG" resource whose export point is "EXPORT POINT".	Pseudo file system is not supported in LifeKeeper v9.6.0 and later. Please delete the resource hierarchy.

## Taking an NFS Resource Out of Service (Remove)

Error Number	Error Message	Description
106008	Unable to unexport the export point "EXPORT POINT" on server "SERVER"	Use the <b>exportfs(8)</b> command to unexport it.
106014	Usage: USAGE STRING	Usage of command run with incorrect arguments. <i>Command line only.</i>
106019	Executing command: "COMMAND"	This message is displayed when LifeKeeper restarts an NFS daemon or exports/unexports an export point. It provides additional information that can be useful if there is a problem.

## Bringing an NFS Resource Back In Service (Recover)

The LifeKeeper core periodically checks the health of every NFS instance In Service on the local server by running an NFS “quickCheck” script. This script verifies the following:

- The file system is exported
- The NFS/HA-NFS daemons are running

If the instance is not fully functional, a "recover" script is invoked to attempt to restart the instance. This simply logs an error message, invokes "restore," prints the final error or success message – depending on error or success of the "restore" script – and returns the same result as "restore." If restore/recover fails, this instance is failed over to another server.

## 6.10.5.5. NFS Hierarchy Delete Messages and Errors

---

Error Number	Error Message	Description
106015	Unable to restore the entry for export point "EXPORT POINT" in <i>/etc/exports</i> on server "SERVER"	Restore the entry manually.
106021	An entry for export point "EXPORT POINT" already exists in <i>/etc/exports</i> . The entry that was being used by the NFS Server Recovery Kit has been placed in the file "FILENAME"	Verify that <i>/etc/exports</i> has the correct export entry.
106049	Restore <i>statedir</i> from <i>/var/lib/.nfs.LK</i> to <i>/var/lib/nfs</i> failed on server SERVER.	Restoring the NFS state directory <i>/var/lib/nfs</i> failed. Try manually restoring the directory by moving <i>/var/lib/.nfs.LK</i> to <i>/var/lib/nfs</i> .

# 6.11. Recovery Kit for Oracle Cloud Infrastructure Administration Guide

---

## Recovery Kit for Oracle Cloud Infrastructure (RK for OCI)

The RK for OCI provides a mechanism for using virtual IP addresses with the IP Recovery Kit on Oracle Cloud Infrastructure (OCI).

For more information on the functionalities, scenarios and operations of the RK for OCI, see [Principles of Operation](#).

### LifeKeeper for Linux Documentation

The following is a list of documents related to LifeKeeper for Linux provided by SIOS Technology Corp.

- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Release Notes](#)
- [SIOS Technical Documentation](#)

\* “Oracle Cloud” is a trademark of Oracle Corporation or its affiliates in the United States and/or other countries. Trademark symbols such as ® and ™ may be omitted from system names and product names in this document.

## 6.11.1. Principles of Operation

---

The Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) manages the assignment of secondary private IP addresses to Virtual Network Interface Cards (VNICs) so that clients on the OCI can connect to the LifeKeeper-protected IP resources (virtual IP addresses).

<https://docs.oracle.com/en-us/iaas/Content/Network/Tasks/managingIPaddresses.htm>

To use a virtual IP address on OCI, you need to assign the target IP address to the VNIC of the active node in addition to assigning the IP address on the OS.

In LifeKeeper, the IP Recovery Kit assigns IP addresses on the OS.

The RK for OCI runs the OCI CLI on the cluster node and manages the assignment of secondary private IP addresses to the VNICs of the active node on the OCI.

The RK for OCI also monitors whether the protected IP address is assigned to the specified VNIC on the active node.

When the active node fails, the RK for OCI assigns an IP address to the VNIC of the standby node and makes the virtual IP address available.

For more details, please refer to [Resource Monitoring and Recovery](#).

## 6.11.2. Resource Monitoring and Local Recovery

### Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) Quick Check

The RK for OCI runs the following command to verify that the secondary private IP address is assigned to the VNIC on the active node.

```
oci network private-ip list --vnic-id <vnicid> --query 'data[?"is-primary" == `false`]."ip-address"'
```

Successful case:

- The protected IP address is included in the return value.

Unsuccessful case:

- Failed to execute the `oci` command.
- The protected IP address is not included in the return value.

If the quick check of the RK for OCI fails, a local recovery of the RK for OCI is performed.

### Local Recovery of RK for OCI

When the quick check of the RK for OCI fails, a local recovery of the RK for OCI is performed.

For local recovery, assign a secondary private IP address to the VNIC of the active node by executing the following command:

```
oci network vnic assign-private-ip --vnic-id <vnicid> --ip-address <ip>
```

If local recovery fails due to a failed `oci` command or because the IP address is already assigned to another instance, LifeKeeper will failover the RK for OCI resources and all resources with dependencies to the standby node. When bringing the resources in service on the standby node during failover, the following command:

```
oci network vnic assign-private-ip --unassign-if-already-assigned --vnic-id <vnicid> --ip-address <ip>
```

is executed to un-assign the IP address. Therefore, the failover will succeed even if the IP address is assigned to an instance outside the cluster node.

## 6.11.3. Requirements

### Oracle Cloud Infrastructure and Software Requirements

Before installing and configuring the Recovery Kit for Oracle Cloud Infrastructure (RK for OCI), be sure that your environment meets the following requirements.

#### Oracle Cloud Infrastructure Compute Instance

- Must be a bare metal instance or virtual machine instance provided by OCI.
- The [Command Line Interface \(CLI\)](#) 3.0.0 or later provided by OCI is available on the instance used as a cluster node.
- The [Instance Metadata Service v1 \(IMDSv1\)](#) provided by OCI is available on the instance used as a cluster node.
  - IDMSv2 is not supported.

#### Oracle Cloud Infrastructure Command Line Interface (CLI)

In order for the RK for OCI to operate OCI, configure the CLI and permissions so that the root user of the cluster node can execute the following commands. Please refer to the OCI documentation for the CLI and permission settings.

```
oci network vnic assign-private-ip --vnic-id <vnicid> --ip-address <ip>
oci network vnic unassign-private-ip --vnic-id <vnicid> --ip-address <ip>
oci network private-ip list --vnic-id <vnicid>
oci network vnic get --vnic-id <vnicid>
oci network private-ip list --subnet-id <subnetid>
```

- \* **<ip>** is the IP address protected by LifeKeeper.  
**<vnicid>** is the OCID of the virtual network interface card (VNIC) to which the IP address is assigned.  
**<subnetid>** is the OCID of the subnet to which the VNIC is related.  
Add the path to the CLI executable file to the parameter PATH in the configuration file */etc/default/LifeKeeper* .

#### IP address to use

- The private IP address to be protected by LifeKeeper RK for OCI must be IPv4.
  - IPv6 is not supported.
- The private IP address to be protected by the LifeKeeper RK for OCI must be an IP address that can be assigned to a VNIC.
  - It is the IP address in the subnet to which the VNIC to be assigned is related.
  - It is not an IP address reserved by OCI.
  - The number of secondary private IP addresses available for assigning has not been exceeded.

## LifeKeeper Software

You must install the same version of LifeKeeper software and any patches on each server.

Please refer to the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.

## LifeKeeper Recovery Kit for Oracle Cloud Infrastructure

You must install the same version of the RK for OCI software and any patches on each server.

## LifeKeeper IP Recovery Kit

You must install the same version of the LifeKeeper for Linux IP Recovery Kit software and any patches on each server.



Please refer to the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information. You should refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper RK for OCI.

SIOS **recommends** using Quorum/Witness when using the RK for OCI. Please refer to [Quorum/Witness](#) for more information.

## 6.11.4. Recovery Kit for Oracle Cloud Infrastructure Notes

---

When you create a Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) resource, the corresponding IP resource is created.

If an IP resource already exists before the RK for OCI resource is created:

- It matches the configuration of the RK for OCI resource and a dependency with the existing IP resource is automatically created.
- It does not match the configurations of the RK for OCI resource, i.e., the device or netmask values do not match, no dependency is created between the IP resource and the RK for OCI resource. Delete the IP resource and then recreate the RK for OCI resource, or modify the IP resource configurations appropriately before creating the dependency.
- You need to disable the broadcast PING check for LifeKeeper IP resources. Configure the NOBCASTPING entry in the */etc/default/LifeKeeper* configuration file as follows:

```
NOBCASTPING=1
```

## 6.11.5. Restrictions when using LifeKeeper on Oracle Cloud Infrastructure (OCI)

---

### Oracle Cloud Infrastructure (OCI) Support Configuration

Oracle provides images for the following Operating Systems support by LifeKeeper.

- Oracle Linux 7 (UEK, RHCK)
- Oracle Linux 8 (UEK, RHCK)
- CentOS 7

The OS configured in BYOI is the same as the OS supported by LifeKeeper (e.g., RHEL). The following Operating Systems are not supported.

- SUSE Linux Enterprise Server (SLES)
- CentOS 8

The following Recovery Kits are not supported at the time of v9.6.1 release.

- WebSphere MQ Recovery Kit
- SAP Recovery Kit
- SAP HANA Recovery Kit
- SAP MaxDB Recovery Kit
- Sybase Recovery Kit

The following Recovery Kits are not available on OCI.

- DB2 Recovery Kit
- Recovery Kit for EC2
- Recovery Kit for Route53
- VMDK as Shared Storage Recovery Kit
- Multipath Recovery Kits



Refer to the [Requirement](#) and [Restrictions](#) pages for Recovery Kit for Oracle Cloud Infrastructure restrictions.

## 6.11.6. Configuration

---

In addition to configuration considerations, this page also includes the steps to configure the Recovery Kit for Oracle Cloud Infrastructure (RK for OCI).

### Specific Configuration Considerations for RK for OCI

Before using the RK for OCI, please review the following topics to ensure that the requirements are met.

- [Requirements](#)
- [Recovery Kit for Oracle Cloud Infrastructure Notes](#)

### RK for OCI Configuration Tasks

- [Creating a Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy](#) – Creates a Recovery Kit for OCI resource in your LifeKeeper cluster.
- [Deleting a Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy](#) – Deletes a Recovery Kit for OCI resource from all nodes that make up your LifeKeeper cluster.
- [Extending Your Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy](#) – Extends a Recovery Kit for OCI resource from a primary node to a secondary node.
- [Unextending Your Recovery Kit for Oracle Cloud Infrastructure Hierarchy](#) – Unextends (removes) a Recovery Kit for OCI resource from a single node in your LifeKeeper cluster.
- [Adjusting RK for OCI Tunable Values](#) – Describes the parameters that can be used with the Recovery Kit for OCI.

All the common tasks across all Recovery Kits are described in the [Administration](#) section of the [LifeKeeper for Linux Technical Documentation](#).

- [Create a Resource Dependency](#) – Creates a parent/child dependency between an existing resource and another resource and propagates the dependency changes to the node where the resource is extended.
- [Delete a Resource Dependency](#) – Deletes a resource dependency and propagates the dependency changes to the node where the resource is extended.
- [In Service](#) – Brings a resource hierarchy into service on a specific node.
- [Out of Service](#) – Takes a resource hierarchy out of service on a specific node.
- [View Properties](#) / [Edit Properties](#) – View or edit the properties of a resource hierarchy on a specific node.

## 6.11.6.1. Creating a Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy

To create a Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) resource, complete the following steps:

1. Launch the **Create Resource Wizard** by referring to [Creating Resource Hierarchies](#).
2. Select **OCIVIP** for the Recovery Kit.
3. Enter the following parameters.

 If you click **Cancel** in the middle of creating a hierarchy, the entire creation process will be canceled.

Field	Tips
Switchback Type	<p>This dictates how the RK for OCI resource will be switched back to this node when the node comes back up after a failover. You can choose either intelligent or automatic.</p> <ul style="list-style-type: none"> <li>• Intelligent switchback requires you to switch back resources manually.</li> <li>• Automatic switchback means the switchback will occur automatically when the LifeKeeper communication path with other nodes is reestablished.</li> </ul> <p><b>Note:</b> The switchback type can be changed later from the General tab of the Resource Properties dialog box.</p>
Server	Select the node where you want to create the RK for OCI resource. All the servers in your cluster are included in the drop down list.
Secondary IP address	Enter the secondary private IP address. The IP address specified here will be assigned to the VNIC.
Network Interface	Select the network interface name to which you want to assign an IP address. You can choose from the network interface names corresponding to the VNICs to which the IP address can be assigned.
OCIVIP Resource Tag	Enter a tag name that is unique to the RK for OCI resource you are creating. The default tag name, ocvip-<ip>, is automatically displayed in this field. <ip> is the specified secondary private IP address.

4. Once all the parameters are entered, the resource will be created and come in service. If LifeKeeper detects a problem, an error will be displayed in the information box or log file.
5. Create a corresponding IP resource and bring it in service. The IP address and network interface name of the IP resource are the same as the RK for OCI resource. Also, the netmask value is the subnet mask (CIDR) value of the subnet with which the VNIC to which the IP address is assigned is associated.
  - a. If an IP resource already exists, skip the creation of the corresponding IP resource.
6. Create a dependency with the IP resource as parent and the RK for OCI resource as child.

- a. If you created a new IP resource in step 5, create a dependency.
- b. If the IP resource already exists in step 5 and the creation of the IP resource is skipped, and when the IP address, network interface name and netmask value of the existing IP resource is:
  - i. an exact match to the setting value of the RK for OCI resource, create a dependency;  
OR
  - ii. inconsistent with the setting value of the RK for OCI resource, the IP resource is not a suitable parent resource for the RK for OCI resource and no dependency will be created. Change the IP resource settings or delete the IP resource first and then recreate a RK for OCI resource.

If the resource is created successfully, the **Pre-Extend configuration task** will be launched.

For more information on how to extend a resource hierarchy to another node, see [Extending Your Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy](#).

## 6.11.6.2. Deleting a Recovery Kit for Oracle Cloud Infrastructure Resource Hierarchy

---

To delete a Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) resource, see [Deleting a Hierarchy from All Servers](#).



When a Recovery Kit for OCI resource is deleted, the resource is taken out of service (i.e. the IP address is unassigned from the VNIC).

## 6.11.6.3. Extending Your OCI Resource Hierarchy

Follow the steps below to extend the Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) resource hierarchy.

1. Launch the **Extend Resource Hierarchy** wizard referring to [Extending Resource Hierarchies](#) .
2. Enter the following information in the **Pre-Extend Wizard**.

Field	Tips
Switchback Type	<p>Select the Switchback Type. This dictates how the RK for OCI resource will be switched back to the extended node when it comes back into service after a failover to the backup server. You can choose either intelligent or automatic.</p> <ul style="list-style-type: none"> <li>• Intelligent switchback means you need to switch back resources manually.</li> <li>• Automatic switchback means the switchback will occur automatically when the LifeKeeper communication path with other nodes is reestablished.</li> </ul> <p><b>Note:</b> The switchback type can be changed later from the General tab of the Resource Properties dialog box.</p>
Template Priority	<p>Enter the priority of the resource hierarchy of the node to extend. Any unused priority value from 1 to 999 is valid. A lower number means a higher priority. Default value is recommended and the number indicates a server's priority in the cascading failover sequence for the resource.</p> <p><b>Note:</b> This selection will appear only for the initial extension of the resource hierarchy.</p>
Target Priority	<p>Enter the priority of the resource hierarchy of the target node. Any unused priority value from 1 to 999 is valid. A lower number means a higher priority. The number indicates a server's priority in the cascading failover sequence for the resource.</p> <p><b>Note:</b> LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities do not need to be consecutive and no two servers can have the same priority for a given resource.</p>

3. **CanExtend** is executed to ensure that the target RK for OCI resource extension requirements are met.
  - a. Make sure that the target node has a VNIC to which an IP address can be assigned.
4. Enter the following parameters to extend the resource.

Field	Tips
Secondary IP address	The IP address of the resource to extend will appear. The value cannot be changed.

Network Interface	Select the network interface name to which you want to assign an IP address. You can choose from the network interface names corresponding to the VNICs to which the IP address can be assigned.
OCIVIP Resource Tag	Enter the RK for OCI resource tag. This is the resource tag name to be used by the RK for OCI resource being extended to the target server. The resource tag name on the source node is displayed by default.

5. Click **Next Server** if you want to extend the same RK for OCI resource to another node in your cluster and repeat the steps. If you click **Finish**, LifeKeeper will verify that the extension of the RK for OCI resource was completed successfully.



Be sure to test the functionality of the RK for OCI resources on all nodes.

## 6.11.6.4. Unextending Your OCI Hierarchy

---

To unextend a hierarchy, see [Unextending a Hierarchy](#).

## 6.11.6.5. Adjusting Recovery Kit for Oracle Cloud Infrastructure Tunable Values

---

For the parameters that can be configured with the Recovery Kit for Oracle Cloud Infrastructure, see the [Recovery Kit for Oracle Cloud Infrastructure Parameters List](#).

## 6.11.7. Troubleshooting

---

There are currently no troubleshooting issues.

## 6.11.7.1. Known Issues / Restrictions

### Known Issues

Description
<p>When creating a Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) resource, if the startup process of the corresponding IP resource fails, the dependency will not be created.</p> <p><b>Solution:</b> After configuring the IP resource to start successfully, manually create a dependency between the RK for OCI resource and the IP resource. Specify the IP resource as the parent resource and the RK for OCI resource as the child.</p>
<p>When executing <code>lkcli import</code>, the following message appears regarding the IP resource <code>&lt;ip-res&gt;</code> corresponding to the RK for OCI resource.</p> <pre>Resource '&lt;ip-res&gt;' already exists. skipping &lt;ip-res&gt;.</pre> <p><b>Cause:</b> This happens because the corresponding IP resource is created together with the RK for OCI resource. The RK for OCI resource is created prior to the IP resource.</p> <p><b>Solution:</b> If all the resources listed in the configuration file (in YAML format) have been created without excess or deficiency, there is no problem. Otherwise, configure the remaining resources manually. Alternatively, perform the import again excluding the RK for OCI resource and manually create the RK for OCI resource only.</p>
<p>When installing LifeKeeper on a BM instance, "Recovery Kit for Oracle Cloud Infrastructure" does not appear in the Recovery Kit selection in the setup menu; the Recovery Kit cannot be selected for installation.</p> <p><b>Solution:</b> Specify the following options when executing setup.</p> <pre>--env oci</pre> <p>Alternatively, <a href="#">create a configuration information file with the create_response_file script</a> and perform a non-interactive installation.</p>

### Restrictions

Description
<p>Due to a bug preventing the assignment of arbitrary IPv6 addresses in the OCI CLI used by the RK for OCI, IPv6 addresses are not available. <b>See bug here:</b> <a href="https://github.com/oracle/oci-cli/issues/421">https://github.com/oracle/oci-cli/issues/421</a></p>

## 6.11.7.2. Error Messages

---

For error messages that you may encounter with the Recovery Kit for Oracle Cloud Infrastructure, see the [Recovery Kit for Oracle Cloud Infrastructure Message Catalog](#).

## 6.12. Oracle Recovery Kit Administration Guide

---

The LifeKeeper for Linux Oracle® Recovery Kit provides fault resilience for Oracle software in a LifeKeeper environment. The Recovery Kit software furnishes a mechanism to tie the data integrity of Oracle databases to the increased availability provided by LifeKeeper.

### Document Contents

This documentation contains the following topics:

[LifeKeeper for Linux Technical Documentation](#) (also available from the Help menu within the LifeKeeper GUI). A list of all the LifeKeeper for Linux documentation and where the information is available.

[Requirements](#). Before you can install and set up the recovery software, your server must meet certain hardware and software requirements. You should refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper Oracle Recovery Kit.

[Configuring Your Recovery Kit](#). To ensure that your LifeKeeper configuration provides the protection and flexibility you require, you need to be aware of the configuration rules. To appropriately plan your configuration, you must understand your network configuration, interface selection, user system setup, hierarchy options and the Oracle configuration tasks. In addition to planning your configuration, this section also includes configuration examples and the specific tasks required to configure your Recovery Kit.

[Troubleshooting](#). This section provides a list of informational and error messages with recommended solutions.

## 6.12.1. Oracle Recovery Kit Hardware and Software Requirements

---

Before attempting to install or remove the LifeKeeper Oracle Recovery Kit, you must understand the hardware and software requirements and the installation and removal procedures.

### Kit Hardware and Software Requirements

Be sure that your configuration meets the following requirements:

**Servers.** The Recovery Kit requires two or more supported computers configured in accordance with LifeKeeper requirements described in the [LifeKeeper Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#).

**Shared Storage.** Oracle databases must reside on a shared disk in a LifeKeeper environment. Depending on your shared storage architecture, the appropriate LifeKeeper shared storage or multipath storage kit will need to be installed on each node in your cluster. In the example of NFS backed database storage, installation of the LifeKeeper NAS Kit is necessary. If you are planning to use LifeKeeper in a data replication (mirroring) environment, see the [SIOS DataKeeper Administration Guide](#). If you are using LifeKeeper in a network attached storage (NAS) environment, see the [LifeKeeper Network Attached Storage Recovery Kit Administration Guide](#).

**LifeKeeper Software.** You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [LifeKeeper for Linux Release Notes](#) and [LifeKeeper for Linux Technical Documentation](#) for specific LifeKeeper requirements.

**LifeKeeper IP Recovery Kit.** This Recovery Kit is required if remote clients will be accessing the Oracle Database. You must have the same version of this Recovery Kit on each server.

**IP Network Interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

**Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation, and so forth.

**TCP/IP Software.** Each server also requires the TCP/IP software.

**Oracle Software.** Each server must have the Oracle software installed and configured before you can configure LifeKeeper and the LifeKeeper Oracle Recovery Kit. The same version should be installed on each server. Consult the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information.

You should refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install

or remove the LifeKeeper Oracle Recovery Kit.

## 6.12.2. Configuring Oracle with LifeKeeper

---

This section contains information you should consider before you start to configure Oracle and examples of typical Oracle configurations.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View Properties](#). View the properties of a resource hierarchy on a specific server.
- [Edit Properties](#). Edit the properties of a resource hierarchy on a specific server.

For instructions on installing Oracle on Linux distributions using the 2.6 kernel, please see your Linux distribution's website.

Also, please refer to your [LifeKeeper for Linux Technical Documentation](#) located on the SIOS Technology website for instructions on configuring your LifeKeeper Core resource hierarchies.

## 6.12.2.1. Specific Configuration Considerations for Oracle

- \* If you plan to use Oracle with Raw I/O, the Raw I/O devices must be properly set up prior to use. See the [Appendix](#) for instructions. (Raw I/O is not an option for LifeKeeper Single Server Protection.)

Before configuring the Oracle Recovery Kit, complete the following preparatory steps to ensure that file systems and disk partitions used by Oracle will be accessible from each server.

1. **Remove personal initialization file prompts.** For the Oracle Recovery Kit to work properly, you must remove (or comment out) all prompts in the personal initialization file (i.e., `.profile`, `.bash_profile`) for the Oracle user. This file is specific to the shell that is being used by the Oracle user. The file cannot be interactive.

- \* If “`stty`” statements are going to be in the personal initialization file, they must be in an “`if`” statement that verifies that an interactive terminal is being used.

2. **Configure Kernel Parameters.** Please refer to the Oracle documentation for information on how linux kernel parameters such as shared memory and other kernel resources should be configured. An example of how to set these parameters is below.

On *each server* in the cluster:

- a. Set the following `ipcs` limits in `/etc/sysctl.conf` before configuring LifeKeeper.

```
# changes for Oracle
kernel.shmmax = <value>
kernel.shmmni = <value>
kernel.shmall = <value>
kernel.sem = <value>
```

- b. Run `sysctl -p` to set the above changes in the kernel.

c. On certain distributions you may need to add `sysctl -p` to the system initialization file (i.e. `boot.local` or `rc.local`) so that these kernel changes are set after each reboot.

3. **`$ORACLE_HOME` directory.** When you configure the `$ORACLE_HOME` directory and associated files on local disks, be sure that the `$ORACLE_HOME` directory and files are identical on all servers. Use the standard Linux utilities to create and copy directories and files to the set of servers.

✿ In certain active/active configurations, the location of `$ORACLE_HOME` are different.

4. **Location.** The `$ORACLE_HOME` directory can be on shared or non-shared disks. The advantage to having the directory on shared media is that you only need to configure files such as the parameter file `Oracle_HOME/dbs/<initSID.ora or spfileSID.ora>` once, if the same shared disk is used for `$ORACLE_HOME` (e.g. in an active/standby configuration). The disadvantage to the shared directory is that direct access to the file system is available to only one server at a time. SCSI reservations permit only one server at a time access to a LifeKeeper protected shared drive. If creating an active-active cluster configuration where two or more Oracle instances (SID) will be protected independently in the cluster, `$ORACLE_HOME` must be installed on local, non-shared storage.
5. **User and Group ID.** An oracle user (oracle) and group (dba) should be created on all servers. The user ID and group ID numbers must be the same on all servers.
6. **Databases, archive files, log files and control files.** All databases, archive files, log files, and control files must be created on shared file systems or disk partitions. These locations are set in the Oracle parameter file `init<SID>.ora` or `spfileSID.ora`. Please refer to the Oracle documentation for information on editing database parameters. The pathnames must be the same for all servers. Oracle internally keeps this information in its control file; therefore, SYSTEMS database space and paths cannot be changed unless Oracle is running.

✿ Oracle log archiving is not enabled by default. If it is enabled prior to the creation of the LifeKeeper Oracle hierarchy, LifeKeeper will detect the location of the archive files and create a separate file system hierarchy if necessary. But if log archiving is enabled after the LifeKeeper Oracle hierarchy has been created, you must manually create and extend a file system hierarchy to protect the shared archive location, and create a dependency from the Oracle resource to this new file system hierarchy.

✿ When using storage applications with locking and following recommendations for the NFS mount options, LifeKeeper requires the additional `noLOCK` option be set, e.g. `rw,noLOCK,bg,hard,nointr,tcp,nfsvers=3,timeo=600,rsize=32768,wsizes=32768,a`

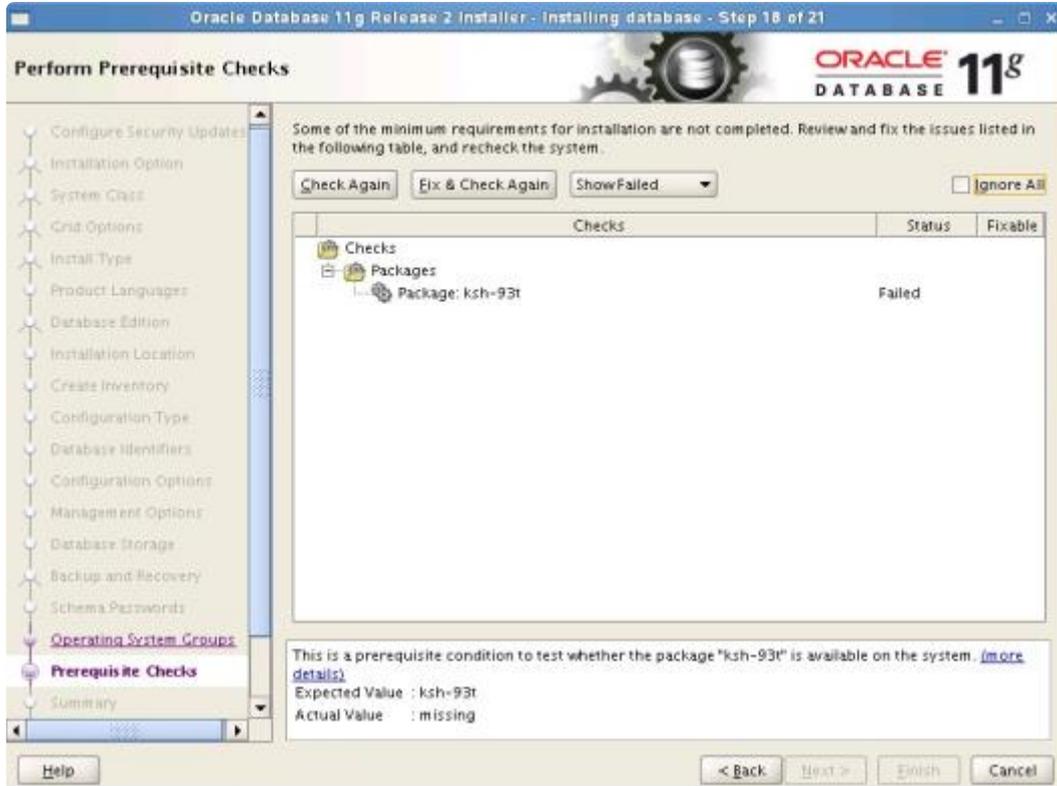
7. On a new installation of Oracle, the final configuration of the database instance is easier if the database installation program is not allowed to create a database. When the installer asks if you want to create a database, select **No**. After the installation is complete, run the **Oracle Database Creation Assistant** (dbca). dbca provides much better control of where database components get created. When running dbca, specify that the Flash Recovery Area gets created on LifeKeeper protected storage.

! The Flash Recovery Destination must be located on a shared drive.

If runInstaller is allowed to create a database, the Flash Recovery Area will have to be relocated

manually. (**Note:** Allowing `runInstaller` to create a database is not recommended.)

- During the installation of Oracle using the "runInstaller" utility, there will be a point where the installer verifies the packages and configuration of Linux before proceeding with the Oracle database installation. If LifeKeeper 7.2 (or higher version) has already been installed, a message complaining about a missing ksh package will appear.



If this message is displayed, check the box in the upper righthand corner, **Ignore All**. The installation of LifeKeeper has removed the ksh package and replaced it with the Public Domain Korn Shell, `pdksh`. Oracle should install fine using `pdksh`.

\* Beginning with Version 8.0, the Oracle ARK no longer requires `pdksh`; however, `pdksh` is still required by the LifeKeeper core and therefore still requires checking the Ignore All setting.

\* Beginning with Version 8.1, LifeKeeper provides its own private `pdksh` package and therefore does not conflict with Oracle's `ksh` requirements.

- Tune the database engine.** Refer to Oracle documentation for guidelines on tuning the database engine for data integrity and performance. In particular, the tuning for memory caching and checkpointing frequency is critical to optimizing the application for fault resilience. The checkpoint interval determines the number of uncommitted database transactions. As a result, it determines the number of database transactions that will be lost in the event of a system failure.
- Database entry in oratab file.** The `/etc/oratab` file must contain an entry for the database.

The LifeKeeper configuration routines use the contents of this file to relate `$ORACLE_HOME` and `$ORACLE_SID` values. Usually, the Oracle installation program creates the required entry. In a configuration in which the Oracle software is installed to a shared file system, however, you must copy the `oratab` file from the server where the Oracle installation was performed to the `/etc` directory of the other servers so that it is available to all the servers.

\* The configuration can have only one `oratab` per server. Refer to the Oracle Product Manual for information on the file format.

\* The `oratab` file can be accommodated in other locations besides `/etc`. By default, the Oracle ARK looks for the `oratab` file in `/etc` followed by `/var/opt/oracle`. If the `oratab` file is not located in one of these default locations, then `ORACLE_ORATABLOC` must be set in `/etc/default/LifeKeeper` to the directory containing `oratab`.

11. **Disable automatic start-ups.** Since LifeKeeper is responsible for starting the databases it controls, be sure to disable any automatic start-up actions. LifeKeeper disables automatic start-up when a hierarchy is created. This is accomplished by modifying the `oratab` file.
12. The Listener configuration file, `listener.ora`. New lines should not be embedded in the entries (e.g., `SID_NAME=xx` should be on one line).
13. **Oracle Database Username and Password.** LifeKeeper will use local session and OS Authentication to control Oracle Database. If you would like to turn off local OS Authentication for security reasons, LifeKeeper can use the specified username and password. The Oracle Database user must be able to connect as `sysdba` authority to the database to be protected, and each server's Oracle Database must have the same username and password. If this configuration is skipped during resource creation, then LifeKeeper will not use username and password to control the Oracle Database resource. This parameter can be added, changed or removed any time after creating the resource.

Once under LifeKeeper protection, the LifeKeeper and database user privileges can be lowered from `sysdba` to `sysoper`. See [Changing Username / Password for the Oracle Database Account](#) for more information.

#### Tips for Creating the Oracle Username and Password.

- a. On the node where the Oracle database is running, log in to Linux with a user that is part of the `dba` group. (The "oracle" account is most common.) Using the `sqlplus` utility, connect to the database as the administrative user by issuing the following command:

```
$ sqlplus / as sysdba
```

- b. Create a new user for this function:

```
SQL> CREATE USER lkdba IDENTIFIED BY "password";
```

c. Then grant this user SYSDBA privileges:

```
SQL> GRANT SYSDBA to lkdba;
```

d. If Oracle has been configured so that each node in the LifeKeeper cluster has a local copy of `$ORACLE_HOME`, execute these commands on each node in the cluster. After creating the LifeKeeper Oracle hierarchy, bring the database in service on the node and then execute the `CREATE` and `GRANT` commands (above) to set up the user in Oracle.

**!** **CAUTION:** Avoid configuring two databases on the same file system. If you must configure two databases on the same file system, exercise great care. In this situation, both databases must be placed under LifeKeeper protection and both hierarchies must have the same primary and backup servers.

## 6.12.2.2. Configuring the Oracle Net Listener for LifeKeeper Protection

If your Oracle database will have remote client connections, you will want to protect the Oracle Listener in addition to the Oracle database server. Please refer to the Oracle documentation for information on using Oracle network configuration utilities to create Oracle network configuration files such as `listener.ora` and `tnsnames.ora`.

**Note:** Refer to the [Creating a Shared Oracle Listener for Multiple Resources](#) section in the appendix in this document for instructions on how to create a shared Oracle Listener for multiple resources.

### Listener Configuration

1. You need to choose a vip address for clients to make connections to. You may want to put this address in DNS. (Refer to the [LifeKeeper IP Recovery Kit Documentation](#) for details on creating an IP resource hierarchy. Refer to the topic [Creating a Resource Dependency](#) under GUI Administration Tasks for details on creating a resource dependency).
2. In the `listener.ora` file, specify this vip address as the HOST for the database service name. (See the Oracle documentation for details about the `listener.ora` file.) Although the DNS name can be used in place of the vip address for the HOST database service name, LifeKeeper best practices does not recommend this. Using the vip address will prevent DNS lookup issues from impacting LifeKeeper's ability to determine the status of a running listener during quickCheck, restore or remove processing. Additionally, a `SID_LIST_LISTENER` must be defined, even though you may have only one listener defined.

Sample format of a `listener.ora`:

```
.
.
.
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = <SID Name>)
    )
  )
.
.
.
<listener name>=
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = <vip>) (PORT = <port number>))
    )
  )
```

```
)
.
.
.
```

Specify the vip address as the HOST in the `tnsnames.ora` file or Oracle Names. (See the Oracle documentation for details about the `tnsnames.ora` file.) Although the DNS name can be used in place of the vip address for the HOST database service name, LifeKeeper best practices does not recommend this. Using the vip address will prevent DNS lookup issues from impacting LifeKeeper's ability to determine the status of a running listener during quickCheck, restore or remove processing:

```
.
.
.
<SID Name>=
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = <vip>) (PORT = <port number>))
    )
    (CONNECT_DATA =
      (SID = <SID Name>)
    )
  )
)
```

These sample files should work with both Oracle 10g and 11g:

`listener.ora`

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ORA11A)
    )
  )
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = 192.0.2.0) (PORT = 1521))
    )
  )
)
```

`tnsnames.ora`

```
ORA01 =
  (DESCRIPTION =
```

```

    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = 192.0.2.0) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SID = ORA01)
    )
  )

```

The normal location of `listener.ora` is in `$ORACLE_HOME/network/admin`. The most common port number is 1521. The global name of the database was defined at creation time. Also keep in mind, if the `$ORACLE_HOME` directory is installed on non-shared storage, a copy of `listener.ora` will need to be on both systems.

**Note:** Oracle Net provides the option of automatically failing over client connections to another listener if the listener for a service should fail. To take advantage of this feature, set the `FAILOVER` parameter to “**ON**” in the `tnsnames.ora` file. If the listener for the LifeKeeper-protected Oracle SID should fail, this allows client connections to continue through another listener until LifeKeeper recovers the protected listener.

## Possible Error

If you encounter the following error please see the solution below.

### Oracle listener fails with TNS-00511: No Listener Linux Error:111: Connection refused

If you have configured the `listener.ora` file the same as in the example below, SIOS LifeKeeper will not be able to start the listener.

```

(ADDRESS_LIST=
  (ADDRESS=(PROTOCOL=tcp) (HOST=<hostname>) (PORT=1521))
  (ADDRESS=(PROTOCOL=ipc) (KEY=EXTPROC1521))
  (ADDRESS=(PROTOCOL=tcp) (HOST=<vip>) (PORT=1521))
)

```

Where `<hostname>` is the hostname of the host and `<vip>` is the virtual IP address,

the actual error which may scroll by quickly is:

TNS-12542: TNS:address already in use TNS-12560: TNS:protocol adapter error

## Solution

SIOS recommends that you change the `listener.ora` as such:

```

(ADDRESS_LIST=

```

```
(ADDRESS=(PROTOCOL=tcp) (HOST=<ip>) (PORT=1521))  
(ADDRESS=(PROTOCOL=ipc) (KEY=EXTPROC1521))  
(ADDRESS=(PROTOCOL=tcp) (HOST=<vip>) (PORT=1521))  
)
```

Where <ip> is the ip address of the host and <vip> is the virtual IP address.

**Note:** if security audits are an issue, this should work as well where vipname is the DNS name of the virtual IP.

```
(ADDRESS_LIST=  
    (ADDRESS=(PROTOCOL=tcp) (HOST=<vipname>) (PORT=1521))  
)
```

This should also pertain to tnsnames.ora file changes.

## 6.12.2.3. Configuring Transparent Application Failover with LifeKeeper

---

When a server failover or an Oracle database failure occurs, users can be severely disrupted. Typically the user's connections to the database will be lost along with most work in progress. Upon the completion of the failover (or recovery of the Oracle database), clients will have to restart their application and reconnect to the database. With the Transparent Application Failover (TAF) feature of Oracle, this disruption can be reduced or eliminated by masking some types of failures. To configure TAF in a LifeKeeper environment, there are tasks that must be performed on both the LifeKeeper server side and the Oracle client side.

For clients to effectively take advantage of the TAF feature, the client application must use failover-aware API calls from the Oracle Call Interface (OCI). The clients must also configure the appropriate TAF support using the Oracle Net parameters in the `tnsnames.ora` file. TAF mode can be configured by including a `FAILOVER_MODE` parameter under the `CONNECT_DATA` section of the `tnsnames.ora` connect descriptor. The TAF mechanism supports several sub-parameters to control and affect the behavior of a client connection during failover. The LifeKeeper for Linux Oracle Recovery Kit supports the following TAF configuration sub-parameters:

### **TYPE= (SELECT or SESSION).**

This value determines how TAF will handle client connection failover. When the type is set to **SELECT**, Oracle keeps track of all select statements issued during transition. Upon establishment of a new connection, the select statements are re-executed, and the cursors repositioned so clients can continue to fetch rows. When type is set to **SESSION** only a new connection is created; work in progress may be lost.

### **METHOD= (BASIC).**

With this method TAF will attempt a reconnect only after the primary connection fails. The alternative method is **PRECONNECT**, LifeKeeper does not currently support the use of **PRECONNECT** as a method.

### **DELAY= (#sec).**

This value is the number of seconds that TAF will wait between attempts to connect following a failure. This value should be carefully determined for your client application and environment.

### **RETRIES= (#number of tries).**

This value is the number of times that TAF will attempt to retry a failed connection before giving up. The combination of **DELAY** and **RETRIES** must allow enough time for a complete recovery of Oracle in the event of a server failure. This will give TAF enough time to restart after the server failover has completed.

An excerpt from a sample `tnsnames.ora` file for a client system is included below.

```
LKproDB=
  (DESCRIPTION=
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL=TCP) (HOST=<switchableIP>) (PORT=<port number>))
    )
    (CONNECT_DATA=
      (SID=LKroDB)
      (SERVER=DEDICATED)
      (FAILOVER_MODE=
        (TYPE=SELECT)
        (METHOD=BASIC)
        (DELAY=5)
        (RETRIES=30)
      )
    )
  )
)
```

The normal location of `tnsnames.ora` is in `$ORACLE_HOME/network/admin`. The most common port number is 1521. The `tnsnames.ora` files can also be located in user's home directories as well. Also, keep in mind, if the `$ORACLE_HOME` directory has been installed on non-shared storage, a copy of `listener.ora` and `tnsnames.ora` will need to be on both systems.

On the LifeKeeper server protecting the Oracle database, the listener should be configured using a LifeKeeper-protected switchable IP address. Refer to the [Configuring the Oracle Net Listener for LifeKeeper Protection](#) section above for details on configuring Oracle Net and listener support.

## 6.12.2.4. Configuring a Pluggable Database with Oracle Multitenant

LifeKeeper can protect the pluggable database (“PDB”) as well as the Oracle database server provided that the Oracle database supports the Oracle Multitenant architecture and protects the container database (“CDB”).

### Checking CDB and PDB

1. Oracle resources must be created to protect the PDB. Also, the protected Oracle resource must be a CDB. You can check whether it is a CDB or not after connecting to the database using the following command.

```
SQL> select CDB from V$DATABASE;
```

2. To protect the PDB, the PDB must be mounted inside the CDB. You can check whether the PDB is mounted by using the following command after connecting to the database.

```
SQL> show pdbs;
```

### Creating Oracle PDB Resources

1. From the LifeKeeper GUI menu, select **Edit**, then select **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

 **Important:** When you create a resource, the Oracle resource must have been created and be in service.

A dialog box appears displaying all the recognized Recovery Kits installed in the cluster in a dropdown list. Select **Oracle Pluggable Database** from the dropdown list. Click **Next** to proceed to the next dialog box.

**Note:** If the **Back** button is active in a dialog box, you can return to the previous dialog box by clicking it. This is especially useful when you encounter errors and need to correct the information you entered earlier.

At any stage of the hierarchy creation process, clicking **Cancel** will cancel the entire creation process.

2. You will be prompted to enter the following information. If the **Back** button is active in a dialog box, you can return to the previous dialog box. This is useful when you encounter errors and need to correct the information you entered earlier. You can click **Cancel** at any time to cancel the entire creation process.

Field	Description
-------	-------------

<b>Server</b>	Select the LifeKeeper server on which Oracle PDB is created.
<b>Switchback Type</b>	Select intelligent or automatic. After the failover, when the Oracle PDB resource is brought back in service (active) on the backup server, how it is switched back to the primary server is determined. Intelligent switchback (intelligent) requires administrator intervention to switch resources back to the primary server, while automatic switchback (automatic) brings the primary server back online and a switchback is performed as soon as the LifeKeeper communication path is reestablished.  <b>Note:</b> The switchback method must be the same as the one for dependent resources used by the Oracle PDB resource.
<b>ORACLE_SID</b>	Specify the SID of the protected Oracle database.
<b>Oracle PDBs</b>	Specify the PDB to protect. This field allows multiple selections.
<b>PDB Tag</b>	A unique tag name for the new Oracle PDB resource on the primary server. The default tag name is "pdb-<ORACLE_SID>". You can also use another unique tag name. You can use letters, numbers, and the special symbols (".", "-", "_", ":", "/") for tag names.

3. Click **Next**. **Create Resource Wizard** appears and the Oracle PDB resource hierarchy is created. LifeKeeper verifies the input data. If a problem is detected, an error message appears in the information box.
4. A message saying that the Oracle PDB resource hierarchy has been successfully created and that the hierarchy must be extended to another server in the cluster to provide failover protection is displayed. Click **Next**.
5. Click **Continue**. The **Pre-extend Wizard** is launched. See Step 2 in "Extending the Oracle PDB Resource Hierarchy" for details on extending the resource hierarchy to another server.

## Extending Oracle PDB Resources

1. Select **Extend Resource Hierarchy** from **Resource** in the **Edit** menu. The **Pre-Extend Wizard** is displayed. If you are not familiar with advanced operations, click **Next**. If you understand the default values for extending the LifeKeeper resource hierarchy and do not need to enter and confirm them, click **Accept Defaults**.
2. Enter the following details in the **Pre-Extend Wizard**.

**Note:** The first two fields appear only when you start the operations from **Extend** in the **Edit** menu.

Field	Description
<b>Template Server</b>	Select the server where the Oracle PDB resource is currently in service.

<b>Tag to Extend</b>	Select the Oracle PDB resource to extend.
<b>Target Server</b>	Enter or select the target server.
<b>Switchback Type</b>	<p>After the failover, when the Oracle PDB is brought in service (active) on the backup server, how it is switched back to the primary server is determined. You can choose intelligent or automatic. The switchback type can be changed later on the <b>General</b> tab of the <b>Resource Properties</b> dialog box when necessary.</p> <p><b>Note:</b> The switchback method must be the same as the one for the dependent resource used by the Oracle PDB resource.</p>
<b>Template Priority</b>	<p>Select or enter a <b>template priority</b>. This is the priority of the Oracle PDB hierarchy which has currently been in service on the server. You can use any unused number between 1 and 999 for the priority, with lower numbers having higher priority (number 1 is the highest priority). During the extension process, priorities that are already in use by another system cannot be specified for this hierarchy. SIOS recommends the default value.</p> <p><b>Note:</b> This field appears only when you extend the hierarchy for the first time.</p>
<b>Target Priority</b>	<p>This is the relative priority which is owned by the newly extending Oracle PDB hierarchy over the equivalent hierarchies on other servers. Any unused priority number between 1 and 999 is available and indicates the server's priority for the resource cascading failover sequence. Note that LifeKeeper assigns "1" by default to the server on which the hierarchy was created. The priorities do not need to be consecutive but two servers cannot have the same priority for a particular resource.</p>

- When the pre-extending checking is successful message is displayed, click **Next**.
- Depending on the hierarchy to extend, a series of information boxes will be displayed showing the resource tags to be extended (some cannot be edited).
- Confirm that the tag name is correct in **Extend Wizard** and click **Extend**.
- When the message "Hierarchy extend operations completed" is displayed, click **Next Server** if you want to extend the hierarchy to another server, or click **Finish**.
- When the message "Hierarchy Verification Finished" is displayed, click **Done**.

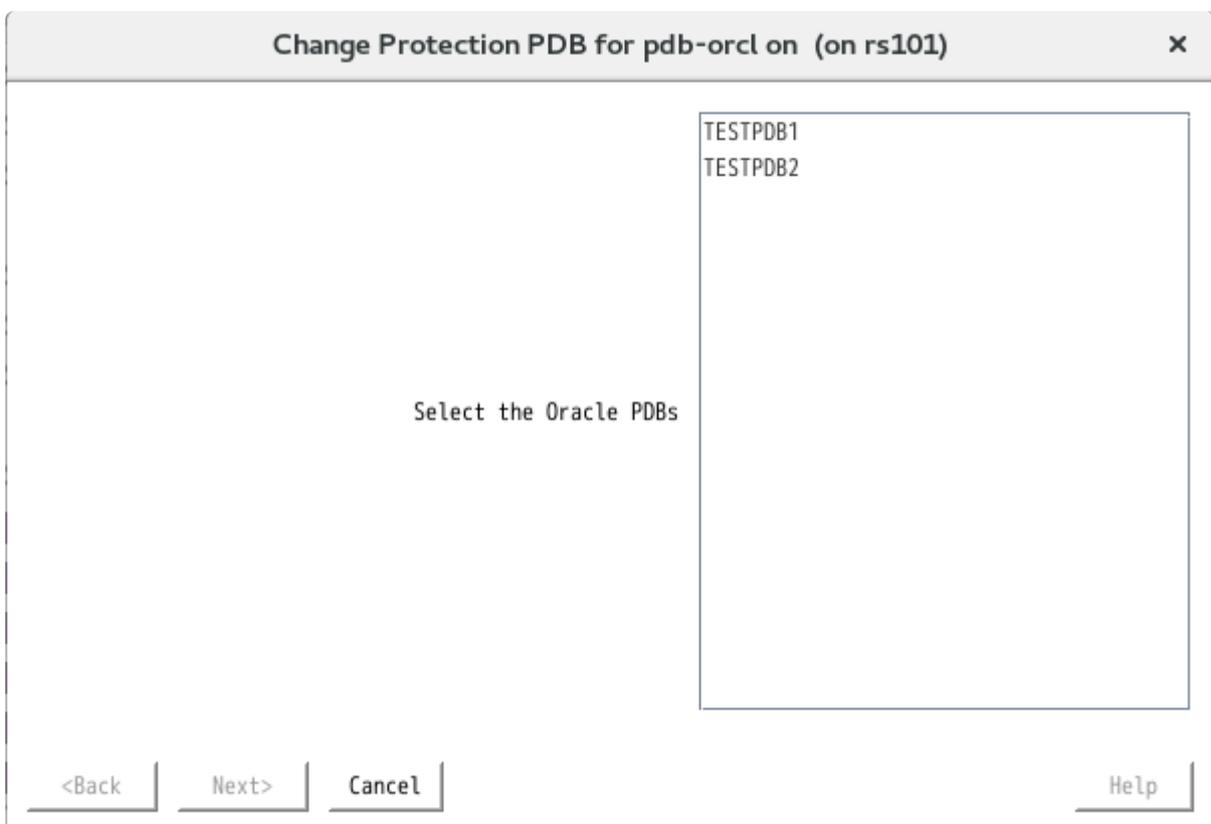
## Changing the PDB to Protect

After creating the hierarchy, change the PDB to protect using the following steps.

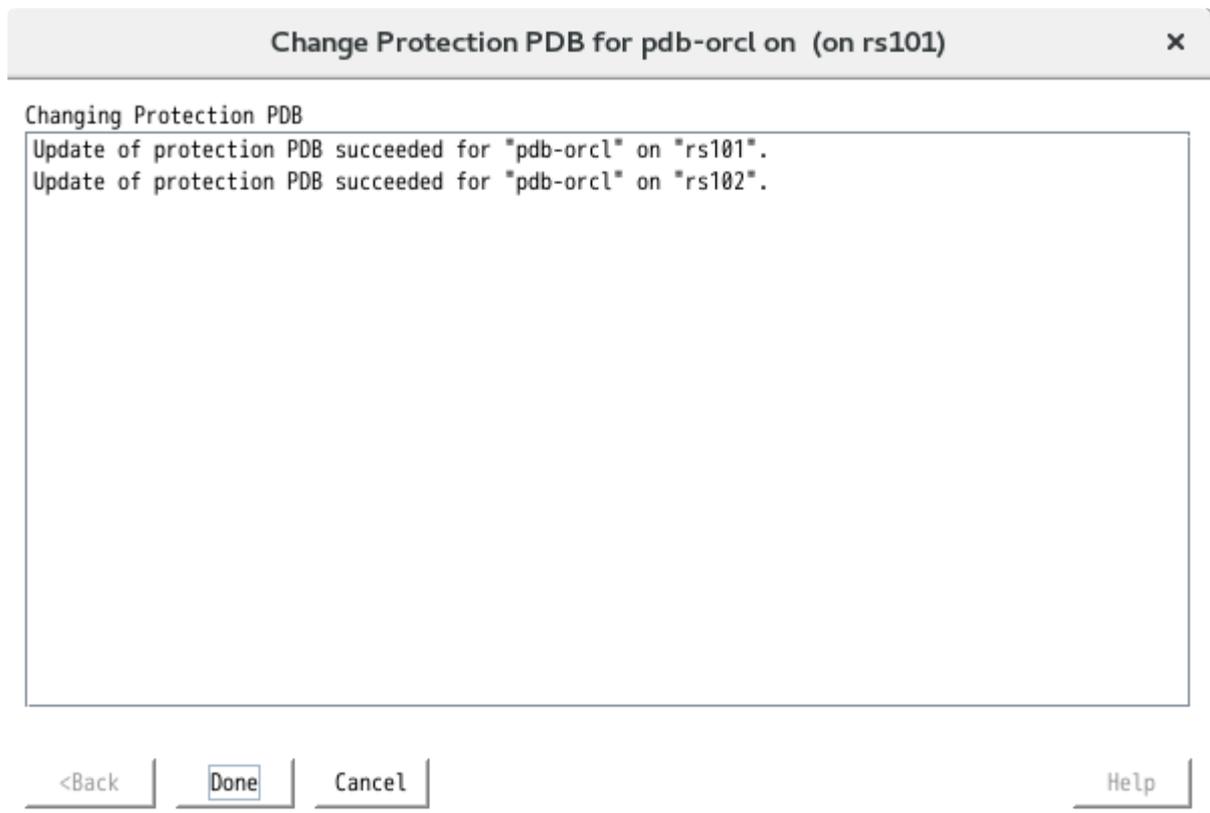
- From the LifeKeeper GUI, right-click the Oracle PDB resource hierarchy and select **Change Protection PDB**.



2. Select the PDB you want to protect (you can select more than one PDB).



3. Click **Next** to change the settings.



4. Click **Done** to finish.

## 6.12.2.5. Oracle Configuration Examples

---

The following figures illustrate examples of both active/standby and active/active Oracle configurations in a LifeKeeper environment.

The examples in this section show how Oracle database instances can be configured on local and shared disks. Each diagram shows the relationship between the type of configuration and the Oracle parameters. Each configuration also adheres to the configuration rules and requirements described in this administration guide that ensure compatibility between the Oracle configuration and the LifeKeeper software.

This section first describes the configuration requirements and then provides these configuration examples:

- [Active/Standby](#)
- [Active/Active](#)

The examples in this section are only a sample of the configurations you could establish, but understanding these configurations and adhering to the configuration rules helps you define and set up workable solutions for your computing environment.

## 6.12.2.5.1. Oracle Configuration Requirements

---

Each of the examples involves one or two databases: **databaseA** and **databaseB**. By default, LifeKeeper offers a tag name matching the Oracle database system identifier (SID). However, the screen examples in the following pages use tag names consisting of the SID and server name such as databaseA-on-server1.

To understand the configuration examples, keep these configuration requirements in mind:

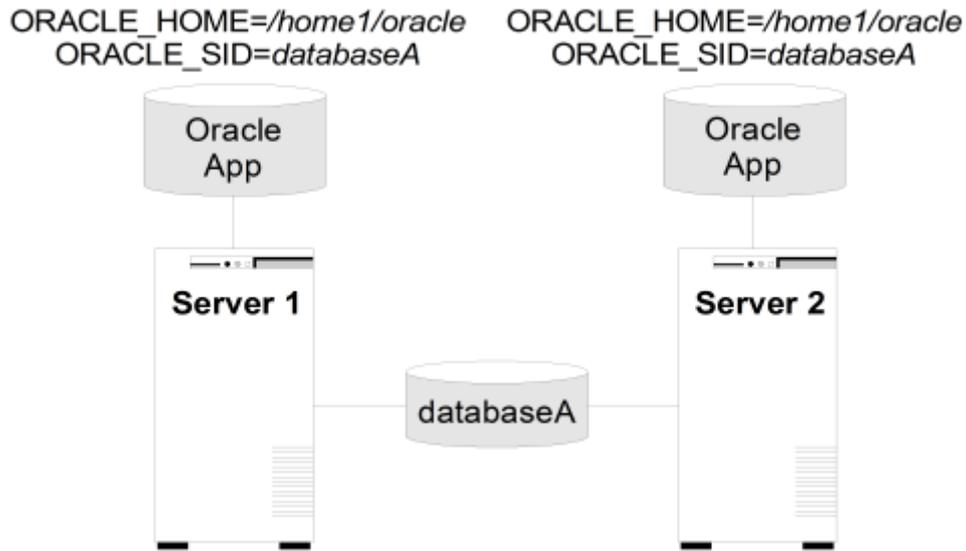
- **LifeKeeper hierarchy.** When performing LifeKeeper administration, the primary server refers to the location where the Oracle instance is currently running. System administration takes place on this server when creating a LifeKeeper hierarchy. For the configuration examples, the primary server is Server 1 and the backup or alternate server is Server 2.
- **Shared disk locked by one server.** When shared storage resources are under LifeKeeper protection, they can only be accessed by one server at a time. If the shared device is a disk array, an entire LUN is protected. If a shared device is a disk, then the entire disk is protected. This prevents inadvertent corruption of the data by other servers in the cluster. When a server fails, the highest priority backup server establishes its own protection, locking out all other servers.
- **Database on shared disk.** In order for the LifeKeeper Oracle Recovery Kit to function properly, the database must always be on a shared device. The database may be on one or more file systems and/or disks.

**Note:** The tables associated with each configuration illustration provide examples of the appropriate information that would be entered in the Create Resource Hierarchy wizard on the primary server (Server 1) and Extend Resource Hierarchy wizard to the backup server (Server 2). For additional detail on what information to enter into the wizards, refer to the [“LifeKeeper Configuration Tasks”](#) section. These tables can be a helpful reference when configuring your Recovery Kit.

## 6.12.2.5.2. Oracle Active/Standby Configurations

This section provides two active/standby configuration examples, shown in Figure 1 and Figure 2. In these configurations, Server 1 is considered active because it has exclusive access to the database. Server 2 does other processing. If Server 1 fails, Server 2 gains access to the database and LifeKeeper re-establishes the database operations.

Figure 1. Active/Standby Configuration, Example 1



### Configuration Notes:

1. Each server has its own \$ORACLE\_HOME directory on a non-shared disk. Each server has the same version of the Oracle application.
2. The \$ORACLE\_HOME path is the same on both servers.
3. The database, databaseA, is on a shared disk.

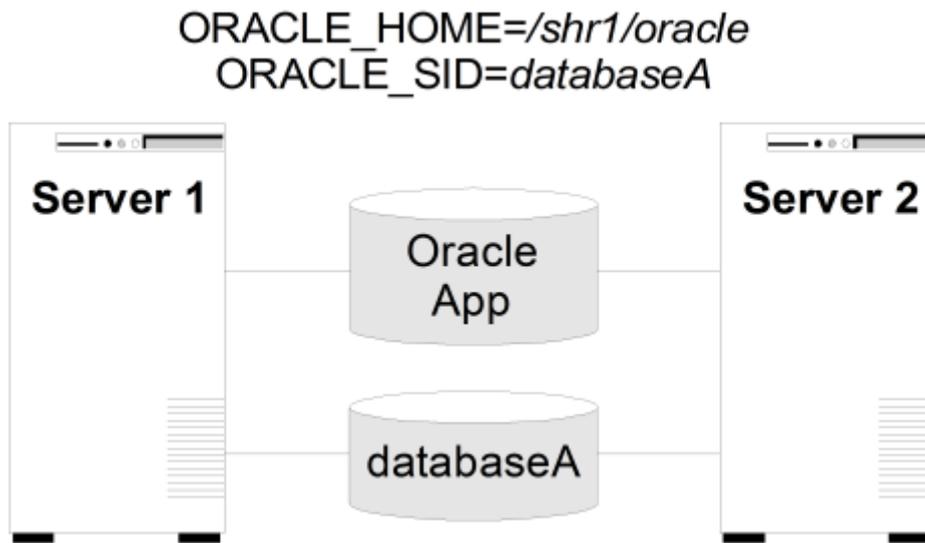
Creating a resource hierarchy on Server 1:

Server:	Server1
ORACLE_SID for Database:	databaseA
Username for Database:	system
Password for Username:	*****
ORACLE_HOME for Database:	/home1/oracle
Database Tag:	databaseA-on-server1

Extending the resource hierarchy to Server 2:

Template Server:	Server1
Tag to Extend:	databaseA-on-server1
Target Server:	Server2
Target Priority:	10
Database Tag:	databaseA-on-server1

Figure 2. Active/Standby Configuration, Example 2



**Configuration Notes:**

1. Both servers use the \$ORACLE\_HOME directory on a shared disk.
2. The \$ORACLE\_HOME path is the same on both servers.
3. The database, databaseA, is on a shared disk.
4. Server 2 can not access files and directories on the shared disk while Server 1 is active.
5. \$ORACLE HOME can be on the same shared disk as the database or on separate disks.

Creating a resource hierarchy on Server 1:

Server:	Server1
ORACLE_SID for Database:	databaseA
Username for Database:	system
Password for Username:	*****
ORACLE_HOME for Database:	/shr1/oracle
Database Tag:	databaseA-on-server1

## Extending the resource hierarchy to Server 2:

Template Server:	Server1
Tag to Extend:	databaseA-on-server1
Target Server:	Server2
Target Priority:	10
Database Tag:	databaseA-on-server1

## 6.12.2.5.3. Oracle Active/Active Configurations

An active/active configuration consists of at least two servers, each running a different database instance. The databases must be on different shared disks.

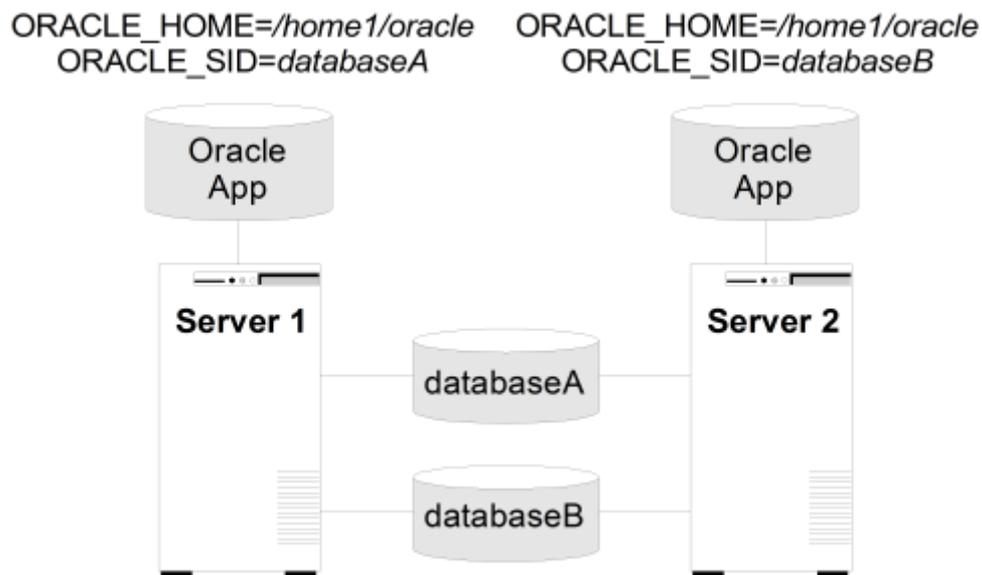
`$ORACLE_HOME` can be on non-shared or on shared disks depending upon the configuration requirements. For example, multiple database instances on any of the servers using a common `$ORACLE_HOME` require `$ORACLE_HOME` to be on non-shared disks. If the `$ORACLE_HOME` directories are on shared disk, they must be on separate shared disks.

This section provides two active/active configuration examples, shown in Figure 3 and Figure 4:

- Databases on shared resources and a common `$ORACLE_HOME` on non-shared resources.
- Databases on shared resources and the appropriate `$ORACLE_HOME` instance on the same shared resource.

**Note:** Multiple database instances on one server using multiple instances of `$ORACLE_HOME` on non-shared resources are not illustrated.

Figure 3. Active/Active Configuration, Example 1



### Configuration Notes:

1. The server has its own `$ORACLE_HOME` directory on a non-shared disk. Each server has the same version of the Oracle application.
2. The `$ORACLE_HOME` path is the same on both servers.
3. The databases, databaseA and databaseB, are on shared disks.

4. The oratab file exists in /etc/ on both servers, containing entries for both Oracle instances.
5. Initially, Server 1 runs databaseA and Server 2 runs databaseB. In a switchover situation, one system can run both databases.
6. See [Creating Oracle Database Hierarchy After Installing Oracle Binaries on Local Storage](#) for further information.

Creating the first resource hierarchy on Server 1:

Server:	Server1
ORACLE_SID for Database:	databaseA
Username for Database:	system
Password for Username:	*****
ORACLE_HOME for Database:	/home1/oracle
Database Tag:	databaseA-on-server1

Extending the first resource hierarchy to Server 2:

Template Server:	Server1
Tag to Extend:	databaseA-on-server1
Target Server:	Server2
Target Priority:	10
Database Tag:	databaseA-on-server1

Creating a second resource hierarchy on Server 2:

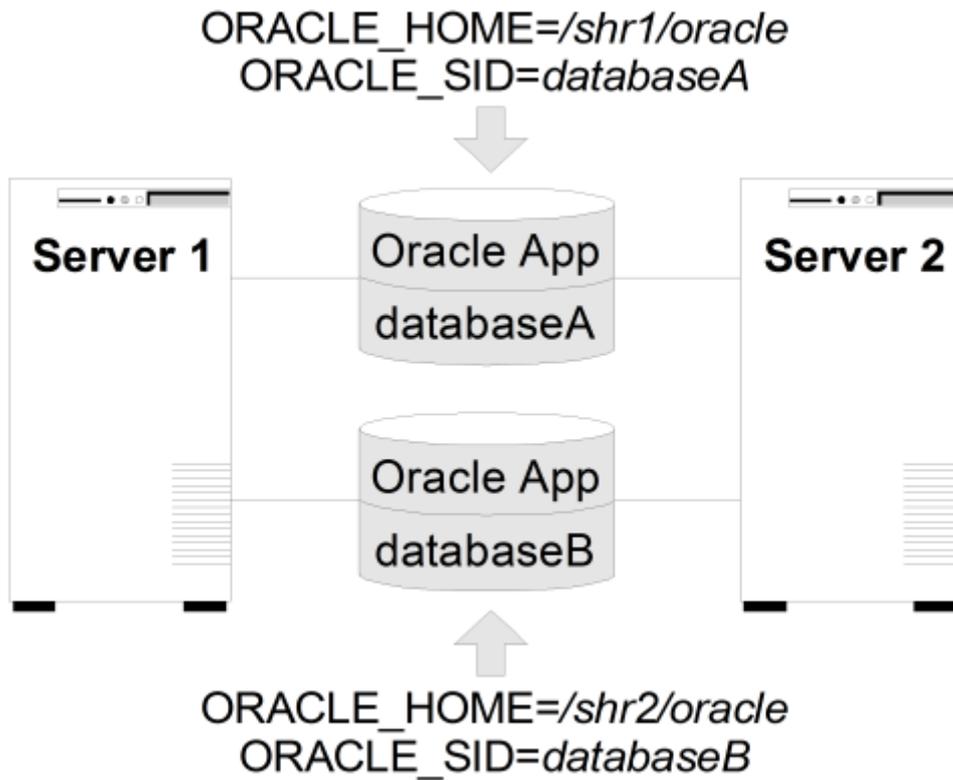
Server:	Server2
ORACLE_SID for Database:	databaseB
Username for Database:	system
Password for Username:	*****
ORACLE_HOME for Database:	/home1/oracle
Database Tag:	databaseB-on-server2

Extending the second resource hierarchy to Server 1::

Template Server:	Server2
Tag to Extend:	databaseB-on-server2
Target Server:	Server1

Target Priority:	10
Database Tag:	databaseB-on-server2

Figure 4. Active/Active Configuration, Example 2



**Configuration Notes:**

1. Both servers use an \$ORACLE\_HOME directory on different shared disks.
2. The Oracle application is the same on both servers. The \$ORACLE\_HOME is different for each instance defined on the server.
3. The databases, databaseA and databaseB, are on shared disks.
4. The oratab file exists in /etc/, containing entries for both Oracle instances.
5. A unique login is required for each Oracle instance. The id and gid for each login should be the same on Server 1 and Server 2.
6. Initially, Server 1 runs databaseA and Server 2 runs databaseB. In a switchover situation, one system can run both databases.

Creating the first resource hierarchy on Server 1:

Server:	Server1
ORACLE_SID for Database:	databaseA

Username for Database	system
Password for Username	*****
ORACLE_HOME for Database:	/shr1/oracle
Database Tag:	databaseA-on-server1

Extending the first resource hierarchy to Server 2:

Template Server:	Server1
Tag to Extend	databaseA-on-server1
Target Server	Server2
Target Priority:	10
Database Tag:	databaseA-on-server1

Creating a second resource hierarchy on Server 2:

Server:	Server2
ORACLE_SID for Database:	databaseB
Username for Database	system
Password for Username	*****
ORACLE_HOME for Database:	/shr2/oracle
Database Tag:	databaseB-on-server2

Extending the second resource hierarchy to Server 1:

Template Server:	Server2
Tag to Extend	databaseB-on-server2
Target Server	Server1
Target Priority:	10
Database Tag:	databaseB-on-server2

## 6.12.3. LifeKeeper Configuration Tasks for Oracle

---

You can perform the following configuration tasks from the LifeKeeper GUI. The following four tasks are described in this guide, as they are unique to an Oracle resource instance and different for each Recovery Kit.

- [Create a Resource Hierarchy](#). Creates an application resource hierarchy in your LifeKeeper cluster. Replicated (SIOS DataKeeper) file system resources must be created before creating the Oracle resource.
- [Delete a Resource Hierarchy](#). Deletes a resource hierarchy from all servers in your LifeKeeper cluster.
- [Extend a Resource Hierarchy](#). Extends a resource hierarchy from the primary server to a backup server.
- [Unextend a Resource Hierarchy](#). Unextends (removes) a resource hierarchy from a single server in the LifeKeeper cluster.
- [View Oracle Configuration Settings](#) – Allows viewing of the Resource Properties dialog.
- [Change Username / Password](#). Change the Username and Password to login to protect Oracle Database.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View Properties](#). View the properties of a resource hierarchy on a specific server.
- [Edit Properties](#). Edit the properties of a resource hierarchy on a specific server.

**Note:** Throughout the rest of this section, we explain how to configure your Recovery Kit by selecting certain tasks from the **Edit** menu of the LifeKeeper GUI. You can also select each configuration task from the toolbar. You may also right-click on a global resource in the Resource Hierarchy Tree (left-hand pane) of the status display window to display the same drop down menu choices as the Edit menu. This,

of course, is only an option when a hierarchy already exists.

You can also right-click on a resource instance in the Resource Hierarchy Table (right-hand pane) of the status display window to perform all the configuration tasks, except *Creating a Resource Hierarchy*, depending on the state of the server and the particular resource.

## 6.12.3.1. Creating an Oracle Resource Hierarchy

---

**Note:** In order to take advantage of Oracle Net remote client access, the IP address used for client connectivity must be under LifeKeeper protection as a dependent of the Oracle hierarchy. (Refer to the section [Configuring the Oracle Net Listener for LifeKeeper Protection](#) for details.)

To create a resource instance from the primary server, you should complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop down menu, select **Create Resource Hierarchy**.

**Important:** The Oracle Application must be running when you create the resource.

A dialog box will appear with a drop down list box with all recognized Recovery Kits installed within the cluster. Select **Oracle Database** from the drop down listing. Click **Next** to proceed to the next dialog box.

**Note:** When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful should you encounter an error that might require you to correct previously entered information.

If you click **Cancel** at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the Oracle instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either intelligent or automatic. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box. Click **Next** to proceed to the next dialog box.

3. Select the **Server** where you want to place the Oracle Database (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop down box. Click **Next** to proceed to the next dialog box.
4. Select the **ORACLE\_SID** for the Database. This is the tag name that specifies the Oracle system identifier of the database being configured. An entry for this database must exist in */etc/oratab*. Click **Next** to proceed to the next dialog box.
5. Input the **Username for ORACLE\_SID**. This is the Oracle Database Username specified during login to ORACLE\_SID. This username must be able to connect as sysdba authority to the database to gain full control. Click **Next** to proceed to the next dialog box. (This field can be left empty. If left empty, LifeKeeper will not use Username and Password to control the Oracle

Database resource, and the next step, **Input Password**, will be skipped.)

6. Input **Password**. This is the password specified during login to ORACLE\_SID. The password will be saved by LifeKeeper with encrypting. Click Next to proceed to the next dialog box.
7. Select the **tag name** of the Listener to be included as a dependency of the Oracle resource. The list displays all the currently protected Listener resource(s) on the server. Select the Listener resource tag that corresponds to the required listener(s) for the Oracle SID. Select **None** if no Listener resource exists.
8. Select or enter the **Database Tag**. This is a tag name that LifeKeeper gives to the Oracle hierarchy. You can select the default or enter your own tag name.

When you click **Create**, the **Create Resource Wizard** will create your Oracle resource.

9. At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your Oracle resource hierarchy. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created.

Creating database/oracle resource...

```
Tue Feb 28 13:38:42 EST 2012 burn create[17114]: INFO:RKBase:create::000000:BEGIN
create of "ORA" on server "burn"
Tue Feb 28 13:38:49 EST 2012 burn create[17114]: INFO:oracle:create::122516:Creating
dependency between Oracle database "ORA (JEF)" and the dependent resource "/oracle"
on "burn".
Tue Feb 28 13:38:49 EST 2012 burn create[17114]:
INFO:oracle:create::122522:Performing in-service of new Oracle resource tag=< ORA >
on "burn".
Tue Feb 28 13:38:49 EST 2012 burn create[17114]: INFO:RKBase:create::000000:END
successful create of "ORA" on server "burn"
```

Click **Next** to proceed to the **Pre-extend dialog box** which is explained later in this documentation. You must extend the hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

## 6.12.3.2. Deleting an Oracle Resource Hierarchy

To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your Oracle resource hierarchy.

**Note:** If you selected the Delete Resource task by right-clicking from the right pane on an individual resource instance, or from the left pane on a global resource where the resource is on only one server this dialog box will not appear.



Target Server

Click **Next** to proceed to the next dialog box.

3. Select the Hierarchy to Delete. Identify the resource hierarchy you wish to delete, and highlight it.

**Note:** If you selected the Delete Resource task by right clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Hierarchy to Delete

ORA

Click **Next** to proceed to the next dialog box.

4. An information box appears confirming your selection of the target server and the hierarchy you

have selected to delete.

```
You have specified the following resource hierarchy for deletion.
Target Server: burn
Target Tags:
ORA
```

Click **Delete** to delete your resource and proceed to the final dialog box.

5. Another information box appears confirming that the Oracle resource was deleted successfully.

```
Deleting resource hierarchy ORA
Removing root resource hierarchy starting at "ORA":
Tue Feb 28 14:33:16 EST 2012 burn
/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete[24962]:
INFO:RKBase:/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete::000000:
BEGIN delete of "ORA" on server "burn"
Tue Feb 28 14:33:24 EST 2012 burn
/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete[24962]:
INFO:RKBase:/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete::000000:
END successful delete of "ORA" on server "burn"
Tue Feb 28 14:33:20 EST 2012 wake1
/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete[26543]:
INFO:RKBase:/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete::000000:
BEGIN delete of "ORA" on server "wake1"
Tue Feb 28 14:33:20 EST 2012 wake1
/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete[26543]:
INFO:RKBase:/opt/LifeKeeper/subsys/database/resources/oracle/actions/delete::000000:
END successful delete of "ORA" on server "wake1"
Tue Feb 28 14:33:28 EST 2012 wake1 delete[26653]: INFO:RKBase:delete::000000:BEGIN
delete of "datarep-oracle" on server "wake1"
Tue Feb 28 14:33:33 EST 2012 wake1 delete[26653]: INFO:RKBase:delete::000000:END
successful delete of "datarep-oracle" on server "wake1"
Tue Feb 28 14:33:36 EST 2012 burn delete[25280]: INFO:RKBase:delete::000000:BEGIN
delete of "datarep-oracle" on server "burn"
Tue Feb 28 14:33:37 EST 2012 burn delete[25280]: INFO:RKBase:delete::000000:END
successful delete of "datarep-oracle" on server "burn"
Successfully removed
```

6. Click **Done** to exit out of the **Delete Resource Hierarchy** menu selection.

**Note:** Refer to the [Creating a Shared Oracle Listener for Multiple Resources](#) section in the appendix of this document for instructions on how to create a shared Oracle Listener for multiple resources.

## 6.12.3.3. Extending Your Oracle Hierarchy

After you have created a hierarchy, you will want to extend that hierarchy to another server in the cluster. There are three possible scenarios to extend your resource instance from the template server to a target server. The first scenario is when you “**Continue**” from creating the resource into extending that resource to another server by clicking **Next** on the information dialog box displayed at the completion of the create. The second scenario is when you enter the **Extend Resource Hierarchy** task from the edit menu as shown below. The third scenario is when you right-click on an unextended hierarchy in either the left- or right-hand pane. Each scenario takes you through the same dialog boxes (with a few exceptions, which are clearly detailed below).

1. If you are entering the **Extend Wizard** from the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop-down menu, select **Extend Resource Hierarchy**. This will launch the **Pre-Extend Resource Hierarchy Wizard**.
2. The first dialog box to appear will ask you to select the **Template Server** where your Oracle resource hierarchy is currently in service. It is important to remember that the **Template Server** you select now and the **Tag to Extend** that you select in the next dialog box represent an in service resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you selected. The drop-down box in this dialog provides the names of all the servers in your cluster.

**Note:** If you are entering the **Pre-Extend Resource Hierarchy** task immediately following the creation of an Oracle resource hierarchy, this dialog box will not appear, since the wizard has already identified the template server in the create stage. This is also the case when you right-click on either the Oracle resource icon in the left hand pane or right-click on the Oracle resource box in the right hand pane the of the GUI window and choose **Extend Resource Hierarchy**.



It should be noted that if you click the **Cancel** button at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extension process to that particular server. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

For example, let's say you have created your resource on Server 1 and extended that resource to Server 2. In the middle of extending the same resource to Server 3, you change your mind and click on the **Cancel** button inside one of the dialog boxes. This will cancel only your action to extend the resource to Server 3, not the extension you created to Server 2. If you want to remove Server 2 from this hierarchy, you must unextend the resource from Server 2.

Click **Next** to proceed to the next dialog box.

3. Select the **Tag to Extend**. This is the name of the Oracle instance you wish to extend from the template server to the target server. The wizard will list in the drop-down box all the resources that

you have created on the template server, which you selected in the previous dialog box.

**Note:** Once again, if you are entering the Pre-Extend Resource Hierarchy task immediately following the creation of an Oracle resource hierarchy, **this dialog box will not appear**, since the wizard has already identified the tag name of your Oracle resource in the create stage. This is also the case when you right-click on either the Oracle resource icon in the left hand pane or on the Oracle resource box in the right hand pane the of the GUI window and choose **Extend Resource Hierarchy**.



A screenshot of a GUI dialog box. It features a label 'Tag to Extend' followed by a text input field containing the text 'ORA'. To the right of the input field is a small downward-pointing arrow icon, indicating a dropdown menu.

Click **Next** to proceed to the next dialog box.

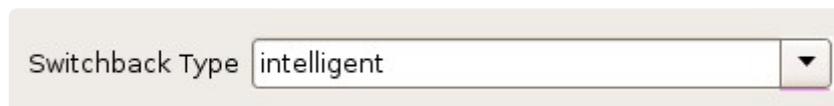
4. Select the **Target Server** where you are extending your Oracle resource hierarchy. The drop-down box provides the names of the servers in your cluster that are not already in the selected hierarchy.



A screenshot of a GUI dialog box. It features a label 'Target Server' followed by a text input field containing the text 'wake1'. To the right of the input field is a small downward-pointing arrow icon, indicating a dropdown menu.

Click **Next** to proceed to the next dialog box.

5. Select the **Switchback Type**. This dictates how the Oracle instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either *intelligent* or *automatic*. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.



A screenshot of a GUI dialog box. It features a label 'Switchback Type' followed by a text input field containing the text 'intelligent'. To the right of the input field is a small downward-pointing arrow icon, indicating a dropdown menu.

The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.

Click **Next** to proceed to the next dialog box.

6. Select or enter a **Template Priority**. This is the priority for the Oracle hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. Note: This selection will appear only for the initial extend of the hierarchy.
7. Select or enter the **Target Priority**. This is the priority for the new extended Oracle hierarchy

relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.

A screenshot of a web form element. It consists of a light gray rounded rectangle containing the text "Target Priority" followed by a text input field with the value "10" and a small downward-pointing arrow icon on the right side of the input field.

Click **Next**.

- An information box will appear explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this Oracle resource have been met. If there were some requirements that had not been met, LifeKeeper would not allow you to select the Next button, and the Back button would be enabled.

```
Executing the pre-extend script...
Building independent resource list
Checking existence of extend and canextend scripts
datarep-oracle is already extended to wake1
Checking extendability for ORA

Pre Extend checks were successful
```

If you click **Back**, you can make changes to your resource extension according to any error messages that may appear in the information box.

If you click **Cancel** now, you will need to come back and extend your Oracle resource hierarchy to another server at some other time to put it under LifeKeeper protection.

When you click **Next**, LifeKeeper will launch you into the Extend Resource Hierarchy configuration task.

- Once the Database Tag displays, click **Extend**.

A screenshot of a web form element. It consists of a light gray rounded rectangle containing the text "Database Tag" followed by a text input field with the value "ORA".

- An information box will appear verifying that the extension is being performed.

**Extending resource hierarchy ORA to server wake1**

Extending resource instances for ORA  
Creating dependencies  
Setting switchback type for hierarchy  
Creating equivalencies  
LifeKeeper Admin Lock (ORA) Released  
  
Hierarchy successfully extended

Click **Next Server** if you want to extend the same Oracle resource instance to another server in your cluster. This will repeat the **Extend Resource Hierarchy** operation.

If you click **Finish**, another dialog box will appear confirming LifeKeeper has successfully extended your Oracle resource.

**Verifying Integrity of Extended Hierarchy...**

Examining hierarchy on wake1  
  
Hierarchy Verification Finished

11. Click **Done** to exit from the **Extend Resources Hierarchy** menu selection.

**Note:** Be sure to test the functionality of the new instance on both servers.

## 6.12.3.4. Unextending Your Oracle Hierarchy

1. From the **LifeKeeper GUI** menu, select **Edit**, then **Resource**. From the drop-down menu, select **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the Oracle resource. It cannot be the server where Oracle is currently in service.

**Note:** If you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance, this dialog box will not appear.



Target Server

Click **Next** to proceed to the next dialog box.

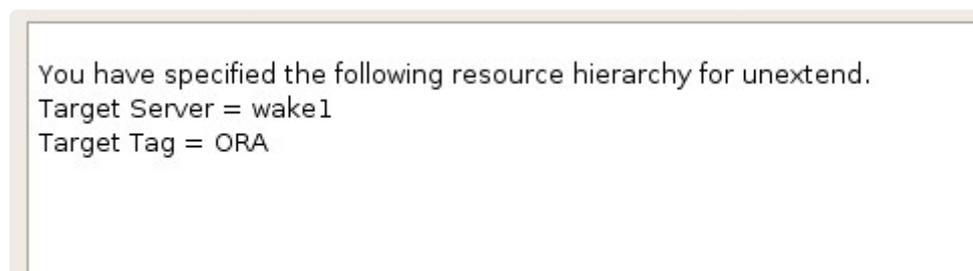
3. Select the **Oracle Hierarchy to Unextend**. Note: If you selected the **Unextend** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Hierarchy to Unextend

Click **Next** to proceed to the next dialog box.

4. An information box appears confirming the target server and the Oracle resource hierarchy you have chosen to unextend.



You have specified the following resource hierarchy for unextend.  
Target Server = wake1  
Target Tag = ORA

Click **Unextend**.

5. Another information box appears confirming that the Oracle resource was unextended successfully.

**Unextending resource hierarchy ORA from wake1**

```
Hierarchy Unextend Manager Initializing
Checking Target Machine Communication Paths
LifeKeeper Admin Lock Flag (ORA) Established
Removing Equivalencies
Removing Resources and Associated Dependencies
Tue Feb 28 14:07:29 EST 2012 wake1 delete[10086]: INFO:RKBase:delete::000000:BEGIN
delete of "datarep-oracle" on server "wake1"
mdadm: stopped /dev/md0
Tue Feb 28 14:07:33 EST 2012 wake1 delete[10086]:
INFO:RKBase:delete::000000:stopping the monitor for /dev/md0 ...
Tue Feb 28 14:07:33 EST 2012 wake1 delete[10086]:
INFO:RKBase:delete::000000:stopping the nbd-client for /dev/nbd0 ...
Tue Feb 28 14:07:33 EST 2012 wake1 delete[10086]:
INFO:RKBase:delete::000000:/dev/md0 has been stopped
Tue Feb 28 14:07:33 EST 2012 wake1 delete[10086]: INFO:DR:delete::104015:The mirror
"/dev/md0" (resource: "datarep-oracle") has been successfully removed from system
"wake1"
Tue Feb 28 14:07:34 EST 2012 wake1 delete[10086]: INFO:DR:delete::104019:The mirror
for resource "datarep-oracle" has been successfully unextended from system "wake1"
Tue Feb 28 14:07:34 EST 2012 wake1 delete[10086]: INFO:RKBase:delete::000000:END
successful delete of "datarep-oracle" on server "wake1"
LifeKeeper Admin Lock Flag (ORA) Released
Synchronizing LifeKeeper Databases
Unextend completed successfully
```

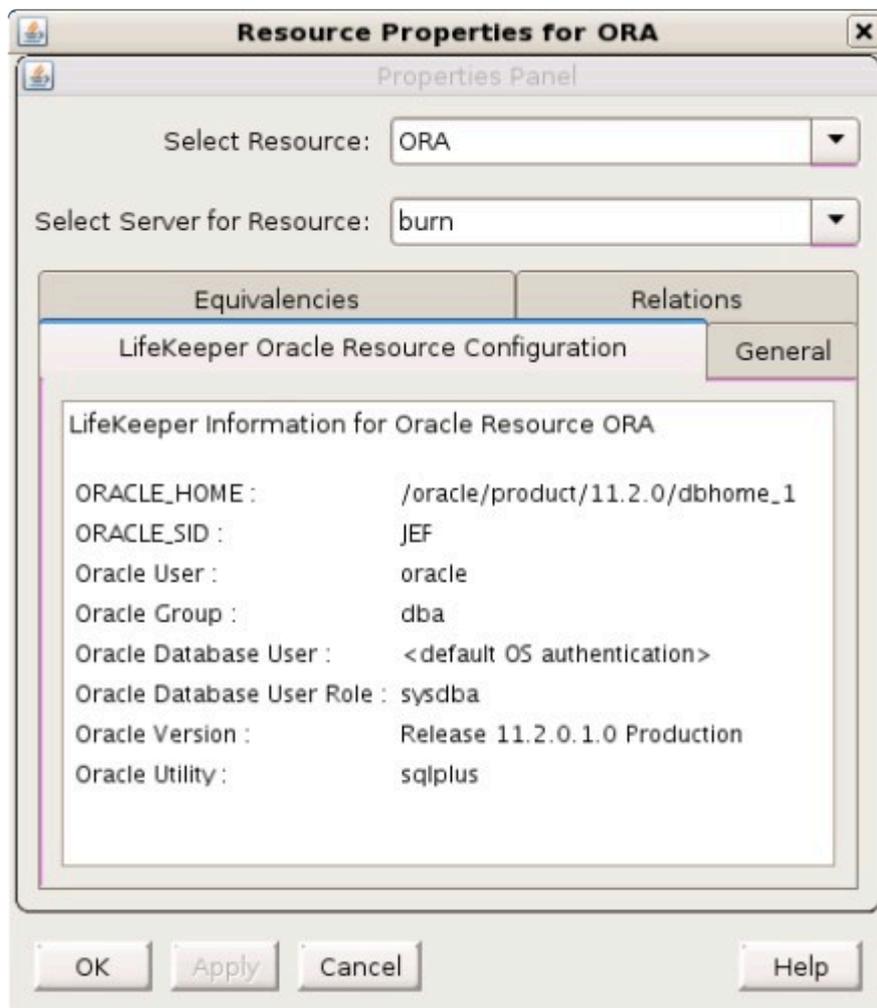
6. Click **Done** to exit out of the **Unextend Resource Hierarchy** menu selection.

## 6.12.3.5. Viewing Oracle Configuration Settings

The **Resource Properties** dialog is available from the **Edit** menu or from a resource context menu. This dialog displays the properties for a particular resource on a server. When accessed from the **Edit** menu, you can select the resource and the server. When accessed from a resource context menu, you can select the server.

From the **Configuration** tab, you can view the following Oracle settings:

- ORACLE\_HOME
- ORACLE\_SID
- Oracle User
- Oracle Group
- Oracle Database User
- Oracle Database User Role
- Oracle Version
- Oracle Utility



## 6.12.3.6. Changing Username / Password for the Oracle Database Account

---

After a hierarchy has been created, change the Username and Password using one of the following procedures.

If \$ORACLE\_HOME is on shared (or replicated) storage (common in active-passive configurations):

1. On the system where the Oracle database resource is operational, edit the LifeKeeper configuration file `/etc/default/LifeKeeper` and add the following line to the file:

```
LK_ORA_NICE=1
```

Do exactly the same on each system in the cluster that has the Oracle resource defined.

2. Use `sqlplus` to change the Oracle user's password.

```
SQL> ALTER USER {username} IDENTIFIED BY {newpassword};
```

3. From the LifeKeeper GUI, right-click on the Oracle Database resource hierarchy, then select **Change Username / Password**.

4. Input **Username** and select **Next**.

5. Input **Password** and select **Next**.

6. Select the **database user role** and click **Apply**. Username and Password will be updated after validating.

7. Select **Done**.

8. Edit the LifeKeeper configuration file on all cluster nodes and make the following change:

```
LK_ORA_NICE=0
```

If \$ORACLE\_HOME is on local storage and each node in the cluster has its own copy of \$ORACLE\_HOME (common in active-active configurations):

9. On the system where the Oracle database resource is operational, edit the LifeKeeper configuration file `/etc/default/LifeKeeper` and add the following line to the file:

```
LK_ORA_NICE=1
```

Do exactly the same on each system in the cluster that has the Oracle resource defined.

10. Use `sqlplus` to change the Oracle user's password.

```
SQL> ALTER USER {username} IDENTIFIED BY {newpassword};
```

11. From the LifeKeeper GUI, right-click on the Oracle Database resource hierarchy, then select **Change Username / Password**.
12. Input user name to use instead temporarily such as sys into **Username** and select **Next**.
13. Input Password and select Next.
14. Select the database user role and click Apply. Username and Password will be updated after validating.
15. Select **Done**.
16. Put the Oracle database resource “In Service” on one of the backup systems.
17. Once the database is running on the backup system, use `sqlplus` to change the Oracle account password.

```
SQL> ALTER USER {username} IDENTIFIED BY {newpassword};
```

When making this password change, use the new password that was set in Step 2. This process resets the security tokens in \$ORACLE\_HOME.

18. Put the database “In Service” on each node in the cluster and repeat Step 8.
19. Once the passwords have been changed on all cluster nodes, put the Oracle database back “In Service” on the desired node.
20. From the LifeKeeper GUI, right-click on the Oracle Database resource hierarchy, then select **Change Username / Password**.
21. Input **Username** and select **Next**.
22. Input **Password** and select **Next**.
23. Select the database user role and click **Apply**. Username and Password will be updated after validating.
24. Select **Done**.
25. Edit the LifeKeeper configuration file on all cluster nodes and make the following change:

```
LK_ORA_NICE=0
```

## 6.12.3.7. Testing Your Oracle Resource Hierarchy

---

You can test your Oracle resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting Edit, then Resource, then finally In Service from the drop down menu. For example, an in service request executed on a backup server causes the application hierarchy to be placed in service on the backup server and taken out of service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the Out of Service request, the application is taken out of service without bringing it in service on the other server.

### Recovery Operations

When the primary server fails, the Oracle Recovery Kit software performs the following tasks:

- Brings Oracle into service on the backup server by bringing in-service the logical interface on one of that server's physical network interfaces. (Note: This occurs only when there is an IP resource instance defined as a dependency of the Oracle hierarchy.)
- Mounts the file system on the shared disk on that server.
- Starts the daemon processes related to Oracle.

Since session context is lost following recovery, after the recovery, Oracle users must reconnect using exactly the same procedures they used to connect originally.

## 6.12.3.8. Patching Oracle Nodes (SAP/Oracle) with DataKeeper

---

Most Oracle updates and patches only require access to the system tables drives and not to the data drives (protected DataKeeper data drives).

The following procedure can be used in general:

Prior to the procedure, set block failover on primary and confirm failover on target.

1. On the standby / target server:
  - a. Apply the appropriate upgrades, patches, etc.
  - b. Reboot your server
2. Then, switchover your resources from the source / active server to the standby / target server:
  - a. Verify that you can access the Oracle database
3. On the other node (which now is a standby after the switchover and the original source / active server):
  - a. Apply the appropriate upgrades, patches, etc.
  - b. Reboot your server
4. (Optional) Switchover again to verify connectivity and access to Oracle on the original / source / active server

In some cases where the patches require access to the DataKeeper volumes (e.g. running catsbp), you need to pause / unlock the mirrors to perform the upgrades on the standby / target system.

Verify that Oracle can be started on the backup, then stop Oracle on the backup node and continue the mirrors. Then repeat the upgrade on the source system.

At the end of the procedure remove the flags 'block failover' on primary and the 'confirm failover' on target.

## 6.12.4. Oracle Troubleshooting

---

The Oracle Recovery Message Catalog below contains listings of all messages that may be encountered while utilizing the Oracle kit.

The [Combined Message Catalog](#) provides a listing of all messages that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received.

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Oracle Message Catalog

### Troubleshooting

The [Message Catalog](#) provides a listing of all messages that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received, or you may go directly to the [Oracle Kit Message Catalog](#) or the [Oracle Listener Message Catalog](#) which contain listings of all messages that may be encountered while utilizing the Oracle Recovery Kit.

- [Oracle Known Issues and Restrictions](#)

## 6.12.4.1. Oracle Known Issues and Restrictions

---

### Control File Switchover Failure

If the \$ORACLE\_HOME directory does not recover, the database control files may not be set up properly. For automatic switchover, the control files need to be configured on a shared device during the database creation. If you keep the control files on separate servers, you must manually update both servers when you need to implement changes.

### Truncated Output

Some versions of Oracle truncate the output when executing the show parameters control\_files in sqldba mode. If your version of Oracle exhibits this behavior, verify the following:

- controlfile parameter. Verify that the controlfile parameter resides in the `$Oracle_HOME/dbs/init#SID.ora` file.
- controlfile devices. Verify that the controlfile devices are on a continuous line, with no new lines, and with each device being separated by a comma.

If an Oracle-related device does not configure properly, then the device can be configured manually using the file system applications available under LifeKeeper Application management.

### Flash Recovery Destination Located on a Shared Drive

As noted in the configuration section of this document, it is important that the Flash Recovery destination is located on a shared drive. To see where Oracle believes the Flash Recovery Area is, issue the following query (as SYSDBA):

```
SQL> SELECT substr(Name,1,30) Name,  
  
       (SPACE_LIMIT/1024/1024/1024) Space_Limit_GB,  
  
       SPACE_USED/1024/1024/1024 Space_Used_GB,  
  
       SPACE_RECLAIMABLE, NUMBER_OF_FILES
```

```
FROM V$RECOVERY_FILE_DEST;
```

```
NAME SPACE_LIMIT_GB SPACE_USED_GB  
SPACE_RECLAIMABLE
```

---

NUMBER\_OF\_FILES

---

/U01/flash\_recovery\_area 3.76171875 .156448364 0

4

Following is an example of how to make a change in `$ORACLE_HOME/dbs/spfile<sid>` to complete this task:

```
SQL> ALTER SYSTEM SET
DB_RECOVERY_FILE_DEST='/oracledb/oracle/flash_recovery_area' scope=both;
```

System altered.

```
SQL> show parameter DB_RECOVERY_FILE_DEST;
```

NAME TYPE VALUE

---

db\_recovery\_file\_dest string /oracledb/oracle/flash\_recovery\_area

db\_recovery\_file\_dest\_size big integer 2G

```
SQL> commit;
```

## Prevent Failover When Unable to Connect to the Database

When resource health checks are performed, the Oracle ARK checks for running database processes and attempts to connect to the database. To prevent a health check failure caused by reaching the maximum allowed connections, set the following in `/etc/default/LifeKeeper`:

```
LK_ORA_NICE=1
```

**Note:** Setting `LK_ORA_NICE` can mask other types of connection errors caused by a non-functioning database. Use this setting with caution.

## Non-Traditional Location for *oratab*

By default, the Oracle ARK looks for the *oratab* file in `/etc` followed by `/var/opt/oracle`. If the *oratab* file is not located in one of these default locations, then `ORACLE_ORATABLOC` must be set in `/etc/default/LifeKeeper` to the directory containing *oratab*.

## NFS Version 4 Not Supported

The Oracle Recovery Kit supports NFSv3 for shared database storage. NFSv4 is not supported at this time due to NFSv4 file locking mechanisms.

## Creating Oracle Database Hierarchy After Installing Oracle Binaries on Local Storage

If you have elected to install the Oracle binaries (`$ORACLE_BASE`) on local storage on each of your LifeKeeper cluster nodes, you will see a message similar to the following when you create your Oracle database hierarchy in the LifeKeeper GUI.

```
BEGIN create of <SID> on server <server1>
Creating resource instance <SID> on server <server1>
Setting resource state for <SID> on server <server1> to "ISP".

ORACLE_HOME "/opt/oracle/app/oracle/product/11.2.0/dbhome_1" does not
reside on a shared file system. Please be sure that the ORACLE_HOME
directory and associated files are identical on all servers. Refer to the
LifeKeeper Oracle Recovery Kit documentation for more information.
Creating dependent file system resource "/u00" on <server1>.
Creating dependency between Oracle database "SID (SID)" and the dependent
resource "/u00" on <server1>.
Creating dependency between Oracle database "SID (SID)" and the listener
resource "LSNR.LISTENER" on <server1>.
Performing in-service of new Oracle resource tag=< SID > on <server1>.
END successful create of on server <server1>
```

You will also find a similar warning in the LifeKeeper log. If this warning is not heeded and you continue on by extending your hierarchy and then try to bring the database resource in service on another node, you will get a message similar to the following in the LifeKeeper GUI dialog:

```
Put resource "OST" in-service
BEGIN restore of "OST" on server "cae-qa-v39"
Begin the "start [ start.normal ]" of the database "OST" on "cae-qa-v39".
The "start [ start.normal ]" attempt of the database "OST" appears to
have failed on "cae-qa-v39".
ORA-01078: failure in processing system parameters
LRM-00109: could not open parameter file '/opt/oracle/app/oracle/product/
11.2.0/dbhome_1/dbs/initOST.ora'
Begin the "start [ start.force ]" of the database "OST" on "cae-qa-v39".
The "start [ start.force ]" attempt of the database "OST" appears to have
failed on "cae-qa-v39".
select 'alter database datafile '||file#||' end backup;' from v\$_backup
where status = 'ACTIVE'
```

It is also possible to get a message similar to the following in the dialog box:

```
Put resource "OST" in-service
BEGIN restore of "OST" on server "cae-qa-v39"
Begin the "start [ start.normal ]" of the database "OST" on "cae-qa-v39".
Initial inspection of "start.normal" failed, verifying failure or success
of received output.
Logon failed with "" for "OST" on "cae-qa-v39". Please check username/
password and privileges.
The "start [ start.normal ]" attempt of the database "OST" appears to
have failed on "cae-qa-v39".
ERROR:

ORA-01031: insufficient privileges

Enter password:

ERROR:

ORA-01005: null password given; logon denied
```

To solve this issue, copy `$ORACLE_BASE/admin` from the primary system where the database instance was created to the backup system (where the hierarchy was extended to) `$ORACLE_BASE/admin`. Also change ownership of this directory to your Oracle username and Oracle group (typically `oracle:oinstall`).

Also copy all `*{$ORACLE_SID}*` (OST in this example) files from the primary system in `$ORACLE_BASE/product/11.2.0/dbhome_1/dbs` to `ORACLE_BASE/product/11.2.0/dbhome_1/dbs` on the backup system.

For example, these were the files that were copied from a primary system to the backup, and the ORACLE SID was OST.

```
-rw-r--- 1 oracle oinstall 1544 2012-05-09 11:02 hc_OST.dat
-rw-r--- 1 oracle oinstall 24 2012-01-31 10:22 lkOST
-rw-r--- 1 oracle oinstall 1536 2012-03-05 09:02 orapwOST
-rw-r--- 1 oracle oinstall 2560 2012-05-09 10:58
spfileOST.ora
```

In another example, where the SID was ORA01, the following files were copied:

```
-rw-r--- 1 oracle oinstall 1536 2010-09-08 18:25 orapwORA01
-rw-r--- 1 oracle oinstall 24 2010-09-08 18:25 lkORA01
-rw-r--- 1 oracle oinstall 2560 2010-09-08 18:30 spfileORA01.ora
-rw-r--- 1 oracle oinstall 1544 2010-09-08 18:30
hc_ORA01.dat
```

and a directory

peshm\_ORA01\_0/:

## Oracle Listener Stays in Service on Primary Server After Failover

Network failures may result in the listener process remaining active on the primary server after an application failover to the backup server.

## 6.12.4.1.1. Oracle Database Creation Problems

<b>Problem:</b>	During DataBase creation using dbca, the following message is received: “ORA-00439 feature not enabled: string”
<b>Action:</b>	Check the value of the environment variable <code>\$ORACLE_SID</code> . Make sure that is the same as the SID that is being created.
<b>Problem:</b>	During Database creation from scripts, the following message is received: “ORA-01092 ORACLE instance terminated. Disconnection forced”
<b>Action:</b>	See the alert in the <code>bdump</code> directory. If you see the message “ORA-12714 invalid national character set specified”, then check the value of the environment variable <code>\$ORA_NLS33</code> . Make sure that it is set to the correct location.
<b>Problem:</b>	If you encounter problems creating the database from the script generated from dbca, then do the following:
<b>Action:</b>	<p>1. Be sure to create the following directories if they do not already exist:</p> <pre>bdump cdump  udump &lt;oracle data base directory&gt;/oradata sid &lt;oracle data base directory&gt;/dbs &lt;oracle data base directory&gt;/admin/&lt;SID&gt;</pre> <p>If you need to determine the path to your <i>bdump</i> and <i>udump</i> directories, you can look in the initialization file (<i>init&lt;SID&gt;.ora</i>).</p> <p>2. Make sure the file <code>\$ORACLE_HOME/dbs/orapw</code> exists; if not, use the <code>orapwd</code> utility to create it.</p>

## 6.12.4.1.2. Oracle Database Startup Problems

---

<b>Problem:</b>	During DataBase start-up using sqlplus, the following message is received: "ORA-03113 end-of-file on communication channel"
<b>Action:</b>	Make sure the initialization file ( <i>init&lt;SID&gt;.ora</i> ) and the password file ( <i>orapw&lt;SID&gt;</i> ) are in the directory <i>\$ORACLE_HOME/dbs</i> .
<b>Problem:</b>	During Database startup, the following message is received: "ORA-01092 ORACLE instance terminated. Disconnection forced"
<b>Action:</b>	See the alert in the <i>bdump</i> directory. If you see the message "ORA-12701 CREATE DATABASE character set is not known", then check the value of the environment variable <i>\$ORA_NLS33</i> . Make sure that it is set to the correct location.

## 6.12.4.1.3. inqfail error in the LifeKeeper Log

If an inqfail error appears in your LifeKeeper error log following a failover, you will need to change the `filesystemio` setting.

**Note:** The disk id and server name will be different for each configuration.

To resolve this problem, you will need to change the setting `filesystemio="SETALL"` to `filesystemio="ASYNCH"`

To locate this setting, query the option with the following SQL command:

1. `SQL> show parameter filesystemio;`
2. Use the following commands to change the settings:

```
SQL> alter system set filesystemio_options=<XXXXXXXX> scope=spfile;
```

<XXXXXXXX> can be set to

```
<XXXXXXXX> = {none | setall | direction | asynch}
```

NONE - no optimization

ASYNCH - enable asynchronous I/O

DIRECTIO - enable direct I/O

SETALL - enables all available features

**IMPORTANT:** Oracle needs to be restarted after resetting the parameter.

## 6.12.5. Oracle Appendix

---

### [Raw I/O](#)

[Adding a Tablespace After Creating Hierarchy](#)

### [Creating Oracle Listener for Multiple Resources](#)

[Updating the Listener Protection Level](#)

[Updating the Listener Recovery Level](#)

[Updating the Protected Listener\(s\)](#)

## 6.12.5.1. Setting up Oracle to Use Raw I/O

Use the following steps to create an Oracle database that uses shared Raw I/O devices instead of files.

1. Determine the minimum number and sizes of files that you will need to create your database, including control files, tablespaces and redologs. You can create a mixed setup with some of those items as files and others on Raw I/O devices. All of the Raw I/O devices must use shared disk partitions.
2. Create a Raw I/O setup with the necessary number of Raw I/O devices.
  - a. Create the raw devices with the same size or larger than you are going to specify for the Oracle database creation.
  - b. Create raw device mappings in the system initialization file (i.e. `boot.local` or `rc.local`) using the `raw` command. You should add meaningful comments to identify which raw device represents which Oracle file. This is done so that the mapping can be re-established in the case of a re-boot of the system. These mappings should be removed from the file manually once the Raw I/O device is under LifeKeeper protection.

3. Make the raw devices writable for the Oracle database using the following command:

```
chown oracle:dba /dev/raw[0-9]*
```

where the owner and group are specific to your Oracle instance's configuration.

4. Activate the raw device settings by executing the file that contains the mappings.
5. If you already have a database creation script, go directly to Step 6. If not, you may use one of the Oracle Java GUI tools, `dbassist` or `dbca`, to generate your database creation scripts. Using either tool, you must choose to "**Save As Script**". Do not choose to create the database.

### Notes:

- In `dbca`, the "**New Database**" template must be selected to generate scripts. Change filenames to shared devices and adjust the values for your configuration if necessary.
  - The DB creation process should not be started at this point! The `dbassist` tool checks to see if the file specified for each tablespace already exists and will not proceed if it does. The `dbca` tool prompts to confirm that it will overwrite the files but fails on raw devices. In either case, you are unable to use raw devices directly from these tools.
6. The database creation scripts (either the existing ones or those created by `dbassist` or `dbca`) must be edited. The desired filename (including the path) must be replaced with the full path name of the Raw I/O device. The affected files should include (at minimum) the file's database creation file (for the `CREATE DATABASE` command) and tablespace creation file (for the `CREATE TABLESPACE` command). Depending on what options you selected in `dbassist` or `dbca`, there may more files to be edited. Also, edit the initialization file to change the control files to Raw I/O

devices, if desired. The initialization file is located in the directory with the creation script. The result looks like this for the data file:

```
. . .

CREATE DATABASE "LK"

    maxdatafiles 254

    maxinstances 8

    maxlogfiles 32

    character set US7ASCII

    national character set US7ASCII

DATAFILE '/dev/raw/raw1' SIZE 260M AUTOEXTEND ON NEXT 10240K

logfile '/ora/LK/redo01.log' SIZE 500K,

    '/ora/LK/redo02.log' SIZE 500K,

    '/ora/LK/redo03.log' SIZE 500K;

. . .
```

The Raw I/O device must be the minimum size required by Oracle for the data that will be stored.

7. Now create the database by running the script that you created in step #5.
8. Be sure to check the create log for any database or tablespace errors that may have occurred.
9. If you have trouble creating the database with the creation scripts, or you want to add tablespaces on raw devices later, you must create the database with the applicable tool (i.e. `dbassist` or `dbca`). Then, add the Raw I/O device data files by executing a command similar to the following from the `sql` utility:

```
tablespace RAWTS
DATAFILE '/dev/raw/raw217' SIZE 50M REUSE
DEFAULT STORAGE (INITIAL 50K NEXT 50K
MINEXTENTS 1 MAXEXTENTS 4) ONLINE
```

10. Add `udev` rules to make the raw device permissions and ownership persistent across reboots, switchovers and failovers.

In the Linux 2.6 kernel, the `udev` system is the default method through which the kernel controls

the creation of the special files such as raw devices. When used by Oracle, raw devices require specific ownership and permission settings. These specific settings conflict with the kernel default settings. Addressing the specific settings requires the use of udev rules to set the ownership of the raw device to the Oracle user and Oracle group used with the LifeKeeper protected Oracle SID.

**Note:** In some OS distributions, rules for creating devices and rules for setting device permissions must be separate. Check your OS distribution udev documentation for more information.

The following are example udev rules that may work for your OS distribution:

```
KERNEL=="raw10", RUN+="/bin/chown oracle:oinstall /dev/raw/raw10"
```

```
KERNEL=="raw[3-5]*", OWNER="oracle", GROUP="oinstall", MODE="660"
```

The udev rules created must be applied to all nodes in the LifeKeeper cluster prior to bringing the resource hierarchy into service.

## 6.12.5.1.1. Adding a Tablespace After Creating Hierarchy

---

If a tablespace is added on a Raw I/O device after the Oracle hierarchy has been created in LifeKeeper, you must create a LifeKeeper Raw I/O hierarchy via the GUI and manually create a dependency between the Oracle resource (as parent) and the Raw I/O resource (as child).

## 6.12.5.2. Creating an Oracle Listener for Multiple Resources

You may want to create an Oracle Listener if any of the following statements are true for your system configuration:

- Multiple Listeners are defined for Multiple Oracle SIDs
- The Oracle Listener is a critical component in your configuration
- A Single Listener is defined for Multiple Oracle SIDs

**\* NOTE:** If multiple listeners are intended to be protected with LifeKeeper, the collection of `LISTENER/SID_LIST_LISTENER` stanzas must all be unique. (The 11g installation's `listener.ora` file would contain the stanza `LISTENER_11G_1/SID_LIST_LISTENER_11G_1.`)

This process will allow protection of listener(s) within LifeKeeper to accommodate various listener(s) and SIDs combinations.

If you are creating a Listener for multiple resources, follow these procedures.

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

**IMPORTANT:** The Oracle Application must be running when you create the resource

2. A dialog box will appear with a drop down list box with all recognized Recovery Kits installed within the cluster. Select **Oracle Database Listener** from the drop down listing. Click **Next** to proceed to the next dialog box.
3. You will be prompted to enter the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful should you encounter an error that might require you to correct previously entered information. If you click **Cancel** at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

Field	Tips
<b>Switchback Type</b>	<p>Choose either intelligent or automatic. This dictates how the Listener resource will be switched back to this server when the server comes back up after a failover. If using data replication, choose intelligent as the switchback type.</p> <p><b>Note:</b> The switchback type must match that of the dependent resources (IP and</p>

	volume resources) used by the Listener resource, or else the create will fail.
<b>Server</b>	Select the server on which you want to create the hierarchy.
<b>Listener Configuration File Path</b>	Select the full path to the Oracle listener configuration file.
<b>Listener Names(s)</b>	Select the name(s) of the Oracle Listener(s) to provide protection for with this resource instance.
<b>Listener Executable(s)</b>	Select the path to the Oracle listener executable. The listener executable is required to start, stop, monitor and recover the specified Oracle listener(s).
<b>Listener Protection Level</b>	Select one of the following levels:  Full Control (Start, Stop, Monitor and Recover)  Intermediate Control (Start, Monitor and Recover)  Minimal Control (Start and Monitor Only)
<b>Listener Recovery Level</b>	Select the level of recovery for the specified listener(s):  Standard (On) – Enable standard LifeKeeper recovery. If Standard (On) is selected, all listener failures will be tried locally, and if necessary, trigger a failover to an available backup server.  Optional (off) – Enable optional LifeKeeper recovery. If Optional (Off) is selected, all listener failures will be tried locally, but will not cause a failover to an available backup server.  <b>Note:</b> Local recovery is performed for both recovery levels when a listener error occurs; however, execution of failover depends on the recovery level.
<b>IP Address Name(s)</b>	Select the IP Address resource name that will be protected as dependents of this resource hierarchy. IP Address associated with the selected listener(s) are displayed in the choice list. Select None if no IP resources are required for this configuration,
<b>Listener Tag</b>	Enter a unique name for the resource on the server. The valid characters allowed for the name are letters, digits, and the following special characters: – _ . /

4. Select the **Create** button to start the hierarchy creation. An information box appears and LifeKeeper will validate that you have provided valid data to create your database listener resource hierarchy. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created.

- The **Pre-Extend Wizard** dialog will appear stating that you have successfully created the resource hierarchy and you will be prompted to select the following information. If you are unfamiliar with the **Extend** operation, click **Next** after making a selection in each dialog box. If you are familiar with the **LifeKeeper Extend Resource Hierarchy** defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.

Field	Tips
<b>Target Server</b>	Select a Target Server to which the hierarchy will be extended. If you select Cancel before extending the resource hierarchy to at least one other server, LifeKeeper will provide no protection for the applications in the hierarchy.
<b>Switchback Type</b>	This dictates how the Oracle Listener instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either intelligent or automatic. The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box. Note: Remember that the switchback strategy must match that of the dependent resources to be used by the Oracle Listener resource.
<b>Template Priority</b>	This field appears only if you did NOT extend directly from the Create function.) Enter a number between 1 and 999 to specify the template server's priority in the cascading failover sequence for this resource. A lower number means a higher priority. LifeKeeper assigns the number "1" to the server on which the hierarchy was created. No two servers can have the same priority for a given resource.
<b>Target Priority</b>	Enter a number between 1 and 999 to specify the target server's priority in the cascading failover sequence for this resource. A lower number means a higher priority. LifeKeeper offers a default of 10 for the first server to which a hierarchy is extended.

- After receiving the message that the pre-extend checks were successful, click **Next** and enter the following information.

Field	Tips
Listener Configuration File Path	Select the full path to the Oracle Listener configuration file.
Listener Executable(s) Path	Select the path to the Oracle Listener executables. The listener executables are required to start, stop, monitor and recover the specified Oracle listener(s).
Listener Tag	This field is automatically populated with a unique name for the new Oracle Listener resource instance on the primary server. The default naming pattern will be displayed for you. You may type in another unique name. The valid characters allowed for the Listener tag are letters, digits and the following special characters – : . /

- Click **Extend**. The **Hierarchy Integrity Verification** window displays with the following message, **Hierarchy Verification Finished**. Click **Next Server** or **Finish**.

## 6.12.5.2.1. Updating the Oracle Listener Protection Level

---

1. Select a **resource** and then the  button from the **Resource** toolbar to update the protection level of the resource.
2. Enter the following information.

Field	Tips
<b>Listener Protection Level</b>	Select one of the following: Full Control (Start, Stop, Monitor and Recover) Intermediate Control (Start, Monitor and Recover) Minimal Control (Start and Monitor Only)

3. Click **Update** to change the **Protection Level** from the current state to the new state. Select **Cancel** to leave the value unchanged.

## 6.12.5.2.2. Updating the Oracle Listener Recovery Level

1. Select a **listener** and then the  button from the **Resource** toolbar to update the recovery level of the resource.
2. Enter the following information.

Field	Tips
<b>Listener Recovery Level</b>	<p>Select the level of recovery for the specified listener(s).</p> <p>Standard (On_ – enables a standard LifeKeeper recovery. If Standard, (On) is selected, all listener failures will be tried locally and if necessary trigger a fail over to an available backup server.</p> <p>Optional (Off) – enables optional LifeKeeper recovery. If Optional (Off) is selected, all listener failures will be tried locally, but will not cause a fail over to an available backup server.</p> <p><b>Note:</b> Local recovery is performed for both recovery levels when a listener error occurs; however, execution of failover depends on the recovery level.</p>
<b>Update Confirmation</b>	<p>Select the Update button to change the Recovery Level from the current state to the new state.</p> <p>Select Cancel to leave the current value unchanged.</p>

3. Click **Update** to change the Recovery Level from the current state to the new state. Select **Cancel** to leave the current value unchanged.

\* **Note:** The **FAILOVER** parameter in `tnsnames.ora` is independent of the LifeKeeper failover or recovery. The moving of connections will be performed by the Oracle software, while the recovery of the listener is a responsibility of the LifeKeeper Oracle Listener Recovery Kit. They are independent of each other.

\* **Note:** Failover is not suppressed in either LifeKeeper or Oracle by setting the “**FAILOVER**” parameter to “on”, there should not be any need to implement special parameters or settings within LifeKeeper to accommodate the failover within oracle. Failover in LifeKeeper does NOT occur in the event of a failed client connection. It occurs after a quickCheck and an event based error.

## 6.12.5.2.3. Updating the Oracle Protected Listener(s)

1. Select a **listener** and then the  button from the **Resource** toolbar to update your protected listener(s).
2. Enter the following information.

Field	Tips
<b>Listener Name(s)</b>	Select the name or names of the Oracle Listener(s) to provide protection for with this resource instance.
<b>IP Address Name(s)</b>	Select the IP Address resource name(s) that will be added as a dependent of this resource hierarchy.  Select None if no additional IP Resources are required for this configuration.

3. Click **Update** to change the **Protected Listener(s)** and **IP assignment** from the current state to the new state. Select **Cancel** to leave the current value unchanged.

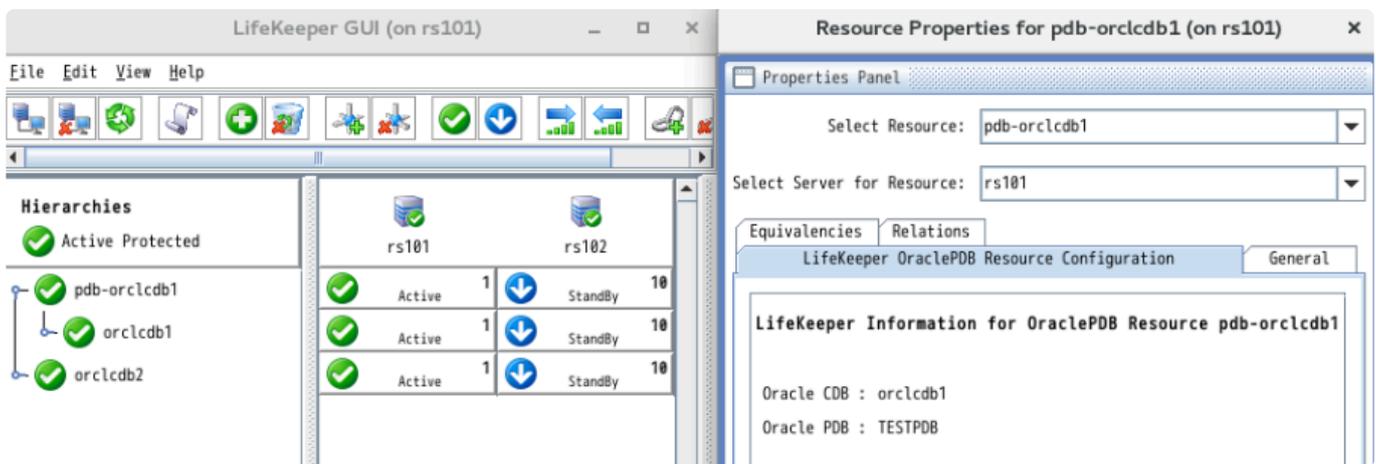
If you select **Update**, a dialog displays stating that the Protected Listeners for the specific resource is being updated. Click **Finish**.

## 6.12.5.3. Migrating a Pluggable Database

This section describes the procedure for migrating a pluggable database (PDB) between container databases (CDB) protected by LifeKeeper. We recommend to make a backup of the database in advance.

### How to Migrate

This section describes how to migrate a PDB using the PDB plug/unplug method. The figure below shows an example of a configuration for migrating a PDB (TESTPDB) from the source CDB (ORCLCDB1) to the target CDB (ORCLCDB2).

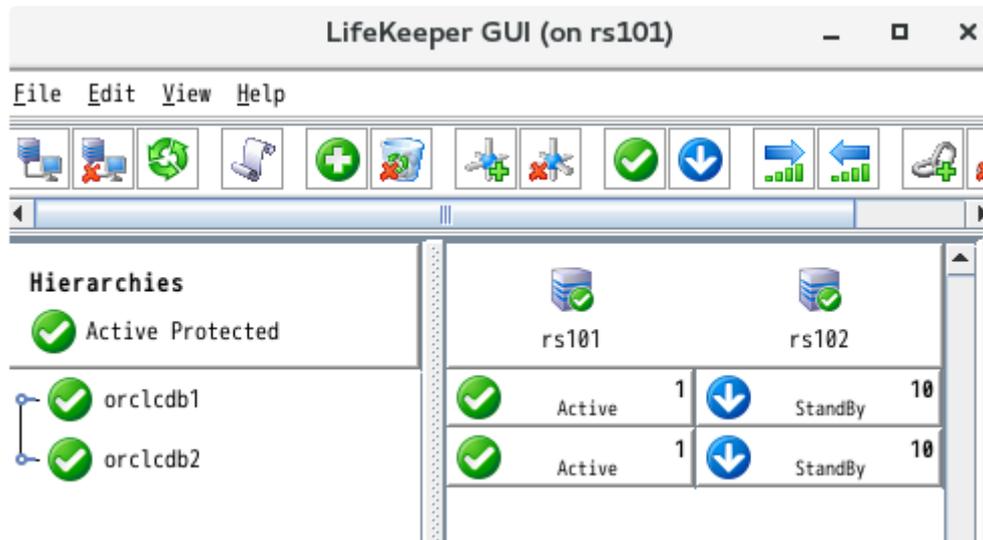


1. Before migration, connect to the source CDB and check the information of the PDB to migrate.

```
SQL> COLUMN NAME FORMAT A8
```

```
SELECT NAME, DBID, GUID FROM V$CONTAINERS WHERE NAME='<PDB>';
```

2. Bring the Oracle PDB resource out of service and delete the resource. If you have an Oracle PDB resource that protects more than one PDB, remove the PDB from protection from the resource setting Change Protection PDB (see [Changing the PDB to Protect](#)).



3. Connect to the source CDB and unplug the PDB. (If the PDB is not stopped, stop the PDB.)

```
SQL> ALTER PLUGGABLE DATABASE <PDB> UNPLUG INTO '/home/oracle/<PDB>.xml';
```

```
SQL> DROP PLUGGABLE DATABASE <PDB>;
```

4. Connect to the target CDB, plug in the PDB and start it.

```
SQL> CREATE PLUGGABLE DATABASE <PDB> USING '/home/oracle/<PDB>.xml'
<COPY|NOCOPY>;
```

```
SQL> ALTER PLUGGABLE DATABASE <PDB> OPEN;
```

5. Verify that the PDB information matches before and after the migration.

```
SQL> COLUMN NAME FORMAT A8
```

```
SELECT NAME, DBID, GUID FROM V$CONTAINERS WHERE NAME='<PDB>';
```

6. Create an Oracle PDB resource specifying the target CDB. If you already have an Oracle PDB resource that protects PDB, you can add the PDB for protection from the resource settings Change Protection PDB (see [Changing the PDB to Protect](#)) to manage it with a single resource.

The screenshot displays the LifeKeeper GUI interface. The main window is titled "LifeKeeper GUI (on rs101)" and features a menu bar (File, Edit, View, Help) and a toolbar with various icons. On the left, a "Hierarchies" pane shows a tree structure: "Active Protected" (with a green checkmark), "orclcdb1" (with a green checkmark), "pdb-orclcdb2" (with a green checkmark), and "orclcdb2" (with a green checkmark). The central area shows a table of resources for servers rs101 and rs102. The table has columns for status (green checkmarks), mode (Active), priority (1), and target mode (StandBy), with a "10" in the final column. The right-hand pane is titled "Resource Properties for pdb-orclcdb2 (on rs101)" and contains a "Properties Panel" with dropdown menus for "Select Resource: pdb-orclcdb2" and "Select Server for Resource: rs101". Below these are tabs for "Equivalencies", "Relations", "LifeKeeper OraclePDB Resource Configuration", and "General". The "General" tab is active, showing "LifeKeeper Information for OraclePDB Resource pdb-orclcdb2" with the following details: "Oracle CDB : orclcdb2" and "Oracle PDB : TESTPDB".

Server	Status	Mode	Priority	Target Mode	Value
rs101	Active	Active	1	StandBy	10
rs102	Active	Active	1	StandBy	10
rs101	Active	Active	1	StandBy	10

## 6.13. PostgreSQL Recovery Kit Administration Guide

---

The LifeKeeper for Linux PostgreSQL Recovery Kit is an SQL compliant, object-relational database management system (ORDBMS) based on POSTGRES. Since its inception, PostgreSQL has become one of the most advanced open source relational database management systems.

The LifeKeeper for Linux PostgreSQL Recovery Kit provides a mechanism for protecting PostgreSQL instances within LifeKeeper. The PostgreSQL software, LifeKeeper Core and PostgreSQL Recovery Kit are installed on two or more servers in a cluster. Once the PostgreSQL database instance is under LifeKeeper protection, clients connect to the database using a LifeKeeper protected IP address. The LifeKeeper protected IP address must be created separately and a dependency made manually between the parent PostgreSQL resource instance and the child IP address resource. In the event that the PostgreSQL server fails, LifeKeeper will first attempt to recover it on the local server. If the local recovery fails, then LifeKeeper will fail over to a backup server.

[PostgreSQL Resource Hierarchy](#)

### LifeKeeper Documentation

The following LifeKeeper product documentation is available from the SIOS Technology Corp. website:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [Optional Recovery Kit Documentation](#)

### PostgreSQL Documentation

You can find the PostgreSQL documentation, including the *Administration Guide*, *User Guide* and *Reference Guide* at the following location on the web:

<http://www.postgresql.org/docs>

## 6.13.1. PostgreSQL Resource Hierarchy

The following example shows a typical PostgreSQL resource hierarchy:



The dependencies in the above example correspond to the following protected resources:

Resource	PostgreSQL Software Component
LKIP.EXAMPLE.COM	Protects the switchable IP address used for client connections
<i>var/lib/pgsql/data</i>	Protects the database data directory (PGDATA)
<i>var/lib/pgsql/exec</i>	Protects the PostgreSQL server and client executables (when executables are installed on a shared file system)
<i>var/lib/pgsql/log</i>	Protects the database log file directory (when the log path is located on a shared file system)
<i>var/lib/pgsql/pg_xlog</i>	Protects the database transaction log directory (PGDATA/pg_xlog) The transaction log directory is also referred to as Write-Ahead-Log directory.
<i>var/lib/pgsql/socket_path</i>	Protects the database socket directory (when the socket path is located on a shared file system).

In the event of failover, LifeKeeper will bring the file system, IP address and database resources (including all the resource dependencies) in service on a backup server. Clients will be disconnected and will need to re-connect to the server. Any SQL statement that has not been committed will need to be re-entered.

## 6.13.2. PostgreSQL Hardware and Software Requirements

---

Your LifeKeeper configuration must meet the following requirements prior to the installation of LifeKeeper for Linux PostgreSQL Recovery Kit. Please refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the installation and configuration of your LifeKeeper hardware and software.

### Hardware Requirements

- **Servers** – Servers should be configured in accordance with the requirements described in the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#).
- **IP Network Interface Cards** – Each server requires at least one Ethernet TCP/IP-supported network interface card. Remember, however, that best practice is for a LifeKeeper cluster to have at least two communication paths. Two separate LAN-based communication paths using dual independent sub-nets are recommended for heartbeats, and at least one of these should be configured as a private network. Using a combination of TCP and TTY heartbeats is also supported.

### Software Requirements

- **TCP/IP Software** – Each server in your LifeKeeper configuration requires TCP/IP Software.
- **PostgreSQL Software** – The same version of the PostgreSQL software must be installed on all servers in the cluster. The PostgreSQL software can be downloaded from one of the mirrors available at <http://www.postgresql.org/download>.
- **LifeKeeper software** – It is imperative that you install the same version of the LifeKeeper software and apply the same versions of the LifeKeeper software patches to each server in your cluster.
- **LifeKeeper for Linux PostgreSQL Recovery Kit** – The PostgreSQL Recovery Kit is provided on the LifeKeeper for Linux Installation Image File (`sps.img`) via ftp download. It is packaged, installed and removed via Red Hat Package Manager, rpm:

```
steeleye-ikPGSQL
```

## 6.13.3. PostgreSQL Configuration Considerations

---

This section contains information that you should consider before you start to configure and administer the PostgreSQL Recovery Kit.

[Using Mirrored File Systems with DataKeeper](#)

[Protecting PostgreSQL: Best Practices](#)

## 6.13.3.1. Protecting PostgreSQL Best Practices

---

In an Active/Standby configuration, the backup server is not actively running the PostgreSQL, but stands by in case the primary server experiences a failure. In an Active/Active configuration, each server is actively running a PostgreSQL instance, while acting as a backup for the other server in case of failure. The following list provides requirements that should be adhered to when protecting a PostgreSQL resource instance in an active/standby or active/active configuration.

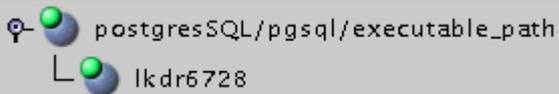
1. The PostgreSQL DataDir and Write-Ahead-LogPath (*PGDATA/pg\_xlog*) must be installed on one or more shared file systems. The paths *DataDir* and *WAL-Path* must be shared between all servers that will protect the resource instance.
  - The PostgreSQL Operating System User must own the data directory and directory containing the Write-Ahead-LogPath.
  - The PostgreSQL database must have been created using the utility *initdb*. The *initdb* utility must be run as the PostgreSQL owner using the *-D <datadir>* option.
  - The automatic startup of the default PostgreSQL instance must either be disabled or the default PostgreSQL instance must be restricted to running on a port other than those intended for use with LifeKeeper.
  - The automatic startup of the PostgreSQL instance to be protected by LifeKeeper must be disabled. LifeKeeper will control the starting and stopping of the protected instance.
  - The PostgreSQL instance must be started manually prior to hierarchy creation. It is required that the instance be started with the backend option *-o "-p <port>"* specified to the *pg\_ctl* utility.
2. The *StartupLogPath*, *SocketPath* and the *ExecutablePath* can be installed to optional shared file systems on the primary server or each local node file system.
  - The PostgreSQL Operating System User must own the directory containing the socket path.
  - The PostgreSQL Operating System User must have write permissions on the directory containing the *StartupLogPath*.
3. It is recommended that each instance use a unique port and socket path when running multiple instances in either an Active/Standby or Active/Active scenario.

## 6.13.3.2. Using Mirrored File Systems with DataKeeper

---

The PostgreSQL Recovery Kit supports the use of SIOS DataKeeper as a shared file system. The mirrored file systems can be used for the PostgreSQL installation path, log path, the data directory and the executable path.

For example, a dependent file system for a PostgreSQL resource would look similar to the following, which shows a file system for the data directory and its dependency, the DataKeeper resource mirror.



```
graph TD; postgresSQL["postgresSQL/pgsql/executable_path"] --- ikdr6728["ikdr6728"]
```

The diagram shows a dependency between two resources. The top resource is labeled "postgresSQL/pgsql/executable\_path" and has a small icon with a green dot. The bottom resource is labeled "ikdr6728" and also has a small icon with a green dot. A vertical line connects the two resources, indicating a dependency.

\* Replicated (SIOS DataKeeper) file system resources must be created before creating the PostgreSQL resource.

## 6.13.4. PostgreSQL Installation

---

### Installing/Configuring PostgreSQL with LifeKeeper

The following sequence is recommended for installing and configuring the PostgreSQL database and LifeKeeper software. Each of these steps links to detailed tasks.

1. [Install the PostgreSQL software.](#)
2. [Create the PostgreSQL database.](#)
3. [Install the LifeKeeper Core and PostgreSQL Recovery Kit.](#)
4. [Configure LifeKeeper Tunable Settings for PostgreSQL Resources.](#)

After you have performed these tasks, you will be ready to create the LifeKeeper resource hierarchy to protect your PostgreSQL database.

### Resource Configuration Tasks

Once you have completed the setup tasks described in the previous section, you are ready to create and extend your PostgreSQL resource hierarchies.

The following tasks are available for configuring the LifeKeeper for Linux PostgreSQL Recovery Kit:

- [Create Resource Hierarchy](#) – Creates a PostgreSQL resource hierarchy.
- [Delete Resource Hierarchy](#) – Deletes a PostgreSQL resource hierarchy.
- [Extend Resource Hierarchy](#) – Extends a PostgreSQL resource hierarchy from the primary server to the backup server.
- [Unextend Resource Hierarchy](#) – Unextends (removes) a PostgreSQL resource hierarchy from a single server in the LifeKeeper cluster.
- [Viewing PostgreSQL Configuration Settings](#) – Allows viewing of the Resource Properties dialog.

Refer to the [GUI Administrative Tasks](#) section of the [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring LifeKeeper Core resource hierarchies, for instance, file system and IP resources.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.

- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View Properties](#). View the properties of a resource hierarchy on a specific server.
- [Edit Properties](#). Edit the properties of a resource hierarchy on a specific server.

**Note:** The configuration tasks throughout this section are performed using the Edit menu. You may also perform most of the tasks:

- from the toolbar.
- by right-clicking on a global resource in the left pane of the status display.
- by right-clicking on a resource in the right pane of the status display.

Using the right-click method allows you to avoid entering information that is required using the **Edit** menu.

## Upgrading

[Upgrading From Previous Version of the PostgreSQL Recovery Kit](#)

## 6.13.4.1. Install the PostgreSQL Software

---

Install the PostgreSQL software on all servers in the cluster using identical parameters/settings. Refer to the [PostgreSQL Administration Guide](#) for details. The following are additional recommendations and reminders to ensure that LifeKeeper will work with PostgreSQL:

- The PostgreSQL client software packages must be installed. These packages must include the PostgreSQL `psql` client utility.
- The PostgreSQL server software packages must be installed. These packages must include the PostgreSQL `pg_ctl` and `initdb` utilities.
- The PostgreSQL client and server packages must be the same version on all servers.
- A PostgreSQL Operating System User must exist on all servers as follows:
  - This PostgreSQL Operating System User should be designated as the owner of the PostgreSQL software installation and subdirectories.
  - This PostgreSQL Operating System User must have authority to use the `pg_ctl` utility. The PostgreSQL Operating System User must be able to start and stop the postmaster instance using the `pg_ctl` commands.
  - The PostgreSQL Operating System User name should contain alphanumeric characters only.
  - The user id and group id of this PostgreSQL Operating System User must be identical on all servers.
- A PostgreSQL Database Administrator User must exist within the PostgreSQL database for LifeKeeper client connections through the `psql` utility.
  - This PostgreSQL Database Administrator User must have the ability to connect to the database (*template1*), as well as obtain the listing of defined databases for the instance.
  - This PostgreSQL Database Administrator User must have the ability to view system tables and make generalized queries.
  - The PostgreSQL Database Administrator User is different from the PostgreSQL Operating System User, although they can have the same name.
  - Example: PostgreSQL Operating System User=`postgres`, and PostgreSQL Database Administrator User=`lkpostgres`; or PostgreSQL Operating System User=`postgres`, and PostgreSQL Database Administrator User=`postgres`.
- Auto Startup at the time of system activation must be disabled because PostgreSQL server daemon is controlled by LifeKeeper.

## 6.13.4.2. Creating a PostgreSQL Database

Follow the instructions in your [PostgreSQL Administration Guide](#) to create your database. In addition, please note the following recommendations:

✿ Replicated (SIOS DataKeeper) file system resources must be created before creating the PostgreSQL resource.

- The PostgreSQL data directory should be initialized using the `initdb` utility, specifying the `-D <data_dir>` option. The `initdb` command must be run as the PostgreSQL Operating System User.
- The PostgreSQL instance data directory must reside on a shared file system.
- The PostgreSQL transaction log directory must reside on a shared file system.
- The PostgreSQL database name should contain alphanumeric characters only.
- After creating your database, you should disable automatic startup of the PostgreSQL database instance. Once under LifeKeeper protection, LifeKeeper will handle the start and stop of the database.
- The PostgreSQL instance must be started manually prior to hierarchy creation. It is required that the instance be started with the backend option `-o "-p <port>"` specified to the `pg_ctl` utility.

No Password Protection (Instance is not Password Protected)

- If the PostgreSQL database instance will not be password protected or will not require a password for local client connections from the PostgreSQL Database Administrator User, then an entry must exist allowing local trust connections. The following is an example of a `pg_hba.conf` entry to enable local client connects for the PostgreSQL Database Administrator User:

```
=====
.
.
Local all postgres trust
.
.
=====
```

Enabling Password Protected (Instance requires a Password for Connections)

- Password Protected database instances require a password entry for the PostgreSQL Database Administrator User to exist in the `.pgpass` credentials file on each server in the cluster where the resource will be protected. The `.pgpass` file must contain a valid and tested entry for each

PostgreSQL Database Administrator User requiring a password.

- The `.pgpass` file must be located in the home directory of the PostgreSQL Operating System User. Please set the appropriate file permissions to restrict access to the file.
- The following is an example of a valid `.pgpass` file with the format

```
<hostname>:<port>:<database>:<user>:<password>
```

```
=====  
*:5443:*:lifekeeper:jh43tmp2009  
=====
```

**Note:** The `.pgpass` file is required for the utility `psql` for unattended (non-terminal or scripted) connections. The `.pgpass` file must exist on each server where the password protected instance will be protected.

## 6.13.4.3. Install the LifeKeeper Software

---

Once you have installed the PostgreSQL software and created your database, you are ready to install the LifeKeeper Core software and any required patches followed by the PostgreSQL Recovery Kit.

Refer to the [LifeKeeper for Linux Installation Guide](#) for details on installing the LifeKeeper packages.

## 6.13.4.4. LifeKeeper Tunable Settings for PostgreSQL

---

The PostgreSQL Recovery Kit provides tunable environment variables to help customize resource protection in certain scenarios. To change the values of these variables, edit the file `/etc/default/LifeKeeper`. No processes need to be restarted for the new settings to take effect. The default values will work for most environments where the PostgreSQL Recovery Kit will be installed.

### LKPGSQL\_CONN\_RETRIES

This tunable controls the amount of time the PostgreSQL Recovery Kit will wait for the database to start. The amount of time is calculated by the Recovery Kit using the following formula:  
(LKPGSQL\_CONN\_RETRIES\* 5) = total time in seconds to wait for a database instance to start.  
The setting of this variable affects both the resource in-service requests and the resource local recovery.

### LKPGSQL\_DISCONNECT\_CLIENT

This tunable controls whether active clients will be disconnected in the event of a postmaster crash. When the value is set to 1 (true), active clients will be disconnected while resource local recovery is in progress. When the value is set to 0 (false), active clients will not be disconnected while resource local recovery is in progress. This variable affects only the resource local recovery events and is only applicable during local recovery events where the postmaster process is not running.

### LKPGSQL\_SDIRS

This tunable controls the client disconnect behavior when the PostgreSQL database is shut down. This comma separated tunable must be added to the defaults file. By setting this option, the specified resource instance or instances corresponding to the protected data directory will not force clients to disconnect during shutdown.

```
LKPGSQL_SDIRS=/protected/pgsql-datadir
```

```
LKPGSQL_SDIRS=/protected/pgsql-datadir,/otherprotected/pgsql-datadir
```

Where `/protected/pgsql-datadir` and `/otherprotected/pgsql-datadir` are the PostgreSQL data directories under LifeKeeper protection.

**Note:** The options LKPGSQL\_SDIRS and LKPGSQL\_IDIRS are exclusive. The value placed in the LKPGSQL\_SDIRS or LKPGSQL\_IDIRS tunable must match exactly with the protected datadir value selected during hierarchy creation.

## LKPGSQL\_IDIRS

This tunable controls the client disconnect behavior when the PostgreSQL database is shut down. This comma separated tunable must be added to the defaults file. By setting this option, the specified resource instance or instances corresponding to the protected data directory will force clients to do an immediate disconnect during shutdown.

```
LKPGSQL_IDIRS=/protected/pgsql-datadir
```

```
LKPGSQL_IDIRS=/protected/pgsql-datadir,/otherprotected/pgsql-datadir
```

Where */protected/pgsql-datadir* and */otherprotected/pgsql-datadir* are the PostgreSQL data directories under LifeKeeper protection.

**Note:** The options LKPGSQL\_SDIRS and LKPGSQL\_IDIRS are exclusive. The value placed in the LKPGSQL\_SDIRS or LKPGSQL\_IDIRS tunable must match exactly with the protected datadir value selected during hierarchy creation.

## 6.13.4.5. Creating a PostgreSQL Resource Hierarchy

Perform the following steps on the primary server:

1. On the **Edit** menu, select **Server**, then **Create Resource Hierarchy**.  
The **Create Resource Wizard** dialog will appear.
2. Select **PostgreSQL Database** from the drop-down list and click **Next**.
3. You will be prompted for the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter any error requiring you to correct the previously entered information. You may click **Cancel** at any time to cancel the entire creation process.

Fields	Tips
<b>Switchback Type</b>	Choose either <b>intelligent</b> or <b>automatic</b> . This determines how the PostgreSQL resource will be switched back to the primary server after it comes in-service (active) on the backup server following a failover. Intelligent switchback requires administrative intervention to switch the resource back to the primary server, while automatic switchback occurs as soon as the primary server is back on line and re-establishes LifeKeeper communication paths.  <b>Note:</b> The switchback strategy must match that of the dependent resources to be used by the PostgreSQL resource.
<b>PostgreSQL Executable Location</b>	This field is used to specify the directory path containing the PostgreSQL executables. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
<b>PostgreSQL Client Executable Location</b>	This field is used to specify the directory path containing the PostgreSQL executable <code>psql</code> . The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
<b>PostgreSQL Administration Executable Location</b>	This field is used to specify the directory path containing the PostgreSQL executable <code>pg_ctl</code> . The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
<b>PostgreSQL Data Directory</b>	This field is used to specify the location of the PostgreSQL data directory ( <i>datadir</i> ) that will be placed under LifeKeeper protection. The specified directory must exist and reside on a shared file system. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
<b>PostgreSQL Port</b>	This field is used to specify the TCP/IP port number on which the postmaster daemon is listening for connections from client applications.

<b>PostgreSQL Socket Path</b>	This field is used to specify the full path to the Unix-domain socket on which the postmaster daemon is listening for connections from client applications. The valid characters allowed for the pathname are letters, digits and the following special characters: – _ . /
<b>PostgreSQL Database Administrator User</b>	This field is used to specify a PostgreSQL Database Administrator User name for the specified database instance with connection and administrator privileges for the instance.
<b>PostgreSQL Logfile</b>	This field is used to specify the log file path used by the -l option of pg_ctl to start and stop PostgreSQL.
<b>PostgreSQL Database Tag</b>	This is a unique tag name for the new PostgreSQL database resource on the primary server. The default tag name consists of the word pgsql followed by the port number for the database instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits and the following special characters: – _ . /

4. Click **Create**. The **Create Resource Wizard** will then create your PostgreSQL resource hierarchy. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.
5. You should see a message indicating that you have successfully created a PostgreSQL resource hierarchy, and you must extend that hierarchy to another server in your cluster to achieve failover protection. Click **Next**.
6. Click **Continue**. LifeKeeper will then launch the **Pre-extend Wizard**. Refer to **Step 2** in the topic [Extending a PostgreSQL Resource Hierarchy](#) for details on how to extend your resource hierarchy to another server.

## 6.13.4.6. Deleting a PostgreSQL Resource Hierarchy

---

To delete a PostgreSQL resource hierarchy from all servers in your LifeKeeper configuration, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your PostgreSQL resource hierarchy.

 **Note:** If you selected the **Delete Resource** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

3. Select the **Hierarchy to Delete**. (This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a resource instance in the left or right pane.) Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Next**.
5. Another information box appears confirming that the PostgreSQL resource was deleted successfully.
6. Click **Done** to exit.

## 6.13.4.7. Extending a PostgreSQL Resource Hierarchy

This operation can be started from the **Edit** menu or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should refer to **Step 2** below.

1. On the **Edit** menu, select **Resource**, then **Extend Resource Hierarchy**. The **Pre-Extend Wizard** appears. If you are unfamiliar with the **Extend** operation, click **Next**. If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information.

**Note:** The first two fields appear only if you initiated the **Extend** from the **Edit** menu.

Field	Tips
<b>Template Server</b>	Select the server where your PostgreSQL resource is currently in service.
<b>Tag to Extend</b>	Select the PostgreSQL resource you wish to extend.
<b>Target Server</b>	Enter or select the server you are extending to.
<b>Switchback Type</b>	<p>This determines how the PostgreSQL resource will be switched back to the primary server after it comes in service (active) on the backup server following a failover. You can choose either <b>intelligent</b> or <b>automatic</b>. The switchback type can be changed later, if desired, from the <b>General</b> tab of the <b>Resource Properties</b> dialog box.</p> <p><b>Note:</b> Remember that the switchback strategy must match that of the dependent resources to be used by the PostgreSQL resource.</p>
• <b>Template Priority*</b>	<p>Select or enter a <b>Template Priority</b>. This is the priority for the PostgreSQL hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.</p> <p><b>Note:</b> This selection will appear only for the initial extend of the hierarchy.</p>
<b>Target Priority</b>	<p>This is the priority for the new extended PostgreSQL hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid indicating a server's priority in the cascading failover sequence for the resource. Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.</p>

3. After receiving the message that the pre-extend checks were successful, click **Next**.
4. Depending upon the hierarchy being extended, LifeKeeper will display a series of information boxes showing the Resource Tags to be extended, some of which cannot be edited.
5. The **Extend Wizard** will prompt you to enter the following information.

Field	Tips
<b>PostgreSQL Executable Location</b>	This field is used to specify the directory path containing the PostgreSQL executables. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
<b>PostgreSQL Database Tag</b>	This is a unique tag name for the new PostgreSQL database resource on the primary server. The default tag name consists of the word pgsq followed by the port number for the database instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits and the following special characters: - _ . /

6. After receiving the message "Hierarchy extend operations completed", click **Next Server** to extend the hierarchy to another server, or click **Finish** if there are no other extend operations to perform.
7. After receiving the message "Hierarchy Verification Finished", click **Done**.

## 6.13.4.8. Unextending a PostgreSQL Resource Hierarchy

---

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. On the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the PostgreSQL resource. It cannot be the server where the resource is currently in service. (This dialog box will not appear if you selected the Unextend task by right-clicking on a resource instance in the right pane.) Click **Next**.
3. Select the PostgreSQL hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the Unextend task by right-clicking on a resource instance in either pane).
4. An information box appears confirming the target server and the PostgreSQL resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the PostgreSQL resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

## 6.13.4.9. Viewing PostgreSQL Configuration Settings

---

The **Resource Properties** dialog is available from the **Edit** menu or from a resource context menu. This dialog displays the properties for a particular resource on a server. When accessed from the **Edit** menu, you can select the resource and the server. When accessed from a **Resource Context** menu, you can select the server.

From the **Configuration** tab, you can view the following PostgreSQL settings:

- Executable Path
- Client Executable Name
- Admin Executable Name
- Bind Setting
- Startup Log Location
- PostgreSQL Operating System User Name
- PostgreSQL Database Administrator User
- Version Number
- Data Directory
- Socket Location
- Port Number
- OS Daemon Name

## 6.13.4.10. Upgrading PostgreSQL

---

During an upgrade from a previous version of the LifeKeeper for Linux PostgreSQL software, the upgrade will make modifications to the existing LifeKeeper PostgreSQL resource instance. When the LifeKeeper software is updated on the server, the following stored values will be added to the internal LifeKeeper information field automatically.

- **Client Executable Location (*psql*)** – the location of the *psql* or equivalent client utility used for connecting to the protected database instance. After an upgrade, this value can be verified from the LifeKeeper GUI properties display. The value can also be verified from the LifeKeeper command line using the `set_value` utility.

`set_value` is the name of a LifeKeeper utility provided for the LifeKeeper PostgreSQL Recovery Kit to update the internal resource information field values. The use of this utility should be limited to issues explained in this topic or at the request and instruction of the SIOS Technology Corp. Support team.

**Note:** The `set_value` utility does not perform rigorous error checking and therefore is not intended for general use.

- **Administration Executable Location (*pg\_ctl*)** – the location of the *pg\_ctl* or equivalent administration utility used for starting, stopping and checking the status of the protected database instance. After an upgrade, this value can be verified from the LifeKeeper GUI properties display. The value can also be verified from the LifeKeeper command line using the `set_value` utility.
- **PostgreSQL Database Administrator User** – the PostgreSQL Database Administrator User for the LifeKeeper protected instance. This user must have connection and administrator privileges for the protected database instance. The default value used following an upgrade is the PostgreSQL Operating System User that owns the PostgreSQL data directory. After an upgrade, this value can be verified from the LifeKeeper GUI properties display. The value can also be verified from the LifeKeeper command line using the `set_value` utility.
- **PostgreSQL Daemon Name (*postmaster*)** – the name of the running backend daemon. This value is determined during the first status check of the database instance. The default value is `postmaster`. After an upgrade, this value can be verified from the LifeKeeper GUI properties display. The value can also be verified from the LifeKeeper command line using the `set_value` utility.
- **Default Test Database (*template1*)** – the database used by LifeKeeper during the database instance monitoring to verify basic connectivity. After an upgrade, the default test database will be set to `template1`.
- **PostgreSQL Maximum Monitoring Hangs ([LKPGSQL\\_QCKHANG\\_MAX](#))** – the setting that provides protection against an unlimited number of connection hangs before a restorative or reparative failover action is initiated. A portion of PostgreSQL Recovery Kit's monitoring requires a connection to the protected database. The number of connection hangs allowed is determined during resource creation by the setting [LKPGSQL\\_QCKHANG\\_MAX](#). The default value previous to

version 8.1.2 was 15. After upgrading to version 8.1.2 (or later), the default value is 2. Since this value is stored with the resource at create time, any resources created prior to upgrading to version 8.1.2 will remain at the default value of 15 unless updated by the user while any resources created after upgrading to 8.1.2 (or later) will contain a default value of 2. The value can also be verified from the LifeKeeper command line using the `set_value` utility.

**\* IMPORTANT NOTE:** Following the upgrade of the LifeKeeper for Linux PostgreSQL Recovery Kit software, you should [test your PostgreSQL resource hierarchy](#) by initiating a manual switchover that will simulate a failover of the resource instance from the primary server to a backup server.

## Important Upgrade Considerations

If a resource does not come into service following the upgrade, check the following conditions:

- **Client Executable name is not found or incorrect**

The value can be updated using the `set_value` utility. The syntax for the Client Executable update is as follows:

```
/opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set_value <tag>
'clientexe' <full path to the psql utility>.
```

**Example:** `/opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set_value pgsq1-5443 'clientexe' '/pgsql/clientutils/psql'.`

- **Administration Executable name is not found or incorrect**

The value can be updated using the `set_value` utility. The syntax for the Administration Executable update is as follows:

```
/opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set_value <tag>
'osexex' <full path to the pg_ctl utility>.
```

**Example:** `/opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set_value pgsq1-5443 'osexex' '/pgsql/adminutils/pg_ctl'.`

### Lowering the interval for recovering from multiple hang events ([LKPGSQL\\_QCKHANG\\_MAX](#))

- **Maximum Monitoring Hangs value is too large in versions prior to 8.1.2**

The value for **Maximum Monitoring Hangs** for existing PostgreSQL resource instances can be viewed or set using the `set_value` utility.

The syntax for setting the value for the **Maximum Monitoring Hangs** ([LKPGSQL\\_QCKHANG\\_MAX](#)) is as follows:

```
/opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set_value <tag>  
'hangmax' <number>.
```

**Example:** /opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set\_value  
pgsql-5443 'hangmax' 3.

**Note:** Include the -c argument to update the value on all nodes in the cluster (set\_value  
-c <tag>...).

The syntax for **viewing** the value is as follows:

```
/opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set_value -l <tag>  
'hangmax'
```

**Example:**/opt/LifeKeeper/lkadm/subsys/database/pgsql/bin/set\_value -l  
pgsql-5443 'hangmax'

## 6.13.5. PostgreSQL Administration

---

### [Updating Database Administrator User](#)

The [Update User](#) option allows the LifeKeeper administrator to change the current PostgreSQL Database Administrator User for the LifeKeeper PostgreSQL resource instance.

### Testing Your PostgreSQL Resource Hierarchy

You can test your PostgreSQL resource hierarchy by initiating a manual switchover that will simulate a failover of the resource instance from the primary server to a backup server.

#### [Performing a Manual Switchover from the LifeKeeper GUI](#)

### EnterpriseDB Postgres Plus Advanced Server Environments

#### [Protecting EnterpriseDB Postgres Plus Advanced Server Resources](#)

### Symfoware Server/Enterprise Postgres Environments

#### [Protecting Symfoware Server/Enterprise Postgres Resources](#)

## 6.13.5.1. Performing a Manual Switchover from the LifeKeeper GUI

---

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource, and In Service**. For example, an in-service request executed on a backup server causes the PostgreSQL resource hierarchy to be placed in service on the backup server and taken out-of-service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the Out of Service request, the resource hierarchy is taken out-of-service without bringing it in service on the other server.

 **Important:** After bringing your resource hierarchy in service on the backup server, you should attempt to connect to the databases. With password protected instances, it is of particular importance that the `.pgpass` file is verified on the backup server. To verify the `.pgpass` file is valid, a client connection to the database should be made using both the `psql` utility and the PostgreSQL Database Administrator User. A valid `.pgpass` file exists if the connection succeeds without prompting for an interactive password.

## 6.13.5.2. Protecting EnterpriseDB Postgres Plus Advanced Server

No additional LifeKeeper configuration settings are needed to protect EnterpriseDB Postgres Plus Advanced Server Resources.

Issue	Solution
<p>During the installation of EnterpriseDB Postgres Plus Advanced Server, if the option <b>PostgreSQL-compatible defaults and samples</b> is chosen in the <b>Configuration Mode</b> dialog, the 'edb' database that is used by LifeKeeper is not created.</p>	<p>Manually add the 'edb' database using the utility 'createdb'.</p> <p>The command 'createdb -p &lt;port&gt; -h &lt;socket path&gt; edb' should be executed as the PostgreSQL Operating System User. The following is an example:</p> <pre>su - postgres  postgres@server1 ~&gt;createdb -p 5435 -h /var/lib/postgres edb</pre>

## 6.13.5.3. Protecting Symfoware Server/Enterprise Postgres

The following table explains the supported features when protecting Symfoware Server/Enterprise Postgres.

The support range
<ul style="list-style-type: none"><li>• Support the compatible functions with PostgreSQL.</li><li>• The mirroring functionality is not supported. Use DataKeeper instead.</li><li>• The native interface of Symfoware Server is not supported. Use Open Interface (Symfoware 12.2), Postgres (Symfoware 12.3 or later).</li><li>• WebAdmin is asymmetry. For the details, refer to Symfoware Server Cluster Operation Guide for Fujitsu Software.</li><li>• Following functions are not supported:<ul style="list-style-type: none"><li>- WAL duplication</li><li>- Encryption (encryption of stored data)</li><li>- Data concealing</li><li>- Parallel search</li><li>- In-memory</li></ul></li></ul>
Notes for the configuration
Set up the environment variable, LD_LIBRARY_PATH required to execute Symfoware commands (pg_ctl,psql, etc.) in the appropriate environment file (.bash_profile, etc.) for the DB Administration User (OS user) log-in.

## 6.13.5.4. Updating Database Administrator User

This **Update User** option will update the stored value for the PostgreSQL Database Administrator User on all systems where the resource is protected. The **Update User** option can be invoked from either the **LifeKeeper resource toolbar** or the **LifeKeeper resource context menu**.

To update the PostgreSQL Database Administrator User, perform the following steps on the primary server:

**Note:** The **Update User** menu and toolbar options will be disabled for any out-of-service resources.

1. On the toolbar, select the **Update User** icon or select **Update User** from the resource context menu.

The **Update User Wizard** dialog will appear.

2. You will be prompted for the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter any error requiring you to correct the previously entered information. You may click **Cancel** at any time to cancel the entire creation process.

Field	Tips
Enter <b>PostgreSQL Database Administrator User</b>	This dialog requests a PostgreSQL Database Administrator User name for the specified database instance with connection and administrator privileges for the instance.  <b>Note:</b> A validation script will verify connectivity using the value specified. A password protected instance will require a valid entry in the <code>.pgpass</code> file for the PostgreSQL Database Administrator User.
<b>Confirm Update Action</b>	This dialog requests confirmation of the update user change of the previous user value to the new user value.

3. Click **Update**. The PostgreSQL Database Administrator User will be updated on all servers where the resource is currently protected.

## 6.13.6. PostgreSQL Troubleshooting

---

This section provides a list of messages that you may encounter while creating and extending a LifeKeeper PostgreSQL resource hierarchy or removing and restoring a resource. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Messages from other LifeKeeper components are also possible. In these cases, please refer to the [Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

- [General Tips](#)
- [Tunables](#)

## 6.13.6.1. PostgreSQL General Tips

The following error messages and conditions may be encountered while using the recovery kit.

Error	Solution
Unable to protect PostgreSQL database using the same port as another LK protected PostgreSQL database.	<p>Verify the version of PostgreSQL includes a <i>postgresql.conf</i> file. In the <i>postgresql.conf</i> file, set the entry <code>listen_address=</code> to the IP address to be used with the database instance.</p> <p><b>Note:</b> The format of the <code>listen_address=</code> in the <i>postgresql.conf</i> file is important as syntax errors can result in a failure to start the database server.</p>
Unable to perform a manual switchover of version 8.X when clients are connected.	The default (smart) shutdown option fails to disconnect clients on a switchover. If shutdown continues to fail with connected clients, verify that the <a href="#">LKPGSQL_SDIRS</a> tunable is not set. If the problem persists, set the LifeKeeper tunable <code>LKPGSQL_IDIRS</code> .
Unable to connect from a remote client to the database server.	To enable remote host login for PostgreSQL, refer to the <i>PostgreSQL Administration Guide</i> on configuring the <i>pg_hba.conf</i> file.
psql: connectDBStart() — connect() failed: No such file or directory. Is the postmaster running at 'localhost' and accepting connections on Unix socket '<port>'?"	Verify that the socket file exists and the instance is currently running. If the socket file resides in <i>/tmp</i> , it may have been removed by a cron job that cleans up the <i>/tmp</i> directory. Take the resource out of service and back in service. Then modify the cron job to leave PostgreSQL socket files.
PostgreSQL resource hierarchy fails to come in service but the database is running.	The database may have failed to respond to the LifeKeeper client request within the specified interval. Adjust the tunable <a href="#">LKPGSQL_CONN_RETRIES</a> in <i>/etc/default/LifeKeeper</i> to increase the number of seconds allowed for the recovery and restart of the PostgreSQL database instance.
PostgreSQL resource hierarchy fails local recovery following a postmaster crash with active client connections.	When a large number of active clients are connected to PostgreSQL, the database may be unable to properly restart until the client connections have terminated. In this scenario, it may be best to force client connections to terminate so that local recovery will be successful. The variable <a href="#">LKPGSQL_DISCONNECT_CLIENT</a> can be set in <i>/etc/default/LifeKeeper</i> to control the behavior of the PostgreSQL resource hierarchy in this scenario. When the value is set to 1( <b>true</b> ), client processes will be sent a SIGTERM signal to force them to disconnect from the database. This action will only be taken if the postmaster process is not running during local recovery.
Unable to connect to database with error "WARNING: password file "/home/<osuser>/.pgpass" has world or group read	The <i>.pgpass</i> file permissions should be <code>u=rw(0600)</code> . Change the permissions and owner of the <i>.pgpass</i> file.

access”	
FATAL: syntax error in file “/datadir/postgresql.conf” line 50, near token “.17”	The <i>postgresql.conf</i> file listen_address= entry does not contain proper quoting. Verify entries are valid and the entry is enclosed in proper quotes.

## 6.13.6.2. PostgreSQL Tunables

Error	Solution
LKPGSQL_KILLPID_TIME	Time to wait after a process id is killed before rechecking for this process.
LKPGSQL_CONN_RETRIES	Replaces LKPGSQLMAXCOUNT – number of times to try a client connection after an action (start or stop)
LKPGSQL_ACTION_RETRIES	Number of times to attempt start or stop action before failing the action command.
LKPGSQL_STATUS_TIME	Timeout for status command.
LKPGSQL_QCKHANG_MAX	Number of quickCheck script hangs allowed before a failover/sendevent is triggered for the database instance.
LKPGSQL_CUSTOM_DAEMON	Allows a user to specify additional aliases for the postgres daemons (default postmaster).
LKPGSQL_IDIRS	Replaces LKPGSQL_IPORTS – Contains datadir entries for instances that will be shutdown using the immediate option only.
LKPGSQL_SDIRS	Contains datadir entries for instances that will be shutdown using the smart option.
LKPGSQL_DISCONNECT_CLIENT	<p>Controls the behavior the PostgreSQL resource hierarchy during a database failure scenario. When the value is set to 1(<b>true</b>), client processes will be sent a SIGTERM signal to force them to disconnect from the database. This action will only be taken if the postmaster process is not running during local recovery.</p> <p><b>Note:</b> This parameter cannot be used for PostgreSQL 8.2 and later.</p>
LKPGSQL_DISCONNECT_CLIENT_BYTAG	<p>Similar to LKPGSQL_DISCONNECT_CLIENT, this setting limits the action to the comma separated list of tags specified by this tunable.</p> <p><b>Note:</b> This parameter cannot be used for PostgreSQL 8.2 and later.</p>
LKPGSQL_RESUME_PROC	Determines if process found in the stopped state (state = ~T) will be resumed when detected or ignored.
LKPGSQLDEBUG	Turns on debug for PostgreSQL database kit as well as for the postgres database. Valid entry range: 0 – 5. Larger numbers produce greater debug information.

	This tunable will be passed on to the postmaster database using the option <code>-d &lt;LKPGSQLDEBUG&gt;</code> .
--	-------------------------------------------------------------------------------------------------------------------

## 6.14. Postfix Recovery Kit Administration Guide

---

Postfix plays a variety of roles, all critical to the proper flow of email. It listens on the network for incoming mail, transports mail messages to other servers, and delivers local mail to a local program.

The LifeKeeper for Linux Postfix Recovery Kit provides a mechanism to recover Postfix from a failed primary server to a backup server in a LifeKeeper environment. Both LifeKeeper and Postfix ensure data integrity throughout the course of the failover process without significant lost time or human intervention.

### Document Contents

This guide contain the following topics:

- [Documentation and References](#). Provides a list of LifeKeeper for Linux documentation and where to find them
- [Requirements](#) A description of the hardware and software necessary to properly setup, install, and operate the Postfix Recovery Kit. Refer to [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove LifeKeeper for Linux software.Postfix Recovery Kit.
- [Configuring the LifeKeeper for Linux Postfix Recovery Kit](#). A description of the procedures required to properly configure the Postfix Recovery Kit.
- [Postfix Configuring Validation](#). Provides steps for validating the Postfix configuration prior to creating the Postfix resource hierarchy.
- [LifeKeeper Configuration Tasks](#). A description of the tasks for creating and managing your Postfix resource hierarchies using the LifeKeeper GUI.
- [Create a Dependency with the Mailbox Spool Resource](#). Describes how to manually create a dependency between the Postfix resource and the Mailbox Spool file system resource.
- [Testing Your Resource Hierarchy](#). Describes steps for testing your Postfix resource hierarchies using the LifeKeeper GUI and command-line interface.
- [Troubleshooting](#). A list of LifeKeeper for Linux error messages including a description for each.

### Documentation and References

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)

- [LifeKeeper for Linux Installation Guide](#)
- [Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is available on the SIOS Technology Corp. website at:

<http://docs.us.sios.com>

## Postfix Documentation and References

The following is a list of reference documents associated with the Postfix application and the LifeKeeper Postfix Recovery Kit:

- Postfix Man Page
- Red Hat Postfix Reference Manual

## 6.14.1. Postfix Hardware and Software Requirements

---

Your LifeKeeper configuration must meet the following requirements **prior** to the installation of the LifeKeeper for Linux Postfix Recovery Kit. Please see [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the configuration of your LifeKeeper hardware and software.

### Hardware Requirements

- **Servers** – LifeKeeper for Linux supported servers configured in accordance with the requirements described in [LifeKeeper for Linux Installation Guide](#) and [LifeKeeper for Linux Release Notes](#).
- **Data Storage** – The Postfix Recovery Kit can be used in conjunction both with shared storage and with replicated storage provided by the LifeKeeper Data Replication product.

### Software Requirements

- **TCP/IP** software. Each server also requires the TCP/IP software.
- **LifeKeeper software**. You must install the same version of LifeKeeper software and any patches on each server.
- **LifeKeeper for Linux IP Recovery Kit**. You must have the same version of the LifeKeeper IP Recovery Kit on each server.
- **IP network interface**. Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

**Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation, and so forth.

- **Postfix software**. Each server must have the Linux distribution version of the Postfix software installed and configured before you can configure LifeKeeper and the Postfix Recovery Kit. The same version should be installed on each server. Consult the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information.

## 6.14.1.1. Postfix Recovery Kit Installation

---

Please refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on the installation and removal of the LifeKeeper for Linux software, including recovery kits.

## 6.14.2. Configuring the LifeKeeper for Linux Postfix Recovery Kit

---

This section describes the LifeKeeper for Linux Postfix Recovery Kit configuration details. It also contains information you should consider before you start to configure and administer the Postfix Recovery Kit. Please refer to [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring LifeKeeper Core resource hierarchies.

## 6.14.2.1. Postfix Protection Objects

---

The Postfix Recovery Kit protects the following objects:

- Postfix service daemon
- Network socket of Postfix

Create one or more virtual IP addresses for monitoring mail and assign them to the `inet_interfaces` parameter in the Postfix configuration file. If “all” is specified for the `inet_interfaces` parameter, then the local loopback address is used for monitoring. The supported SMTP / SMTPS service ports supported by the Postfix `smtpd` daemon are 25 and 465 respectively.

- The queue directory (filesystem) Postfix uses

If you need the mailbox spool area on another file system and need to protect it, you must create the file system hierarchy for it and create a dependency between the Postfix resource and this resource. Please refer to [Create Dependency with Mailbox Spool Resource](#).

## 6.14.2.2. Postfix Configuration Requirements

---

- If the IP address used by the SMTP service is specified, the IP address should be a virtual IP address that is protected by Lifekeeper.
- main.cf

The Postfix Recovery Kit refers to the value of the following parameters:

- mail\_owner
- setgid\_group
- daemon\_directory
- command\_directory
- process\_id\_directory
- inet\_interfaces

Specify the virtual IP addresses to be monitored. One or more may be specified. Use “all” to specify all IP addresses.

- queue\_directory
- mail\_spool\_directory

- master.cf

You must specify the following:

- A smtp(s) service entry to start smtpd.
- The directory specified for the queue\_directory value must be on shared storage. This is necessary so that the file system of this directory can be LifeKeeper protected.
- If the system has a mailbox spool, the directory specified for the mail\_spool\_directory value has to be on shared storage.
- Owner id of postfix has to be the same id on all cluster servers.
- Group id of postdrop (setgid\_group) has to be the same id on all cluster servers.
- Auto startup at the time of the system activation must be disabled because Postfix service is

controlled by LifeKeeper.

## 6.14.2.3. Port and TCP Interface Definition and the Postfix Recovery Kit

---

The Postfix Recovery Kit listens to the port specified in the SMTP entry in the Postfix configuration file (master.cf). If the port is specified as a service name (e.g., smtp) then the port number is looked up in the /etc/services file (smtp is "25" and smtps is "465").

smtp	inet	n	-	n	-	-	smtpd
------	------	---	---	---	---	---	-------

## 6.14.2.4. DNS, Postfix and LifeKeeper

---

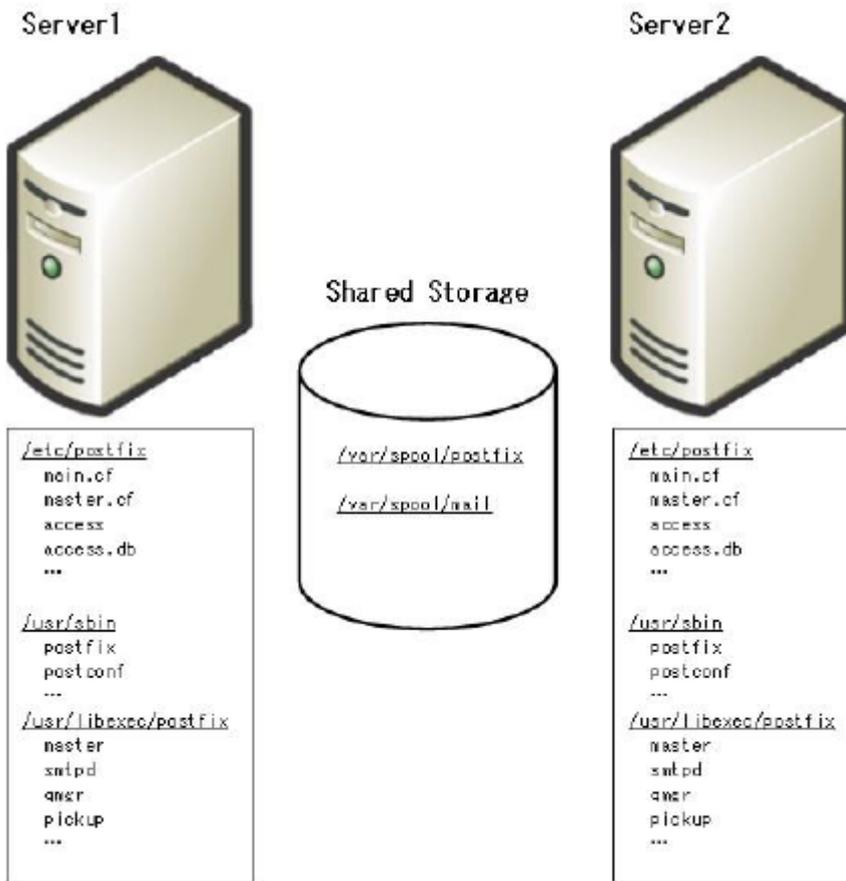
DNS offers a mechanism (MX Records) for specifying backup or alternate hosts for mail delivery. This mechanism also allows hosts to assume mail-handling responsibilities for other hosts that are not configured to accept mail, such as a null client. MX records also provide a mechanism of forcing all mail to go to the hub machine or mail server. MX records specify a mail exchanger for a domain name (i.e. a host that will process and/or forward mail for the specified hostname). As an example, this is done by adding entries into the DNS server as follows:

```
himalaya.sc.steeleye.com IN  MX  10 relay.steeleye.com.
```

In the example, the server `himalaya.sc.steeleye.com` has an MX record that will cause mail for this server to be delivered to `relay.steeleye.com`. The server which is to be LifeKeeper protected should not have any MX records. The LifeKeeper protected alias IP address that is used during the Postfix resource hierarchy creation should be used for MX records instead.

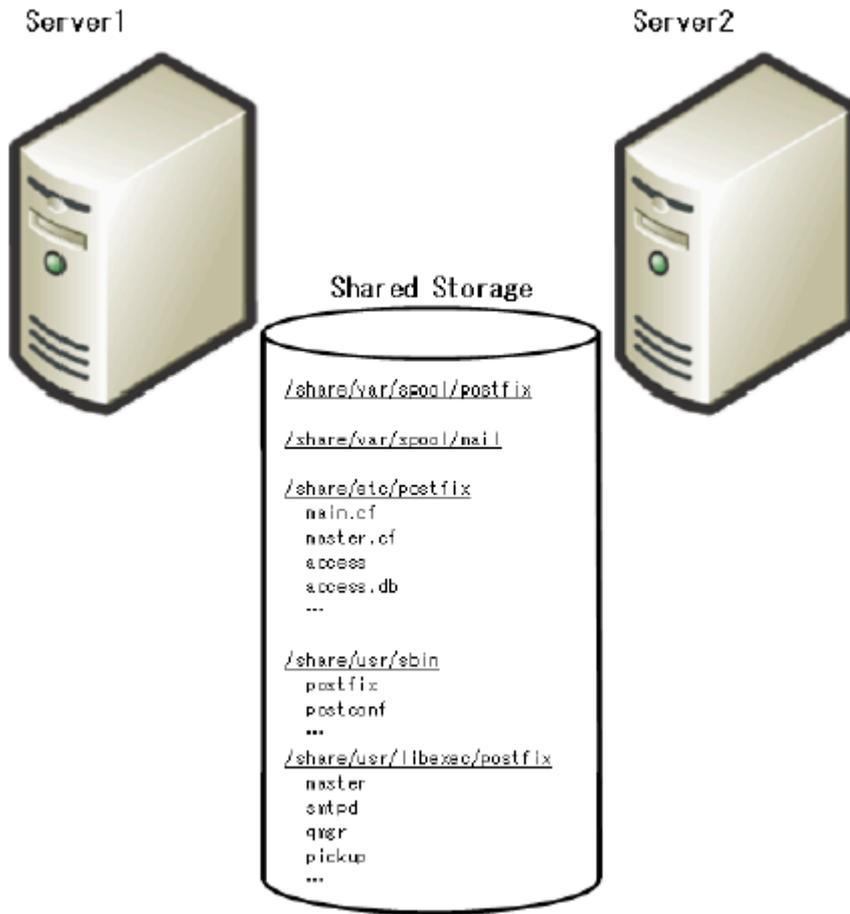
## 6.14.2.5. Postfix Configuration Examples

### Configuration 1: Active/Standby Configuration Example



**Figure 1: Typical LifeKeeper Active/Standby Postfix Environment 1**

- The Postfix configuration files are on both servers
- The Postfix executable files are on both servers.
- The queue area (e.g. /var/spool/postfix) is on shared storage.
- The spool area (e.g. /var/spool/mail) is on shared storage.



**Figure 2: Typical LifeKeeper Active/Standby Postfix Environment 2**

- The Postfix configuration files are on shared file system.
- The Postfix executable files are on shared file system.
- The queue area (e.g. /var/spool/postfix) is on shared storage.
- The spool area (e.g. /var/spool/mail) is on shared storage.

## Configuration 2: Active/Active Configuration Example

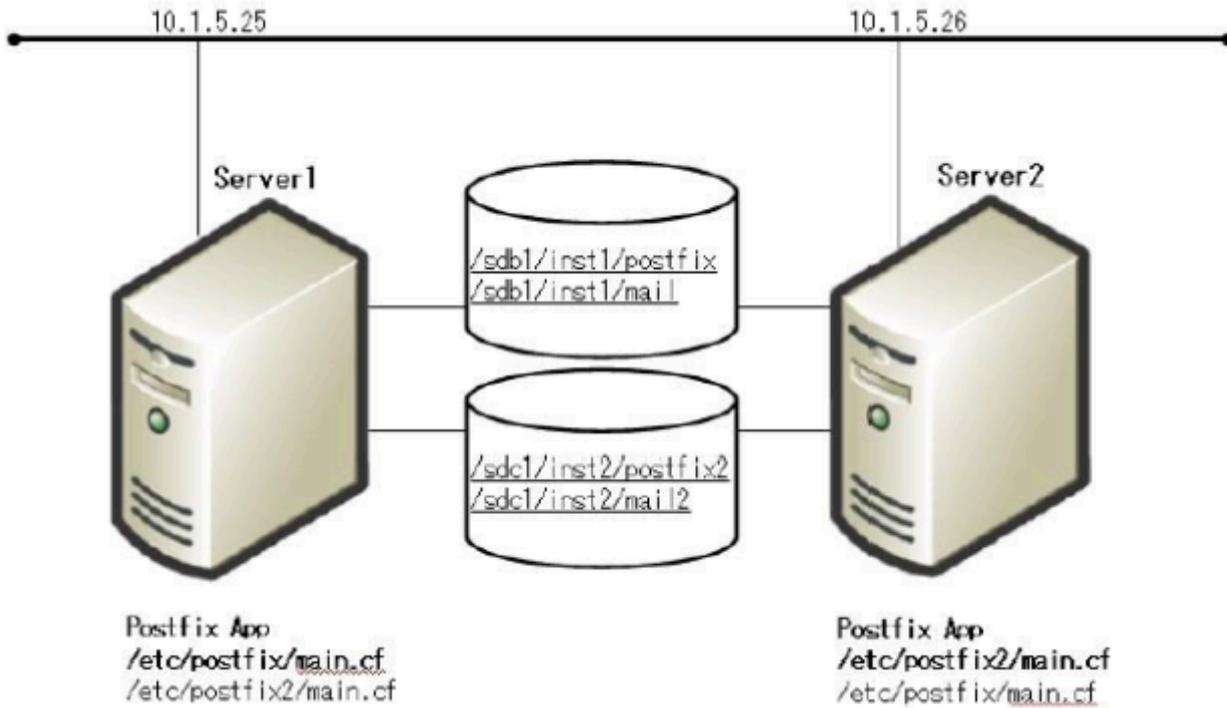


Figure 3: Typical LifeKeeper Active/Active Postfix Environment

[Server1 (Instance 1 is active)]

The Postfix configuration file: /etc/postfix

The Postfix executable files: /usr/sbin

The queue area: /sdb1/inst1/postfix

The spool area: /sdb1/inst1/mail

<main.cf>

inet\_interfaces = 10.1.5.25, localhost

[Server2 (Instance 2 is active)]

The Postfix configuration file: /etc/postfix2

The Postfix executable files: /usr/sbin

The queue area: /sdc1/inst2/postfix2

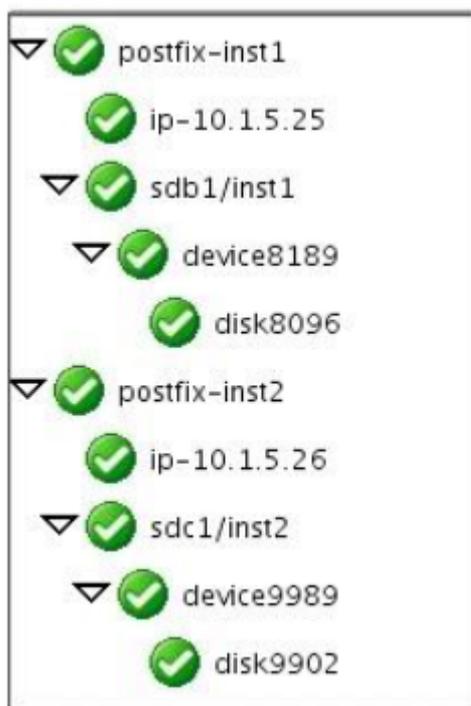
The spool area: /sdc1/inst2/mail2

<main.cf>

```
inet_interfaces = 10.1.5.26
```

```
alternate_config_directories = /etc/postfix2
```

The following figure shows the Postfix resource hierarchies displayed in the LifeKeeper GUI:



## 6.14.3. Postfix Configuration Validation

This section shows a method to check the systems by using the Typical LifeKeeper Postfix Environment 1 as an example before you start to create resources in LifeKeeper.

### Postfix Configuration Validation Steps

#### 1. Postfix Configuration

The Postfix configuration files are on both servers.

```
main.cf (extract)

daemon_directory = /usr/libexec/postfix
command_directory = /usr/sbin
process_id_directory = pid
inet_interfaces = localhost , 192.168.0.10
mail_spool_directory = /var/spool/mail
queue_directory = /var/spool/postfix
```

```
master.cf (extract)

smtp      inet  n       -       n       -       -       smtpd
```

#### 2. Bring up virtual IP address for SMTP

You must bring up virtual IP address for SMTP. You can configure it by using the “ifconfig” command or creating a LifeKeeper IP resource.

```
# ifconfig eth0:1 192.168.0.10 netmask 255.255.255.0 up
```

#### 3. Mount the shared filesystem for queue area

```
# mkfs.ext3 /dev/sda1
# mount -t ext3 /dev/sda1 /mnt/queue
# mkdir -p /mnt/queue/postfix
```

```
# cp -rp /var/spool/postfix/* /mnt/queue/postfix/
# mv /var/spool/postfix /var/spool/postfix.org
# ln -s /mnt/queue/postfix /var/spool/postfix
# postfix check
```

4. Mount the shared filesystem for spool area

```
# mkfs.ext3 /dev/sdb1
# mv /var/spool/mail /var/spool/mail.org
# mkdir -p /var/spool/mail
# mount -t ext3 /dev/sdb1 /var/spool/mail
```

5. Start Postfix

```
# postfix -c /etc/postfix start
postfix/postfix-script: starting the Postfix mail system
```

6. Verify processes and socket for Postfix

```
# netstat -pltn | grep master
tcp    0  0  127.0.0.1:25      0.0.0.0:*        LISTEN  15931/master
tcp    0  0  192.168.0.10:25  0.0.0.0:*        LISTEN  15931/master

# ps -ef | grep -v grep | grep postfix
root    15931      1  0  16:11  ?  00:00:00  /usr/libexec/postfix/master
postfix 15932    15931  0  16:11  ?  00:00:00  pickup -l -t fifo -u
postfix 15933    15931  0  16:11  ?  00:00:00  qmgr -l -t fifo -u
```

7. Stop Postfix

```
# postfix -c /etc/postfix stop
postfix/postfix-script: stopping the Postfix mail system
```

If you cannot start or stop Postfix in steps 5-7, please check the Postfix error messages. Once there are no error messages in the log file, the configuration is correct. Next, repeat steps 1-7 on all systems in the cluster and confirm that the configuration is correct.

## 6.14.3.1. LifeKeeper Configuration Tasks for Postfix

---

You can perform all LifeKeeper for Linux Postfix Recovery Kit administrative tasks via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer, and monitor Postfix resources.

The following tasks are available for configuring the LifeKeeper for Linux Postfix Recovery Kit:

- [Create a Resource Hierarchy](#) – Creates a Postfix resource hierarchy.
- [Delete a Resource Hierarchy](#) – Deletes a Postfix resource hierarchy.
- [Extend a Resource Hierarchy](#) – Extends a Postfix resource hierarchy from the primary server to the backup server.
- [Unextend a Resource Hierarchy](#) – Unextends (removes) a Postfix resource hierarchy from a single server in the LifeKeeper cluster.
- [Create Dependency](#) – Creates a child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete Dependency](#) – Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#) – Activates a resource hierarchy.
- [Out of Service](#) – Deactivates a resource hierarchy.
- [View /Edit](#) Properties – View or edit the properties of a resource hierarchy.

**Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You can also perform most of these tasks:

1. From the toolbar, right – click on a global resource in the left pane of the status display.
2. Right – click on a resource instance in the right pane of the status display.

**Note:** Using the right-click method allows you to avoid entering information that is required when using the **Edit** menu.

## 6.14.3.1.1. Creating a Postfix Resource Hierarchy

After you have completed the necessary setup tasks, use the following steps to define the Postfix resource hierarchy.

**IMPORTANT:** The alias IP address should be under LifeKeeper protection before creating the Postfix resource instance.

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the menu, select **Create Resource Hierarchy**.

The **Create Resource Wizard** dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

2. Select **Postfix Mail Server** and click **Next**.
3. You will be prompted to enter the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter an error requiring you to correct previously entered information. You may click Cancel at any time to cancel the entire creation process.

Field	Tips
Switchback Type	Choose either <i>intelligent</i> or <i>automatic</i> . This dictates how the Postfix instance will be switched back to this server when the server comes back up after a failover. The switchback type can be changed later from the General tab of the Resource Properties dialog box.  <b>Note:</b> The switchback strategy should match that of the IP or File System resource to be used by the Postfix resource. If they do not match the Postfix resource, creation will attempt to reset them to match the setting selected for the Postfix resource.
Server	Select the <b>Server</b> on which you want to create the hierarchy.
Postfix Binary Location	Enter the directory path name where the Postfix daemon resides.
Postfix server Config File Location	Enter the directory path name where the Postfix configuration file (main.cf) resides.
Queue Root	Enter the directory path name of the Postfix queue directory. The default is decided from the configuration file, which you selected in the previous dialog box. The Postfix queue directory

Directory	must be on a shared disk. If the Postfix queue directory is a symbolic link, the dialog box will show the root directory of the symbolic link pointing to the directory's original location.
Root Tag	Either select the default root tag offered by LifeKeeper, or enter a unique name for the resource instance on this server. The default is postfix-on- <queue directory path>. You may use letters, numbers and the following special characters: – _ . /

4. Click **Create**. The Create Resource Wizard will then create your Postfix resource hierarchy. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.
5. An information box will appear indicating that you have successfully created a Postfix resource hierarchy, and you must extend that hierarchy to another server in your cluster in order to achieve failover protection. Click **Next**.
6. Click **Continue**. LifeKeeper will then launch the *Pre-Extend Wizard*. Refer to Step 2 under [Extending Your Hierarchy](#) (below) for details on how to extend your resource hierarchy to another server.

## 6.14.3.1.2. Extending a Postfix Resource Hierarchy

This operation can be started from the **Edit** menu, or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should refer to Step 2 below.

1. On the **Edit** menu, select Resource, then Extend Resource Hierarchy. The *Pre-Extend Wizard* appears. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper **Extend Resource Hierarchy** defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The *Pre-Extend Wizard* will prompt you to enter the following information.

**Note:** The first two fields appear only if you initiated the Extend from the **Edit** menu.

Field	Tips
Template Server	Enter the server where your Postfix resource is currently in service.
Tag to Extend	Select the Postfix resource you wish to extend.
Target Server	Enter or select the server you are extending to.
Switchback Type	Select either <i>intelligent</i> or <i>automatic</i> . The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.  <b>Note:</b> Remember that the switchback strategy must match that of the dependent resources to be used by the Postfix resource.
Template Priority	Select or enter a priority for the template hierarchy. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (the number 1 indicates the highest priority). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.  <b>Note:</b> This selection will appear only for the initial extend of the hierarchy.
Target Priority	Either select or enter the priority of the hierarchy for the target server.

3. After receiving the message that the pre-extend checks were successful, click **Next**.
4. LifeKeeper will display a series of dialog boxes for the Postfix resource to be extended, some of which cannot be edited.

Field	Tips
Root Tag	LifeKeeper will provide a default tag name for the new Postfix resource instance on the target server. The default tag name is the same as the tag name for this resource on the template server. If you enter a new name, be sure it is unique on the target server. You may use letters, numbers and the following special characters: – _ . /
Binary Directory (Information Only)	This dialog box is for informational purposes only. You cannot change the Binary Directory that appears in the box.
Configuration Directory (Information Only)	This dialog box is for informational purposes only. You cannot change the Configuration Directory that appears in the box.

If the IP and Filesystem dependent resource are also being extended, LifeKeeper will display a series of dialog box for the resources, some of which cannot be edited.

Click **Extend**

5. After receiving the message “Hierarchy extend operations completed” click **Next Server** to extend the hierarchy to another server, or click **Finish** if there are no other extend operations to perform.
6. After receiving the message “Hierarchy Verification Finished”, click **Done**.

## 6.14.3.1.3. Unextending a Postfix Resource Hierarchy

---

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. On the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the Postfix resource. It cannot be the server where the Postfix resource is currently in service. (This dialog box will not appear if you selected the Unextend task by right clicking on a resource instance in the right pane.) **Click Next**.
3. Select the Postfix hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the Unextend task by right clicking on a resource instance in either pane).
4. An information box appears confirming the target server and the Postfix resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the Postfix resource was unextended successfully. Click **Done** to exit the Unextend Resource Hierarchy menu selection.

## 6.14.3.1.4. Deleting a Postfix Resource Hierarchy

---

It is important to understand what happens to dependencies and protected services when a Postfix hierarchy is deleted.

- **Dependencies:** When you choose to delete the Postfix hierarchy, only the Postfix resource will be deleted. Dependent IP and file system resources will not be removed.
- **Protected Services:** If you take the Postfix resource hierarchy out of service before deleting it, the Postfix daemons will be stopped. If you delete a hierarchy while it is in service, the Postfix daemons will continue running and offering services (without LifeKeeper protection) after the hierarchy is deleted.

To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the **Target Server** where you will be deleting your Postfix resource hierarchy and click **Next**. (This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in either pane.)
3. Select the **Hierarchy to Delete**. (This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in the left or right pane.) Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete**.
5. Another information box appears confirming that the Postfix resource was deleted successfully.
6. Click **Done** to exit.

## 6.14.3.1.5. Create Dependency with Mailbox Spool Resource

---

If the Postfix queue directory and Mailbox Spool directory are on the same file system (LUN) on the shared disk, both directories are protected by creating the Postfix resource hierarchy and extending the Postfix resource hierarchy to another server in your cluster. If your spool directory is on another file system (LUN), you must create a file system resource for Mailbox Spool and create a dependency for the resource.

To create a resource instances and create dependencies for your Mailbox Spool directory, you should complete the following step:

1. Mount file system for your Mailbox Spool Directory.
2. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop down menu, select **Create Resource Hierarchy**.
3. Select File System from the drop down listing.
4. Select Switchback Type.
5. Select the Primary Server.
6. Select the Mount Point for the file system resource hierarchy.
7. Select or enter Root Tag.

Through this process, the file system resource is created on the primary server, and you must extend it to backup servers. Next, create dependencies for each file system resources to the Postfix resource. You should refer [Creating Resource Dependency](#) section of LifeKeeper for Linux Technical Documentation for specific instructions on how to create dependencies.

## 6.14.3.1.6. Testing Your Postfix Resource Hierarchy

---

You can test your Postfix resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit**, then **Resource**, then finally **In Service** from the drop down menu. For example, an in service request executed on a backup server causes the application hierarchy to be placed in service on the backup server and taken out of service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the **Out of Service** request, the application is taken out of service without bringing it in service on the other server.

### Performing a Manual Switchover from the Command-Line Interface

You can initiate a manual switchover from the LifeKeeper command-line interface by the following steps on the server:

```
# /opt/LifeKeeper/bin/perform_action -t [tag-name] -a [restore|remove]
```

- **-t**
- This specifies the last resource instance that the action will be performed on. “tag-name” are the information elements that may be used to describe the resources in the hierarchy, the name can be checked from LifeKeeper GUI, or “lcdstatus” command.
- **-a**

This specifies the resource action that will be performed. To bring the resource instance into service, specify restore, to take a resource out of service, specify remove.

Please refer to man pages of *perform action* for more details.

### Recovery Operations

When the following failure occurs on the in service server, the Postfix Recovery Kit software performs Recovery:

- Failure in the Postfix resource
- Failure in IP resource relative to the Postfix resource
- Failure in file system resource relative to the Postfix resource

- Node Failure

When the primary server fails, the Postfix Recovery Kit software performs the following tasks:

- Brings the alias IP address into service on the backup server by bringing *in service* a logical interface on one of that server's physical network interfaces
- Mounts the file system(s) on the shared disk on that server
- Starts the daemon processes related to Postfix

Since session context is lost following recovery, after the recovery, Postfix users must reconnect using exactly the same procedures they used to connect originally.

## 6.14.4. Postfix Troubleshooting

---

This section provides a list of messages that you may encounter during the process of creating, extending, removing and restoring a LifeKeeper Postfix hierarchy, and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. Other messages from other LifeKeeper scripts and utilities are also possible. In these cases, please refer to the documentation for the specific script or utility. Messages in this section fall under these categories:

[Hierarchy Creation](#)

[Hierarchy Extend](#)

[Hierarchy Remove, Restore and Recovery Error Messages](#)

## 6.14.4.1. Postfix Hierarchy Creation Error Messages

Error	Error Message
No config path	The Postfix configuration path was not found. Please enter the configuration path.
main.cf not found in the configuration path	The file main.cf does not exist in the path specified. Please enter the correct path.
master.cf not found in the configuration path	The file master.cf does not exist in the configuration path. Please enter the correct path.
A value of inet_interfaces must be IPv4 or "all"	Please specify an IPv4 address or "all" for the inet_interfaces parameter in the main.cf file.
No execute path	Must specify the absolute path to the Postfix executables. Please enter the correct path.
Postfix command invalid	The Postfix command is invalid. Please verify the Postfix installation or command and enter the correct command.
<queue directory> is not found. This directory must exist on a shared filesystem	The mail queue directory(s) must be located on a shared filesystem. Please make sure your configuration is correct.
<tag name> not in service on the server	The tag name is not in service. Please create the IP resource and verify that the virtual IP address is active on the server.
Could not find IP resource for "<IP address>"	The LifeKeeper IP resource for the IP address specified for the inet_interfaces parameter in main.cf is missing. Please create the LifeKeeper IP resource.

## 6.14.4.2. Postfix Hierarchy Extend Error Messages

---

Error	Error Message
postfix id does not match between servers	The Postfix uid does not match on the servers in the cluster. Please set the same uid for the user "postfix" on the cluster servers.
postdrop gid does not match between servers	The Postfix postdrop gid does not match on the servers in the cluster. Please set the same gid for the group "postdrop" on the cluster servers.

## 6.14.4.3. Postfix Resource In-Service / Out-of-Service / Health Monitoring Error Messages

Error	Error Message
Master process of postfix is not running	The master process of Postfix is not running. Please check the Postfix error log.
Failed in a check by postfix command	Postfix command check option has failed. Please check the Postfix configuration file or Postfix environment.
Couldn't start postfix resource	The Postfix resource could not start. Please check the Postfix error log.
Failed in a stop process by kill command	The kill command failed to stop Postfix. Please check the Postfix error log.
PID <pid> does not exist. postfix may have already stopped	The Postfix pid does not exist. Please check the Postfix error log and Postfix processes. The Postfix process may have been stopped and then restarted and assigned another pid.
Check script was not able to be connected to a socket ( <i>vip:port</i> )	The check script was not able to connect to the socket for service. Please check the Postfix configuration file and the Postfix owner.
Execute files (postfix or postconf command) is not an executable file	The files postfix or postconf does not exist or are not executable. The files are located in the executable path that was specified when the resource was created. Please check these files.
Configuration files (main.cf or master.cf) does not exist	The Postfix configuration files main.cf or master.cf does not exist or is not readable. The files and located in the configuration path that was specified when the resource was created. Please check these files.
The postfix owner <owner name> does not exist	The Postfix owner does not exist. Please check the Postfix configuration and Postfix owner.
The postdrop group id does not match and attribute of queue directory	The postdrop group id does not match the group id associates with the files in the mail queue directory. Please check the Postfix configuration file.

# 6.15. Recovery Kit for Route 53™

## Administration Guide

---

The Recovery Kit for Route 53™ provides a mechanism for updating Amazon Route 53 DNS information corresponding to a virtual IP address and an actual IP address information of IP resources that are in dependency relation when switching to a failed primary server to a backup server

### LifeKeeper Documentation

The following is a list of LifeKeeper for Linux related information available from SIOS Technology Corp.

- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Release Note](#)
- [SIOS Technology Corp. Documentation](#)

For the details, please refer to [Amazon Route 53 Documentation](#).

 **Note:** Amazon Web Services, the “Powered by Amazon Web Services” logo, “AWS”, “Amazon EC2”, “EC2”, “Amazon Elastic Compute Cloud”, “Amazon Route 53”, “Amazon Virtual Private Cloud”, and “Amazon VPC” are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries.

## 6.15.1. Recovery Kit for Route 53™ Requirements

---

Prior to installing and configuring the Recovery Kit for Route 53™, be sure your configuration meets the following requirements.

- AWS Command Line Interface (AWS CLI) must be installed on all EC 2 instances. Refer to [Installing AWS Command Line Interface](#) for installation information.
- Instances need to have an access to Amazon Route 53 service endpoint, route53.amazonaws.com, with HTTPS protocol. Please configure EC2 and the OS properly.
- Register an appropriate domain name for Amazon Route 53.
- In order for LifeKeeper to operate AWS, an IAM user or an IAM role with the following access privilege is required. Please configure [IAM roles for Amazon EC2](#) or the [AWS CLI](#) appropriately so that it can be accessed from root user of the Amazon EC2 instance.
  - route53:GetChange
  - route53:ListHostedZones
  - route53:ChangeResourceRecordSets
  - route53:ListResourceRecordSets

 **Note:** If the path name of AWS CLI executable files is not specified on the “PATH” parameter in the LifeKeeper defaults file `/etc/default/LifeKeeper`, you must append the path name of AWS CLI executable files to the “PATH” parameter.

### LifeKeeper Software:

You need to install the same version of LifeKeeper software and patches on each server. For the specific LifeKeeper requirements, please refer to the [Technical Documentation](#) or the [LifeKeeper for Linux Release Notes](#).

SIOS recommends using Quorum/Witness when using the Recovery Kit for Route 53™. Please refer to [Quorum/Witness](#) for more information.

## 6.15.2. Recovery Kit for Route 53™ Configuration

---

To configure LifeKeeper to provide the required protection capability and flexibility, you need to know the configuration requirements. You also need to understand Amazon, Amazon Virtual Private Cloud (VPC), Amazon Elastic Compute Cloud (EC2), Amazon 53 and hierarchy configuration options of the user system. In addition to the configuration planning, this section also describes the specific tasks required to set up the Recovery Kit.

### Specific Configuration Considerations for Route53 Resources

The following configuration tasks for Route53 resources are described in this section. They are unique to a Route53 resource instance and different for each recovery kit.

- [Creating a Resource Hierarchy](#): Creates an application resource hierarchy in your LifeKeeper cluster.
- [Deleting a Resource Hierarchy](#): Deletes a resource hierarchy from all servers in your LifeKeeper cluster.
- [Extending Your Hierarchy](#): Extends a resource hierarchy from the primary server to a backup server.
- [Unextending Your Hierarchy](#): Unextends (removes) a resource hierarchy from a single server in your LifeKeeper cluster.
- [Adjusting Recovery Kit for Route 53™ Tunable Values](#): Tunes characteristics of the overall behavior of the Recovery Kit for Route 53™.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#). They are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#): Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#): Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#): Brings a resource hierarchy into service on a specific server.
- [Out of Service](#): Takes a resource hierarchy out of service on a specific server.
- [View Properties](#) / [Edit Properties](#): View or edit the properties of a resource hierarchy on a specific

**server.**

## 6.15.2.1. Creating a Route53™ Resource Hierarchy

To create a resource instance from the primary server, complete the following steps.

1. From the LifeKeeper GUI menu, select **Edit** then **Server**. From the drop down menu, select **Create Resource Hierarchy**.
2. A dialog box will appear with a drop down list showing all of the recognized recovery kits installed within the cluster. Select **Amazon Route53** from the drop down list and click **[Next]**
3. You will be prompted to enter the following information. (When the Back button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful in the event that you need to correct previously entered information.)

**Note:** you click the Cancel button at any time when creating your hierarchy, LifeKeeper will cancel the entire creation process.

Field	Tips
Switchback Type	<p>This dictates how the Route53 instance will be switched back to this server when the server recovers after a failover. You can choose either intelligent or automatic.</p> <p>Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server.</p> <p>Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.</p> <p><b>Note:</b>The switchback type can be changed later from the General tab of the Resource Properties dialog box.</p>
Server	Select the Server for the Route53 resource (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop down list.
Domain name (Route53 hosted zone)	Route53 hosted zones are listed in the drop down list. Select the domain name to use.
Host Name (Not FQDN)	Enter the host name.
IP resource	Select the IP resource. This is the virtual IP address or the actual IP address that is protected by LifeKeeper.
Route53 Resource Tag	Select or enter a unique Route53 Resource Tag name for the Route53 resource instance you are creating. This field is populated automatically with a default tag name,

	route53-<host name>.
--	----------------------

4. Click **Create**. The Create Resource Wizard will then create your Route53 resource
5. At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your Route53 resource hierarchy. If LifeKeeper detects a problem an ERROR will appear in the information box. If the validation is successful your resource will be created. Click **Next**

Another information box will appear confirming that you have successfully created a Route53 resource hierarchy. You must extend that hierarchy to another server in your cluster in order to place it under LifeKeeper protection

When you click **Continue**, LifeKeeper will launch the Pre-Extend configuration task. Refer to Extending Your Hierarchy for details on how to extend your resource hierarchy to another server.

If you click **[Cancel]** now, another dialog box will appear alerting you that you will need to manually extend your Route53 resource hierarchy to another server at some other time to put it under LifeKeeper protection.

## 6.15.2.2. Deleting a Route53™ Resource Hierarchy

---

To delete a resource hierarchy from all of the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the dropdown menu, select **Delete Resource Hierarchy**.
2. Select the name of the Target Server that you are deleting from your Route53 resource hierarchy and click **Next**

**Note:**This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in either pane.

3. Select the Hierarchy to Delete. Identify the resource hierarchy you wish to delete, highlight it then click **Next**

**Note:**This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in the left or right pane.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to **Delete** to proceed.
5. An information box appears confirming that the Route53 resource was deleted successfully.
6. Click **Done** to exit.

## 6.15.2.3. Extending Your Route53™ Resource Hierarchy

After you have created a hierarchy, you must extend that hierarchy to another server in the cluster. There are three possible scenarios to extend your resource instance from the template server to a target server.

- Continue from creating the resource into extending that resource to another server.
- Enter the Extend Resource Hierarchy task from the edit menu as shown below.
- Right click on an unextended hierarchy in either the left or right hand pane.

Each scenario takes you through the same dialog boxes (with a few exceptions, detailed below).

1. If you are entering the Extend wizard from the LifeKeeper GUI menu, select **Edit** , then **Resource** . From the drop down menu, select **Extend Resource Hierarchy**. This will launch the Extend Resource Hierarchy wizard. If you are unfamiliar with the Extend operation, click **Next** . If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The Pre-Extend Wizard will prompt you to enter the following information

**Note:** The first two fields appear only if you initiated the Extend from the Edit menu. It should be noted that if you click Cancel at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extension process to that particular server. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Field	Tips
Switchback Type	<p>Select the Switchback Type. This dictates how the Route53 instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either intelligent or automatic.</p> <p>* Intelligent switchback requires administrative intervention to switchback the instance to the primary/original server.</p> <p>* Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths.</p> <p>The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.</p>

Template Priority	<p>Select or enter a Template Priority. This is the priority for the Route53 hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.</p> <p><b>Note:</b> This selection will appear only for the initial extend of the hierarchy.</p>
Target Priority	<p>Select or enter the Target Priority. This is the priority for the new extended Route53 hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest).</p> <p><b>Note:</b> LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities do not need to be consecutive and no two servers can have the same priority for a given resource.</p>

- An information box will appear explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this Route53 resource have been met. If there were some requirements that have not been met, LifeKeeper will not allow you to select the **Next** button, and the **Back** button will be enabled. If you click **Back**, you can make changes to your resource extension according to any error messages that may appear in the information box. If you click **Cancel** now, you will need to manually extend your Route53 resource hierarchy to another server to put it under LifeKeeper protection. When you click **Next**, LifeKeeper will launch you into the Extend Resource Hierarchy configuration task.
- The Extend Resource Hierarchy configuration task will prompt you to enter the following information

Field	Tips
Route53 Resource Tag	<p>Select or enter the Route53 Resource Tag. This is the resource tag name to be used by the Route53 resource being extended to the target server.</p> <p><b>Note:</b>The field is not editable.</p>

- An information box will appear verifying that the extension is being performed. Click **Next Server** if you want to extend the same Route53 resource instance to another server in your cluster. This will repeat the Extend Resource Hierarchy operation. If you click **Finish**, LifeKeeper will verify that the extension of the Route53 resource was completed successfully.
- Click **Done** to exit from the Extend Resources Hierarchy menu selection.

**Note:** Be sure to test the functionality of the new instance on all servers.

## 6.15.2.4. Unextending Your Route53™ Resource Hierarchy

---

To unextend a hierarchy complete the following steps:

1. From the **LifeKeeper GUI menu**, select **Edit**, then **Resource**. From the dropdown menu, select **Unextend Resource Hierarchy**.
2. Select the Target Server that you are unextending from the Route53™ resource. It cannot be the server that the Route53™ resource is currently in service on. Click **Next**.

**Note:** If you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance, the dialog box will not appear.

3. Select the Route53™ Hierarchy to unextend. Click **Next**

**Note:** If you selected the **Unextend** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, the dialog will not appear.

4. An information box will appear confirming the target server and the Route53™ resource hierarchy you have chosen to unextend. Click **Unextend**.
5. An information box will appear confirming the Route53™ resource hierarchy you have chosen to unextend.
6. Click **Done** to exit.

## 6.15.2.5. Adjusting Recovery Kit for Route 53™ Tunable Values

---

For the parameters that can be configured in the Recovery Kit for Route 53™, refer to [Parameters List](#).

## 6.15.2.6. Route53™ Resource Monitoring and Recovery

---

The Route53 resource monitors the normality of the retrieval of the DNS A record registered at the time of creation and the association with the virtual IP address. The monitoring process is as follows.

1. Obtain the address set in the Route 53 A record with API. If it fails to obtain the record, it will retry 3 additional times waiting 2 seconds between attempts (by default). After the third unsuccessful attempt it will stop the monitoring and record the failure in the log.
2. Obtain an IP address from the dependent IP resource and compare it with the IP address in the DNS A record information. If the IP address information matches, then exit with a success as no errors exist. If the IP addresses do not match then exit with a failure to initiate a local recovery.

## 6.15.2.7. Route53™ User System Setup

---

An IP resource is required for creating the Route53 resource and can be either the virtual IP resource or the actual IP resource (resource for the primary IP address that is configured for the Network Interface).

### When Using the Virtual IP Resource

When using the virtual IP resource for a child resource of the Route53 resource, you need to reconfigure the route table so that the communication with the virtual IP address to the backup server is enabled when switching over the resource. Please use the Recovery Kit for EC2 along with the Recovery Kit for Route 53™. For details, please refer to the [Recovery Kit for EC2™ document](#)

### When Using the Actual IP Resource

No additional information needs to be configured when using the actual IP resource for a child resource for the Route53 resource. However, because the destination IP address will be changed every time the switch over occurs, please note that the connection should be established with the host name that is protected by the Route53 resource.

## 6.15.3. Recovery Kit for Route 53™ Troubleshooting

---

The [Message Catalog](#) provides a listing of all messages that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received.

### Updating the record associated with the Route53 resource startup may take time

Amazon provides the following information regarding the propagating speed of changes made to DNS record.

Amazon Route 53 FAQs

Q: How quickly will changes I make to my DNS settings on Amazon Route 53 propagate globally?

[https://aws.amazon.com/route53/faqs/?nc1=h\\_ls](https://aws.amazon.com/route53/faqs/?nc1=h_ls)

The Route53 resource checks the status of updates to the DNS record using the Route53 API. It considers that updates are completed when receiving INSYNC status, and retries the status checking when it receives PENDING status. As a result, the Route53 resource may result in a startup failure when it takes a long time to propagate updates to the DNS even though the record is updated successfully for the Route53 resource startup process.

If the startup of the Route53 resource fails, check the Route53 management console to make sure that the A record is updated correctly. If it is updated, updates to the relevant DNS service have been completed. Updates to LifeKeeper is required to propagate the updates to the DNS service. Restart the Route53 resource from LifeKeeper GUI

If you encounter the startup failure of the Route53 resource all the time due to the above mentioned reason, increase the number of the value of "ROUTE53\_CHANGEID\_TRY\_COUNT" in /etc/default/LifeKeeper to 6 or 7 (the default value is 5). Restart of LifeKeeper or the OS is not required for this change.

### Correctly set TTL value of the DNS record

An access from a client after a switchover or a failover uses the DNS information cache that each client holds until the time set as TTL is passed. If the longer TTL value is set, access attempts to the address before switching increase and unexpected problems may occur. If the shorter TTL value is set, DNS resolution often occurs and network load increases. Please set the TTL value as short as possible according to your environment.

Set the "ROUTE53\_TTL" for the TTL value in /etc/default/LifeKeeper. The unit should be in seconds.

# 6.16. Samba Recovery Kit Administration Guide

---

The LifeKeeper for Linux Samba Recovery Kit provides fault resilient protection for Samba file and print shares on a Linux server existing in a heterogeneous network. This enables a failure on the primary Samba server to be recovered on a designated backup server without significant lost time or human intervention.

## Document Contents

This guide contain the following topics:

- [Requirements](#). A description of the hardware and software necessary to properly setup, install, and operate the Samba Recovery Kit. Refer to LifeKeeper for Linux Installation Guide for specific instructions on how to install or remove LifeKeeper for Linux software.Samba Recovery Kit .
- [Samba Recovery Kit Overview](#). Provides a brief description of the Samba Recovery Kit's features and functionality.
- [Configuring the LifeKeeper for Linux Samba Recovery Kit](#). A description of the procedures required to properly configure the Samba Recovery Kit.
- [LifeKeeper Configuration Tasks](#). A description of the tasks for creating and managing your Samba resource hierarchies using the LifeKeeper GUI.
- [Samba Hierarchy Administration](#). Provides information about tasks that may be required after your Samba resources are created.
- [Troubleshooting](#). A list of LifeKeeper for Linux error messages including a description for each.

## Documentation and References

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Installation Guide](#)
- [Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is available on the SIOS Technology Corp. website at:

<http://docs.us.sios.com>

## 6.16.1. Samba Recovery Kit Requirements

Your LifeKeeper configuration must meet the following requirements **prior** to the installation of the Samba Recovery Kit. Please see the [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the configuration of your LifeKeeper hardware and software.

- **Servers.** The Recovery Kit requires two or more servers configured in accordance with the requirements described in the [LifeKeeper for Linux Installation Guide](#). See the [LifeKeeper for Linux Release Notes](#) for supported Linux distributions.
- **LifeKeeper software.** You must install the same version of LifeKeeper software and any patches on each server.
- **LifeKeeper IP Recovery Kit.** You must have the same version of the LifeKeeper IP Recovery Kit on each server.
- **IP network interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

 **Note:** Even though each server requires only a single network interface, you should use multiple interfaces for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation, and so forth.

- **TCP/IP Software.** Each server also requires the TCP/IP software.
- **Samba Software.** Samba is delivered with all Linux distributions that LifeKeeper for Linux supports. Please use the delivered version of Samba.

## 6.16.2. Samba Recovery Kit Installation

---

Please refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on the installation and removal of the LifeKeeper for Linux software.

## 6.16.3. Samba Recovery Kit Overview

Samba is a suite of applications that speak the Server Message Block (SMB) protocol, allowing a Linux server to communicate in a heterogeneous network with servers and clients running Microsoft Windows products.

The Samba Recovery Kit enables LifeKeeper to protect Samba file and print shares on a Linux server. While Samba provides other services such as client authentication, Network Neighborhood browsing assistance and WINS name server resolution, this release of LifeKeeper does not protect these additional services. These other Samba services may coexist on a LifeKeeper server running as an unprotected instance of Samba as long as they adhere to the rules specified in the section [Running Multiple Instances of Samba](#).

The Samba Recovery Kit provides a mechanism to recover protected Samba file and print shares from a failed primary server onto a backup server. LifeKeeper can detect failures either at the server level (via heartbeat) or resource level (by monitoring the Samba daemons) so that control of the Samba resources is transferred to a backup server.

### Samba Resource Hierarchies

A typical Samba hierarchy will be comprised of a Samba resource, one or more file system resources, one or more IP resources, and possibly a print services resource. An example of a resource hierarchy protecting a Samba file share is shown below:



This *Samba-smb.conf* hierarchy protects one fileshare *filesys7531* (which is dependent upon the partition *device-nfs7457*), and one IP address 172.17.101.131. The following sections describes how the Samba resources are configured.

## 6.16.4. Configuring Samba with LifeKeeper

---

There are a number of Samba configuration considerations that need to be made before attempting to create LifeKeeper for Linux Samba resource hierarchies. Samba services on a Linux server are provided by two daemon processes, **smbd** and **nmbd**. These daemon processes are controlled by the values defined in the Samba configuration file which is described below.

## 6.16.4.1. The Samba Configuration File

---

While a Samba configuration file can contain many different directives, this description focuses on those aspects of the configuration file that affect your LifeKeeper configuration. Here are some key points about the configuration file:

- The configuration file is comprised of sections which correspond to the share (or service) they provide. Each section of the configuration file contains individual configuration options (or directives) unique to that share.
- The directives that are specified are sanity checked by the Samba Recovery Kit. Failure to set the directives properly will cause Samba resource creation to fail.
- The default configuration file for Samba is typically named *smb.conf* and resides in */etc* or */etc/samba* depending on the Linux distribution.
- Configuration file names must be unique within the cluster, or must reside in a different directory on each server for Active/Active configurations. The unique naming or location is required as the Samba Recovery Kit replicates a copy of the configuration file during extension to the same location on the backup server.
- Default set up can be used (and recommended) in the case to execute only one Samba daemon instance with active/standby set up. In this case, Samba daemon automatic startup must be disabled.
- If more than one version of Samba will be running in an Active/Standby configuration or if you use an Active/Active configuration, unique Samba configuration file names are required. See [Running Multiple Instances of Samba](#) for more requirements and information on running multiple versions of Samba.

The following sections of this document describe the sections of the configuration file, including the options required for LifeKeeper to protect a Samba share.

## 6.16.4.2. [Global] Section of the Configuration File

---

The [global] section is a special section in the configuration file that must appear in every configuration file used in a LifeKeeper Samba resource hierarchy. As the name implies, any options set in this section apply to all other sections unless that directive is called out specifically in the other sections. LifeKeeper requires that certain directives be defined in the [global] section. Some of these directives may not exist in a default configuration file and will therefore need to be added. They are:

- **netbios name** – The unique name given to the set of resources that comprise a LifeKeeper Samba hierarchy. This is the name used by clients to connect to the shares via the IP addresses defined in the interfaces directive (e.g. NetBIOS name = server1\_print1).
- **interfaces** – The list of network addresses for the Linux Samba server to recognize and respond. Here are the requirements for properly configuring the interfaces directive:
  - All subnets that are serviced by the Samba server must be listed. These must be LifeKeeper protected addresses and they must be unique within the cluster (no other Samba configuration file should use the same IP addresses).
  - The interfaces directive can have multiple formats, IP addresses (dot version or host name), and network interface names and can make use of wild cards. However, the Samba Recovery Kit requires the use of the IP address in dot format (100.25.104.25) without wild cards.
  - The subnet mask may be used in conjunction with the IP address but it is not used by LifeKeeper.
  - LifeKeeper IP resources for the address specified in this directive must be created prior to the creation of the Samba resource hierarchy. Additionally, if the network mask is applied to the addresses in this directive it must match the mask used on the IP resource when it was created.
  - Other non-protected instances of Samba should also use the interfaces directive, being sure to specify IP addresses different than those used by LifeKeeper.

**Note:** Because of the use of the bind interfaces only directive discussed below, the interfaces directive may need to contain the localhost address of 127.0.0.1 to ensure proper operation of the utility **smbpasswd**. See [Running Multiple Instances of Samba](#) for information to help you determine whether the localhost address is needed.

- **lock directory** (or lock dir) – The name and location of a unique lock file location for the Samba instance on all servers. This directory must already exist on all servers in the cluster.

**Note:** This directive is sometimes call **lock dir**. The Samba Recovery Kit will handle both directive names.

- **bind interfaces only** – This directive tells **smbd** and **nmbd** processes to serve SMB requests on the addresses defined in the interfaces directive only. It must be set to **Yes**. Other non-protected instances of Samba running on the system must also have this directive set to yes. When set to yes, Samba will not service requests on subnets that are not listed in the interfaces directive nor will it service requests for other instances of Samba that may be running on the server.

## 6.16.4.3. [Homes] Section of the Configuration File

---

The [homes] section is a special section in the configuration file to handle connection attempts to a user's home directory on a Samba server if it is not specifically defined as a share. LifeKeeper does not protect users' home directories via this special share; therefore **it should be removed or commented out**. In order for the LifeKeeper Samba Recovery kit to protect a Samba share it must have a path directive specified. The path directive is used to determine the file system that the LifeKeeper Samba hierarchy protects. The [homes] section does not have a path specified by default because the path is determined at the time a user makes a connection to the Samba server. It is for this reason that this special share must be removed or comment out.

## 6.16.4.4. [Printers] Section of the Configuration File

---

The [printers] section handles connection attempts to printers on a Samba server if it is not specifically defined as a share. LifeKeeper does not protect printer shares via this special section nor through the global directive load printers. Each LifeKeeper-protected printer share must be defined in its own share section in the configuration file.

## 6.16.4.5. Share Definition Sections of the Configuration File

---

All other sections in the configuration file define the file and/or print shares that clients can attempt to access for this instance. A configuration file must have one or more shares defined. The Samba configuration file can contain file shares only, print shares only or a combination of both file and print shares. LifeKeeper does not limit the number of shares that can be defined, but one must realize that a failure relating to any one share could cause the entire hierarchy to be switched over to the backup server. The following directives must be defined for each share:

- **path** – This directive identifies the pathname at the root of the file or print share. The value determines the File System resource to be protected as part of the Samba hierarchy. If the LifeKeeper File System resource does not already exist when the Samba resource is created, LifeKeeper will create it for you.

**Note:** This directive is sometimes called directory. The recovery kit will handle both directive names.

- **printable** – A Yes value indicates that the Samba share is used as a print spool repository for printing to Linux printers. If the share is to be a regular file share then set this directive to No or do not specify it, as it is No by default unless set to Yes in the [global] section. If this directive is set to Yes, then creation of a Samba hierarchy will require the existence of LifeKeeper Print Services resource that protects the printer defined via the printer name directive listed below.

**Note:** This directive is sometimes called print ok. The recovery kit will handle both directive names.

- **printer name** – This directive defines the printer name used by the share and is used to find a Print Services instance that protects the named printer. The Print Services instance will become a child resource in the Samba hierarchy. If this directive is not defined for a printer share, the Samba Recovery Kit will use the share name as the printer name.

**Note:** This directive is sometimes called printer. The kit will handle both directive names.

 **Note:** The Samba configuration file allows the use of variable substitution for a number of directives. Variable substitution should not be used for any of the directives specified above unless the variable is resolved by the Samba utility *testparm*.

## 6.16.4.6. Running Multiple Instances of Samba

---

Running multiple instances of Samba in a LifeKeeper cluster introduces additional configuration requirements and restrictions. The following Samba configuration scenarios may involve multiple instances of Samba:

- Active/Standby configuration with multiple LifeKeeper Samba instances on one server
- Active/Active configuration with multiple LifeKeeper Samba instances on more than one server

Either of these configurations could include a non-LifeKeeper protected version of Samba.

As previously noted in [Configuring the LifeKeeper for Linux Samba Recovery Kit](#), when running multiple instances of Samba each version must have a uniquely named configuration file, or the files must reside in different directories. Within each configuration file a number of directives are required and must be unique – in particular, netbios name, lock directory, pid directory, interfaces and log file. If these directives are not unique, Samba may not startup and therefore will not be available for client connections. Additionally, the lock, log file, and pid directories specified for each instance must exist on all servers in the cluster.

### smbpasswd Utility and Multiple Instances of Samba

Although not required by LifeKeeper, some Samba utilities used by the Samba Recovery Kit expect to be able to open *smb.conf* in its default location. The Recovery Kit uses the **smbclient** and **nmblookup** utilities to connect to **smbd** and **nmbd** (respectively) in order to determine the health of the daemon processes while under LifeKeeper protection. These two utilities will not error out if they do not find *smb.conf* in its default location. However, *smb.conf* is required by the **smbpasswd** utility to be in its default location.

**smbpasswd** is used to maintain the *smbpasswd* file for authentication of users on client connection requests when the security level is set to share or user. If the default configuration file is missing, any attempt to change Samba passwords will fail. To avoid this problem, one of the instances of Samba must use the default configuration file if the security level is set to share or user, or if the server is acting as the **smbpasswd** server for those systems with Samba security level set to server. The reason for this is that **smbpasswd** uses the default configuration file to obtain the location of the *smbpasswd* file. Because of this requirement only one location for the *smbpasswd* file can exist within the LifeKeeper cluster. The configuration files for all instances of Samba in the cluster must have the directive `smb passwd file` set to the same value. Additionally, the *smbpasswd* file must be kept in sync on all servers in the cluster.

The **smbpasswd** utility is also affected by the use of the `bind interfaces only` directive, which is required by the LifeKeeper Samba Recovery Kit. With the `bind interfaces only` directive set to Yes, a regular user changing his Samba password will attempt to connect to a **smbd** daemon process using the localhost address of 127.0.0.1. If that address has been added to the `interfaces` directive in the configuration file used by the **smbd** daemon, and if **smbd** has connected to and is listening on that address, then the

password change will be successful. If the daemon does not have that address in its configuration file `interfaces` directive, then the password change will fail. In a multiple instance environment, if the `localhost` is specified in more than one configuration file, only one instance will be able to start up and run. Using the `-r netbios_name` option to **smbpasswd** will work in place of adding the `localhost` address to the `interfaces` list (for example: **smbpasswd -r server1 print1...**).

 **Note:** As previously stated, non-protected Samba instances running on a LifeKeeper server with protected Samba instances must also have the `bind interfaces only` directive set to “Yes”.

## Samba and User Authentication Considerations

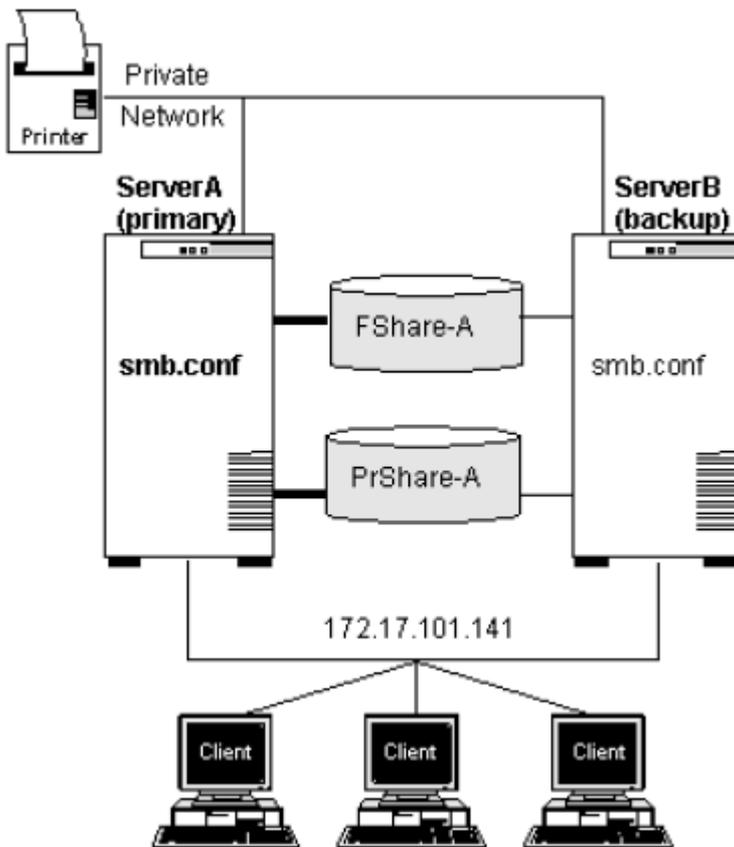
Samba supports several methods for user authentication via the security parameter (e.g. `share`, `user`, `domain`, ...) which must be considered when protecting Samba via LifeKeeper to ensure data files such as `/etc/samba/smbpasswd` or `/etc/samba/secrets.tdb` are kept in sync on all servers in the cluster. So when using security methods such as `user`, you must ensure that the `smbpasswd` file is kept in sync on all servers in the cluster. Additionally, security methods such as `domain` require synchronization of the `secrets.tdb` file. A LifeKeeper `active/active` configuration with the `secrets.tdb` file requires the use of the `private dir` parameter to specify the location of the file. The value for this parameter must be unique for each LifeKeeper Samba instance.

## 6.16.4.7. Samba Configuration Examples

This section contains definitions and examples of typical Samba configurations. Each example includes the configuration file entries that apply to LifeKeeper.

### Configuration 1 : Active/Standby Configuration

In the Active/Standby configuration, ServerA is the primary LifeKeeper server. It exports the file and print shares that reside on a shared storage device. While ServerB may be handling other applications/ services, it acts only as a backup for the Samba resources in LifeKeeper's context.



#### Configuration Notes:

- The clients connect to the Samba servers using the NetBIOS name LKServerA over the protected IP address (172.17.101.141), which is defined by the interfaces directive of the configuration file.
- The configuration file smb.conf has been copied to ServerB upon extension of the Samba resource hierarchy. It contains the following directives:

```
[global] netbios name = LKServerA bind interfaces only = yes
```

```
lock directory = /var/lock/samba interfaces = 172.17.101.141 127.0.0.1
```

```
log file = /var/log/sambaServA/log
```

```
[FShare-A]
```

path = /FShare-A

read only = no

public = yes

valid users =

printable = no

create mode = 0664

directory mode = 0775

[PRShare-A]

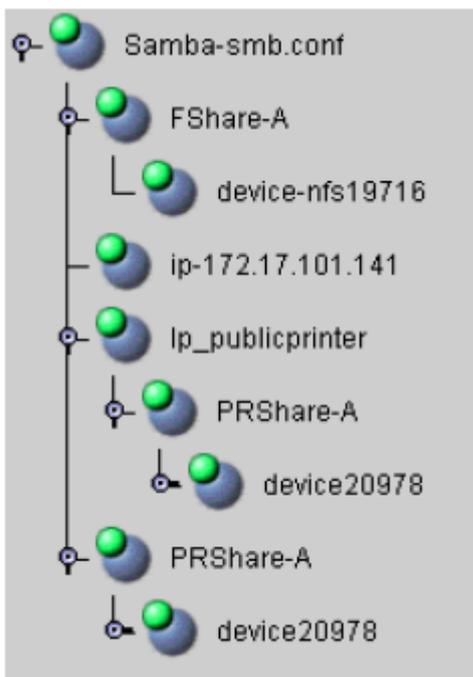
path = /PRShare-A

printer = publicprinter

printable = yes

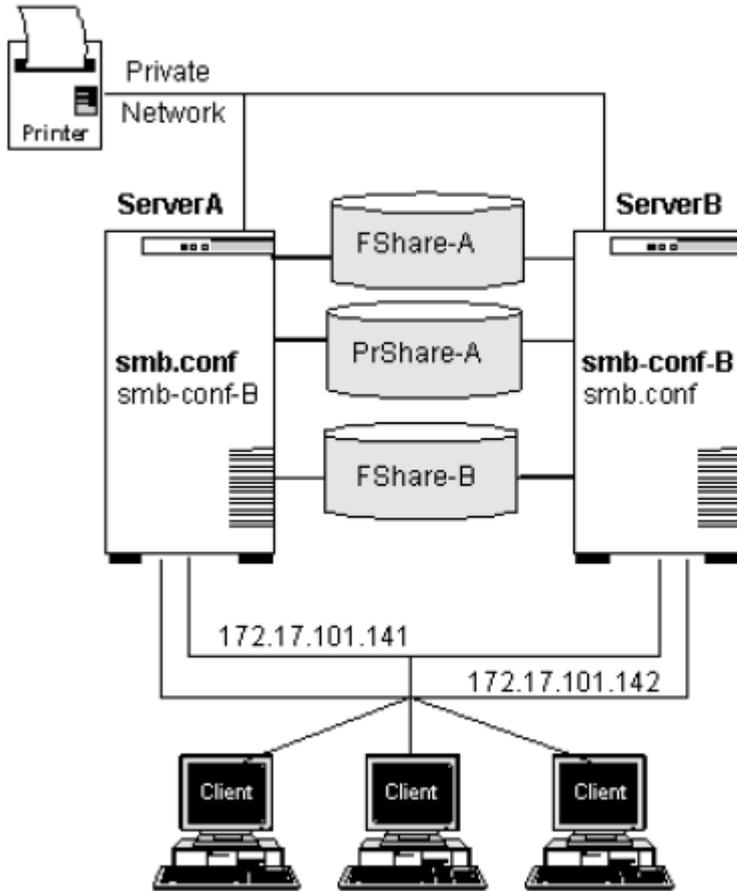
browseable = no

- The Samba resource hierarchy would look like the following:



## Configuration 2 : Active/Active Configuration

In the Active/Active configuration below, both ServerA and ServerB are primary LifeKeeper servers for Samba resources. Each server is also the backup server for the other.



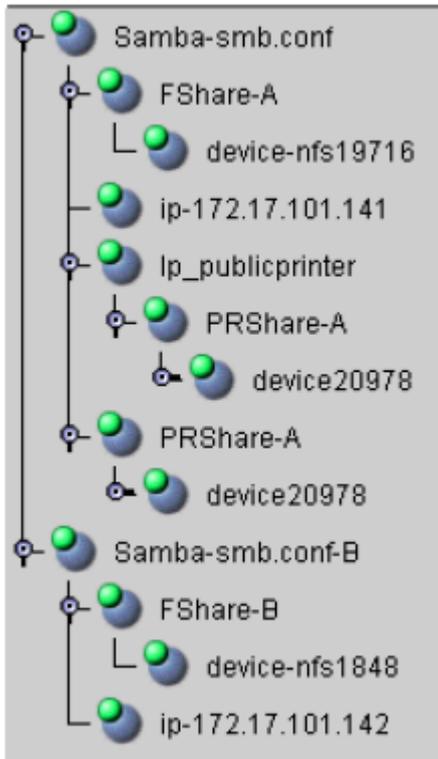
**Configuration Notes:**

- The clients connect to the Samba servers using the NetBIOS name LKServerA and LKServerB over the protected IP addresses (172.17.101.141 and 172.17.101.142 respectively), which are defined by the interfaces directive of the configuration files.
- The configuration file *smb.conf* was copied to ServerB upon extension of the Samba resource hierarchy. Likewise, the configuration file *smb.conf-B* was copied to ServerA upon extension of the Samba resource hierarchy.
- ServerA protects the file share */Fshare-A*; ServerB protects the file share */Fshare-B*.
- ServerA protects the print share */publicprinter*; ServerB does not protect a print share.
- The two configuration files contain the following directives:

smb.conf	smb-conf-B
<pre>[global] netbios name = LKServerA bind interfaces only = yes</pre>	<pre>[global] netbios name = LKServerB bind interfaces only = yes</pre>

lock directory = /var/lock/sambaServA	lock directory = /var/lock/
pid directory = /var/run/sambaServA	pid directory = /var/run/sambaServB
interfaces = 172.17.101.141 127.0.0.1	sambaServB interfaces = 172.17.101.142
log file = /var/log/sambaServA/log	log file = /var/log/sambaServB/log
[FShare-A]	[FShare-B]
path = /FShare-A	path = /FShare-B
read only = no	read only = no
public = yes	public = yes
valid users =	valid users =
printable = no create mode = 0664 directory mode = 0775	printable = no create mode = 0664 directory mode = 0775
PRShare-A	
path = /PRShare-A	
printer = publicprinter	
printable = yes browseable = no	

- The Samba resource hierarchies would look like the following:



## 6.16.5. Samba Configuration Steps

---

This section provides steps that you should take to configure your Samba resources.

1. Plan your Samba configuration. This includes the following:
  - NetBIOS name(s) to be used
  - The interfaces that will be protected and allowed access to the shares
  - The file systems to be used for the Samba shares and thus protected
  - The location of the lock and log directory (or directories)

Consideration should be given to the number of configuration files to be used and the type of configuration (Active/Standby vs. Active/Active). For example, if you have four Samba shares to protect, you could list all four shares in one configuration file, with the disadvantage that a failure of any one file system will cause the failover of the entire Samba hierarchy, including all four file shares. Alternatively, you could create four separate configuration files, each protecting one file share, which requires that four NetBIOS names be defined and managed.

2. Setup your Samba configuration file(s) based on the plan made in step 1. This includes the required directives in the [global] section as well as those for the file and print shares to be used. See [Configuring the LifeKeeper for Linux Samba Recovery Kit](#) for a discussion of the global and share directives required for LifeKeeper Samba hierarchies.
3. Create protected IP addresses under LifeKeeper, which will be used for client connections to the Samba server via the NetBIOS name. The protected IP address(es) should match the value(s) placed in the interface directive in the configuration file. (Refer to the [LifeKeeper for Linux IP Recovery Kit Administration Guide](#) for details on setting up IP resources.) Test the protected IP addresses by pinging them from all clients and other cluster servers. A protected IP resource for the local host (127.0.0.1) is not required.
4. Start the Samba daemons and test client connections.

- a. The commands to start the daemons are as follows:

```
s/nmbd -D -s ConfigurationFile
```

- b. Use the Samba utility smbclient to test connections to the smbd daemon as follows. This should be done for each address defined in the interfaces directive.

```
smbclient -L netbios_name -U% -I Protected_IP_Address
```

- c. Use **nmblookup** to test connection to the nmbd daemon process. This should be done for each broadcast address. Use the associated broadcast address for each address defined in the interfaces directive. (The broadcast address can be obtained by running **ip addr show**).

**nmblookup -B broadcast\_address netbios\_name**

5. Stop the Samba daemons started in the previous step. This is accomplished via the **kill** command. Find the running daemon processes via the **ps** command and issue a **kill pid** which will cause them to exit.
6. Create protected file system(s) under LifeKeeper that will host the Samba file and print shares as defined in the above steps. (Refer to [Creating a file system resource hierarchy](#) in the LifeKeeper for Linux Technical Documentation for information on creating a File System resource hierarchy.) This step may be skipped since File System resources will be created automatically when creating a Print Services resource or Samba resource. Replicated (SIOS DataKeeper) file system resources must be created before creating the Samba resource.
7. Create directories on the protected file systems for the shares should one file system be used for multiple Samba shares.
8. Create protected Print Services hierarchies under LifeKeeper, which will be used for client printing should any printer shares be defined in the configuration file.
9. Create the Samba resource hierarchy in LifeKeeper and extend it to at least one backup server (see the [LifeKeeper Configuration Tasks](#) section below). The extend script will copy the Samba configuration file from the template server to the same location on the target server.
10. On the primary server, test client connections to the shares that are protected by the Samba hierarchy which is in service. For instance, map the shared directory from a Windows client and ensure that it can access files on the share. Repeat the test for all servers in the cluster. You should also test your Samba resource by performing a manual switchover to a backup server. (See [Testing Your Resource Hierarchy](#).)
11. Automatic startup must be disabled at the time of system boot because Samba daemon protected by LifeKeeper is controlled by LifeKeeper.

## 6.16.6. LifeKeeper Configuration Tasks for Samba

---

You can perform all LifeKeeper for Linux Samba Recovery Kit administrative tasks via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer, and monitor Samba resources.

The following tasks are available for configuring the LifeKeeper for Linux Samba Recovery Kit:

- [Create a Resource Hierarchy](#) – Creates a Samba resource hierarchy.
- [Delete a Resource Hierarchy](#) – Deletes a Samba resource hierarchy.
- [Extend a Resource Hierarchy](#) – Extends a Samba resource hierarchy from the primary server to the backup server.
- [Unextend a Resource Hierarchy](#) – Unextends (removes) a Samba resource hierarchy from a single server in the LifeKeeper cluster.
- [Create Dependency](#) – Creates a child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete Dependency](#) – Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#) – Activates a resource hierarchy.
- [Out of Service](#) – Deactivates a resource hierarchy.
- [View](#) / [Edit](#) Properties – View or edit the properties of a resource hierarchy.

 **Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You can also perform most of these tasks:

1. From the toolbar, right – click on a global resource in the left pane of the status display.
2. Right – click on a resource instance in the right pane of the status display.

 **Note:** Using the right-click method allows you to avoid entering information that is required when using the **Edit** menu.

## 6.16.6.1. Creating a Samba Resource Hierarchy

After you have completed the necessary setup tasks, use the following steps to define the Samba resource hierarchy.

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the menu, select **Create Resource Hierarchy**.

The *Create Resource Wizard* dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

2. Select Samba Share and click **Next**.
3. You will be prompted to enter the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter an error requiring you to correct previously entered information. You may click **Cancel** at any time to cancel the entire creation process.

Field	Tips
Switchback Type	<p>Choose either <i>intelligent</i> or <i>automatic</i>. This dictates how the Samba instance will be switched back to this server when the server comes back up after a failover. The switchback type can be changed later from the General tab of the Resource Properties dialog box.</p> <p><b>Note:</b> The switchback strategy should match that of the Print Server, IP or File System resource to be used by the Samba resource. If they do not match the Samba resource creation will attempt to reset them to match the setting selected for the Samba resource.</p>
Server	Select the Server on which you want to create the hierarchy.
Location of Configuration File	Select the directory where the Samba configuration file is located.
Config File Name	<p>Enter the name of the Samba configuration file to be used for this resource creation. The default is <i>smb.conf</i>.</p> <p><b>Note:</b> LifeKeeper will read the selected configuration file, and if the file does not specify the required directives, LifeKeeper will generate an error message. It does minimal checking of the configuration file (to verify that shares exist, that they have a path directive, that a lock directory has been specified, and that the directory exists). Additional checking is done during the creation process.</p>

Root Tag	Either select the default root tag offered by LifeKeeper, or enter a unique name for the resource instance on this server. The default is Samba- <i>configfilename</i> , where <i>configfilename</i> is the name of the associated configuration file. You may use letters, numbers and the following special characters: – _ . /
----------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

4. Click **Create**. The Create Resource Wizard will then create your Samba resource hierarchy. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.
5. An information box will appear indicating that you have successfully created a Samba resource hierarchy, and you must Extend that hierarchy to another server in your cluster in order to achieve failover protection. Click **Next**.
6. Click Continue. LifeKeeper will then launch the Pre-Extend Wizard. Refer to Step 2 under [Extending a Resource Hierarchy](#) (below) for details on how to extend your resource hierarchy to another server.



**Note:** See [Failure Restoring Samba Hierarchy](#) in the Samba Troubleshooting section for tips to follow in the case that the creation of the Samba hierarchy fails.

## 6.16.6.2. Extending Your Samba Resource Hierarchy

This operation can be started from the **Edit** menu, or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should refer to Step 2 below.

1. On the **Edit** menu, select **Resource**, then Extend Resource Hierarchy. The Pre-Extend Wizard appears. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper **Extend Resource Hierarchy** defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The *Pre-Extend Wizard* will prompt you to enter the following information.

 **Note:** The first two fields appear only if you initiated the Extend from the **Edit** menu.

Field	Tips
Template Server	Enter the server where your Samba resource is currently in service.
Tag to Extend	Select the Samba resource you wish to extend.
Target Server	Enter or select the server you are extending to.
Switchback Type	Select either intelligent or automatic. The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box. <b>Note:</b> Remember that the switchback strategy must match that of the dependent resources to be used by the Samba resource.
Template Priority	Select or enter a priority for the template hierarchy. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (the number 1 indicates the highest priority). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. <b>Note:</b> This selection will appear only for the initial extend of the hierarchy.
Target Priority	Either select or enter the priority of the hierarchy for the target server.
Root Tag	LifeKeeper will provide a default tag name for the new Samba resource instance on the target server. The default tag name is the same as the tag name for this resource on the template

	server. If you enter a new name, be sure it is unique on the target server. You may use letters, numbers and the following special characters: – _ . /
--	--------------------------------------------------------------------------------------------------------------------------------------------------------

3. After receiving the message that the pre-extend checks were successful, click **Next**.
4. Depending upon the hierarchy being extended, LifeKeeper will display a series of information box showing the Resource Tags to be extended, which cannot be edited. Click **Extend**.
5. After receiving the message "Hierarchy extend operations completed" click **Next Server** to extend the hierarchy to another server, or click Finish if there are no other extend operations to perform.
6. After receiving the message "Hierarchy Verification Finished", click **Done**.

## 6.16.6.3. Unextending Your Samba Resource Hierarchy

---

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. On the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the Samba resource. It cannot be the server where the Samba resource is currently in service. (This dialog box will not appear if you selected the Unextend task by right clicking on a resource instance in the right pane.) Click **Next**.
3. Select the Samba hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the Unextend task by right clicking on a resource instance in either pane)
4. An information box appears confirming the target server and the Samba resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the Samba resource was unextended successfully. Click **Done** to exit the Unextend Resource Hierarchy menu selection.

## 6.16.6.4. Deleting a Samba Resource Hierarchy

---

It is important to understand what happens to dependencies and protected services when a Samba hierarchy is deleted.

- **Dependencies:** Before removing a resource hierarchy, you may wish to remove the dependencies. Dependent file systems will be removed unless they are used in another hierarchy. Dependent IP and Print Services resources will not be removed as long as the delete is done via the LifeKeeper GUI or the Samba delete script.
- **Protected Services:** If you take the Samba resource hierarchy out of service before deleting it, the Samba daemons will be stopped. If you delete a hierarchy while it is in service, the Samba daemons will continue running and offering services (without LifeKeeper protection) after the hierarchy is deleted.

To delete a resource hierarchy from **all** the servers in your LifeKeeper environment, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the **Target Server** where you will be deleting your Samba resource hierarchy and click **Next**. (This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in either pane.)
3. Select the **Hierarchy to Delete**. (This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in the left or right pane.) Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete**.
5. Another information box appears confirming that the Samba resource was deleted successfully.
6. Click **Done** to exit.

## 6.16.6.5. Testing Your Samba Resource Hierarchy

---

You can test your Samba resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

Selecting **Edit**, then **Resource**, then **In Service**. For example, an In Service request executed on a backup server causes the application hierarchy to be taken out of service on the primary server and placed in service on the backup server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the *Out of Service* request, the application is taken out of service without bringing it in service on the other server.

## 6.16.7. Samba Hierarchy Administration

---

Once your Samba resource hierarchies are created, follow these guidelines for ongoing administration of your Samba resources.

[Modifying the Samba Configuration File](#)

[Maintaining the smbpasswd File](#)

## 6.16.7.1. Modifying the Samba Configuration File

---

When changes are required to a Samba configuration file that is used in a LifeKeeper Samba instance, perform these procedures on the server that is In Service, Protected (ISP). There are three types of configuration file changes:

- Those that do not directly impact the Samba hierarchy
- Those that directly impact the hierarchy but do not require a delete and recreation of hierarchy
- Those that directly impact the hierarchy and require a delete and recreation of the hierarchy

### Modifications that do not directly impact the Samba Hierarchy

Any changes to configuration file directives not used by LifeKeeper fall into this category. (See [Configuring the LifeKeeper for Linux Samba Recovery Kit](#) for a list the directives used by the kit.) Example directives not used by LifeKeeper would include security, hosts allow, hosts deny and valid users to name a few. The procedures are as follows:

1. Take the Samba resource for the configuration file out of service. This step is required to stop the Samba daemons.
2. Make the necessary updates to the Samba configuration file.
3. Synchronize the configuration within the cluster. Use the utility `synccfg` to perform this task:

```
LKROOT/ikadm/subsys/gen/samba/bin/synccfg -t TargetSys -c ConfigFile
```

where *LKROOT* is the install location of LifeKeeper (*/opt/LifeKeeper* by default), *TargetSy* is the node to update and *ConfigFile* is the full path to the configuration file to copy.

4. Repeat the previous step for all servers in the hierarchy.
5. Bring the hierarchy back in service to restart the Samba daemons.

### Modifications that directly impact the Samba Hierarchy

Any changes to configuration file directives used by LifeKeeper (see [Configuring the LifeKeeper for Linux Samba Recovery Kit](#) for a list), with the exception of the netbios name or the physical movement of the configuration file, fall into this category. Depending on the extent of the changes, it may be quicker and easier to proceed to the third category and just recreate the hierarchy. The typical types of changes expected in this category include the addition of new file and print shares, removal of file and print shares or the addition or removal of IP interfaces.

1. Take the Samba resource for the configuration file out of service. This step is required to stop the Samba daemons.
2. Make the necessary updates to the Samba configuration file.
3. Make the necessary updates to the Samba hierarchy. This varies depending on the type of change made to the configuration file. For example:

° If an additional IP address has been added to the interfaces directive, then a new IP resource needs to be created, extended and then added as a dependent child to the Samba resource hierarchy. See *Creating a Resource Dependency in the LifeKeeper for Linux Technical Documentation* for information on how to create dependencies.

° If a new file share has been added to the configuration file, then a File System resource may need to be created, extended and added as a dependent child to the Samba resource hierarchy. If the File System resource already exists as a child in the hierarchy (e.g. the path directive defined for the new share has the same file system mount point as another file or print share) then it does not need to be created and added as a dependent child.

° If a new print share is added, then File System and Print Services resources need to be created, extended and added as dependent children in the Samba hierarchy. If a print services resource does not exist that protects the printer as defined by the print share name or printer/printer name directive, then one must be created. See file share above to determine if a file system resource needs to be added.

° If a file or print share is removed, or if an IP address is removed from the interfaces directives, delete the dependency in the Samba hierarchy and then delete the individual resource.

° If a print share name is changed, follow the delete of print share followed by the addition of new print share.

4. Synchronize the configuration within the cluster. Use the utility **synccfg** to perform this task:

**LKROOT/lkadm/subsys/gen/samba/bin/synccfg -t TargetSys -c ConfigFile**

where *LKROOT* is the install location of LifeKeeper , *\_TargetSys* is the server to update and *ConfigFile* is the full path to the configuration file to copy.

5. Repeat the previous step for all servers in the hierarchy.
6. Bring the hierarchy back in service to restart the Samba daemons.

**Note:** If you are making a number of changes that require numerous resource creations and dependency additions or deletions, you may wish to create all the new resources before you take the Samba hierarchy out of service so that downtime is minimized.

## Modifications that directly impact the Samba Hierarchy, requiring a deletion and recreation of the Hierarchy

If the netbios name directive is changed or the physical location of the configuration file is changed, then you must:

1. Delete the hierarchy. (See [Deleting a Resource Hierarchy](#) for details.)
2. Change the NetBIOS name or move the configuration file.
3. Create a new Samba hierarchy and extend to all backup servers.

## 6.16.7.2. Maintaining the `smbpasswd` File

---

Samba provides four different authentication methods via the security directive. The share and user security settings both require access to the local `smbpasswd` file to determine if access will be granted. As noted in the section [Running Multiple Instances of Samba](#) there can only be one `smbpasswd` file, and this presents a potential administration problem in a LifeKeeper cluster. If share or user level security is selected, the file must be kept in sync on all servers so that authentication will succeed after a failover.

In a cluster with only one Samba hierarchy, the use of share or user level security can be accomplished by placing the `smbpasswd` file on a file share defined in the configuration file. The access to this share should be such that only administrators have access. In a multiple instances scenario, either server level or domain level security is suggested.

 **Note:** If firewalls are in use, ensure that the firewall will allow connections to the `smbd` daemon, and that the `nmblookup` will work.

## 6.16.8. Samba Troubleshooting

---

### Failure restoring Samba hierarchy

Failure of a Samba hierarchy restore can leave the daemon processes **smbd** and **nmbd** running. (The restore operation is initiated via the completion of a create, failover, or manual switchover, or via a local recovery caused by a connection failure to **smbd** and **nmbd**.) If the daemons are not stopped and restarted after the problem is corrected, and the restore is attempted again, the restore could fail again.

#### Suggested Action:

Correct the cause of the connection failure (for instance, an incorrect mask setup for the interfaces directive). Next, manually stop the **smbd** and/or **nmbd** daemons. Then bring the hierarchy in service, or re-create the hierarchy if the failure occurred during creation.

Stopping the daemons ensures that a re-read of the configuration file occurs during the restore.

## 6.16.8.1. Common Samba Error Messages

This section provides a list of messages that you may encounter with the use of the LifeKeeper Samba Recovery Kit. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Because the Samba Recovery Kit relies on other LifeKeeper components to drive the creation and extension of hierarchies, messages from these other components are also possible. In these cases, please refer to the **Message Catalog** (located on our Technical Documentation site under **Search for an Error Code**) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

### Common Error Messages

Error Number	Error Message	Description
109009	Error getting netbios name from the instance information field for tag "Samba-smb.conf" on server "ServerA".	Extracting the NetBIOS name from the instance info failed. Check that the info field contains the configuration file and NetBIOS name.
109015	The Samba utility testparm failed. Unable to parse Samba configuration file.	The Samba utility <b>testparm</b> that is used to parse the configuration file failed. Run <b>testparm</b> from the command line specifying the configuration file used for the hierarchy to determine the failure.
109019	Failed to initialize for reading of the Configuration File "/tmp/smb.ini.1234".	Attempts to read the generated output of the configuration file failed. The utility <b>testparm</b> generated a bad file.
109022	Error getting configuration name from the instance information field for tag "Samba- smb.conf" on server "ServerA".	Extracting the NetBIOS name from the instance info failed. Check that the info field contains the configuration file and NetBIOS name.
109030	Failure opening "/var/lock/samba/smbd.pid" on server "ServerA": "File Not Found"	The attempted open of the daemon process ID file failed for the listed reason. Correct the problem based on the listed error code.
109050	Open of the testparm output file failed.	The open of output file created by running testparm failed because the file does not exist or does not contain any data. Run <b>testparm</b> from the command line specifying the configuration file used for the hierarchy to determine the failure.

## 6.16.8.2. Hierarchy Creation

Error Number	Error Message	Description
109001	Usage: "valid_cf" CfgPath CfgName TemplateSys	The <b>valid_cf</b> script requires three arguments, the directory containing the configuration file, the name of the configuration file and the name of the template system on which to validate the configuration file. You must specify all three.
109002	Must specify an absolute path to the "smb.conf" configuration file.	You must specify the absolute path to the configuration file when running the scripts <b>choice_cf</b> and <b>valid_cf</b> .
109003	The file "smb.conf" does not exist in "/etc".	You must specify the correct path to the configuration file when running <b>valid_cf</b> to validate the select configuration file.
109004	The path "/export/fs" in "/etc/samba/smb.conf" does not reside on a shared file system.	The Samba configuration shares must contain a path directive that can be protected by LifeKeeper via a File System resource. Edit the configuration file and change the path to a file system that LifeKeeper can protect.
109005	Usage: "choices_cf" CfgFile	The <b>choice_cf</b> script requires the full path to an existing configuration file. Please specify the correct path.
109006	Samba Configuration file not specified.	No configuration file was specified for the creation of the Samba resource hierarchy.
109007	Cannot bring hierarchy "Samba-smb.conf" in service on server "ServerA". <b>Action:</b> After correcting the problem, try bringing the hierarchy in service manually.	The in service attempt at the end of creation failed. View the log file for possible reason for the failure.
109010	The Samba configuration file does not have the interfaces directive defined. This directive is required to create Samba File Share hierarchies.	The configuration file is missing the interfaces directive or the directive does not contain any IP addresses other than the localhost (127.0.0.1).
109011	The Samba configuration file does not have a correctly formatted interfaces directive. The interfaces directive must be in full dotted decimal IP address format with or without the mask parameter.	The interfaces directive must contain one or more IP addresses in the format of aaa.bbb.ccc.ddd or aaa.bbb.ccc.ddd/mask separated by a space.

109012	The Samba configuration file section "FileShare1" does not have a path directive defined. All Samba shares must have a path directive.	The specified share in the configuration file does not have a path directive or the directive does not contain a value.
109014	No IP resources defined on server "ServerA". <b>Action:</b> Create IP resources for the IP addresses defined in the Samba configuration file interfaces directive.	The specified server does not contain any LifeKeeper protected IP resources needed for the creation of the Samba resource hierarchy.
109016	The IP(s) "100.25.104.25,100.35.104.26" defined in the interfaces directive are not under LifeKeeper protection. <b>Action:</b> Create IP resources for the unprotected addresses defined in the Samba configuration file interfaces directive.	LifeKeeper protected IP resources must exist for all of the IP addresses listed in the interfaces directive. Those listed are not protected by LifeKeeper.
109017	Missing configuration file name.	No configuration file exists when attempting to run <b>testparm</b> during Samba resource creation.
109018	No Samba shares found in "/etc/samba/smb.conf".	The Samba configuration specified for the resource must contain at least one file or print share.
109020	Bad configuration file. No section information found in the file "/tmp/smb.ini.1234".	The Samba configuration specified for the resource must contain at least one file or print share.
109021	Creation of Samba hierarchy with tag "Samba-smb.conf" on server "ServerA" failed.	The create of the Samba hierarchy failed. Examine the other error messages to determine the cause of the failure.
109023	The file system resource "filesys1328" is not in-service on server "ServerA".	The File System resource needed as a dependent child in the Samba hierarchy is not in service on the template server. Bring the resource in service and retry the resource creation.
109024	Selected IP resource "ip-100.25.104.26" does not exist on server "ServerA". <b>Action:</b> Retry the operation.	The specified IP resource tag no longer exists and is needed for the creation of the Samba hierarchy. Recreate the IP resource and retry the resource creation.
109025	LifeKeeper was unable to create a	The dependency creation attempt between the

	dependency between the Samba hierarchy "Samba-smb.conf" and the IP resource "ip-100.25.104.26" on server "ServerA".	Samba resource and the IP resource failed. Examine the other messages to determine the cause of the failure.
109026	The Samba configuration file does not have a netbios name directive defined. <b>Action:</b> Add a netbios name directive in the global section of the configuration file.	All configuration files must contain a NetBIOS name. Add a NetBIOS name directive to the configuration file.
109029	The Samba configuration file "%s" directive defines a directory that does not exist. The %s can contain pid directory,lock dir, or lock directory.	All configuration files must contain an existing directory as specified by the directive. Add the missing directory.
109035	The Samba directive "bind interfaces only" must be set to "Yes". <b>Action:</b> Change "bind interfaces only" to "Yes" and recreate the hierarchy.	All configuration files must have the "bind interfaces only" directive set to Yes. Correctly set the directive to Yes.
109036	Selected Printer resource "lp-admin" does not exist on server "ServerA". <b>Action:</b> Retry the operation.	The specified Print Services resource tag no longer exists and is needed for the creation of the Samba hierarchy. Recreate the Print Services resource and retry the resource creation.
109037	LifeKeeper was unable to create a dependency between the Samba hierarchy "Samba-smb.conf" and the Printer resource "lp-admin" on server "ServerA".	The dependency creation attempt between the Samba resource and the Print Services resource failed. Examine the other messages to determine the cause of the failure.
109038	The Printers(s) "lpadmin" defined in the configuration file are not under LifeKeeper protection. <b>Action:</b> Create Printer instances for the unprotected printers defined in the Samba configuration file.	LifeKeeper protected Print Services resources must exist for all of the printers defined in the configuration file. Those listed are not LifeKeeper protected.
109041	The selected configuration file "/etc/samba/smb.conf" is in use by Samba	A Samba configuration file or netbios name can only be protected once in the cluster. Rename the configuration file or select a new NetBIOS name.

	<p>resource "Samba-smb.conf".</p> <p>Or</p> <p>The selected netbios name "LKServer" is in use by Samba resource "Samba-smb.conf".</p>	
--	---------------------------------------------------------------------------------------------------------------------------------------	--

## 6.16.8.3. Hierarchy Extension

Error Number	Error Message	Description
109008	<p>Replication of config file to target server "ServerB" failed. The "mkdir" of "/etc/samba/config_files" failed.</p> <p>Or</p> <p>Replication of config file to target server "ServerB" failed. The "/etc/samba/config_files" failed.</p>	The attempted copy of the Samba configuration file on the extend failed. Either the <b>mkdir</b> or remote copy failed.
109042	<p>WARNING: The configuration file "/etc/samba/smb.conf" currently exists on server "ServerB" and will be overwritten if this resource is extended.</p>	The configuration file used for the resource hierarchy already exists on the backup server and will be overwritten. Cancel the extension to abort the overwriting of the file.

## 6.16.8.4. Restore

---

Error Number	Error Message	Description
109027	Failed start of smbd as daemon process.	The attempt to start <b>smbd</b> as a daemon failed. Check the Samba log files for additional information.
109028	Failed start of nmbd as daemon process.	The attempt to start <b>nmbd</b> as a daemon failed. Check the Samba log files for additional information.

## 6.16.8.5. Remove

---

Error Number	Error Message	Description
109033	Failed start of smbd as daemon process. Attempting to stop via SIGKILL.	The normal termination of the <b>smbd</b> daemon process failed so the daemon will be forcibly terminated.
109034	Failed to stop nmbd daemon process. Attempting to stop via SIGKILL.	The normal termination of the <b>nmbd</b> daemon process failed so the daemon will be forcibly terminated.

## 6.16.8.6. Resource Monitoring

Error Number	Error Message	Description
109031	Connection attempt to <b>smbd</b> daemon on address "100.25.104.26" failed.	The health check of the <b>smbd</b> daemon process failed when a connection attempt on the listed IP address failed. A local recovery will be attempted. If the local recovery fails, the hierarchy will be switch over to the backup server.
109032	Connection attempt to <b>nmbd</b> daemon failed for broadcast address "100.25.107.255" using netbios name "FILESERV1".	The health check of the <b>nmbd</b> daemon process failed when a connection attempt on the listed broadcast address failed. A local recovery will be attempted. If the local recovery fails, the hierarchy will be switched over to the backup server.
109039	No dependent IP resources were found for tag "Samba-smb.conf" on server "ServerA".	No IP children were found for the Samba hierarchy when attempting to ascertain the health of the <b>nmbd</b> daemon. Examine the logs to determine the failure.
109040	Unable to determine IP address and/or mask for IP resource "ip-100.25.104.25".	Extracting the IP address and Mask from the info for the IP resource failed. Check that the info field contains the IP address and mask.

## 6.16.8.7. Configuration File Synchronization Utility

Error Number	Error Message	Description
109013	The hierarchy for the specified configuration file does not have an equivalency with the target system. Select another target system or configuration file.	The -t TargetSys argument to the <b>synccfg</b> utility specified a system that does not contain an equivalency for the Samba hierarchy on the local system.
109044	Usage: synccfg -t TargetSys -c ConfigFile	The <b>synccfg</b> utility requires two arguments, -t TargetSys where the configuration file will be copied and -c ConfigFile for the name of the configuration to copy. You must specify both arguments.
109045	Specified Configuration file does not exist on this server.	The -c ConfigFile argument to the <b>synccfg</b> utility specified a configuration that does not exist on this system.
109046	The target system specified for updating is the same as the template system. Select a new target system.	The -t TargetSys argument to the <b>synccfg</b> utility specified the local server as the target of the copy. The target server must not be the local system.
109047	The specified configuration file is not used in any Samba instances. Select another configuration file.	The -c ConfigFile argument to the <b>synccfg</b> utility specified a configuration file that is not used in any LifeKeeper Samba instance. Only configuration file protected by LikeKeeper can be copied with this utility.
109048	The synchronization of the Samba configuration file failed with a status of "%s".	The attempted synchronization of the configuration files between the local system and the target system failed. The status provides the reason for the failure.
109049	The hierarchy for the specified configuration file is ISP on the target node. Select another target system or configuration file.	The -t TargetSys argument to the <b>synccfg</b> utility specified a system where the Samba hierarchy is currently ISP. The target system must not be ISP.

## 6.17. SAP Recovery Kit Administration Guide

---

The LifeKeeper for Linux SAP Recovery Kit provides a mechanism to recover SAP NetWeaver from a failed primary server onto a backup server in a LifeKeeper environment. The SAP Recovery Kit works in conjunction with other LifeKeeper Recovery Kits (the [IP Recovery Kit](#), [NFS Server Recovery Kit](#), [NAS Recovery Kit](#) and a database recovery kit, e.g. [Oracle Recovery Kit](#)) to provide comprehensive failover protection.

This documentation provides information critical for configuring and administering your SAP resources. You should follow the configuration instructions carefully to ensure a successful LifeKeeper implementation. You should also refer to the documentation for the related [recovery kits](#).

[SAP Recovery Kit Overview](#)

### Documentation

The following is a list of LifeKeeper related information available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#) (also available from the Help menu within the LifeKeeper GUI)
- [SIOS Technology Corp. Documentation and Support](#)
- [Abbreviations and Definitions](#) – Contains a list of abbreviations and terms that are used throughout this documentation along with their meaning.
- [LifeKeeper/SAP Icons](#) – Contains a list of icons being used and their meanings.

### Reference Documents

The following are documents associated with SAP that are referenced throughout this documentation:

- *SAP R/3 in Switchover Environments* (SAP document 50020596)
- *R/3 Installation on UNIX: (Database specific)*
- *SAP Web Application Server in Switchover Environments*
- *Component Installation Guide SAP Web Application Server (Database Specific)*
- *SAP Notes 7316, 14838, 201144, 27517, 31238, 34998 and 63748*

## 6.17.1. SAP Abbreviations and Definitions

The following abbreviations are used throughout this documentation:

Abbreviation	Meaning
<b>AS</b>	SAP Application Server. Although AS typically refers to any application server, within the context of this document, it is intended to mean a non-CI, redundant application server. Thus, the application server is not required for protection by LifeKeeper.
<b>ASCS</b>	ABAP SAP Central Services Instance. This is the SAP instance that contains the Message and Enqueue services for the NetWeaver ABAP environment. This instance is a single point of failure and must be protected by LifeKeeper.
<b>(ASCS)</b>	The backup ABAP SAP Central Services Instance server. This is the server that hosts the ASCS when the primary ASCS server fails.
<b>DB</b>	The SAP Database instance. This database may be Oracle or any other database supported by SAP. This instance is a single point of failure and must be protected by LifeKeeper. Note that the CI and DB may be located on the same server or different servers. DB is also used to denote the Primary DB Server.
<b>(DB)</b>	The backup Database server. This is the server that hosts the DB when the primary DB server fails. Note that a single server might be a backup for both the Database and Central Instance.
<b>ENSAv1</b>	Standalone Enqueue Server Version 1. This is the version of the enqueue server available in SAP kernel versions prior to 7.51.
<b>ENSAv2</b>	Standalone Enqueue Server Version 2. This version of the enqueue server is available in SAP kernel versions 7.51 and later.
<b>ERS</b>	Enqueue Replication Server.
<b>ERSv1</b>	Enqueue Replication Server Version 1. This is the version of the enqueue replication server available in SAP kernel versions prior to 7.51.
<b>ERSv2</b>	Enqueue Replication Server Version 2. This version of the enqueue replication server is available in SAP kernel versions 7.51 and later.
<b>HA</b>	Highly Available; High Availability.
<b>ID or &lt;ID&gt;</b>	Two digit numerical identifier for an SAP instance.
<b>&lt;INST&gt;</b>	Directory for an SAP instance whose name is derived from the services included in the instance and the instance number, for example a CI <INST> might be <i>DVEBMGS00</i> .
<b>PAS</b>	Primary Application Server Instance.
<b>SAP Instance</b>	A group of processes that are started and stopped at the same time.
<b>SAP System</b>	A group of SAP Instances.

<b>&lt;sapmnt&gt;</b>	SAP home directory which is <i>/sapmnt</i> by default but may be changed by the user during installation.
<b>SCS</b>	SAP Central Services Instance. This is the SAP instance that contains the Message and Enqueue services for the NetWeaver Java environment. This instance is a single point of failure and must be protected by LifeKeeper.
<b>(SCS)</b>	The backup SAP Central Services Instance server. This is the server that hosts the SCS when the primary SCS server fails.
<b>SID or &lt;SID&gt;</b>	System ID.
<b>sid or &lt;sid&gt;</b>	Lower case version of SID.
<b>SPOF</b>	Single Point of Failure.

## 6.17.2. LifeKeeper – SAP Icons

The following icons are significant on how to interpret the status of SAP resources in a LifeKeeper environment. These icons will show up in the LifeKeeper UI.

	<b>Active</b> – Resource is active and in service (Normal state).
	<b>Standby</b> – Resource is on the backup node and is ready to take over if the primary resource fails (Normal state).
	<b>Failed</b> – Resource has failed; you can try to put the resource in service (right-click on the resource and scroll down to “In Service” and enter). If the resource fails again, then recovery has failed (Failure state).
	<b>Attention needed</b> – SAP resource has failed or is in a caution state. If it has failed and automatic recovery is enabled (Protection Level Full or Standard), then LifeKeeper will try to automatically recover the resource. Right-click on the SAP resource and choose <b>Properties</b> . This will show which resource has a caution state. An SAP state of <b>Yellow</b> may be normal, but it signifies that SAP resources are running slow or have performance bottlenecks.
	<b>Update protection level</b> – Action icon to allow changing the SAP resource protection level.
	<b>Update recovery level</b> – Action icon to allow changing the SAP resource recovery level.
	<b>Handle warnings</b> – Action icon to set failure on warnings.
	Action icon to select different actions, such as <b>start</b> , <b>stop</b> , <b>migrate</b> or <b>setting maintenance mode</b> .

## 6.17.3. SAP Recovery Kit Overview

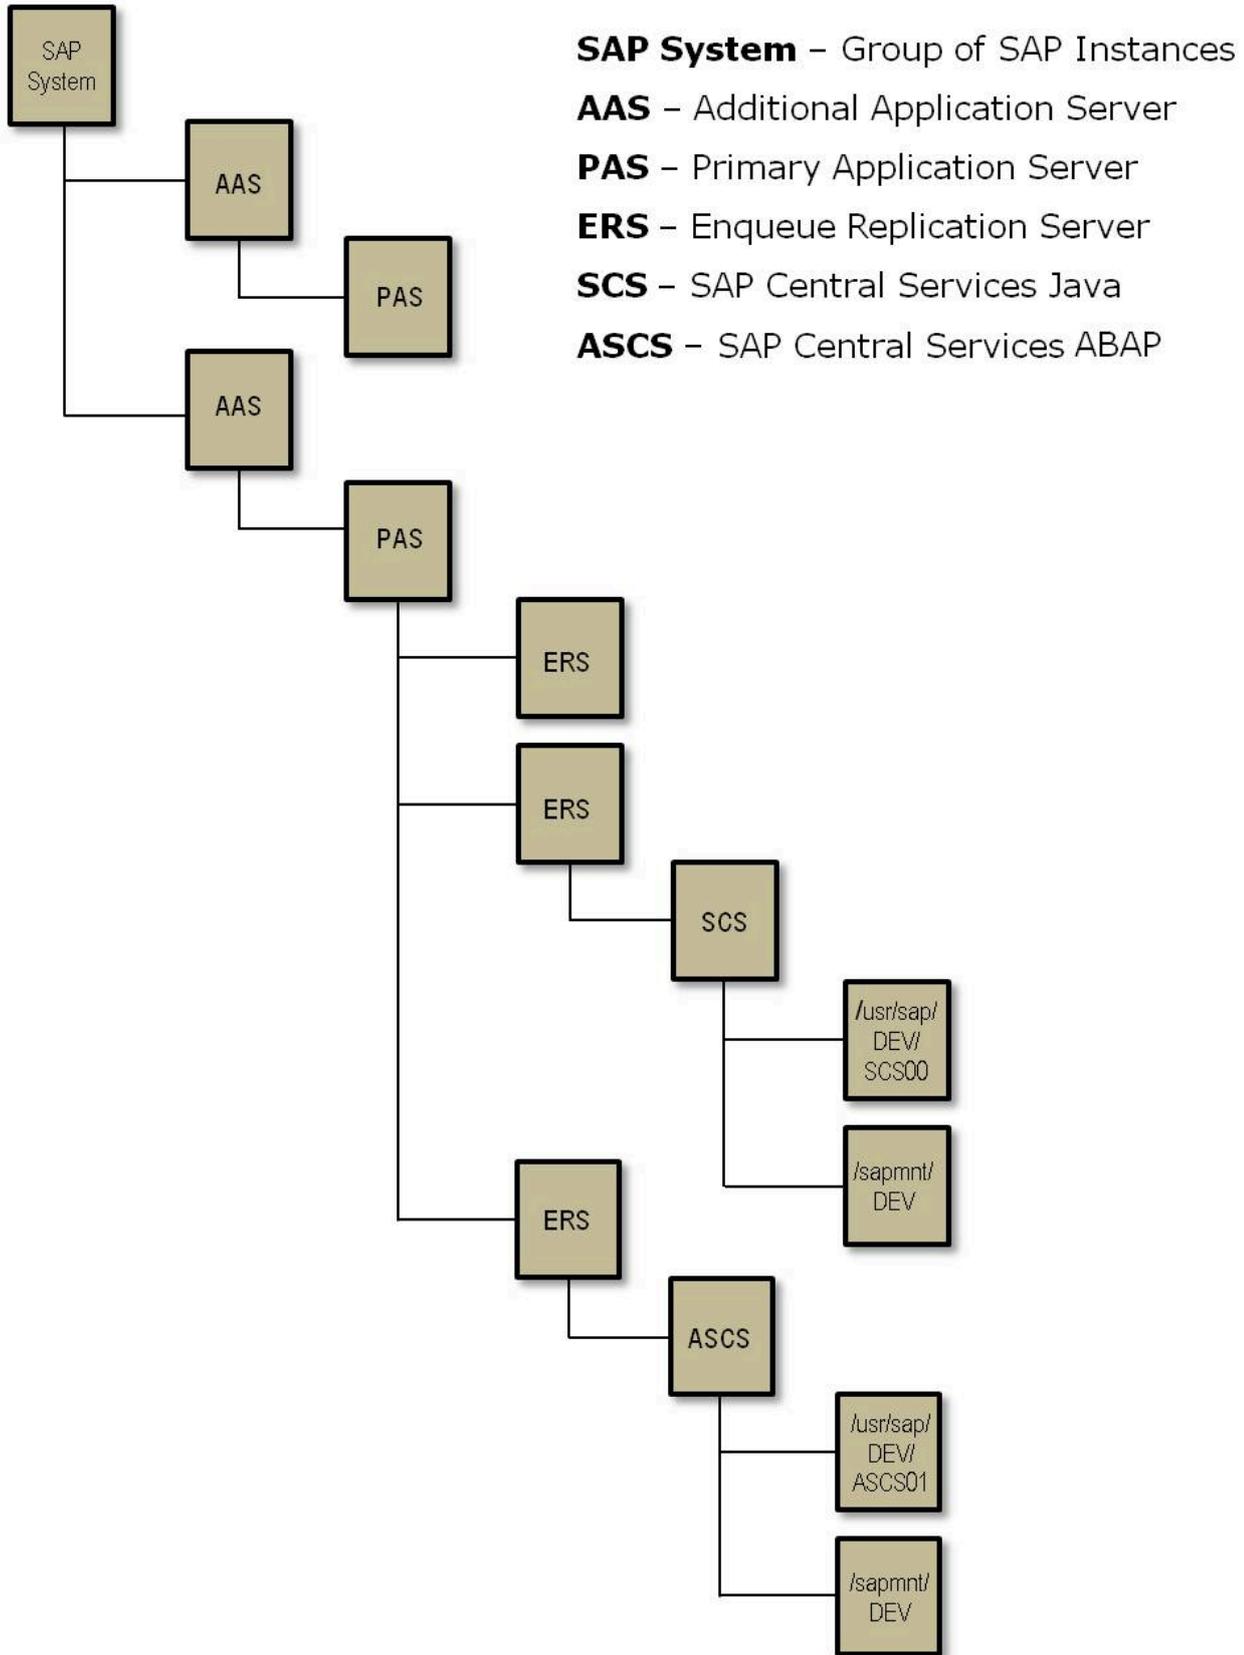
---

There are some services in the SAP NetWeaver framework that cannot be replicated. They cannot exist more than once for the same SAP system, therefore, they are single points of failure. The LifeKeeper SAP Recovery Kit provides protection for these single points of failure with standard LifeKeeper functionality. In addition, the kit provides the ability to protect, at various levels, the additional pieces of the SAP infrastructure. The protection of each infrastructure component will be represented in a single resource within the hierarchy.

The SAP Recovery Kit provides monitoring and switchover for different SAP instances; the SAP Primary Application Server (PAS) Instance, the ABAP SAP Central Service (ASCS) Instance and the SAP Central Services (SCS) Instance (the Central Service Instances protect the enqueue and message servers). The SAP Recovery Kit works in conjunction with the appropriate database recovery kit to protect the database, and with the Network File System (NFS) Server Recovery Kit to protect the NFS mounts. The IP Recovery Kit is also used to provide a virtual IP address that can be moved between network cards in the cluster as needed. The Network Attached Storage (NAS) Recovery Kit can be used to protect the local NFS mounts. The various recovery kits are used to build the SAP resource hierarchy which provides protection for all of the components of the application environment.

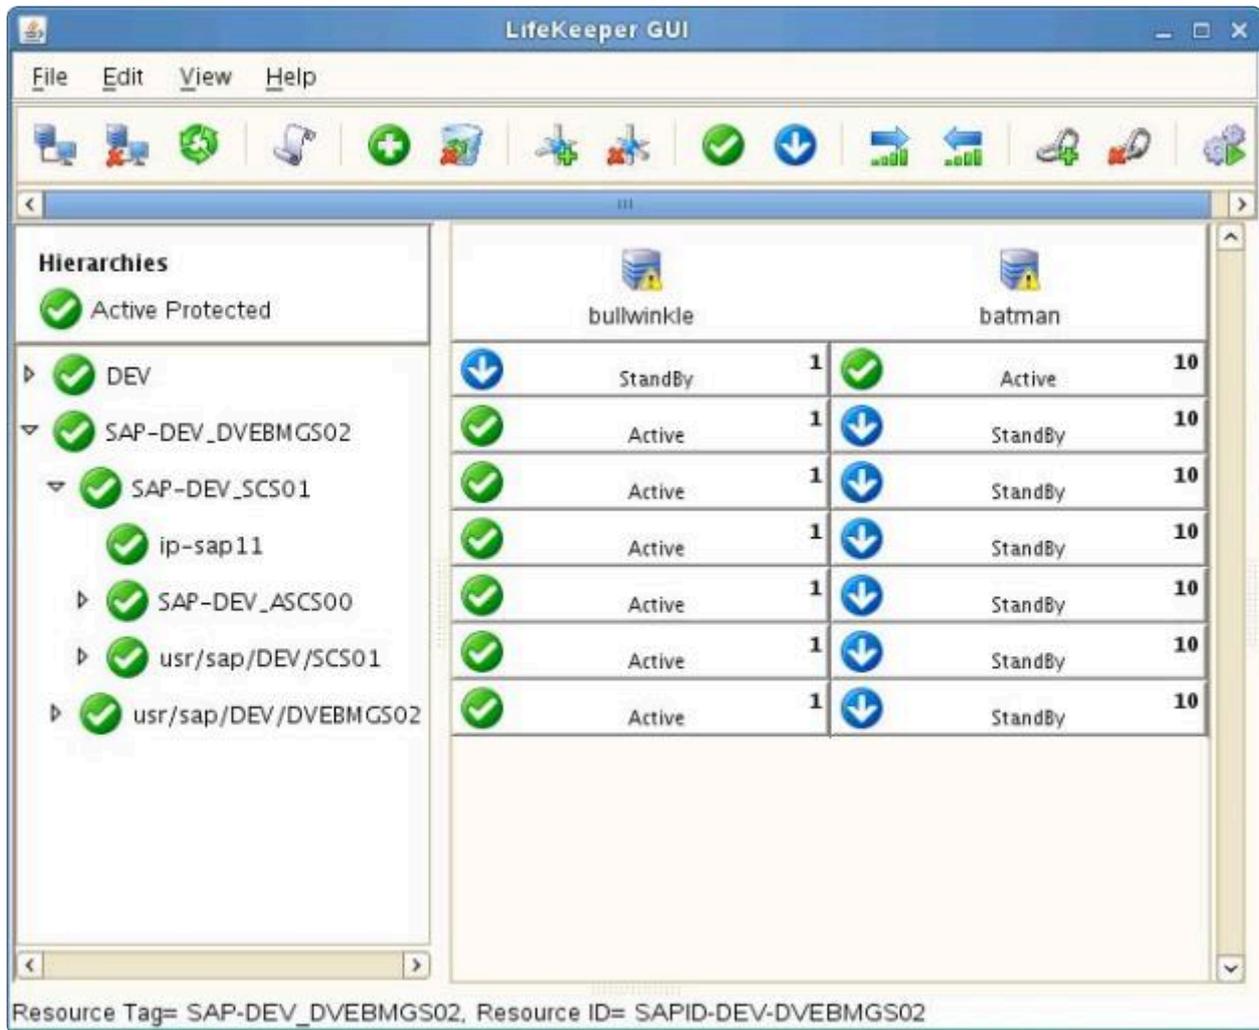
Each recovery kit monitors the health of the application under protection and is able to stop and restart the application both locally and on another cluster server.

### **Map of SAP System Hierarchy**



A typical SAP resource hierarchy as it appears in the LifeKeeper GUI is shown below.

## SAP Resource Hierarchy



\* **Note:** The directory `/usr/sap/trans` is optional in SAP environments. The directory does not exist in the SAP NetWeaver Java only environments.

\* **Note:** An ERS resource created in LifeKeeper for Linux v9.4.0 or later will operate in its own independent hierarchy. See [ERS Resource Types in LifeKeeper](#) for more details.

# 6.17.4. LifeKeeper SAP Solution Page

## SAP High Availability Interface 7.73

\* **Note:** The operating system and configuration used in your SAP deployment must be supported by SAP, SIOS, and your infrastructure provider (public/private cloud or on-premise). Consult SAP’s Product Availability Matrix (PAM) in order to determine which operating systems are supported for various versions of SAP NetWeaver and S/4HANA.

Package	Version	OS/Application Version Support
<a href="#">LifeKeeper Core (SAP)</a>	9.6.1	Red Hat Enterprise Linux 7 (Up to 7.9) Red Hat Enterprise Linux 8 (Up to 8.4) SUSE LINUX Enterprise Server (SLES) 12 (SP1 to SP5) SUSE LINUX Enterprise Server (SLES) 15 (Up to SP3) Oracle Linux 7.0 to 7.6 (including UEK R4)
<a href="#">LifeKeeper SAP Recovery Kit</a>	9.6.1	SAP NetWeaver 7.3 including Enhancement Package 1 SAP NetWeaver 7.4 SAP NetWeaver 7.5 SAP NetWeaver AS for ABAP 7.51 innovation package SAP NetWeaver AS for ABAP 7.52 innovation package SAP S/4HANA 1809 Platform SAP S/4HANA 1909 Platform SAP S/4HANA 2020 Platform SAP S/4HANA 2021 Platform
<a href="#">LifeKeeper NFS Server Recovery Kit</a>	9.6.1	NFS exported file systems on Linux distributions with a kernel version of 2.6 or later
LifeKeeper Supported Databases for SAP	9.6.1	Oracle, SAP ASE (Sybase), SAP HANA, SAP MaxDB, IBM DB2
SAP HA Interface Connector	7.73	SAP High Availability Interface 7.73, 7.77, 7.81, 7.85

\* Beginning in v9.5.0 SIOS has released the new SAP HANA Application Recovery Kit. SIOS will continue to support the SAP HANA gen/app based Recovery Kit with the 9.4.x releases until **March 31, 2022**. If you are using LifeKeeper for Linux v9.5 or later you must use the new (built-in) SAP HANA Application Recovery Kit.

! The existing SAP HANA gen/app based Recovery Kit is **not** supported with LifeKeeper for Linux v9.5.0 or later. Users who wish to upgrade to the LifeKeeper for Linux v9.5.0 or later **must** convert their existing SAP HANA gen/app based Recovery Kit to the new [SAP HANA Recovery Kit](#). Refer to [Upgrading from the SAP HANA Gen/App to the SAP HANA Recovery Kit](#) for details.

\* **NOTE:** Operating system versions built for enhanced SAP support (such as Red Hat Enterprise Linux for SAP Business Applications, Red Hat Enterprise Linux for SAP Solutions, and SUSE Linux Enterprise Server for SAP Applications) are also supported as long as the running Linux kernel version is the same as one of the supported OS versions listed above.

## 6.17.5. SAP Hardware and Software Requirements

---

Before installing and configuring the LifeKeeper SAP Recovery Kit, be sure that your configuration meets the following requirements:

- **Servers.** This recovery kit requires two or more computers configured in accordance with the requirements described in the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#).
- **Shared Storage.** SAP Primary Application Server (PAS) Instance, ABAP SAP Central Service (ASCS) Instance, SAP Central Services (SCS) Instance and program files must reside on shared disk(s) in a LifeKeeper environment. The file system for the ERS instance may also reside on a shared disk in the case of an ERSv2 configuration.
- **IP Network Interface.** Each server requires at least one Ethernet TCP/IP-supported network interface. In order for IP switchover to work properly, user systems connected to the local network should conform to standard TCP/IP specifications.

**Note:** Even though each server requires only a single network interface, multiple interfaces should be used for a number of reasons: heterogeneous media requirements, throughput requirements, elimination of single points of failure, network segmentation and so forth. (See [IP Local Recovery and Configuration Considerations](#) for additional information.)

- **Operating System.** Linux operating system. (See the [LifeKeeper for Linux Release Notes](#) for a list of supported distributions and kernel versions.)
- **TCP/IP software.** Each server requires the TCP/IP software.
- **SAP Software.** Each server must have the SAP software installed and configured before configuring LifeKeeper and the LifeKeeper SAP Recovery Kit. The same version should be installed on each server. Consult the [LifeKeeper for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information.
- **LifeKeeper software.** You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [LifeKeeper for Linux Release Notes](#) for specific LifeKeeper requirements.
- **LifeKeeper IP Recovery Kit.** This recovery kit is required if remote clients will be accessing the SAP PAS, ASCS or SCS instance. You must use the same version of this recovery kit on each server.
- **LifeKeeper for Linux NFS Server Recovery Kit.** This recovery kit is required for most configurations. You must use the same version of this recovery kit on each server.
- **LifeKeeper for Linux Network Attached Storage (NAS) Recovery Kit.** This recovery kit is required for some configurations. You must use the same version of this recovery kit on each

server.

- **LifeKeeper for Linux Database Recovery Kit.** The LifeKeeper recovery kit for the database being used with SAP must be installed on each database server. Please refer to the [LifeKeeper for Linux Release Notes](#) for information on supported databases. A LifeKeeper database hierarchy must be created for the SAP PAS, ASCS or SCS Instance prior to configuring SAP.

#### Important Notes:

- If running an SAP version prior to Version 7.3, please consult your SAP documentation and notes on how to download and install **SAPHOSTAGENT** (see [Important Note](#) in the Plan Your Configuration topic).
- Refer to the [LifeKeeper for Linux Installation Guide](#) for instructions on how to install or remove the Core software and the SAP Recovery Kit.
- The installation steps should be performed in the order recommended. The SAP installation will fail if LifeKeeper is installed first.
- For details on configuring each of the required LifeKeeper Recovery Kits, you should refer to the [documentation for each kit](#) (IP, NFS Server, NAS, and Database Recovery Kits).
- Please refer to SAP installation documentation for further installation requirements, such as swap space and memory requirements.

## 6.17.6. SAP Configuration Considerations

---

This section contains information to consider before starting to configure SAP and contains examples of typical SAP configurations. It also includes a step-by-step process for configuring and protecting SAP with LifeKeeper.

Also, refer to your [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring your LifeKeeper Core resources (for example, file system resources).

### Supported Configurations

There are many possible configurations of database and application servers in an SAP Highly Available (HA) environment. The specific steps involved in setting up SAP for LifeKeeper protection are different for each configuration, so it is important to recognize the configuration that most closely fits your requirements. Some supported configuration examples are:

[ABAP+Java Configuration](#) (ASCS and SCS) **Note:** An ERS resource created in LifeKeeper for Linux 9.4.0 or later will operate in its own independent hierarchy. See [ERS Resource Types in LifeKeeper](#) for more details.

[ABAP Only Configuration](#) (ASCS) **Note:** An ERS resource created in LifeKeeper for Linux 9.4.0 or later will operate in its own independent hierarchy. See [ERS Resource Types in LifeKeeper](#) for more details.

[Java Only Configuration](#) (SCS) **Note:** An ERS resource created in LifeKeeper for Linux 9.4.0 or later will operate in its own independent hierarchy. See [ERS Resource Types in LifeKeeper](#) for more details.

The configurations pictured in the above examples consist of two servers hosting the Central Services Instance(s) with an ERS Instance, Database Instance, Primary Application Server Instance and zero or more additional redundant Application Server Instances (AS). Although it is possible to configure SAP with no redundant Application Servers, this would require users to log in to the ASCS Instance or SCS Instance which is not recommended by SAP. The ASCS Instance, SCS Instance and Database servers have access to shared file storage for database and application files.

While Central Services do not use a lot of resources and can be switched over very fast, Databases have a significant impact on switchover speeds. For this reason, it is recommended that the Database Instances and Central Services Instances (ASCS and SCS) be protected through two distinct LifeKeeper hierarchies. They can be run on separate servers or on the same server.

### Configuration Notes

The following are technical notes related to configuring SAP to run in an HA environment. Please see subsequent topics for step-by-step instructions on protecting SAP with LifeKeeper.

[Directory Structure](#)

[Virtual Server Name](#)

[SAP Health Monitoring](#)

[SAP License](#)

[Automatic Switchback](#)

[Other Notes](#)

## 6.17.6.1. ABAP+Java Configuration (ASCS and SCS)

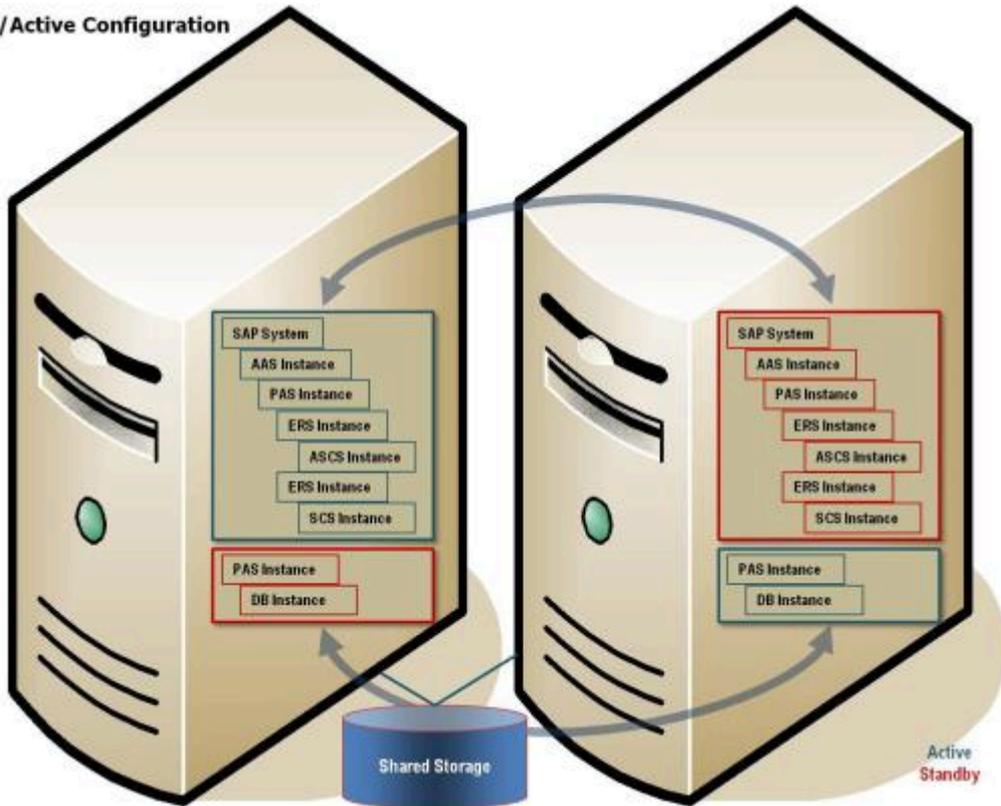
---

The ABAP+Java Configuration comprises the installation of:

- Central Services Instance for ABAP (ASCS Instance)
- Enqueue Replication Server Instance (ERS instance) for the ASCS Instance (optional) (Both the ASCS Instance and the SCS Instance must each have their own ERS instance)
- Central Services Instance for Java (SCS Instance)
- Enqueue Replication Server Instance (ERS Instance) for the SCS Instance (optional)
- Database Instance (DB Instance) – The ABAP stack and the Java stack use their own database schema in the same database
- Primary Application Server Instance (PAS)
- Additional Application Server Instances (AAS) – It is recommended that Additional Application Server Instances (AAS) be installed on hosts different from the Primary Application Server Instance (PAS) Host

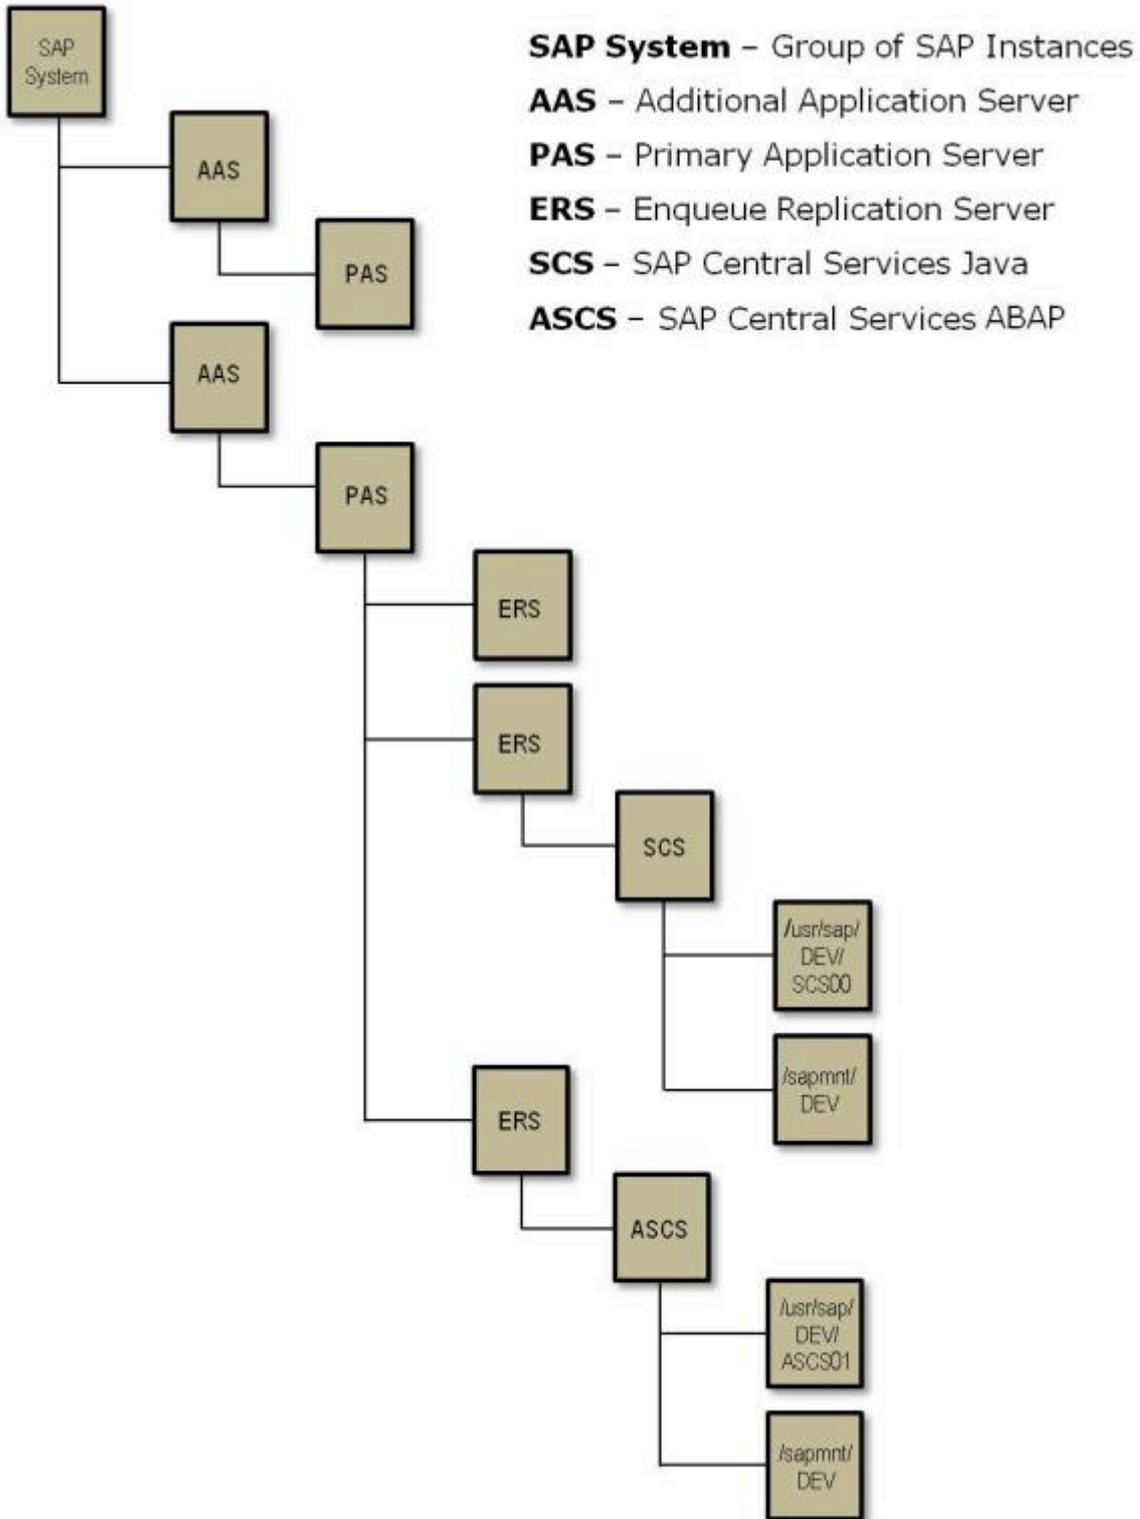
# Switchover Cluster for an SAP Dual-stack (ABAP+Java) System

Active/Active Configuration



In the above example, ASCS and SCS are in a separate LifeKeeper hierarchy from the Database and these Central Services Instances are active on a separate server from the Database.

## Example SAP Hierarchy



\* **Note:** An ERS resource created in LifeKeeper for Linux v9.4.0 or later will operate in its own independent hierarchy. See [ERS Resource Types in LifeKeeper](#) for more details.

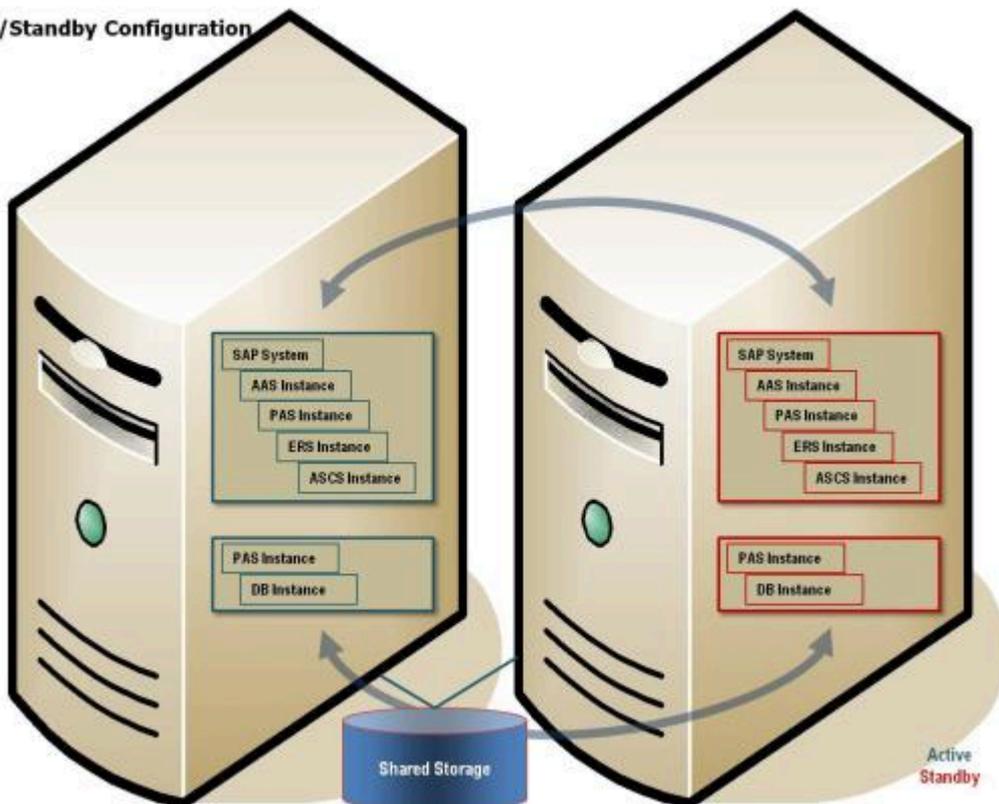
## 6.17.6.2. ABAP SCS (ASCS)

The ABAP Only Configuration comprises the installation of:

- Central Services Instance for ABAP (ASCS Instance)
- Enqueue Replication Server Instance (ERS instance) for the ASCS Instance (optional)
- Database Instance (DB Instance)
- Primary Application Server Instance (PAS)
- Additional Application Server Instances (AAS) – It is recommended that Additional Application Server Instances (AAS) be installed on hosts different from the Primary Application Server Instance (PAS) Host

### Switchover Cluster for an SAP ABAP Only (ASCS) System

Active/Standby Configuration



In the above example, ASCS is in a separate LifeKeeper hierarchy from the Database. Although it is active on the same server as the Database, they can fail over separately.

\* **Note:** An ERS resource created in LifeKeeper for Linux v9.4.0 or later will operate in its own independent hierarchy. See [ERS Resource Types in LifeKeeper](#) for more details.

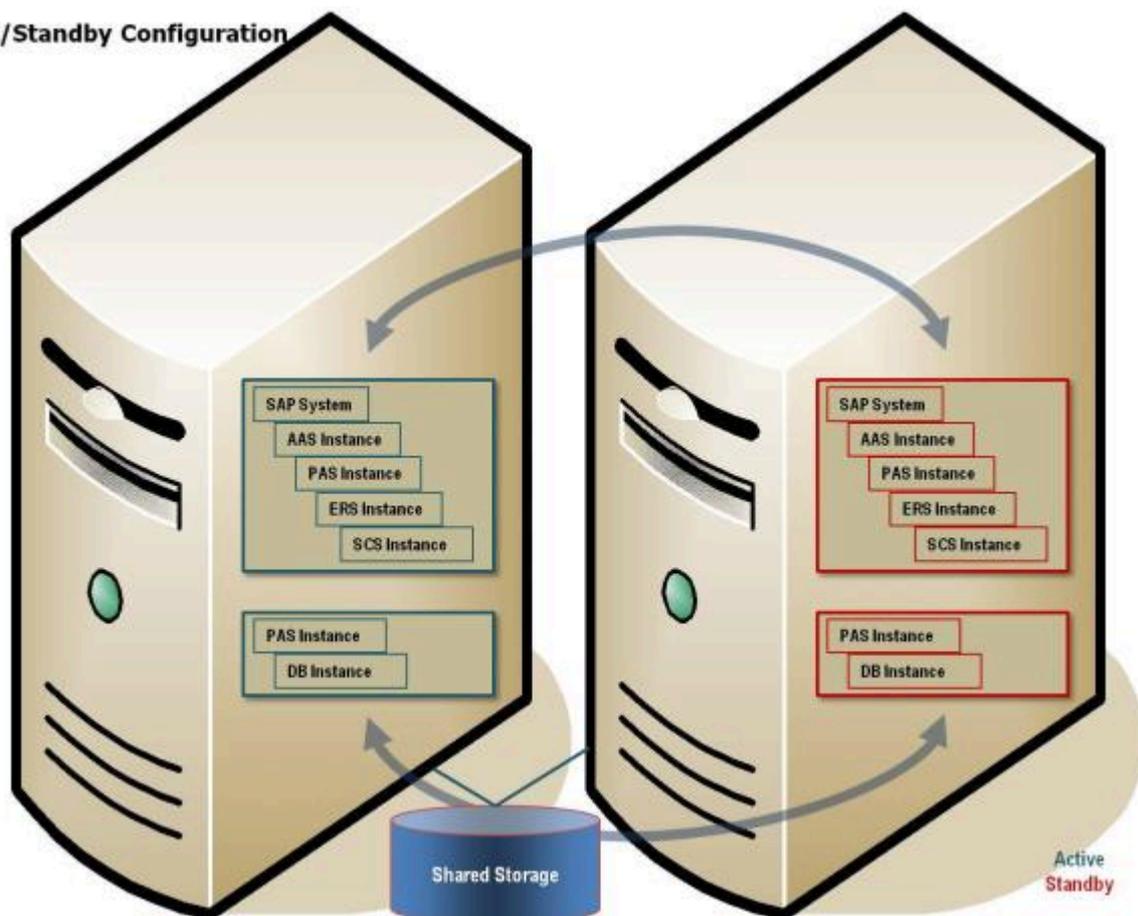
## 6.17.6.3. Java Only Configuration (SCS)

The Java Only Configuration comprises the installation of:

- Central Services Instance for Java (SCS Instance)
- Enqueue Replication Server Instance (ERS Instance) for the SCS Instance (optional)
- Database Instance (DB Instance)
- Primary Application Server Instance (PAS)
- Additional Application Server Instances (AAS) – It is recommended that Additional Application Server Instances (AAS) be installed on hosts different from the Primary Application Server Instance (PAS) Host

### Switchover Cluster for a Java Only System (SCS)

Active/Standby Configuration



In the above example, SCS is in a separate LifeKeeper hierarchy from the Database. Although it is active on the same server as the Database, they can fail over separately.

\* **Note:** An ERS resource created in LifeKeeper for Linux v9.4.0 or later will operate in its own independent hierarchy. See [ERS Resource Types in LifeKeeper](#) for more details.

## 6.17.6.4. SAP Directory Structure

---

The directory structure for the database will be different for each database management system that is used with the SAP system. Please consult the SAP installation guide specific to the database management system for details on the directory structure for the database. All database files must be located on shared disks to be protected by the LifeKeeper Recovery Kit for the database. Consult the database specific [Recovery Kit Documentation](#) for additional information on protecting the database.

See the [Directory Structure Diagram](#) below for a graphical depiction of the SAP directories described in this section.

The following types of directories are created during installation:

**Physically shared directories** (reside on global host and shared by NFS):

*/<sapmnt>/<SAPSID>* – Software and data for one SAP system (should be mounted for all hosts belonging to the same SAP system)

*/usr/sap/trans* – Global transport directory (has to have an export point)

**Logically shared directories** that are bound to a node such as */usr/sap* with the following local directories (reside on the local host with symbolic links to the global host):

*/usr/sap/<SAPSID>*

*/usr/sap/<SAPSID>/SYS*

*/usr/sap/hostctrl*

**Local directories** (reside on the local host and shared) that contain the SAP instances such as:

*/usr/sap/<SAPSID>/DVEBMGS<No.>* — Primary application server instance directory

*/usr/sap/<SAPSID>/D<No.>* — Additional application server instance directory

*/usr/sap/<SAPSID>/ASCS<No.>* — ABAP central services instance (ASCS) directory

*/usr/sap/<SAPSID>/SCS<No.>* — Java central services instance (SCS) directory

*/usr/sap/<SAPSID>/ERS<No.>* — Enqueue replication server instance (ERS) directory for the ASCS and SCS

The SAP directories: */sapmnt/<SAPSID>* and */usr/sap/trans* are mounted from NFS; however, SAP instance directories (*/usr/sap/<SAPSID>/<INSTTYPE><No.>*) should always be mounted on the cluster node currently running the instance. Do not mount such directories with NFS. The required directory structure depends on the chosen configuration. There are several issues that dictate the required directory structure.

## NFS Mount Points and Inodes

LifeKeeper maintains NFS share information using inodes; therefore, every NFS share is required to have a unique inode. Since every file system root directory has the same inode, NFS shares must be at least one directory level down from root in order to be protected by LifeKeeper. For example, referring to the information above, if the `/usr/sap/trans` directory is NFS shared on the SAP server, the `/trans` directory is created on the shared storage device which would require mounting the shared storage device as `/usr/sap`. It is not necessarily desirable, however, to place all files under `/usr/sap` on shared storage which would be required with this arrangement. To circumvent this problem, it is recommended that you create an `/exports` directory tree for mounting all shared file systems containing directories that are NFS shared and then create a soft link between the SAP directories and the `/exports` directories, or alternately, locally NFS mount the NFS shared directory. (**Note:** The name of the directory that we refer to as `/exports` can vary according to user preference; however, for simplicity, we will refer to this directory as `/exports` throughout this documentation.) For example, the following directories and links/mounts for our example on the SAP Primary Server would be:

For the <code>/usr/sap/trans</code> share	
Directory	Notes
<code>/trans</code>	created on shared file system and shared through NFS
<code>/exports/usr/sap</code>	mounted to <code>/</code> (on shared file system)
<code>/usr/sap/trans</code>	soft linked to <code>/exports/usr/sap/trans</code>

Likewise, the following directories and links for the `<sapmnt>/<SAPSID>` share would be:

For the <code>&lt;sapmnt&gt;/&lt;SAPSID&gt;</code> share	
Directory	Notes
<code>/&lt;SAPSID&gt;</code>	created on shared file system and shared through NFS
<code>/exports/sapmnt</code>	mounted to <code>/</code> (on shared file system)
<code>&lt;sapmnt&gt;/&lt;SAPSID&gt;</code>	NFS mounted to <code>&lt;virtual SAP server&gt;:/exports/sapmnt/&lt;SAPSID&gt;</code>

Detailed instructions are given for creating all directory structures and links in the configuration steps later in this documentation. See the [NFS Server Recovery Kit Documentation](#) for additional information on inode conflicts and for information on using the new features in NFSv4.

## Local NFS Mounts

The recommended directory structure for SAP in a LifeKeeper environment requires a locally mounted NFS share for one or more SAP system directories. If the NFS export point for any of the locally mounted NFS shares becomes unavailable, the system may hang while waiting for the export point to become available again. Many system operations will not work correctly, including a system reboot. You should be aware that the NFS server for the SAP cluster should be protected by LifeKeeper and should not be manually taken out of service while local mount points exist.

To avoid accidentally causing your cluster to hang by inadvertently stopping the NFS server, please follow the recommendations listed in the [NFS Considerations](#) topic.

When NFS shares are not accessible the unmount can fail. LifeKeeper will attempt to unmount the filesystem multiple times. These multiple attempts will typically succeed in eventually taking the resource out of service. However, this will cause delays in taking the resource out of service. To avoid these retries, use 'nfsvers=3, proto=udp' mount options.

**!** Note the usage of udp; this is important for failover and recovery.

## NFS Mounts and su

LifeKeeper accomplishes many database and SAP tasks by executing database and SAP operations using the `su - <admin name> -c <command>` command syntax. The `su` command, when called in this way, causes the login scripts in the administrator's home directory to be executed. These login scripts set environment variables to various SAP paths, some of which may reside on NFS mounted shares. If these NFS shares are not available for some reason, the `su` calls will hang, waiting for the NFS shares to become available again.

Since hung scripts can prevent LifeKeeper from functioning properly, it is desirable to configure your servers to account for this potential problem. The LifeKeeper scripts that handle SAP resource remove, restore and monitoring operations have a built-in timer that prevents these scripts from hanging indefinitely. No configuration actions are therefore required to handle NFS hangs for the SAP Application Recovery Kit.

Note that there are many manual operations that unavailable NFS shares will still affect. You should always ensure that all NFS shares are available prior to executing manual LifeKeeper operations.

**\*** To avoid any delay in handling inaccessible NFS shares please follow the recommendations listed in [NFS Considerations](#).

## Location of <INST> directories

Since the `/usr/sap/<SAPSID>` path is not NFS shared, it can be mounted to the root directory of the file system. The `/usr/sap/<SAPSID>` path contains the `SYS` subdirectory and an `<INST>` subdirectory for each SAP instance that can run on the server. For certain configurations, there may only be one `<INST>` directory, so it is acceptable for it to be located under `/usr/sap/<SAPSID>` on the shared file system. For other configurations, however, the backup server may also contain a local AS instance whose `<INST>` directory should not be on a shared file system since it will not always be available. To solve this problem, it is recommended that for certain configurations, the PAS's, ASCS's or SCS's `/usr/sap/<SAPSID>/<INST>`, `/usr/sap/<SAPSID>/<ASCS-INST>` or `/usr/sap/<SAPSID>/<SCS-INST>` directories should be mounted to the shared file system instead of `/usr/sap/<SAPSID>` and the `/usr/sap/<SAPSID>/SYS` and `/usr/sap/<SAPSID>/<AS-INST>` for the AS should be located on the local server.

 The ERS filesystem should be mounted locally in the case of ERSv1 and mounted on shared storage in the case of ERSv2.

For example, the following directories and mount points should be created for the ABAP+Java Configuration:

Directory	Notes
<code>/usr/ sap/&lt;SAPSID&gt;/DVEBMGS&lt;Instance #&gt;</code>	mounted to / (on shared file system)
<code>/usr/sap/&lt;SAPSID&gt;/SCS&lt;Instance #&gt;</code>	mounted to / (on shared file system)
<code>/usr/sap/&lt;SAPSID&gt;/ERS&lt;Instance #&gt; (for SCS instance)</code>	should be locally mounted on all cluster nodes or mounted from a NAS share ( <i>should not be mounted on shared file system</i> )
<code>/usr/sap/&lt;SAPSID&gt;/ASCS&lt;Instance #&gt;</code>	mounted to / (on shared file system)
<code>/usr/sap/&lt;SAPSID&gt;/ERS&lt;Instance #&gt; (for ASCS instance)</code>	should be locally mounted on all cluster nodes or mounted from a NAS share ( <i>should not be mounted on shared file system</i> )
<code>/usr/sap/&lt;SAPSID&gt;/AS&lt;Instance #&gt;</code>	created for AS on backup server

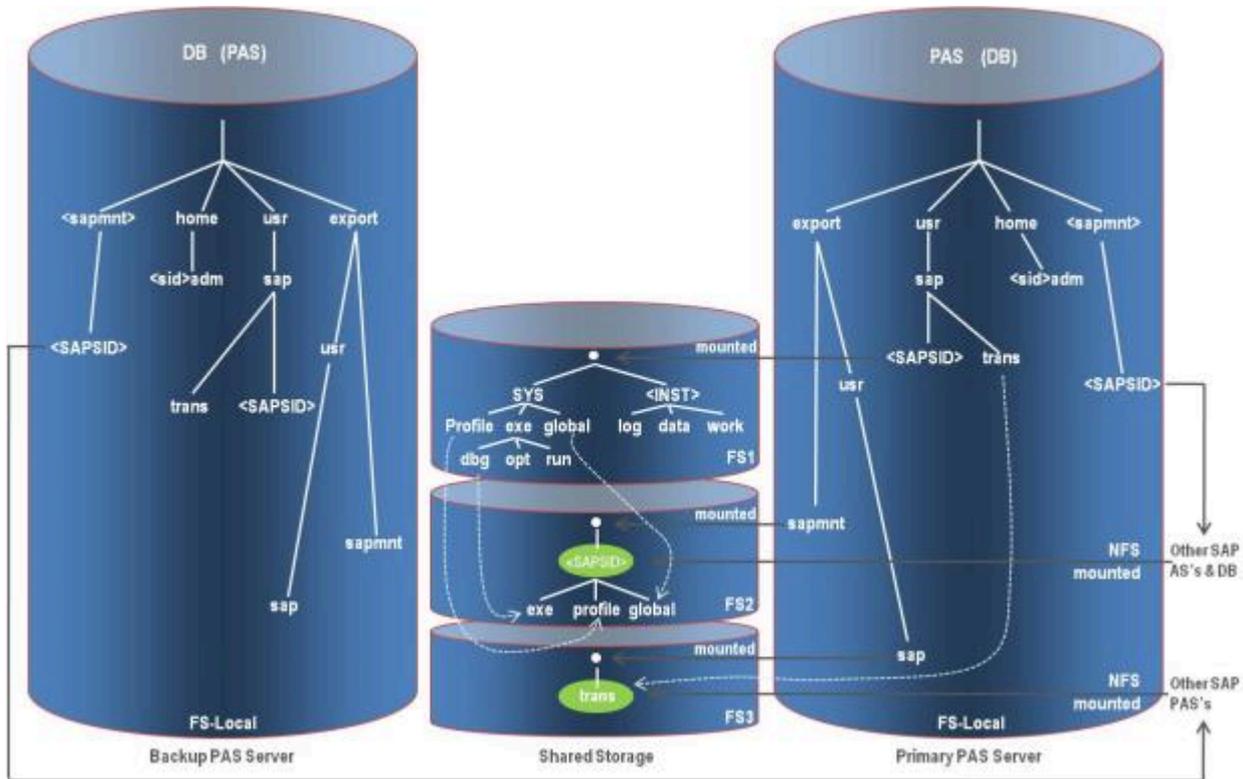
**Note:** The Enqueue Replication Server (ERS) resource will be in-service (ISP) on the primary node in your cluster. However, the architecture and function of the ERS requires that the actual processes for the instance run on the backup node. This allows the standby server to hold a complete copy of the lock table information for the primary server and primary enqueue server instance. When the primary server running the enqueue server fails, it will be restarted by LifeKeeper on the backup server on which the ERS process is currently running. The lock table (replication table) stored on the ERS is transferred to the enqueue server process being recovered and the new lock table is created from it. Once this process is complete, the active replication server is then deactivated (it closes the connection to the enqueue server and deletes the replication table). LifeKeeper will then restart the ERS processes on the new current backup node (formerly the primary) which has been inactive until now. Once the ERS process becomes active, it connects to the enqueue server and creates a replication table. For more information on the ERS process and SAP architecture features, visit <http://help.sap.com> and search for **Enqueue Replication Service**.

Since the replication server is always active on the backup node, it cannot reside on a LifeKeeper protected file system as the file system would be active on the primary node while the replication server process would be active on the backup node. Therefore, the file systems that ERS uses should be locally mounted on all cluster nodes or mounted from a NAS share.

# Directory Structure Diagram

The directory structure required for LifeKeeper protection of ABAP only environments is shown graphically in the figure below. See the [Abbreviations and Definitions](#) section for a description of the abbreviations used in the figure.

## Directory Structure Example



## Legend

-  exe soft link
-  mount point
-  NFS shared file system
- PAS Primary PAS
- (PAS) Backup PAS
- AS Application Server
- DB Primary DB Server
- (DB) Backup DB Server
- FSn File System on shared storage
- FS-local File System on local disk

## Directory Structure Options

The configuration steps presented in this documentation are based on the directory structure and diagrams described above. This is the recommended directory structure as tested and certified by SIOS Technology Corp.

There are other directory structure variations that should also work with the SAP Recovery Kit, although not all of them have been tested. For configurations with directory structure variations, you should follow the guidelines below.

- The `/usr/sap/trans` directory can be hosted on any server accessible on the network and does not have to be the PAS server. If you locate the `/usr/sap/trans` directory remotely from the PAS, you will need to decide whether access to this directory is mission critical. If it is, then you may want to protect it with LifeKeeper. This will require that it be hosted on a shared or replicated file system and protected by the [NFS Server Recovery Kit](#). If you have other methods of making the `/usr/sap/trans` directory available to all of the SAP instances without NFS, this is also acceptable.
- The `/usr/sap/trans` directory does not have to be NFS shared regardless of whether it is located on the PAS server.
- The `/usr/sap/trans` directory does not have to be on a shared file system if it is not NFS shared or protected by LifeKeeper.
- The directory structure and path names used to export NFS file systems shown in the diagrams are examples only. The path `/exports/usr/sap` could also be `/exports/sap` or just `/sap`.
- The `/usr/sap/<SAPSID>/<INST>` path needs to be on a shared file system at some level (except in the case of an ERSv1 instance). It does not matter which part of this path is the mount point for the file system. It could be `/usr`, `/usr/sap`, `/usr/sap/<SAPSID>` or `/usr/sap/<SAPSID>/<INST>`.
- The `/sapmnt/<SAPSID>` path needs to be on a shared file system at some level. The configuration diagrams show this path as NFS mounted, although this is an SAP requirement and not a LifeKeeper requirement.

## 6.17.6.5. SAP Virtual Server Name

---

SAP Application Servers and SAP clients communicate with the SAP Primary Application Server (PAS) using the name of the server where the PAS Instance is running. Likewise, the SAP PAS communicates with the Database (DB) using the name of the DB server. In a high availability (HA) environment, the PAS may be running on either the Primary Server or Backup Server at any given time. In order for other servers and clients to maintain a seamless connection to the PAS regardless of which server it is active on, a virtual server name is used for all communication with the PAS. This virtual server name is also mapped to a switchable IP address that can be active on whichever server the PAS is running on.

The switchable IP address is created and handled by LifeKeeper using the IP Recovery Kit. The virtual server name is configured manually by adding a virtual server name/switchable IP address mapping in DNS and/or in all of the servers' and clients' host files. See the [IP Local Recovery](#) topic for additional information on how this works.

Additionally, SAP configuration files must be modified so that the virtual server name is substituted for the physical server name. This is covered in detail in the [Installation](#) section where additional instructions are given for configuring SAP with LifeKeeper.

 **Note:** A separate switchable IP address is recommended for use with SAP Application Server hierarchies and the NFS Server hierarchies. This allows the IP address used for NFS clients to remain separate from the IP used for SAP clients.

## 6.17.6.6. SAP Health Monitoring

LifeKeeper monitors the health of the Primary Application Server (PAS) Instance and initiates a recovery operation if it determines that SAP is not functioning correctly.

The status will be returned to the user, via the GUI Properties Panel and CLI, as **Gray** (unknown/inactive/offline), **Red** (failed), **Yellow** (issue) or **Green** (healthy).

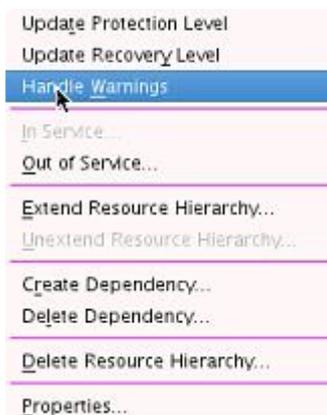
If the status of the instance is **Gray**, state is unknown; no information is available.

If the status of the instance is **Red**, the resource will be considered in a failed state and LifeKeeper will initiate the appropriate recovery handling operations.

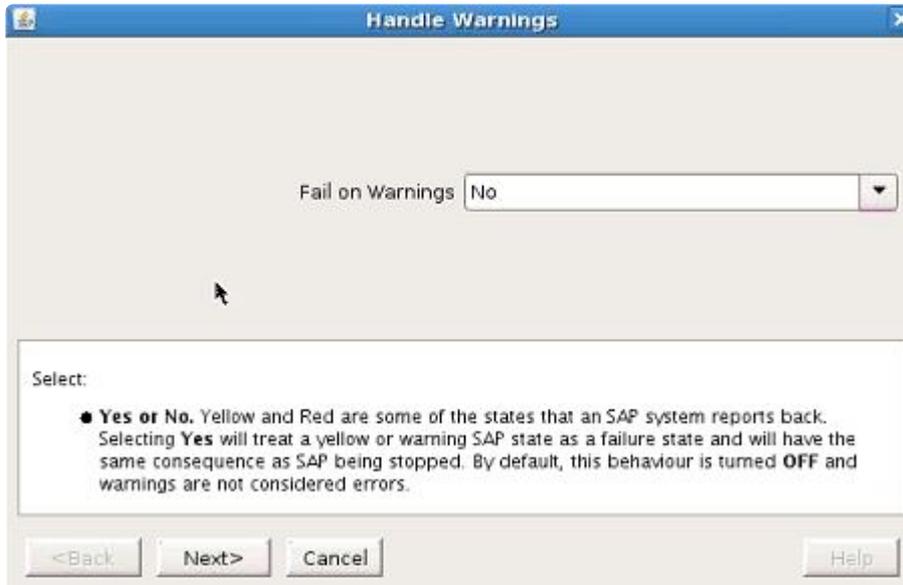
If the status of the instance is **Yellow**, it indicates that there may be an issue with the SAP processes for the defined Instance. The default behavior for a yellow status is to continue monitoring without initiating recovery.

This default behavior can be changed by configuring this option via the GUI resource menu.

1. Right-click the **Instance**.
2. Select **Handle Warnings**.



3. The following screen will appear, prompting you to select whether to **Fail on Warnings**.



Selecting **Yes** will cause a **Yellow Warning** to be treated as an error and will initiate recovery.

\* **Note:** It is highly recommended that this setting be left on the default selection of **No** as Yellow is a transient state that most often does not indicate a failure.

## 6.17.6.7. SAP License

---

In a high availability (HA) environment, SAP is configured to run on both a Primary and a Backup Server. Since the SAP licensing scheme is hardware dependent, a separate license is required for each server where SAP is configured to run. It will, therefore, be necessary to obtain and install an SAP license for both the Primary and Backup Servers.

## 6.17.6.8. SAP Automatic Switchback

---

In Active/Active configurations, the SAP Primary Application Server Instance (PAS), ABAP SAP Central Services Instance (ASCS) or SAP Central Services Instance (SCS) and Database (DB) hierarchies are separate and are in service on different servers during normal operation. There are times, however, when both hierarchies will be in service on the same server such as when one of the servers is being taken down for maintenance. If both hierarchies are in service on one of the servers and both servers go down, then when the servers come back up, it is important that the database hierarchy come in service before the SAP hierarchy in-service operation times out. Since LifeKeeper brings hierarchies in service during startup serially, if it chooses to bring SAP up first, the database in-service operation will wait on the SAP in-service operation to complete and the SAP in-service operation will wait on the database to become available, which will never happen because the DB restore operation can only begin after the PAS, ASCS or SCS restore completes. This deadlock condition will exist until the PAS, ASCS or SCS restore operation times out. (**Note:** SAP will time out and fail after 10 minutes.)

To prevent this deadlock scenario, it is important for this configuration to set the switchback flag for both hierarchies to **Automatic Switchback**. This will force LifeKeeper to restore each hierarchy on its highest priority server during LifeKeeper startup, which in this case is two different servers. Since LifeKeeper restore operations on different servers can occur in parallel, the deadlock condition is prevented.

## 6.17.6.9. Notes – Special Configuration Steps

---

The following items require special configuration steps in a high availability (HA) environment. Please consult the document [SAP Web Application Server in Switchover Environments](#) for additional information on configuration requirements for each:

- Login Groups
- SAP Spoolers
- Batch Jobs
- SAP Router
- SAP System Upgrades

## 6.17.7. SAP Installation

---

### Configuration/Installation

Before using LifeKeeper to create an SAP resource hierarchy, perform the following tasks **in the order recommended** below. Note that there are additional non-HA specific configuration tasks that must be performed that are not listed below. Consult the appropriate SAP installation guide for additional details.

The following tasks refer to the “**SAP Primary Server**” and “**SAP Backup Server.**” The **SAP Primary Server** is the server on which the Central Services will run during normal operation, and the **SAP Backup Server** is the server on which the Central Services will run if the SAP Primary Server fails.

Although it is not necessarily required, the steps below include the recommended procedure of protecting all shared file systems with LifeKeeper prior to using them. Prior to LifeKeeper protection, a shared file system is accessible from both servers and is susceptible to data corruption. Using LifeKeeper to protect the file systems preserves single server access to the data.

### Before Installing SAP

The tasks in the following topic are required before installing your SAP software. Perform these tasks in the order given. Please also refer to the SAP document *SAP Web Application Server in Switchover Environments* when planning your installation in NetWeaver Environments.

[Plan Your Configuration](#)

### Installing SAP Software

These tasks are required to install your SAP software for high availability. Perform the tasks below in the order given. Click on each task for details. Please refer to the appropriate SAP Installation Guide for further SAP installation instructions.

#### Primary Server Installation

[Install the Core Services, ABAP and Java Central Services](#)

[Install the Database](#)

[Install the Primary Application Server Instance](#)

[Install Additional Application Server Instances](#)

#### Backup Server Installation

[Install on the Backup Server](#)

# Installing LifeKeeper

[Install LifeKeeper](#)

[Create File Systems and Directory Structure](#)

[Move Data to Shared Disk and LifeKeeper](#)

[Upgrading From a Previous Version of the SAP Recovery Kit](#)

## Configuring SAP with LifeKeeper

### Resource Configuration Tasks

The following tasks explain how to configure your recovery kit by selecting certain options from the **Edit** menu of the LifeKeeper GUI. Each configuration task can also be selected from the toolbar or you may right-click on a global resource in the **Resource Hierarchy Tree** (left-hand pane) of the status display window to display the same drop down menu choices as the **Edit** menu. This, of course, is only an option when a hierarchy already exists.

Alternatively, right-click on a resource instance in the **Resource Hierarchy Table** (right-hand pane) of the status display window to perform all the configuration tasks, except creating a resource hierarchy, depending on the state of the server and the particular resource.

[IP Resources](#)

[Creating an SAP Resource Hierarchy](#)

[Deleting a Resource Hierarchy](#)

[Extending Your Hierarchy](#)

[Unextending Your Hierarchy](#)

[Common Recovery Kit Tasks](#)

[Setting Up SAP from the Command Line](#)

To enable the SAP SIOS HA Cluster Connector for an SAP instance, see [Activating the SAP SIOS HA Cluster Connector \(SSHCC\)](#).

For proper administration of the ERS instance in LifeKeeper, the ERS profile must use the `Start_Program` parameter instead of `Restart_Program` for starting the ERS process. See the [ASCS + ERS Restart\\_Program Parameter](#) page for details on how to modify this parameter in the ERS instance profile.

## Test the SAP Resource Hierarchy

You should thoroughly test the SAP hierarchy after establishing LifeKeeper protection for your SAP software. Perform the tasks in the order given.

[Test Preparation](#)

[Perform Tests](#)

## 6.17.7.1. Plan Your SAP Configuration

---

1. Determine which [configuration](#) you wish to use. The required tasks vary depending on the configuration.
2. Determine whether the SAP system-wide `/usr/sap/trans` directory will be hosted on the SAP Primary Application Server or on a file server. It can be hosted either place as long as it is NFS shared and fully accessible. If it is hosted on the SAP Primary Application Server and located on a shared file system, it should be protected by LifeKeeper and included in the SAP hierarchy.
3. Consider the storage requirements for SAP and DB as listed in the *SAP Installation Guide*. Most of the SAP files will have to be installed on shared storage. Consult the [LifeKeeper for Linux Technical Documentation](#) for the database-specific recovery kit for information on which database files are installed on shared storage and which are installed locally. Note that in an SAP environment, SAP requires local access to the database binaries, so they will have to be installed locally. Determine how to best use your shared storage to meet these requirements.

Also note that when shared storage resources are under LifeKeeper protection, they can only be accessed by one server at a time. If the shared device is a disk array, an entire LUN is protected. If a shared device is a disk, then the entire disk is protected. All file systems located on a single volume will therefore be controlled by LifeKeeper together. This means that you must have at least two logical volumes (LUNs), one for the database and one for SAP.

4. Virtual host names will be needed in order to identify your systems for failover. A new IP address is required for each virtual host name used. Make sure that the virtual host name can be correctly resolved in your Domain Name System (DNS) setup, then proceed as follows:

- a. Create the new virtual ip addresses by using the command:

```
ip addr add {IPADDRESS}/{NETMASK} dev eth0 (use the right netmask for your configuration)
```

**Note:** To verify these new virtual ip addresses use the `ip addr show` command.

- b. A separate virtual IP will also be needed for the ERS instance if using ERSv2.

In order to associate the switchable IP addresses with the virtual server name, edit `/etc/hosts` and add the new virtual ip addresses.

**Note:** This step is optional if the Primary Application Server and the Database are always running on the same server and communication between them is always local. But it is advisable to have separate switchable IP addresses and virtual server names for the Primary Application Server and the Database in case you ever want to run them on different servers and associated virtual hostnames.

5. Stop the caching daemon on both machines.

```
rcnscd stop
```

6. Mount the software.

```
mount //{path of software} (no password needed)
```

7. Run an X session (either an ssh -X or a VNC session — for Microsoft Windows users, Hummingbird Exceed X Windows can be used).

**Note:** When `sapinst` is run, the directory will be extracted under `/tmp`

8. When installing SAP (specifically ASCS or SCS) you need to specify the `SAPINST_USE_HOSTNAME={virtual hostname}` option. This is also required for the ERS instance if using ERSv2, but not if using ERSv1.

**Note:** Specify `sapinst SAPINST_USE_HOSTNAME={virtual hostname}` where **virtual hostname** is the hostname that resolves to the virtual IP that will float between the nodes.

## Important Note

The LifeKeeper SAP Recovery Kit relies on the SAP Host Agent being installed. If this software is not installed, then the LifeKeeper SAP Recovery Kit will not install. With SAP Netweaver version 7.3 and higher, this host agent is supplied; however, prior versions require a download from SAP. It is recommended that you consult your SAP help notes for your specific version. You can also refer to the Help Forum ([help.sap.com](http://help.sap.com)) for further documentation.

- The `saphostexec` module, either in RPM or SAR format, can be downloaded from SAP.
- To make sure that the modules are installed properly, there are a few modules to search for (`saposcol`, `saphostexec`, `saphostctrl`). These modules are typically found where SAP is installed (typically `/usr/sap` directory).

## 6.17.7.2. Installation of the Core Services

---

Before installing software, make sure that the date/time is synchronized on all servers. This is important for both LifeKeeper and SAP.

The Core Services, ABAP and Java Central Services (ASCS and SCS), are single points of failure (SPOFs) and therefore must be protected by LifeKeeper. Install these core services on the SAP Primary Server using the appropriate SAP Installation Guide.

### Installation Notes

- To be able to use the required virtual host names that were created in the [Plan Your Configuration](#) topic, set the `SAPinst` property `SAPINST_USE_HOSTNAME` to specify the required virtual host names before starting `SAPinst`. (**Note:** Document the `SAPINST_USE_HOSTNAME` virtual IP address as it will be used later during creation of the SAP resources in LifeKeeper.)

```
Run ./sapinst SAPINST_USE_HOSTNAME={hostname}
```

- In seven phases, the **Core Services** should be created and started. If permission errors occur on `jdbccconnect.jar`, go to `/sapmnt/STC/exe/uc/linuxx86_64` and make that directory as well as file `jdbccconnect.jar` writeable (`chmod 777 ---`).
- Installation completes with a success message.

## 6.17.7.3. Installation of the Database

---

1. Note the group id for dba and oinstall as this will be needed for the backup machine.
2. Change to the software directory and run the following:

```
./sapinst SAPINST_USE_HOSTNAME={database connectivity ip address}
```

3. Run SAPinst to install the Database Instance using the appropriate SAP Installation Guide.

### Installation Notes

- SIOS recommends removing the orarun package, if it is already installed, prior to installation of the Database Instance (see **SAP Note 1257556**).
- The database installation option in the SAPinst window assumes that the database software is already installed, except for Oracle. For Oracle databases, SAPinst stops the installation and prompts you to install the database software.
- The <DBSID> identifies the database instance. SAPinst prompts you for the <DBSID> when you are installing the database instance. The <DBSID> can be the same as the <SAPSID>.
- If you install a database on a host other than the SAP Global host, you must mount global directories from the SAP Global host.
- If you run into an issue where the Listener was started, kill it using the command

```
(ps -ef | grep lsnrctl)
```

- To reset passwords for SAPR3 and SAPR3DB users, use the command

```
brtools
```

After Database installation is complete, close the original dialog and continue with SAP installation, [Installing Application Services](#).

## 6.17.7.4. Installation of the Primary Application Server Instance

---

1. To install the Primary Application Server instance, rerun `sapinst` from the previously mentioned directory.

```
./sapinst SAPINST_USE_HOSTNAME=<vip>
```

2. When prompted, Select **Primary Application Server Instance** and continue with installation using the appropriate SAP Installation Guides.

### Installation Notes

- The Primary Application Server Instance does not need to be part of the cluster because it is no longer a single point of failure (SPOF). The SPOF is now in the central services instances (SCS instance and ASCS instance), which are protected by the cluster.
- The directory of the Primary Application Server Instance is called `DVEBMGS<No>`, where `<No>` is the instance number.
- Installation of application server is complete when the **OK** message is received.
- When installing replicated enqueue on 7.1, run `sapinst as-is`.

## 6.17.7.5. Installation of Additional Application Server Instances

---

It is recommended that Additional Application Server Instances be installed to create redundancy. Application Server Instances are not SPOFs, therefore, they do not need to be included in the cluster.

On every additional application server instance host, do the following:

1. Run `SAPinst` to install the Additional Application Server Instance.
2. When prompted, select **Additional Application Server Instance** and continue with installation using the appropriate SAP Installation Guide.

## **6.17.7.6. Installation on the Backup Server**

On the backup server, repeat the Installation procedures that were performed on the primary server:

1. [Install the Core Services, ABAP and Java Central Services](#)
2. [Install the Database](#)
3. [Install the Application Services](#)

## 6.17.7.7. Install LifeKeeper

---

On both the **Primary** and the **Backup** servers, LifeKeeper software will now be installed including the following recovery kits:

- SAP
- appropriate database (i.e. Oracle, SAP MaxDB)
- IP
- NFS
- NAS

1. Stop **Oracle Listener** and **SAP** on both machines.

For Example, if the Oracle user is *orastc*, the Oracle listener is *LISTENER\_STC*, and the SAP user is *stcadm*:

a. su to user *orastc* and run command `lsnrctl stop LISTENER_STC`

b. su to user *stcadm* and run command `stopsap sap{No.}`

c. From root user, make sure there are no SAP or Oracle user processes; if there are, enter `killall sapstartsrv`; even after this command, if there are processes, run `ps -ef` and kill each process

2. Go into */etc/hosts* on both machines and ensure that host and/or DNS entries are properly specified.
3. **Stop** and **remove** the IP addresses from the current interfaces. **Note:** This step is required before the IP addresses can be protected by the LifeKeeper IP Recovery Kit.

```
ip addr delete {A/SCS VIRTUAL IP ADDRESS}/{NETMASK} dev eth0
```

```
ip addr delete {ERSv2 VIRTUAL IP ADDRESS}/{NETMASK} dev eth0
```

4. Verify the IP addresses have been removed by performing a connection attempt, for example using ping.
5. Following the steps in the [LifeKeeper Installation Guide](#), install LifeKeeper on both the Primary server and the Backup server (DE, core, DataKeeper, LVM as well as the licenses). When prompted to select Recovery Kits, make sure you select the following:

SAP, appropriate database (i.e. Oracle), IP, NFS and NAS

The installation script does certain checking and might fail if the environment is not set up correctly as shown in this example:

```
SAP Services file /usr/sap/sapservices not found

SAP Installation is not valid; please check environment and retry

error: %pre(steeleye-lkSAP-7.3.1-1.noarch) scriptlet failed, exit
status 2

error: install: %pre scriptlet failed (2), skipping steeleye-
lkSAP-7.3.1-1
```

In the above example, the expected SAP file, `/usr/sap/sapservices`, is missing. It is very important for the environment to be in the right state before installation can continue.

Refer to the [Recovery Kit Documentation](#) for additional information about installing the recovery kits and configuring the servers for protecting resources.

## 6.17.7.8. Create File Systems and Directory Structure

---

While there are many different configurations depending upon which database management system is being used, below is the basic layout that should be adhered to.

\* Replicated (SIOS DataKeeper) file system resources must be created before creating the SAP resource.

- Set up comm paths between the primary and secondary servers
- Add virtual ip resources to *etc/hosts*
- Create virtual ip resources for instance and database hosts
- Set up shared disks
- Create file systems for SAP (located on shared disk)
- Create file systems for database (located on shared disk)
- Mount the main SAP file systems
- Create mount points
- Mount the PAS, ASCS/SCS, and ERS directories (if applicable) as well as any additional Application Servers

Please consult the SAP installation guide specific to the database management system for details on the directory structure for the database. All database files must be located on shared disks to be protected by the LifeKeeper Recovery Kit for the database. Consult the database specific [Recovery Kit Documentation](#) for additional information on protecting the database.

The following example is only a sample of the many configurations than can be established, but understanding these configurations and adhering to the configuration rules will help define and set up workable solutions for your computing environment.

1. From the UI of the primary server, set up comm paths between the primary server and the secondary server.
2. Add an entry for the actual primary and secondary virtual ip addresses in */etc/hosts*.
3. Log in to LifeKeeper on the primary server and create virtual ip resources for your host and your database (ex. *ip-db10* and *ip-sap10*).

4. Set up shared disks between the two machines.

**Note:** One lun for database and another for SAP data is recommended in order to enable independent failover.

5. For certain configurations, the following tasks may need to be completed:

- Create the physical devices
- Create the volume group
- Create the logical volumes for SAP
- Create the logical volumes for Database

6. Create the file systems on shared storage for SAP (these are *sapmnt*, *saptrans*, *ASCS{No}*, *SCS{No}*, *DVEBMGS{No}*). **Note:** SAP must be stopped in order to get everything on shared storage.

7. Create all file systems required for your database (Example: *mirrlogA*, *mirrlogB*, *origlogA*, *origlogB*, *sapdata1*, *sapdata2*, *sapdata3*, *sapdata4*, *oraarch*, *saparch*, *sapreorg*, *saptrace*, *oraflash* — `mkfs -t ext3 /dev/oracle/mirrlogA`).

**Note:** Consult the [LifeKeeper for Linux Technical Documentation](#) for the database-specific recovery kit and the *Component Installation Guide SAP Web Application Server* for additional information on which file systems need to be created and protected by LifeKeeper.

8. Create mount points for the main SAP file systems and then mount them (required). For additional information, see the [NFS Mount Points and Inodes](#) topic. (Note: */exports* directory was used to mount the file systems.)

```
mount /dev/sap/sapmnt /exports/sapmnt
```

```
mount /dev/sap/saptrans /exports/saptrans
```

9. Create temporary mount points using the following command.

```
mkdir /tmp/m{No}
```

10. Mount the three SAP directories (the following mount points are necessary for each Application Server present whether using external NFS or not).

```
mount /dev/sap/ASCS00 /tmp/m1
```

```
mount /dev/sap/SCS01 /tmp/m2
```

```
mount /dev/sap/DVEBMGS02 /tmp/m3
```

Proceed to [Moving Data to Shared Disk and LifeKeeper](#).

## 6.17.7.9. Move Data to Shared Disk and LifeKeeper

The following steps are an example using Oracle. **Note:** In this example STC is being used as the DB ID as well as the SAP SID. All occurrences of “STC” or “stc” in these commands (e.g., /usr/sap/STC or user orastc) should be replaced with the actual DB ID or SAP SID being used in the user’s cluster configuration.

**\* Before Beginning:** Primary and backup have been specified for the two servers. At the end of this procedure, the roles will be reversed. It is recommended that you first read through the steps, plan out which machine will be the desired primary and which will be the intended backup. At the end of this procedure, the role of primary and backup should become interchangeable. Understanding that in certain environments some machines are intended to be primary and some the backups, it is important to understand how this is structured.

1. Change directory to /usr/sap/STC, then change to each subdirectory and copy the data.

- cd ASCS{No.}
- cp -a \* /tmp/m1
- cd ../SCS{No.}
- cp -a \* /tmp/m2
- cd ../DVEBMGS{No.}
- cp -a \* /tmp/m3

2. Change the temporary directories to the correct user permission.

```
chown stcadm:sapsys /tmp/m1 (repeat for m2 and m3)
```

3. Unmount the three temp directories using `umount /tmp/m1` and repeat for m2 and m3.

4. Re-mount the device over the old directories.

```
mount /dev/sap/ASCS{No.} /usr/sap/STC/ASCS{No.}
```

```
mount /dev/sap/SCS{No.} /usr/sap/STC/SCS{No.}
```

```
mount /dev/sap/DVEBMGS{No.} /usr/sap/STC/DVEBMGS{No.}
```

5. Mount the thirteen temp directories for Oracle.

```
mount /dev/oracle/sapdata1 /tmp/m1
```

```
mount /dev/oracle/sapdata2 /tmp/m2
```

```
mount /dev/oracle/sapdata3 /tmp/m3
```

```
mount /dev/oracle/sapdata4 /tmp/m4
```

```
mount /dev/oracle/mirrlogA /tmp/m5
```

```
mount /dev/oracle/mirrlogB /tmp/m6
```

```
mount /dev/oracle/origlogA /tmp/m7
```

```
mount /dev/oracle/origlogB /tmp/m8
```

```
mount /dev/oracle/saparch /tmp/m9
```

```
mount /dev/oracle/sapreorg /tmp/m10
```

```
mount /dev/oracle/saptrace /tmp/m11
```

```
mount /dev/oracle/oraarch /tmp/m12
```

```
mount /dev/oracle/oraflash /tmp/m13
```

6. Change the directory to */oracle/STC* and copy the data.

a. Change to each subdirectory (`cd /tmp/m1`) and perform `cp -a * /oracle/STC`

7. Repeat this previous step for each subdirectory as shown in the relationship above.

8. Change the temporary directories to the correct user permission.

```
chown orastc:dba /tmp/m1 (repeat for m2 to m12)
```

9. Unmount all the temp directories.

```
umount /tmp/m*
```

10. Re-mount the device over the old directories.

```
mount /dev/oracle/sapdata1 /oracle/STC/sapdata1
```

11. Repeat the above for all the listed directories.

12. Edit the */etc/exports* file and insert the mount points for SAP's main directories.

```
/exports/sapmnt *(rw, sync, no_root_squash)
```

```
/exports/saptrans *(rw, sync, no_root_squash)
```

13. Start the NFS server using `systemctl start nfs-server.service`. If the NFS server is already active, you may need to do an “`exportfs -va`” to export those mount points.
14. Execute the following mount commands (**note the usage of udp; this is important for failover and recovery**).

```
mount { virtual ip }:/exports/sapmnt/< SID > /sapmnt/< SID > -o
rw, sync, bg, udp
```

```
mount { virtual ip }:/exports/saptrans /usr/sap/trans -o
rw, sync, bg, udp
```



**Note:** The example above uses the udp protocol. Please refer to the support matrix and release notes to insure that udp is supported for your OS version.

15. Log in to Oracle and start Oracle (after su to orastc).

```
lsnrctl start LISTENER_STC
```

```
sqlplus / as sysdba
```

```
startup
```

16. Log in to SAP and start SAP (after su to stcadm).

```
startsap sap{No.}
```

17. Make sure all processes have started.

```
ps -ef | grep en.sap (2 processes)
```

```
ps -ef | grep ms.sap (2 processes)
```

```
ps -ef | grep dw.sap (17 processes)
```

SAP MCC (Microsoft Management Console Snap-In for SAP) is an SAP-supplied Windows client that can be used to administer SAP instances. A corresponding version for Unix/Linux called SAP MC (SAP Management Console) is also available.

18. Stop SAP and the Oracle Listener. (**Note:** Note in Step c, we use the SQL\*Plus utility from Oracle to log in to Oracle and shut down the database.)

- a. su to stcadm and enter command `"stopsap sap{No.}"`
- b. su to orastc and enter command `"lsnrctl stop LISTENER_STC"`
- c. su to orastc and enter `"sqlplus sys as SYSDBA"` and enter `"shutdown"` at the command prompt
- d. enter command `"stopsap sap{No.}"`
- e. `killall sapstartsrv` as root
- f. kill any leftover processes still associated to the stcadm and orastc users (e.g., `ps -u stcadm` and `ps -u orastc`)

19. Unmount all the file systems.

```
umount /usr/sap/trans  
umount /sapmnt/STC  
umount /oracle/STC/*  
umount /usr/sap/STC/DVEBMGS{No.}  
umount /usr/sap/STC/SCS{No.}  
umount /usr/sap/STC/ASCS{No.}
```

20. Stop the NFS server using the `systemctl stop nfs-server.service` command and perform the unmounts.

```
umount /exports/sapmnt  
umount /exports/saptrans
```

21. Copy `/etc/exports` to the backup system.

```
scp /etc/exports (backup ip):/etc/exports
```

22. Deactivate the logical volumes on the primary.

```
lvchange -an oracle  
lvchange -an sap
```

23. Create the corresponding SAP directories on the backup system.

```
mkdir -p /exports/sapmnt
```

```
mkdir -p /exports/saptrans
```

24. Activate the logical volumes on the backup system.

```
lvchange -ay oracle
```

```
lvchange -ay sap
```

**Note:** Problems may occur on this step if any rearranging of storage occurred on the primary when the volume groups were built. A reboot of the backup will clear this up.

25. Mount the directories on the backup machine.

```
mount /dev/sap/sapmnt /exports/sapmnt
```

```
mount /dev/sap/saptrans /export/saptrans
```

```
mount /dev/sap/ASCS00 /usr/sap/STC/ASCS{No.}
```

```
mount /dev/sap/SCS01 /usr/sap/STC/SCS{No.}
```

```
mount /dev/sap/DVEBMGS02 /usr/sap/STC/DVEBMGS{No.}
```

```
mount /dev/oracle/sapdata1 /oracle/STC/sapdata1
```

```
mount /dev/oracle/sapdata2 /oracle/STC/sapdata2
```

```
mount /dev/oracle/sapdata3 /oracle/STC/sapdata3
```

```
mount /dev/oracle/sapdata4 /oracle/STC/sapdata4
```

```
mount /dev/oracle/origlogA /oracle/STC/origlogA
```

```
mount /dev/oracle/origlogB /oracle/STC/origlogB
```

```
mount /dev/oracle/mirrlogA /oracle/STC/mirrlogA
```

```
mount /dev/oracle/mirrlogB /oracle/STC/mirrlogB
```

```
mount /dev/oracle/oraarch /oracle/STC/oraarch
```

```
mount /dev/oracle/saparch /oracle/STC/saparch
```

```
mount /dev/oracle/saptrace /oracle/STC/saptrace
```

```
mount /dev/oracle/sapreorg /oracle/STC/sapreorg
```

26. Switch over the IP addresses to the backup system via LifeKeeper.

## 27. Mount the NFS exports on the backup

```
mount sap{No.}:/exports/sapmnt/STC /sapmnt/STC
```

```
mount sap{No.}:/exports/saptrans/trans /usr/sap/trans
```

## 28. Log in to Oracle and start Oracle (after su to orastc).

```
lsnrctl start LISTENER_STC
```

```
sqlplus / as sysdba
```

```
startup
```

## 29. Log in to SAP and start SAP (after su to stcadm).

```
startsap sap{No.}
```

## 30. Log in to LifeKeeper and switch primary and backup priority instances (make backup higher priority).

## 31. On the original primary, save the original directories as such:

```
mv /exports /exports-save
```

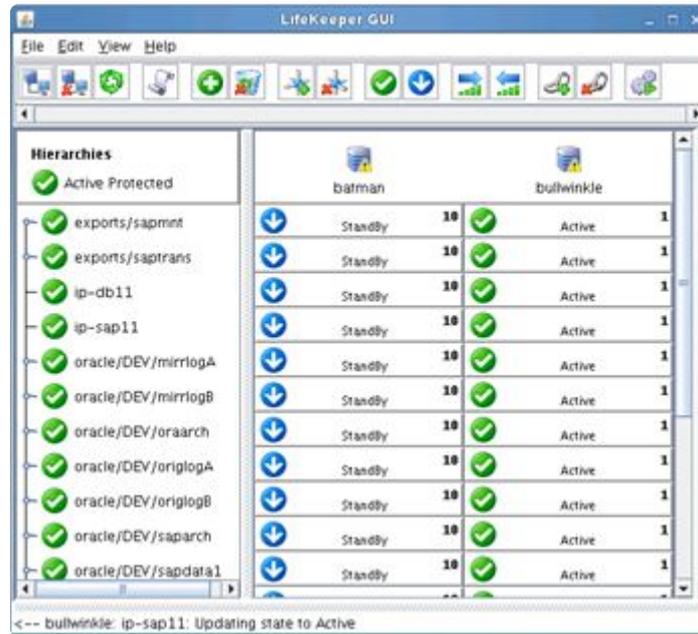
```
mv /usr/sap/STC/DVEBMGS{No.} /usr/sap/STC/DVEBMGS{No.}-save (repeat  
for SCS{No.} and ASCS{No.})
```

```
mv /oracle/STC/sapdata1 /oracle/STC/sapdata1-save (repeat for  
sapdata2, sapdata3, sapdata4, mirrlogA, mirrlogB, origlogA,  
origlogB, sapreorg, saptrace, saparch, oraarch)
```

## 32. Create "file system" resources, all the 17 mount points (5 for SAP and 12 for Oracle) one by one.

## 33. Extend to the original primary.

LifeKeeper resource hierarchy and SAP cluster are set up. (**Note:** This is a screen shot from the DEV instance.)



## 6.17.7.10. Modify ASCS and ERS Instance Profile Settings

Change the Restart\_Program parameter to Start\_Program for the enqueue server and enqueue replicator processes (in the ASCS and ERS instance profiles, respectively) to prevent the sapstart utility from automatically restarting them. Also change the entry 'Autostart = 1' to 'Autostart = 0' in both instance profiles to prevent each instance from being automatically restarted when the system reboots.

Before performing the following steps, identify the locations of the active ASCS and ERS instance profiles. This can be found from the /usr/sap/sapservices file by looking at the path of the file provided to sapstartsrv as the 'pf=' parameter:

**Note:** The following is only an example, where we are using ASCS instance number, 00 and ERS on 10, change the numbers accordingly for your environment. Also replace the <SID> to be the actual System ID.

```
LD_LIBRARY_PATH=/usr/sap/SID/ASCS00/exe:$LD_LIBRARY_PATH; export
```

```
LD_LIBRARY_PATH; /usr/sap/SID/ASCS00/exe/sapstartsrv
```

```
pf= /usr/sap/SID/SYS/profile/SID_ASCS00_sap1 -D -u SIDadm
```

```
LD_LIBRARY_PATH=/usr/sap/SID/ERS10/exe:$LD_LIBRARY_PATH; export
```

```
LD_LIBRARY_PATH; /usr/sap/SID/ERS10/exe/sapstartsrv
```

```
pf= /usr/sap/SID/SYS/profile/SID_ERS10_sap2 -D -u SIDadm
```

In this example, the active ASCS instance profile is located at /usr/sap/SID/SYS/profile/SID\_ASCS00\_sap1 and the active ERS instance profile is located at /usr/sap/SID/SYS/profile/SID\_ERS10\_sap2.

### Steps

1. Edit the ASCS instance profile as follows:
  - a. Change 'Autostart = 1' to 'Autostart = 0', or manually add the line 'Autostart = 0'.
  - b. The exact format of the line in the profile that starts the enqueue server process will vary depending on whether version 1 or 2 of the Standalone Enqueue Server Framework is being used. Modify the profile according to the following table:

If the following line appears...	Change it to...
Restart_Program_01 = local \$(_EN) pf=\$(_PF)	Start_Program_01 = local \$(_EN) pf=\$(_PF)

Restart_Program_01 = local \$_ES2 pf=\$_PF)	Start_Program_01 = local \$_ES2 pf=\$_PF)
Restart_Program_01 = local \$_ENQ) pf=\$_PF)	Start_Program_01 = local \$_ENQ) pf=\$_PF)

 **Note:** The numbers xx in the Start\_Program\_xx or Restart\_Program\_xx entries may be different on your system. They do not need to be changed to match the lines shown above.

2. Verify that all entries are correct and save the changes to the ASCS instance profile.
3. Edit the ERS instance profile as follows:
  - a. Change 'Autostart = 1' to 'Autostart = 0', or manually add the line 'Autostart = 0'.
  - b. The exact format of the line in the profile that starts the enqueue replicator process will vary depending on whether version 1 or 2 of the Standalone Enqueue Server Framework is being used. Modify the profile according to the following table:

If the following line appears...	Change it to...
Restart_Program_00 = local \$_ER) pf=\$_PFL) NR=\$(SCSID)	Start_Program_00 = local \$_ER) pf=\$_PFL) NR=\$(SCSID)
Restart_Program_00 = local \$_ER2) pf=\$_PF)	Start_Program_00 = local \$_ER2) pf=\$_PF)
Restart_Program_00 = local \$_ENQR) pf=\$_PF)	Start_Program_00 = local \$_ENQR) pf=\$_PF)

 **Note:** The numbers xx in the Start\_Program\_xx or Restart\_Program\_xx entries may be different on your system. They do not need to be changed to match the lines shown above.

5. Verify that all entries are correct and save the changes to the ERS instance profile.
6. SIOS recommends restarting the system to ensure the updated profile is read and no caching is in effect.
7. SIOS also recommends to always verify changes in a test environment before applying to production workload.

 Changing autostart from one value to another requires a reboot in order for the parameter change to take effect.

**Notes:**

1. The entries 'Autostart = 0' and 'Start\_Program\_xx ...' must be on separate lines in each instance

profile.

2. The numbers xx in the Start\_Program\_xx or Restart\_Program\_xx entries may be different on your system. They do not need to be changed to match the numbers given in this solution.
3. In a Java-based or dual stack Java+ABAP deployment there will be an SCS central services instance. In this case, steps 1-3 also need to be performed for the SCS instance.
4. See SAP Note 768727 (Automatic restart functions in sapstart for processes) for more details on the differences between Start\_Program and Restart\_Program.

## 6.17.7.11. Upgrading from a Previous Version of the SAP Recovery Kit

---

To upgrade from a previous version of the SAP Recovery Kit, perform the following steps.

1. Prior to upgrading, please review the [Plan Your Configuration](#) topic to make sure you understand all the implications of the new software.

**Note:** If running a version prior to SAP Netweaver 7.3, the SAPHOST agent will need to be installed. See the [Important Note](#) in the Plan Your Configuration topic for more information.

It is recommended that you take a snapshot of your current hierarchy using the `lkbackup` utility.

2. Follow the instructions in the “[Upgrading LifeKeeper](#)” topic in the [SPS for Linux Installation Guide](#).

A backup will be performed of the existing hierarchy. The upgrade will then destroy the old hierarchy and recreate the new hierarchy. If there is a failure, see In Case of Failure below.

3. At the end of the upgrade, stop and restart the LifeKeeper GUI in order to load the updated GUI client.

The LifeKeeper GUI server caches pages, so a restart is needed for it to refresh the new pages. As root, enter the command “`lkGUIserver restart`” which should stop and restart the GUI server. Exit all clients before attempting such a restart.

**Note:** Restarting your entire LifeKeeper system is not necessary, but it would be advisable in a production setting to schedule some down time and go through an orderly system preparation time even though testing has not required a system recycle.

4. Log in to the LifeKeeper UI, note the hierarchy and make sure the hierarchy is correct.

### In Case of Failure

It is possible to retry the upgrade. The upgrade script is kept intact in `/tmp` directory (`lkcreatesaptmp`). This is a temporary file that is used during the upgrade. The commands are written here and can be executed to create the hierarchy.

If there is a failure, an error or you suspect the hierarchy is not correct, the following steps are recommended:

1. Stop LifeKeeper by running `$LKROOT/bin/lkcli stop`.
2. Remove the new rpm “`rpm -e steeleye-lkSAP`”.
3. Install the old rpm “`rpm -i steeleye-lkSAP-<previous version>.noarch.rpm`”.

4. Restore the old hierarchy using `lkbbackup -x`
5. Restart LifeKeeper.
6. [Contact SIOS Support](#) for help. Prior to contacting Support, please have on hand the logs, the previous snapshot of the hierarchy, the hierarchy that was created and failed and the error messages received during the upgrade.

## 6.17.7.12. SAP IP Resources

---

Before continuing to set up the LifeKeeper hierarchy, determine the IP address that the SAP resource will use for failover or switchover. This is typically the virtual IP address used during the installation of SAP using the parameter `SAPINST_USE_HOSTNAME`. This IP address is a virtual IP address that is shared between the nodes in a cluster and will be active on one node at a time. This IP address will be different than the IP address used to protect the database hierarchy. Please note these IP addresses so they can be utilized when creating the SAP resources.

## 6.17.7.13. Creating an SAP Resource Hierarchy

---



<https://fast.wistia.net/embed/iframe/3j1g7ubr4a>

To protect the SAP System, an SAP Hierarchy will be needed. This SAP Hierarchy consists of the Core (Central Services) Resource, the ERS Resource, the Primary Resource and Secondary Resources. To create this hierarchy, perform the following tasks from the Primary Server.

**Note:** The below example is meant to be a guideline for creating your hierarchy. Tasks will vary somewhat depending upon your configuration.

### Create the Core Resource

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

A dialog box will appear with a drop-down list box with all recognized recovery kits installed within the cluster. Select **SAP** from the drop-down listing.

Please Select Recovery Kit

Click **Next**.

When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful should you encounter an error that might require you to correct previously entered information.

If you click **Cancel** at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the SAP instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either **intelligent** or **automatic**. **Intelligent switchback** requires administrative intervention to switch the instance back to the primary/original server. **Automatic switchback** means the switchback will occur as soon as the primary server comes back on line and re-establishes LifeKeeper communication paths.



A screenshot of a dropdown menu labeled "Switchback Type". The menu is open, showing the selected option "intelligent". The dropdown arrow is visible on the right side of the menu.

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box.

Click **Next**.

3. Select the Server where you want to place the SAP PAS, ASCS or SCS (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop-down list box.



A screenshot of a dropdown menu labeled "Server". The menu is open, showing the selected option "ip-12-0-0-20". The dropdown arrow is visible on the right side of the menu.

Click **Next**.

4. Select the **SAP SID**. This is the system identifier of the SAP PAS, ASCS or SCS system being protected.



A screenshot of a dropdown menu labeled "SAP SID". The menu is open, showing the selected option "EXM". The dropdown arrow is visible on the right side of the menu.

Click **Next**.

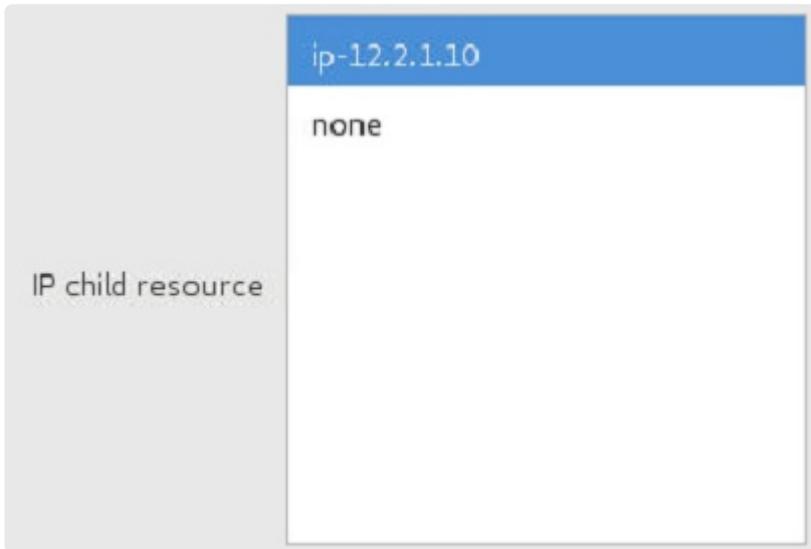
5. Select the SAP Instance Name (ex. ASCS<No.>) (Core Instance first) for the SID being protected.



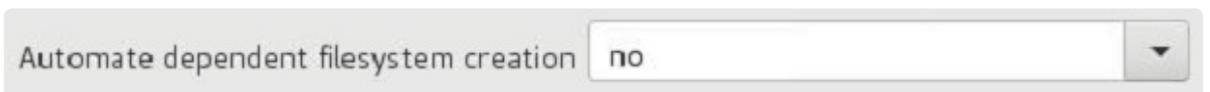
A screenshot of a dropdown menu labeled "SAP Instance for EXM". The menu is open, showing the selected option "ASCS02". The dropdown arrow is visible on the right side of the menu.

**Note:** Additional screens may appear related to customization of Protection and Recovery Levels.

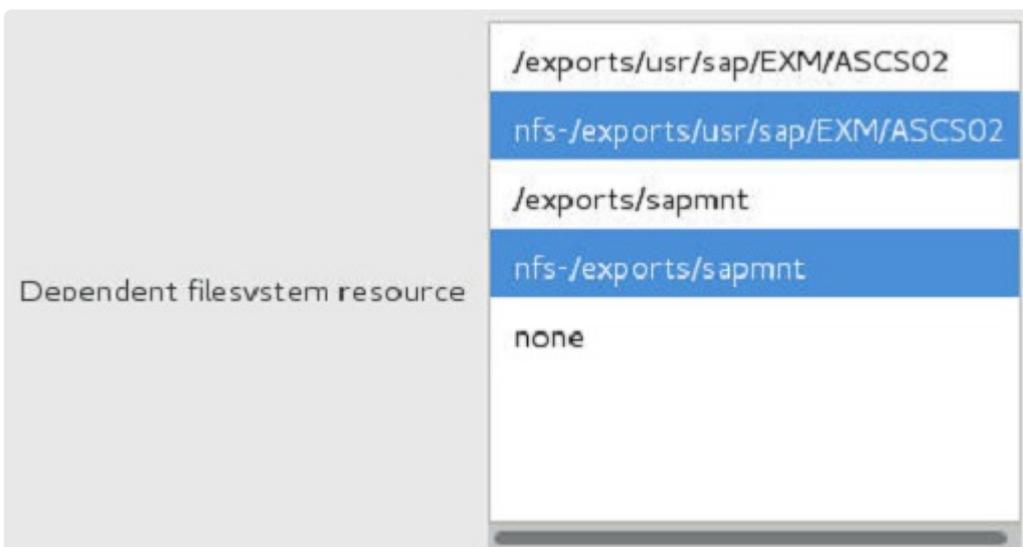
6. Select the **IP Child Resource**. This is typically either the Virtual Host IP address noted during SAP installation (SAPINST\_USE\_HOSTNAME) or the IP address needed for failover.



a. Select whether LifeKeeper should attempt to **automate creation of dependent filesystems** for the instance. If **yes** is selected, LifeKeeper will attempt to create the necessary filesystem resources and add them as dependencies under the SAP resource. (Note that replicated filesystems cannot be created automatically.) If **no** is selected, the following dialog box will prompt the user to select existing LifeKeeper filesystem resources to be added as dependencies.



b. If **no** was chosen in the previous dialog, the option will be provided to select the filesystem resource(s) which should be added as a dependency under the SAP resource. Multiple resources can be selected by holding CTRL and clicking each resource separately. **Note:** Filesystem resources must be in-service (ISP) on this server in LifeKeeper in order to appear as choices in this dialog.



7. Select or enter the **SAP Tag**. This is a tag name that LifeKeeper gives to the SAP resource. You can select the default or enter your own tag name. The default tag is *SAP-<SID>\_<INST>*.

SAP Tag

When you click **Create**, the **Create SAP Resource Wizard** will create your SAP resource.

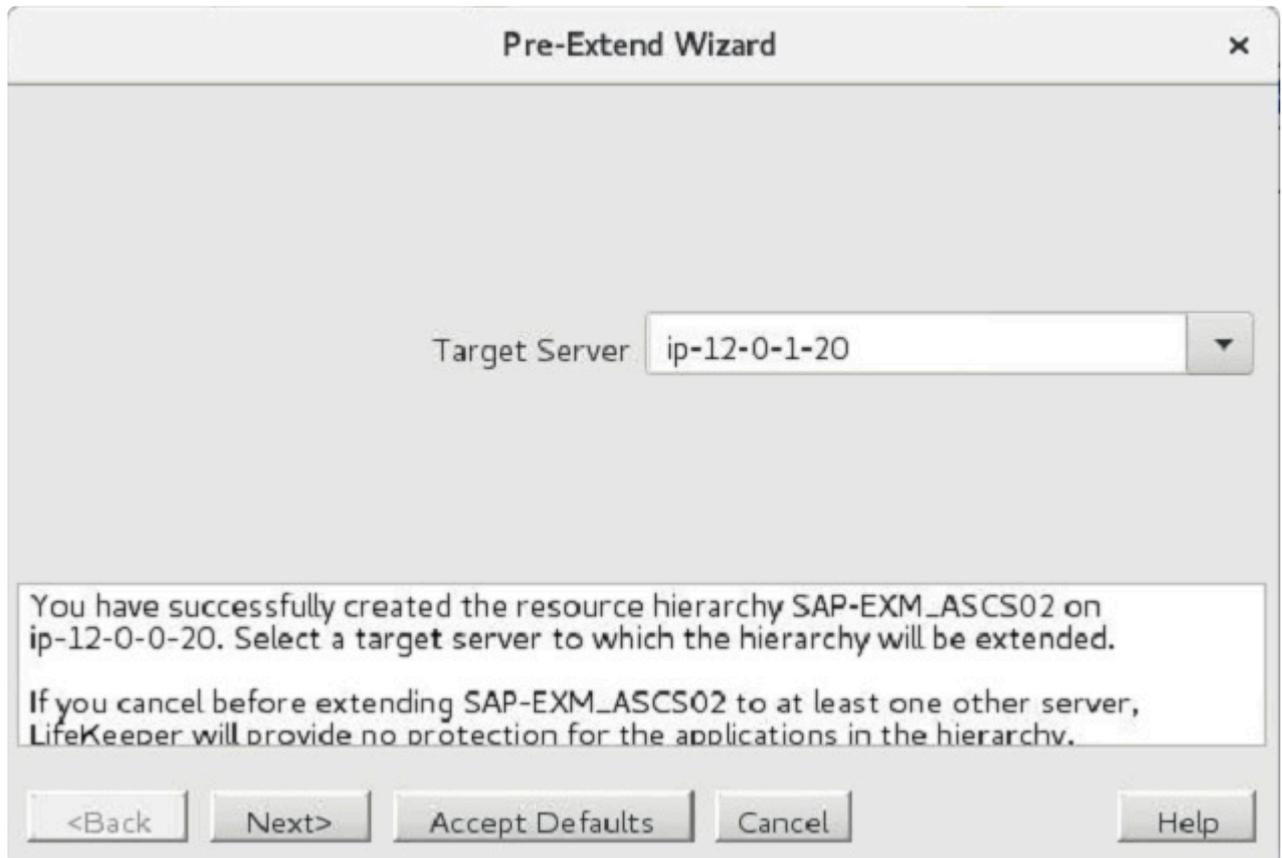
- At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your SAP resource hierarchy. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created. There may also be errors or messages output from the SAP startup scripts that are displayed in the information box.



Click **Next**

- Another information box will appear explaining that you have successfully created an SAP resource hierarchy, and you must **Extend** that hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

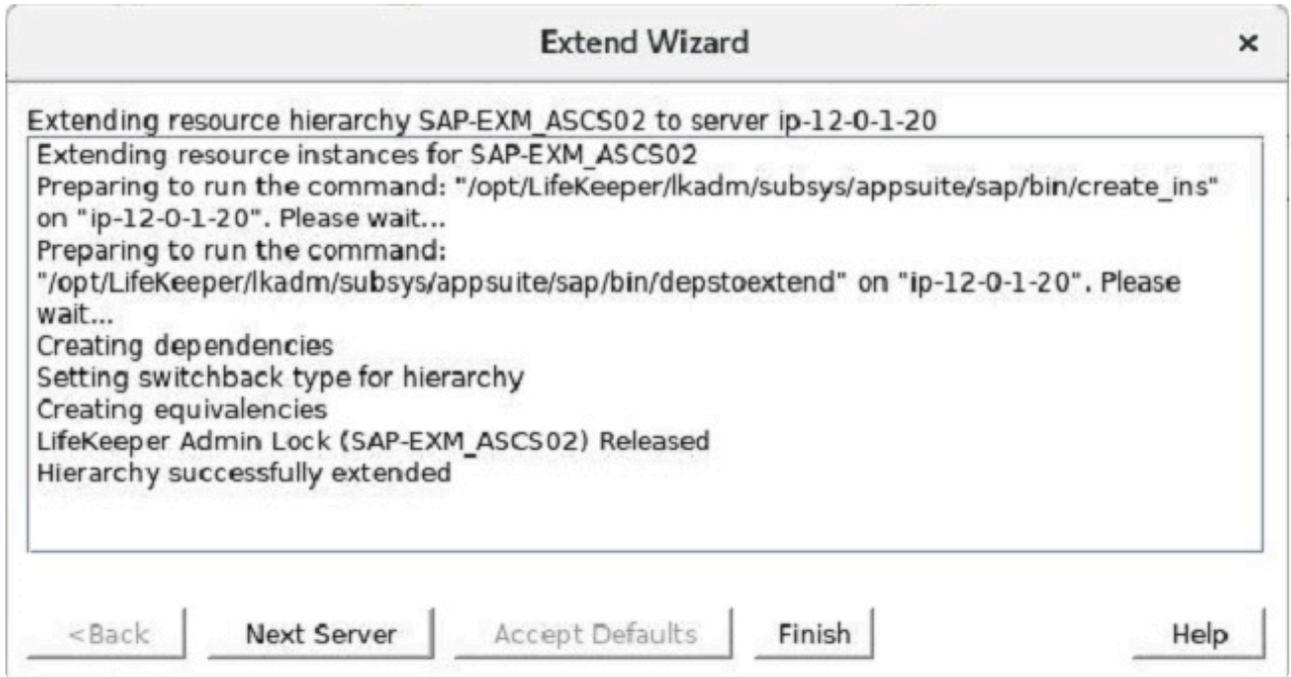
When you click **Next**, LifeKeeper will launch the **Pre-Extend Wizard** that is explained later in this section.



If you click **Cancel** now, a dialog box will appear warning you that you will need to come back and extend your SAP resource hierarchy to another server at some other time to put it under LifeKeeper protection.



10. The **Extend Wizard** dialog will appear stating **Hierarchy successfully extended**. Click **Finish**.



11. The **Hierarchy Integrity Verification** dialog appears. Once Hierarchy Verification finishes, click **Done** to exit the **Create Resource Hierarchy** menu selection.

**Hierarchy with the Core as the Top Level**



## Create the ERS Resource

The ERS resource provides additional protection against a single point of failure of a Core Instance (Central Services Instance) or enqueue server process. When a Core Instance (Central Services Instance) fails and is restarted, it will retrieve a backup copy of the enqueue lock table (i.e., the *replication table*) from the enqueue replication server. The result is that, in the event of the enqueue server failure, no transactions or updates are lost and the service for the SAP system continues.

For a discussion of the differences in the implementation of ERS resources in LifeKeeper for Linux in versions 9.4.0 and later versus the implementation prior to version 9.4.0, see [ERS Resource Types in LifeKeeper](#).

**! Important Note:** The creation and extension of the ERS hierarchy **must occur on a server where the corresponding ASCS/SCS instance is not in-service (ISP) in LifeKeeper**. In the case of ERSv2, the underlying ERS filesystem resources must first be brought in-service on the backup server where the resource will be created.

Perform the following steps to create the ERS Resource.

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**. A dialog box will appear with a drop-down list box with all recognized recovery kits installed within the cluster. Select **SAP** from the drop-down listing. Click **Next**.



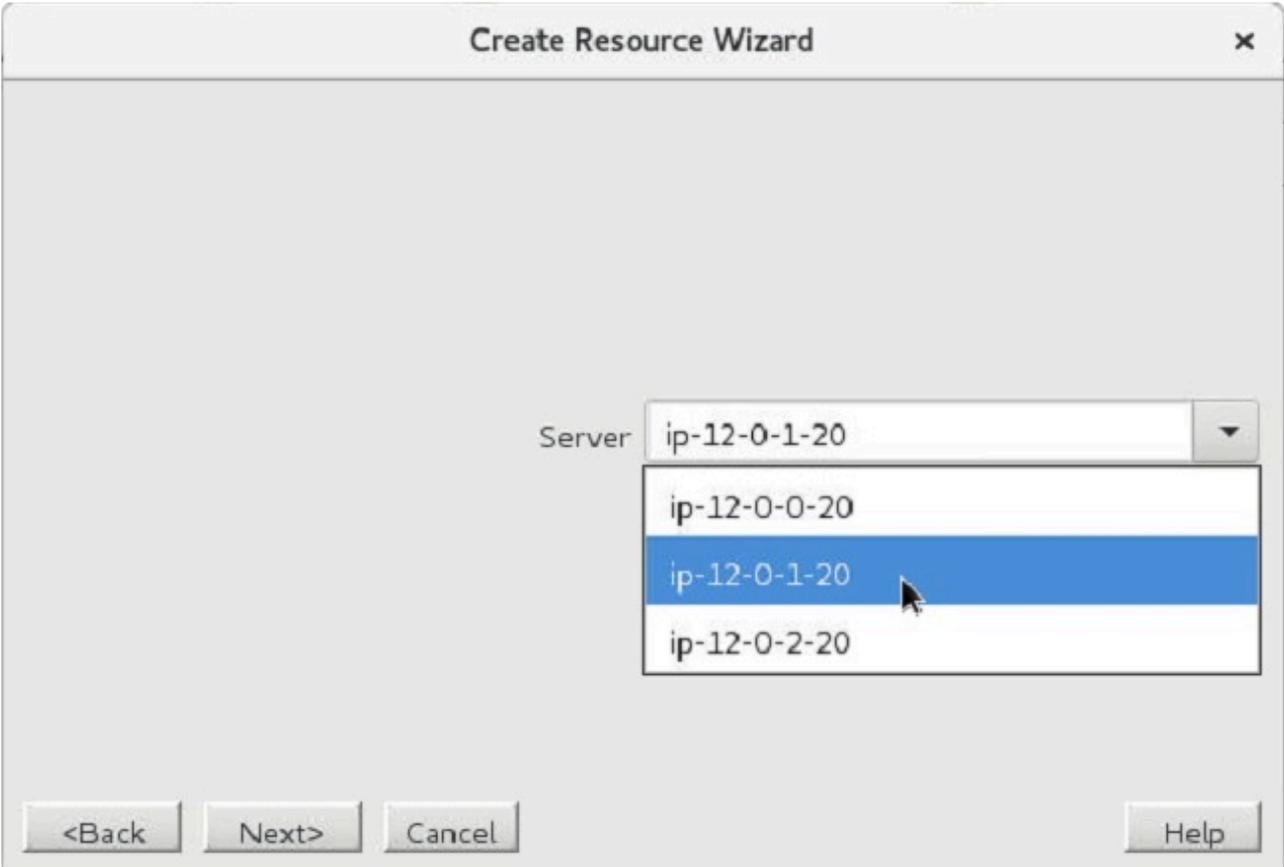
Please Select Recovery Kit

2. Select the **Switchback Type**. Click **Next**.



Switchback Type

3. **Important:** Select a **Server** in the cluster where the corresponding ASCS/SCS instance is **not ISP**. Click **Next**.



**Create Resource Wizard**

Server

- ip-12-0-0-20
- ip-12-0-1-20**
- ip-12-0-2-20

<Back   Next>   Cancel   Help

- Select the **SAP SID** for the ERS instance. Click **Next**.

A screenshot of a dropdown menu. The label 'SAP SID' is on the left. The selected value is 'EXM'. A downward arrow is on the right side of the dropdown box.

- Select the **SAP Instance Name** (ex. ERS<No.>). Click **Next**.

A screenshot of a dropdown menu. The label 'SAP Instance for EXM' is on the left. The selected value is 'ERS12'. A downward arrow is on the right side of the dropdown box.

- If creating a resource to represent an ERSv2 instance, select the **IP Child Resource**. This choice will not appear when creating a resource to represent an ERSv1 instance.

A screenshot of a dropdown menu titled 'IP child resource'. The selected option is 'ip-12.2.2.10'. Another option, 'none', is visible below it.

- If creating a resource to represent an ERSv2 instance, select the **Dependent Filesystem Resource** to be added as a dependency under the ERS resource. Note that filesystem resources must be in-service (ISP) on this server in LifeKeeper in order to appear on this list. This choice will not appear when creating a resource to represent an ERSv1 instance.

A screenshot of a dropdown menu titled 'Dependent filesystem resource'. The selected option is 'nfs-/exports/usr/sap/EXM/ERS12'. Other options include '/exports/usr/sap/EXM/ERS12' and 'none'.

- Select or enter the **SAP Tag**.

SAP Tag

9. Follow prompts to **extend resource hierarchy**. **Note:** A resource representing an ERSv1 instance may only be extended to one backup server. The hierarchy extension will fail if attempting to extend an ERSv1 hierarchy to a third cluster node.
10. Once **Hierarchy Successfully Extended** displays, select **Finish**.
11. Select **Done**.

### Separate ASCS and ERSv1 Hierarchies



### Separate ASCS and ERSv2 Hierarchies



In the case where the ASCS instance is running ENSAv2, the ERS instance is running ERSv2, and the hierarchies have been extended to the same systems in a cluster with three or more nodes, the ASCS and ERS resource hierarchies can be forced to attempt to avoid each other on switchover and failover by using special “hierarchy avoidance” generic application resources. See [Enforcing ASCS/ERS Avoidance Behavior When Using ENSAv2/ERSv2](#) for more details on how to create these resources and add them as dependencies in the ASCS and ERS hierarchies. **Note:** These hierarchy avoidance resources should not be used in an ENSAv1/ERSv1 configuration or in a two node cluster.

## Create the Primary Application Server Resource

1. Again, for this same SAP SID, repeat the above steps to create the Primary Application Server Resource selecting **DVEBMGS{XX}** (where {XX} is the instance number) when prompted.
2. Select the **Level of Protection** when prompted (default is **FULL**). Click **Next**.



3. Select the **Level of Recovery** when prompted (default is **FULL**). Click **Next**.



4. When prompted for **Dependent Instances**, select the “parent” instance, which would be the **ERS instance** created above.
5. Select the **IP Child Resource**.
6. Follow prompts to extend resource hierarchy.
7. Once **Hierarchy Successfully Extended** displays, select **Finish**.
8. Select **Done**.

### Hierarchy with Primary Application Server as Top Level



## Create the Secondary Application Server Resources

If necessary, create the Secondary Application Server Resources in the same manner as above.

**Note:** For command line instructions, see [Setting Up SAP from the Command Line](#).

## 6.17.7.14. Deleting an SAP Resource Hierarchy

---

To delete a resource from all servers in your LifeKeeper configuration, complete the following steps.

 **Note:** Each resource should be deleted separately in order to delete the entire hierarchy.

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the name of the **TargetServer** where you will be deleting your resource hierarchy.

**Note:** If you selected the **Delete Resource** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

Click **Next**.

3. Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete and highlight it.

**Note:** This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a resource instance in the left or right pane.

Click **Next**.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete**.
5. Another information box appears confirming that the resource was deleted successfully. Click **Done** to exit.

## 6.17.7.15. Common SAP Recovery Kit Tasks

---

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View](#) / [Edit](#) Properties. View or edit the properties of a resource hierarchy on a specific server.

## 6.17.7.16. Setting Up SAP from the Command Line

---

You can set up the SAP Recovery Kit through the use of the command line.

### Creating an SAP Resource from the Command Line

From the Primary Server, execute the following command:

```
$LKROOT/lkadm/subsys/appsuite/sap/bin/create <primary sys> <tag> <SAP SID> <SAP Instance> <switchback type> <IP Tag> <Protection Level> <Recovery Level> <Additional SAP Dependents>
```

#### Example:

```
$LKROOT/lkadm/subsys/appsuite/sap/bin/create liono SAP-STC_SCS00 STC SCS00 intelligent ip-sap10 Full Full none
```

#### Notes:

- **Switchback Type** – This dictates how the SAP instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either **Intelligent** or **Automatic**. Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. [Automatic switchback](#) means the switchback will occur as soon as the primary server comes back on line and re-establishes LifeKeeper communication paths.
- **IP Tag** – This represents the IP resource that will become a dependent of the SAP resource hierarchy.
- **Protection Level** – The **Protection Level** represents the actions that are allowed for each resource.
- **Recovery Level** – The **Recovery Level** provides instruction for the resource in the event of a failure.
- **Additional SAP Dependents** – This value represents the LifeKeeper SAP resource tag that will become a dependent of the current SAP resource being created.

### Extending the SAP Resource from the Command Line

Extending the SAP Resource copies an existing hierarchy from one server and creates a similar hierarchy on another LifeKeeper server. To extend your resource via the command line, execute the following command:

```
system "$LKROOT/lkadm/bin/extmgrDoExtend.pl -p 1 -f, \"$tag\" \"$backupnode\""
```

```
\ "$priority\" \"$switchback\" \\\"$sapbundle\\\"";
```

**Example:** Using a simple script for usability and ease.

```
#!/etc/default/LifeKeeper-perl
require "/etc/default/LifeKeeper.pl";
my $lkroot="$ENV{LKROOT}";
my $tag="SAP";
my $backupnode="snarf";
my $switchback="INTELLIGENT";
my $priority=10;
$sapbundle = "\"$tag\", \"$tag\"";
system "$lkroot/lkadm/bin/extmgrDoExtend.pl -p 1 -f,
 \"$tag\" \"$backupnode\"
 \"$priority\" \"$switchback\" \\\"$sapbundle\\\"";
```

## 6.17.7.17. Activating the SAP SIOS HA Cluster Connector (SSHCC)

---

The SAP SIOS HA Cluster Connector (SSHCC) provides an interface between the SAP Start Service (sapstartsrv) and LifeKeeper. While the HA Cluster Connector is active for an SAP instance, calls through sapcontrol which affect the state of the instance will be routed through LifeKeeper in order to keep the status of the resource in the cluster up-to-date.

In order to activate the SAP SIOS HA Cluster Connector for an SAP instance, follow these steps:

1. Identify the location of the active profile for the SAP instance. This can be found from the /usr/sap/sapservices file by looking for the line corresponding to the instance. For example, consider the following line corresponding to an ASCS instance:

```
LD_LIBRARY_PATH=/usr/sap/STC/ASCS00/exe:$LD_LIBRARY_PATH; export
LD_LIBRARY_PATH; /usr/sap/STC/ASCS00/exe/sapstartsrv pf=/usr/sap/STC/SYS/
profile/STC_ASCS00_sap1 -D -u stcadm
```

In this example, the active ASCS instance profile is located at /usr/sap/STC/SYS/profile/STC\_ASCS00\_sap1.

2. Edit the instance profile found in step 1 and add the following lines to the bottom of the file:

```
#SIOS
#-----
service/halib = saphascriptco.so
service/halib_cluster_connector = /opt/LifeKeeper/lkadm/subsys/appsuite/sap/
bin/sap_sios_cluster_connector
service/halib_debug_level = 1
```

3. Add the following line to the /etc/sudoers file on each system in the cluster to allow the SAP administrator user to run the sap\_sios\_cluster\_connector script. Replace <sid> in the following line with the lower-case SAP SID for your system.

```
<sid>adm ALL=NOPASSWD:/opt/LifeKeeper/lkadm/subsys/appsuite/sap/bin/
sap_sios_cluster_connector-main
```

4. In order for the profile change to take effect, the sapstartsrv process for the instance must be restarted. This can be accomplished with the following command (replacing <sid> with the lower-case SAP SID and <SID> with the upper-case SAP SID):

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function RestartService <SID>"
```

5. To verify that the HA Cluster Connector was successfully activated for the SAP instance, run the following command:

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function HAGetFailoverConfig"
```

- If the HA Cluster Connector has been successfully configured for the instance, the output from this command should show

HAActive: TRUE, as in the following sample output.

```
HAGetFailoverConfig
OK
HAActive: TRUE
HAProductVersion: "LifeKeeper for Linux" steeleye-lk 9.3.2-6863
HASAPInterfaceVersion: "LifeKeeper for Linux" steeleye-lkHACONNECTOR-for-SAP
7.5.0-6855
HADocumentation: docs.us.sios.com/Linux/current/LK4L/SAP/index.htm
HAActiveNode: ip-12-0-0-20
HANodes: ip-12-0-0-20, ip-12-0-1-20, ip-12-0-2-20
```

- If the HA Cluster Connector did not initialize successfully, relevant error messages can be found in the sapstartsrv.log file for the instance, typically located at /usr/sap/<

SID>/<INST>/work/sapstartsrv.log.

## 6.17.7.18. SAP Test Preparation

---

1. Set up an SAP GUI on an SAP client to log in to SAP using the virtual SAP server name.
2. Set up an SAP GUI on an SAP client to log in to the redundant AS.
3. If desired, install additional AS's on other servers in the cluster and set up a login group among all application servers, excluding the PAS. For every AS installed, the profile file will have to be modified as previously described.

## 6.17.7.19. Perform SAP Tests

---

Perform the following series of tests. The test steps are different for each configuration. Some steps call for verifying that SAP is running correctly but do not call out specific tests to perform. For a list of possible tests to perform to verify that SAP is configured and running correctly, refer to the appendices of the SAP document, *SAP R/3 in Switchover Environments*.

### Tests for Active/Active Configurations

1. When the SAP hierarchy is created, the SAP and DB will in-service on different servers. From an SAP GUI, log in to SAP. Verify that you can successfully log in and that SAP is running correctly.
2. Log out and re-log in through a redundant AS. Verify that you can successfully log in.
3. If you have set up a login group, verify that you can successfully log in through this group.
4. Using the LifeKeeper GUI, bring the SAP hierarchy in service on the SAP Backup Server. Both the SAP and DB will now be in service on the same server.
5. Again, verify that you can log in to SAP using the SAP virtual server name, a redundant AS and the login group. Verify that SAP is running correctly.
6. Using the LifeKeeper GUI, bring the DB hierarchy in service on the DB Backup Server. Each hierarchy will now be in service on its backup server.
7. Again, verify that you can log in to SAP using all login methods and that SAP is running correctly. If you execute transaction SM21, you should be able to see in the logs where the PAS lost then regained communication with the DB.
8. While logged in to SAP, shut down the SAP Backup server where SAP is currently in service by pushing the power supply switch. Verify that the SAP hierarchy comes in service on the SAP Primary Server, and that after the failover, you can again log in to the PAS, and that it is running correctly.
9. Restore power to the failed server. Using the LifeKeeper GUI, bring the DB hierarchy back in service on the DB Primary Server. Again, while logged in to SAP, shut down the DB Primary server where the DB is currently in service by pushing the power supply switch. Verify that the DB hierarchy comes in-service on the DB Backup Server and that, after the failover, you are still logged in to SAP and can execute transactions successfully.
10. Restore power to the failed server. Using the LifeKeeper GUI, bring the DB hierarchy back in service on the DB Primary Server.

### Tests for Active/Standby Configurations

1. When the hierarchy is created, both the SAP and DB will be in service on the Primary Server. The redundant AS will be started on the Backup Server. From an SAP GUI, log in to SAP. Verify that

you can successfully log in and that SAP is running correctly. Execute transaction SM51 to see the list of SAP servers. This list should include both the PAS or ASCS and AS.

2. Log out and re-log in through the redundant AS on the Backup Server. Verify that you can successfully log in.
3. If you have set up a login group, verify that you can successfully log in through this group.
4. Using the LifeKeeper GUI, bring the SAP/DB hierarchy in service on the Backup Server.
5. Again, verify that you can log in to SAP using the SAP virtual server name, a redundant AS and the login group. Verify that SAP is running correctly.
6. While logged in to SAP, shut down the SAP/DB Backup server where the hierarchy is currently in service by pushing the power supply switch. Verify that the SAP/DB hierarchy comes in service on the Primary Server, and after the failover, you can again log in to the PAS and that it is running correctly (you will lose your connection when the server goes down and will have to re-log in).
7. Restore power to the failed server. Again, while logged in to SAP, shut down the SAP/DB Primary server where the DB is currently in service by pushing the power supply switch. Verify that the SAP/DB hierarchy comes in service on the Backup Server and that, after the failover, you can again log in to SAP, and that it is running correctly.
8. Again, restore power to the failed server. Using the LifeKeeper GUI, bring the SAP/DB hierarchy in service on the Primary Server.

## 6.17.8. SAP Administration

---

This section provides tips and other information that may be helpful for administration and maintenance of certain configurations.

[NFS Considerations](#)

[Client Reconnect](#)

[LifeKeeper SAP Tunable Parameters](#)

[Separation of SAP and NFS Hierarchies](#)

[Update Protection Level](#)

[Update Recovery Level](#)

[View Properties](#)

## Oracle Database

[Special Considerations for Oracle](#)

## 6.17.8.1. NFS Considerations

As previously described in the [Configuration Considerations](#) topic, if the file system has been configured on either the PAS Primary or Backup server to locally mount NFS shares, an NFS hierarchy out-of-service operation will hang the system and prevent a clean reboot. To avoid causing your cluster to hang by inadvertently stopping the NFS server, we make the following recommendations:

\* **Note:** Please refer to the support matrix and release notes to insure that udp is supported for your OS version. **For systems not supporting UDP** please see [Setting up NFS](#). There you will find new tunable values for improving TCP performance with NFSv4.

- Do not take your NFS hierarchy out of service on a server that contains local NFS mount points to the protected NFS share. You may take your SAP resource in and out of service freely so long as the NFS child resources stay in service. You may also bring your NFS hierarchies in service on a different server prior to shutting a server down.
- If you must stop LifeKeeper on a server where the NFS hierarchy protecting locally mounted NFS shares is in service, always use the `-f` option. Stopping LifeKeeper using the command `lkstop -f` stops LifeKeeper without taking the hierarchies out of service, thereby preventing a server hang due to local NFS mounts. See the `lkstop` man page for additional information.
- If you must reboot a server where the NFS hierarchy protecting locally mounted NFS shares is in service, you should first stop LifeKeeper using the `-f` option as described above. A server reboot will cause the system to stop LifeKeeper without the `-f` option, thereby taking the NFS hierarchies out-of-service and hanging the system.
- If you need to uninstall the SAP package, do not do so when there are SAP hierarchies containing NFS resources that are in-service protected (ISP) on the server. Delete the SAP hierarchy prior to uninstalling the package.
- If you are upgrading LifeKeeper or if you need to run the LifeKeeper Installation setup scripts, it is recommended that you follow the upgrade instructions included in the [LifeKeeper for Linux Installation Guide](#). This includes switching all applications away from the server to be upgraded before running the setup script on the LifeKeeper Installation image file and/or updating your LifeKeeper packages. Specifically, the setup script on the LifeKeeper Installation image file should not be run on a server where LifeKeeper is protecting active NFS shares, since upgrading the `nfsd` kernel module requires stopping NFS on that server which may cause the server to hang with locally mounted NFS file systems. For additional information, refer to the [NFS Server Recovery Kit Documentation](#).
- Using TCP can lead to hangs during out-of-service operations during the `forceumount` call. When NFS shares are not accessible the `umount` can fail. LifeKeeper will attempt to unmount the filesystem multiple times. These multiple attempts will typically succeed in eventually taking the resource out of service. However, this will cause delays in taking the resource out of service. To avoid these retries, use `'nfsvers=3, proto=udp'` mount options.

! Note the usage of udp; this is important for failover and recovery.

- If the /sapmnt (or /sapmnt/<SID>) filesystem is shared via NFS, 'SAP\_NFS\_CHECK\_DIRS=/sapmnt' should be added to /etc/default/LifeKeeper on each node in the cluster to help prevent hangs in SAP resource administration actions due to a loss of the NFS shares.

! SAP\_NFS\_CHECK\_DIRS should not be used if the filesystems are being shared with EFS on AWS since the pingnfs check does not apply in that case.

- If IT restrictions do not allow you to add mount entries to /etc/fstab to mount NFS shared file systems at system boot, these mount entries may instead be added to the LifeKeeper "critical NFS mounts" file for each SAP resource on each server in the cluster. These files are located at /opt/LifeKeeper/subsys/appsuite/resources/sap/critical\_nfs\_mounts\_<Tag>. File systems with mount entries listed in these files will be automatically mounted when the corresponding SAP resource is brought in-service with LifeKeeper. See [Automatic Mounting of Critical NFS Shares](#) for more details.

## 6.17.8.2. SAP Client Reconnect

---

An SAP client can either be configured to log on to a specific SAP instance or a logon group. If configured to log on through a logon group, SAP determines which running instance the client actually connects to.

If the instance to which the client is connected goes down, the client connection is lost and the client must re-log on. If the database is temporarily lost, but the instance to which the client is connected stays up, the client will be temporarily unavailable until the database comes back up but the client does not have to re-log on.

For performance reasons, clients should log on to redundant Application instances and not the PAS, ASCS or SCS. Administrators may wish, however, to be able to log on to the PAS to view logs, etc. After protecting SAP with LifeKeeper, a client login can be configured using the virtual SAP server name so the client can log on regardless of whether the SAP Instance is active on the SAP Primary or Backup server.

## 6.17.8.3. Adjusting SAP Recovery Kit Tunable Values

Several of the SAP scripts have been written with a timeout feature to allow hung scripts to automatically kill themselves. This feature is required due to potential problems with unavailable NFS shares. This is explained in greater detail in the [NFS Mounts and su](#) topic. Each script equipped with this feature has a default timeout value in seconds that can be overridden if necessary. To reduce timeout wait times add required NFS file systems to 'SAP\_NFS\_CHECK\_DIRS' as a comma-separated list. This enables ping to verify the file system is exported and if not it returns immediately.

SAP\_CREATE\_NAS can be enabled or disabled. The default for SAP\_CREATE\_NAS is 1 (enabled). This is used for automatically including a NAS resource for NAS mounted file systems. To disable, set this parameter to 0.

The default for the tunable SAP\_CONFIG\_REFRESH is LKCHECKINTERVAL/2. The user can change the number of seconds between calls to refresh the properties panel for an SAP resource in the LifeKeeper GUI.

The table below shows the script names, variable names, and default values. To override a default value, simply add a line to the `/etc/default/LifeKeeper` file with the desired value for that script. For example, to allow the remove script to run for a full minute before being killed, add the following line to `/etc/default/LifeKeeper`:

```
SAP_REMOVE_TIMEOUT=60
```

 **Note:** The script may actually run for slightly longer than the timeout value before being killed.

 **Note:** It is not necessary to stop and restart LifeKeeper when changing these values.

Script Name	Variable Name	Default Value
remove	SAP_REMOVE_TIMEOUT	804 seconds
restore	SAP_RESTORE_TIMEOUT	516 seconds
recover	SAP_RECOVER_TIMEOUT	1320 seconds
quickCheck	SAP_QUICKCHECK_TIMEOUT	60 seconds
create NAS	SAP_CREATE_NAS	1 (to disable, set to 0)
GUI Properties Panel refresh	SAP_CONFIG_REFRESH	1/2 the value of LKCHECKINTERVAL
NFS shares to check	SAP_NFS_CHECK_DIRS	empty

 **Note:** In a NetWeaver Java Only environment, if you choose to start the Java PAS in addition to the SCS Instance, you may need to increase the values for `SAP_RESTORE_TIMEOUT` and `SAP_RECOVER_TIMEOUT`.

## 6.17.8.4. Separation of SAP and NFS Hierarchies

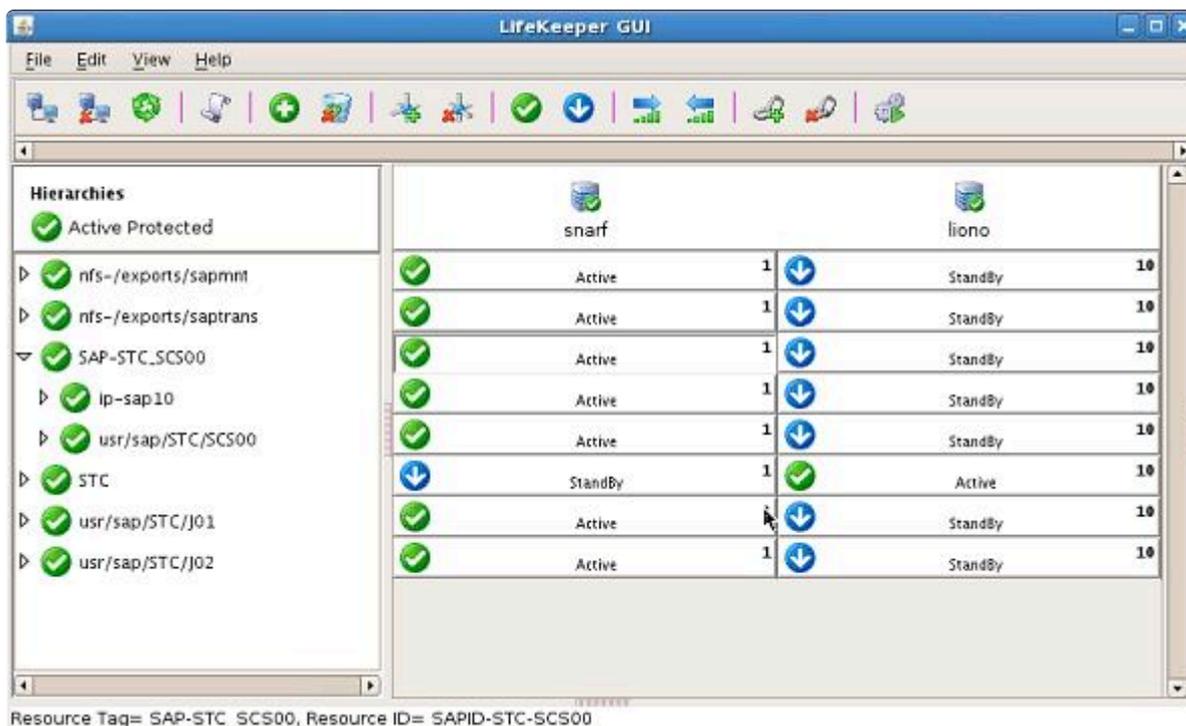
Although the LifeKeeper SAP hierarchy described in this section implements the SAP NFS hierarchies as child dependencies to the SAP resource, it is possible to detach and maintain the NFS hierarchies separately after the SAP hierarchy is created. You should consider the advantages and disadvantages of maintaining these as separate hierarchies as described below prior to removing the dependency. Note that this is only possible if the NFS shares being detached are hosted on a logical volume (LUN) that is separate from other SAP filesystems.

To maintain these two hierarchies separately, simply create the SAP hierarchy as described in this documentation and then manually break the dependency between the SAP and NFS resources through the LifeKeeper GUI.

**Advantage to maintaining SAP and NFS hierarchies separately:** If there is a problem with NFS, the NFS hierarchy can fail over separately from SAP. In this situation, as long as SAP handles the temporary loss of NFS mounted directories transparently, an SAP failover will not occur. The NFS mounts need to be in both /etc/mtab and listed in the tunable value SAP\_NFS\_CHECK\_DIRS in /etc/default/LifeKeeper in order to avoid LifeKeeper hanging while the NFS shares are unavailable.

**Disadvantage to maintaining SAP and NFS hierarchies separately:** NFS shares are not guaranteed to be hosted on the same server where the PAS, ASCS, SCS or ERS instance is running. **Note:** Consult your *SAP Installation Guide* for SAP’s recommendations.

The diagram below shows the SAP and NFS hierarchies after the dependency has been deleted.



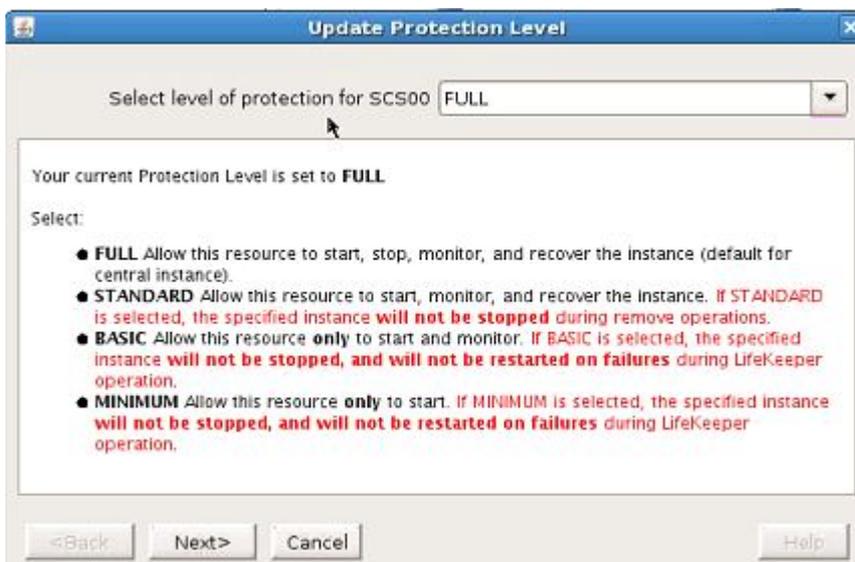
## 6.17.8.5. Update SAP Protection Level

The **Protection Level** represents the actions that are allowed for each resource. To find out what your current protection level is or to change this option, go to **Update Protection Level**. The level of protection can be set to **FULL**, **STANDARD**, **BASIC** or **MINIMUM**.

1. Right-click your instance.
2. Select **Update Protection Level**.



3. The following screen will appear, prompting you to select the Level of Protection.



**FULL.** This is the default level which provides full protection, allowing the instance to be started, stopped, monitored and recovered.

**STANDARD.** Selecting this level will allow the resource to start, monitor and recover the instance, but it will not be stopped during a remove operation.

**BASIC.** Selecting this level will allow the resource to start and monitor only. It will not be stopped or restarted on failures.

**MINIMUM.** Selecting this level will only allow the resource to start the instance. It will not be

stopped or restarted on failures.

 **Note:** The **BASIC** and **MINIMUM** Protection Levels are for placing the LifeKeeper protected application in a temporary maintenance mode. The use of **BASIC** or **MINIMUM** as an ongoing state for the Protection Level of a resource is not recommended. See [Hierarchy Remove Errors](#) in the Troubleshooting Section for further information.

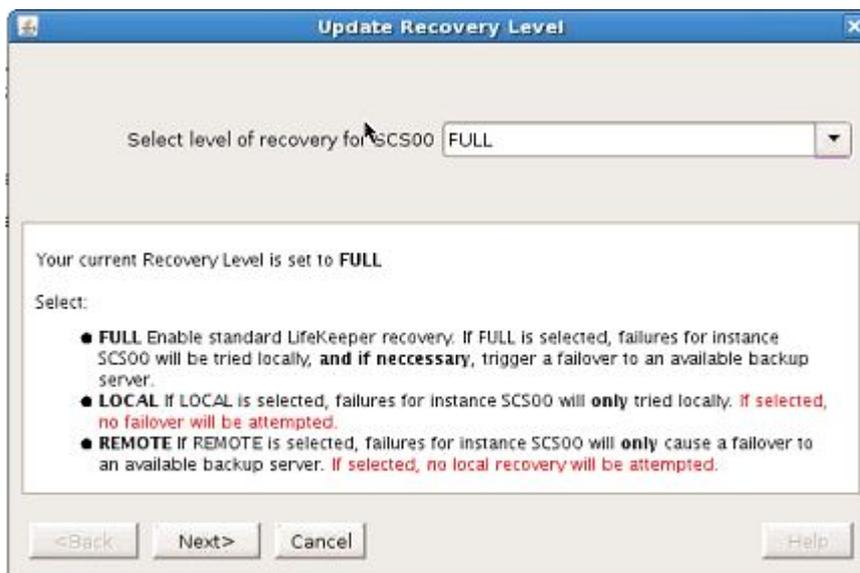
## 6.17.8.6. Update SAP Recovery Level

The **Recovery Level** provides instruction for the resource in the event of a failure. To find out what your current recovery level is or to change this option, go to **Update Recovery Level**. The recovery level can be set to **FULL**, **LOCAL** or **REMOTE**.

1. Right-click your instance.
2. Select **Update Recovery Level**.



3. The following screen will appear, prompting you to select the **Level of Recovery**.



**FULL.** When recovery level is set to **FULL**, the resource will try to recover locally. If that fails, it will try to recover remotely until it is successful.

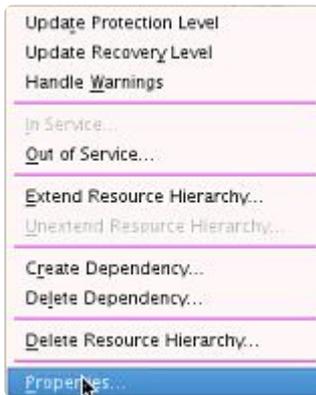
**LOCAL.** When recovery level is set to **LOCAL**, the resource will only try to restart locally; it will not fail over.

**REMOTE.** When recovery level is set to **REMOTE**, the resource will only try to restart remotely. It will not attempt to restart locally first.

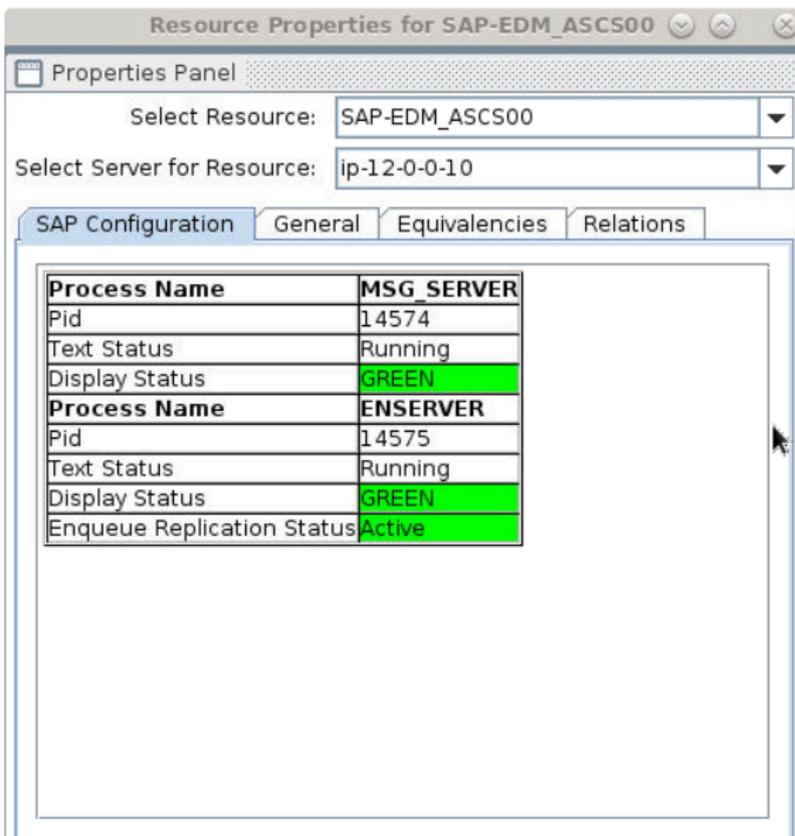
## 6.17.8.7. View SAP Properties

The **Resource Properties** page allows you to view the configuration details for a specific SAP resource. To view the properties of a resource on a specific server or display the status of SAP processes, view the **Properties** Screen:

1. Right-click your instance.
2. Select **Properties**.



3. The following **Properties** screen will appear.



The resulting **Properties** page contains four tabs. The first of those tabs, labeled SAP Configuration, contains configuration information that is specific to SAP resources. The remaining three tabs are available for all LifeKeeper resource types.

## 6.17.8.8. Special Considerations for Oracle

---

Once the SAP processes are functioning on the systems in the LifeKeeper cluster, resources will need to be created in LifeKeeper for the major SAP functions. These include the ASCS system, the DVEBMSG system, the SCS system and the Oracle database.

This topic will discuss some special considerations for protecting Oracle in a LifeKeeper environment.

- Make sure that the LifeKeeper for Linux Oracle Application Recovery Kit is installed.
- Consult the [Oracle Recovery Kit documentation](#).
- During the installation of SAP, the SAPinst process normally assumes that the database software has already been installed and configured. However, if Oracle is the database to be used with SAP, the SAPinst process will prompt the installer to start the Oracle installation tool (RUNINSTALLER) and complete the Oracle install.
- While installing Oracle during the installation of SAP, an Oracle SID was created. This SID is needed by the Oracle Recovery Kit, so be prepared to supply it when creating the Oracle resource in LifeKeeper.
- When creating a standard SAP installation with Oracle, thirteen separate file systems are created that the Oracle instance will use. Commonly, each of these file systems is built on top of an LVM logical volume and each may contain many separate physical volumes. For LifeKeeper to properly represent these file systems, a separate resource is created for each physical and logical volume and volume group. Since this large collection of resources needs to be assembled into a LifeKeeper hierarchy, it may take some time to complete the creation and extension of the Oracle hierarchy. Do not be surprised if it takes at least an hour for the creation process to complete, and another 10 to 20 minutes for the extension to complete.
- Building the necessary Oracle (and SAP) file systems on top of LVM is not required, and the SAP and Oracle recovery kits in LifeKeeper will work fine with standard Linux file systems.
- The LifeKeeper Oracle Recovery Kit can identify ten of the thirteen file systems the Oracle SAP installation uses as standard Oracle dependencies, and the kit will automatically create dependencies in the hierarchy for these file systems. The Oracle Recovery Kit does not recognize the *saptrace*, *sapreorg* and *saparch* file systems automatically. The administrator setting up LifeKeeper will need to manually [create resource dependencies](#) for these additional file systems.

## 6.17.8.9. SSHCC HA Actions

---

The SAP SIOS HA Cluster Connector (SSHCC) Actions provide a list of advanced configuration operations that work in conjunction with the SAP SIOS HA Cluster Connector. The following advanced configuration operations are available:

- **Start Instance** – performs an SAP SIOS HA Cluster Connector start action on the specified resource tag on the current node
- **Stop Instance** – performs an SAP SIOS HA Cluster Connector stop action on the specified resource tag on the current node
- **Migrate Instance** – performs an SAP SIOS HA Cluster Connector migrate action on the specified resource tag on the current node
- **Maintenance Mode** – performs an SAP SIOS HA Cluster Connector maintenance mode action on the specified resource tag on all cluster nodes

 **Note:** These advanced configuration operations should not be used as a replacement for the standard LifeKeeper for Linux in-service or out-of-service operations.

To perform these operations:

1. Right click on your instance
2. Select SSHCC HA Actions from the available menu
3. Select the desired operation from the provided choices
4. Confirm the selected SSHCC HA operation for the specified tag
5. Select Update to continue with the selected operation

## 6.17.8.10. ERS Resource Types in LifeKeeper

**✿ Important:** Please see the [SAP Recovery Kit – Known Issues / Restrictions](#)

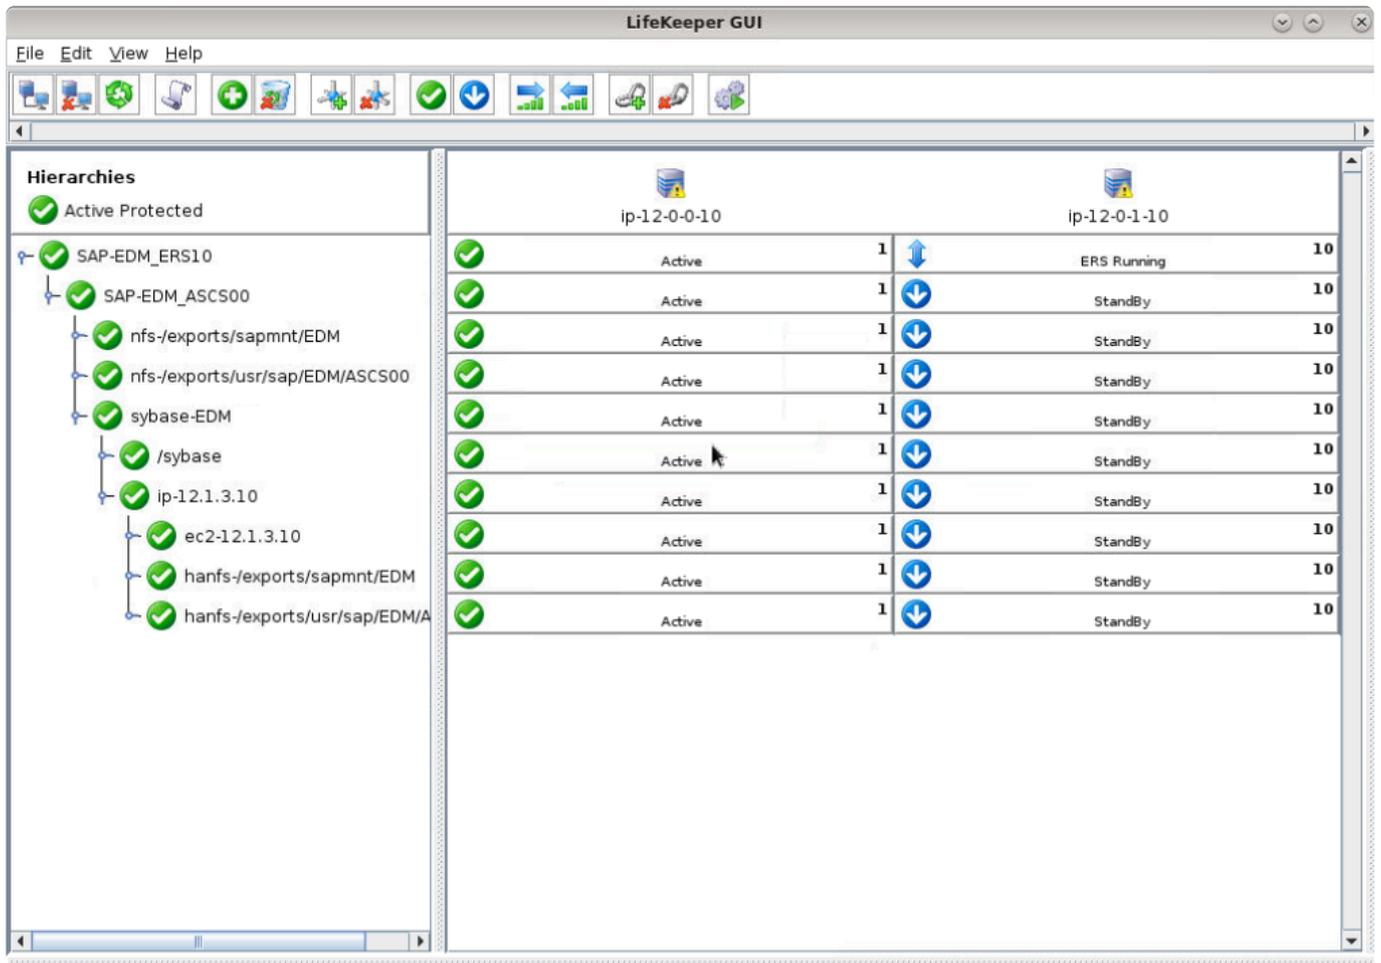
The design and implementation of the ERS resource type was modified in LifeKeeper 9.4.0. This page describes the differences between the implementation prior to version 9.4.0 and the implementation in versions 9.4.0 and later, how to determine which version exists on your system, and how to upgrade to the new resource type.

### ERS Resources Prior to LifeKeeper 9.4.0

In versions of LifeKeeper prior to 9.4.0, the ERS resource was designed to sit at the top of an SAP hierarchy with a dependency on the Central Services (ASCS/SCS) resource that it provides lock table redundancy for.



The behavior of this resource was designed such that it would start the ERS instance on the backup node (where the ERS resource was listed as Standby in the LifeKeeper GUI). Upon switchover or failover of the SAP hierarchy, the ASCS/SCS instance would be started on the backup node and would obtain the backup copy of the lock table (i.e., the replication table) from shared memory on that system. Once the enqueue server successfully obtained the lock table, it would send a signal to notify the ERS instance to terminate itself. Once the ERS resource became Active (ISP) in LifeKeeper on the backup node, the ERS instance would be started on the original primary node when it became available. At that point the replication server would reconnect to the enqueue server and resume lock table replication. When the ERS instance is running on the backup node the LifeKeeper GUI status will change from 'StandBy' to 'ERS Running'.



## ERS Resources in LifeKeeper 9.4.0 and Later

In LifeKeeper 9.4.0 and later, the ERS resource was redesigned to operate in its own independent hierarchy.

ERSv1 Resource in an Independent Hierarchy



ERSv2 Resource in an Independent Hierarchy with Virtual IP and Highly Available Dependent Filesystem

## Hierarchies

✓ Active Protected



This newer ERS resource design supports ERSv1 instances in two-node clusters and ERSv2 instances in clusters with any number of nodes. When using this resource type, the ERS instance will be started on the same node where the LifeKeeper is currently Active (ISP). The design change was made to facilitate the ability for an ERSv2 hierarchy (including its dependent virtual IP and filesystem resources) to failover independently of its corresponding Central Services resource hierarchy.

In order to attempt to keep the ERS resource from being Active (ISP) on the same node as its corresponding Central Services resource, this newer ERS resource type will check during each quickCheck interval (default: two minutes) whether:

1. The ERS resource is Active on the same node as its corresponding Central Services resource,
2. The lock table replication is in-sync between the enqueue server and the replication server, and
3. There is a different node available in the cluster that all resources in the ERS hierarchy could successfully relocate to.

If all three of these conditions are met, LifeKeeper will automatically relocate the ERS hierarchy to a different cluster node in order to provide redundancy of the enqueue server lock table data across cluster nodes. This automatic relocation behavior can be disabled by setting the flag 'sap\_no\_ers\_relocation\_<ERS Tag>' in LifeKeeper. This can be accomplished with a command similar to the following (where SAP-EXM\_ERS12 is the example tag for the ERS resource):

```
/opt/LifeKeeper/bin/flg_create -f
"sap_no_ers_relocation_SAP-EXM_ERS12"
```

Creating this flag will not disable failover due to a failed local recovery of the ERS instance.

**Note when using NFS:** Since this design eliminates the LifeKeeper dependency of the ERS resource on the sapmnt filesystem, it is important to make appropriate use of the tunable value SAP\_NFS\_CHECK\_DIRS to help prevent action scripts from hanging due to the sapmnt NFS share being inaccessible. See [NFS Considerations](#) for more details.

## Which ERS Resource Type Do I Have in My Hierarchy?

If you are not sure in which version of LifeKeeper your existing ERS resource was created, run the following command (replacing <ERS Tag> with the tag name of your ERS resource):

```
/opt/LifeKeeper/bin/ins_list -f: -t <ERS Tag> | cut -d: -f6
```

The output of this command should be similar to:

```
EXMERS12sap2/sapmntFULLFULL5
```

- If the last digit in the output is 2, then this resource was created in LifeKeeper 9.3.2 or earlier. This type of resource may only be used to represent an ERSv1 instance in a two-node cluster and should sit at the top of the SAP hierarchy with a dependency on its corresponding Central Services (ASCS/SCS) instance below it. See the **Upgrading the ERS Resource Type** section below for instructions on how to switch to the newer design that operates in an independent hierarchy.
- If the last digit in the output is 5, then this resource was created in LifeKeeper 9.4.0 or later. This type of resource can be used to represent either an ERSv1 instance in a two-node cluster or an ERSv2 instance in a cluster with any number of nodes. It should operate in a hierarchy independent of its corresponding Central Services resource.

## Upgrading the ERS Resource Type

To upgrade to the newer ERS resource type in LifeKeeper 9.4.0 or later, complete the following steps.

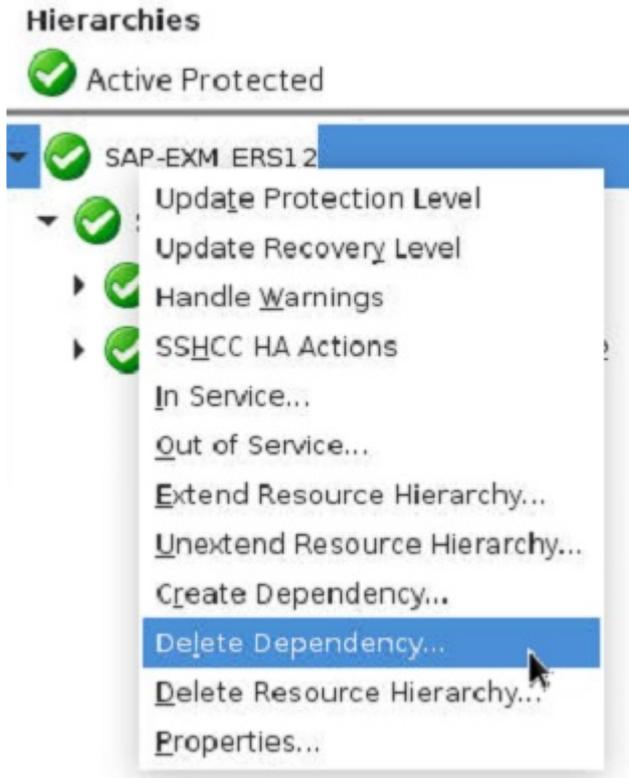
1. Before attempting the upgrade process, create a backup of your LifeKeeper hierarchies on all cluster nodes by running the command:

```
/opt/LifeKeeper/bin/lkbackup -c --cluster
```

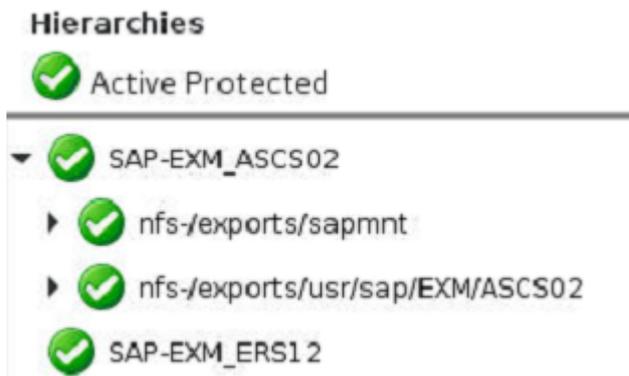
- If you make a mistake at any point, the saved hierarchy configuration can be restored by stopping LifeKeeper on all nodes and running the command:

```
/opt/LifeKeeper/bin/lkbackup -x--cluster
```

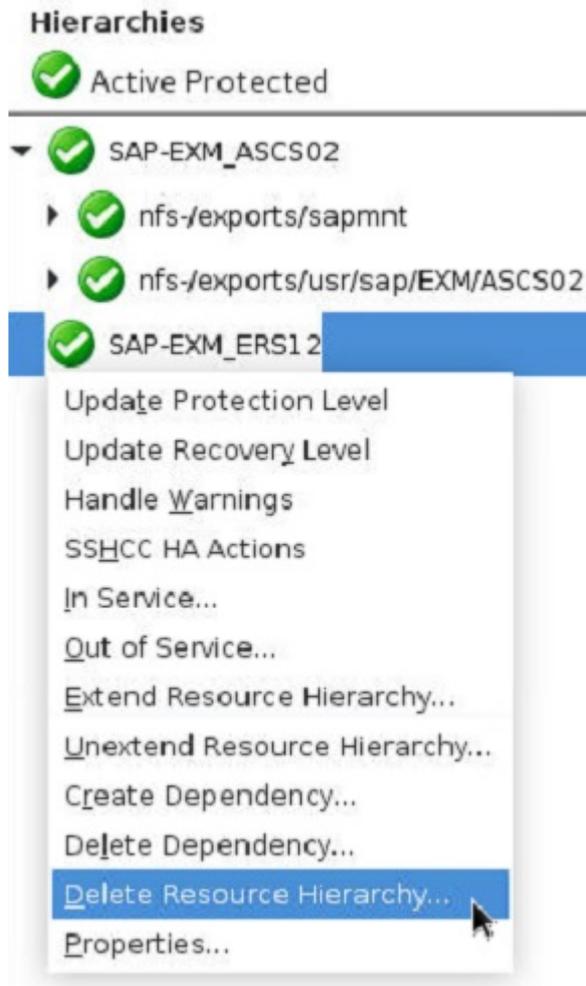
2. Right-click the ERS resource in the LifeKeeper hierarchy panel and select **Delete Dependency...**



3. Delete all dependencies of the ERS resource. You may need to select **Delete Dependency...** multiple times in order to delete all dependencies. When you are finished, the ERS resource should exist in a hierarchy by itself with no child dependencies and no other resources dependent on it, as shown in the following image.



4. Right-click the ERS resource in the LifeKeeper hierarchy panel and select **Delete Resource Hierarchy...** Select any server as the **Target Server** and click **Next**. Click **Delete** to delete the ERS resource on all nodes. **Warning:** If you did not successfully delete all dependencies between the ERS resource and other resources in the SAP hierarchy in the previous step, then this step could delete your entire SAP hierarchy.



5. Once the ERS resource has been successfully deleted on all nodes, follow the instructions in the “Create the ERS Resource” section in [SAP Installation → Creating an SAP Hierarchy](#) to create a new ERS resource and extend it to the desired cluster nodes.

## 6.17.8.11. Upgrading from ENSAv1 to ENSAv2

---

In order to upgrade from Standalone Enqueue Server version 1 to Standalone Enqueue Server version 2, first ensure that your SAP kernel version supports ENSAv2 then complete the following steps:

1. Set the following parameters in the default profile (typically located at `/usr/sap/<SID>/SYS/profile/DEFAULT.PFL`). These parameters must be the same for all instances:

```
enq/enable=TRUE
enq/serverhost=<ASCS instance host>
enq/serverinst=<ASCS instance number>
enque/process_location=REMOTESA
```

2. In the ASCS instance profile (typically located at `/usr/sap/<SID>/SYS/profile/<SID>_ASCS<No>_<VIP>`), set the following parameters:

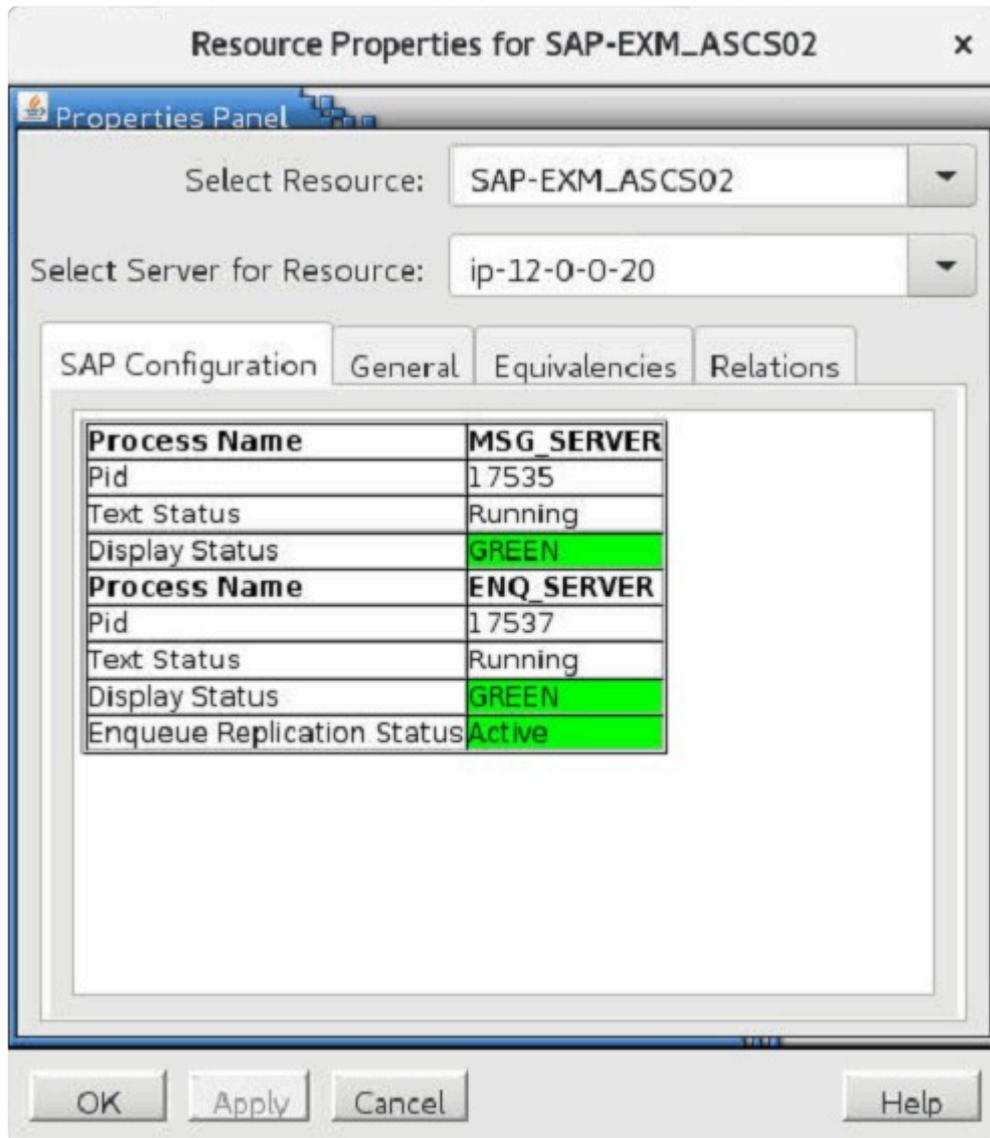
```
_ENQ = enq.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
Execute_01 = local rm -f $_ENQ
Execute_02 = local ln -s -f $(DIR_EXECUTABLE)/enq_server$(FT_EXE)
Start_Program_01 = local $_ENQ pf=$_PF
```

**Note:** The number used for the `Execute_*` and the `Start_Program_*` parameters should be the first number not yet used for that parameter in this profile.

3. After setting the parameters in the default and ASCS instance profiles, restart the SAP Start Service for the ASCS instance by running the following command (replacing `<sid>` with your lower-case SAP SID and `<SID>` with your upper-case SAP SID):

```
su - <sid>adm -c "sapcontrol -nr <ASCS Inst#> -function RestartService <SID>"
```

4. **On all cluster nodes that the ASCS resource has been extended to**, edit the file `/opt/LifeKeeper/subsys/appsuite/resources/sap/INFO_<ASCS Tag>` and ensure that `SAPENQ_VERSION=2`. This may require you to add the line `"SAPENQ_VERSION=2"` if the INFO file does not yet contain a value for `SAPENQ_VERSION`.
5. Restart the SAP system. Once the ASCS instance is restarted, it will be using the `enq_server` (ENSAv2) process instead of the `enserver` (ENSAv1) process. This can be verified by right-clicking the ASCS resource in LifeKeeper and selecting **Properties...**



**Note:** The Enqueue Replication Status on your system will not show Active until you have also completed the upgrade from ERSv1 to ERSv2 for the corresponding ERS resource. See [Upgrading from ERSv1 to ERSv2](#) for more information.

## 6.17.8.12. Upgrading from ERSv1 to ERSv2

In order to upgrade from Enqueue Replication Server version 1 to Enqueue Replication Server version 2, first ensure that your SAP kernel version supports ERSv2 then complete the following steps:

1. Upgrade the ASCS instance to use ENSAv2 by following the instructions on [Upgrading from ENSAv1 to ENSAv2](#). The same version of the Standalone Enqueue Server and Enqueue Replication Servers must be used since mixed version configurations are not supported by SAP. See the SAP documentation on ENSAv2/ERSv2 for more details.
2. Set up the virtual IP and shared file system for the ERS instance. Also create the corresponding virtual IP and filesystem resource hierarchies in LifeKeeper. They will be used in step 9 during the recreation of the ERS resource. **Note:** These LifeKeeper resources should be Active (ISP) on a node where the corresponding ASCS instance is currently Standby (OSU).
3. Set the following parameters in the default profile (typically located at /usr/sap/<SID>/SYS/profile/DEFAULT.PFL). These parameters must be the same for all instances:

```
enq/enable=TRUE
enq/serverhost=<ASCS instance host>
enq/serverinst=<ASCS instance number>
enq/replicatorhost=<ERS instance host>
enq/replicatorinst=<ERS instance number>
enque/process_location=REMOTESA
```

4. In the ASCS instance profile (typically located at /usr/sap/<SID>/SYS/profile/<SID>\_ASCS<No>\_<VIP>), set the following parameters:

```
enq/server/replication/enable = true
_ENQ = enq.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
Execute_01 = local rm -f $_ENQ
Execute_02 = local ln -s -f $(DIR_EXECUTABLE)/enq_server$(FT_EXE) $_ENQ
Start_Program_01 = local $_ENQ pf=$_PF
```

**Note:** The number used for the Execute\_\* and the Start\_Program\_\* parameters should be the first number not yet used for that parameter in this profile.

5. In the ERS instance profile (typically located at /usr/sap/<SID>/SYS/profile/<SID>\_ERS<No>\_<VIP>), set the following parameters:

```
_ENQR = enqr.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
Execute_01 = local rm -f $_ENQR
Execute_02 = local ln -s -f $(DIR_EXECUTABLE)/enq_replicator$(FT_EXE)
$_ENQR
Start_Program_00 = local $_ENQR pf=$_PF
```

**Note:** The number used for the `Execute_*` and the `Start_Program_*` parameters should be the first number not yet used for that parameter in this profile.

- After setting the parameters in the default, ASCS, and ERS instance profiles, restart the SAP Start Service for the ASCS and ERS instances by running the following commands (replacing `<sid>` with your lower-case SAP SID and `<SID>` with your upper-case SAP SID):

```
su - <sid >adm -c "sapcontrol -host <ASCS VIP> -nr <ASCS Inst#> -function
RestartService <SID>"
su - <sid >adm -c "sapcontrol -host <ERS VIP> -nr <ERS Inst#> -function
RestartService <SID>"
```

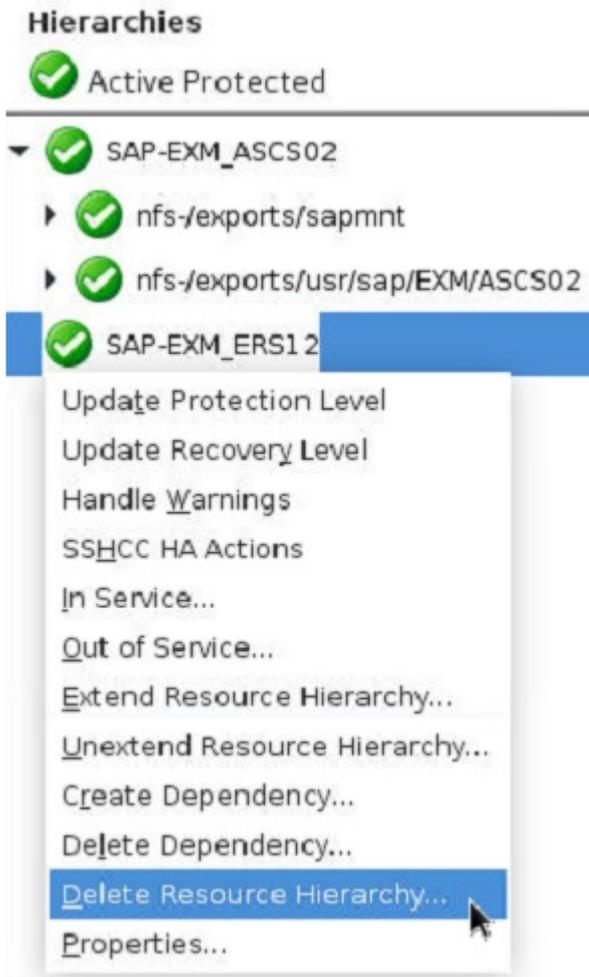
- If you do not have an existing ERS resource in LifeKeeper that needs to be upgraded, skip to step 9. Otherwise, delete all dependencies of the ERS resource in LifeKeeper by right-clicking the resource and selecting **Delete Dependency....** You may need to select **Delete Dependency...** multiple times in order to delete all dependencies. When you are finished, the ERS resource should exist in a hierarchy by itself with no child dependencies and no other resources dependent on it, as shown in the following image.

### Hierarchies

 Active Protected

- ▼  SAP-EXM\_ASCS02
  - ▶  nfs-/exports/sapmnt
  - ▶  nfs-/exports/usr/sap/EXM/ASCS02
-  SAP-EXM\_ERS12

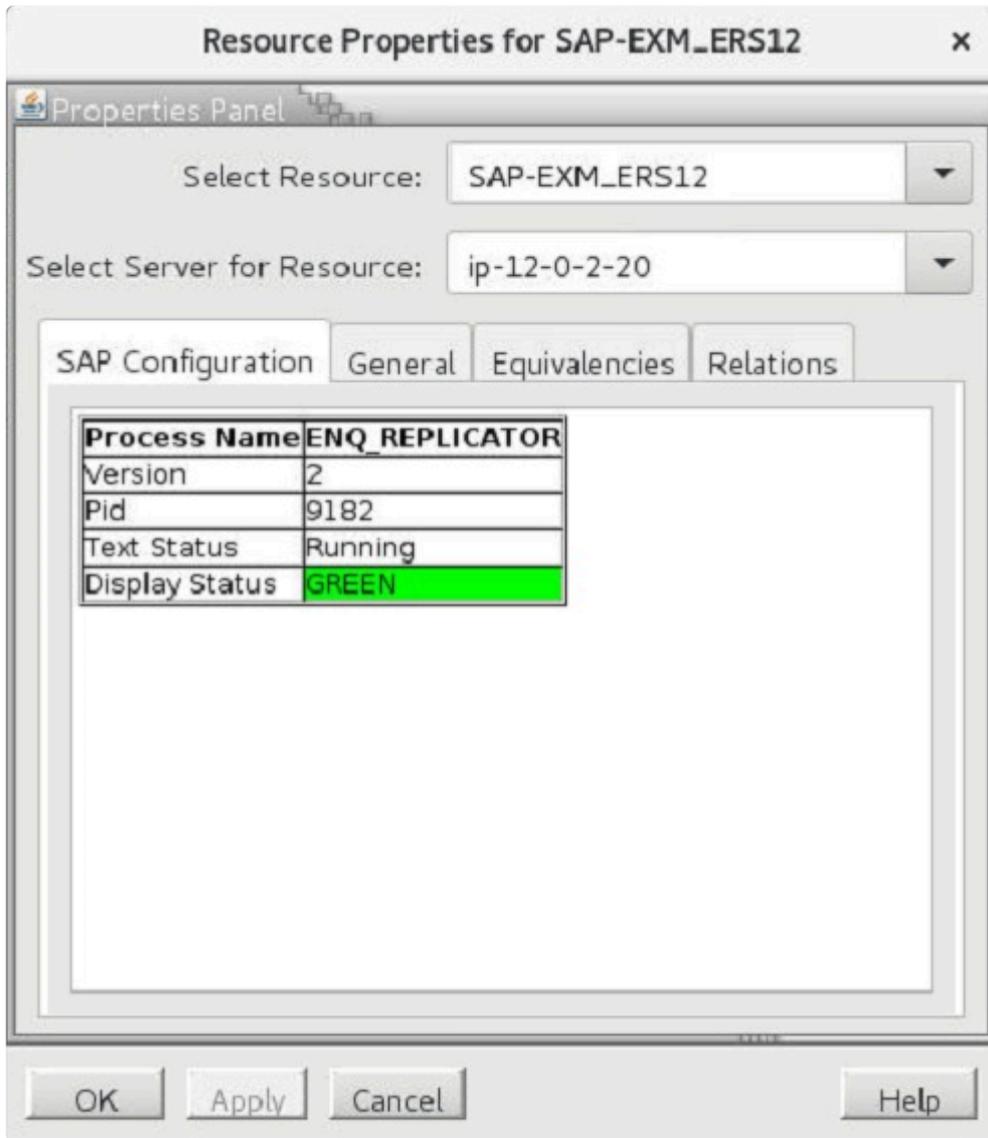
- Right-click the ERS resource in the LifeKeeper hierarchy panel and select **Delete Resource Hierarchy....** Select any server as the **Target Server** and click **Next**. Click **Delete** to delete the ERS resource on all nodes. **Warning:** If you did not successfully delete all dependencies between the ERS resource and other resources in the SAP hierarchy in the previous step, then this step could delete your entire SAP hierarchy.



9. Once the ERS resource has been successfully deleted on all nodes, follow the instructions in the “Create the ERS Resource” section in [SAP Installation → Creating an SAP Hierarchy](#) to create a new ERS resource and extend it to the desired cluster nodes. During the ERS resource creation, select the virtual IP and dependent filesystem resources created in step 2, if applicable. If selected, these resources will be automatically added as child dependencies in the ERS hierarchy.



10. On all cluster nodes that the ERS resource has been extended to, edit the file /opt/LifeKeeper/subsys/appsuite/resources/sap/INFO\_<ERS Tag> and ensure that SAPENQREP\_VERSION=2. This may require you to add the line "SAPENQREP\_VERSION=2" if the INFO file does not yet contain a value for SAPENQREP\_VERSION.
11. Restart the SAP system. Once the ERS instance is restarted, it will be using the enq\_replicator (ERSv2) process instead of the enrepsvr (ERSv1) process. This can be verified by right-clicking the ERS resource in LifeKeeper and selecting **Properties....**



## 6.17.8.13. Automatic Mounting of Critical NFS Shares

---

In a highly-available SAP environment, some important file systems are shared with and mounted on each server that hosts a LifeKeeper-protected SAP instance. Examples include the /sapmnt (or /sapmnt/<SID>) and /usr/sap/trans file systems. Typically these shared file systems are mounted at system boot by adding a mount entry for the file system to the /etc/fstab file. For example:

```
sapnfs:/export/sapmnt/SID /sapmnt/SID nfs rw,sync,bg 0 0
```

However, for compliance reasons, some administrators may be unable to add mount entries directly to the /etc/fstab file on their SAP cluster servers.

In this case an administrator may instead add fstab-style mount entries to the LifeKeeper “critical NFS mounts” file for each SAP resource, which is located on each server at **/opt/LifeKeeper/subsys/appsuite/resources/sap/critical\_nfs\_mounts\_<Tag>**. Before attempting any administrative actions for an SAP resource, LifeKeeper will verify that each file system present in the critical\_nfs\_mounts file for that resource is mounted, and will attempt to mount any listed file system which is not currently mounted.

### Example

For this example we will assume that we have a two-node cluster (with hostnames node-a and node-b) with protected SAP instances ASCS10 (protected by LifeKeeper resource **SAP-SPS\_ASCS10**) and ERS20 (protected by LifeKeeper resource **SAP-SPS\_ERS20**) installed under SAP system ID ‘SPS’. The SAP Mount file system for the SAP installation with SID ‘SPS’ is being shared by a highly-available NFS server cluster using virtual hostname ‘sapnfs’ and export point sapnfs:/export/sapmnt/SPS. This file system must be mounted at /sapmnt/SPS on each server before either protected instance (ASCS10 or ERS20) can be started there.

If the /etc/fstab file cannot be modified to mount this file system at boot, then the following entries can be added to the LifeKeeper critical\_nfs\_mounts files for both SAP resources on both servers.

#### **/opt/LifeKeeper/subsys/appsuite/resources/sap/critical\_nfs\_mounts\_SAP-SPS\_ASCS10**

```
# critical_nfs_mounts_SAP-SPS_ASCS10
# NFS shared file system mounts added to this file will be
# automatically mounted by the SIOS LifeKeeper
# SAP Recovery Kit before performing any SAP administrative actions
# for this resource.
#
# Duplicate entries in /etc/fstab with the same mount point take
# precedence over this file.
sapnfs:/export/sapmnt/SPS /sapmnt/SPS nfs rw,sync,bg 0 0
```

#### **/opt/LifeKeeper/subsys/appsuite/resources/sap/critical\_nfs\_mounts\_SAP-SPS\_ERS20**

```
# critical_nfs_mounts_SAP-SPS_ERS20
# NFS shared file system mounts added to this file will be
automatically mounted by the SIOS LifeKeeper
# SAP Recovery Kit before performing any SAP administrative actions
for this resource.
#
# Duplicate entries in /etc/fstab with the same mount point take
precedence over this file.
sapnfs:/export/sapmnt/SPS /sapmnt/SPS nfs rw,sync,bg 0 0
```

**!** The `critical_nfs_mounts` files are server and resource-specific, so the desired mount entries must be added to the appropriate files for all intended SAP resources on all servers in the cluster.

Now suppose that we wish to bring the `SAP-SPS_ERS20` resource in-service on node-b, but the `/sapmnt/SPS` file system is not currently mounted there. As long as the mount entry shown above exists in the `critical_nfs_mounts_SAP-SPS_ERS20` file on node-b, LifeKeeper will automatically attempt to mount the file system while bringing the `SAP-SPS_ERS20` resource in-service. When it performs the mount operation, a message similar to the following is logged:

```
Jan 1 00:00:00 node-b restore[16995]:
INFO:sap:restore:SAP-SPS_ERS20:112184:Mounting NFS shared file system
'sapnfs:/export/sapmnt/SPS' at mount point '/sapmnt/SPS' with mount
options 'rw,sync,bg' on server node-b.
```

As long as the NFS server is available and the shared file system is properly exported, the file system should mount successfully, allowing the `SAP-SPS_ERS20` resource to come in-service successfully on node-b.

## 6.17.9. SAP Troubleshooting

---

This section provides a list of messages that you may encounter during the process of creating and extending a LifeKeeper SAP resource hierarchy, removing and restoring a resource and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition.

Messages from other LifeKeeper components are also possible. In these cases, please refer to the [Message Catalog](#) which provides a listing of all error codes, including operational, administrative and GUI, that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the error code and necessary action to resolve the issue. This full listing may be searched for any error code received, or you may go directly to one of the individual Message Catalogs for the appropriate LifeKeeper component.

Messages in this section fall under these topics:

[Changing ERS Instances](#)

[ASCS + ERS Restart\\_Program Parameter](#)

[Hierarchy Remove Errors](#)

[SAP Error Messages During Failover or In-Service](#)

[SAP Installation Errors](#)

[Troubleshooting sapinit](#)

[‘tset’ Errors Appear in the LifeKeeper Log File](#)

## 6.17.9.1. Disable Autostart in ERS Profile

---

### Symptom:

A status check of an ERS Instance causes a `sapstart` for the selected instance.

ERS instance is always running on both systems.

### Cause:

When an ERS Instance has `Autostart=1` set in the profile, certain `sapcontrol` calls will cause the instance to be started as a part of the running command.

### Action:

Stop the running ERS Instances in the cluster and modify the profile for the ERS Instances and set `Autostart=0`.

This profile change will require a restart of `sapstartsrv` in order to take effect. This can be accomplished by running the following command (replacing `<sid>` by the lower-case SAP SID and `<SID>` by the upper-case SAP SID):

```
su - <sid>adm -c "sapcontrol -nr <ERS Inst#> -function RestartService <SID>"
```



Changing autostart from one value to another requires a reboot in order for the parameter change to take effect.

## 6.17.9.2. ASCS + ERS Restart\_Program Parameter

---

### Symptom:

The enqueue server or enqueue replicator process is running on both the primary and backup servers at the same time.

### Cause:

If the ERS instance profile is configured to start the instance with Restart\_Program instead of Start\_Program, the sapstart process will automatically restart the ERS instance when it is terminated for any reason. This will cause unexpected behavior and could lead to a loss of the enqueue server lock table during failover or switchover.

### Action:

Modify the ASCS and ERS instance profile parameters to use Start\_Program instead of Restart\_Program when starting the enqueue server and enqueue replicator processes. See [Modify ASCS and ERS Instance Profile Settings](#) for more details.

Once the ASCS and ERS profiles have been updated to use Start\_Program instead of Restart\_Program and SAP Start Service has been restarted, SIOS recommends restarting the system to ensure the updated profile is read and no caching is in effect.

If successful, the ERS instance should only be running on the backup system where the ASCS/SCS instance is not currently running.

 For more details refer to [Installing SAP](#).

## 6.17.9.3. SAP Hierarchy Remove Errors

---

### Symptom:

File system remove fails with file system in use.

### Cause:

1. **Resource Protection Level** was set to **Basic** or **Minimum** after the create or extend. When the resource Protection Level is set to Basic or Minimum, the SAP resource hierarchy will not be stopped during the remove operation. This leaves the processes running for that instance when remove is called. If the processes are also accessing the protected file system, LifeKeeper may be unable to unmount the file system.
2. **Resource Protection Level** was set to **Standard** for a non-replicated enqueue resource. When the resource Protection Level is set to Standard, the SAP resource hierarchy will not be stopped during the remove operation. This leaves the processes running for that instance when remove is called. If the processes are also accessing the protected file system, LifeKeeper may be unable to unmount the file system.

### Action:

1. The **Basic** and/or **Minimum** settings should be used to place a resource in a temporary maintenance mode. It should not be used as an ongoing Protection Level. If the resource in question will require Basic or Minimum as the ongoing Protection Level, the Instance should be configured to use local storage and/or the entire resource hierarchy should be configured without the use of the LifeKeeper NAS Recovery Kit for local NFS mounts.
2. The **Standard** setting should be used for replicated enqueue resources only. **Note: Standard** is used for ERS resources that were created in LifeKeeper 9.3.2 or earlier and reside at the top of the SAP hierarchy with a dependency on the corresponding central services resource. ERS instances created in LifeKeeper 9.4.0 or later that reside in a hierarchy independent of the central services resource should have their Protection Level set to **Full**.

## 6.17.9.4. SAP Hierarchy Restore Errors

### Symptom:

When attempting to bring an SAP resource hierarchy which does not contain a dependent file system resource back in-service on a previous host node, the in-service (restore) of the SAP resource fails.

Messages similar to the following appear in the SAP instance trace logs indicating that a port required by one of the instance processes is already in-use (this example was taken from `/usr/sap/<SID>/ASCS<##>/work/dev_ms.new`):

```
[Thr 140225937864512] ***LOG Q0I=> NiIBindSocket: bind (98: Address already in use) [/bas/781_REL/src/base/ni/nixxi.cpp 3946]
[Thr 140225937864512] *** ERROR => NiIBindSocket: SiBind failed for hdl 1/sock 6 (SI_EPORT_INUSE/98; I4; ST; 0.0.0.0:3610) [nixxi.cpp 3946]
[Thr 140225937864512] *** ERROR => MsSCommInit: NiBuf2Listen(sapmsSHC) (rc=NIESERV_USED) [msxxserv.c 12838]
[Thr 140225937864512] *** ERROR => MsSInit: MsSCommInit [msxxserv.c 2732]
[Thr 140225937864512] *** ERROR => MsSInit failed, see dev_ms.new for details [msxxserv.c 7363]
```

### Cause:

**Resource Protection Level** was set to **Basic**, **Minimum**, or **Standard** after the create or extend. When the resource Protection Level is set to Basic, Minimum, or Standard, the protected SAP instance will not be stopped when taking the SAP resource out of service. This leaves the processes for that instance running after the remove script is called. Since the SAP resource hierarchy does not contain a dependent file system resource, the file systems that the SAP instance depends on will stay mounted after the SAP hierarchy is taken out of service. This means that the SAP instance processes will continue running, whereas they would have been killed with a `fuser -k` call if LifeKeeper had been unmounting an underlying file system.

When LifeKeeper attempts to bring the resource back in-service on the server where the SAP instance processes were left running, the protected instance is unable to start due to the required ports already being in-use by these running processes.

### Action:

1. The **Basic** and/or **Minimum** settings should only be used to place a resource in a temporary maintenance mode. They should not be used as ongoing Protection Level settings.
2. The **Standard** setting may be used for **replicated** enqueue resources only. In particular, it should not be used when the SAP resource does not have a dependent file system under LifeKeeper protection.

\* **Note: Standard** is also used for ERS resources that were created in LifeKeeper 9.3.2 or earlier and reside at the top of the SAP hierarchy with a dependency on the corresponding central services resource. ERS instances created in LifeKeeper 9.4.0 or later that reside in a hierarchy independent of the central services resource should have their Protection Level set to **Full**.

3. To resolve the issue, follow these steps:
  - a. Reboot all servers in the cluster that the SAP resource has been extended to. After rebooting, all processes in the protected SAP instance will be stopped and the required ports will be available.
  - b. Bring the SAP resource hierarchy in-service on one of the servers. Change the SAP resource Protection Level to **Full** by right-clicking the resource in the LifeKeeper GUI and selecting "Update Protection Level".
  - c. Repeat step (b) for each server in the cluster that the SAP resource has been extended to.

## 6.17.9.5. SAP Error Messages During Failover or In-Service

---

After a failover of a SAP, there will be error messages in the SAP logs. Many of these error messages are normal and can be ignored.

### On Failure of the DB

**BVx: Work Process is in reconnect status** – This error message simply states that a work progress has lost the connection to the database and is trying to reconnect.

**BVx: Work Process has left reconnect status** – This is not really an error, but states that the database is back up and the process has reconnected to it.

**Other errors** – There could be any number of other errors in the logs during the period of time that the database is down.

### On Startup of the CI

**E15: Buffer SCSA Already Exists** – This error message is not really an error at all. It is simply telling you that a previously created shared memory area was found on the system which will be used by SAP.

**E07: Error 00000 : 3No such process in Module rslgsmcc (071)** – See *SAP Note 7316* – During the previous shutdown, a lock was not released properly. This error message can be ignored.

### During a LifeKeeper In-Service Operation

The following messages may be displayed in the LifeKeeper In Service Dialog during an in-service operation:

```
error: permission denied on key 'net.unix.max_dgram_qlen'
```

```
error: permission denied on key 'kernel.cap-bound'
```

These errors occur when `saposcol` is started and can be ignored (see *SAP Note 201144*).

## 6.17.9.6. SAP Installation Errors

---

### Incorrect Name in *tnsnames.ora* or *listener.ora* Files

#### Cause:

When using the Oracle database, if the SAP installation program complains about the incorrect server name being in the *tnsnames.ora* or *listener.ora* file when you do the PAS Backup Server installation, then you may not have installed the Oracle binaries on local file systems.

#### Action:

The Oracle binaries in */oracle//920<32 or 64>\_* must be installed on a local file system on each server for the configuration to work properly.

## 6.17.9.7. Troubleshooting sapinit

---

### Symptom:

`sapstartsrv` processes and additional SAP instance processes started by init script fail or cause processes to run on LifeKeeper Standby Node.

### Cause:

SAP provides an init script for automatically starting SAP instances on a local node. When a resource is added to LifeKeeper protection, the init script (`sapinit` ) may attempt to start SAP Instance processes that should not be running on the current node.

### Action:

Disable the `sapinit` script or modify `sapinit` to skip over LifeKeeper protected Instances. To disable this behavior, the user must stop `sapinit` (Example: `/etc/init.d/sapinit stop` ). The `sapinit` script should also be disabled using `chkconfig` or similar tool (Example: `chkconfig sapinit off`).

## 6.17.9.8. tset Errors Appear in the LifeKeeper Log File

---

### Cause:

The su commands used by the SAP and Database Recovery Kits cause a 'tset' error message to be output to the LK log that appears as follows:

```
tset: standard error: Invalid argument
```

This error comes from one of the profile files in the SAP administrator's and Database user's home directory and it is only in a non-interactive shell.

### Action:

If using the c-shell for the Database user and SAP Administrator, add the following lines into the `.sapenv_<hostname>.sh` in the home directory for these users. This code should be added around the code that determines if 'tset' should be executed:

```
if ( $?prompt ) then

    tty -s

    if ( $status == 0 ) then

        .

        .

        .

    endif

endif

endif
```

**Note:** The code from "tty -s" to the inner "endif" already exists in the file.

If using the bash shell for the Database user and SAP Administrator, add the following lines into the `.sapenv_<hostname>.sh` in the home directory for the users.

Before the code that determines if 'tset' should be executed add:

```
case $- in

*i*) INTERACTIVE ="yes" ;;
```

```
*) INTERACTIVE ="no";;
```

```
esac
```

Around the code that determines if 'tset' should be executed add:

```
if [ $INTERACTIVE == "yes" ]; then
```

```
  tty -s
```

```
  if [ $? -eq 0 ]; then
```

```
    .
```

```
    .
```

```
    .
```

```
  fi
```

```
fi
```

**Note:** The code from "tty -s" to the inner "endif" already exists in the file.

## 6.17.10. Maintenance Mode

---

### Maintaining a LifeKeeper Protected System

When performing shutdown and maintenance on a LifeKeeper-protected server, you must put that system's resource hierarchies in service on the backup server before performing maintenance. This process stops all activity for shared disks on the system needing maintenance.

Perform these actions in the order specified, where *Server A* is the primary system in need of maintenance and *Server B* is the backup server:

1. **Bring hierarchies in service on Server B.** On the backup, *Server B*, use the LifeKeeper GUI to bring in service any resource hierarchies that are currently in service on *Server A*. This will unmount any file systems currently mounted on *Server A* that reside on the shared disks under LifeKeeper protection. See [Bringing a Resource In Service](#) for instructions.
2. **Stop LifeKeeper on Server A.** Use the command `$LKROOT/bin/lkcli stop` to stop LifeKeeper.
3. **Shut down Linux and power down Server A.** Shut down the Linux operating system on *Server A*, then power off the server.
4. **Perform maintenance.** Perform the necessary maintenance on *Server A*.
5. **Power on Server A and restart Linux.** Power on *Server A*, then reboot the Linux operating system.
6. **Start LifeKeeper on Server A.** Use the command `$LKROOT/bin/lkcli start` to start LifeKeeper.
7. **Bring hierarchies back in-service on Server A, if desired.** On *Server A*, use the LifeKeeper GUI to bring in service all resource hierarchies that were switched over to *Server B*.

## 6.17.10.1. SAP Maintenance Mode

### Placing LifeKeeper Protected Resources in Maintenance Mode During SAP Software Update

LifeKeeper has the ability to place the resources in an SAP hierarchy into **maintenance mode** to allow the user to upgrade their SAP software in-place without having to first bring the hierarchy in-service on a backup server.

#### When maintenance mode is enabled for the SAP hierarchy:

- Resource health monitoring, local recovery, and failover are disabled for all resources in the hierarchy.

 **Note:** To minimize the potential for data corruption during the upgrade process, SCSI reservation errors will still be detected and acted on by LifeKeeper. This may cause a resource hierarchy failover or may halt the system in the case of a lost SCSI reservation. To modify the behavior of LifeKeeper when a SCSI device cannot be accessed, see the `SCSIERROR` parameter in the [Core Parameters List](#).

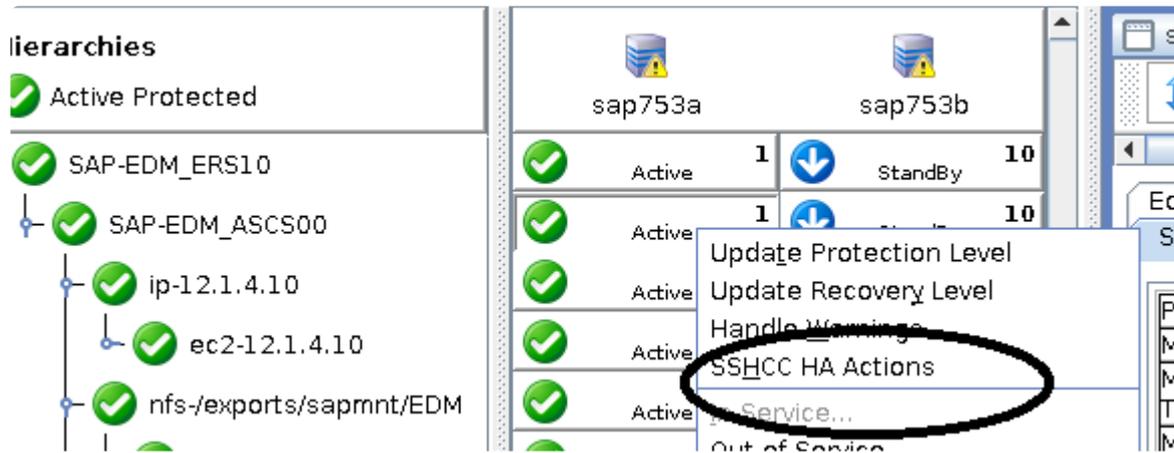
 **Note:** When using Quorum/Witness features, the default behavior of LifeKeeper is to halt the node upon loss of quorum. This is beneficial if the user has critical application hierarchies for which a loss of quorum needs to be acted on, but could possibly lead to an unexpected halt while performing maintenance activities if a network communication error occurs. Steps to manually disable/enable Quorum/Witness features can be found on the [Quorum/Witness](#) documentation page.

- Any DataKeeper mirrors in the hierarchy will be paused until either maintenance mode is disabled or the user manually resumes the mirror.

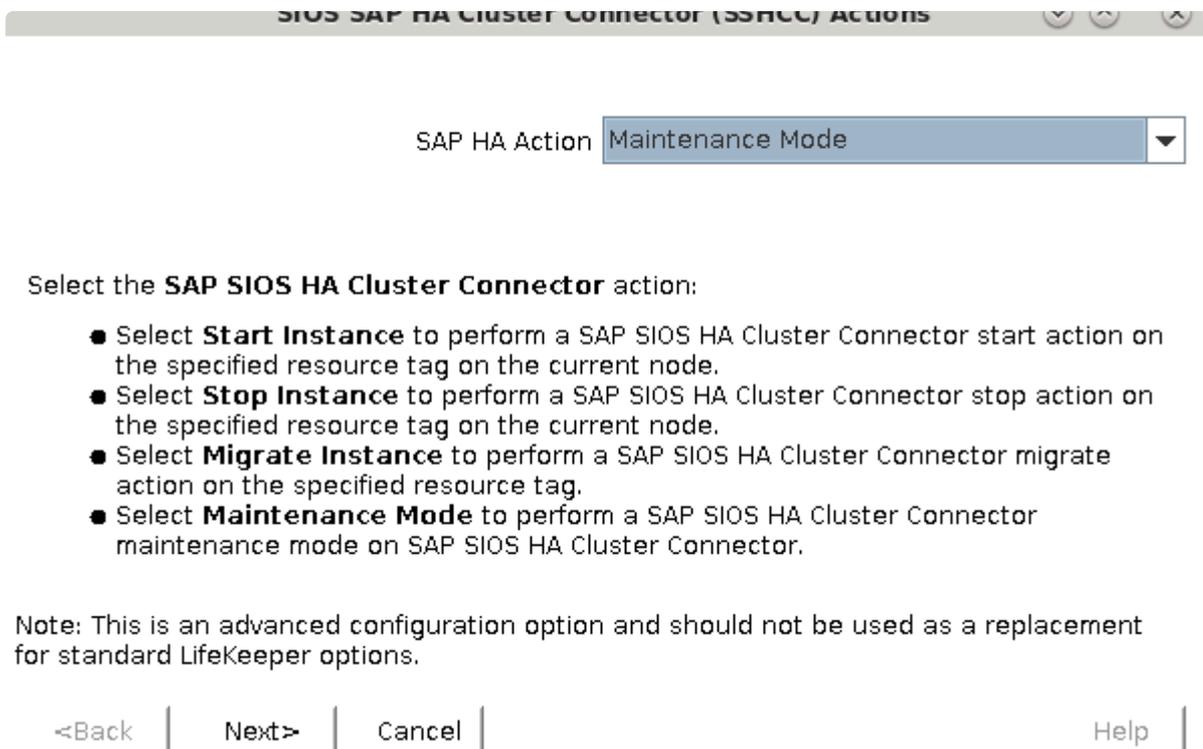
 **Note:** It is not recommended to manually resume any mirror while maintenance mode is enabled for the hierarchy. Since health monitoring, local recovery, and failover are disabled, any failure of the mirror will not be acted on by LifeKeeper.

To place the SAP hierarchies into maintenance mode:

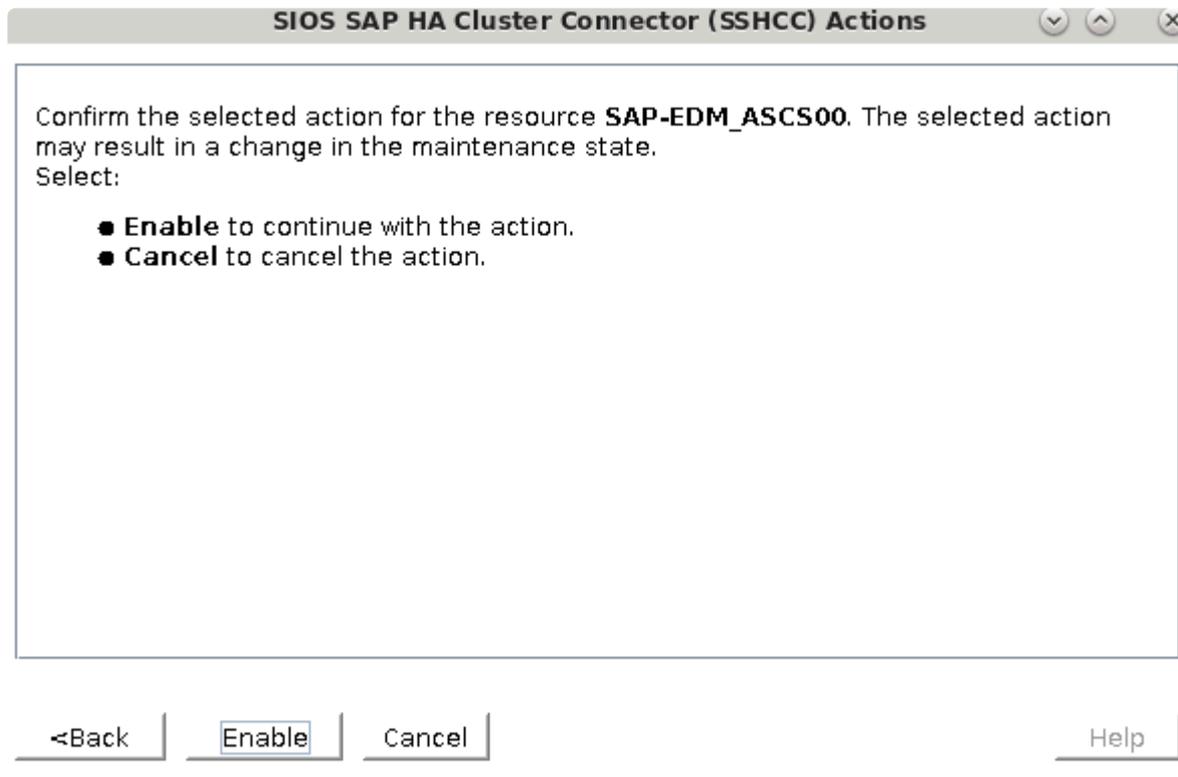
1. **Open the LifeKeeper GUI.**
2. **Right-click** each **SAP** resource on the node where it is currently in-service and click **SSHCC HA Actions**.



3. Select **Maintenance Mode** in the drop-down box.



4. Select **Enable** or **Disable** maintenance mode for all resources in the chosen hierarchy. This action will be attempted on all nodes in the cluster.



\* **Note:** Maintenance mode is enabled or disabled at the hierarchy level, so maintenance mode must be enabled or disabled for each independent SAP resource hierarchy (e.g., the A/SCS and ERS hierarchies) separately.

\* **Note:** In SAP kernel 7.49 PL 200 and later, placing the SAP software into maintenance mode via SAP Update Manager will automatically place the corresponding LifeKeeper SAP hierarchy into maintenance mode via the HA connector function `HASetMaintenanceMode` as long as the SAP SIOS HA Cluster Connector is active for the SAP instance. If LifeKeeper resources are placed into maintenance mode manually via the LifeKeeper GUI, the user will need to refer to the documentation for SAP Update Manager to determine how to place the SAP software itself into maintenance mode before upgrading.

## Enabling and Disabling SAP Maintenance Mode from the Command Line

To enable SAP Maintenance Mode for all resource hierarchies on all servers in the cluster, execute the following command on one of the servers that is hosting LifeKeeper-protected SAP resources:

```
sudo /opt/LifeKeeper/lkadm/subsys/appsuite/sap/bin/lk_maintenance_mode --mode=enable --cluster
```

The optional `--tag` parameter may also be used to specify that only particular resource hierarchies should be placed into maintenance mode. For example, to place only the (A)SCS and ERS resource hierarchies into maintenance mode, execute the following command on one of the servers that is hosting

### LifeKeeper-protected SAP resources:

```
sudo /opt/LifeKeeper/lkadm/subsys/appsuite/sap/bin/lk_maintenance_mode --mode=enable --tag=<A/SCS Resource Tag>,<ERS Resource Tag> --cluster
```

To disable SAP Maintenance Mode for all resources hierarchies on all servers in the cluster, execute the following command on one of the servers that is hosting LifeKeeper-protected SAP resources:

```
sudo /opt/LifeKeeper/lkadm/subsys/appsuite/sap/bin/lk_maintenance_mode --mode=disable --cluster
```

As described above, the optional `--tag` parameter may also be used to disable maintenance mode for specific resource hierarchies. For example:

```
sudo /opt/LifeKeeper/lkadm/subsys/appsuite/sap/bin/lk_maintenance_mode --mode=disable --tag=<A/SCS Resource Tag>,<ERS Resource Tag> --cluster
```

## Checking Maintenance Mode Status for LifeKeeper Protected Resources

The LifeKeeper GUI does not currently show the maintenance mode status for each resource hierarchy. To check the maintenance mode status of an SAP hierarchy in LifeKeeper, run the following command from the command line:

```
sudo /opt/LifeKeeper/lkadm/subsys/appsuite/sap/bin/lk_maintenance_mode --mode=check --tag=<Resource Tag> --cluster
```

where `<Resource Tag>` is the LifeKeeper tag of the SAP resource (e.g., `SAP-SID_ASCS00` or `SAP-SID_ERS10`). The output will show whether maintenance mode is fully enabled, partially enabled, or fully disabled for the hierarchy containing the given resource.

## 6.17.10.2. Custom and Maintenance-Mode Behavior via Policies

---

\* **Note:** As of LifeKeeper version 9.3.2, HA Maintenance Mode can be enabled for an SAP hierarchy via the SSHCC Actions menu in the LifeKeeper UI. See [SAP Maintenance Mode](#) for more details.

### Resource Policy Management

#### Overview

Resource Policy Management in LifeKeeper for Linux provides behavior management of resource local recovery and failover. Resource policies are managed with the **ikpolicy** command line tool (CLI).

#### LifeKeeper

LifeKeeper is designed to monitor individual applications and groups of related applications, periodically performing local recoveries or notifications when protected applications fail. Related applications, by example, are hierarchies where the primary application depends on lower-level storage or network resources. When an application or resource failure occurs, the default behavior is:

1. **Local Recovery:** First, attempt **local** recovery of the resource or application. An attempt will be made to restore the resource or application on the local server without external intervention. If local recovery is successful, then LifeKeeper will not perform any additional action.
2. **Failover:** Second, if a local recovery attempt *fails* to restore the resource or application (or the *recovery kit* monitoring the resource has no support for local recovery), then a **failover** will be initiated. The failover action attempts to bring the application (and all dependent resources) into service on another server within the cluster.

Please see [LifeKeeper Fault Detection and Recovery Scenarios](#) for more detailed information about our recovery behavior.

### Custom and Maintenance-Mode Behavior via Policies

LifeKeeper Version 7.5 and later supports the ability to set additional policies that modify the default recovery behavior. There are four policies that can be set for individual resources (see the section below about precautions regarding individual resource policies) *or* for an entire server. ***The recommended approach is to alter policies at the server level.***

The available policies are:

## Standard Policies

- **Failover** This policy setting can be used to turn on/off resource failover. (**Note:** In order for reservations to be handled correctly, **Failover** cannot be turned off for individual scsi resources.)
- **LocalRecovery** – LifeKeeper, by default, will attempt to recover protected resources by restarting the individual resource or the entire protected application prior to performing a failover. This policy setting can be used to turn on/off local recovery.
- **TemporalRecovery** – Normally, LifeKeeper will perform local recovery of a failed resource. If local recovery fails, LifeKeeper will perform a resource hierarchy failover to another node. If the local recovery succeeds, failover will not be performed.

There may be cases where the local recovery succeeds, but due to some irregularity in the server, the local recovery is re-attempted within a short time; resulting in multiple, consecutive local recovery attempts. This may degrade availability for the affected application.

To prevent this repetitive local recovery/failure cycle, you may set a temporal recovery policy. The temporal recovery policy allows an administrator to limit the number of local recovery attempts (successful or not) within a defined time period.

*Example:* If a user sets the policy definition to limit the resource to three local recovery attempts in a 30-minute time period, LifeKeeper will fail over when a third local recovery attempt occurs *within* the 30-minute period.

Defined temporal recovery policies may be turned *on* or *off*. When a temporal recovery policy is *off*, temporal recovery processing will continue to be done and notifications will appear in the log when the policy *would* have fired; however, no actions will be taken.

**Note:** It is possible to disable failover and/or local recovery with a temporal recovery policy also in place. This state is illogical as the temporal recovery policy will **never** be acted upon if failover or local recovery are disabled.

## Meta Policies

The “meta” policies are the ones that can affect more than one other policy at the same time. These policies are usually used as shortcuts for getting certain system behaviors that would otherwise require setting multiple standard policies.

- **NotificationOnly** – This mode allows administrators to put LifeKeeper in a “monitoring only” state. **Both** local recovery **and** failover **of a resource (or all resources in the case of a server-wide policy) are affected**. The user interface will indicate a Failure state if a failure is detected; *but no recovery or failover action will be taken*. **Note:** The administrator will need to correct the problem that caused the failure manually and then bring the affected resource(s) back in service to continue normal LifeKeeper operations.

## Important Considerations for Resource-Level Policies

Resource level policies are policies that apply to a specific resource only, as opposed to an entire resource hierarchy or server.

*Example :*

```
app
- IP
- file system
```

In the above resource hierarchy, app depends on both IP and file system. A policy can be set to disable local recovery or failover of a specific resource. This means that, for example, if the IP resource's local recovery fails and a policy was set to *disable* failover of the IP resource, then the IP resource will not fail over or cause a failover of the other resources. However, if the file system resource's local recovery fails and the file system resource policy does not have failover disabled, then the entire hierarchy will fail over.

**Note:** It is important to remember that resource level policies apply only to the specific resource for which they are set.

This is a simple example. Complex hierarchies can be configured, so care must be taken when setting resource-level policies.

## The lkpolicy Tool

The `lkpolicy` tool is the command-line tool that allows management (querying, setting, removing) of policies on servers running LifeKeeper for Linux. `lkpolicy` supports setting/modifying policies, removing policies and viewing all available policies and their current settings. In addition, defined policies can be set on or off, preserving resource/server settings while affecting recovery behavior.

The general usage is :

```
lkpolicy [-list-policies | -get-policies | -set-policy | -remove-policy] <name value pair data...>
```

The `<name value pair data...>` differ depending on the operation and the policy being manipulated, particularly when setting policies. For *example*: Most on/off type policies only require `-on` or `-off` switch, but the temporal policy requires additional values to describe the threshold values.

## Example lkpolicy Usage

### Authenticating With Local and Remote Servers

The `lkpolicy` tool communicates with LifeKeeper servers via an API that the servers expose. This API requires authentication from clients like the `lkpolicy` tool. The first time the `lkpolicy` tool is asked to

access a LifeKeeper server, if the credentials for that server are not known, it will ask the user for credentials for that server. These credentials are in the form of a username and password and:

1. Clients must have LifeKeeper admin rights. This means the username must be in the *lkadmin* group according to the operating system's authentication configuration (via pam). It is **not** necessary to run as **root**, but the root user can be used since it is in the appropriate group by default.
2. The credentials will be stored in the *credential store* so they do not have to be entered manually each time the tool is used to access this server.

See [Configuring Credentials for LifeKeeper](#) for more information on the credential store and its management with the `credstore` utility.

An example session with `lkpolicy` might look like this:

```
[root@thor49 ~]# lkpolicy -l -d v6test4
Please enter your credentials for the system 'v6test4'.
Username: root
Password:
Confirm password:
Failover
LocalRecovery
TemporalRecovery
NotificationOnly
[root@thor49 ~]# lkpolicy -l -d v6test4
Failover
LocalRecovery
TemporalRecovery
NotificationOnly
[root@thor49 ~]#
```

## Listing Policies

```
lkpolicy --list-policy-types
```

## Showing Current Policies

```
lkpolicy --get-policies
```

```
lkpolicy --get-policies tag=\*
```

```
lkpolicy --get-policies --verbose tag=mysql\* # all resources starting with mysql
```

```
lkpolicy --get-policies tag=mytagonly
```

## Setting Policies

```
lcpolicy —set-policy Failover —off
```

```
lcpolicy —set-policy Failover —on tag=myresource
```

```
lcpolicy —set-policy Failover —on tag=\*
```

```
lcpolicy —set-policy LocalRecovery —off tag=myresource
```

```
lcpolicy —set-policy NotificationOnly —on
```

```
lcpolicy —set-policy TemporalRecovery —on recoverylimit=5 period=15
```

```
lcpolicy —set-policy TemporalRecovery —on —force recoverylimit=5 period=10
```

## Removing Policies

```
lcpolicy —remove-policy Failover tag=steve
```

**Note:** *NotificationOnly* is a policy alias. Enabling *NotificationOnly* is the equivalent of disabling the corresponding *LocalRecovery* and *Failover* policies.

# 6.18. SAP HANA Recovery Kit Administration Guide

---

\* Beginning in v9.5.0 SIOS has released the new SAP HANA Application Recovery Kit. SIOS will continue to support the SAP HANA gen/app based Recovery Kit with the 9.4.x releases until **March 31, 2022**. If you are using LifeKeeper for Linux v9.5 or later you must use the new (built-in) SAP HANA Application Recovery Kit.

! The existing SAP HANA gen/app based Recovery Kit is **not** supported with v9.5.0. Users who wish to upgrade to the LifeKeeper for Linux v9.5.0 **must** convert their existing SAP HANA gen/app based Recovery Kit to the new SAP HANA Recovery Kit. Refer to [Upgrading from the SAP HANA Gen/App to the SAP HANA Recovery Kit](#) for details.

SAP HANA is an enterprise-class in-memory database system that can be deployed for a wide array of purposes. It forms the basis for the SAP S/4HANA Enterprise Resource Planning platform. The SAP HANA Recovery Kit provides fault resilient protection for SAP HANA databases in a LifeKeeper for Linux environment.

## Document Contents

This guide includes the following topics to help you successfully create and manage your SAP HANA hierarchy:

- [SAP HANA Recovery Kit Requirements](#). Lists the hardware and software necessary to properly set up, install and operate the SAP HANA Recovery Kit.
- [Overview](#). Describes the SAP HANA Recovery Kit's features and functionality.
- [Configuring SAP HANA with LifeKeeper](#). Provides instructions for installing and configuring the SAP HANA software.
- [Resource Configuration Tasks](#). Describes the various functions you may perform on your hierarchies using the LifeKeeper GUI: create, extend, delete and unextend.
- [Hierarchy Administration](#). Provides important recommendations for ongoing administration of the SAP HANA hierarchy.
- [Troubleshooting](#). Lists and describes the error messages associated with the SAP HANA Recovery Kit.

## LifeKeeper Documentation

The following LifeKeeper product documentation is available from the SIOS Technology Corp. website:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)

- [LifeKeeper for Linux Installation Guide](#)
- [Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is provided on the SIOS Technology Corp. website at:

<http://docs.us.sios.com/>

and from the Help menu in the LifeKeeper GUI.

## **SAP HANA Documentation**

Documentation for SAP HANA can be found at:

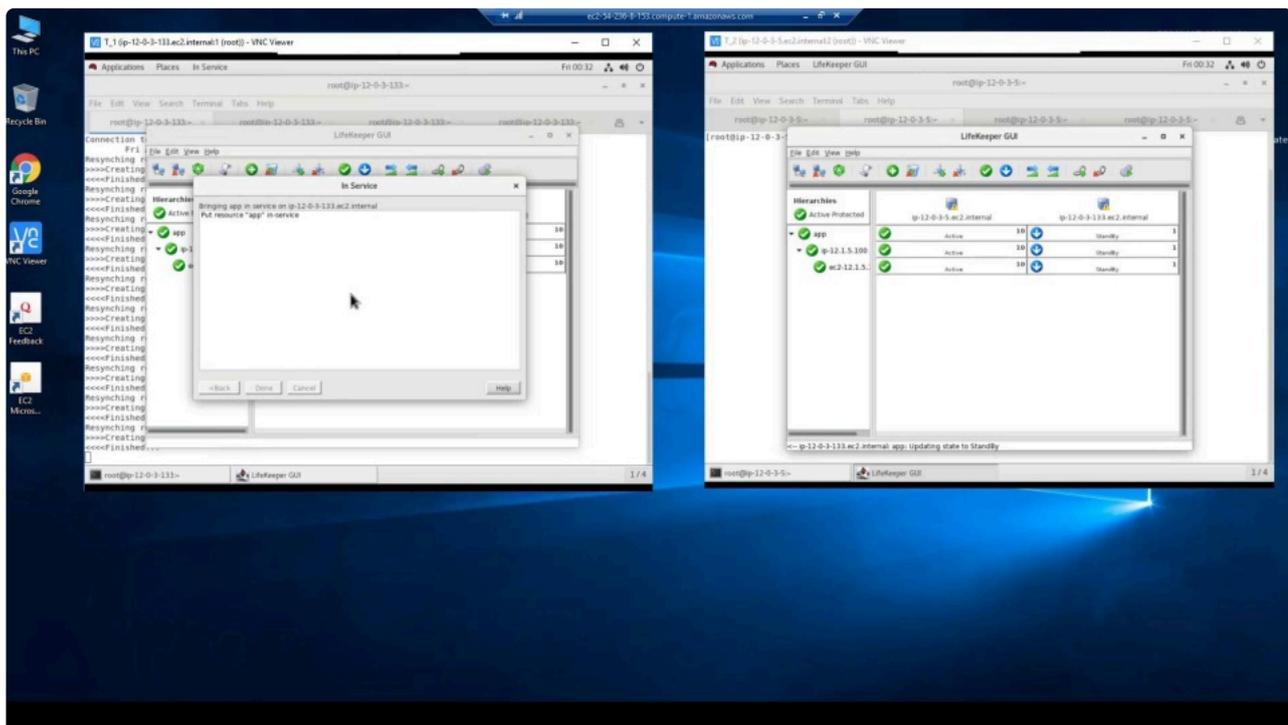
<http://help.sap.com>

# 6.18.1. Upgrading from the SAP HANA Gen/ App to the SAP HANA Recovery Kit

✿ Beginning in v9.5.0 SIOS has released the new SAP HANA Application Recovery Kit. SIOS will continue to support the SAP HANA gen/app based Recovery Kit with the 9.4.x releases until **March 31, 2022**. If you are using LifeKeeper for Linux v9.5 or later you must use the new (built-in) SAP HANA Application Recovery Kit.

! The existing SAP HANA gen/app based Recovery Kit is **not** supported with v9.5.0. Users who wish to upgrade to the LifeKeeper for Linux v9.5.0 **must** convert their existing SAP HANA gen/app based Recovery Kit to the new SAP HANA Recovery Kit. Please follow the upgrade steps below.

✿ Please refer to the [Support Matrix](#) before upgrading.



<https://fast.wistia.net/embed/iframe/aggjiny6n1>

## Upgrade LifeKeeper

1. Upgrade LifeKeeper to the latest version on all nodes in the cluster (refer to [Upgrading LifeKeeper](#)).
2. Install the **SAP HANA ARK** and the **required** corresponding license on all nodes in the cluster.

## Remove SAP HANA Gen/App Resources

1. LifeKeeper should be running on all nodes and the gen/app resource hierarchy should be in service.
2. Backup the LifeKeeper configuration in case the resource hierarchies need to be restored to the previous settings. To perform the backup execute the following command:

```
/opt/LifeKeeper/bin/lkbackup -c --cluster
```

By default the backup file is located at `/opt/LifeKeeper/config/archive.<date – yyyymmddhhmm>.tar.gz`.

3. Right click the **SAP HANA gen/app** resource hierarchy in the left panel.
4. Choose **Delete Dependency...** if there are any dependencies attached to the gen/app resource (e.g., Virtual IP or AWS EC2 EIP).
5. After the dependencies are removed, right click the **SAP HANA gen/app** resource.
6. Choose **Delete Resource Hierarchy...** to remove the old SAP HANA gen/app resource.

 Please note that after the SAP HANA genapp resource is removed, there will be no LifeKeeper Failover or Monitoring of the SAP HANA instance until a new SAP HANA resource is created. HANA System Replication will remain active unless changes outside of LifeKeeper cause it to terminate.

## Create and Extend SAP HANA Resource

Create/Extend a SAP HANA resource using the newly installed SAP HANA ARK.

1. To create and extend a **SAP HANA resource**, follow the steps in [Creating an SAP HANA Resource Hierarchy](#) and [Extending an SAP HANA Resource Hierarchy](#).
2. Once an **SAP HANA resource** hierarchy is created and extended, if you want to create or re-create any dependencies that were removed in Step 4 above, perform the following steps:
  - a. Right click the **SAP HANA resource** hierarchy in the left panel.
  - b. Choose **Create Dependency...** to create dependencies in the new SAP HANA resource (e.g., Virtual IP or AWS EC2 EIP) that were attached to the SAP HANA gen/app resource.
3. Verify/test the new SAP HANA resource hierarchy by performing the tests found [here](#).

## Remove HANA Gen/App Files

To remove the gen/app package from each system, execute the following command:

```
rpm -e steeleye-lkHOTFIX-HANA-SP1-9.1.0-6538.noarch
```

## 6.18.2. SAP HANA Supported Configurations

**\* Note:** The operating system and configuration used in your SAP HANA deployment must be supported by SAP, SIOS, and your infrastructure provider (public/private cloud or on-premise). Consult SAP’s Product Availability Matrix (PAM) and SAP Note 2235581 (SAP HANA: Supported Operating Systems) in order to determine which operating systems are supported for various versions of SAP HANA.

Package	Version	OS/Application Version Support
<a href="#">LifeKeeper Core (SAP HANA)</a>	9.6.1	RHEL 7.6 RHEL 7.7 (HANA 2.0 SPS04 rev 48 and newer) RHEL 7.9 (HANA 2.0 SPS05 rev 54 and newer) RHEL 8.0 (HANA 2.0 SPS04 only, rev 40 and newer) RHEL 8.1 (HANA 2.0 SPS04 rev 45 and newer) RHEL 8.2 (HANA 2.0 SPS04 rev 48.02, HANA 2.0 SPS05 rev 52 and newer) RHEL 8.4 (HANA 2.0 SPS05 rev 55 and newer) SLES 12 SP4 SLES 12 SP5 (HANA 2.0 SPS04 rev 45 and newer) SLES 15 (GA) SLES 15 SP1 (HANA 2.0 SPS04 rev 44 and newer) SLES 15 SP2 (HANA 2.0 SPS04 rev 48.01 and newer) SLES 15 SP3 (HANA 2.0 SPS05 rev 55 and newer)
<a href="#">LifeKeeper SAP HANA Recovery Kit</a>	9.6.1	SAP HANA 2.0 SPS04, SPS05, SPS06  Versions prior to SAP HANA 2.0 SPS04 are not supported by the SAP HANA Recovery Kit.

**\* NOTE:** Operating system versions built for enhanced SAP support (such as Red Hat Enterprise Linux for SAP Business Applications, Red Hat Enterprise Linux for SAP Solutions, Red Hat Enterprise Linux for SAP HANA, and SUSE Linux Enterprise Server for SAP Applications) are also supported as long as the running Linux kernel version is the same as one of the supported OS versions listed above.

Supported configurations for protecting application server components of SAP NetWeaver or S/4HANA can be found on the [LifeKeeper for SAP Solution Page](#).

## 6.18.3. SAP HANA Recovery Kit Hardware and Software Requirements

\* Beginning in v9.5.0 SIOS has released the new SAP HANA Application Recovery Kit. SIOS will continue to support the SAP HANA gen/app based Recovery Kit with the 9.4.x releases until **March 31, 2022**. If you are using LifeKeeper for Linux v9.5 or later you must use the new (built-in) SAP HANA Application Recovery Kit.

! The existing SAP HANA gen/app based Recovery Kit is **not** supported with v9.5.0. Users who wish to upgrade to the LifeKeeper for Linux v9.5.0 **must** convert their existing SAP HANA gen/app based Recovery Kit to the new [SAP HANA Recovery Kit](#). Refer to [Upgrading from the SAP HANA Gen/App to the SAP HANA Recovery Kit](#) for details.

### Hardware Requirements

- **Servers:** Servers should be configured in accordance with the requirements described in the SAP HANA Master Guide, SAP notes mentioned in this guide, LifeKeeper for Linux Documentation, and the LifeKeeper for Linux Release Notes.
- **Storage:** For SAP HANA databases utilizing SAP HANA System Replication, no shared storage is necessary. Special storage requirements are given in the SAP HANA Master Guide and the aforementioned SAP notes.

### Software Requirements

- **LifeKeeper Software:** It is imperative that you install the same version of LifeKeeper software and apply the same versions of the LifeKeeper software patches to each server in your cluster.
- **SAP HANA Software:** SAP HANA Platform Edition 2.0 SP04 (or later) is required for use of the SAP HANA Recovery Kit. The same version of the SAP HANA software must be installed and configured on each server before configuring LifeKeeper and the LifeKeeper SAP HANA Recovery Kit. Since SAP HANA licenses are tied to hardware, each server will require its own license from SAP. SAP HANA System Replication must be enabled and fully configured and the database must be running on all cluster nodes before creating the SAP HANA resource hierarchy in LifeKeeper.
- **LifeKeeper SAP HANA Recovery Kit:** The SAP HANA Recovery Kit is bundled as an optional recovery kit with the core installation in LifeKeeper for Linux v9.5.0 and later.

\* The SAP HANA Recovery Kit supports only two-node configurations. It does not support three-node configurations.

- **Witness Server:** SIOS recommends adding a 3rd node to a 2-node configuration as a witness server.

## 6.18.4. SAP HANA Recovery Kit Overview

---

The SAP HANA Recovery Kit is compatible with SAP HANA Platform 2 (SP04 or later).

SAP HANA provides three different mechanisms to increase availability.

- **Host Auto-Failover** – At least one standby node added to a SAP HANA system. These nodes are configured to work in standby mode. If the required processes or databases are not active, LifeKeeper will attempt to restart them. In case of an unsuccessful restart of the processes on the primary node, LifeKeeper will attempt to bring the database in-service on the backup node, register that node as primary master in SAP HANA System Replication, and register the previous primary node as the secondary replication site. If the previous primary node cannot be configured as the secondary SAP HANA System Replication site, the resource will be marked as Failed (OSF) on that node until the problem is corrected and it can be successfully registered. Once the previous primary node has been successfully registered as a secondary replication site, LifeKeeper will update the state of the SAP HANA resource on the node to Standby (OSU).
- **Storage Replication** – The storage used on the primary SAP HANA node replicates all data to another SAP HANA node. This replication works without a control process from the SAP HANA system. The storage replication is provided by hardware partners.
- **System Replication** – SAP HANA replicates all data from the primary SAP HANA node to a backup node by use of SAP's own built-in replication framework. Data is constantly pre-loaded on the secondary SAP HANA node.

With the SAP HANA Recovery Kit, SAP HANA systems, utilizing System Replication, can be protected and administered through SIOS LifeKeeper.

The Recovery Kit is able to start the SAP HANA system on all nodes and perform the takeover and replication site registration processes of SAP HANA System Replication. To ensure the functionality of the SAP HANA system, the following processes and states are continuously monitored:

- SAP Host Agent on all nodes
- SAP Start Service (sapstartsrv) of HDB instance on all nodes
- State of SAP HANA database on all nodes
- State of SAP HANA System Replication mode (primary on active node, sync|syncmem|async on secondary node)

If the required processes or databases are not active, LifeKeeper will attempt to restart them. In case of an unsuccessful restart of the processes on the primary node, LifeKeeper will attempt to bring the database in-service on the backup node, register that node as primary master in SAP HANA System Replication, and register the previous primary node as the secondary replication site. If the previous primary node cannot be configured as the secondary SAP HANA System Replication site, the resource will be marked as Failed (OSF) on that node until the problem is corrected and it can be successfully registered. Once the previous primary node has been successfully registered as a secondary replication site, LifeKeeper will update the state of the SAP HANA resource on the node to Standby (OSU).

In case of an invalid state of the SAP HANA System Replication, the SAP HANA resource is also placed in the state “Out of Service – Faulty” (OSF). It has to be decided with the help of a database administrator whether a takeover is to be performed or how the SAP HANA System Replication mode should be corrected.

When carrying out the “Out of Service” action for an SAP HANA resource in LifeKeeper, only the database on the primary node is stopped by default. The database on the secondary node remains active and retains its SAP HANA System Replication mode.

## 6.18.4.1. SAP HANA GUI States

The active (ISP) resource can have the following states. Some of these warning and failure states are transient and may appear while LifeKeeper is attempting to recover required processes or before LifeKeeper initiates a failover of the HANA resource hierarchy.

GUI Text	GUI Properties State	Description	Icon
Active	Active	Sapstartsrv and the HDB instance are running properly. (State Name: ISP)	
Active – sapstartsrv Failure	Sapstartsrv is not running, HDB status unknown	Sapstartsrv is not running. The HDB instance status is unknown since sapstartsrv is used to check the status of HDB. (State Name: ISPSapStartSrvNotRunning)	
Active – HSR Disabled	HSR is running in replication mode 'none'	Sapstartsrv and HDB instance are running properly. The replication mode is 'none', indicating that SAP HANA Replication is disabled. (State Name: ISPHSRDisabled)	
Active – HDB Stopped	Sapstartsrv is running, HDB is not running properly	Sapstartsrv is running. The HDB instance is not running properly. (State Name: ISPHDBNotRunning)	
Active – Secondary	HSR is running in a secondary replication mode	Sapstartsrv and the HDB instance are running properly but the active node is registered as a secondary replication site. In normal operation, the active node is expected to be the primary replication site. (State Name: ISPSecondary)	
Active – Unknown Repl Mode	Replication mode cannot be determined	Sapstartsrv and HDB instance are running properly. The replication mode cannot be determined using 'hdbnsutil -sr_state' or it returns an unsupported mode. (State Name: ISPUknownReplMode)	
Active – Suspended	Database is primary and suspended	The database is suspended on the server where the SAP HANA resource is currently in-service. This is an unexpected state, and typically occurs when a user has performed a "takeover with handshake" outside of LifeKeeper. (State Name: ISPSuspended)	

The state shown in the GUI for the standby (OSU) resource is affected by the active (ISP) resource restore and quickCheck processes. The standby resource can take several minutes to get into its final state during the in-service operation, during which transient intermediate resource states may appear. This is normal and expected operation for SAP HANA resources. Some of the warning and failure states

shown below are transient and may appear while LifeKeeper is attempting to recover required processes or re-register the standby server as a secondary replication site. If the SAP HANA database cannot be successfully started or registered as a secondary replication site on the backup server, the resource state on the backup server will transition to Failed (OSF). The standby resource can have the following states:

GUI Text	GUI Properties State	Description	Icon
Standby	No cluster HANA resource ISP, HDB is not running	This is the normal status for a resource that is out-of-service when an equivalent resource is not in-service on another node. (State Name: OSUStopped)	
Standby – HDB Running	HDB is running, HSR monitoring inactive	Sapstartsrv and the HDB instance are running on the standby (OSU) node. There is either no active (ISP) node, or quickCheck has not determined the HSR status on the active node. (State Name: OSUHDBRunning)	
Standby – In Sync	HSR active and in-sync	Sapstartsrv and the HDB instance are running with HSR configured and reporting in-sync. (State Name: SecondaryActive)	
Standby – sapstartsrv Failure	sapstartsrv is not running, HDB status unknown	Sapstartsrv is not running. There is an active (ISP) node. The HDB instance status is unknown since sapstartsrv is used to check the status of HDB. (State Name: OSUSapStartSrvNotRunning)	
Standby – Unknown Repl. Mode	Replication mode cannot be determined	Sapstartsrv and the HDB instance are running. There is an active (ISP) node. The replication mode cannot be determined using 'hdbnsutil -sr_state'. (State Name: OSUUnknownReplMode)	
Standby – HSR Disabled	Replication mode is 'none'	Sapstartsrv and the HDB instance are running. There is an active (ISP) node. Either the replication mode is 'none' according to 'hdbnsutil -sr_state', or the HANA utility systemReplicationStatus.py returned '10'. (State Name: OSUHSRDisabled)	
Standby – HSR Error	Active node HSR reports error (11)	Sapstartsrv and the HDB instance are running. There is an active (ISP) node. The HANA utility systemReplicationStatus.py returned '11'. (State Name: OSUHSRError)	
Standby – Initializing	HSR Initializing	Sapstartsrv and the HDB instance are running. There is an active (ISP) node. HANA utility systemReplicationStatus.py returned '13'. (State Name: SecondaryInitializing)	
Standby – Syncing	HSR Synchronizing	Sapstartsrv and the HDB instance are running. There is an active (ISP) node. The HANA utility systemReplicationStatus.py returned '14'. (State Name: SecondarySyncing)	

Standby – HDB Stopped	HANA resource ISP in cluster, HDB is not running	Sapstartsrv is running. The HDB instance is not running. There is an active (ISP) node. (State Name: OSUHDBNotRunning)	
Standby – Primary	HSR is running in replication mode 'Primary'	Sapstartsrv and the HDB instance are running properly, but the standby node is registered as a primary replication site. In normal operation, the standby node is expected to be a secondary replication site. (State Name: OSUPrimary)	
Standby – Unknown HSR Status	Active node HSR reports error (12)	Sapstartsrv and the HDB instance are running. There is an active (ISP) node. The HANA utility systemReplicationStatus.py returned '12' . (State Name: SecondaryUnknownHSRStatus)	
Standby – Unexpected HSR state	HSR status from ISP node not in expected range 10 – 15	Sapstartsrv and the HDB instance are running. There is an active (ISP) node. The HANA utility systemReplicationStatus.py returned a value that was not in the expected range of 10 to 15 . (State Name: OSUUnexpectedHSRState)	
Standby – Suspended	Database is primary and suspended	The database is suspended on a server where the SAP HANA resource is currently out of service. This is an expected transient state which typically appears during the “takeover with handshake” process while the suspended database is being stopped and re-registered as a secondary replication site. (State Name: OSUSuspended)	

## 6.18.4.2. SAP HANA Resource Hierarchy

The following example shows a typical SAP HANA resource hierarchy:



The child resource **vip-sps-hana** in this example is protecting the switchable IP address used for client connections to the database.

In the event of failover, LifeKeeper will bring the IP address and database resource in service on a backup server. Any database transaction which has not yet been committed will need to be run again.

## 6.18.4.3. Multitenant Database Support

---

The SAP HANA Recovery Kit is able to function in an environment with multiple tenant databases. However, please be aware of the following behaviors and special considerations:

- SAP HANA System Replication (HSR) must be able to successfully initialize replication for **all** tenant databases on the system in order for the HSR monitoring in the recovery kit to function correctly. This means, for example, that backups must be taken of all tenant databases before enabling HSR. If HSR is out-of-sync for any of the tenant databases contained within the protected instance, the SAP HANA Recovery Kit will leave the !HANA\_DATA\_OUT\_OF\_SYNC flag on the standby system, preventing switchover or failover of the SAP HANA resource hierarchy.
- The SAP HANA Recovery Kit currently does not restart individual tenant databases if they are manually stopped (e.g., with an HDBSQL 'alter system stop database DB\_NAME' command). In this case, the SAP HANA software removes the corresponding indexserver process from the list of processes for the HDB instance, and the SAP HANA Recovery Kit will treat the manual stoppage of the tenant database as intentional.
- The SAP HANA Recovery Kit currently does not monitor the tenant databases at an individual level. It monitors the entire HDB instance (which will contain indexserver processes for each running tenant database) as a whole. If any process in the HDB instance has failed, the recovery kit will attempt to restart the entire instance. This means, for example, that the indexserver process for TENANT1 may be restarted unexpectedly in the case that the indexserver process for TENANT2 fails.

## 6.18.5. Configuring SAP HANA with LifeKeeper

---

The following sequence is recommended for installing and configuring the SAP HANA database and LifeKeeper software. Each of these steps links to detailed tasks that follow.

[Install the SAP HANA Software](#)

[Configure SAP HANA System Replication](#)

[Modify the SAP HANA Instance Profile](#)

[Install the LifeKeeper Software](#)

After you have performed these tasks, you will be ready to create the LifeKeeper resource hierarchy to protect your SAP HANA database.

## 6.18.5.1. Install the SAP HANA Software

---

Install the SAP HANA software on all servers in the cluster using identical parameters/settings. In particular, the same SAP System ID (SID) and instance number must be used on all systems. Refer to the SAP HANA Master Guide for installation details.

## 6.18.5.2. Configure SAP HANA System Replication

---

Configure SAP HANA System Replication according to the instructions provided in the SAP HANA System Replication Guide (available at <http://help.sap.com>). Once SAP HANA System Replication has been successfully enabled on the intended primary replication site and the backup server has been successfully registered as a secondary replication site, continue with the rest of the steps in this guide.

 **Note:** In an environment with multiple tenant databases, SAP HANA System Replication must be able to successfully initialize replication for all tenant databases in order for the SAP HANA Recovery Kit to function correctly. In particular, backups must be taken of all tenant databases contained within the protected instance before enabling system replication.

## 6.18.5.3. Modify the SAP HANA Instance Profile

---

### Disable Autostart for the HDB Instance

In order for LifeKeeper to successfully manage the SAP HANA resource during failover or system reboot, the Autostart feature which automatically starts the HDB instance on system boot must be disabled. In order to disable this feature, edit the instance profile for your HDB instance (typically located at `/usr/sap/<SID>/SYS/profile/<SID>_HDB<##>_<HostName>`) **on all cluster nodes** and ensure that the line

```
Autostart = 0
```

is present in the profile. Either add or modify this line, as necessary, and save the changes to the profile.

**!** Warning: If Autostart is not disabled for the HDB instance, then a machine failure of the primary SAP HANA System Replication site will result in a “System Replication split brain” scenario once the original primary node comes back online. In this scenario, the HDB instance is running and registered as primary master concurrently on multiple cluster nodes. As a result, LifeKeeper is unable to determine which site the user intends to be registered as the primary SAP HANA System Replication site. The SAP HANA resource is placed in the OSF (“Out of Service – Failed”) state on the original standby node and a warning message is broadcast to all open consoles until the situation is manually resolved by a database administrator. See [Resolving Split Brain Scenarios](#) for more details.

**\* Note:** In order to ensure continued successful management of the SAP HANA resource, LifeKeeper will monitor the value of the Autostart parameter for the HDB instance on all cluster nodes each quickCheck interval (defined by the LKCHECKINTERVAL parameter, default 2 minutes) and will automatically disable Autostart if necessary.

## 6.18.5.4. Install the LifeKeeper Software

---

Once you have installed the SAP HANA software and configured SAP HANA System Replication, you are ready to install the LifeKeeper Core software and any required patches followed by the SAP HANA Recovery Kit.

Refer to the [LifeKeeper for Linux Installation Guide](#) for details on installing the LifeKeeper packages.

## 6.18.6. SAP HANA Resource Configuration Tasks

---

Once you have completed the setup tasks described in the previous section, you are ready to create and extend your SAP HANA resource hierarchies.

The following tasks are available for configuring the LifeKeeper for Linux SAP HANA Recovery Kit:

- **Create Resource Hierarchy** – Creates an SAP HANA resource hierarchy
- **Delete Resource Hierarchy** – Deletes an SAP HANA resource hierarchy
- **Extend Resource Hierarchy** – Extends an SAP HANA resource hierarchy from the primary server to the backup server
- **Unextend Resource Hierarchy** – Unextends (removes) an SAP HANA resource hierarchy from a single server in the LifeKeeper cluster

Please refer to your [LifeKeeper for Linux Technical Documentation](#) for additional instructions on configuring your LifeKeeper resource hierarchies.

The following tasks are described in the [Administration](#) section within the LifeKeeper for Linux Technical Documentation because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View / Edit Properties](#). View or edit the properties of a resource hierarchy on a specific server.

**Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu.

You may also perform most of the tasks:

- from the toolbar
- by right-clicking on a global resource in the left pane of the status display
- by right-clicking on a resource in the right pane of the status display

Using the right-click method allows you to avoid entering certain information that is required when using the **Edit** menu.

## 6.18.6.1. Creating an SAP HANA Resource Hierarchy

**! Important Note:** Before creating your SAP HANA resource hierarchy, SAP HANA System Replication must be enabled and fully configured and the database must be running on all servers in the cluster. See [Configure SAP HANA System Replication](#) for details.

Perform the following steps on the primary server:

1. From the **Edit** menu, select **Server**, then **Create Resource Hierarchy**. The **Create Resource Wizard** dialog will appear.
2. Select **SAP HANA** from the drop-down list and click **Next**.



Please Select Recovery Kit

3. Select *intelligent* as the **Switchback Type** to be used for the SAP HANA resource. Click **Next**.



Switchback Type

Intelligent Switchback means that after a failover to the backup server, an administrator must manually switch the SAP HANA resource back to the primary server. **CAUTION:** This release of LifeKeeper does not support Automatic Switchback for SAP HANA resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource with a dependency on a SAP HANA resource.

4. Select the **Server** on which the SAP HANA resource will be created. When creating the initial SAP HANA resource, the server chosen at this step must be the one on which the database is currently registered as primary master in SAP HANA System Replication. Click **Next**.



Server

5. Select the **SAP HANA SID** under which the SAP HANA database is installed. Click **Next**.



SAP HANA SID

- 6. Select the **SAP HANA** Instance to be protected by LifeKeeper. Click **Next**.

SAP HANA Instance for SPS

- 7. Select the **Dependent Virtual IP** resource to be protected by LifeKeeper. The IP resource must already exist and be in-service on the selected server in order to appear in the list. Select *none* if switching over of the virtual IP/host name on failover or switchover will be managed without using a LifeKeeper IP resource. Click **Next**.

IP child resource

none

- 8. Select whether to **Enable or Disable Local Recovery** for the SAP HANA resource on the chosen server. In some situations when using SAP HANA System Replication, especially when protecting a large database instance, after detection of a failure on the primary database host it may be faster to switch over to the running standby database instance rather than attempt to stop and restart the database on the current primary host. In order to disable all local recovery attempts for the SAP HANA resource on the chosen server, select **Disable**. Otherwise, select **Enable**.

**Create SAP HANA Database Resource** [X]

Enable/Disable Local Recovery

Select whether to enable or disable local recovery for SID **SPS** and instance **HDB00** on server **hana2-1**.

When using SAP HANA System Replication and protecting a large database instance, it is possible that the time required to fail over to the running secondary database may be shorter than the time required to repair and restart the database locally. In situations like this it may be preferable to skip any local recovery attempts and immediately initiate failover when a database failure is detected on the primary host.

In order to disable all local recovery attempts for the protected database on server hana2-1, select **Disable**. Otherwise, select **Enable**.

<Back    Next>    Cancel    Help

- 9. Enter the **SAP HANA Resource Tag** to be used to identify the SAP HANA resource which will be created on the chosen server. The default tag name has the form HANA-<SID>\_<HDB Instance>, but can be modified to use a different tag name as long as it is not currently being used to identify another LifeKeeper resource. Valid characters in tag names are letters, digits, and the following

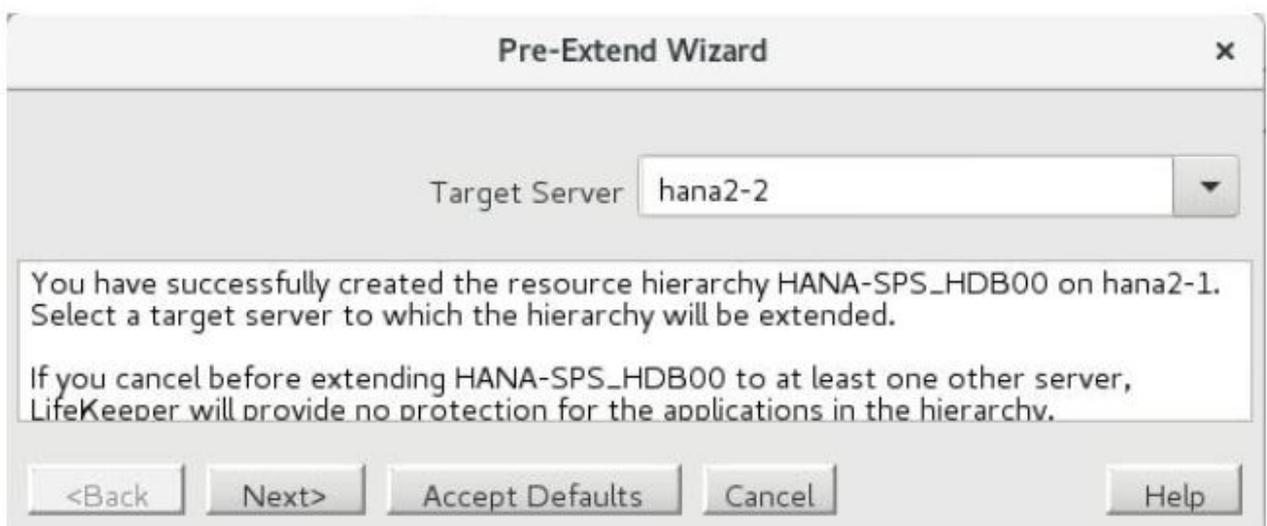
special characters: “-”, “\_”, “.”, and “/”.



- 10. Click **Create**. The **Create Resource Wizard** will then create your SAP HANA resource hierarchy. LifeKeeper will validate the data entered as well as data obtained from the SAP HANA System Replication framework. If a problem is detected, an error message will appear in the information box. Click **Next**.



- 11. You should see a message indicating that you have successfully created an SAP HANA resource hierarchy and you must extend that hierarchy to another server in your cluster to achieve failover protection. Click **Next**.



- 12. Click **Continue**. LifeKeeper will then launch the **Pre-Extend Wizard**. Refer to [Extending an SAP HANA Resource Hierarchy](#) for details on how to extend your resource hierarchy to another server.

## 6.18.6.2. Extending an SAP HANA Resource Hierarchy

**! Important Note:** Before extending your SAP HANA resource hierarchy, SAP HANA System Replication must be enabled and fully configured and the database must be running on all servers in the cluster. See [Configure SAP HANA System Replication](#) for details.

This operation can either be performed from the **Edit** menu or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should begin with Step 2 below.

1. From the **Edit** menu, select **Resource**, then **Extend Resource Hierarchy**. The **Pre-Extend Wizard** appears. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. Select the **Template Server** from which you want to extend an existing SAP HANA resource. (*This dialog box will not appear if you selected the **Extend** task by right-clicking on a resource instance in the right pane.*) Click **Next**.

Template Server

3. Select the **Tag to Extend** for the SAP HANA resource that you would like to extend. (*This dialog box will not appear if you selected the **Extend** task by right-clicking on a resource instance in either pane.*) Click **Next**.

Tag to Extend

4. Select *intelligent* for the **Switchback Type** to be used for the SAP HANA resource. Click **Next**.

Switchback Type

Intelligent Switchback means that after a failover to the backup server, an administrator must manually switch the SAP HANA resource back to the primary server. **CAUTION:** This release of LifeKeeper does not support Automatic Switchback for SAP HANA resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource with a dependency on a SAP HANA resource.

5. Select or enter the **Template Priority**. Click **Next**.

A screenshot of a web form element labeled "Template Priority". It consists of a text input field containing the number "1" and a small downward-pointing arrow icon on the right side, indicating a dropdown menu.

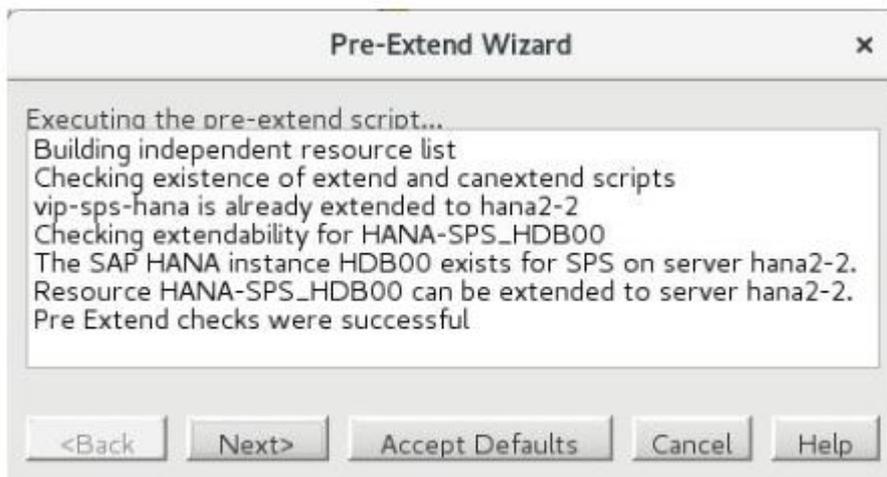
This is the priority for the SAP HANA hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. **Note:** This selection will appear only for the initial extend of the hierarchy.

- 6. Select or enter the **Target Priority**. Click **Next**.

A screenshot of a web form element labeled "Target Priority". It consists of a text input field containing the number "10" and a small downward-pointing arrow icon on the right side, indicating a dropdown menu.

This is a priority for the newly extended SAP HANA hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. Note that LifeKeeper assigns priority 1 to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.

- 7. After receiving the message that the pre-extend checks were successful, click **Next**.

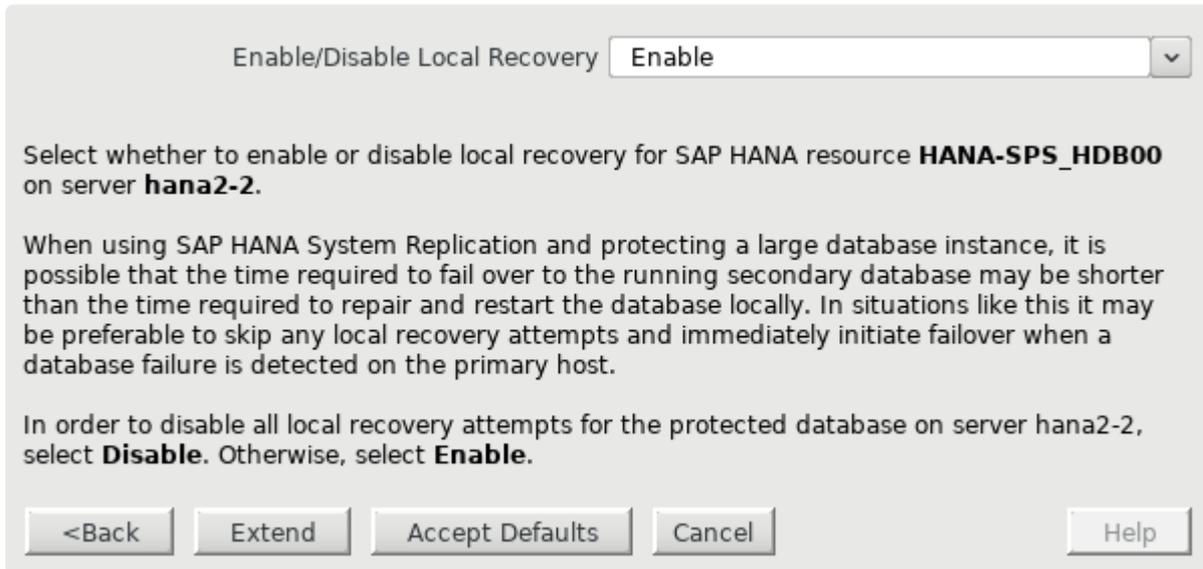


- 8. The **Extend Wizard** will prompt you to enter the **Root Tag** for the SAP HANA resource. This is the unique name used by LifeKeeper to identify the equivalent SAP HANA resource being created on the target server. **Note:** The SAP HANA resource tag name is required to be the same across all cluster servers, so it cannot be edited in this dialog box.

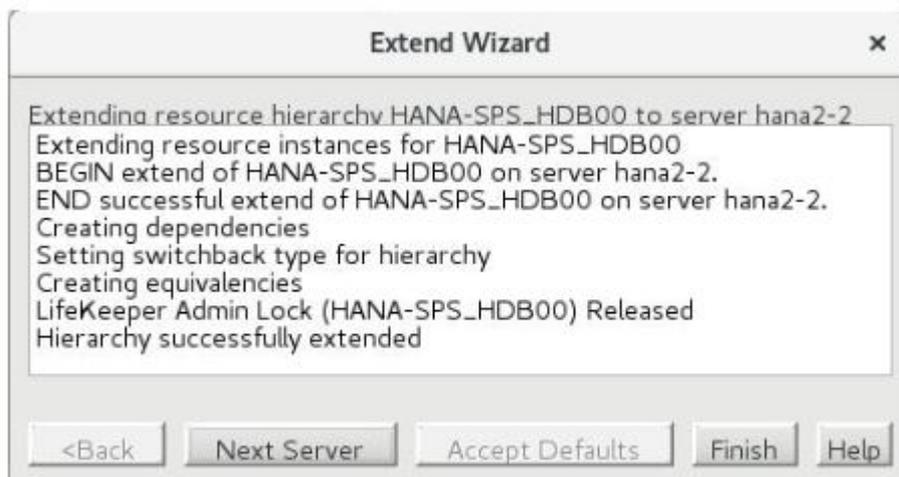
A screenshot of a web form element labeled "Root Tag". It consists of a text input field containing the text "HANA-SPS\_HDB00".

- 9. Select whether to **Enable or Disable Local Recovery** for the SAP HANA resource on the chosen

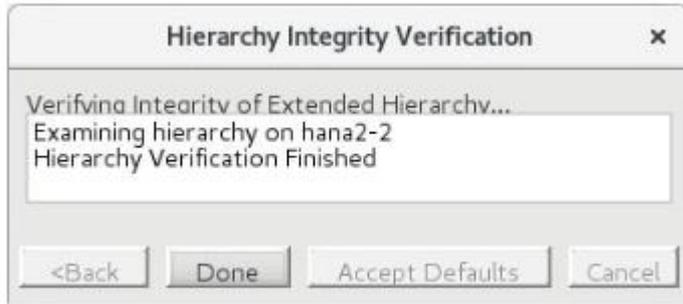
server. In some situations when using SAP HANA System Replication, especially when protecting a large database instance, after detection of a failure on the primary database host it may be faster to switch over to the running standby database instance rather than attempt to stop and restart the database on the current primary host. In order to disable all local recovery attempts for the SAP HANA resource on the chosen server, select **Disable**. Otherwise, select **Enable**.



- Click **Extend**. The **Extend Wizard** will then extend your SAP HANA resource hierarchy to the target server. If a problem is detected, an error message will appear in the information box.



- After receiving the message "Hierarchy successfully extended", click **Next Server** to extend the hierarchy to another server, or click **Finish** if there are no other extend operations to perform.
- After receiving the message "Hierarchy Verification Finished", click **Done**.



## 6.18.6.3. Unextending an SAP HANA Resource Hierarchy

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. From the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the SAP HANA resource. It cannot be the server where the resource is currently in service. (*This dialog box will not appear if you selected the **Unextend** task by right-clicking on a resource instance in the right pane.*) Click **Next**.



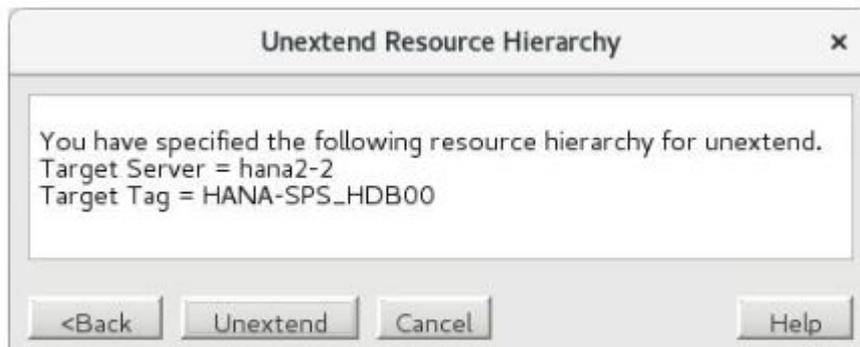
Target Server hana2-2

3. Select the SAP HANA hierarchy to unextend and click **Next**. (*This dialog box will not appear if you selected the **Unextend** task by right-clicking on a resource instance in either pane.*)



Hierarchy to Unextend HANA-SPS\_HDB00

4. An information box appears confirming the target server and the SAP HANA resource hierarchy you have chosen to unextend. Click **Unextend**.



**Unextend Resource Hierarchy** [X]

You have specified the following resource hierarchy for unextend.  
Target Server = hana2-2  
Target Tag = HANA-SPS\_HDB00

<Back Unextend Cancel Help

5. Another information box appears confirming that the SAP HANA resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.



## 6.18.6.4. Deleting an SAP HANA Resource Hierarchy

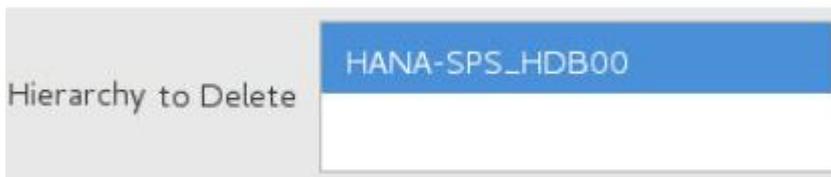
To delete SAP HANA resource from all servers in your LifeKeeper cluster environment, complete the following steps:

1. From the Edit menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your SAP HANA resource hierarchy. (*This dialog will not appear if you selected the **Delete Resource** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance.*)



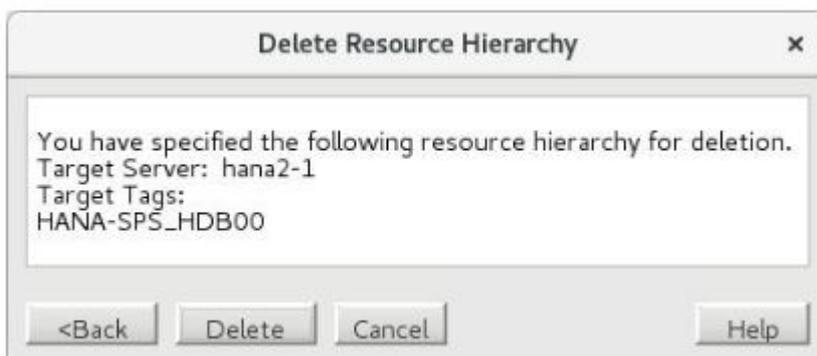
Target Server hana2-1

3. Select the **Hierarchy to Delete**. (*This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a resource instance in the left or right pane.*) Click **Next**.



Hierarchy to Delete HANA-SPS\_HDB00

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Next**.



**Delete Resource Hierarchy** x

You have specified the following resource hierarchy for deletion.  
Target Server: hana2-1  
Target Tags:  
HANA-SPS\_HDB00

<Back Delete Cancel Help

5. Another information box appears confirming that the SAP HANA resource was deleted successfully.



6. Click **Done** to exit.

## 6.18.6.5. Testing your SAP HANA Resource Hierarchy

---

### Test Scenarios

To understand the behavior of the SAP HANA Recovery Kit, perform the following tests. The following prerequisites must be completed before performing any test:

- LifeKeeper and the SAP HANA database must be installed and configured according to the installation instructions provided by SIOS and SAP.
- SAP HANA System Replication must be enabled and active on all servers in the cluster, with the secondary replication site registered using one of the valid replication modes (sync, syncmem, or async) and operation modes (delta\_datashipping, logreplay, or logreplay\_readaccess). See [Configure SAP HANA System Replication](#) for more details.
- If managing the switchable IP address associated with the SAP HANA database with a LifeKeeper IP resource, there must exist a dependency of the SAP HANA resource on the IP resource. See Step 7 in [Creating an SAP HANA Resource Hierarchy](#) for more details.

### Test Recovery of SAP Host Agent

Determine the status and the process ID's of the SAP Host Agent processes by using:

```
font face="Courier New"># /usr/sap/hostctrl/exe/saphostexec -status
```

```
saphostexec running (pid = 3818)
```

```
sapstartsrv running (pid = 3867)
```

```
saposcol running (pid = 3965)
```

Either manually kill one of the processes listed in the output or execute

```
/usr/sap/hostctrl/exe/saphostexec -stop
```

to impair the functionality of SAP Host Agent. The SAP HANA Recovery Kit will recognize that SAP Host Agent is not working properly and restart it on that node. The behavior can be observed by monitoring the LifeKeeper log with the following command:

```
tail -f /var/log/lifekeeper.log
```

During this recovery process, the SAP HANA resource does not change its state. After a successful recovery, SAP Host Agent is fully functional again. If the recovery kit is unable to restart SAP Host Agent, the HANA database and the resource remains in its current state. SAP Host Agent will be checked again and if possible restarted later.

## Test Recovery of sapstartsrv for the SAP HANA Instance

To test the recovery of the SAP Start Service (sapstartsrv) for the SAP HANA instance, the service must be stopped. One method to stop sapstartsrv is by executing the sapcontrol StopService webmethod:

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function StopService"
```

where <sid> is the lower-case SAP System ID for the HANA installation and <Inst#> is the HDB instance number. Another method is to kill the sapstartsrv process directly. In either case, sapstartsrv will be restarted by SAP HANA Recovery Kit. The resource does not change its state as long as sapstartsrv is able to be restarted successfully.

## Test Recovery of the Secondary SAP HANA DB (Replication Target)

In the event of a failure of the secondary database instance (replication target) or if the secondary replication site is unregistered in SAP HANA System Replication, the recovery kit will re-register the secondary site with the previous replication and operation modes and restart the secondary database instance.

To induce such a failure, execute one of the following commands on the secondary replication site:

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function Stop"
```

```
su - <sid>adm -c "hdbnsutil -sr_unregister"
```

The behavior can be observed by monitoring the log file /var/log/lifekeeper.log. After the recovery, the state of the database instance and SAP HANA System Replication can be tested by running the following commands on the secondary replication site:

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function GetProcessList"
```

```
su - <sid>adm -c "hdbnsutil -sr_state"
```

In the event that the secondary database instance cannot be started by the recovery kit, the SAP HANA resource is flagged as Failed (OSF) on the corresponding node.

Hierarchies		hana2-1		hana2-2	
! Unprotected					
! HANA-SPS_HDB00	Active	1	✖	Failed	10
✓ vip-sps-hana	Active	1	↓	StandBy	10

Once the cause of an unsuccessful start is fixed by an administrator, the SAP HANA Recovery Kit will start the database instance in the subsequent quickCheck cycle. Once started successfully, the resource state will be updated to Standby (OSU) on the corresponding node.

## Test Recovery of the Primary SAP HANA DB

In the event of a failure of the primary database instance (replication source), the replication mode of the database instance on the primary node is determined. If the replication mode is set to primary, the database instance will be started again. If the mode is not set to primary, the recovery kit will log a warning stating that the replication mode has been changed outside of LifeKeeper and suspend all monitoring of the SAP HANA resource until the issue is resolved. In the latter case, manual intervention is required to bring the HANA resource hierarchy in-service on the correct primary system. The behavior in this case can be observed in the LifeKeeper GUI, which will show the state “Active – HSR Disabled”, “Active – Unknown Repl Mode”, or “Active – Secondary” for the resource on the primary node, or by monitoring the log file `/var/log/lifekeeper.log`.

A failure of the primary database instance can be induced by running the following command on the primary replication site:

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function Stop"
```

After the recovery, the state of the database and the replication can be tested by using:

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function GetProcessList"
```

```
su - <sid>adm -c "hdbnsutil -sr_state"
```

In the event that the primary database instance cannot be started by the recovery kit on that node, LifeKeeper will initiate a failover of the entire hierarchy to the secondary node. On this node, the HANA Recovery Kit performs a takeover of SAP HANA System Replication and the previous secondary node becomes the new primary node for replication. LifeKeeper will attempt to re-register the faulty node as the secondary replication site using the previous replication and operation modes. If this is successful, the secondary database is also restarted. In the event that either the secondary node cannot be successfully registered as the secondary replication site or that the database cannot be successfully restarted on the secondary node, the HANA resource will be flagged as Failed (OSF) on the corresponding node. At this point, manual intervention is typically necessary to eliminate the cause of the failure. If the failover of the primary database instance failed, the resource is flagged as faulty Failed (OSF) and remains in this state until a manual in-service operation is performed by an administrator.

## Test Machine Failure of the Secondary Node (reboot -f, power off)

If an error causes the secondary node to fail, the resource remains Active (ISP) on the primary node but SAP HANA System Replication is disrupted. Once the secondary node is restarted and LifeKeeper is active, the secondary database instance is automatically restarted as a replication target.

## Test Machine Failure of the Primary Node (reboot -f, power off)

If an error causes the primary node to fail, a failover of the HANA resource hierarchy to the secondary node is initiated. A takeover of SAP HANA System Replication is performed on the secondary node and the previous secondary replication site becomes the new primary replication site. Once the faulty node is restarted and LifeKeeper is active, the node is registered as a secondary replication site and the database instance is automatically restarted as a replication target.

## Additional Test Cases

Before putting a highly-available SAP HANA cluster into production, it is very important that common failure and recovery scenarios have been thoroughly tested. The test cases provided on this page are meant to be used as a starting point when developing a comprehensive test plan for your highly-available SAP HANA cluster deployment. The following example values will be used throughout:

<b>Primary Server Host Name</b>	node1
<b>Standby Server Host Name</b>	node2
<b>SAP SID</b>	SPS
<b>SAP HANA Database Instance</b>	HDB00
<b>SAP HANA LifeKeeper Resource Tag Name</b>	HANA-SPS_HDB00
<b>HANA System Replication Primary Site Name</b>	SiteA
<b>HANA System Replication Secondary Site Name</b>	SiteB
<b>HANA System Replication Mode</b>	sync
<b>HANA System Replication Operation Mode</b>	logreplay

When testing, these sample values must be adapted to fit the environment where the tests are being performed.

 **Note:** The following test cases assume that the SAP HANA resource hierarchy is in-service on the primary server (node1). **For full coverage, the same tests should also be performed while the resource hierarchy is in-service on node2.** To test these scenarios with the server roles reversed, make the substitutions “primary” ↔ “standby”, “node1” ↔ “node2”, and “SiteA” ↔ “SiteB” in each of the test cases given below.

### Manual Switchover

The test cases in this section ensure that manual switchovers can be performed successfully.

<b>Manual Switchover Test</b>
<b>Description</b>
The SAP HANA resource hierarchy can be manually switched over from the primary server to the standby server.
<b>Preconditions</b>
Before performing this test, ensure that the following conditions are met: <ul style="list-style-type: none"> <li>Both servers (node1 and node2) are operational,</li> </ul>

- The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and
- HANA system replication is in-sync.

**Test Steps**

1. Bring the SAP HANA resource hierarchy in-service on node2 by using either of the following methods:
  - a. In the LifeKeeper GUI, right-click the HANA-SPS\_HDB00 resource on node2 and select “In Service...” from the context menu. On the resulting confirmation dialog, click “In Service” to begin the switchover process.
  - b. From a terminal window on node2, execute the following command as a user with lkadmin group permissions (e.g., as the root user):

```
[root@node2 ~]# sudo /opt/LifeKeeper/bin/lkcli resource restore --tag HANA-SPS_HDB00
```

**Expected Results**

1. The SAP HANA resource hierarchy is successfully taken out of service on node1. During this process, the database is stopped on node1 and any dependent resources (such as an associated virtual IP address, if applicable) are also removed on node1.
2. The SAP HANA resource hierarchy is successfully brought in-service on node2. During this process:
  - a. Any dependent resources (such as an associated virtual IP resource, if applicable) are brought in-service on node2.
  - b. The running database instance on node2 is promoted to the primary replication role in HANA system replication.
  - c. The database instance on node1 is registered as a secondary replication site in HANA system replication using the appropriate HSR parameters (e.g., site name ‘SiteA’, replication mode ‘sync’, and operation mode ‘logreplay’).
  - d. The database instance is started on node1.

**Note:** The process of registering and restarting the database on node1 may take several minutes. Once this process is complete, the LifeKeeper GUI will show the state of the HANA-SPS\_HDB00 resource on node1 as ‘Standby – In Sync’.

**Handshake Takeover Test**

**Note:** This test case requires LifeKeeper v9.5.2 or later.

**Description**

The SAP HANA resource hierarchy can be manually switched over from the primary server to the standby server by using the SAP HANA “Takeover with Handshake” feature.

**Preconditions**

Before performing this test, ensure that:

- Both servers (node1 and node2) are operational,
- The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and
- HANA system replication is in-sync.

### Test Steps

1. Bring the SAP HANA resource hierarchy in-service on node2 by using either of the following methods:
  - a. In the LifeKeeper GUI, right-click the HANA-SPS\_HDB00 resource on node2 and select “In Service – Takeover with Handshake...” from the context menu. On the resulting confirmation dialog, click “Perform Takeover” to begin the switchover process.
  - b. From a terminal window on either node1 or node2, execute the following command as a user with lkadmin group permissions (e.g., as the root user):

```
#sudo /opt/LifeKeeper/bin/lkcli resource config hana --tag
HANA-SPS_HDB00 --takeover_with_handshake node2
```

### Expected Results

1. The SAP HANA resource hierarchy is successfully taken out of service on node1. During this process the database is not stopped on node1, but any dependent resources (such as an associated virtual IP address, if applicable) are removed on node1.
2. The SAP HANA resource hierarchy is successfully brought in-service on node2. During this process:
  - a. Any dependent resources (such as an associated virtual IP resource, if applicable) are brought in-service on node2.
  - b. The running database instance on node2 is promoted to the primary replication role in HANA system replication. During the HSR takeover process, the database instance on node1 is suspended.
  - c. The suspended database instance on node1 is stopped.
  - d. The database instance on node1 is registered as a secondary replication site in HANA system replication using the appropriate HSR parameters (e.g., site name ‘SiteA’, replication mode ‘sync’, and operation mode ‘logreplay’).
  - e. The database instance is started on node1.

**Note:** The process of stopping, registering, and restarting the database on node1 may take several minutes. Once this process is complete, the LifeKeeper GUI will show the state of the HANA-SPS\_HDB00 resource on node1 as ‘Standby – In Sync’.

## Graceful Shutdown

The test cases in this section verify the expected behavior of the SAP HANA resource hierarchy when each server is gracefully rebooted.

 **Note:** The expected results in the first test depend on whether the “Switchover on Shutdown” strategy is enabled or disabled for the primary server. See [Setting Server Shutdown Strategy](#) for more details.

## Primary Server Reboot Test

### Description

The SAP HANA resource hierarchy behaves as expected when the primary server is gracefully rebooted.

### Preconditions

Before performing this test, ensure that:

- Both servers (node1 and node2) are operational,
- The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and
- HANA system replication is in-sync.

### Test Steps

1. Gracefully reboot node1:

```
[root@node1 ~]# reboot now
```

### Expected Results

1. While node1 shuts down:
  - a. The SAP HANA resource hierarchy is successfully taken out of service on node1. During this process, the database is stopped on node1 and any dependent resources (such as an associated virtual IP address, if applicable) are also removed on node1.
  - b. **[If “Switchover on Shutdown” is enabled on node1]** The SAP HANA resource hierarchy is successfully brought in-service on node2. During this process:
    - i. Any dependent resources (such as an associated virtual IP resource, if applicable) are brought in-service on node2.
    - ii. The running database instance on node2 is promoted to the primary replication role in HANA system replication. Since node1 has been shut down, HANA system replication is currently inactive.
2. After node1 is back online:
  - a. **[If “Switchover on Shutdown” is disabled on node1]** The SAP HANA resource hierarchy automatically comes back in-service on node1. During this process:
    - i. Any dependent resources (such as an associated virtual IP resource, if applicable) are brought in-service on node1.
    - ii. The database instance on node1 is started and HANA system replication resumes to the secondary replication site on node2.
  - b. **[If “Switchover on Shutdown” is enabled on node1]** The SAP HANA resource hierarchy remains in-service on node2.
    - i. During the first quickCheck cycle on node2 after node1 is back online, the SAP

HANA Recovery Kit detects that the database instance is not running on node1 and fires a 'remoteregisterdb' event.

- ii. The 'remoteregisterdb' event script registers node1 as a secondary HANA system replication site using the appropriate HSR parameters (e.g., site name 'SiteA', replication mode 'sync', and operation mode 'logreplay') and starts the database instance on node1.

**Note:** The process of registering and restarting the database on node1 may take several minutes. Once this process is complete, the LifeKeeper GUI will show the state of the HANA-SPS\_HDB00 resource on node1 as 'Standby – In Sync'.

<b>Standby Server Reboot Test</b>
<b>Description</b>
The SAP HANA resource hierarchy behaves as expected when the standby server is gracefully rebooted.
<b>Preconditions</b>
<p>Before performing this test, ensure that:</p> <ul style="list-style-type: none"> <li>• Both servers (node1 and node2) are operational,</li> <li>• The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and</li> <li>• HANA system replication is in-sync.</li> </ul>
<b>Test Steps</b>
<p>1. Gracefully reboot node2:</p> <pre>[root@node2 ~]# reboot now</pre>
<b>Expected Results</b>
<ol style="list-style-type: none"> <li>1. While node2 shuts down, the database instance is stopped on node2. HANA system replication will be inactive until node2 reboots and the secondary database instance is restarted.</li> <li>2. After node2 is back online:             <ol style="list-style-type: none"> <li>a. During the first quickCheck cycle on node1 after node2 is back online, the SAP HANA Recovery Kit detects that the database instance is not running on node2 and fires a 'remoteregisterdb' event.</li> <li>b. The 'remoteregisterdb' event script starts the database instance on node2.</li> </ol> </li> </ol>
<p><b>Note:</b> The process of restarting the database on node2 may take several minutes. Once this process is complete, the LifeKeeper GUI will show the state of the HANA-SPS_HDB00 resource on node2 as 'Standby – In Sync'.</p>

## Machine Failover

The test case in this section verifies the expected behavior of the SAP HANA resource hierarchy when the primary server is forcefully rebooted.

<b>Machine Failover Test</b>
<b>Description</b>
The SAP HANA resource hierarchy behaves as expected when the primary server is forcefully rebooted.
<b>Preconditions</b>
<p>Before performing this test, ensure that:</p> <ul style="list-style-type: none"> <li>• Both servers (node1 and node2) are operational,</li> <li>• The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and</li> <li>• HANA system replication is in-sync.</li> </ul>
<b>Test Steps</b>
<p>1. Forcefully reboot node1:</p> <pre>[root@node1 ~]# echo b &gt; /proc/sysrq-trigger</pre>
<b>Expected Results</b>
<ol style="list-style-type: none"> <li>1. Once LifeKeeper on node2 detects that node1 is down (the exact time will vary depending on the values being used for the <a href="#">LifeKeeper heartbeat parameters</a>), the SAP HANA resource hierarchy is successfully brought in-service on node2. During this process: <ol style="list-style-type: none"> <li>a. Any dependent resources (such as an associated virtual IP resource, if applicable) are brought in-service on node2.</li> <li>b. The running database instance on node2 is promoted to the primary replication role in HANA system replication. Since node1 has been shut down, HANA system replication is currently inactive.</li> </ol> </li> <li>2. After node1 is back online, the SAP HANA resource hierarchy remains in-service on node2. <ol style="list-style-type: none"> <li>a. During the first quickCheck cycle on node2 after node1 is back online, the SAP HANA Recovery Kit detects that the database instance is not running on node1 and fires a 'remoteregisterdb' event.</li> <li>b. The 'remoteregisterdb' event script registers node1 as a secondary HANA system replication site using the appropriate HSR parameters (e.g., site name 'SiteA', replication mode 'sync', and operation mode 'logreplay') and starts the database instance on node1.</li> </ol> </li> </ol> <p><b>Note:</b> The process of registering and restarting the database on node1 may take several minutes. Once this process is complete, the LifeKeeper GUI will show the state of the HANA-SPS_HDB00 resource on node1 as 'Standby – In Sync'.</p>

## SAP Host Agent Failure

The test cases in this section verify the expected behavior of the SAP HANA resource hierarchy when supporting SAP Host Agent-related processes fail on each server.

Primary Server SAP Host Exec Failure Test
<b>Description</b>
<p>The SAP HANA resource hierarchy behaves as expected when the <code>saphostexec</code> process is killed on the primary server.</p>
<b>Preconditions</b>
<p>Before performing this test, ensure that:</p> <p>Both servers (node1 and node2) are operational,  The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and  HANA system replication is in-sync.</p>
<b>Test Steps</b>
<p>Stop the <code>saphostexec</code> process on node1:</p> <pre>[root@node1 ~]# /usr/sap/hostctrl/exe/saphostexec -stop</pre>
<b>Expected Results</b>
<ol style="list-style-type: none"> <li>1. During the next quickCheck interval, the SAP HANA Recovery Kit detects that the <code>saphostexec</code> process is no longer running on node1 and fires a 'recover' event to restart it.</li> <li>2. The 'recover' event script restarts the <code>saphostexec</code> process on node1.</li> </ol>

Standby Server SAP Host Exec Failure Test
<b>Description</b>
<p>The SAP HANA resource hierarchy behaves as expected when the <code>saphostexec</code> process is killed on the standby server.</p>
<b>Preconditions</b>
<p>Before performing this test, ensure that:</p> <ul style="list-style-type: none"> <li>• Both servers (node1 and node2) are operational,</li> <li>• The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and</li> <li>• HANA system replication is in-sync.</li> </ul>
<b>Test Steps</b>

<p>1. Stop the <code>saphostexec</code> process on node2:</p> <pre>[root@node2 ~]# /usr/sap/hostctrl/exe/saphostexec -stop</pre>
<p><b>Expected Results</b></p>
<ol style="list-style-type: none"> <li>1. During the next quickCheck interval, the SAP HANA Recovery Kit detects that the <code>saphostexec</code> process is no longer running on node2 and fires a 'remoteregisterdb' event to restart it.</li> <li>2. The 'remoteregisterdb' event script restarts the <code>saphostexec</code> process on node2.</li> </ol>

<p><b>Primary Server SAP OS Collector Failure Test</b></p>
<p><b>Description</b></p> <p>The SAP HANA resource hierarchy behaves as expected when the <code>saposcol</code> process is killed on the primary server.</p>
<p><b>Preconditions</b></p> <p>Before performing this test, ensure that:</p> <ul style="list-style-type: none"> <li>• Both servers (node1 and node2) are operational,</li> <li>• The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and</li> <li>• HANA system replication is in-sync.</li> </ul>
<p><b>Test Steps</b></p> <p>Stop the <code>saposcol</code> process on node1:</p> <pre>[root@node1 ~]# /usr/sap/hostctrl/exe/saposcol -k</pre>
<p><b>Expected Results</b></p> <ol style="list-style-type: none"> <li>1. During the next quickCheck interval, the SAP HANA Recovery Kit detects that the <code>saposcol</code> process is no longer running on node1 and fires a 'recover' event to restart it.</li> <li>2. The 'recover' event script restarts the <code>saposcol</code> process on node1.</li> </ol>

<p><b>Standby Server SAP OS Collector Failure Test</b></p>
<p><b>Description</b></p> <p>The SAP HANA resource hierarchy behaves as expected when the <code>saposcol</code> process is killed on the standby server.</p>
<p><b>Preconditions</b></p> <p>Before performing this test, ensure that:</p>

<ul style="list-style-type: none"> <li>• Both servers (node1 and node2) are operational,</li> <li>• The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and</li> <li>• HANA system replication is in-sync.</li> </ul>
<p><b>Test Steps</b></p>
<p>Stop the <code>saposcol</code> process on node2:</p> <pre>[root@node2 ~]# /usr/sap/hostctrl/exe/saposcol -k</pre>
<p><b>Expected Results</b></p>
<ol style="list-style-type: none"> <li>1. During the next quickCheck interval, the SAP HANA Recovery Kit detects that the <code>saposcol</code> process is no longer running on node2 and fires a 'remoteregisterdb' event to restart it.</li> <li>2. The 'remoteregisterdb' event script restarts the <code>saposcol</code> process on node2.</li> </ol>

## Database Failure

The test cases in this section verify the expected behavior of the SAP HANA resource hierarchy when processes within the protected HANA database instance fail.

**Notes:**

- Some of the expected results in this section depend on whether local recovery is enabled or disabled for the SAP HANA resource. Local recovery is enabled by default. See [Setting Local and Temporal Recovery Policies for SAP HANA Resources](#) for more details.
- If the database instance is stopped gracefully when performing these tests (e.g., with an `HDB stop` command), the background process which gracefully stops the database may conflict with the SAP HANA Recovery Kit's attempt to restart the database locally and may lead to a failover of the SAP HANA resource hierarchy. For this reason, we recommend simulating a crash of the HDB instance by forcefully and immediately killing the database processes at the operating system level, for example with the `HDB kill-9` command. See [SAP HANA – Known Issues / Restrictions](#) for more details.

<p><b>Primary Server Database Instance Failure Test</b></p>
<p><b>Description</b></p> <p>The SAP HANA resource hierarchy behaves as expected when the HDB instance processes are killed on the primary server.</p>
<p><b>Preconditions</b></p> <p>Before performing this test, ensure that:</p> <ul style="list-style-type: none"> <li>• Both servers (node1 and node2) are operational,</li> <li>• The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on</li> </ul>

- node2, and
- HANA system replication is in-sync.

### Test Steps

As the HANA administrative user (<sid>adm), forcefully kill the HDB instance processes (at the operating system level) on node1:

```
[root@node1 ~]# su - spsadm -c "HDB kill-9"
```

### Expected Results

- [If local recovery is enabled for the HANA-SPS\_HDB00 resource on node1]** The SAP HANA Recovery Kit detects the failure and restarts the database instance on node1:
  - During the next quickCheck interval, the SAP HANA Recovery Kit detects that the HDB instance processes are no longer running on node1 and fires a 'recover' event to restart them.
  - The 'recover' event script restarts the HDB instance processes on node1.
- [If local recovery is disabled for the HANA-SPS\_HDB00 resource on node1]** LifeKeeper immediately initiates a failover of the SAP HANA resource hierarchy to node2:
  - The SAP HANA resource hierarchy is successfully taken out of service on node1. During this process, any dependent resources (such as an associated virtual IP address, if applicable) are removed on node1.
  - The SAP HANA resource hierarchy is successfully brought in-service on node2. During this process:
    - Any dependent resources (such as an associated virtual IP resource, if applicable) are brought in-service on node2.
    - The running database instance on node2 is promoted to the primary replication role in HANA system replication.
    - The database instance on node1 is registered as a secondary replication site in HANA system replication using the appropriate HSR parameters (e.g., site name 'SiteA', replication mode 'sync', and operation mode 'logreplay').
    - The database instance is started on node1.

**Note:** The process of registering and restarting the database on node1 may take several minutes. Once this process is complete, the LifeKeeper GUI will show the state of the HANA-SPS\_HDB00 resource on node1 as 'Standby – In Sync'.

## Standby Server Database Instance Failure Test

### Description

The SAP HANA resource hierarchy behaves as expected when the HDB instance processes are killed on the standby server.

### Preconditions

Before performing this test, ensure that:

- Both servers (node1 and node2) are operational,
- The SAP HANA resource hierarchy is in-service (ISP) on node1 and out-of-service (OSU) on node2, and
- HANA system replication is in-sync.

### Test Steps

As the HANA administrative user (<sid>adm), forcefully kill the HDB instance processes (at the operating system level) on node2:

```
[root@node2 ~]# su - spsadm -c "HDB kill-9"
```

### Expected Results

1. During the next quickCheck interval, the SAP HANA Recovery Kit detects that the HDB instance processes are no longer running on node2 and fires a 'remoteregisterdb' event to restart them.
2. The 'remoteregisterdb' event script restarts the HDB instance processes on node2.

## Appendix: Useful SAP HANA Administrative Commands

While the status of the SAP HANA environment may be monitored through SAP-provided dashboards (e.g., HANA Studio or HANA Cockpit), the following commands may also be useful while testing. Throughout, <sid> denotes the lowercase SAP SID for the protected SAP HANA database installation and <InstNum> denotes the instance number of the protected HDB instance (e.g., for instance HDB00, <InstNum> is 00).

Command	Description
su - <sid>adm -c "sapcontrol -nr <InstNum> -function StopService"	Stop the sapstartsrv process for the HDB instance.
su - <sid>adm -c "sapcontrol -nr <InstNum> -function StartService <SID>"	Start the sapstartsrv process for the HDB instance.
su - <sid>adm -c "sapcontrol -nr <InstNum> -function GetProcessList"	View current status of HDB instance processes.
su - <sid>adm -c "HDB stop"	Gracefully stop the HDB instance.
su - <sid>adm -c "HDB kill-9"	Forcefully kill the HDB instance processes.
su - <sid>adm -c "HDB start"	Start the HDB instance.
su - <sid>adm -c "hdbnsutil -sr_state"	Check the HANA system replication state on the local server.
su - <sid>adm -c "python /hana/shared/<SID>/HDB<InstNum>/exe/python_support/systemReplicationStatus.py"	Check the current HANA system replication status. This command must be executed on the server which is the primary HANA system replication site.

<pre>su - &lt;sid&gt;adm -c "hdbsql -n &lt;HANA virtual hostname&gt; -i &lt;InstNum&gt; -u SYSTEM -p &lt;SYSTEM user password&gt; -d &lt;SID&gt; '\s'"</pre>	<p>Test the connection to the &lt;SID&gt; tenant database through the associated virtual hostname.</p>
<pre>/usr/sap/hostctrl/exe/saphostexec -status</pre>	<p>Check the status of saphostexec.</p>
<pre>/usr/sap/hostctrl/exe/saphostexec -stop</pre>	<p>Stop saphostexec.</p>
<pre>/usr/sap/hostctrl/exe/saphostexec -restart</pre>	<p>Restart saphostexec.</p>
<pre>/usr/sap/hostctrl/exe/saposcol -s</pre>	<p>Check the status of saposcol.</p>
<pre>/usr/sap/hostctrl/exe/saposcol -k</pre>	<p>Stop saposcol.</p>
<pre>/usr/sap/hostctrl/exe/saposcol -l</pre>	<p>Start saposcol.</p>
<pre>top -U &lt;sid&gt;adm ps -ef   grep &lt;sid&gt;adm</pre>	<p>View information for processes owned by the &lt;sid&gt;adm user.</p>

## 6.18.7. SAP HANA Resource Hierarchy Administration

\* **Important Note:** Unless otherwise noted in this guide, all administrative tasks for a LifeKeeper-protected SAP HANA Database should be performed through LifeKeeper. Performing administrative actions such as stopping the database or disabling SAP HANA System Replication outside of LifeKeeper while the SAP HANA resource is Active/ISP in LifeKeeper will result in LifeKeeper taking action to place the database back into its expected running state.

### Switchover of the SAP HANA Resource

When a switchover of the primary database instance is initiated, the SAP HANA Recovery Kit performs the following steps:

- The database instance is stopped on the previous primary node
- A takeover of SAP HANA System Replication is executed on the new primary node (i.e., the previous secondary node)
- The new secondary node (i.e., the previous primary node) is re-registered as the secondary SAP HANA System Replication site
- The database instance is started on the new secondary node

If a message similar to the following:

```
ERROR:hana:restore:HANA-SPS_HDB00:136266:The resource HANA-SPS_HDB00
protecting SAP HANA database HDB00 is not in sync. To protect the data
LifeKeeper will not restore the resource on $me. Please restore the resource
on the previous source server to allow the resync to complete.
```

is displayed while bringing the SAP HANA resource in-service, this means that SAP HANA System Replication was not in-sync when the primary database instance was stopped. Therefore data may exist on the primary database server which has not yet been replicated to the secondary database server. For this reason, LifeKeeper will not allow the secondary server to take over the primary replication role. The recommendation in this scenario is to bring the SAP HANA resource hierarchy back in-service on the previous primary server and allow the resynchronization to complete.

If the previous primary server is down and cannot be recovered, the SAP HANA resource can be forced online on server where the data is out-of-sync, but **this will result in a loss of all data that has not yet been replicated from the previous primary server**. If it is determined by a database administrator that this data loss is acceptable or unavoidable, the out-of-sync data flag can be manually removed with the following command:

```
/opt/LifeKeeper/bin/flg_remove -f '!HANA_DATA_OUT_OF_SYNC_<Tag>'
```

where <Tag> is the SAP HANA resource tag name in LifeKeeper (e.g., HANA-SPS\_HDB00). After removing the out-of-sync data flag, reattempt the in-service operation for the HANA resource. Once brought in-service, the database will take over the primary SAP HANA System Replication role and all data on the previous primary server that has not been replicated will be lost. Once the previous primary server is repaired and brought back online, it will be registered as the secondary system replication site and the database will be restarted as the replication target.

## Takeover with Handshake of the SAP HANA Database

The “takeover with handshake” feature, available in SAP HANA 2.0 SPS04 and later, allows for reduced downtime of the primary database during switchover by suspending the primary database (rather than completely stopping it) before performing a takeover of SAP HANA System Replication on the new database host. For more information on how to perform this type of takeover with the SAP HANA Recovery Kit, see [Takeover with Handshake](#).

## Stopping the SAP HANA Database

When the SAP HANA resource is taken out of service in LifeKeeper, only the primary database instance is stopped. The secondary database instance is kept running to minimize downtime during switchover or failover of the HANA resource hierarchy.

There are two special cases that cause exceptional behavior:

1. If the `!volatile!hana_leave_db_running_<HANA Tag>` LifeKeeper flag is set on the system where the SAP HANA resource is being taken out of service, LifeKeeper will not stop the database instance during the out of service operation. If this flag has been set unintentionally, it can be removed with the following command:

```
/opt/LifeKeeper/bin/flg_remove -f '!volatile!hana_leave_db_running_<HANA Tag>'
```

2. If the database is suspended on the system where the SAP HANA resource is being taken out of service, LifeKeeper will not stop the database instance during the out of service operation. Keeping the suspended database instance running in this scenario preserves the option to resume it, should that action be required. This scenario most commonly occurs when a “takeover with handshake” has been manually performed by a database administrator outside of LifeKeeper. If necessary, the suspended database instance can be stopped manually with the following command:

```
su - <sid>adm -c "sapcontrol -nr <InstNum> -function StopWait 600 5"
```

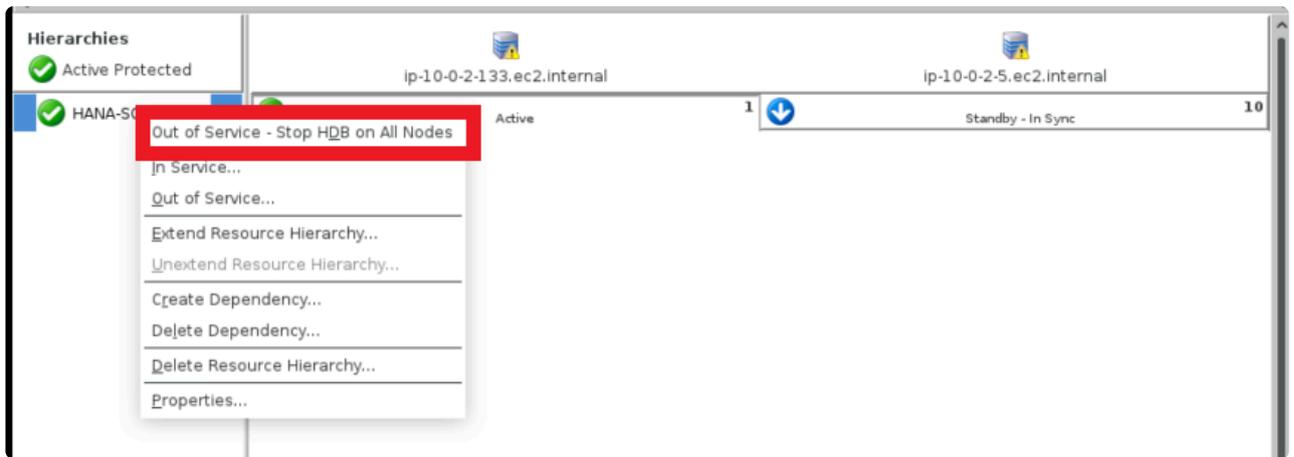
where <sid> is the lower-case SID for the SAP HANA installation and <InstNum> is the HDB instance number.

## Stopping all SAP HANA Databases (Maintenance Mode)

When this option is chosen the primary HANA resource is taken out of service and all of the HANA database instances in a HANA resource cluster will be stopped. This option must be executed with utmost care, as it brings the possibility of a quick failover/switchover to the backup machine. **Note:** This option should only be chosen in the event that the secondary database instance must also be stopped (e.g. during the maintenance window).

To use this option perform the following steps on HANA resource hierarchy:

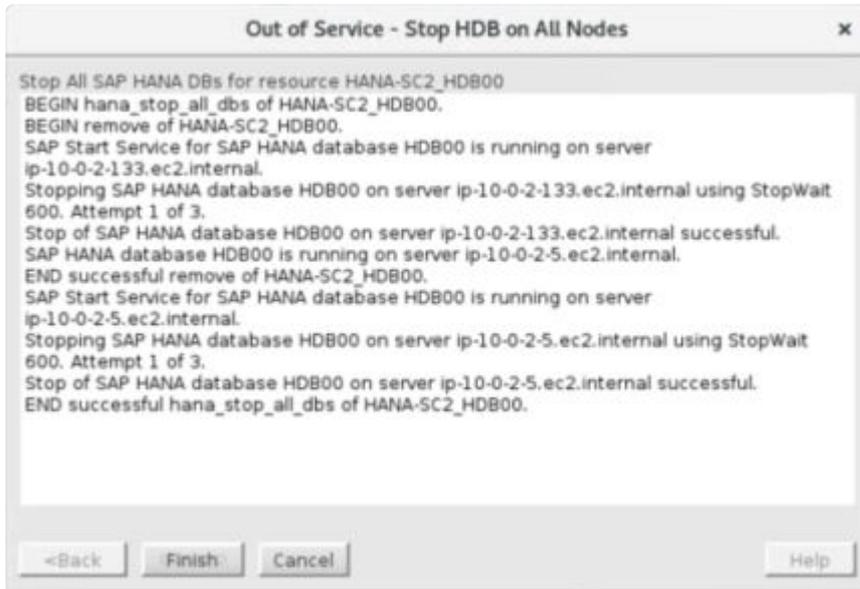
1. Right click either on HANA resource under the left hand panel or an in-service server and choose the option Out of Service – Stop HDB on All Nodes.



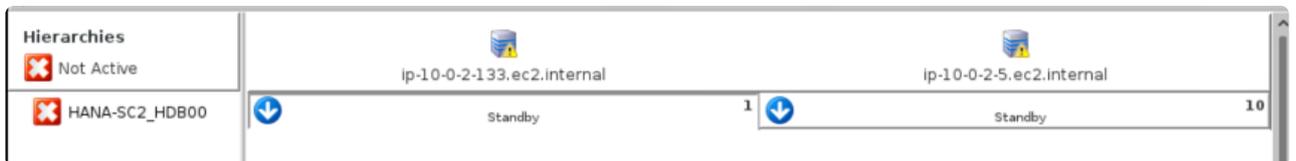
2. Verify the HANA resource and follow the instructions given in the dialog box. Click on **Stop All SAP HANA DBs** to start the process.



3. Once the process finishes, click **Finish** to complete the process.



4. The final state of SAP HANA resource will appear as shown below:



5. Once all of the maintenance activities are complete, bring the SAP HANA resource hierarchy in service on the last primary system.

The protected SAP HANA database may also be stopped on all servers in the cluster by executing the following command on the server where the SAP HANA resource is currently in-service (ISP):

```
/opt/LifeKeeper/bin/hana_stop_all_dbs -t <HANA Resource Tag>
```

## 6.18.7.1. Changing Replication and Operation Modes

! While the replication and operation modes may be changed for the secondary replication site, the user should never perform a manual takeover or switchover of SAP HANA System Replication outside of LifeKeeper. Doing so will result in an error state in which the SAP HANA System Replication modes do not align with the modes that LifeKeeper is expecting on the Active/Standby servers. This error state is represented in the LifeKeeper GUI with the following resource states:

Hierarchies		hana2-1		hana2-2	
Not Active					
HANA-SPS_HDB00		Active - Secondary	1	Standby - Primary	10
vip-sps-hana		Active	1	StandBy	10

While in this state, all monitoring for the SAP HANA resource will be suspended until the resource is brought in-service in LifeKeeper on the server that is intended to be the primary replication site.

### Changing the Replication Mode

The SAP HANA System Replication mode can be changed by using the “hdbnsutil -sr\_changemode” command, even while the database is running on the secondary replication site.

1. Execute the following command on the secondary replication site:

```
su - <sid>adm -c "hdbnsutil -sr_changemode --mode=[sync|syncmem|async]"
```

where <sid> is the lower-case SAP System ID for the SAP HANA installation and the desired replication mode (sync, syncmem, or async) is provided in the -mode option. The database can be running or stopped on the secondary site when this command is executed.

2. To confirm that the replication mode was changed successfully, execute the following command on the secondary replication site:

```
su - <sid>adm -c "hdbnsutil -sr_state"
```

and verify that the “mode” parameter has been updated to the new replication mode.

After the next quickCheck interval (default 2 minutes), LifeKeeper will automatically detect the change and update the info fields for all equivalent SAP HANA resources to contain the new replication mode. This can be verified by inspecting the info field for the HANA resource using the following command:

```
/opt/LifeKeeper/bin/ins_list -d <HostName> -t <Tag> -f: | cut -d: -f6 | tr
```

```
\\002' \:'
```

where `<HostName>` is the host name of the server to obtain information on and `<Tag>` is the tag name of the SAP HANA resource in LifeKeeper (e.g., HANA-SPS\_HDB00) on the given server.

The currently stored secondary replication mode is the third field in the output of this command:

```
SPS:HDB00:sync:shiba:logreplay
```

## Changing the Operation Mode

 **Note:** Changing the operation mode from `logreplay` or `logreplay_readaccess` to `delta_datashipping` will require a full data shipping from the primary site to the secondary site when replication resumes. This full data shipping will be performed automatically by SAP HANA when one of these operation mode changes is detected.

Changing the operation mode in SAP HANA System Replication typically requires stopping the database on the secondary site, re-registering it with the new operation mode, and restarting the database on the secondary site. In order to prevent LifeKeeper from automatically restarting the database on the secondary site during this process, it is recommended to suspend monitoring for the HANA resource until the database has been successfully re-registered with the new operation mode.

1. Suspend monitoring of the HANA resource by creating the corresponding `nomonitor` flag on the server where the resource is currently Active (ISP):

```
/opt/LifeKeeper/bin/flg_create -f 'nomonitor_<Tag>'
```

where `<Tag>` is the HANA resource tag in LifeKeeper on that server (e.g., HANA-SPS\_HDB00). Once finished with this process, it is very important to remember to resume monitoring of the resource by removing the `nomonitor` flag with the command:

```
/opt/LifeKeeper/bin/flg_remove -f 'nomonitor_<Tag>'
```

**Warning:** Failure to remove this flag after performing the maintenance operations will cause any failures of the SAP HANA database to go undetected by LifeKeeper.

2. Ensure that the SAP HANA database is stopped on the backup server by executing the following command:

```
su - <sid>adm -c "sapcontrol -nr <HDB Inst# > -function StopSystem HDB"
```

where `<HDB Inst#>` is the instance number for the SAP HANA database instance (e.g., for an instance named HDB00, the instance number is 00).

3. Execute the following command on the backup server with the desired replication and operation modes to re-register it as a secondary replication site:

```
su - <sid>adm -c "hdbnsutil -sr_register --name=<SecondarySiteName>
--remoteHost=<PrimaryHost> --remoteInstance=<PrimaryInst#>
--replicationMode=[sync|syncmem|async]
--operationMode=[delta_datashipping|logreplay|logreplay_readaccess]"
```

where `<SecondarySiteName>` is the alias to be used by SAP HANA System Replication to identify the secondary replication site, `<PrimaryHost>` is the host name of the server which is currently registered as the primary replication site, and `<PrimaryInst#>` is the instance number for the HDB instance on the primary replication site. **Note:** It is not necessary to unregister the secondary site with the “`hdbnsutil -sr_unregister`” command before re-registering it with a new replication or operation mode.

4. To verify that the backup server was successfully re-registered as a secondary replication site with the new operation mode, execute the following command on the backup server:

```
su - <sid>adm -c "hdbnsutil -sr_state"
```

and verify that “is secondary/consumer system” is true and that the parameter “operation mode” has been updated to the new operation mode.

5. Once the backup server has been successfully registered as a secondary replication site, start the database by executing the following command on the backup server:

```
where su - <sid>adm -c "sapcontrol -nr <HDB Inst#> -function StartSystem
HDB"
```

where `<HDB Inst#>` is the instance number for the SAP HANA database instance.

6. Once the process is complete, remove the *nomonitor* flag that was created in Step 1 in order to resume LifeKeeper monitoring of the HANA resource:

```
where /opt/LifeKeeper/bin/flg_remove -f 'nomonitor_<Tag>'
```

**Warning:** Failure to remove this flag will cause any failures of the SAP HANA database to go undetected by LifeKeeper.

After the next quickCheck interval (default 2 minutes), LifeKeeper will automatically detect the change and update the info fields for all equivalent SAP HANA resources to contain the new operation mode. This can be verified by inspecting the info field for the HANA resource using the following command:

```
/opt/LifeKeeper/bin/ins_list -d <HostName> -t <Tag> -f: | cut -d: -f6 | tr
'\002' \:'
```

where `<HostName>` is the host name of the server to obtain information on and `<Tag>` is the tag name of the SAP HANA resource in LifeKeeper (e.g., HANA-SPS\_HDB00) on the given server.

The currently stored secondary replication mode is the fifth field in the output of this command:

SPS:HDB00:sync:shiba:**logreplay**

## 6.18.7.2. Resolving Split Brain Scenarios

A “split brain” scenario occurs when the SAP HANA database is running and configured as the primary SAP HANA Replication site on multiple cluster nodes. In this situation, LifeKeeper will suspend all monitoring of the HANA database until the issue is manually resolved by a database administrator.

There are two common types of split brain scenarios which may occur for an SAP HANA resource hierarchy.

- **LifeKeeper HANA Resource Split Brain:** The HANA resource is Active (ISP) in LifeKeeper on multiple cluster nodes. This situation is typically caused by a temporary network outage affecting the communication paths between cluster nodes.
- **SAP HANA System Replication Split Brain:** The HANA resource is Active (ISP) on the primary node and Standby (OSU) on the backup node in LifeKeeper, but the database is running and registered as the primary replication site on both nodes. This situation is typically caused by either a failure to stop the database on the previous primary node during failover, having Autostart enabled for the database, or a database administrator manually running “hdbnsutil -sr\_takeover” on the secondary replication site outside of LifeKeeper.

Recommendations for resolving each type of split brain scenario are given below.

### LifeKeeper HANA Resource Split Brain Resolution

Hierarchies	hana2-1		hana2-2	
Unknown				
HANA-SPS_HDB00	Active	1	Active	10
vip-sps-hana	Active	1	Active	10

While in this split brain scenario, a message similar to the following is logged and broadcast to all open consoles every quickCheck interval (default 2 minutes) until the issue is resolved.

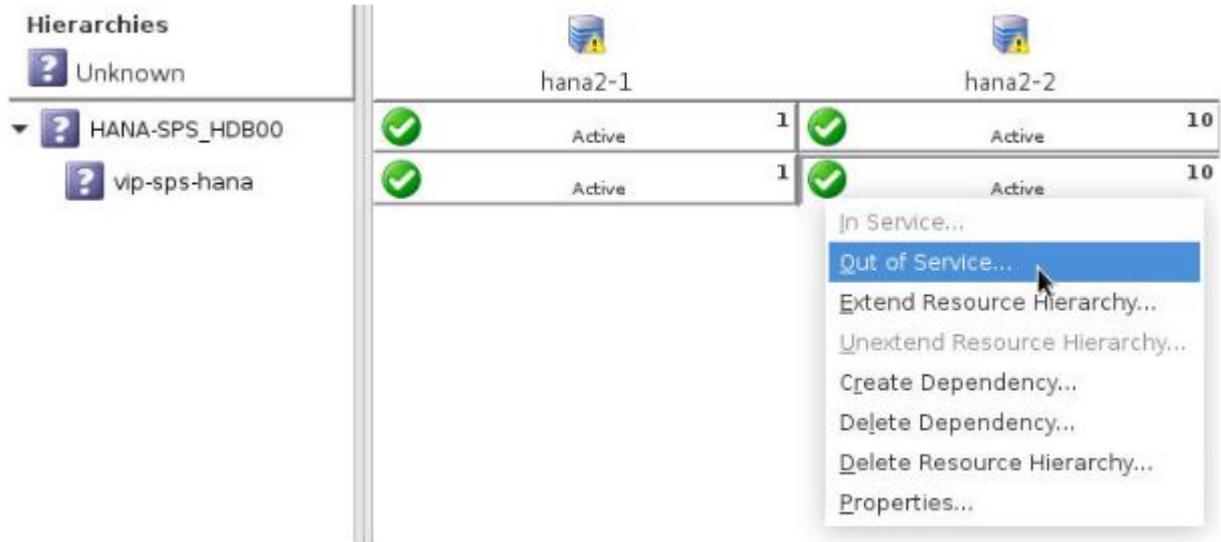
```
EMERG:hana:quickCheck:HANA-SPS_HDB00:136363:WARNING: A temporary communication failure has occurred between servers hana2-1 and hana2-2. Manual intervention is required in order to minimize the risk of data loss. To resolve this situation, please take one of the following resource hierarchies out of service: HANA-SPS_HDB00 on hana2-1 or HANA-SPS_HDB00 on hana2-2. The server that the resource hierarchy is taken out of service on will become the secondary SAP HANA System Replication site.
```

#### Recommendations for resolution:

1. Investigate the database on each cluster node to determine which instance contains the most up-

to-date or relevant data. This determination must be made by a qualified database administrator who is familiar with the data.

- The HANA resource on the node containing the data that needs to be retained will remain Active (ISP) in LifeKeeper, and the HANA resource hierarchy on the node that will be re-registered as the secondary replication site will be taken entirely out of service in LifeKeeper. Right-click on each leaf resource in the HANA resource hierarchy on the node where the hierarchy should be taken out of service and click **Out of Service ...**



\* It is important in this step that the entire SAP HANA resource hierarchy (including the virtual IP resource, if one exists) is taken out of service on the node that will be re-registered as the secondary replication site.

- Once the SAP HANA resource hierarchy has been successfully taken out of service, LifeKeeper will re-register the Standby node as the secondary replication site during the next quickCheck interval (default 2 minutes). Once replication resumes, any data on the Standby node which is not present on the Active node will be lost. Once the Standby node has been re-registered as the secondary replication site, the SAP HANA hierarchy has returned to a highly-available state.



## SAP HANA System Replication Split Brain Resolution

While in this split brain scenario, a message similar to the following is logged and broadcast to all open

consoles every quickCheck interval (default 2 minutes) until the issue is resolved.

```
EMERG:hana:quickCheck:HANA-SPS_HDB00:136234:WARNING: SAP HANA database HDB00
is running and registered as primary master on the following servers: hana2-2,
hana2-1. Manual intervention is required in order to minimize the risk of data
loss. To resolve this situation, please stop database HDB00 on the standby
server by running the command `su - spsadm -c "sapcontrol -nr 00 -function
StopWait 600 5"' on that server, allow LifeKeeper to register the standby
server as a secondary replication site, then use LifeKeeper to bring resource
HANA-SPS_HDB00 in-service on the intended primary replication site.
```

### Recommendations for resolution:

1. Investigate the database on each cluster node to determine whether important data exists on the Standby node which does not exist on the Active node. If important data has been committed to the database on the Standby node while in the split brain state, the data will need to be manually copied to the Active node. This determination must be made by a qualified database administrator who is familiar with the data.
2. Once any missing data has been copied from the database on the Standby node to the Active node, stop the database on the Standby node by running the command given in the LifeKeeper warning message:

```
su - <sid>adm -c "sapcontrol -nr <Inst#> -function StopWait 600 5"
```

where <sid> is the lower-case SAP System ID for the HANA installation and <Inst#> is the instance number for the HDB instance (e.g., the instance number for instance HDB00 is 00).

3. Once the database has been successfully stopped, LifeKeeper will re-register the Standby node as the secondary replication site during the next quickCheck interval (default 2 minutes). Once replication resumes, any data on the Standby node which is not present on the Active node will be lost. Once the Standby node has been re-registered as the secondary replication site, the SAP HANA hierarchy has returned to a highly-available state and may be brought in-service on any server in the cluster.

## 6.18.7.3. Takeover with Handshake

---

### Takeover with Handshake of the SAP HANA Database

The “takeover with handshake” feature, available in SAP HANA 2.0 SPS04 and later, allows for reduced downtime of the primary database during switchover by suspending the primary database (rather than completely stopping it) before performing a takeover of SAP HANA System Replication on the new database host.

When a “takeover with handshake” is initiated, the SAP HANA Recovery Kit performs the following steps:

- The database instance is left running on the previous primary node.
- A “takeover with handshake” of SAP HANA System Replication is executed on the new primary node (i.e., the previous secondary node). This puts the database on the original primary node into a suspended state before registering the new primary system replication site.
- The suspended database on the new secondary node (i.e., the previous primary node) is stopped and re-registered as the secondary system replication site.
- The database instance is started on the new secondary node.

### Performing a Takeover with Handshake

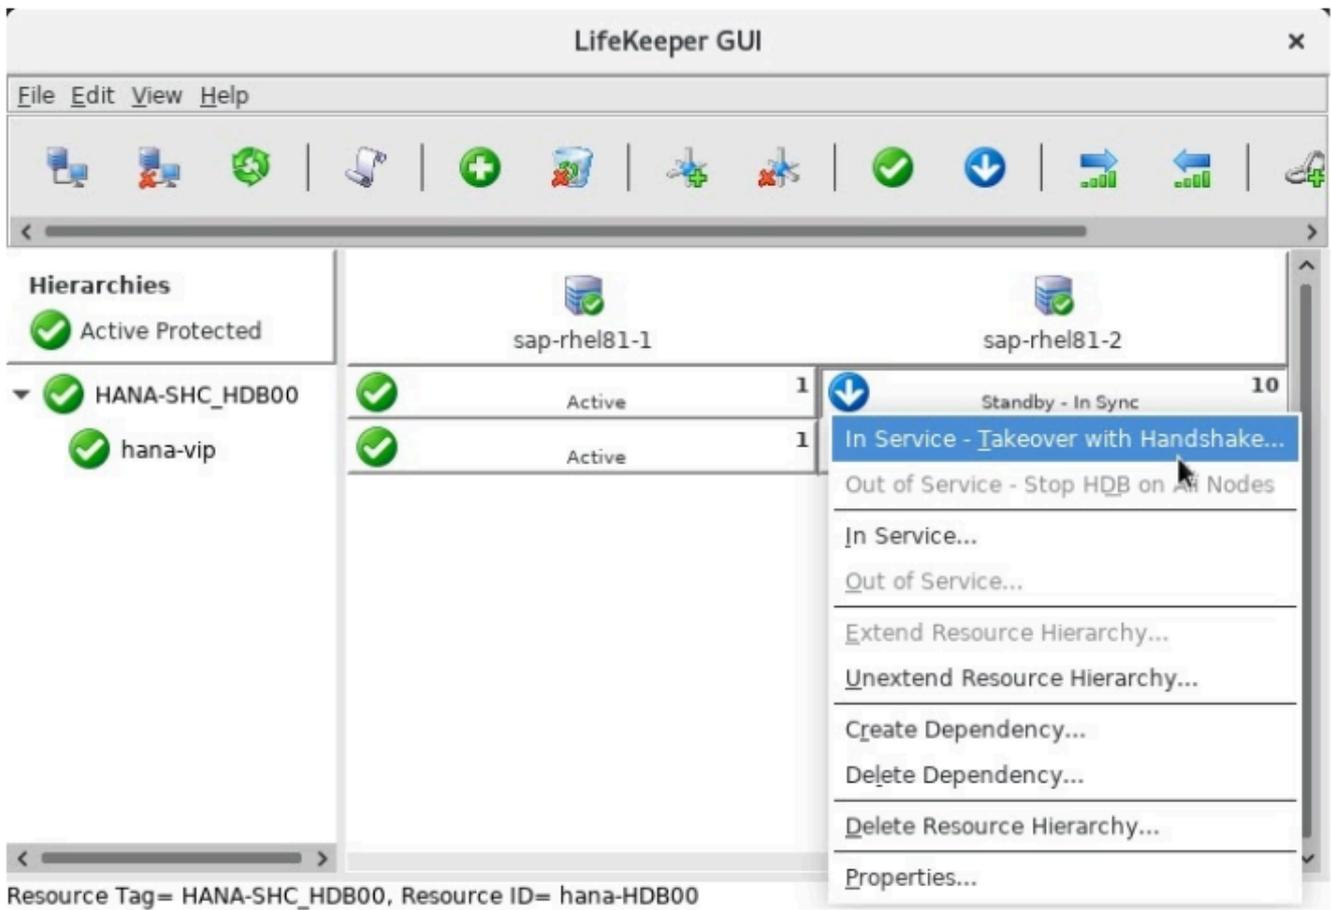
#### Prerequisites

Before performing a takeover with handshake, first ensure that the following prerequisites are met:

1. SAP HANA 2.0 SPS04 or later is installed on each server,
2. The protected SAP HANA database is in-service on a server in the cluster, and
3. SAP HANA System Replication is in-sync.

#### Takeover with Handshake in the LifeKeeper GUI

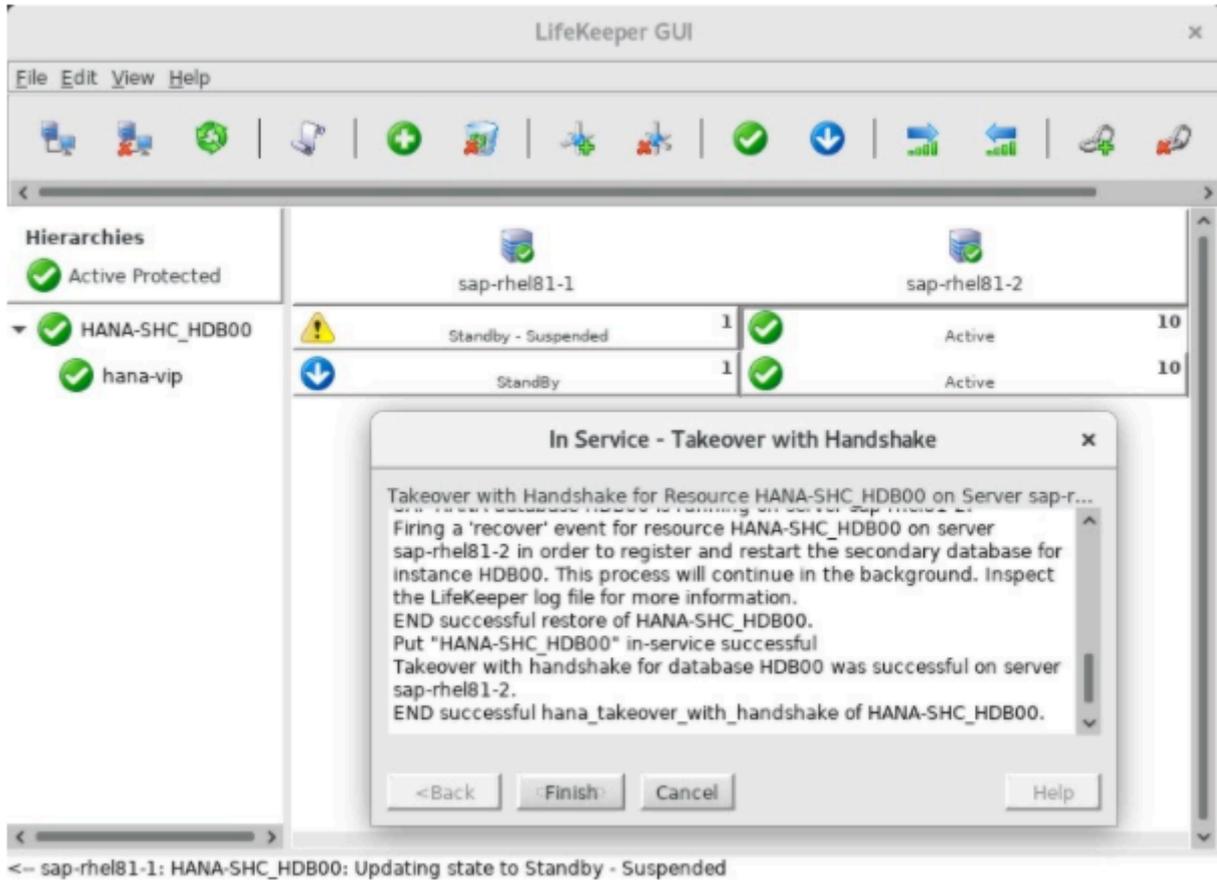
To perform a takeover with handshake in the LifeKeeper GUI, right-click the resource on the standby server and select the “In Service – Takeover with Handshake...” action.



This will bring up a confirmation dialog describing the takeover with handshake feature:



Click **Perform Takeover** to initiate the takeover. Once the takeover process is complete, the states of the SAP HANA resources on each node will transition to **Standby – Suspended** and **Active**.



Once the SAP HANA resource has successfully been brought in-service on the target server, a 'remoteregisterdb' event will automatically fire in the background to stop, register, and restart the secondary database instance. If you would prefer this process of restarting the secondary database to occur in the foreground while bringing the SAP HANA resource in-service, set `HANA_REGISTER_SECONDARY_DURING_RESTORE=true` in `/etc/default/LifeKeeper`.

Once the 'remoteregisterdb' event has successfully restarted the secondary database instance, the resource state will transition to **Standby – HDB Running**. After LifeKeeper spawns its HSR monitoring process during the subsequent quickCheck cycle and determines that HSR is in-sync, the resource state will transition to **Standby – In Sync**. At this point, the SAP HANA database is highly-available.

### Takeover with Handshake on the Command Line

Takeover with handshake may be performed on the command line by executing either of the following commands:

- `/opt/LifeKeeper/bin/hana_takeover_with_handshake -t <HANA Tag> -s <Target Server> [-b]`
- `lkcli resource config hana -tag <HANA Tag> -takeover_with_handshake <Target Server>`

The optional `-b` parameter in the first command controls how much of the SAP HANA resource hierarchy is brought in-service on the target server. Without the `-b` option, the entire hierarchy (including all parent resources and resources with shared dependencies) will be brought in-service. With the `-b` option, only

the given SAP HANA resource and its dependencies will be brought in-service.

## Controlling Failback Behavior

By default, if the SAP HANA resource cannot be successfully brought in-service during a takeover with handshake attempt, it will be left in the **Out of Service – Failed (OSF)** state on the target server and will require manual intervention to be brought back in-service. In this scenario, LifeKeeper may also be configured to attempt an automated failback to bring the SAP HANA resource hierarchy back in-service on the previous host. This automated failback behavior can be enabled by setting `HANA_HANDSHAKE_TAKEOVER_FAILBACK=true` in `/etc/default/LifeKeeper`.

## Resuming a Suspended Primary Database

If the protected SAP HANA database becomes suspended on the node where the resource is currently in-service (for example due to an administrator performing a takeover with handshake manually outside of LifeKeeper), the **Active – Suspended** and **Standby – Primary** resource states will be displayed.

Hierarchies	sap-rhel81-1		sap-rhel81-2			
<ul style="list-style-type: none"> <li> Not Active</li> <li> HANA-SHC_HDB00</li> <li> hana-vip</li> </ul>		Active - Suspended	1		Standby - Primary	10
		Active	1		StandBy	10

While in this state, all SAP HANA resource monitoring is suspended and a message similar to the following is logged and broadcast to all open terminals on the system where the SAP HANA resource is currently in-service until the issue is resolved:

```
EMERG:hana:quickCheck:HANA-SHC_HDB00:136377:SAP HANA database HDB00
corresponding to resource HANA-SHC_HDB00 is currently suspended on server sap-
rhel81-1 due to actions performed outside of LifeKeeper. Please take the SAP
HANA resource out of service on server sap-rhel81-1 and bring it in-service on
the server where the database should be registered as primary master. Bringing
resource HANA-SHC_HDB00 back in-service on sap-rhel81-1 will resume the
suspended database. Resource monitoring for HANA-SHC_HDB00 will be suspended
until the issue is resolved.
```

If necessary, the database can be manually resumed by executing the following command on the server where it is suspended:

```
su - <sid>adm -c "hdbnsutil -sr_resumeSuspendedPrimary"
```

After resuming the database with this command, it must also be manually stopped on the standby server where it is registered as primary by executing the following command:

```
su - <sid>adm -c "sapcontrol -nr <InstNum> -function StopWait 600 5"
```

Alternatively, if the standby server is intended to become the new primary server, then the issue can be resolved by switching over the SAP HANA resource to the server where the SAP HANA resource state is **Standby – Primary**.

## 6.18.7.4. Setting Local and Temporal Recovery Policies for SAP HANA Resources

### Local Recovery Policies

In a highly-available SAP HANA deployment utilizing HANA System Replication (HSR), the data from the primary database instance is continuously replicated to a running secondary database instance on a standby server. This running secondary instance acts as a warm standby, providing the ability to switch or fail over the SAP HANA resource to the standby server without the time-consuming requirement of starting the database during the in-service operation. Because of this, there are some situations, especially when protecting very large database instances, where it may be faster to immediately fail over to the standby server instead of restarting the database on the primary server when a failure is detected.

This behavior can be enforced by setting a LifeKeeper local recovery policy for the SAP HANA resource. Several ways of setting a local recovery policy are described below.

### Setting the Local Recovery Policy During Resource Creation

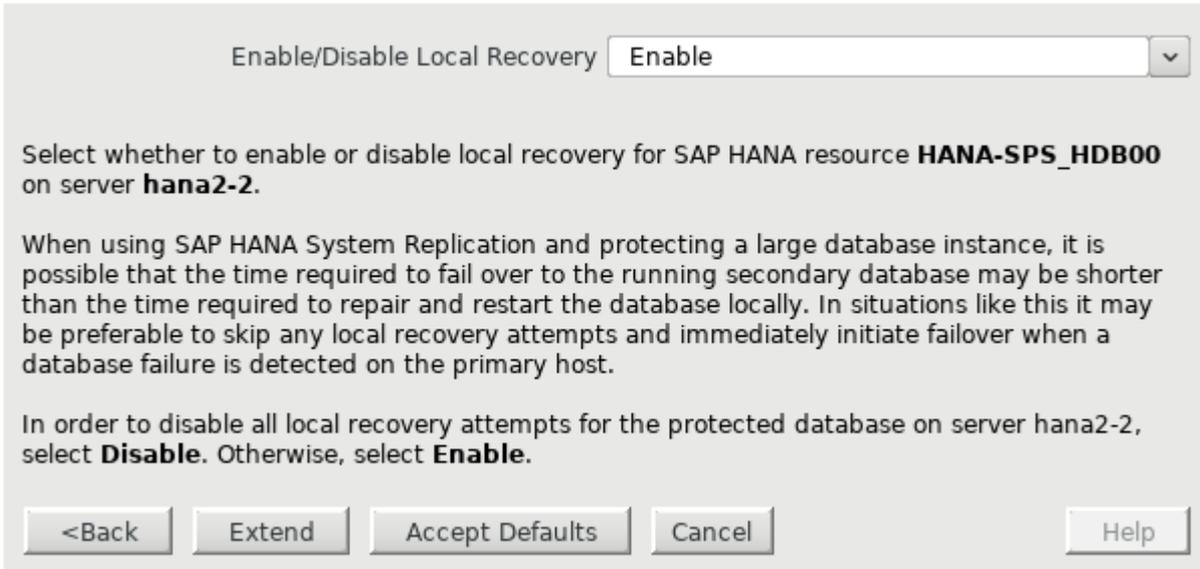
During creation of a SAP HANA resource, a dialog is presented which allows the user to either enable or disable local recovery for the resource. Select **Enable** to allow local recovery attempts on the server where the resource is being created, or **Disable** to skip all local recovery attempts and immediately fail the resource hierarchy over to the standby server when a database failure is detected on this server.



### Setting the Local Recovery Policy During Resource Extension

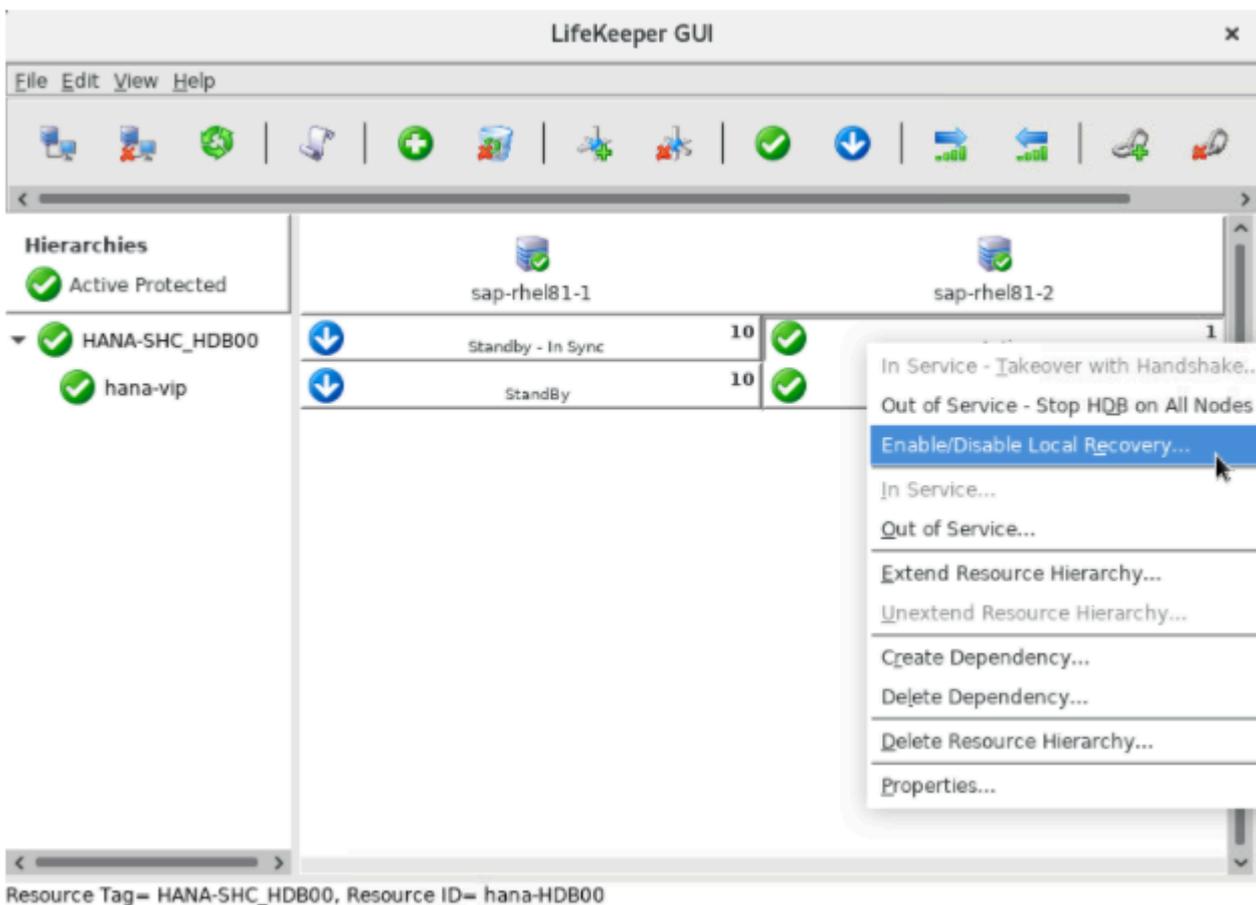
During extension of a SAP HANA resource, a dialog is presented which allows the user to either enable or disable local recovery for the resource. Select **Enable** to allow local recovery attempts on the server that the resource is being extended to, or **Disable** to skip all local recovery attempts and immediately fail

the resource hierarchy to the standby server when a database failure is detected on this server.



### Setting the Local Recovery Policy Using the LifeKeeper GUI

Right-clicking a SAP HANA resource in the LifeKeeper GUI and selecting “**Enable/Disable Local Recovery...**” will bring up a dialog allowing the user to enable or disable local recovery for the resource on the chosen server.



## Setting the Local Recovery Policy Using LKCLI

The following [LKCLI](#) command may be used to enable or disable local recovery for an SAP HANA resource on the server where the command is executed. Replace <HANA Tag> in the command by the tag name of the SAP HANA resource to be configured.

```
# /opt/LifeKeeper/bin/lkcli resource config hana --tag <HANA Tag> --
set_local_recovery_policy <enable|disable>
```

## Setting the Local Recovery Policy Using LKPolicy

The [LKPolicy](#) utility may be used to create a local recovery policy.

Execute the following command to disable local recovery for a given resource on the server where the command is run:

```
# /opt/LifeKeeper/bin/lkpolicy --set-policy LocalRecovery --off tag=<Resource
Tag>
```

Execute the following command to enable local recovery for a given resource on the server where the command is run:

```
# /opt/LifeKeeper/bin/lkpolicy --set-policy LocalRecovery --on tag=<Resource
Tag>
```

 When using the LKPolicy utility for the first time, user credentials must first be configured using the LifeKeeper credential store. See [Configuring Credentials](#) for more information.

## Temporal Recovery Policies

In some situations it may be the case that even though all necessary database processes are able to start successfully, they begin to fail shortly afterwards due to some underlying server issue. If local recovery is enabled for the SAP HANA resource, this sort of situation could lead to a potentially endless cycle of quickCheck failures followed by successful local recovery attempts. Therefore a sequence of several local recovery attempts within a short period of time may be indicative of a server issue, even if all of the attempts are all successful.

To help avoid endless local recovery cycles like this, the [LKPolicy](#) utility may be used to establish a temporal recovery policy on each server. With a temporal recovery policy, LifeKeeper will immediately fail all resource hierarchies from the faulty server to a standby server when it experiences X local recovery attempts within Y minutes (where X and Y are parameters set by the policy).

For example, the following command can be executed on a server to set a temporal recovery policy which will trigger failover after 3 recovery attempts on that server within 10 minutes:

```
# /opt/LifeKeeper/bin/lkpolicy --set-policy TemporalRecovery --on
```

```
recoverylimit=3 period=10
```

The following command may be executed to remove an existing temporal recovery policy:

```
# /opt/LifeKeeper/bin/lkpolicy --remove-policy TemporalRecovery
```

## Local Recovery Enhancements Added To The HANA Recovery Kit

DEMO

# 9.6.0 HANA Feature

LifeKeeper for Linux



<https://fast.wistia.net/embed/iframe/l8v21odnmf>

## 6.18.8. SAP HANA Troubleshooting

---

The [Message Catalog](#) provides a list of all messages that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received, or you may go directly to the [SAP HANA Recovery Kit Message Catalog](#) which contains a list of all messages that may be encountered while utilizing the SAP HANA Recovery Kit.

For information on known issues and restrictions when using the SAP HANA Recovery Kit, see [SAP HANA – Known Issues / Restrictions](#).

# 6.19. SAP MaxDB Recovery Kit Administration Guide

---

SAP MaxDB is a SQL-based, industrial-strength database system that can be deployed for a wide array of purposes. It is highly scalable, platform-independent and provides full transaction support. The database system was originally owned by SAP but has since been released to the Open Source community.

The SAP MaxDB Recovery Kit provides fault resilient protection for SAP MaxDB databases in a LifeKeeper for Linux environment.

## Document Contents

This guide includes the following topics to help you successfully define and manage your SAP MaxDB hierarchy:

- [SAP MaxDB Recovery Kit Requirements](#). Lists the hardware and software necessary to properly set up, install and operate the SAP MaxDB Recovery Kit.
- [Overview](#). Describes the SAP MaxDB Recovery Kit's features and functionality.
- [Configuration Considerations](#). Contains information to consider before you install and configure the SAP MaxDB Recovery Kit.
- [Configuring SAP MaxDB with SAP](#). Provides instructions for installing and configuring the SAP MaxDB software and SAP software.
- [Resource Configuration Tasks](#). Describes the various functions you may perform on your hierarchies using the LifeKeeper GUI: create, extend, delete and unextend.
- [Hierarchy Administration](#). Provides important recommendations for ongoing administration of the SAP MaxDB hierarchy.
- [Troubleshooting](#). Lists and describes the error messages associated with the SAP MaxDB Recovery Kit.
- [Appendix](#). Provides requirements and instructions for setting up raw devices for use with the SAP MaxDB Recovery Kit.

## LifeKeeper Documentation

The following LifeKeeper product documentation is available from the SIOS Technology Corp. website:

[LifeKeeper for Linux Release Notes](#)

[LifeKeeper for Linux Technical Documentation](#)

[LifeKeeper for Linux Installation Guide](#)

[Optional Recovery Kit Documentation](#)

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is provided on the SIOS Technology Corp. website at:

<http://docs.us.sios.com/>

and from the Help menu in the LifeKeeper GUI.

## **SAP MaxDB Documentation**

You can find the SAP MaxDB documentation, including the Installation Guide, User Manual and Reference Manual, at the following locations on the web:

<http://maxdb.sap.com/documentation/>

## 6.19.1. SAP MaxDB Recovery Kit Hardware and Software Requirements

---

Your LifeKeeper configuration must meet the following requirements prior to the installation of LifeKeeper for Linux SAP MaxDB Recovery Kit. Please see the [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the installation and configuration of your LifeKeeper hardware and software.

### Hardware Requirements

- **Servers** – Servers should be configured in accordance with the requirements described in the [LifeKeeper for Linux Installation Guide](#) and the [LifeKeeper for Linux Release Notes](#).
- **IP Network Interface Cards** – Each server requires at least one Ethernet TCP/IP-supported network interface card. Remember, however, that a LifeKeeper cluster requires at least two communications paths. Two separate LAN-based communication paths using dual independent sub-nets are recommended for heartbeats, and at least one of these should be configured as a private network. Using a combination of TCP and TTY heartbeats is also supported.

### Software Requirements

- **TCP/IP Software** – Each server in your LifeKeeper configuration requires TCP/IP Software.
- **SAP MaxDB/MaxDB Software** – Supported versions of SAP MaxDB/MaxDB Software listed in “LifeKeeper for Linux Support Matrix” should be installed.

**Note:** The same version of the SAP MaxDB software must be installed on all servers in the cluster.

- **LifeKeeper Software** – It is imperative that you install the same version of the LifeKeeper software and apply the same versions of the LifeKeeper software patches to each server in your cluster.
- **LifeKeeper for Linux SAP MaxDB / Max DB Recovery Kit** – The SAP MaxDB / Max DB Recovery Kit is provided on the LifeKeeper for Linux Installation Image File (Steeleye-lkSAPDB). It is packaged, installed and removed via Red Hat Package Manager, rpm.

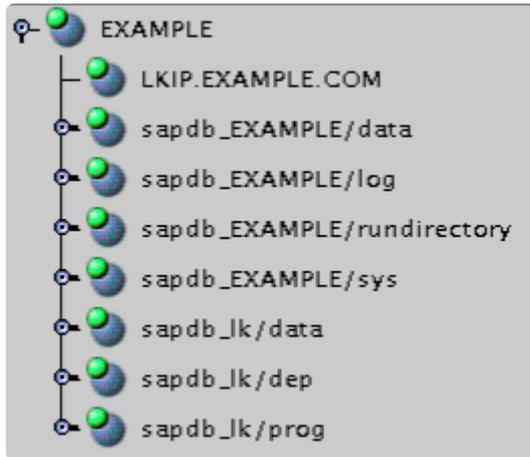
## 6.19.2. SAP MaxDB Recovery Kit Overview

---

The LifeKeeper for Linux SAP MaxDB Recovery Kit provides a mechanism for protecting SAP MaxDB instances within LifeKeeper. The SAP MaxDB software, LifeKeeper Core and SAP MaxDB Recovery Kit are installed on two or more servers in a cluster. Once the SAP MaxDB database instance is under LifeKeeper protection, clients connect to the database using a LifeKeeper protected IP address. The LifeKeeper protected IP address must be created separately and a dependency made manually between the parent SAP MaxDB resource instance and the child IP address resource. In the event that the SAP MaxDB server fails, LifeKeeper will first attempt to recover it on the local server. If the local recovery fails, then LifeKeeper will fail over to a backup server.

## 6.19.2.1. SAP MaxDB Resource Hierarchy

The following example shows a typical SAP MaxDB resource hierarchy:



The dependencies in the above example correspond to the following protected resources:

Resource	SAP MaxDB Software Component
LKIP.EXAMPLE.COM	protects the switchable IP address used for client connections
sapdb_EXAMPLE /data	protects the database data device space for the EXAMPLE database
sapdb_EXAMPLE /log	protects the database log device space for the EXAMPLE database
sapdb_EXAMPLE /rundirectory	protects the database RUNDIRECTORY for the EXAMPLE database
sapdb_EXAMPLE /sys	protects the database system device space for the EXAMPLE database
sapdb_1k/data	protects the independent data path
sapdb_1k/dep	protects the dependent path
sapdb_1k/prog	protects the independent program path

In the event of failover, LifeKeeper will bring the file system, IP address and database resources (including all the resource dependencies) in service on a backup server. Clients will be disconnected and will need to re-connect to the server. Any SQL statement that has not been committed will need to be re-entered.

## 6.19.3. SAP MaxDB Configuration Considerations

---

This section contains information that you should consider before you start to configure and administer the SAP MaxDB Recovery Kit.

[Using Raw I/O](#)

[Using Mirrored File Systems with DataKeeper](#)

[Using Internal Load Balancer](#)

[Active/Standby Considerations](#)

[Active/Standby Configuration Example](#)

[Active/Active Considerations](#)

[Active/Active Configuration Example](#)

## 6.19.3.1. Using Raw I/O with SAP MaxDB

---

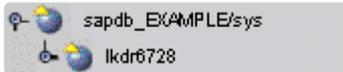
If you plan to use SAP MaxDB with raw devices, you must install the LifeKeeper Raw I/O Recovery Kit from the LifeKeeper Installation Image file. You must also properly set up the raw I/O devices prior to use. See the [Appendix](#) for instructions.

## 6.19.3.2. Using SAP MaxDB Mirrored File Systems with DataKeeper

---

The SAP MaxDB Recovery Kit supports the use of DataKeeper as a device space. In addition, the SAP MaxDB software can be installed on mirrored file systems.

For example, a dependent file system for an SAP MaxDB resource would look similar to the following, which shows a file system for the system device space and its dependency, the DataKeeper resource mirror.



\* Replicated (SIOS DataKeeper) file system resources must be created before creating the SAP MaxDB resource.

## 6.19.3.3. Using Internal Load Balancer

---

The SAP MaxDB Recovery Kit supports the use of Internal Load Balancers (ILB) like [Azure](#) supports. To enable support for ILB, set MAXDB\_ILB\_ENABLED to 1 in the LifeKeeper defaults file (/etc/default/LifeKeeper).

The ILB requires X server ports to be active on one server at a time. This requires all SAP MaxDB resources to be in the same hierarchy so they will be active on the same node. When configured where IndepDataPath and IndepProgPath are on shared file systems all resources will already be forced to the same server. When these are not shared file systems a gen/app resource can be created as a leaf node and/or a root node to keep the hierarchy active on the same server.

 Due to the ILB restriction that the X server must be active on one server, active/active configurations are not supported with ILB.

To place multiple SAP MaxDB resources in a single hierarchy create a gen/app resource as a leaf node at the bottom of the hierarchy and/or as a root node on the top of the hierarchy ([Creating a Generic Application Resource Hierarchy](#) where the restore/remove scripts can be “/bin/true”).

## 6.19.3.4. SAP MaxDB Active/Standby Considerations

---

In an Active/Standby configuration, the backup server is not actively running the SAP MaxDB but stands by in case the primary server experiences a failure. The following scenarios provide specific requirements that must be adhered to when protecting an SAP MaxDB resource instance in Active/Standby configurations.

### Active/Standby Scenarios

The typical Active/Standby configurations are explained below in Scenarios 1 and 2.

#### Scenario 1

The SAP MaxDB *IndepDataPath*, *IndepProgPath* and *DependPath* are installed to **one or more shared file systems on the primary server**.

- The paths *IndepDataPath*, *IndepProgPath*, and *DependPath* must be shared between all servers that will protect the resource instance.
- The registry file */etc/opt/sdb* must exist on each server in the cluster. This file should not be located on shared storage since it must be accessible from the target server during resource extend operations.
- The database instance data device spaces (*data devspaces*), log device spaces (*log devspaces*) and system device spaces (*sys devspaces*) must reside on a shared disk (either shared file system or shared raw device).

#### Scenario 2

The SAP MaxDB *IndepDataPath* and *IndepProgPath* are installed locally on both servers. The SAP MaxDB *DependPath* can be installed locally or on a shared file system on the primary server.

- The registry file */etc/opt/sdb* must exist on each server in the cluster. This file should not be located on shared storage since it must be accessible from the target server during resource extend operations.
- The database instance data device spaces (*data devspaces*), log device spaces (*log devspaces*) and system device spaces (*sys devspaces*) must reside on a shared disk (either shared file system or shared raw device).
- The database instance run directory (*RUNDIRECTORY*) must be located on shared storage. The value of *RUNDIRECTORY* can be modified via the DBMCLI command *param\_directput*. If the value of *RUNDIRECTORY* is modified after the database is created, the database instance must be stopped and restarted to complete the parameter update.

- The database instance config (<IndepDataPath>/config) directory structure must exist in the same location on all servers in the cluster where the database instance will be protected. In addition, the parameter files for the database instance must be copied from the template (or primary) server to all backup servers in the cluster. The parameter files must be redistributed to all servers in the cluster after any parameter has been updated. The required files are:

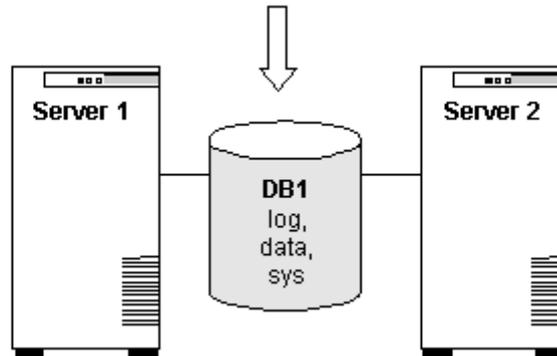
```
config/Databases.ini
```

```
config/Installations.ini
```

## 6.19.3.4.1. Active/Standby Configuration Example

---

IndepData = /shr1/data  
IndepPrograms = /shr1/programs  
DependPath = /shr1/depend



### Configuration Notes:

- Both servers use the *IndepProgPath*, *DependPath* and *IndepDataPath* on the shared storage.
- The database instance *DB1* is located on the shared storage. This includes all log device spaces, data device spaces and system device spaces.
- *Server 2* cannot access files and directories on the shared disk while *Server 1* is active.

## 6.19.3.5. SAP MaxDB Active/Active Considerations

---

In an Active/Active configuration, each server is actively running one SAP MaxDB instance while acting as a backup for the other server in case of failure. The following scenario provides specific requirements that must be adhered to in sequential order when protecting an SAP MaxDB resource instance in an Active/Active configuration.

✿ Internal Load Balancer is not supported in Active/Active configurations.

### Active/Active Scenario

The SAP MaxDB *IndepDataPath*, *IndepProgPath* and *DependPath* are installed locally on both servers.

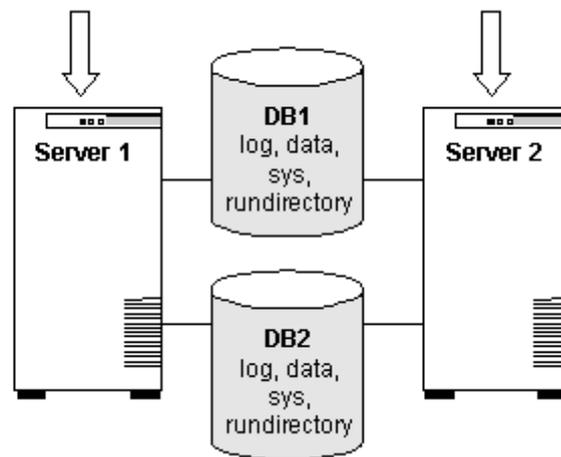
- The registry file */etc/opt/sdb* must exist on each server in the cluster. This file should not be located on shared storage since it must be accessible from the target server during resource extend operations.
- The database instance data device spaces (*data devspaces*), log device spaces (*log devspaces*) and system device spaces (*sys devspaces*) must reside on a shared disk (either shared file system or shared raw device).
- The database instance run directory (*RUNDIRECTORY*) must be located on shared storage. The value of *RUNDIRECTORY* can be modified via the DBMCLI command *param\_directput*. If the value of *RUNDIRECTORY* is modified after the database is created, the database instance must be stopped and restarted to complete the parameter update.
- The database instance config (*<IndepDataPath>/config*) directory structure must exist in the same location on all servers in the cluster where the database instance will be protected. In addition, the parameter files for the database instance must be copied from the template (or primary) server to all backup servers in the cluster. The parameter files must be redistributed to all servers in the cluster after any parameter has been updated. The required files are:

```
config/Databases.ini
```

```
config/Installations.ini
```

## 6.19.3.5.1. Active/Active Configuration Example

IndepData = /usr/sapdb/data IndepPrograms = /usr/sapdb/prog DependPath = /usr/sapdb/dep	IndepData = /usr/sapdb/data IndepPrograms = /usr/sapdb/prog DependPath = /usr/sapdb/dep
-----------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------



### Configuration Notes:

- The *IndepDataPath*, *IndepProgPath* and *DependPath* are locally installed on both servers.
- Each database is configured on separate shared disks. The database instance includes all log device spaces, system device spaces and data device spaces.
- The `RUNDIRECTORY` for each database instance is also on a shared disk.
- The database configuration files for DB1 have been copied to *Server 2* and the database configuration files for DB2 have been copied to *Server 1*. The configuration files are located at `<IndepDataPath>/config/<db instance>`.
- Initially, *Server 1* runs DB1 and *Server 2* runs DB2. In a switchover situation, one server can run both databases.

## 6.19.4. Configuring SAP MaxDB with LifeKeeper

---

The following sequence is recommended for installing and configuring the SAP MaxDB database and LifeKeeper software. Each of these steps links to detailed tasks that follow.

[Install the SAP MaxDB Software](#)

[Create the SAP MaxDB Database](#)

[Create the User\\_Key](#)

[Install the LifeKeeper Core and SAP MaxDB Recovery Kit](#)

After you have performed these tasks, you will be ready to create the LifeKeeper resource hierarchy to protect your SAP MaxDB database.

## 6.19.4.1. Install the SAP MaxDB Software

---

Install the SAP MaxDB software on all servers in the cluster using identical parameters/settings. Refer to the *SAP MaxDB Installation Guide* for details. The following are additional recommendations to ensure that LifeKeeper will work with SAP MaxDB:

- A non-root system user (OS User) must exist on all servers as follows:
  - This OS User should be designated as the owner of the SAP MaxDB software installation and subdirectories or have adequate permissions on the software installation path and subdirectories as required to manage a database instance.
  - This OS User must have authority to use the DBMCLI and x\_server utilities. The OS User must be able to start and stop the vserver using the x\_server commands.
  - The OS User name should contain alpha-numeric characters only.
  - The User ID and Group ID of this OS User must be identical on all servers.
- The SAP MaxDB client software packages must be installed. These packages must include the SAP MaxDB DBMCLI client utility, and the SAP MaxDB x\_server utility.
- Each LifeKeeper server containing an SAP MaxDB resource hierarchy must have identical service entries in the */etc/services* file for the SAP MaxDB instance.

## 6.19.4.2. Create the SAP MaxDB Database

---

Follow the instructions in your *SAP MaxDB User Manual* to create your database. In addition, please note the following recommendations:

- There must be a DBM operator with authority for starting, stopping, obtaining status and obtaining database parameters via client utilities.
- The database instance data device spaces (`data devspaces`), log device spaces (`log devspaces`) and system device spaces (`sys devspaces`) must reside on a shared disk (either shared file system, or shared raw device).
- The SAP MaxDB database name should contain alphanumeric characters only.
- A `User_Key` is required for use by the SAP MaxDB Recovery Kit during operation with the DBMCLI utility. See [Create the User\\_Key](#) for required parameters.
- After creating your database, you should disable automatic startup of the SAP MaxDB database instance. Once under LifeKeeper protection, LifeKeeper will handle the start and stop of the database.
- The `sdb` file must exist on all servers in `/etc/opt`. If this file does not exist, several SAP MaxDB utilities may return erroneous results. This will also affect the behavior of LifeKeeper during resource create and extension. In an Active/Standby configuration, you must manually copy this file to the backup server.
- Verify the `databases.ini` file exists on all servers in the `IndepDataPath/config` directory.
- Copy each database entry in the `databases.ini` file that will be configured with LifeKeeper to the backup server.

## 6.19.4.3. Create the User\_Key

The SAP MaxDB instance requires several options for a user to successfully access a database instance. These required pieces of information must be passed in to the SAP MaxDB tool being used to access the database instance. The SAP MaxDB software includes the **xuser** tool for simplifying the specification of many required options to SAP MaxDB tools. The **xuser** tool allows you to predefine and save user data. Once this data has been saved, it can be used when you call the **DBMCLI** or other tools requiring user options. This predefined user data is stored in a user key (User\_Key). An individual user can manage and maintain several user keys for the same or multiple databases. Each key includes a combination of options including username/password, database name as well as database server name.

The SAP MaxDB Recovery Kit requires a valid User\_Key for each database instance under protection. This User\_Key must be created and accessible by the OS User that owns the database instance. The user information specified for each User\_Key must be for a DBM operator with the following permissions:

- DBStart
- DBStop
- DBInfoRead
- ParamRead

The User\_Key can be generated using the command:

```
xuser -b <file name>
```

where <file name> is the name of a file containing the valid XUSER entries as follows:

Parameter	Parameter Definition
USERKEY	Unique name for the User_Key
USERID	User name of the dbm operator
PASSWORD	Password of the user
SERVERDB	Name of the database instance that this key will refer to
SERVERNODE	The name of the server where the database is running (this should be the DNS or host file entry for the LifeKeeper protected IP)
SQLMODE	This determines what SQL dialects are compatible
CACHELIMIT	This determines the cache limits for a given session
TIMEOUT	Time in seconds before terminating an inactive session (-1 is the default)
ISOLATION	Determines the isolation level used for locks that affect the user (-1 is the default)
DB_LOCAL	Specifies the database locale

Refer to the *SAP MaxDB User Manual* for more information on parameters. Once proper entries have been specified, use the **xuser** tool to generate the `.XUSER.62` file in the OS User home directory. A sample XUSER file is included below containing two keys (an entry must exist for the DEFAULT User\_Key).

```
DEFAULT
NULLDB
NULLDB
NULLDB
LKIP.example.com
INTERNAL
-1
-1
-1
my_locale
LK_USERKEY
LKDBMOPER
LKDBMPASSWD
DB1
LKIP.example.com
INTERNAL
-1
-1
-1
en_US
```

This example XUSER file specifies that two user keys be created, `DEFAULT` and `LK_USERKEY`. Once the **xuser** tool has been run to generate the User\_Key(s), the file specified for use by the **xuser** tool should be deleted.

## 6.19.4.4. Install the LifeKeeper Software

---

Once you have installed the SAP MaxDB software, created your database and created the User\_Key, you are ready to install the LifeKeeper Core software and any required patches followed by the SAP Max DB Recovery Kit. Also, if you plan to use SAP MaxDB with raw devices, you must install the LifeKeeper Raw I/O Recovery Kit from the LifeKeeper Installation Image file. See the [Appendix](#) for requirements and instructions on setting up raw devices.

Refer to the [LifeKeeper for Linux Installation Guide](#) for details on installing the LifeKeeper packages.

## 6.19.5. SAP MaxDB Resource Configuration Tasks

---

Once you have completed the setup tasks described in the previous section, you are ready to create and extend your SAP MaxDB resource hierarchies.

The following tasks are available for configuring the LifeKeeper for Linux SAP MaxDB Recovery Kit:

- **Create Resource Hierarchy** – Creates an SAP MaxDB resource hierarchy
- **Delete Resource Hierarchy** – Deletes an SAP MaxDB resource hierarchy
- **Extend Resource Hierarchy** – Extends an SAP MaxDB resource hierarchy from the primary server to the backup server
- **Unextend Resource Hierarchy** – Unextends (removes) an SAP MaxDB resource hierarchy from a single server in the LifeKeeper cluster

Please refer to your [LifeKeeper for Linux Technical Documentation](#) located on the SIOS Technology website for instructions on configuring your LifeKeeper Core resource hierarchies.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View / Edit Properties](#). View or edit the properties of a resource hierarchy on a specific server.

**Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You may also perform most of the tasks:

- from the toolbar
- by right-clicking on a global resource in the left pane of the status display
- by right-clicking on a resource in the right pane of the status display

Using the right-click method allows you to avoid entering information that is required using the **Edit**

menu.

## 6.19.5.1. Creating an SAP MaxDB Resource Hierarchy

Perform the following steps on the primary server:

1. On the **Edit** menu, select **Server**, then **Create Resource Hierarchy**.

The **Create Resource Wizard** dialog will appear.

2. Select **SAP MaxDB Database** from the drop-down list and click **Enter**.
3. You will be prompted for the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter any error requiring you to correct the previously entered information. You may click **Cancel** at any time to cancel the entire creation process.

Field	Tips
<b>Switchback Type</b>	<p>Choose either intelligent or automatic. This determines how the SAP MaxDB resource will be switched back to the primary server after it comes in-service (active) on the backup server following a failover. Intelligent switchback requires administrative intervention to switch the resource back to the primary server, while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.</p> <p><b>Note:</b> The switchback strategy must match that of the dependent resources to be used by the SAP MaxDB resource.</p>
<b>SAP MaxDB Programs Directory</b>	<p>This field contains by default the SAP MaxDB Program Path found in the SAP /etc/opt/sdb file on the corresponding server. You may type in another directory path. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /</p>
<b>SAP MaxDB Instance</b>	<p>This field contains by default the name of the first SAP MaxDB instance found on the system for which no LifeKeeper hierarchy exists. The drop-down list shows other database instances that may be available on your system.</p>
<b>SAP MaxDB System User</b>	<p>This is the System User that owns or has permission to execute SAP MaxDB commands. This user must exist on the corresponding server. Enter a valid user name in the selection window.</p>
<b>User_Key</b>	<p>This field contains a default value for the XUSER User_Key. The User_Key is used to store database user data for use with SAP MaxDB Tools. Enter a valid User_Key for the corresponding server, OS User and database instance combination.</p>
<b>SAP</b>	<p>This is a unique tag name for the new SAP MaxDB database resource on the primary server.</p>

<b>MaxDB Database Tag</b>	The default tag name consists of the SAP MaxDB instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits and the following special characters: – _ . /
---------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

4. Click **Create**. The **Create Resource Wizard** will then create your SAP MaxDB resource hierarchy. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.
5. You should see a message indicating that you have successfully created an SAP MaxDB resource hierarchy and you must extend that hierarchy to another server in your cluster to achieve failover protection. Click **Next**.
6. Click **Continue**. LifeKeeper will then launch the **Pre-Extend Wizard**. Refer to Step 2 under [Extending an SAP MaxDB Resource Hierarchy](#) for details on how to extend your resource hierarchy to another server.

## 6.19.5.2. Extending an SAP MaxDB Resource Hierarchy

This operation can be started from the **Edit** menu or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should refer to Step 2 below.

1. On the **Edit** menu, select **Resource**, then **Extend Resource Hierarchy**. The **Pre-Extend Wizard** appears. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information.

**Note:** The first two fields appear only if you initiated the Extend from the **Edit** menu.

Field	Tips
<b>Template Server</b>	Select the server where your SAP MaxDB resource is currently in service.
<b>Tag to Extend</b>	Select the SAP MaxDB resource you wish to extend.
<b>Target Server</b>	Enter or select the server you are extending to.
<b>Switchback Type</b>	<p>This determines how the SAP MaxDB resource will be switched back to the primary server after it comes in-service (active) on the backup server following a failover. You can choose either intelligent or automatic. The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.</p> <p>Note: Remember that the switchback strategy must match that of the dependent resources to be used by the SAP MaxDB resource.</p>
<b>Template Priority</b>	<p>Select or enter a Template Priority. This is the priority for the SAP MaxDB hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.</p> <p>Note: This selection will appear only for the initial extend of the hierarchy.</p>
<b>Target Priority</b>	<p>This is the priority for the new extended SAP MaxDB hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given</p>

	resource.
--	-----------

3. After receiving the message that the pre-extend checks were successful, click **Next**.
4. Depending upon the hierarchy being extended, LifeKeeper will display a series of information boxes showing the Resource Tags to be extended, some of which cannot be edited.
5. The **Extend Wizard** will prompt you to enter the following information.

<b>SAP MaxDB Programs Directory</b>	This field contains by default the SAP MaxDB Program Path found in the SAP /etc/opt/sdb file on the corresponding server. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . /
<b>User_Key</b>	This field contains a default value for the XUSER User_Key. The User_Key is used to store database user data for use with SAP MaxDB Tools. Enter a valid User_Key for the corresponding server, OS User and database instance combination.
<b>SAP MaxDB Database Tag</b>	This is a unique tag name for the new SAP MaxDB database resource on the target server. The default tag name consists of the SAP MaxDB instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits and the following special characters: - _ . /

6. After receiving the message "Hierarchy extend operations completed", click **Next Server** to extend the hierarchy to another server, or click **Finish** if there are no other extend operations to perform.
7. After receiving the message "Hierarchy Verification Finished", click **Done**.

## 6.19.5.3. Unextending an SAP MaxDB Resource Hierarchy

---

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. On the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the SAP MaxDB resource. It cannot be the server where the resource is currently in service. (*This dialog box will not appear if you selected the **Unextend** task by right-clicking on a resource instance in the right pane.*) Click **Next**.
3. Select the SAP MaxDB hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the **Unextend** task by right-clicking on a resource instance in either pane).
4. An information box appears confirming the target server and the SAP MaxDB resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the SAP MaxDB resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

## 6.19.5.4. Deleting an SAP MaxDB Resource Hierarchy

---

To delete an SAP MaxDB resource from all servers in your LifeKeeper configuration, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your SAP MaxDB resource hierarchy. *(This dialog will not appear if you selected the **Delete Resource** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance.)*
3. Select the **Hierarchy to Delete**. *(This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a resource instance in the left or right pane.)* Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Next**.
5. Another information box appears confirming that the SAP MaxDB resource was deleted successfully.
6. Click **Done** to exit.

## 6.19.5.5. Testing Your SAP MaxDB Resource Hierarchy

---

You can test your SAP MaxDB resource hierarchy by initiating a manual switchover that will simulate a failover of the resource instance from the primary server to a backup server.

### Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource** and **InService**. For example, an in-service request executed on a backup server causes the SAP MaxDB resource hierarchy to be placed in service on the backup server and taken out of service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the Out of Service request, the resource hierarchy is taken out of service without bringing it in service on the other server.

**!** **IMPORTANT:** After bringing your resource hierarchy in service on the backup server, you should attempt to connect to the databases, especially when using raw devices as device spaces. This is necessary to ensure that all disk partitions are visible on the backup servers and the raw bindings are being established correctly.

If the raw bindings have not been established on the backup servers, it is most likely caused by the fact that new partitions were created on the primary server and added to the configuration, but the partition tables have not yet been updated on the backup servers.

The solution is to reboot the backup servers so that the partition tables are updated correctly.

## 6.19.6. SAP MaxDB Resource Hierarchy Administration

---

The following tasks may be required after your resource hierarchies have been created.

[Modifying User\\_Keys](#)

[Modifying OS User](#)

[Updating Parameters](#)

## 6.19.6.1. Modifying User\_Keys

---

If the User\_Key for an existing hierarchy needs to be changed, the hierarchy must be deleted and recreated.

## 6.19.6.2. Modifying OS User

---

If the OS User that owns the database instance needs to be changed, the hierarchy must be deleted and recreated.

## 6.19.6.3. Updating SAP MaxDB Parameters

---

When database parameters are updated for an SAP MaxDB instance, it is necessary to ensure that the updated parameter files are redistributed to all servers protecting the instance. If the `IndepDataPath` is on a shared disk, then all servers protecting the instance will be updated automatically.

# 6.19.7. SAP MaxDB Troubleshooting

## General Tips

The following error messages may be encountered while using the SAP MaxDB Recovery Kit.

Error Message	Solution
<p>The tunable MAXDB_ILB_ENABLED is set in the LifeKeeper defaults file for Internal Load Balancer support. Protecting multiple SAP MaxDB resources in this configuration requires all SAP MaxDB resource hierarchies to share a common resource.</p>	<p>Create a dependency with a common resource for each SAP MaxDB resource or disable internal Load Balancer support. Some Internal Load Balancers (ILB) like <a href="#">Azure's</a> require X server ports to be active on one server at a time. The LifeKeeper in-service operation for the SAP MaxDB resource will start the X server processes. The normal out-of-service operation of a SAP MaxDB resource will not stop the X server processes. The tunable MAX_ILB_ENABLED can be set in the LifeKeeper defaults file (/etc/default/LifeKeeper) to change the out-of-service behavior so it will stop X server processes when the last SAP MaxDB resource is removed on a system. This requires that all SAP MaxDB resources must be in the same hierarchy to allow LifeKeeper to keep the X Server active on only one server. To place multiple SAP MaxDB resources in a single hierarchy create a gen/app resource as a root node at the top of the hierarchy:</p> <ul style="list-style-type: none"> <li>• On the <b>Edit</b> menu, select <b>Server</b>, then <b>Create Resource Hierarchy</b>. The <b>Create Resource Wizard</b> dialog will appear.</li> <li>• Select <b>Generic Application</b> from the drop-down list and click <b>Enter</b>.</li> <li>• Follow the instruction in <a href="#">Creating a Generic Application Resource Hierarchy</a> where the restore/remove scripts can be “/bin/true”.             <ul style="list-style-type: none"> <li>◦ Once the root resource is created, create a dependency between the root resource with each SAP MaxDB resource.</li> <li>◦ On the <b>Edit</b> menu, select <b>Resource</b>, then <b>Create Dependency</b>.</li> <li>◦ Select the <b>Server</b> where the resources are in-service.</li> <li>◦ Select the <b>Parent Resource Tag</b> for the root resource created above in the drop-down list.</li> <li>◦ Select the <b>Child Resource Tag</b> for the SAP MaxDB resource.</li> </ul> </li> </ul> <p>It is also recommended to have a terminal leaf node. Once the single root resource is created above, the utility create_terminal_leaf can be used to create the gen/app resource on each branch of the hierarchy. If the tag for the root resource created above is “MaxDB_root” then to create the leaf nodes run, “create_terminal_leaf MaxDB_root”. New SAP MaxDB resources will need to be added to the hierarchy by creating dependencies with the root resource and the terminal leaf resource.</p>
<p>Unable to create pipe</p>	<p>The directory <code>/usr/spool/sql</code> must have proper permissions to allow access</p>

<p><code>/usr/spool/sql/ fifo/&lt;db instance&gt;</code></p>	<p>for system user that owns the database instance.</p>
<p>Open device space &lt;dev&gt; permission denied</p>	<p>The device spaces on the backup and primary must have the same owner as well as the same user and group permissions.</p>
<p>Unable to set uid on startup</p>	<p>The setuid bit on <code>&lt;DependPath&gt;/pgm/dbmsrv</code> must be set and the owner of the file must be the SAP MaxDB system user.</p>
<p>runtime environment error</p>	<p>There are several possible causes with different solutions:</p> <ul style="list-style-type: none"> <li>• The database instance parameter and configuration files do not exist. Create the database parameter files or copy the files from the template server.</li> <li>• The database has encountered a library problem. The server and software installation combination may require the use of the library <code>libpthread-0.8.so</code>. Consult the SAP MaxDB documentation for instructions.</li> <li>• The database instance environment has been corrupted. The processes must be manually killed. Then attempt to restore the resource to the in-service state.</li> </ul>
<p>open Registry: Permission denied</p>	<p>The directory <code>/usr/spool/sql/ini</code> should be owned by the system user and group that owns the SAP MaxDB software. In addition, the user and group must also have read/write permissions on the directory.</p>
<p><b>ERR_USRREAD:</b> could not read user data</p>	<p>The config files from <code>&lt;IndepDataPath&gt;/config/&lt;db instance&gt;</code> do not exist on the server or do not have the correct permissions. Verify that the files exist with the correct permissions for the system user that owns the database instance.</p>

## 6.19.7.1. SAP MaxDB Recovery Kit Error Messages

Error Number	Message
111000	Usage: %s independent_program_path <validate:value_1:...:value_n:>
111001	Usage: %s %s %s
111002	No value specified to script %s for input argument %s.
111003	User %s with User_Key %s cannot access instance %s. <b>Action: Specify a User_Key for the given user with database access rights.</b>
111004	The user %s does not exist on the server %s.
111005	The SAP DB instance %s is not running on server %s.
111006	Database Manager Utilities were not found in the specified path %s.
111007	A LifeKeeper internal error occurred in utility %s. <b>Action: Retry operation.</b>
111008	Unable to obtain %s device space information for SAP DB instance %s for user %s and User_Key %s. <b>Action: Verify that the user and User_Key are valid for the corresponding database instance.</b>
111009	Unable to create raw resource hierarchy for %s. <b>Action: Verify that the underlying device is a shared device.</b>
111010	Unable to create filesystem resource hierarchy for %s. <b>Action: Verify that the underlying device is a shared device.</b>
111011	Unable to determine the type of the dev space or install path %s. <b>Action: Valid dev space types include file system and/or raw devices.</b>
111012	The path %s is not on a shared filesystem .

111013	The SAP DB instance %s is already under LifeKeeper protection on server %s.
111014	The SAP DB instance %s has been successfully started on server %s.
111015	The SAP DB instance %s has been successfully stopped on server %s.
111016	Unable to start SAP DB instance %s on server %s.
111017	Unable to stop SAP DB instance %s on server %s.
111018	Attempting db_warm for database instance %s after db_start failure.
111019	The SAP DB x_server has been successfully started on server %s.
111020	The SAP DB x_server has been successfully stopped on server %s.
111021	The SAP DB x_server is not running on server %s
111022	Unable to start SAP DB x_server on server %s. <b>Action: A problem has occurred using the x_server utility, check the SAP DB logs and correct the problem.</b>
111023	Unable to stop SAP DB x_server on server %s. <b>Action: A problem has occurred using the x_server utility; check the SAP DB logs and correct the problem.</b>
111024	The SAB DB file SAP_DBTech.ini was not found on server %s. <b>Action: Verify that SAP DB is installed correctly on the specified server.</b>
111025	The user id for user %s is not the same on server %s and %s.
111026	The group id for user %s is not the same on server %s and %s.
111027	The service file entries for are not the same on server %s and %s.
111028	One or more of the SAP DB service file entries do not exist on server %s.
111029	No dependents were found for resource %s on server %s.
111064	The tunable MAXDB_ILB_ENABLED is set in the LifeKeeper defaults file for Internal Load Balancer support. Protecting multiple SAP DB resources in this configuration requires all SAP DB resource hierarchies to share a common resource. <b>Action: Please refer to <a href="#">SAP MaxDB Troubleshooting</a>.</b>

## 6.19.8. Appendix – Creating Device Spaces Using Raw I/O with SAP MaxDB

---

If you plan to use SAP MaxDB with raw devices, you must install the LifeKeeper Raw I/O Recovery Kit from the LifeKeeper Installation Image file. You must also properly set up the raw I/O devices prior to use.

### Requirements

In order to use the SAP MaxDB Recovery Kit with raw I/O, the following requirements must be met:

- The Linux OS must support raw I/O devices. For most distributions, this support was included in the 2.4 kernel, but there are some distributions that support raw I/O on a 2.2 kernel.
- All raw I/O devices must be bound to a shared disk partition. The number of device spaces (*devspaces*) that will be located on raw I/O devices determines the exact number of raw devices and shared disk partitions required. Refer to the *SAP MaxDB Manual* for guidelines for creating *devspaces* on raw devices.
- The version of the SAP MaxDB software must support the use of raw I/O devices.

## 6.19.8.1. Raw I/O Setup Steps

---

1. Select a shared disk partition of appropriate size for the SAP MaxDB device space.
2. Bind an unused raw device node to this partition. Since this needs to be done every time the machine is rebooted and requires root access, you may want to add the raw bindings to a system initialization file (i.e. *rc.local* or *boot.local*). These bindings must be removed from the file once the hierarchy is under LifeKeeper protection. LifeKeeper will re-establish the raw bindings for raw I/O devices that are under LifeKeeper protection. Use the command *raw -qa* to see which raw device nodes are already in use. For example:

```
# raw -qa
```

```
# raw /dev/raw/raw1 /dev/sda1
```

3. Set global read permissions on both the raw device controller (*/dev/raw/rawctl*) and the disk partition on all servers that will protect the database instance.

```
# chmod a+r /dev/raw/rawctl
```

4. Set group and user read/write permissions on the raw device on all servers that will protect the database instance.

```
# chmod 664 /dev/raw/raw1
```

5. Change the owner of the raw device to the SAP MaxDB OS User for the given database instance on all servers that will protect the database instance.

```
# chown -R sapdb:sapdb /dev/raw/raw1
```

6. Add the device space to the database using *param\_adddevspace* or *db\_adddevspace*. Refer to the *SAP MaxDB User Manual* and/or the *Database Manager CLI Manual*.

## 6.19.8.2. Adding a Device Space after Creating a Hierarchy

---

If a tablespace is added on a raw I/O device or shared file system after the SAP MaxDB hierarchy has been created in LifeKeeper, you must manually create a resource hierarchy for the raw device or file system via the LifeKeeper GUI. The newly created resource hierarchy must then be made a dependent (child) of the SAP MaxDB resource hierarchy. The updated parameter files must be redistributed if necessary to all servers that protect the database instance (*this is not required if the `IndepDataPath` is located on a shared disk*).

## 6.20. Sybase ASE Recovery Kit Administration Guide

---

Sybase Adaptive Server Enterprise is a powerful data management platform for high performance business applications. Sybase ASE is a versatile, enterprise-class RDBMS that is especially good at handling OLTP workloads. Sybase ASE is used widely in financial, E-commerce, and other technology arenas. The Sybase ASE platform includes many standard components, such as the Adaptive Server, Monitor Server, and Backup Server, as well as other plug-in components. The Adaptive Server component is the relational database server. The Monitor Server is a separate server from the database server that monitors the Adaptive Server. The Monitor Server can provide real time or historical data to client applications. The Backup Server is an Open Server-based application that manages all database backup (dump) and restore (load) operations for Adaptive Server.

The LifeKeeper for Linux Sybase ASE Recovery Kit will provide LifeKeeper resource protection for the Sybase ASE components Adaptive Server, Monitor Server, and Backup Server.

### LifeKeeper Documentation

The following is a list of LifeKeeper for Linux related information available from the [SIOS Technology Corp. Documentation](#) site:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [Optional Recovery Kit Documentation](#)

### Sybase ASE Documentation

You can find Sybase ASE documentation, including the *Installation Guide Adaptive Server for Linux*, *User Manual*, *Monitor Server User Manual*, *Troubleshooting Guide* and *Reference Manual(s)* at the following location on the web:

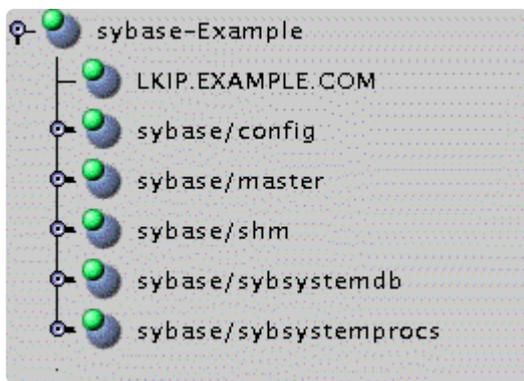
[Sybase Product Documentation](#)

## 6.20.1. Sybase ASE Recovery Kit Overview

The LifeKeeper for Linux Sybase ASE Recovery Kit provides a mechanism for protecting Sybase ASE Server instances within LifeKeeper. The Sybase ASE software, LifeKeeper Core and Sybase ASE Recovery Kit are installed on two or more servers in a cluster. Once the Sybase ASE Server instance is under LifeKeeper protection, clients connect to the database using a LifeKeeper protected IP address. The LifeKeeper protected IP address must be created separately prior to the creation of the Sybase ASE resource hierarchy. The Sybase ASE resource hierarchy creation will create the dependency between the parent Sybase ASE resource instance, and the child IP address resource. In the event that the Sybase ASE Server instance fails, LifeKeeper will first attempt to recover it on the local server. If the local recovery fails, then LifeKeeper will fail over to a backup server.

### Sybase ASE Resource Hierarchy

The following example shows an example Sybase ASE resource hierarchy:



The dependencies in the above example correspond to the following protected resources:

Resource	Sybase ASE Software Component
LKIP.EXAMPLE.COM	Protects the switchable IP address used for client connections
sybase/config	Protects the file system containing the Sybase Adaptive Server, Monitor Server, and Backup Server configuration files
sybase/master	Protects the Sybase ASE master device
sybase/shm	Protects the Sybase Adaptive Server, and Monitor Server shared memory path
sybase/sybsystemdb	Protects the Sybase ASE sybsystemdb device
sybase/sybsystemprocs	Protects the Sybase ASE sybsystemprocs device

In the event of failover, LifeKeeper will bring the file system, IP address and database resources (including all the resource dependencies) in service on a backup server. Clients will be disconnected, and will need to re-connect to the server. Any SQL statement that has not been committed will need to be re-entered.

## 6.20.2. Sybase ASE Recovery Kit Hardware and Software Requirements

---

Your LifeKeeper configuration must meet the following requirements prior to the installation of the LifeKeeper for Linux Sybase Recovery Kit. Please refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the installation and configuration of your LifeKeeper hardware and software.

### Hardware Requirements

- **Servers** – Servers should be configured in accordance with the requirements described in the [LifeKeeper for Linux Technical Documentation](#) and the [LifeKeeper for Linux Release Notes](#).
- **IP Network Interface Cards** – Each server requires at least one Ethernet TCP/IP-supported network interface card. Remember, however, that a LifeKeeper cluster requires at least two communication paths. Two separate LAN-based communication paths using dual independent sub-nets are recommended for heartbeats, and at least one of these should be configured as a private network. Using a combination of TCP and TTY heartbeats is also supported.
- **Storage** – Servers should be configured to use LifeKeeper supported shared storage or the DataKeeper for Linux storage.

### Software Requirements

- **TCP/IP Software** – Each server in your LifeKeeper configuration requires TCP/IP Software.
- **Sybase ASE Software** – LifeKeeper supports version 15.5 and later of the Sybase ASE software. This version can be obtained from Sybase Inc. at <http://www.sybase.com/products/databaseservers/ase>. **Note:** The same version of the Sybase ASE software must be installed on all servers in the cluster. In addition, only one version of the Sybase ASE software may be installed on the LifeKeeper protected servers.
- **LifeKeeper Software** – It is imperative that you install the same version of the LifeKeeper software and apply the same versions of the LifeKeeper software patches to each server in your cluster.
- **LifeKeeper for Linux IP Recovery Kit** – The LifeKeeper for Linux IP Recovery Kit is required by the LifeKeeper for Linux Sybase ASE Recovery Kit. The LifeKeeper for Linux IP Recovery Kit is provided on the LifeKeeper for Linux image file (*sps.img*) via ftp download.
- **LifeKeeper for Linux Sybase ASE Recovery Kit** – The Sybase ASE Recovery Kit (*steel-eye-lkSYBASE*) is provided on the LifeKeeper for Linux Installation image file (*sps.img*) via ftp download. It is installed and removed via this image file.

## 6.20.3. Sybase ASE Recovery Kit Configuration Considerations

---

### Configuration Considerations

Contains information to consider before you install and configure the Sybase ASE Recovery Kit.

---

[Using Raw I-O](#)

[Using Mirrored File Systems with DataKeeper](#)

[Interfaces File Considerations](#)

[Sybase Software Asset Manager](#)

[Active-Standby Considerations](#)

[Active-Active Considerations](#)

[Sybase Monitor Server and Backup Server](#)

[Using Network Attached Storage](#)

## 6.20.3.1. Using Raw I/O with Sybase

---

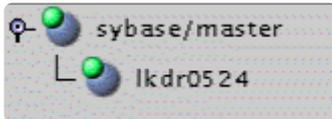
If you plan to use Sybase ASE with raw devices, you must install the LifeKeeper Raw I/O Recovery Kit from the LifeKeeper Core image file. You must also properly set up the raw I/O devices prior to use. See the topic [Creating Database Devices Using Raw I/O](#) for instructions.

## 6.20.3.2. Using Sybase ASE Mirrored File Systems with DataKeeper

---

The Sybase ASE Recovery Kit supports the use of SIOS DataKeeper as a device space. In addition, the Sybase ASE software can be installed on mirrored file systems.

For example, a dependent file system for a Sybase ASE resource would look similar to the following, which shows a file system for the system device space and its dependency, the DataKeeper resource mirror.



- \* Replicated (SIOS DataKeeper) file system resources must be created before creating the Sybase ASE resource.

## 6.20.3.3. Sybase Interfaces File Considerations

---

The Sybase ASE Recovery Kit uses the Sybase ASE interfaces file for the detection of the client IP addresses and ports. This file is located under \$SYBASE and is typically called interfaces. This file is updated whenever an Adaptive Server, Monitor Server or Backup Server instance is created using the `srvbuild` or similar configuration utility. The LifeKeeper for Linux Sybase ASE Recovery Kit requires this file to exist with entries for each Sybase ASE component to be protected. Comment lines are not allowed. All server names that appear in the interfaces file must be resolvable to a valid virtual IP address. All servers that will protect the Sybase ASE resource hierarchy must be able to resolve the server names that appear in the interfaces file. In addition, it is recommended that the virtual IP address be used instead of the server name.

### Example

```
master tcp ether  
example.com 4100
```

```
query tcp ether  
example.com 4100
```

### Example\_back

```
master tcp ether  
example.com 4200
```

```
query tcp ether  
example.com 4200
```

### Example\_mon

```
master tcp ether  
example.com 4200
```

```
query tcp ether
```

example.com 4200

*Sample Interfaces File*

## 6.20.3.4. Sybase ASE Software Asset Manager (SySAM)

---

The Sybase Software Asset Management (SySAM) is used to manage licensed Sybase products. At Sybase ASE server startup, each ASE server component checks the license file in its environment for permission to run specific features. In order for the ASE server to do this, a license manager and vendor module must be running. The LifeKeeper for Linux Sybase ASE Recovery Kit does not provide protection for the SySAM license manager. It is recommended that the license manager be configured in a redundant server system. In the redundant server system, the redundant license allows you to specify local servers as the first license server in the queue, and make remote servers available as backup license servers. The SySAM application attempts to check out a license from a license-file list, starting with the first server. If that server fails for any reason, the second server in the list is contacted, and so on. The `LM_LICENSE_FILE` variable must be set properly in the user profile for the redundant license server environment.

## 6.20.3.5. Sybase ASE Active/Standby Considerations

---

In an Active/Standby configuration, the backup server is not actively running the Sybase ASE but stands by in case the primary server experiences a failure. The following scenarios provide specific requirements that must be adhered to when protecting a Sybase ASE resource instance in active/standby configurations.

### Scenario 1

The Sybase ASE product is installed **locally on all servers in the cluster**.

- All Sybase Adaptive Server, Monitor Server and Backup Server devices are configured on shared storage.
- The Sybase Adaptive Server, Monitor Server and Backup Server configuration files are stored on a shared file system.
- The Sybase Adaptive Server and Monitor Server shared memory directory is located on a shared file system.
- The interfaces file must be manually updated on all servers to contain common entries for each instance to be protected.
- All interfaces file entries must be resolvable by all servers where the resource will be protected.
- The RUN files for each Adaptive Server, Monitor Server and Backup Server must exist on all servers in the cluster.
- The RUN files for each Adaptive Server, Monitor Server and Backup Server must be executable on all servers in the cluster.
- The RUN files for each Adaptive Server, Monitor Server and Backup Server must contain the same options on all servers in the cluster.

### Scenario 2

The Sybase ASE product is installed to **one or more shared file systems on the primary server**.

- All Sybase Adaptive Server, Monitor Server and Backup Server devices are configured on shared storage.
- The Sybase Adaptive Server, Monitor Server and Backup Server configuration files are stored on a shared file system.
- The Sybase Adaptive Server and Monitor Server shared memory directory is located on a shared

file system.

- The interfaces file does not have to be updated on the target servers.
- All interfaces file entries must be resolvable by all servers where the resource will be protected.
- On the LifeKeeper backup server, */etc/ld.so.conf* must be updated to add entries for the Sybase product libraries.

- Add an entry for `$SYBASE/ASE/lib`

- Add an entry for `$SYBASE/OCS/lib`

- Mount the shared file system containing the Sybase ASE installed products and run `ldconfig`

## 6.20.3.6. Sybase ASE Active/Active Considerations

---

In an Active/Active configuration, each server is actively running one or more Sybase ASE Servers, while acting as a backup for the other LifeKeeper server in case of failure. The following scenario provides specific requirements that must be adhered to in sequential order when protecting a Sybase ASE resource instance in an active/active configuration.

### Scenario 1

The Sybase ASE product is installed locally on all servers in the cluster.

- All Sybase Adaptive Server, Monitor Server, and Backup Server devices are configured on shared storage.
- The Sybase Adaptive Server, Monitor Server, and Backup Server configuration files are stored on a shared file system.
- The Sybase Adaptive Server and Monitor Server shared memory directory is located on a shared file system.
- The interfaces file must be manually updated on all servers to contain common entries for each instance to be protected.
- All interfaces file entries must be resolvable by all servers where the resource will be protected.
- The RUN files for each Adaptive Server, Monitor Server and Backup Server must exist on all servers in the cluster.
- The RUN files for each Adaptive Server, Monitor Server and Backup Server must be executable on all servers in the cluster.
- The RUN files for each Adaptive Server, Monitor Server and Backup Server must contain the same options on all servers in the cluster.

### Scenario 2

The Sybase ASE product is installed to one or more shared file systems on the primary server.

- All Sybase Adaptive Server, Monitor Server and Backup Server devices are configured on shared storage.
- The Sybase Adaptive Server, Monitor Server and Backup Server configuration files are stored on a shared file system.
- The Sybase Adaptive Server and Monitor Server shared memory directory is located on a shared

file system.

- The interfaces file does not have to be updated on the target servers.
- All interfaces file entries must be resolvable by all servers where the resource will be protected.
- On the LifeKeeper backup server, */etc/ld.so.conf* must be updated to add entries for the Sybase product libraries.

- Add an entry for `$SYBASE/ASE/lib`

- Add an entry for `$SYBASE/OCS/lib`

- Mount the shared file system containing the Sybase ASE installed products and run `ldconfig`

## 6.20.3.7. Sybase ASE Monitor Server and Backup Server

---

The LifeKeeper for Linux Sybase ASE Recovery Kit provides resource protection for the Adaptive Server, Backup Server, and Monitor Server components. However, the Backup Server and Monitor Server components are not required components of a resource hierarchy. The Sybase Backup Server, and the Sybase Monitor Server can be excluded from the resource protection. During the resource hierarchy creation users that do not wish to protect the Sybase Monitor Server, and/or the Sybase Backup Server can choose none for the respective component choices. Selecting none during the GUI resource creation will exclude the selected component from protection in the resource hierarchy. **Note:** 'none' is a reserved word in the Sybase ASE Recovery Kit, therefore neither the Sybase Backup Server nor the Sybase Monitor Server can be named 'none'.

When choosing whether to protect these components it is important to note that the configuration files that share a common file system with the Adaptive Server configuration files, device paths, log paths, or shared memory directories will be protected by LifeKeeper. If one or more components will not be protected with LifeKeeper, considerations for file placement should be made to prevent sharing between the protected components and the non-protected components.

 **NOTE:** The Sybase Monitor Server is no longer supported with the Sybase ASE ARK v9.0.2 and later.

## 6.20.3.8. Using Network Attached Storage with Sybase ASE

---

There are a couple of special considerations to take into account when configuring LifeKeeper to use an NFS file server (Network Attached Storage) as cluster storage.

### Use the NAS Recovery Kit

The optional Network Attached Storage (NAS) Recovery Kit is required when using an NFS server as a shared storage array with LifeKeeper for Linux. Install the NAS Recovery Kit (and a license) on each cluster node. See the [NAS Recovery Kit](#) documentation for more details.

### Possible Error Message

When using Network Attached Storage (NAS) with Sybase ASE, you may experience Sybase not restarting following a failover due to a system crash. The Sybase error log should indicate the cause of the error.

#### Sybase ASE 15.x

```
00:00:00000:00000:2011/05/09 16:08:51.66 kernel Adaptive Server
Enterprise(Developer Edition)
00:00:00000:00000:2011/05/09 16:08:51.66 kernel basis_dlock: file
\'s10/sybase-data155/data/master.dat\' already in use by an ASE
00:00:00000:00000:2011/05/09 16:08:51.66 kernel kdconfig: unable to
read primary master device
00:00:00000:00000:2011/05/09 16:08:51.66 server kiconfig: read of
config block failed
```

This indicates that the Sybase dataserver has set an NFS lock on the file "*master.dat*" on the NFS file system that is being controlled by LifeKeeper. The lock was not cleared by the system crash, so LifeKeeper is unable to bring the dataserver back into service. Sybase thinks that some other process is using the *master.dat* file.

### Solution

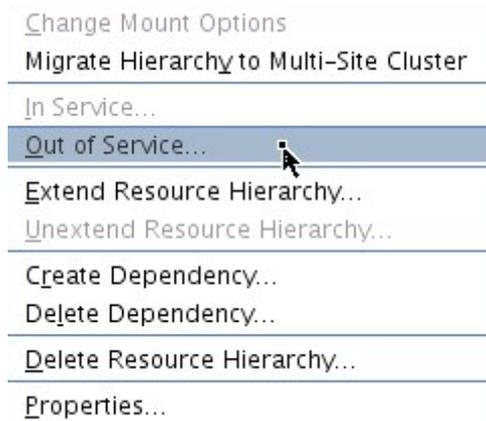
To fix this, mount the NFS file system that will hold *master.dat* with the "nolock" NFS option before the File System resource is created. By default, NFS allows file locks to be set. If the "nolock" option is used before resource creation, LifeKeeper will pick up this option and use it each time it brings the file system resource in service. Since LifeKeeper will be controlling access (from the cluster nodes) to the file system containing *master.dat*, the lock is not typically critical. The NFS mount options used during testing were "rw, sync, tcp, nfsvers=3, noac, nolock".

It is not necessary to use the "nolock" on other file systems used by the Sybase resource hierarchy

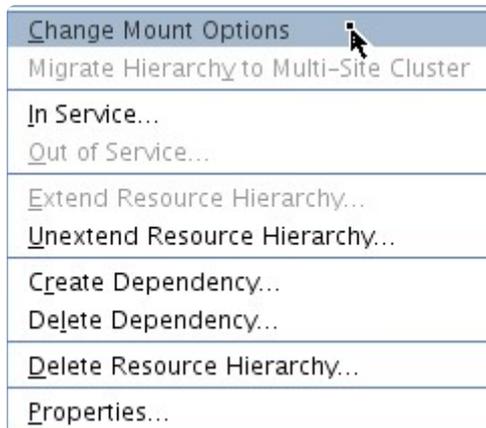
such as the file system where the Sybase ASE binaries are located.

If the NAS File System resource has already been created without the "noLock" option set, use the following procedure to change the mount option:

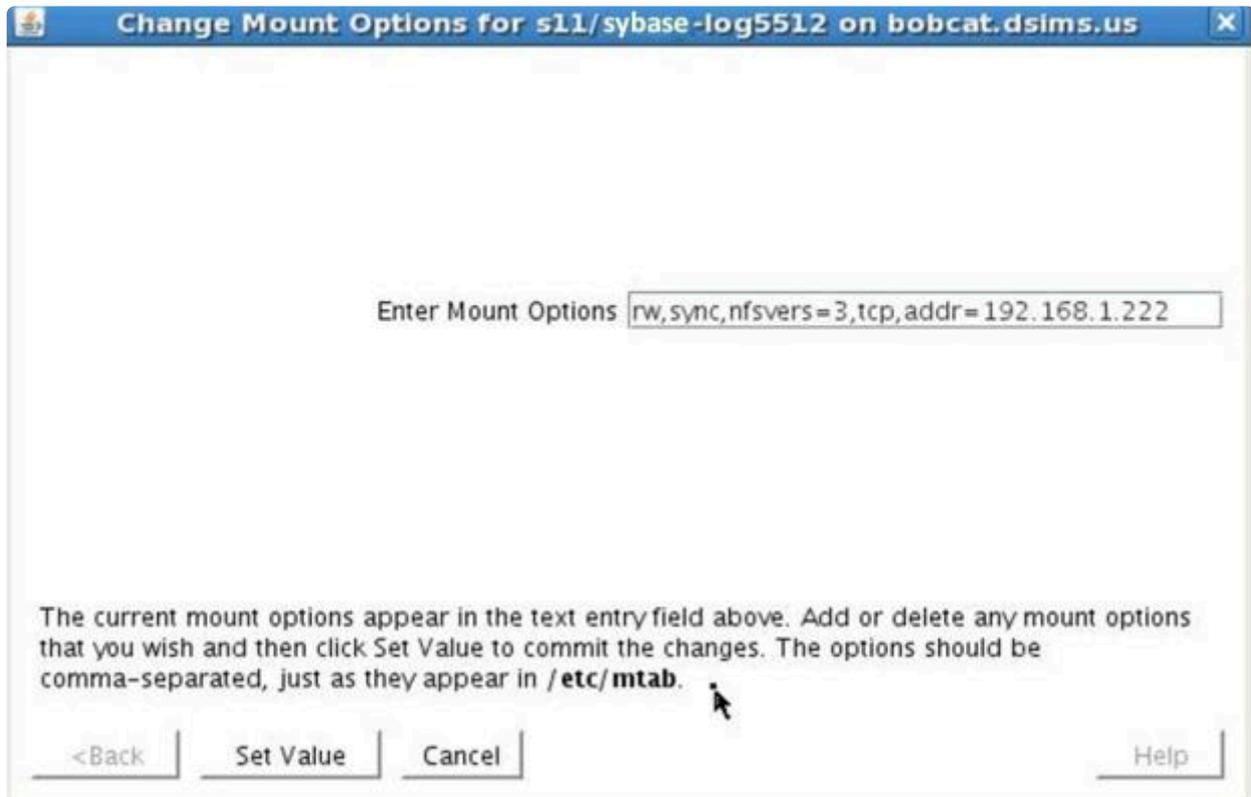
1. Using the LifeKeeper GUI, take the file system resource that needs to be changed out of service. This can be done from the LifeKeeper GUI putting the pointer on the file system resource and doing a right mouse click, and select **"Out of Service"** from the drop-down menu. This action may take parent resources out of service as well.



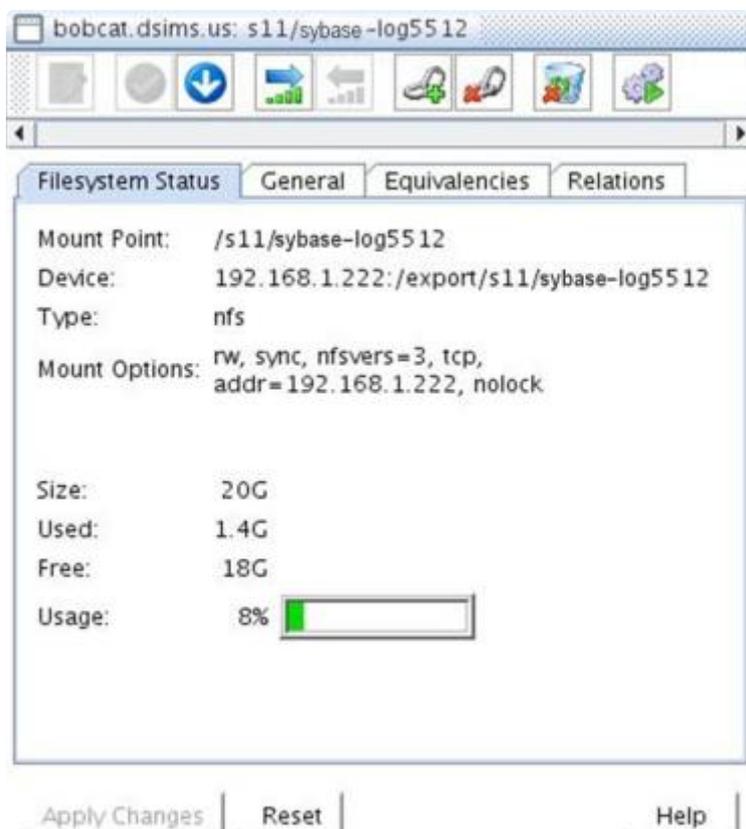
2. Confirm the **"Out of Service"** action and allow the process to complete.
3. Once the file system resource is out of service, you can put the pointer on the resource and do another right mouse click, and from the drop-down menu, select **"Change Mount Options"**.



4. In the popup window, add "noLock" to the line of options, and click **"Set Value."** You will need to repeat steps 3 and 4 for each node in the cluster.



5. Bring the NAS File System resource back in service by doing a right mouse click, and selecting **"In Service"**.
6. The File System resource's property panel should now reflect that "nolock" is one of the current mount options.



## 6.20.4. Installing and Configuring Sybase ASE with LifeKeeper

---

The following sequence is recommended for installing and configuring the Sybase ASE product and LifeKeeper software. Each of these steps links to detailed tasks.

[Install the Sybase ASE Software](#)

[Create the Sybase ASE Servers](#)

[Install the LifeKeeper Software](#)

After you have performed these tasks, you will be ready to create the LifeKeeper resource hierarchy to protect your Sybase ASE Server(s).

### Resource Configuration Tasks

Describes the various functions you may perform on your hierarchies using the LifeKeeper GUI: **create**, **extend**, **delete** and **unextend**.

Once you have completed the setup tasks described in the previous section, you are ready to create and extend your Sybase ASE resource hierarchies.

The following tasks are available for configuring the LifeKeeper for Linux Sybase ASE Recovery Kit:

- [Create Resource Hierarchy](#) – Creates a Sybase ASE resource hierarchy
- [Delete Resource Hierarchy](#) – Deletes a Sybase ASE resource hierarchy
- [Extend Resource Hierarchy](#) – Extends a Sybase ASE resource hierarchy from the primary server to the backup server
- [Unextend Resource Hierarchy](#) – Unextends (removes) a Sybase ASE resource hierarchy from a single server in the LifeKeeper cluster
- [Testing Your Resource Hierarchy](#) – Tests your Sybase ASE resource hierarchy

Refer to the [GUI Administrative Tasks](#) section of the [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring LifeKeeper Core resource hierarchies, for instance, file system and IP resources.

The following tasks are described in the [Administration](#) section within the [LifeKeeper for Linux Technical Documentation](#) because they are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all

applicable servers in the cluster.

- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View /Edit Properties](#). View or edit the properties of a resource hierarchy on a specific server.

 **Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You may also perform most of the tasks:

- from the toolbar
- by right-clicking on a global resource in the left pane of the status display
- by right-clicking on a resource in the right pane of the status display

Using the right-click method allows you to avoid entering information that is required using the **Edit** menu.

## 6.20.4.1. Install the Sybase ASE Software

---

Install the Sybase ASE software on all servers in the cluster using identical parameters/settings. Refer to the *Installation Guide Adaptive Server for Linux* for details. The following are additional recommendations to ensure that LifeKeeper will work with Sybase ASE:

- A non-root system user (Sybase OS User) must exist on all servers. The user must have the same user id, group id, and home directory on all servers where the resource(s) will be protected.
- The Sybase ASE common software packages must be installed. This package provides both the Sybase *srvbuild* and Sybase *isql* utilities.
- Each LifeKeeper server containing a Sybase ASE resource hierarchy must have identical service entries in the `$SYBASE/interfaces` file for the Sybase ASE Server(s).
- Verify that a link exists between `$SYBASE/ASE-<version>` and `$SYBASE/ASE`. If the link does not exist, it must be manually created. See the topic [Creating Links for ASE and OCS](#) for additional information.
- Verify that a link exists between `$SYBASE/OCS-<version>` and `$SYBASE/OCS`. If the link does not exist, it must be manually created. See the topic [Creating Links for ASE and OCS](#) for additional information.
- Refer to the *Installation Guide Adaptive Server for Linux* for details on configuring shared memory parameters for the Adaptive Server, Monitor Server and Backup Server.
- The database device must be protected by LifeKeeper as shared storage. In addition, configuration files and other files must be on the shared storage protected by LifeKeeper. See [Creating the Sybase ASE Servers](#) for details.
- The Sybase ASE common software package should be installed on the shared storage protected by LifeKeeper or on the same path in the local area on all servers in the cluster.

## 6.20.4.2. Create the Sybase ASE Servers

 **NOTE:** The Sybase Monitor Server is no longer supported with the Sybase ASE ARK v9.0.2 and later.

Follow the instructions in your *Installation Guide Adaptive Server for Linux* for configuring the Sybase Adaptive Server, Monitor Server and Backup Server. The following considerations should be followed:

- Use the `srvbuild` utility or other Sybase ASE utility to create the Sybase Adaptive Server instance
  - Configure all system devices on shared storage
  - Configure the Adaptive Server configuration files on shared storage
  - Configure the Adaptive Server shared memory directory on shared storage
  - Configure the interface to use a LifeKeeper switchable IP address
  - Optionally configure the logs on shared storage
- If required, create the Sybase Monitor Server instance
  - Configure all system devices on shared storage
  - Configure the Monitor Server configuration files on shared storage
  - Configure the Monitor Server shared memory directory on shared storage
  - Configure the interface to use a LifeKeeper switchable IP address
  - Optionally configure the logs on shared storage
- If required, create the Sybase Backup Server instance
  - Configure all system devices on shared storage
  - Configure the Monitor Server configuration files on shared storage
  - Configure the Monitor Server shared memory directory on shared storage
  - Configure the interface to use a LifeKeeper switchable IP address
  - Optionally configure the logs on shared storage

## 6.20.4.3. Install the LifeKeeper Software with Sybase

---

Once you have installed the Sybase ASE software and created your database servers, you are ready to install the LifeKeeper Core software, LifeKeeper for Linux IP Recovery Kit and any required patches followed by the Sybase ASE Recovery Kit. Also, if you plan to use Sybase ASE with raw devices, you must install the LifeKeeper Raw I/O Recovery Kit from the LifeKeeper Core image file. See [Creating Device Spaces Using Raw I/O](#) for requirements and instructions on setting up *raw* devices.

Refer to the [LifeKeeper for Linux Installation Guide](#) for details on installing the LifeKeeper packages.

## 6.20.4.4. Creating a Sybase ASE Resource Hierarchy

**Note:** Make sure that the Sybase ASE is running on the primary server.

Perform the following steps on the primary server:

1. On the **Edit** menu, select **Server**, then **Create Resource Hierarchy**.



The **Create Resource Wizard** dialog will appear.

2. Select **Sybase ASE Database** from the drop-down list and click **Next**.
3. You will be prompted for the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter any error requiring you to correct the previously entered information. You may click **Cancel** at any time to cancel the entire creation process.
4. Click **Next**. The **Create Resource Wizard** will then create your Sybase ASE resource hierarchy. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.

Field	Tips
<b>Server</b>	Select the LifeKeeper server where the Sybase ASE resource is to be created.
<b>Switchback Type</b>	Choose either <b>intelligent</b> or <b>automatic</b> . This determines how the Sybase ASE resource will be switched back to the primary server after it comes in-service (active) on the backup server following a failover. Intelligent switchback requires administrative intervention to switch the resource back to the primary server, while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.  <b>Note:</b> The switchback strategy must match that of the dependent resources to be used by the Sybase ASE resource.
<b>Sybase</b>	This field is used to specify the installation location of the Sybase ASE product. You may type

<b>Install Directory</b>	in another directory path. The valid characters allowed for the pathname are letters, digits and the following special characters: – _ . /
<b>Sybase Instance Directory</b>	This field is used to specify the directory path that contains the Sybase data directory. The data directory will typically contain the ASE-<version>/RUN_* files for the instance.
<b>Sybase Instance</b>	This field contains by default the name of the first Sybase instance found on the system, for which no LifeKeeper hierarchy exists. The drop down list shows other Sybase instances that may be available on your LifeKeeper server. This field is used to specify the Sybase ASE Database instance that will be placed under LifeKeeper protection. The specified instance must exist and must be running.
<b>Sybase Username</b>	This field is used to enter the user name for the Sybase System Administrator. By default the user name is sa. This System Administrator must have login and full privileges on any database on the Sybase Adaptive Server being protected.
<b>Sybase Login Password</b>	This field is used to specify the password for the Sybase System Administrator.
<b>Sybase Backup Server</b>	This field is used to specify the Sybase Backup server for the specified Adaptive Server instance. This Sybase Backup will be placed under LifeKeeper protection. The user may select 'none' if the Sybase Backup Server does not need to be included under LifeKeeper protection.
<b>Sybase ASE Database Tag</b>	This is a unique tag name for the new Sybase ASE database resource on the primary server. The default tag name consists of the word sybase followed by the name of the Adaptive Server instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits, and the following special characters: – _ . /

5. You should see a message indicating that you have successfully created a Sybase ASE resource hierarchy, and you must extend that hierarchy to another server in your cluster to achieve failover protection. Click **Next**.
6. Click **Continue**. LifeKeeper will then launch the **Pre-extend Wizard**. Refer to **Step 2** under [Extending a Sybase ASE Resource Hierarchy](#) for details on how to extend your resource hierarchy to another server.

## 6.20.4.5. Extending a Sybase ASE Resource Hierarchy

This operation can be started from the Edit menu or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should refer to Step 2 below.

1. On the Edit menu, select **Resource**, then **Extend Resource Hierarchy**. The **Pre-Extend Wizard** appears. If you are unfamiliar with the **Extend** operation, click **Next**. If you are familiar with the LifeKeeper **Extend Resource Hierarchy** defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information.

 **Note:** The first two fields appear only if you initiated the **Extend** from the **Edit** menu.

3. After receiving the message that the pre-extend checks were successful, click **Next**.

Field	Tips
<b>Template Server</b>	Select the server where your Sybase ASE resource is currently in service.
<b>Tag to Extend</b>	Select the Sybase ASE resource you wish to extend.
<b>Target Server</b>	Enter or select the server you are extending to.
<b>Switchback Type</b>	<p>This determines how the Sybase ASE resource will be switched back to the primary server after it comes in-service (active) on the backup server following a failover. You can choose either intelligent or automatic. The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.</p> <p><b>Note:</b> Remember that the switchback strategy must match that of the dependent resources to be used by the Sybase ASE resource.</p>
<b>Template Priority</b>	<p>Select or enter a Template Priority. This is the priority for the Sybase ASE hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.</p> <p><b>Note:</b> This selection will appear only for the initial extend of the hierarchy.</p>

<b>Target Priority</b>	This is the priority for the new extended Sybase ASE hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.
------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

4. Depending upon the hierarchy being extended, LifeKeeper will display a series of information boxes showing the Resource Tags to be extended, some of which cannot be edited.
5. The **Extend Wizard** will prompt you to enter the following information.
6. After receiving the message "Hierarchy extend operations completed". click **Next Server** to extend the hierarchy to another server, or click **Finish** if there is no other extend operations to perform.

<b>Sybase ASE Install Directory</b>	This field contains by default the Sybase ASE install path of the Template Resource. The valid Sybase ASE installation path should be specified. The valid characters allowed for the pathname are letters, digits, and the following special characters: - _ . /
<b>Sybase ASE Database Tag</b>	This is a unique tag name for the new Sybase ASE database resource on the primary server. The default tag name consists of the word sybase followed by the name of the Adaptive Server instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits, and the following special characters: - _ . /

7. After receiving the message "Hierarchy Verification Finished", click **Done**.

## 6.20.4.6. Unextending a Sybase ASE Resource Hierarchy

---

To remove a resource hierarchy from a single server in the LifeKeeper cluster, do the following:

1. On the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the Target Server where you want to unextend the Sybase ASE resource. It cannot be the server where the resource is currently in service. (This dialog box will not appear if you selected the **Unextend** task by right-clicking on a resource instance in the right pane.) Click **Next**.
3. Select the Sybase ASE hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the Unextend task by right-clicking on a resource instance in either pane.)
4. An information box appears confirming the target server and the Sybase ASE resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the Sybase ASE resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

## 6.20.4.7. Deleting a Sybase ASE Resource Hierarchy

---

To delete a Sybase ASE resource from all servers in your LifeKeeper configuration, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the name of the Target Server where you will be deleting your Sybase ASE resource hierarchy.

 **Note:** If you selected the **Delete Resource** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

3. Select the Hierarchy to Delete. (This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a resource instance in the left or right pane.) Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Next**.
5. Another information box appears confirming that the Sybase ASE resource was deleted successfully.
6. Click **Done** to exit.

## 6.20.4.8. Testing Your Sybase ASE Resource Hierarchy

---

You can test your Sybase ASE resource hierarchy by initiating a manual switchover that will simulate a failover of the resource instance from the primary server to a backup server.

### Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource** and **In Service**. For example, an in-service request executed on a backup server causes the Sybase ASE resource hierarchy to be placed in service on the backup server and taken out-of-service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the **Out of Service** request, the resource hierarchy is taken out-of-service without bringing it in service on the other server.

-  **IMPORTANT:** After bringing your resource hierarchy in service on the backup server, you should attempt to connect to the databases, especially when using raw devices as device spaces. This is necessary to ensure that all disk partitions are visible on the backup servers and the raw bindings are being established correctly.
- If the raw bindings have not been established on the backup servers, it is most likely caused by the fact that new partitions were created on the primary server and added to the configuration, but the partition tables have not yet been updated on the backup servers.
  - The solution is to reboot the backup servers so that the partition tables are updated correctly.

## 6.20.5. Sybase ASE Recovery Kit Administration

---

### Resource Hierarchy Administration

Provides important recommendations for ongoing administration of the Sybase ASE hierarchy.

The following tasks may be required after your resource hierarchies have been created.

[Modifying Protection for the Sybase Backup Server](#)

[Modifying Protection for the Sybase Monitor Server](#)

[Updating Parameters](#)

## 6.20.5.1. Modifying Protection for the Sybase Backup Server

---

The Sybase Backup Server is an Open Server-based application that manages all database backups (dump) and restores (load) operations for Adaptive Server. The Sybase Backup Server can be protected by the LifeKeeper Sybase ASE resource hierarchy during the resource creation, or added to the LifeKeeper protection after the resource hierarchy creation. In addition, the Sybase Backup Server can be removed from LifeKeeper protection after the hierarchy has been created.

### Adding a Sybase Backup Server

To add a Sybase Backup Server to an existing Sybase ASE resource hierarchy, the Sybase `srvbuild` or other configuration utility must have created one.

1. On the **Edit** menu, select **Resource**, then select **Properties**. A **Resource Properties Wizard** will appear.
2. Select the resource tag from the **Select Resource** drop-down. This is the resource tag for the LifeKeeper protected Sybase ASE resource to modify.
3. Select the LifeKeeper Server from the **Select Server for Resource** drop-down. This will be the server to update the Sybase ASE resource instance on. If changes are required on more than one LifeKeeper server, then this process should be repeated for each server in the cluster.
4. Select the **Resource Configuration** button on the **Resource Properties** page. This will launch a **Reconfiguration Wizard** for the protected resource selected in Step 3. The first screen of the wizard will display the current configuration settings for the resource under LifeKeeper protection. Select **Next**.
5. If a valid Sybase Backup Server exists on the specified server, the next screen will display a drop-down for the Sybase Backup Server to add or remove. Select the Sybase Backup Server to add from the list. Select **Next**. **Note:** For Sybase ASE installations where the Sybase software is installed on shared storage, the file system containing the installation must be in service on the server where the reconfiguration will take place.
6. If a valid Sybase Monitor Server exists, the next screen will allow you to configure it now. Refer to [Modifying Protection for the Sybase Monitor Server](#) for considerations regarding modifying the Monitor Server protection.
7. Select **Reconfigure**. If any errors are displayed they must be corrected before proceeding. Otherwise, select **Done**.
8. Any Sybase Backup Server configuration file paths or associated database devices should be manually protected with a LifeKeeper file system resource and made a dependent of the parent resource hierarchy.

9. The virtual IP address associated with the Sybase Backup Server must be made a dependent of the parent resource hierarchy. To find the associated IP address, look for the master and query lines following the Sybase Backup Server name in the interfaces file.

## Removing a Sybase Backup Server

The following steps outline the process for removing a Sybase Backup Server from an existing Sybase ASE resource hierarchy.

1. On the **Edit** menu, select **Resource**, select **Properties**. A **Resource Properties Wizard** will appear.
2. Select the resource tag from the **Select Resource** drop-down. This is the resource tag for the LifeKeeper protected Sybase ASE resource to modify.
3. Select the LifeKeeper Server from the **Select Server for Resource** drop-down. This will be the server to update the Sybase ASE resource instance on. If changes are required on more than one LifeKeeper server, then this process should be repeated for each server in the cluster.
4. Select the **Resource Configuration** button on the **Resource Properties** page. This will launch a **Reconfiguration Wizard** for the protected resource selected in Step 3. The first screen of the wizard will display the current configuration settings for the resource under LifeKeeper protection. Select **Next**.
5. If a valid Sybase Backup Server exists on the specified server, the next screen will display a drop-down for the Sybase Backup Server to add or remove. Select 'none' from the list to remove protection for the Sybase Backup Server. Select **Next**.
6. If a valid Sybase Monitor Server exists, the next screen will allow you to configure it now. Refer to [Modifying Protection for the Sybase Monitor Server](#) for considerations regarding modifying the Monitor Server protection.
7. Select **Reconfigure**. If any errors are displayed, they must be corrected before proceeding. Otherwise, select **Done**.
8. Any Sybase Backup Server configuration file paths or associated database devices that are no longer in use should be removed from the Sybase ASE resource dependency and deleted from LifeKeeper.
9. Any Sybase Backup Server virtual IP resources that are no longer in use should be removed from the Sybase ASE resource dependency and deleted from LifeKeeper.

## 6.20.5.2. Modifying Protection for the Sybase Monitor Server

---

\* **NOTE:** The Sybase Monitor Server is no longer supported with the Sybase ASE ARK v9.0.2 and later.

The Monitor Server is a separate server from the database server that monitors the Adaptive Server. The Monitor Server can provide real time or historical data to client applications. The Sybase Monitor Server can be protected by the LifeKeeper Sybase ASE resource hierarchy during the resource creation, or added to the LifeKeeper protection after the resource hierarchy creation. In addition, the Sybase Monitor Server can be removed from LifeKeeper protection after the hierarchy has been created.

### Adding a Sybase Monitor Server

To add a Sybase Monitor Server to an existing Sybase ASE resource hierarchy, the Sybase srvbuild or other configuration utility must have created one.

1. On the **Edit** menu, select **Resource**, select **Properties**. A **Resource Properties Wizard** will appear.
2. Select the resource tag from the **Select Resource** drop-down. This is the resource tag for the LifeKeeper protected Sybase ASE resource to modify.
3. Select the LifeKeeper Server from the **Select Server for Resource** drop-down. This will be the server to update the Sybase ASE resource instance on. If changes are required on more than one LifeKeeper server, then this process should be repeated for each server in the cluster.
4. Select the **Resource Configuration** button on the **Resource Properties** page. This will launch a **Reconfiguration Wizard** for the protected resource selected in Step 3. The first screen of the wizard will display the current configuration settings for the resource under LifeKeeper protection. Select **Next**.
5. If a valid Sybase Backup Server exists, the next screen will allow you to configure it now. Refer to [Modifying Protection for the Sybase Backup Server](#) for considerations regarding modifying the Backup Server protection.
6. If a valid Sybase Monitor Server exists on the specified server, the next screen will display a drop-down for the Sybase Monitor Server to add or remove. Select the Sybase Monitor Server to add from the list. Select **Next**. **Note:** For Sybase ASE installations where the Sybase software is installed on shared storage, the file system containing the installation must be in-service on the server where the reconfiguration will take place.
7. Select **Reconfigure**. If any errors are displayed they must be corrected before proceeding. Otherwise, select **Done**.

8. Any Sybase Monitor Server configuration file paths or associated database devices should be manually protected with a LifeKeeper file system resource and made a dependent of the parent Sybase ASE resource hierarchy.
9. The virtual IP address associated with the Sybase Monitor Server must be made a dependent of the parent Sybase ASE resource hierarchy. To find the associated IP address, look for the master and query lines following the Sybase Monitor Server name in the interfaces file.

## Removing a Sybase Monitor Server

The following steps outline the process for removing a Sybase Monitor Server from an existing Sybase ASE resource hierarchy.

1. On the **Edit** menu, select **Resource**, select **Properties**. A **Resource Properties Wizard** will appear.
2. Select the resource tag from the **Select Resource** drop-down. This is the resource tag for the LifeKeeper protected Sybase ASE resource to modify.
3. Select the LifeKeeper Server from the Select Server for Resource pull down. This will be the server to update the Sybase ASE resource instance on. If changes are required on more than one LifeKeeper server, then this process should be repeated for each server in the cluster.
4. Select the **Resource Configuration** button on the **Resource Properties** page. This will launch a **Reconfiguration Wizard** for the protected resource selected in Step 3. The first screen of the wizard will display the current configuration settings for the resource under LifeKeeper protection. Select **Next**.
5. If a valid Sybase Backup Server exists, the next screen will allow you to configure it now. Refer to [Modifying Protection for the Sybase Backup Server](#) for considerations regarding modifying the Backup Server protection
6. If a valid Sybase Monitor Server exists on the specified server, the next screen will display a pull down for the Sybase Monitor Server to add or remove. Select 'none' from the list to remove protection for the Sybase Monitor Server. Select **Next**.
7. Select **Reconfigure**. If any errors are displayed they must be corrected before proceeding. Otherwise, select **Done**.
8. Any Sybase Monitor Server configuration file paths or associated database devices that are no longer in use should be removed from the Sybase ASE resource dependency and deleted from LifeKeeper.
9. Any Sybase Monitor Server virtual IP resources that are no longer in use should be removed from the Sybase ASE resource dependency and deleted from LifeKeeper.

## **6.20.5.3. Updating Sybase ASE Parameters**

---

When database parameters are updated for a Sybase ASE instance, it is necessary to check that all changes will allow the instance to function on all LifeKeeper servers in the cluster. If changes require the addition or deletion of LifeKeeper resources, such as file systems, raw devices or virtual IP addresses, these must be added manually and made a dependency of the parent Sybase ASE resource hierarchy.

## 6.20.6. Troubleshooting Sybase ASE Error During Resource Creation

### Sybase ASE Error During Resource Creation

**Symptom:** Unable to create the resource instance during the resource creation

**Cause:** If the instance is running, it could be because the profile is not located in the default \$Sybase directory

**Solution:** In /etc/default/LifeKeeper set the tunable SYBASE\_PROFILE to the location of the correct SYBASE.sh profile

For example: Add to /etc/default/LifeKeeper

```
SYBASE_PROFILE=/opt/my-non-standard-path/SYBASE.sh
```

**Symptom:** LifeKeeper Sybase Resource fails to come in-service but the database instance is started.

**Cause:** The instance took longer than the default start up time to complete its startup and recovery process.

**Solution:** Increase the start wait tunable via the /etc/default/LifeKeeper file

For example: Add to /etc/default/LifeKeeper

```
SYBASE_STARTWAIT=120
```

### Sybase ASE Recovery Kit Error Messages

Lists and describes the error messages associated with the Sybase ASE Recovery Kit.

114000	Usage: %s
114001	The Sybase Install Directory cannot be empty. <b>ACTION: Please specify a value for this field.</b>
114002	The path %s is not a valid directory
114003	The Sybase Product was not found in the directory %s on server %s. <b>ACTION: Verify that a supported version of Sybase is installed in the specified location.</b>
114004	The specified instance %s is not a valid Sybase ASE Server on %s.
114005	Unable to verify that the Sybase ASE Server %s is running.
114006	The Sybase Monitor Server %s will be protected.
114007	The Sybase Backup Server %s will be protected.
114008	The Sybase ASE Server %s is already under LifeKeeper protection on %s.

114009	An unknown error has occurred in utility %s on server %s. <b>ACTION: View the LifeKeeper logs for details and retry the operation.</b>
114010	Unable to get the version for the Sybase Server %s installed under %s on %s.
114011	The device %s for Sybase ASE Server %s is not a valid device.
114012	An error has occurred while trying to obtain the devices for Sybase ASE Server %s.
114013	Unable to create raw resource hierarchy for %s.
114014	Unable to create file system resource hierarchy for %s.
114015	The path %s is not on a shared file system.
114016	Unable to create resource dependency for parent %s and child %s.
114017	Information: LifeKeeper will not protect the path %s because it is not located on a shared file system.
114018	Unable to get the owner for the Sybase ASE Server %s installed under %s on %s.
114019	Unable to open file %s on server %s due to error %s.
114020	There are no hosts defined for the Sybase ASE Server %s in the file %s.
114021	There are no ports defined for the Sybase ASE Server %s in the file %s.
114022	The specified host name %s defined for the Sybase ASE Server %s in the file %s cannot be resolved.
114023	Unable to detect the host and ports for the Sybase ASE Server %s.
114024	A LifeKeeper resource hierarchy does not exist for the IP address %s on server %s. <b>ACTION: Create a LifeKeeper resource hierarchy for the specified IP address</b>
114025	The values specified for the target and the template servers are the same. <b>ACTION: Please specify the correct values for the target and template servers.</b>
114026	The system user %s does not exist on the server %s.
114027	The group id for user %s is not the same on template server %s and target server %s.
114028	The user id for user %s is not the same on template server %s and target server %s.
114029	There are no IP dependent resources defined for the Sybase resource %s on %s. <b>ACTION: Create the required dependent IP resource hierarchy.</b>
114030	The interfaces defined for Sybase ASE Server %s differ on template server %s and target server

	%s
114031	The ports defined for Sybase ASE Server %s differ on template server %s and target server %s
114032	The port %s used by the Sybase resource hierarchy %s on the server %s is in use by another application on server%s.
114033	The startup of the Sybase ASE Server(s) on %s failed for the following Sybase ASE Server(s): %s.
114034	Unable to stop the Sybase ASE Server(s) %s on %s.
114035	The Sybase ASE resource hierarchy %s does not contain any valid gen/filesys or scsi/raw resource dependents on server %s. <b>ACTION: The hierarchy does not contain any valid dependents, you must delete and recreate the hierarchy.</b>
114036	There are no Sybase ASE Servers available for protection with LifeKeeper.
114037	Unable to obtain the pid of the backupserver process corresponding to instance %s.
114038	The pid detected for Sybase Backup Server %s in the LifeKeeper pidfile %s.LK on server %s exists in another LifeKeeper pidfile on this server. <b>ACTION: The duplicate pid entry in the pid files should be resolved. The pid file for the instance that is not running should be removed.</b>
114039	Unable to update the resource instance %s on server %s.
114040	The update of the resource instance %s failed on server %s. All attempts to rollback the instance information field have failed. <b>ACTION: Manual intervention is required.</b>
114041	The interfaces file %s on %s contains an invalid comment line. <b>ACTION: Please correct the interfaces file to remove any comment lines.</b>
114042	One or more of the Sybase ASE Servers is missing from the file %s.
114043	The file %s does not exist on server %s.
114044	The reconfiguration of the Sybase ASE resource hierarchy %s on server %s was successful
114045	The update of the resource instance %s failed on server %s. The instance information field has not been modified. <b>ACTION: Retry the reconfiguration operation</b>

114046	The home directory for user %s is not the same on template server %s and target server %s
114047	The file %s on server %s is a link that does not resolve to a dependent shared resource on the template server %s.
114048	The link %s and its resolved path %s are not on a protected shared filesystem.

## 6.20.7. Appendix – Creating Device Spaces Using Raw I/O with Sybase ASE

---

### Creating Device Spaces Using Raw I/O

[Requirements](#)

[Naming Conventions](#)

[Raw I-O Setup Steps](#)

[Adding a Database Device After Creating Hierarchy](#)

[Creating Links for ASE and OCS](#)

## 6.20.7.1. Requirements for Using Sybase ASE with Raw I/O

---

In order to use the Sybase ASE Recovery Kit with raw I/O, the following requirements must be met:

- The Linux OS must support raw I/O devices. For most distributions this support was included in the 2.4 kernel, but there are some distributions that support raw I/O on a 2.2 kernel.
- All raw I/O devices must be bound to a shared disk partition. The number of database devices (devspaces) that will be located on raw I/O devices determines the exact number of raw devices and shared disk partitions required. Refer to the *Installation Guide Adaptive Server for Linux* for guidelines for creating database devices on raw devices.
- The version of the Sybase ASE software must support the use of raw I/O devices.

## 6.20.7.2. Naming Conventions

---

The naming of raw devices and controller varies by Linux distribution.

- On Red Hat, the device name is `/dev/raw/raw<number>` and the controller is `/dev/rawctl`
- On SuSE SLES 11 versions, the device name is `/dev/raw/raw<number>` and the controller is `/dev/raw/rawctl`

## 6.20.7.3. Using Raw I/O with Sybase Setup Steps

---

1. Select a shared disk partition of appropriate size for the Sybase ASE database device.
2. Bind an unused raw device node to this partition. Since this needs to be done every time the machine is rebooted, and requires root access, you may want to add the raw bindings to a system initialization file (i.e. rc.local or boot.local). These bindings must be removed from the file once the hierarchy is under LifeKeeper protection. LifeKeeper will re-establish the raw bindings for raw I/O devices that are under LifeKeeper protection. Use the command `raw -qa` to see which raw device nodes are already in use. For example:

```
# raw -qa
```

```
# raw /dev/raw/raw1 /dev/sda1
```

3. Set global read permissions on both the raw device controller (`/dev/rawctl` or `/dev/raw/rawctl`), and the disk partition on all servers that will protect the database instance.

```
# chmod a+r /dev/rawctl (or chmod a+r /dev/raw/rawctl)
```

4. Set group and user read/write permissions on the raw device on all servers that will protect the database instance.

```
# chmod 664 /dev/raw/raw1
```

5. Change the owner of the raw device to the Sybase ASE owner for the given database instance on all servers that will protect the database instance.

```
# chown -R sybase:sybase /dev/raw/raw1
```

6. Refer to the *Installation Guide Adaptive Server for Linux* for information on adding the raw device to the database server(s).

## 6.20.7.4. Adding a Device Space after Creating a Sybase Hierarchy

---

If a database device is added on a raw I/O device or shared file system after the Sybase ASE hierarchy has been created in LifeKeeper, you must manually create a resource hierarchy for the raw device or file system via the LifeKeeper GUI. The newly created resource hierarchy must then be made a dependent (child) of the Sybase ASE resource hierarchy.

## 6.20.7.5. Creating Links for ASE and OCS

The LifeKeeper for Linux Sybase ASE Recovery Kit requires that the path `$$SYBASE/ASE-<version>` be symbolically linked to `$$SYBASE/ASE`. In addition, the path `$$SYBASE/OCS-<version>` must be symbolically linked to `$$SYBASE/OCS`. The LifeKeeper for Linux Sybase ASE Recovery Kit uses these links to access various Sybase utilities and files. To create the links follow the steps below.

1. From the command line, change directories into the `$$SYBASE` directory.

Example:

```
server1 # cd $$SYBASE

server1 # pwd

/opt/sybase-15.5
```

2. Locate the `ASE-<version>` directory

Example:

```
server1 # ls -ld ASE*

drwxrwxr-x 16 sybase sybase 4096 Nov 18 09:08 ASE-15_5

lrwxrwxrwx 1 sybase sybase 8 Nov 17 11:35 ASEP -> ASEP-1_0

drwxrwxr-x 4 sybase sybase 4096 Nov 17 11:35 ASEP-1_0p
```

 **Note:** If a link already exists between `ASE-15_5` and `ASE`, proceed to Step 5.

3. Verify that the `ASE-<version>` directory contains the `bin/srvbuild` utility.

Example:

```
server1 # ls ASE-15_5/bin/srvbuild

srvbuild
```

 **Note:** If a “no such file or directory” error occurs, then you have chosen the wrong path.

4. From the command line, create a link between the identified `ASE-<version>` directory and `ASE`.

Example:

```
server1 # pwd

/opt/sybase-15.5

server1 # ln -s ASE-15_5 ASE
```

5. Verify the link was properly created.

Example:

```
server1 # ls -ld ASE*

lrwxrwxrwx 1 sybase sybase 8 Nov 17 11:20 ASE -> ASE-15_5

drwxrwxr-x 16 sybase sybase 4096 Nov 18 09:08 ASE-15_5

lrwxrwxrwx 1 sybase sybase 8 Nov 17 11:35 ASEP -> ASEP-1_0

drwxrwxr-x 4 sybase sybase 4096 Nov 17 11:35 ASEP-1_0

server1 # ls ASE/bin/srvbuild

srvbuild
```

6. From the command line, change directories into the *\$SYBASE* directory.

Example:

```
server1 # cd $SYBASE

server1 # pwd

/opt/sybase-15.5
```

7. Locate the *OCS-<version>* directory

Example:

```
server1 # ls -ld OCS*

drwxrwxr-x 16 sybase sybase 4096 Nov 18 09:08 OCS-15_5
```

 **Note:** If a link already exists between *OCS-15\_5* and *OCS*, proceed to Step 5.

8. Verify that the *OCS-<version>* directory contains the *bin/isql* utility.

Example:

```
server1 # ls OCS-15_5/bin/isql

isql
```

 **Note:** If a “no such file or directory” error occurs, then you have chosen the wrong path.

9. From the command line, create a link between the identified *OCS-<version>* directory and *OCS*.

Example:

```
server1 # pwd

/opt/sybase-15.5

server1 # ln -s OCS-15_5 OCS
```

10. Verify the link was properly created.

Example:

```
server1 # ls -ld OCS*

lrwxrwxrwx 1 sybase sybase 8 Nov 17 11:20 OCS -> OCS-15_5

drwxrwxr-x 16 sybase sybase 4096 Nov 18 09:08 OCS-15_5

server1 # ls OCS/bin/isql

isql
```

 **Note:** Version *15\_5* is used only as an example.

# 6.21. VMDK Shared Storage Recovery Kit Administration Guide

---

## VMDK Shared Storage Recovery Kit Technical Documentation

LifeKeeper for Linux VMDK Shared Storage Recovery Kit (hereafter referred to as the VMDK Recovery Kit) provides a VMware virtual hard disk as shared storage. The VMDK Recovery Kit allows LifeKeeper users to employ virtual hard disks as the storage basis for their LifeKeeper hierarchies.

This guide contains the following topics:

- [Documentation and References](#). Provides a list of LifeKeeper for Linux documentation.
- [Requirements](#). A description of the hardware and software necessary to properly setup, install, and operate the VMDK Recovery Kit. Refer to the [LifeKeeper for Linux Installation Guide](#) for specific instructions on how to install or remove LifeKeeper for Linux software.
- [Overview](#). A description of the VMDK Recovery Kit's features and functionality.
- [Configuring LifeKeeper for Linux VMDK Recovery Kit](#). A description of the procedures required to properly configure the VMDK Recovery Kit.
- [LifeKeeper Configuration Tasks](#). A description of the tasks for creating and managing your VMDK resource hierarchies using the LifeKeeper GUI.
- [Troubleshooting](#). A list of LifeKeeper for Linux error messages including a description for each.

## 6.21.1. VMDK Documentation and References

---

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- [LifeKeeper for Linux Release Notes](#)
- [LifeKeeper for Linux Technical Documentation](#)
- [LifeKeeper for Linux Installation Guide](#)
- [Optional Application Recovery Kit Documentation](#)

This documentation along with documentation associated with the optional LifeKeeper Application Recovery Kits is available at [docs.us.sios.com](https://docs.us.sios.com).

## 6.21.2. VMDK Hardware and Software Requirements

---

Your LifeKeeper configuration must meet the following requirements prior to the installation of the LifeKeeper for Linux VMDK Recovery Kit. See the [LifeKeeper for Linux Installation Guide](#) for specific instructions regarding the configuration of your LifeKeeper hardware and software.

### Hardware Requirements

- **Servers** – LifeKeeper for Linux supported VMware guests are configured in accordance with the requirements described in the [LifeKeeper for Linux Release Notes](#) and [LifeKeeper for Linux Installation Guide](#).
- **IP Network Interface Cards** – Each server requires at least one Ethernet TCP/IP supported network interface card. LifeKeeper clusters require two communications paths. Two separate LAN-based communication paths using dual independent sub-nets are recommended for heartbeats and at least one of these should be configured as a private network. Using a combination of TCP and TTY heartbeats is also supported.

### Software Requirements

- **TCP/IP Software** – Each server in your LifeKeeper configuration requires TCP/IP software.
- **LifeKeeper Software** – It is imperative that you install the same version of the LifeKeeper for Linux software and apply the same versions of the LifeKeeper for Linux software patches to each server in your cluster.
- **LifeKeeper for Linux VMDK Recovery Kit** – The VMDK Recovery Kit is included in the LifeKeeper installation image. It will be installed when selected on the Recovery Kit Selection screen of the setup script.
- **Linux Software** – Additional software required to run the VMDK Recovery Kit is included in the LifeKeeper installation image. Run the LifeKeeper setup script to install the VMDK Recovery Kit.

See the [LifeKeeper for Linux Installation Guide](#) for specific instructions on the installation and removal of the LifeKeeper for Linux software.

## 6.21.3. VMDK Recovery Kit Overview

---

The primary focus of the LifeKeeper for Linux VMDK Recovery Kit is to offer LifeKeeper users an alternative storage method for shared storage and data replication. The VMDK Recovery Kit enables the creation of LifeKeeper resource hierarchies on LifeKeeper protected servers. A virtual hard disk provided by the VMware hypervisor is connected to this resource hierarchy and the file system created on that disk is mounted. When a failure is detected on a node in the cluster where the virtual hard disk is connected, the VMDK Recovery Kit initiates a failover to the predetermined backup node and connects the same virtual hard disk to the backup node.

Once the file system configured on the virtual disk is mounted on a LifeKeeper server, it can be fully utilized as additional storage for LifeKeeper hierarchies. Resource hierarchies for the VMDK Recovery Kit are created using the existing File System Recovery Kit available with the LifeKeeper Core product (**steeleye-ik** package).

### VMDK Recovery Kit Restrictions

- This version of the VMDK Recovery Kit does not include support for a local recovery when access to the virtual hard disk fails. When a failure is detected the default action is to initiate a transfer of the hierarchy to a backup server. Depending on the makeup of the resource hierarchy, this action can result in hung processes. To avoid hung processes, the default action can be changed to halt the server and force a failover to a backup server. To change the default switchover behavior, alter the VMDK\_ERROR setting in /etc/default/LifeKeeper. See [Configuring the LifeKeeper for Linux VMDK Recovery Kit](#) for more information on VMDK\_ERROR.
- All guests participating in the cluster must have the same SCSI controller configuration. The VMDK Recovery Kit reconnects the virtual hard disk to the SCSI controller on the virtual hard disk that was connected when the resource was created.
- Snapshots cannot be created or restored on a LifeKeeper protected VMDK. When the virtual hard disk is switched between nodes, consistency of the snapshots cannot be guaranteed.

## 6.21.4. Configuring the VMDK Recovery Kit

---

This section describes the VMDK Recovery Kit configuration details. It also contains information you should consider before you start to configure and administer the VMDK Recovery Kit. Refer to the [LifeKeeper for Linux Technical Documentation](#) for instructions on configuring LifeKeeper Core resource hierarchies.

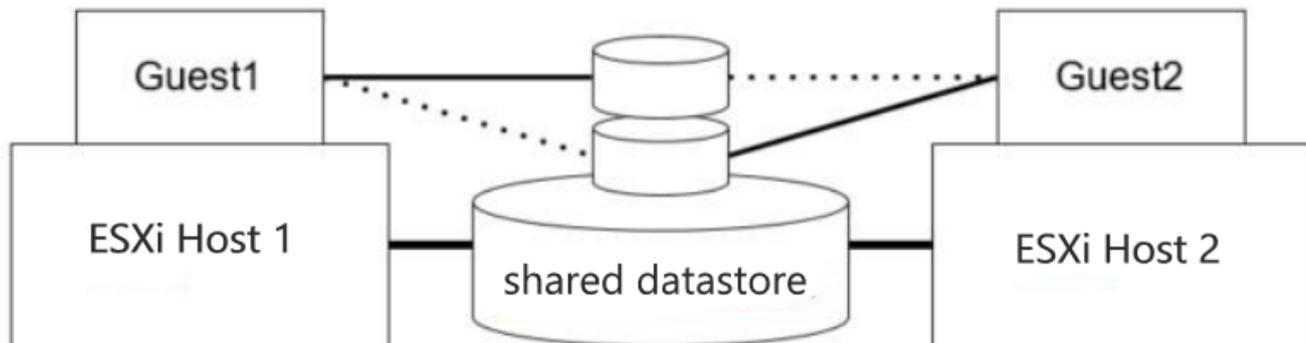
## 6.21.4.1. VMDK Configuration Considerations

---

1. Install the VMDK Recovery Kit on the servers in the cluster where you want to share your virtual hard disks. Create a virtual hard disk in VMDK format.
2. Virtual hard disks must be created on a datastore that is shared with the guests that make up the cluster.
3. Since exclusive control of the virtual hard disk depends on the hypervisor, the sharing setting of the connected SCSI controller must be set to “**None**”.
4. Because this kit operates the virtual hard disks using APIs provided by VMware, it is necessary to be able to access all VMware ESXi hosts running the guests that are participating in the cluster or managed vCenter Server via https.
5. The built-in file system recovery kit used to build the VMDK hierarchy detects and removes processes that are not under LifeKeeper protection using file systems mounted in a failover condition. **It is highly recommended that only processes that are under LifeKeeper protection be configured to use a file system under VMDK protection.**
6. The VMDK\_ERROR tunable controls the actions the VMDK Recovery Kit takes when access to the virtual hard disk fails. The tunable has two values, halt and event with halt being the default.
  - If the value is set to **halt** and an access failure is detected, the VMDK Recovery Kit will immediately halt the system and force a failover to the backup server.
  - If the value is set to **event**, the VMDK Recovery Kit notifies LifeKeeper with an abnormal status of the disk when access is lost. LifeKeeper will then attempt to initiate a switchover to a backup node. It is possible that the switchover process may hang, due to unkillable processes running on the shared VMDK.

## 6.21.4.2. VMDK Configuration Examples

### Configuration 1: Active/Standby Configuration Example

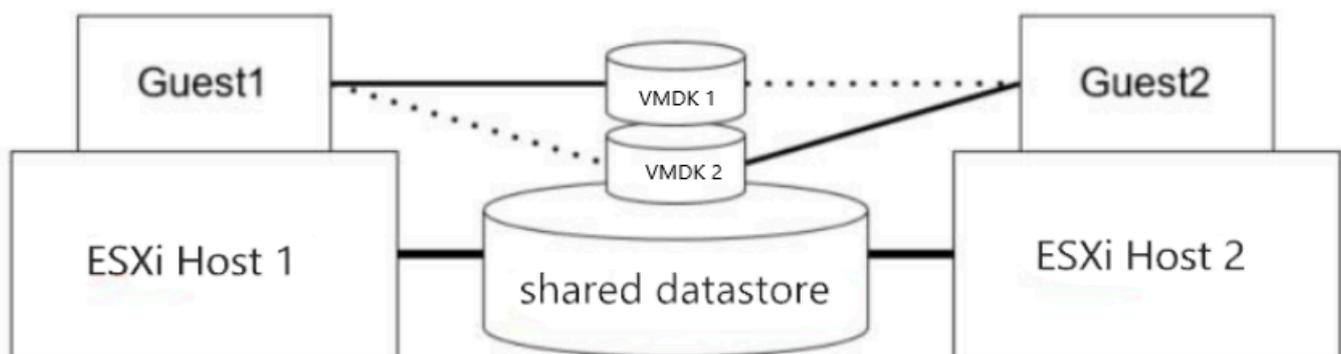


In this configuration, Guest 1 on ESXi Host 1 is considered active because it is able to access the virtual hard disk with the VMDK Recovery Kit software. If Guest 1 fails, Guest 2 gains access to the VMDK and the file system.

#### Configuration Notes

- The VMDK Recovery Kit must be installed on both servers.
- Create the file system on the shared VMDK virtual hard disk.
- Guest 2 should not access files and directories on the shared virtual disk while Guest 1 is active.

### Configuration 2: Active/Active Configuration Example



An active/active configuration consists of two or more systems actively running the VMDK Recovery Kit software and connecting different virtual hard disks.

#### Configuration Notes:

- The VMDK Recovery Kit must be installed on both servers.
- Initially, Guest 1 imports a file system and Guest 2 imports a different file system. In a switchover situation, one system can import both file systems.

## 6.21.5. LifeKeeper VMDK Recovery Kit Configuration Tasks

---

You can perform all LifeKeeper for Linux VMDK Recovery Kit administrative tasks via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer, and monitor VMDK resources.

The following tasks are available for configuring the LifeKeeper for Linux VMDK Recovery Kit:

- [Register an ESXi Host](#) – Register the information on the ESXi host that manages the virtual hard disk.
- [Change the VM Options](#) – Set the options required for the VMDK Recovery Kit.
- [Create a Resource Hierarchy](#) – Creates a VMDK resource hierarchy.
- [Delete a Resource Hierarchy](#) – Deletes a VMDK resource hierarchy.
- [Extend a Resource Hierarchy](#) – Extends a VMDK resource hierarchy from the primary server to the backup server.
- [Unextend a Resource Hierarchy](#) – Unextends (removes) a VMDK resource hierarchy from a single server in the LifeKeeper cluster.
- Create Dependency – Creates a child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- Delete Dependency – Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- In Service – Activates a resource hierarchy.
- Out of Service – Deactivates a resource hierarchy.
- View / Edit Properties – View or edit the properties of a resource hierarchy.

**Note:** Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You can also perform most of these tasks by:

1. From the toolbar, right-click on a global resource in the left pane of the status display.
2. Right-click on a resource instance in the right pane of the status display.

\*Using the right-click method allows you to avoid entering information that is required when using the Edit menu.

## 6.21.5.1. Register ESXi Host

---

Before creating a VMDK resource, register the ESXi host information. Follow the steps below:

1. Execute the following command on the console screen.

```
# /opt/LifeKeeper/lkadm/subsys/scsi/vmdk/bin/esxi_register -a <ESXi host name>
```

- When you execute the command, you will be asked for the username and password. Enter the username and password used to log in to the

ESXi host. Failing to login will lead to an error and you will not be able to register.

2. Register each ESXi host in the cluster the same way.
3. Once all of the ESXi hosts have been registered, make sure that they have been registered correctly with the following command:

```
# /opt/LifeKeeper/lkadm/subsys/scsi/vmdk/bin/esxi_register -l
```

- A list of registered hosts will be provided.
4. Register the ESXi host in the same way on all nodes in the cluster.

For details of the `esxi_register` command, see [VMDK Maintenance](#).

## 6.21.5.2. Changing the Virtual Machine Option Settings

The VMDK Recovery Kit requires the following options to be set:

Key	Value
disk.enableUUID	TRUE

Open the edit dialog and add the above parameters for each VM:

- Edit settings
- VM Options
- Advanced
- Configuration Parameters

**Edit settings - RHEL7.6 (ESXi 6.7 virtual machine)**

Virtual Hardware | **VM Options**

- ▶ General Options: VM Name:
- ▶ VMware Remote Console Options:  Lock the guest operating system when the last remote user disconnects
- ▶ VMware Tools: Expand for VMware Tools settings
- ▶ Power management: Expand for power management settings
- ▶ Boot Options: Expand for boot options
- ▶ Advanced: Expand for advanced settings
- ▶ Fiber Channel NPIV: Expand for fiber channel NPIV

Save Cancel

### Edit settings - RHEL7.6 (ESXi 6.7 virtual machine)

Expand for best options

- Advanced
  - Settings
    - Disable acceleration
    - Enable logging
  - Debugging and statistics
    - Run normally
  - Swap file location
    - Default  
Use the settings of the cluster or host containing the virtual machine.
    - Virtual machine directory  
Store the swap file in the same directory as the virtual machine.
    - Datastore specified by host  
Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.
  - Configuration Parameters
    - Edit Configuration...
  - Latency Sensitivity
    - Normal

Save Cancel

### Configuration Parameters

+ Add parameter -x Delete parameter

Q Search

Key	Value
tools.guest.desktop.autolock	FALSE
nvram	RHEL7.6.nvram
pciBridge0.present	TRUE
svga.present	TRUE
pciBridge4.present	TRUE
pciBridge4.virtualDev	pcieRootPort
pciBridge4.functions	8
pciBridge5.present	TRUE

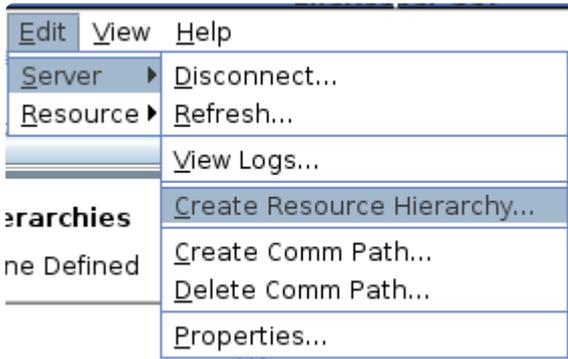
64 items

OK Cancel

## 6.21.5.3. Creating a VMDK Resource Hierarchy

Perform the following on your primary server and initiate the **Create Resource Wizard**.

1. Select **Edit > Server > Create Resource Hierarchy**



2. The **Select Recovery Kit** dialog appears. Select the **File System** option from the dropdown list. (**Note:** A VMDK Resource Hierarchy is a File System Hierarchy created on a shared virtual disk.)

Please Select Recovery Kit

Click **Next** to continue.

\* If you click the **Cancel** button at any time during the process of creating your hierarchy, LifeKeeper will cancel the entire creation process.

3. The **Switchback Type** dialog appears. The switchback type determines how the VMDK resource will be switched back to the primary server when it becomes in-service (active) on the backup server following a failover. Switchback types are either *intelligent* or *automatic*.

**Intelligent switchback** requires administrative intervention to switch the resource back to the primary server while **automatic switchback** occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.

Switchback Type

Click **Next** to continue.

4. The **Server** dialog appears. Select the name of the server where the VMDK resource will be created (typically this is your primary server). All servers in your cluster are included in the dropdown list.

Server

Click **Next** to continue.

- Select the **Mount Point** path to be protected by the VMDK (File System) Resource Hierarchy. All "local" (i.e. file systems using shared storage) and mount points of the virtual hard disk that can be managed with the VMDK Recovery Kit are listed. Select the desired mount point from the dropdown list.

Mount Point

Click **Next** to continue.

- The **Root Tag** dialog is automatically populated with a unique name for the resource instance on the target server (i.e. the server selected above). You may accept the default or enter a unique tag consisting of letters, numbers and the following special characters: -, \_, ., or /.

Root Tag

Click **Create Instance**.

- An information box appears indicating the start of the hierarchy creation.

```

Creating gen/filesys resource /test on a110.yo-satoh.localdomain
/opt/LifeKeeper/lkadm/subsys/gen/filesys/bin/creFShier a110.yo-satoh.localdomain /test
/test intelligent
devicehier: Using /opt/LifeKeeper/lkadm/subsys/scsi/vmdkp/bin/devicehier to construct the
hierarchy
BEGIN create of "vmdk30822"
END successful create of "vmdk30822"
BEGIN create of "vmdkp30815"
END successful create of "vmdkp30815"
Creating dependency "vmdkp30815"-"vmdk30822" on machine
"a110.yo-satoh.localdomain".

```

Click **Next** to continue.

- An information box appears after the successful creation of your VMDK resource hierarchy. You must **Extend** the hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

You have successfully created a resource hierarchy on one server. You may select continue in order to extend this resource hierarchy to another server, or you may cancel at this point.

If you cancel, the resource hierarchy provides no protection for your applications until it is extended to at least one other server in the cluster.

Click **Continue** to extend the resource.

Click **Cancel** if you want to extend your resource at a later time.

#### Verifying Integrity of Extended Hierarchy...

##### Hierarchy Verification Finished

WARNING: Your hierarchy exists on only one server. Your  
WARNING: application has no protection until you extend it  
WARNING: to at least one other server.

9. Click **Done** to exit the Create Resource Hierarchy menu.

## 6.21.5.4. Deleting a VMDK Resource Hierarchy

To delete a VMDK resource from all servers in your LifeKeeper configuration, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the dropdown menu, select **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your VMDK resource hierarchy.

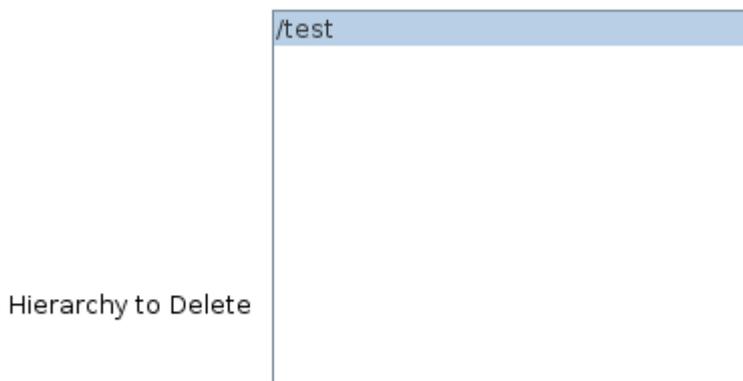
**Note:** If you selected the **Delete Resource Hierarchy** by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

Target Server

Click **Next** to continue.

3. Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete, and highlight it.

**Note:** If you selected the **Delete Resource Hierarchy** by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete.

You have specified the following resource hierarchy for deletion.  
 Target Server: all10.yo-satoh.localdomain  
 Target Tags:  
 /test

Click **Delete** to continue.

5. An information box appears confirming that the VMDK resource instance was deleted successfully.

```
Deleting resource hierarchy /test
```

```
Removing root resource hierarchy starting at "/test":  
BEGIN delete of "vmdkp30815"  
END successful delete of "vmdkp30815"  
BEGIN delete of "vmdk30822"  
END successful delete of "vmdk30822"  
Hierarchies successfully removed
```

6. Click **Done** to exit the Delete Resource Hierarchy menu selection.

## 6.21.5.5. Extending Your VMDK Hierarchy

After you have created a hierarchy, you should extend that hierarchy to another server in the cluster. There are three possible ways to extend your resource instance:

1. When you successfully create your VMDK resource hierarchy you will have an opportunity to select **Continue** which will allow you to proceed with extending your resource hierarchy to your backup server.
2. Right-click on an unextended hierarchy in either the left or right pane on the LifeKeeper GUI.
3. Select the **Extend Resource Hierarchy** task from the LifeKeeper GUI by selecting **Edit, Resource, Extend Resource Hierarchy** from the dropdown menu. This sequence of selections will launch the Extend Resource Hierarchy wizard. The **Accept Defaults** button that is available for the **Extend Resource Hierarchy** option is intended for the user who is familiar with the LifeKeeper Extend Resource Hierarchy defaults and wants to quickly extend a LifeKeeper resource hierarchy without being prompted for input or confirmation. Users who prefer to extend a LifeKeeper resource hierarchy using the interactive, step-by-step interface of the GUI dialogs should use the **Next** button.

a. The first dialog box to appear will ask you to select the **Template Server** where your VMDK resource hierarchy is currently in service. Remember that the **Template Server** you select now and the **Tag to Extend** that you select in the next dialog box represent an in-service (activated) resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you have selected. The dropdown list in this dialog provides the names of all the servers in your cluster.

**Note:** If you are entering the Extend Resource Hierarchy task by continuing from the creation of a VMDK resource hierarchy, this dialog box will not appear because the wizard has already identified the template server in the create stage. This is also the case when you right-click on either the VMDK resource icon in the left pane or right-click on the VMDK (File System) resource box in the right pane of the GUI window and choose **Extend Resource Hierarchy**.

Template Server

**!** **CAUTION:** If you click the **Cancel** button at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extend hierarchy process. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Click **Next** to continue.

b. Select the **Tag to Extend**. This is the name of the VMDK instance you want to extend from the template server to the target server. All of the resources that you have created on

the template server will be listed in the dropdown.

**Note:** If you are entering the Extend Resource Hierarchy task immediately following the creation of a VMDK hierarchy, this dialog box will not appear because the wizard has already identified the tag name of your resource in the create stage. This is also the case when you right-click on either the VMDK (File System) resource icon in the left pane or on the VMDK (File System) resource box in the right pane of the GUI window and choose **Extend Resource Hierarchy**.

Tag to Extend

Click **Next** to continue.

c. Select the **Target Server** where you will extend your VMDK resource hierarchy.

Target Server

Click **Next** to continue.

d. The **Switchback Type** dialog appears. The switchback type determines how the VMDK resource will be switched back to the primary server when it becomes in service (active) on the backup server following a failover. Switchback types are either *intelligent* or *automatic*. **Intelligent switchback** requires administrative intervention to switch the resource back to the primary server while **automatic switchback** occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.

Switchback Type

Click **Next** to continue.

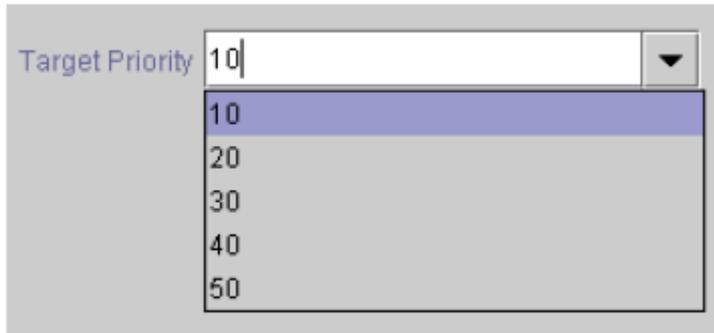
e. Select or enter a **Template Priority**. This is the priority for the VMDK hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.

**Note:** This selection will appear only for the initial extending of the hierarchy.

Click **Next** to continue.

f. Select or enter the **Target Priority**. This is the priority for the new extended VMDK hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns

the number “1” to the server on which the hierarchy is created by default. The priorities do not need to be consecutive, but no two servers can have the same priority for a given resource.



The image shows a user interface element labeled 'Target Priority'. It consists of a text input field containing the number '10' and a dropdown arrow on the right. Below the input field is a list box containing the numbers 10, 20, 30, 40, and 50. The number 10 is currently selected and highlighted in blue.

Click **Next** to continue.

g. An information box appears confirming that LifeKeeper has successfully checked your environment and that all requirements for extending this resource have been met. If there are requirements that have not been met, LifeKeeper will disable the **Next** button and enable the **Back** button.

```
Executing the pre-extend script...
Building independent resource list
Checking existence of extend and canextend scripts
Checking extendability for /test
Pre Extend checks were successful
```

Click **Back** to make changes to your resource extension.

Click **Cancel** to extend your resource another time.

Click **Next** to launch the Extend Resource Hierarchy configuration task.

Click **Finish** to confirm the successful extension of your VMDK resource instance.

4. Click **Done** to exit the Extend Resources Hierarchy menu selection.

**Note:** Be sure to test the functionality of the new instance on both servers.

## 6.21.5.6. Unextending Your VMDK Hierarchy

1. From the LifeKeeper GUI menu, select **Edit > Resource > Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the VMDK resource. It cannot be the server where the resource is currently in-service (active).

**Note:** If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear

Target Server

Click **Next** to continue.

3. Select the **Hierarchy to Unextend**.

**Note:** If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.

Hierarchy to Unextend

Click **Next** to continue.

4. An information box appears confirming the target server and the VMDK resource hierarchy you have chosen to unextend.

```
You have specified the following resource hierarchy for unextend.  
Target Server = all1.yo-satoh.localdomain  
Target Tag = /test
```

Click **Unextend**.

5. Another information box appears confirming that the VMDK resource was unextended successfully.
6. Click **Done** to exit the Unextend Resource Hierarchy menu selection.

## 6.21.5.7. Testing Your VMDK Resource Hierarchy

---

You can test your VMDK resource hierarchy by initiating a manual switchover that will simulate a failover of the resource instance from the primary server to the backup server.

### Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit > Resource > In Service**. For example, an in-service request executed on a backup server causes the VMDK resource hierarchy to be placed in-service on the backup server and taken out-of-service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the Out of Service request, the resource hierarchy is taken out-of-service without bringing it in-service on the other server.

## 6.21.5.8. VMDK Maintenance

---

### Recovery from a Failover Caused by a Node Failure

If a failover occurs due to an ESXi host failure, the virtual hard disk cannot be disconnected. Therefore, the virtual hard disk remains connected to the stopped guest. For this reason, multiple guests try to access the virtual hard disk when returning, but the operation is restricted by the hypervisor and therefore the guest cannot be started. In this case, use the vSphere client to manually disconnect the virtual hard disk from the guest.

### Changing ESXi Login Information

To change the username and password of the ESXi host perform the following steps:

1. Stop LifeKeeper or all VMDK resources.
2. Execute the following command from the command line:  

```
# /opt/LifeKeeper/lkadm/subsys/scsi/vmdk/bin/esxi_register -u <ESXi host name>
```

Enter a new username and password interactively. Once you log in successfully the information will be updated. If the VMDK resource is running or you cannot log in an error occurs and the information is not updated.
3. Start the stopped LifeKeeper or VMDK resources.
4. Repeat the same steps for all nodes and update the login information.

### Deleting ESXi Host Information

To delete registered ESXi host information perform the following steps:

1. Stop LifeKeeper or all the VMDK resources.
2. Execute the following command from the command line:  

```
# /opt/LifeKeeper/lkadm/subsys/scsi/vmdk/bin/esxi_register -d <ESXi host name>
```
3. Start the stopped LifeKeeper or VMDK resources.
4. Repeat the same steps for all nodes and update the login information.

### Esxi\_register Details

Registering a host	<code>esxi_register -a &lt;ESXi host name&gt;</code>
--------------------	------------------------------------------------------

Deleting a host	<code>esxi_register -d &lt;ESXi host name&gt;</code>
Updating login information	<code>esxi_register -u &lt;ESXi host name&gt;</code>
A list of registered hosts	<code>esxi_register -l</code>

## 6.21.6. VMDK Troubleshooting

Symptom	Possible Cause
<b>Mount point is not included in the selection when creating resources</b>	<p>Possible causes are as follows:</p> <ul style="list-style-type: none"> <li>• PowerShell/PowerCLI is not installed</li> <li>• An ESXi host is not registered</li> <li>• disk.enableUUID parameter is not set</li> <li>• The virtual hard disk is on a datastore that is not shared</li> <li>• SCSI controller sharing is configured as “virtual” or “physical”</li> </ul> <p>Error details are recorded in <code>/var/log/lifekeeper.log</code>. Check the log and review the settings.</p>
<b>It takes longer to bring the VMDK resource in service.</b>	<p><b>Cause:</b> The processing performed during bringing the resource in service takes more time in proportion to the number of virtual machines running on the ESXi host.</p> <p><b>Action:</b> A fundamental fix is under consideration for a future release. As an immediate workaround for this issue, please consider the following:</p> <ul style="list-style-type: none"> <li>• Reduce the number of VMDK resources. Since the process is performed for each VMDK resource, the time required for the process increases in proportion to the number of VMDK resources. <b>If multiple partitions or file systems are required for a single resource hierarchy</b>, create multiple partitions or file systems on a single VMDK resource rather than using multiple VMDK resources.</li> <li>• If ESXi hosts that are not related to the cluster node are registered with the VMDK resource, unregister them. <a href="#">How to manage ESXi host information</a></li> <li>• Reduce the number of running virtual machines.</li> </ul>

## 6.21.6.1. VMDK Error Messages

This section provides a list of messages that you may encounter while creating and extending a **LifeKeeper VMDK** resource hierarchy or removing and restoring a resource. Where appropriate, it provides an additional explanation of the cause of an error and necessary action to resolve the error condition.

Messages from other LifeKeeper components are also possible. In these cases, refer to the appropriate LifeKeeper component documentation.

The messages are grouped into the following topics:

- Common error message
- Creating a hierarchy
- Extending a hierarchy
- Deleting, restoring, recovering a hierarchy

### Common Error Message

Error Number	Error Message
000002	Usage error
000010	Error getting resource information
000011	Both Tag and ID name not specified
000019	Resource not found on local server
000022	END failed hierarchy <tag name> in service on server <server name>
000026	END failed ACTION for <tag name> on server <server name> due to <signal> signal

### Creating a Hierarchy

Error Number	Error Message
000012	Usage error
000013	Error getting resource information
000014	Resource with either matching tag <tag name> or ID exists
000015	ins_create failed on server <server name>
000018	Error creating resource <tag name> on server <server name>
000021	Removing resource instance <tag name> from server <server name> due to an error during creation
000023	Error bringing resource <tag name> in service on server <server name>

000024	Failed resource creation of resource <tag name> on server <server name>
000027	Removing file system dependency from <parent tag> to <child tag> on server <server name> due to an error during creation
000028	Removing file system hierarchy <filesys tag> created by <parent tag> on server <server name> due to an error during creation
000029	Switchback type mismatch between parent <parent tag> and child <child tag> on server <server name>  <b>Action:</b> Switchback type mismatch can cause unexpected behavior. You can eliminate this by manually changing the switchback type using the ins_setas command.
000030	create: tag name not specified  or  extend: tag name not specified

### Extending a Hierarchy

Error Number	Error Message
000003	Template resource <tag name> on server <server name> does not exist
000004	Template resource <tag name> cannot be extended to server <server name> because it already exists there
000005	Cannot access canextend script on server <server name>
000006	Cannot access extend script <path to extend> on server <server name>
000007	Cannot access depstoextend script <path to depstoextend> on server <server name>
000008	Cannot extend resource <tag name> to server <server name>
000009	Either <templatesys> or <templatetag> argument missing
000014	Resource with either matching tag <tag name> or ID exists
000015	ins_create failed on server <server name>
000018	Error creating resource <tag name> on server <server name>
000025	END failed resource extension of <tag name> on server <server name> due to a "<signal>" signal - backing out changes made to server
000030	create: tag name not specified  or  extend: tag name not specified

## Restore

Error Number	Error Message
000023	Error bringing resource <tag name> in service on server <server name>

## Resource Monitoring

Error Number	Error Message
000001	Calling sendevent for resource <tag name> on server <server name>

## VMDK Recovery Kit

Error Number	Error Message
137000	PowerShell is not installed.
137001	PowerCLI is not installed.
137002	A valid network interface was not found.
137003	Attaching VMDK \$vmdkfile.
137004	VMDK already attached.
137005	Failed to attach VMDK.
137006	Attach success.
137008	Detaching VMDK \$vmdkfile.
137009	VMDK Already detached.
137010	Failed to detach VMDK.
137011	Detach success.
137012	Restarting VMDK status checker daemon.
137016	Stopping VMDK status checker daemon.
137020	Failed to execute VMDK status checker daemon.
137026	Flushing \$dev.
137027	Skipping flush for \$dev.
137030	Disk not specified.
137031	Cannot get disk uuid for \$Disk. Please check your ESXi settings.
137032	PowerCLI failed. %s

137034	Cannot bring VMDK resource \"%s\" in service on server \"%s\".
137035	Error detected conflict in expected tag name \"%s\" on target machine \"%s\".
137036	Template resource \"%s\" on server \"%s\" does not exist.
137037	This system is not a VMware guest.
137038	Unable to find shared device on \"%s\" for \"%s\".
137039	Unable to find SCSI controller on \"%s\" for \"%s\".
137050	Failed to connect to ESXi server \$addr.
137051	There is no ESXi server connected.
137055	Cannot determine ESXi VM ID because multiple network interfaces were found with the MAC address \$MAC_ADDR.
137056	Failed to \$action VMDK \$VMDK_FILENAME. Retrying \$action in \$wait_sec milliseconds.
137057	Usable SCSI controller not found.
137058	Cannot find VMDK with ID \$UUID.
137059	This guest has snapshots present.
137060	The VMDK with ID \$UUID cannot be attached to this guest.
137061	The virtual storage controller has an incompatible sharing mode configured.
137062	VMDK_TIMER too short. Using default value.
137063	Calling sendevent for resource \"\$Tag\" on server \"\$me\""
137064	skip quickcheck for \"\$Tag\" on server \"\$me\", sendevent pending.
137065	sendevent issued for tag \"\$Tag\" has not finished, halt server \"\$me\""
137066	skip quickcheck for \"\$Tag\" on server \"\$me\", sendevent pending.
137068	The VMDK detection failed. Retry count exceeded.
137070	Connect failed.
137071	Get-LocalVM failed.
137072	The VMDK is remote detached. This server has lost ownership.
137073	The VMDK quickCheck daemon has been stopped.
137074	VMDK_RETRY too small. Using default value.
137075	Cannot find virtual SCSI controller \$CONTROLLER.
137076	The virtual storage controller has an incompatible sharing mode configured.
137077	Cannot find VM with MAC address \$MAC_ADDR.
137078	Cannot find VM with MAC address \$MAC_ADDR.

137100	Re-reading partition table on %s.
137101	Partition information not defined for %s on %s. Retry.
137102	Partition information not defined for %s on %s.
137103	Resource %s is OSF, Skip flushing buffers.
137104	Flushing buffers on %s.
137105	Device not specified.
137106	Cannot get device uuid for \$Device. Please check your ESXi settings.
137107	%s is not shareable with any machine.
137109	Creating dependency \"%s\"-\"%s\" on machine \"%s\".
137110	Dependency \"%s\"-\"%s\" on machine \"%s\" exists.
137111	Failed to create dependency \"%s\"-\"%s\" on machine \"%s\".
137112	Cannot bring VMDKP resource \"%s\" in service on server \"%s\".
137113	detected conflict in expected tag name \"%s\" on target machine \"%s\".

## 7. Parameters List

### Core Parameters List

The table below lists and explains names and meanings of the Core parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value	Wh Ap
FS_KERNEL_RETRIES	The maximum number of times that the forceumount script, after exhausting all attempts to kill processes accessing the file system via SIGTERM and SIGKILL signals, will continue to attempt to unmount the file system. These unmount attempts occur every three seconds until the maximum number of attempts is exceeded.	Integers	60	As rec (takes imme
FS_UMOUNT_RETRIES	The maximum number of times that the forceumount script will attempt to use a SIGTERM signal to kill all processes currently accessing the file system which is being unmounted. If this number of attempts is exceeded, the forceumount script will attempt at most three times to kill the processes using a SIGKILL signal.	Integers	1	As rec (takes imme
REMOTETIMEOUT	Number of seconds between when a process sends a request through the "lcdsendremote" function to another machine before it expects a response. If no response is received in this time interval, the function will try an alternate path if available.	Integers	900	LifeKe startu (takes when restar LifeK
CONFIRMSODEF	The default action to take during machine failover processing when failover confirmation is configured. The default action is only taken when no manual response is received from the administrator within the timeout period (see CONFIRMSOTO).	0: proceed with failover 1: block the failover	0	As rec (takes imme
CONFIRMSOTO	The time in seconds to wait for administrator action when failover confirmation is configured. When the timeout period expires the default action for CONFIRMSODEF is taken. Otherwise, the administrator action is taken.	Integers	600	As rec (takes imme

FAILFASTTIMER	Number of seconds between verifying that a reserved device is still reserved by the local system. If the device is not reserved then the system will halt and reboot.	Integers	5	LifeKeeper startup (takes when restart LifeKeeper)
SCSIERROR	Determines the action to take when a SCSI device cannot be opened, accessed, or another SCSI error occurs (e.g., timeout).	event: LifeKeeper's core should be informed that a device needs to be switched over to a backup system  halt: The system should immediately be halted and rebooted to avoid data corruption	event	LifeKeeper startup (takes when restart LifeKeeper)
LKCHECKINTERVAL	Application health monitoring wait time (in seconds) between checks. Set to zero to disable health monitoring.	Integers (0, 1 and over)	120	LifeKeeper startup (takes when restart LifeKeeper)
FILESYSFULLWARN	The file system full threshold at which time warning messages will start appearing in the LifeKeeper log. Setting to 0 will disable monitoring.	Integers	90	As required (takes immediate)

FILESYSFULLERROR	<p>The file system full threshold at which time error messages will start appearing in the LifeKeeper log. Additionally the LKROOT/events/filesys/diskfull/notify script will be called when this threshold is reached.</p> <p>Setting to 0 will disable monitoring.</p>	Integers	95	As rec (takes imme
LK_TRAP_MGR	<p>One or more network managers (separated by commas) to receive SNMP traps. No traps are sent if this variable is not set.</p>	String	(not set)	As rec (takes imme
LK_NOTIFY_ALIAS	<p>Email address or address list used to receive notification messages when certain events occur in a LifeKeeper cluster. A null value indicates no notification will occur. The expected format is:</p> <p>LK_NOTIFY_ALIAS=</p> <ul style="list-style-type: none"> <li>- no notification is sent</li> </ul> <p>LK_NOTIFY_ALIAS=user1@domain1</p> <ul style="list-style-type: none"> <li>- mail sent to user1 at domain1</li> </ul> <p>LK_NOTIFY_ALIAS=user1@domain1,user2@domain1</p> <ul style="list-style-type: none"> <li>- mail sent to user1 and user2 at domain1</li> </ul>	String	(not set)	As rec (takes imme
LKSYSLOGTAG	<p>Tag for syslog.</p>	String	LifeKeeper	LifeKe startu (takes when restar LifeKe
LKSYSLOGSELECTOR	<p>Level for syslog.</p>	user,	local6	LifeKe

		daemon, local0, local1, ...or local7		startup (takes when restar LifeK
LCMHBEATTIME	The interval, in seconds, used to send heartbeats signal Failing to receive the LCM signal, which includes heartbeat signal, from another server within the interval time is determined as heartbeat stop.	Integers	5	LifeKe startu (takes when restar LifeK
LCMNUMHBEATS	Number of consecutive missed heartbeats to mark a communication path down. In the real implemented system, it is not the number, but LCMHBEATTIME x LCMNUMHBEATS seconds missing communication is determined as communication path disconnection.	Integers	3	LifeKe startu (takes when restar LifeK
LC_MESSAGES	Changes the language environment.	String	C	LifeKe startu (takes when restar LifeK
GUI_WEB_PORT	Specifies the port to use for LifeKeeper Management Web servers (lkGUI).	Integers	81	Resta stele lighttp

API_SSL_PORT	Specifies the port used for the LifeKeeper API.	Integers	778	Restar steele lighttp
LOGMGR_LOGLEVEL	Specifies the log level of Generic Applications.	LK_INFO or LK_ERROR	LK_ERROR	LifeKe startu restar lk_log proces (takes when restar LifeK

## 7.1. EC2 Parameters List

The table below lists the EC2 parameters. These values are set by adding them to the `/etc/default/LifeKeeper` configuration file. Because none of the components of the Recovery Kit for EC2 are memory resident, changes to these particular values become effective immediately after they are changed in `/etc/default/LifeKeeper` without requiring a LifeKeeper restart.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	Notes
EC2_RESTORE_TIMEOUT	Timeout for the resource restore, in seconds.	Integers	300	As required <b>(takes effect immediately)</b>	
EC2_REMOVE_TIMEOUT	Timeout for the resource remove, in seconds.	Integers	300	As required <b>(takes effect immediately)</b>	
EC2_RECOVER_TIMEOUT	Timeout for the local recovery, in seconds.	Integers	300	As required <b>(takes effect immediately)</b>	
EC2_QUICKCHECK_TIMEOUT	Timeout for the quickCheck, in seconds.	Integers	100	As required <b>(takes effect immediately)</b>	
EC2_AWS_REGION	Specifies the region where EC2 resources reside.	String	(not set)	As required <b>(takes effect immediately)</b>	
IP_NOLINKCHECK	Disables the link check for the protected network interface.	0: enabled  1: disabled	0	As required <b>(takes effect immediately)</b>	This value only applies when protecting an Elastic IP.
IP_WAIT_LINKDOWN	Number of seconds to wait in between taking the protected network interface down and back up. A delay between these two actions is necessary in some environments.	Integers	5	As required <b>(takes effect immediately)</b>	This value only applies when protecting an Elastic IP.
IP_MAX_LINKCHK	The maximum number of seconds to wait for the link to come back up after it has been repaired. In some	Integers	5	As required <b>(takes effect immediately)</b>	This value only applies when

	environments, it may be necessary to increase this value.				protecting an Elastic IP.
AWSCLI_CONNECT_TIMEOUT	The connection timeout value in seconds used when running “AWS” commands. It is specified via <code>—cli-connect-timeout</code> argument.	Integers	10	As required <b>(takes effect immediately)</b>	This is the same parameter as used in <a href="#">Route53</a> .
AWSCLI_READ_TIMEOUT	The read timeout value in seconds used when running “AWS” commands. It is specified via <code>—cli-read-timeout</code> argument.	Integers	5	As required <b>(takes effect immediately)</b>	This is the same parameter as used in <a href="#">Route53</a> .
HTTP_PROXY HTTPS_PROXY NO_PROXY	Set these parameters when using a HTTP proxy for accessing the service endpoint. The value set here is passed to AWS CLI. Please refer to AWS Documentation for details. <a href="https://docs.aws.amazon.com/cli/latest/userguide/cli-http-proxy.html">https://docs.aws.amazon.com/cli/latest/userguide/cli-http-proxy.html</a>	String	(not set)	As required <b>(takes effect immediately)</b>	This is the same parameter as used in <a href="#">Route53</a> and <a href="#">Quorum</a> .

## 7.2. IP Parameters List

The table below lists and explains names and meanings of the IP parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	Notes
IP_PINGTRIES	Number of ping retries that will be performed during an IP health check.	Integers	3	As required <b>(takes effect immediately)</b>	
IP_PINGTIME	Time in seconds that LifeKeeper waits for one packet ping reply during IP health check.	Integers	1	As required <b>(takes effect immediately)</b>	When using a manually configured <i>Ping List</i> rather than the broadcast ping mechanism, any value greater than 3 for this tunable is ineffective, because the Linux TCP/IP implementation always returns a "Destination Host Unreachable" error after 3 seconds with no reply, regardless of the timeout value specified in the ping command.
IP_QUICKCHECK_TIMEOUT	Increases the time before quickCheck is marked as failed when added to <code>/etc/default/LifeKeeper</code> .  The default time allowed for the IP quickcheck to complete is 12 seconds.	Integers	12	As required <b>(takes effect immediately)</b>	Setting this value does not require stopping or restarting any LifeKeeper processes.

<p>NOIPUNIQUE</p>	<p>Disables the IP uniqueness checking done when an IP resource is brought in-service. By default LifeKeeper will ensure the IP address is not in use somewhere else on the network.</p>	<p>0: enabled 1: disabled</p>	<p>0</p>	<p>As required <b>(takes effect immediately)</b></p>	
<p>NOBCASTPING</p>	<p>Disables the broadcast ping mechanism for checking the health of IP resources.</p>	<p>0: enabled 1: disabled</p>	<p>0</p>	<p>As required <b>(takes effect immediately)</b></p>	
<p>IP_NOLINKCHECK</p>	<p>Disables the link status check portion of the IP health check.</p>	<p>0:enabled 1: disabled</p>	<p>0</p>	<p>As required <b>(takes effect immediately)</b></p>	<p>This setting may need to be disabled on virtual environments, specifically after this message in the logs "Link check failed for virtual IP".</p>
<p>IP_MAX_LINKCHK</p>	<p>The maximum number of seconds to wait for the link to come back up after it has been repaired. In some environments, it may be necessary to</p>	<p>Integers</p>	<p>5</p>	<p>As required <b>(takes effect immediately)</b></p>	

	increase this value.				
IP_WAIT_LINKDOWN	Number of seconds to wait in between taking the protected network interface down and back up. A delay between these two actions is necessary in some environments.	Integers	5	As required <b>(takes effect immediately)</b>	

## 7.3. MQ Parameters List

The table below lists and explains names and meanings of the MQ parameters. These values are tuned by editing the /etc/default/LifeKeeper configuration file.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	Notes
MQS_QUICKCHECK_TIMEOUT_SC	Timeout in seconds for the server connect check.	Integers	10	As required <b>(takes effect immediately)</b>	
MQS_QUICKCHECK_TIMEOUT_CC	Timeout in seconds for the client connect check.	Integers	10	As required <b>(takes effect immediately)</b>	
MQS_QUICKCHECK_TIMEOUT_PUTGET	Timeout in seconds for the PUT/GET check.	Integers	10	As required <b>(takes effect immediately)</b>	
MQS_QUICKCHECK_TIMEOUT_PS	Timeout in seconds for checking whether publish/subscribe is in use.	Integers	5	As required <b>(takes effect immediately)</b>	
MQS_QUICKCHECK_TIMEOUT_CLUSTER	Timeout in seconds for checking whether the queue manager is part of an WebSphere MQ cluster or not.	Integers	5	As required <b>(takes effect immediately)</b>	
MQS_QUICKCHECK_TIMEOUT	Timeout in seconds for the quickCheck	Integers	40	As required <b>(takes effect immediately)</b>	If the value is less than 10 seconds, it

	script.				will be set to the default.
MQS_QMGR_START_TIMEOUT	Timeout in seconds for the queue manager start command to complete.	Integers	60	As required <b>(takes effect immediately)</b>	
MQS_CMDMS_START_TIMEOUT	Timeout in seconds for the command server start command to complete.	Integers	30	As required <b>(takes effect immediately)</b>	
MQS_LISTENER_START_TIMEOUT	Timeout in seconds for the listener start command to complete.	Integers	30	As required <b>(takes effect immediately)</b>	
MQS_LISTENER_LIST_TIMEOUT	Timeout in seconds for the listener list command to complete.	Integers	10	As required <b>(takes effect immediately)</b>	
MQS_CHECK_TIMEOUT_ACTION	The action in case a server connect check or client connect check times out.	ignore: a message about the timeout is logged, but no recovery is initiated sendevent: local recovery is initiated in case a	ignore	As required <b>(takes effect immediately)</b>	

		server connect check timed out			
MQS_LISTENER_CHECK_DELAY	Time in seconds between the start of the listener and the check for the successful listener start. The default of 2 seconds should be sufficient to detect port in use conditions.	Integers	2	As required <b>(takes effect immediately)</b>	If the value is less than 2 seconds, it will be set to the default.
NO_AUTO_STORAGE_DEPS	Determines if the shared storage checks and file system resource creation step are performed for the queue manager and log storage directories during MQ resource hierarchy creation. A value of 0 indicates these tasks will be perform. A value of 1	0 or 1	0	As required <b>(takes effect immediately)</b>	

	will bypass these tasks.				
MQS_DSPMQVER_TIMEOUT	Timeout in seconds for the dspmqver command (needed to find out the version of WebSphere MQ), must be at least 2 seconds.	Integers	5	As required <b>(takes effect immediately)</b>	
MQS_SKIP_CRT_MISSING_Q	Determines if missing test queue is automatically created. A value of 0 indicates missing test queues will automatically be created. A value of 1 indicates this process will be skipped.	0 or 1	0	As required <b>(takes effect immediately)</b>	
MQS_ALT_USER_NAME	The alternate user name to use for all WebSphere MQ commands. By default the user mqm is used. If set the alternate user must have its primary group set to	Character string	mqm if not set or the user does not have membership in the "mqm" group.	As required <b>(takes effect immediately)</b>	Should only be set if a WebSphere MQ addon package requires a user other than "mqm"

	the group mqm or must have secondary membership in that group.				
--	----------------------------------------------------------------------------------	--	--	--	--

## 7.4. NFS Parameters List

---

The table below lists and explains names and meanings of the NFS parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	Notes
RESTARTMOUNTD	Enables the stop and restart of <code>rpc.mount</code> on all NFS restores.	true: enabled  false: disabled	true	As required ( <b>takes effect immediately</b> )	

## 7.5. Recovery Kit for Oracle Cloud Infrastructure Parameters List

The table below lists and explains names and meanings of the Recovery Kit for Oracle Cloud Infrastructure (RK for OCI) parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

The parameters of RK for OCI will take effect immediately when the value in `/etc/default/LifeKeeper` is changed, without restarting LifeKeeper.

Parameter Name	Meaning	Setting Value	Default Value
OCIVIP_RESTORE_TIMEOUT	Specifies the restore timeout value of the resource in seconds.	Integers	60
OCIVIP_REMOVE_TIMEOUT	Specifies the timeout value in seconds to stop the resource.	Integers	60
OCIVIP_RECOVER_TIMEOUT	Specifies the timeout value for local recovery in seconds.	Integers	60
OCIVIP_QUICKCHECK_TIMEOUT	Specifies the quickCheck timeout value in seconds.	Integers	60
OCI_CLI_PROFILE	Specifies the profile name used by the OCI CLI.  See the OCI documentation for details. <a href="https://docs.oracle.com/en-us/iaas/Content/API/SDKDocs/cliconfigure.htm">https://docs.oracle.com/en-us/iaas/Content/API/SDKDocs/cliconfigure.htm</a>	String	None

## 7.6. Oracle Parameters List

The table below lists and explains names and meanings of the Oracle parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	Notes
ORACLE_ORATABLOC	<code>/etc/oratab</code> is used by default. Specifies the directory where the alternative <code>oratab</code> file is located. <code>/etc/oratab</code> is referenced, but if the file does not exist, <code>oratab</code> in the directory specified with this parameter is read.	String	<code>/var/opt/oracle</code>	As required <b>(takes effect immediately)</b>	
LK_ORA_NICE	Determines whether a recovery attempt will occur on a database connection failure caused when the maximum number of allowed connections has been reached. A recovery attempt when the maximum number has been reached can cause a failover to the standby node.	0: execute the recovery attempt 1: prevent the recovery attempt	0	As required <b>(takes effect immediately)</b>	
ORACLE_RESTORE_TIMEOUT	Specifies the timeout value in seconds when starting Oracle resources.	Integers	300	As required <b>(takes effect immediately)</b>	
ORACLE_REMOVE_TIMEOUT	Specifies the timeout value in seconds when stopping Oracle resources.	Integers	300	As required <b>(takes effect immediately)</b>	

## 7.7. PostgreSQL Parameters List

The table below lists and explains names and meanings of the PostgreSQL parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value
LKPGSQL_KILLPID_TIME	Time in seconds to wait after a process id is killed before rechecking for this process.	Integers	3
LKPGSQL_CONN_RETRIES	Replaces LKPGSQLMAXCOUNT – number of times to try a client connection after an action (start or stop).	Integers	12
LKPGSQL_ACTION_RETRIES	Number of times to attempt start or stop action before failing the action command.	Integers	4
LKPGSQL_STATUS_TIME	Timeout in seconds for the status command.	Integers	17 + (3 * LKPGSQL_KILLPID)
LKPGSQL_QCKHANG_MAX	Number of quickCheck script hangs allowed before a failover/sendevent is triggered for the database instance.	Integers	2
LKPGSQL_CUSTOM_DAEMON	Allows a user to specify additional aliases for the postgres daemons (postgres.bin,postmaster,postgres,edb-postgres).	String	(not set)
LKPGSQL_IDIRS	Replaces LKPGSQL_IPORTS – contains datadir entries for instances that will be shutdown using the	String	(not set)

	immediate option only.		
LKPGSQL_SDIRS	Contains datadir entries for instances that will be shutdown using the smart option.	String	(not set)
LKPGSQL_DISCONNECT_CLIENT	Controls whether active clients will be disconnected in the event of a postmaster crash. When the value is set to 1, client processes will be sent a SIGTERM signal to force them to disconnect from the database. This action will only be taken if the postmaster process is not running during local recovery.	0: enabled 1: disabled	1
LKPGSQL_DISCONNECT_CLIENT_BYTAG	Similar to LKPGSQL_DISCONNECT_CLIENT, this setting limits the action to the comma separated list of tags specified by this tunable.	String	(not set)
LKPGSQL_RESUME_PROC	Determines if a process found in the stopped state (state = ~T) will be resumed when detected or ignored.	0: ignore 1: resume	1
LKPGSQLDEBUG	Turns on debug for PostgreSQL database kit as well as for the postgres database. Valid entry range: 0 – 5.  This parameter will be passed on to the postmaster database using the option <code>-d &lt;LKPGSQLDEBUG&gt;</code> .	Integers (0 – 5)	0

## 7.8. Quorum Parameters List

The table below lists the Quorum parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value
QUORUM_MODE	Specifies the quorum mode.	majority tcp_remote storage none or off	majority
QUORUM_HOSTS	Specifies a host and port name combination to be used for determining quorum. The format for entries is "host:port". When specifying more than one host:port combination, the entries must be comma separated (do not include a space).  (Example) QUORUM_HOSTS=myhost:80,router1:443,router2:22	String	(not set)
WITNESS_MODE	Specifies the witness mode.	remote_verify storage none or off	remote_v
QUORUM_TIMEOUT_SECS	The time allowed for tcp/ip witness connections to complete. Connections that don't complete within this time are treated as failed/unavailable. This only applies when the QUORUM_MODE is tcp_remote.	Integers	20
QUORUM_LOSS_ACTION	Specifies the action when quorum is lost.	fastkill fastboot osu	fastkill
QWK_STORAGE_TYPE	Specifies the type of shared storage.  Must be specified when QUORUM_MODE is storage.	block file aws_s3	(not set)

QWK_STORAGE_HBEATTIME	Specifies the interval in seconds between reading and writing the QWK objects.	An integer between 5 and 10	6
QWK_STORAGE_NUMHBEATS	Specifies the number of consecutive heartbeat checks that when reached indicates the target node has failed. A missed heartbeat occurs when the QWK object has not been updated since the last check.	An integer of 3 or more	4
QWK_STORAGE_OBJECT_<Host name>  <b>Note:</b> If the host name contains a "-" or ".", replace them with an underscore "_" (e.g. lksios-1 → lksios_1).  <b>Note:</b> The host name used by LifeKeeper can be checked via the lcduname command	Specifies the path to the QWK objects. You must specify paths for all nodes in the cluster. <b>[When QWK_STORAGE_TYPE is block]</b> Specify the device file path. <b>Note:</b> Use WWID (/dev/disk/by-id/) to specify a permanent path. (Example) QWK_STORAGE_OBJECT_nodeA=/dev/disk/by-id/xxxxx QWK_STORAGE_OBJECT_nodeB=/dev/disk/by-id/yyyyy  <b>[When QWK_STORAGE_TYPE is file]</b> Specify the regular file path. (Example) QWK_STORAGE_OBJECT_nodeA=/quorum/nodeA QWK_STORAGE_OBJECT_nodeB=/quorum/nodeB	String (Maximum length is 256 characters)	(not set)

	<p><b>[When QWK_STORAGE_TYPE is aws_s3]</b>                  Specify the S3uri for the Amazon S3 object. Use an S3 object from a different region than the one where LifeKeeper is running. It is also recommended that 2 different S3 objects be used and that they reside in different regions. When specifying 2 regions, make sure to separate them with commas(do not include spaces).                  (Example 1)                  QWK_STORAGE_OBJECT_nodeA=s3://bucket1/nodeA,s3://bucket2/nodeA                  QWK_STORAGE_OBJECT_nodeB=s3://bucket1/nodeB,s3://bucket2/nodeB                  (Example 2)                  QWK_STORAGE_OBJECT_nodeA=s3://bucket/quorum/nodeA                  QWK_STORAGE_OBJECT_nodeB=s3://bucket/quorum/nodeB</p> <p><b>Note:</b> NodeA and nodeB must be regular files and not directories.</p>		
<p>HTTP_PROXY                  HTTPS_PROXY                  NO_PROXY</p>	<p>Set this parameter when using HTTP proxy for accessing the service endpoint. The value set here will be passed to AWS CLI.                  See <a href="#">AWS Documentation</a> for details.</p>	<p>String</p>	<p>(not set)</p>
<p>QUORUM_DEBUG</p>	<p>Specifies the debug mode.</p>	<p>0: disabled                  1: enabled</p>	<p>0</p>

## 7.9. Route53 Parameters List

The table below lists and explains the tunable values that are available for modifying the behavior of the Route53 Recovery Kit. These values are set by adding the `/etc/default/LifeKeeper` configuration file. Because none of the components of the Route53 Recovery Kit are memory resident, changes to these particular values become effective immediately after they are changed in `/etc/default/LifeKeeper` without requiring a restart of LifeKeeper or the OS.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	
ROUTE53_TTL	The default setting value for TTL (Time To Live) of the A record created for the Route53 resource, in seconds.	Integers	10	After switchover	
ROUTE53_CHANGEID_INTERVAL	The interval of Route 53 API communications when checking the status, in seconds.	Integers	20	As required <b>(takes effect immediately)</b>	
ROUTE53_CHANGEID_TRY_COUNT	The number of trials of Route 53 API communications when checking the status.	Integers	5	As required <b>(takes effect immediately)</b>	
ROUTE53_QUICKCHECK_TIMEOUT	Timeout for the quickCheck in seconds.	Integers	25	As required <b>(takes effect immediately)</b>	To help pro out, you ca ROUTE53  You may v quickcheck 25 to some AWS CLI c would be t (ROUTE53
AWSCLI_CONNECT_TIMEOUT	The connection timeout value in seconds used when running "AWS" commands. It is specified via <code>—cli-connect-timeout</code> argument.	Integers	10	As required <b>(takes effect immediately)</b>	This is the <a href="#">EC2</a> .
AWSCLI_READ_TIMEOUT	The read timeout value in seconds used when running "AWS" commands. It is specified via <code>—cli-read-timeout</code> argument.	Integers	5	As required <b>(takes effect immediately)</b>	This is the <a href="#">EC2</a> .
HTTP_PROXY	Set these parameters when	String	(not	As required	This is the

<p>HTTPS_PROXY NO_PROXY</p>	<p>using HTTP proxy for accessing the service endpoint. The value set here is passed to AWS CLI. Please refer to AWS Documentation for details. <a href="https://docs.aws.amazon.com/cli/latest/userguide/cli-http-proxy.html">https://docs.aws.amazon.com/cli/latest/userguide/cli-http-proxy.html</a></p>		<p>set)</p>	<p><b>(takes effect immediately)</b></p>	<p>and <a href="#">Quoru</a></p>
---------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--	-------------	------------------------------------------	----------------------------------

## 7.10. SAP Parameters List

The table below lists and explains names and meanings of the SAP parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	Notes
SAP_CONFIG_REFRESH	Refresh time in seconds of the Configuration properties page.	Integers	LKCHECKINTERVAL/2	As required <b>(takes effect immediately)</b>	If the value is less than 5 seconds, it will be set to the default.
SAP_CREATE_NAS	Automatically includes a NAS resource for NAS mounted file systems.	0: disabled 1: enabled	1	As required <b>(takes effect immediately)</b>	
SAP_NFS_CHECK_DIRS	Comma-separated list of NFS mount points to check	String	empty	As required when using NFS shared filesystems	Do not use for Amazon EFS mount points
SAP_QUICKCHECK_TIMEOUT	Timeout in seconds for the quickCheck process.	Integers	60 seconds	As required <b>(takes effect immediately)</b>	
SAP_RESTORE_TIMEOUT	Timeout in seconds for the restore process.	Integers	516 seconds	As required <b>(takes effect immediately)</b>	
SAP_REMOVE_TIMEOUT	Timeout in seconds for the remove process.	Integers	804 seconds	As required <b>(takes effect immediately)</b>	
SAP_RECOVER_TIMEOUT	Timeout in	Integers	1320 seconds	As required	If the

	seconds for the recover process.			<b>(takes effect immediately)</b>	value is less than the default, it will be set to the default.
--	----------------------------------	--	--	-----------------------------------	----------------------------------------------------------------

# 7.11. DataKeeper Parameters List

The table below lists and explains names and meanings of the DataKeeper parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	
LKDR_CHUNK_SIZE	Sets the chunk size of bitmap in kilobits.	Integers	256	Creating a resource	This set t 0 to mirro be a not s or is size crea will f zero 256K too h muc bitm chur bein ever If the the b will b will l the r
LKDR_CONNECT_NBD_DURING_RESTORE	Specifies if the NBD connection should be established when the mirror is restored (brought in-service).	True or False	True	To improve switchover and failover performance on a cluster with multiple mirrors and multiple targets	The ("tru max integ that resto imm data avail in-se perfo setti conn esta first

					inter LKC 120 defa decr requ mirro also when repli targe perfo impr notic mult mult
LKDR_SPEED_LIMIT	Specifies the maximum bandwidth that a resync will ever take – this should be set high enough to allow resyncs to go at the maximum speed possible.	Integers	500000	Resyncing a resource	
LKDR_SPEED_LIMIT_MIN	Specifies how fast the resync should be allowed to go when there is other I/O going on at the same time. As a rule of thumb, this should be set to half or less of the drive’s maximum write throughput in order to avoid starving out normal I/O activity when a resync occurs.	Integers	20000	Resyncing a resource	This less for LKD
LKDR_ASYNC_LIMIT	Specifies the number of	Integers	4096	Creating a resource	The this

	outstanding target writes that can be in flight at a given time. Increase this value for higher performance with asynchronous mirrors.				1638 above mirro it is s limit 4096 will c othe 2 – 1 an a that Valu not r
LKDR_NO_FULL_SYNC	Suppresses a full resync of newly added targets.	0: not suppress 1: suppress	0	As required <b>(takes effect immediately)</b>	For c infor “ <a href="#">Avo</a> <a href="#">Resu</a>
LKDR_WAIT_TO_RESYNC	Specifies which parent resource(s) must be in-service before initializing mirror resynchronization.	False, hierarchy, <resource type>	filesystem	Before resuming replication to verify replicated data is consistent.	Fals beha earli wait hiera “<res ‘files pare that hiera serv
LKDR_WAIT_FOR_PREVIOUS_SOURCE_TIMEOUT	Specifies how long to wait for the previous source to join the cluster so that its bitmap can be merged. A full resync is required if replication is resumed to a target before the previous source’s bitmap is merged. This setting applies to all netraid devices;	0: do not wait -1: wait indefinitely for the previous source to rejoin the cluster > 0 number of seconds to wait	-1	As required <b>(takes effect immediately)</b>	<a href="#">How</a> <a href="#">Data</a>

	individual devices can NOT be configured with a different value.				
--	---------------------------------------------------------------------------	--	--	--	--

## 7.12. Standby Node Health Check Parameters List

It is necessary to enable/disable each functionality and configure the value for the [Standby Node Health Check](#). These settings can be customized in the `/etc/default/LifeKeeper` configuration file.

Parameter	Description	Value to Set	Default Value	When to Apply	Notes
<code>SNHC</code>	Enables the overall functionality of the <a href="#">Standby Node Health Check</a> . To enable, set this parameter to 1; it is disabled by default. <code>SNHC_XX</code> for individual functions must be enabled.	0: Disabled 1: Enabled	0	When <code>lkcheck</code> is started	After setting restart <code>lkcheck</code> . See <a href="#">Standby Node Health Check</a> for details.
<code>SNHC_CPUCHECK</code>	Enables CPU monitoring with <a href="#">Node Monitoring</a> . To enable, set this parameter to 1; it is disabled by default.	0: Disabled 1: Enabled	0	Anytime	
<code>SNHC_CPUCHECK_THRESHOLD</code>	Specifies the threshold for CPU utilization that is considered to be abnormal in CPU monitoring. If not specified, 99 is used.	Integer value between 10 and 99	99	Anytime	If not configured, users are prompted to configure with an ERROR message in the log.
<code>SNHC_CPUCHECK_TIME</code>	Specifies the number of consecutive times that CPU utilization must be over the threshold ( <code>SNHC_CPUCHECK_THRESHOLD</code> ) to be considered an error. If not specified, 1 is used.	Integer value between 1 and 100	1	Anytime	If not configured, users are prompted to configure with an ERROR message in the log.
<code>SNHC_MEMCHECK</code>	Enables memory monitoring with <a href="#">Node Monitoring</a> . To enable, set this parameter to 1; it is disabled by default.	0: Disabled 1: Enabled	0	Anytime	

<i>SNHC_MEMCHECK_THRESHOLD</i>	Specifies the threshold for memory utilization that is considered to be abnormal in memory monitoring. If not specified, 99 is used.	Integer value between 10 and 99	99	Anytime	If not configured, users are prompted to configure with an ERROR message in the log.
<i>SNHC_MEMCHECK_TIME</i>	Specifies the number of consecutive times that memory utilization must be over the threshold ( <i>SNHC_MEMCHECK_THRESHOLD</i> ) to be considered an error. If not specified, 1 is used.	Integer value between 1 and 100	1	Anytime	If not configured, users are prompted to configure with an ERROR message in the log.
<i>SNHC_IPCHECK</i>	Enables <a href="#">OSU resource monitoring</a> for IP resources. It is disabled by default.	0: Disabled 1: Enabled	0	Anytime	
<i>SNHC_IPCHECK_SLEEPTIME</i>	Specifies the wait time to check the link after starting the NIC, when the relevant NIC is down during OSU monitoring for IP resources. If the NIC was down before monitoring, the NIC will be left down after monitoring.	Wait time (sec)	1	Anytime	
<i>SNHC_DISKCHECK</i>	Enables <a href="#">OSU resource monitoring</a> for disk or DMMP resources . It is disabled by default.	0: Disabled 1: Enabled	0	Anytime	
<i>SNHC_DISKCHECK_&lt;Tag name&gt;</i>	Enables <a href="#">OSU resource monitoring</a> for only disk or DMMP resources of the specified tag name. It is disabled by default.	0: Disabled 1: Enabled	0	Anytime	
<i>SNHC_NASCHECK</i>	Enables <a href="#">OSU resource monitoring</a> for NAS resources.	0: Disabled 1: Enabled	0	Anytime	
<i>SNHC_NASCHECK_&lt;Tag name&gt;</i>	Enables OSU resource monitoring		0	Anytime	If the tag

	for only NAS resources of the specified tag name.	0: Disabled 1: Enabled			name includes hyphen (-), replace it with underscore to configure
--	---------------------------------------------------	---------------------------------	--	--	-------------------------------------------------------------------

## 7.13. SAP HANA Parameters List

The table below lists and explains names and meanings of SAP HANA parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value
HANA_HANDSHAKE_TAKEOVER_FAILBACK	Whether to perform an automatic failback of the SAP HANA resource in the case of a failed Takeover with Handshake attempt. When set to 'true', LifeKeeper will automatically bring the SAP HANA resource back in-service on the previous host in the case of a failed Takeover with Handshake attempt. When set to 'false' (default), the SAP HANA resource remains in the Out-of-Service Failed (OSF) state after a failed Takeover with Handshake attempt and requires manual intervention to bring in-service.	True or False	False
HANA_HSR_POLL_INTERVAL	Specifies how often (in seconds) LifeKeeper will poll the status of SAP HANA System Replication.	Positive integers	10
HANA_QUICKCHECK_TIMEOUT	Specifies the timeout value (in seconds) for the SAP HANA quickCheck script.	Positive integers, Minimum value 30	LKCHECKINTERVAL – 10
HANA_REGISTER_SECONDARY_DURING_RESTORE	Whether the SAP	True or	False

	<p>HANA Recovery Kit should register and restart the secondary database before completing the in-service operation for the primary database instance. When set to 'true', the SAP HANA resource is not marked in-service until the recovery kit has attempted to register and restart the secondary database. When set to 'false' (default), the SAP HANA resource is marked in-service as soon as the primary database instance is running, and a local recovery event is initiated immediately afterward to register and restart the secondary database.</p>	False	
HANA_RECOVER_TIMEOUT	<p>Specifies the timeout value (in seconds) for the SAP HANA recover script.</p>	Positive integers	1800
HANA_START_WAIT	<p>Specifies the amount of time (in seconds) to wait when starting the SAP HANA database instance.</p>	Positive integers	2700
HANA_STOP_WAIT	<p>Specifies the amount of time (in seconds) to wait when stopping the SAP HANA database instance.</p>	Positive integers	600
HANA_STOP_COUNT	<p>Specifies the number of attempts to issue the sapcontrol StopWait command</p>	Positive integers	3

	<p>when stopping the SAP HANA database instance.</p>		
<p>HANA_STOP_FINAL_WAIT</p>	<p>After HANA_STOP_COUNT failed attempts to stop the SAP HANA database, this specifies the amount of additional time (in seconds) to wait to allow the database to fully stop.</p>	<p>Positive integers</p>	<p>60</p>

## 7.14. SAP MaxDB Parameters List

The table below lists and explains names and meanings of the SAP MaxDB Recovery Kit parameters. These values are tuned by editing the `/etc/default/LifeKeeper` configuration file.

Parameter Name	Meaning	Setting Value	Default Value	When to Apply	Notes
MAXDB_START_TIMEOUT	Number of seconds to wait before aborting an in-service action that is hung.	Integers	300	As required ( <b>takes effect immediately</b> )	
MAXDB_PID_CLEANUP	Specifies whether child processes of a hung in-service action will be cleaned up (killed).	y: clean up n: no cleanup	n	As required ( <b>takes effect immediately</b> )	
MAXDB_STOP_COUNT	Number of times LifeKeeper will attempt to offline the database.	Integers	5	As required ( <b>takes effect immediately</b> )	
MAXDB_WAIT	Number of seconds to wait between database offline attempts.	Integers	5	As required ( <b>takes effect immediately</b> )	
MAXDB_DEBUG	Enables or disables debug logging.	0: disabled 1: enabled	0	As required ( <b>takes effect immediately</b> )	
MAXDB_OFFLINE_ENABLED	Specifies whether the database will be taken offline when the LifeKeeper MaxDB resource is taken out of service.	0: do not offline 1: do offline	1	Use with caution, if you need to temporarily stop LifeKeeper (e.g., to upgrade or perform maintenance) while not stopping the database.	Please ensure you re-enable database offlines after maintenance is complete.

## 8. Search for an Error Code

---

Please look at the [Combined Message Catalog](#)

## 8.1. Combined Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
000200	ERROR	pam_start() failed	
000201	ERROR	pam_authenticate failed (user %s, retval %d	
000202	ERROR	pam_end() failed?!?!	
000203	ERROR	Did not find expected group 'lkguest'	
000204	ERROR	Did not find expected group 'lkoper'	
000205	ERROR	Did not find expected group 'lkadmin'	
000208	ERROR	pam_setcred establish credentials failed (user %s, retval %d	<p><b>Cause:</b> Unable to establish valid login credentials for user {user}. The pam_setcred call returned: {retval}.</p> <p><b>Action:</b> Check /var/log/security and /var/log/messages for more information.</p>
000209	ERROR	pam_setcred delete credentials failed (user %s, retval %d	<p><b>Cause:</b> Unable to clear login credentials for user {user}. The pam_setcred call returned: {retval}.</p> <p><b>Action:</b> Check /var/log/security and /var/log/messages for more information.</p>
000902	ERROR	Error removing system name from loopback address line in /etc/hosts file. You must do this manually before starting the GUI server.	<p><b>Cause:</b> System name did not get removed from /etc/hosts file.</p> <p><b>Action:</b> Remove system name manually then restart the GUI server, then enter the following: run &lt;action name&gt;</p>
000918	ERROR	LifeKeeper GUI Server error during Startup	<p><b>Cause:</b> The GUI server terminated due</p>

Code	Severity	Message	Cause/Action
			<p>to an abnormal condition.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
001052	FATAL	Template resource "%s" on server "%s" does not exist	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p>
001053	ERROR	Cannot access canextend script "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to run pre-extend checks because it was unable to find the script CANEXTEND on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
001054	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p>
001055	ERROR	Cannot access extend script "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource hierarchy because it was unable to find the script EXTEND on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
001057	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p>
001059	ERROR	Resource with tag "%s" already exists	<p><b>Cause:</b> The name provided for a resource is already in use.</p> <p><b>Action:</b> Either choose a different name for the resource, or use the existing resource.</p>

Code	Severity	Message	Cause/Action
001060	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server "%s" for App "%s" and Type "%s"	<p><b>Cause:</b> The name or id provided for a resource is already in use.</p> <p><b>Action:</b> Either choose a different name or id for the resource or use the existing resource.</p>
001061	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected failure occurred while creating a resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
001081	WARN	IP address \"%\$ip\" is neither v4 nor v6	<p><b>Cause:</b> The IP address provided is neither an IPv4 nor an IPv6 address.</p> <p><b>Action:</b> Please check the name or address provided and try again. If a name was provided, verify that name resolution returns a valid IP address.</p>
004024	ERROR		<p><b>Cause:</b> LCD failed to fetch resource information for resource id {id} during resource recovery.</p> <p><b>Action:</b> Verify the input resource id and retry the recovery operation.</p>
004028	ERROR	%s occurred to resource \"%s\"	<p><b>Cause:</b> Local recovery failed for resource {resource}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004055	ERROR	attempt to remote-remove resource \"%s\" that can't be found	<p><b>Cause:</b> Remotely removing a resource from service failed while attempting to fi</p>

Code	Severity	Message	Cause/Action
			<p>nd the resource by tag name {tag}.</p> <p><b>Action:</b> Check input tag name and retry the recovery operation.</p>
004056	ERROR	attempt to remote-remove resource \"%s\" that is not a shared resource	<p><b>Cause:</b> Remotely removing a resource from service failed given that the tag name {tag} is not a shared resource.</p> <p><b>Action:</b> Check input tag name and retry the recovery operation.</p>
004060	ERROR	attempt to transfer-restore resource \"%s\" that can't be found	<p><b>Cause:</b> Remote transfer of an in service resource failed given tag name {tag}.</p> <p><b>Action:</b> Check input tag name and retry the recovery operation.</p>
004061	ERROR	attempt to transfer-restore resource \"%s\" that is not a shared resource with machine \"%s\"	<p><b>Cause:</b> LifeKeeper failed to find a shared resource by {tag} name during remote transfer of a resource in service from a remote {machine}.</p> <p><b>Action:</b> Check input tag name and retry the recovery operation.</p>
004089	ERROR	ERROR: Parallel recovery initialization failed.\n	<p><b>Cause:</b> Parallel recovery failed to initialize the list of resources in the hierarchy.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004091	ERROR	ERROR: fork failed. continuing to next resource\n	<p><b>Cause:</b> Parallel recovery failed to fork a new process attempting to restore a single resource.</p>

Code	Severity	Message	Cause/Action
004093	ERROR	ERROR: reserve failed. continuing to next resource\n	<p><b>Cause:</b> Parallel recovery failed to reserve a single resource from the collective hierarchy.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004096	ERROR	ERROR: clone %d is hung, attempting to kill it\n	<p><b>Cause:</b> A single sub process of a resource recovery is hung during parallel recovery of a whole resource hierarchy.</p> <p><b>Action:</b> A kill of the hanging sub process will be executed automatically.</p>
004097	ERROR	ERROR: Could not kill clone %d\n	<p><b>Cause:</b> Failed to kill the hung sub process.</p>
004116	ERROR	%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed attempting to create the intermediate folder. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the intermediate folder is not created.</p>
004117	ERROR	open(%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while attempting to open a temporary file. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file is not successfully opened.</p>
004118	ERROR	write(%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while atte</p>

Code	Severity	Message	Cause/Action
			<p>tempting to write to a temporary file. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file writing is failed.</p>
004119	ERROR	fsync(%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while attempting to fsync a temporary file. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the "fsync" is failed.</p>
004120	ERROR	close(%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while attempting to close a temporary file. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file closing is failed.</p>
004121	ERROR	rename(%s, %s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while attempting to rename a temporary file {file} to original file {file}. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file renaming is failed.</p>
004122	ERROR	open(%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while attempting to open an intermediate directory {directory}. This is a system error."</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check log for the error information in detail and determine why the directory open is failed.</p>
004123	ERROR	fsync(%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while attempting to fsync an intermediate directory {directory}. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the directory "fsync" is failed.</p>
004124	ERROR	close(%s	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed while attempting to close an intermediate directory {directory}. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the directory close is failed.</p>
004125	ERROR	wrote only %d bytes of requested %d\n	<p><b>Cause:</b> Writing an on-disk version of an in-memory data object failed as the final size {size} bytes of written data is less than the requested number {number} of bytes.</p> <p><b>Action:</b> Check log for the related error information in detail and determine why the data writing is failed.</p>
004126	ERROR	open(%s	<p><b>Cause:</b> Attempting to open a data file failed during the reading of an on-disk version of the data object into the buffer. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file o</p>

Code	Severity	Message	Cause/Action
			pen is failed.
004127	ERROR	open(%s	<p><b>Cause:</b> Attempting to open a temporary data file failed during the reading of an on-disk version of the data object into the buffer. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file open is failed.</p>
004128	ERROR	read(%s	<p><b>Cause:</b> Reading a data file failed during the loading of an on-disk version of the data object into the buffer. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file reading is failed.</p>
004129	ERROR	read buffer overflow (MAX=%d)\n	<p><b>Cause:</b> The read buffer limit {max} was reached while attempting to read an on-disk version of the data object into the buffer.</p> <p><b>Action:</b> Check the LifeKeeper configuration and restart the LifeKeeper.</p>
004130	ERROR	close(%s	<p><b>Cause:</b> Failed to close a data file during reading an on-disk version of the data object into the buffer. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file close is failed.</p>
004131	ERROR	rename(%s, %s	<p><b>Cause:</b> Failed to rename a temporary d</p>

Code	Severity	Message	Cause/Action
			<p>ata file during reading an on-disk version of the data object into the buffer. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the file renaming is failed.</p>
004132	ERROR	Can't open %s : %s	<p><b>Cause:</b> Failed to open a directory {directory} with error {error} during reading an on-disk version of the application and resource type information into the buffer. This is a system error.</p> <p><b>Action:</b> Check log for the error information in detail and determine why the open directory is failed.</p>
004133	ERROR	path argument may not be NULL	<p><b>Cause:</b> The command "lcdrpc" failed during a file copy because the input source path is missing.</p> <p><b>Action:</b> Check the input source path and retry "lcdrpc".</p>
004134	ERROR	destination path argument may not be NULL	<p><b>Cause:</b> The "lcdrpc" command failed during a file copy because the input destination path is missing.</p> <p><b>Action:</b> Check the input destination path and retry "lcdrpc".</p>
004135	ERROR	destination path can't be zero length string	<p><b>Cause:</b> Input destination path was empty during file copy when using "lcdrpc".</p> <p><b>Action:</b> Check the input destination path and retry "lcdrpc".</p>

Code	Severity	Message	Cause/Action
004136	ERROR	open(%s	<p><b>Cause:</b> Failed to open source file path during file copy using "lcdrp". This is a system error.</p> <p><b>Action:</b> Check the existence/availability of the input source path and retry "lcdrp". Also check the related log for error information in detail.</p>
004137	ERROR	fstat(%s	<p><b>Cause:</b> Failed to fetch file attributes using "fstat" during file copy by "lcdrp". This is a system error.</p> <p><b>Action:</b> Check the log for error information in detail.</p>
004138	ERROR	file \"%s\" is not an ordinary file (mode=0%o	<p><b>Cause:</b> Detected source file as a non ordinary file during file copy using "lcdrp".</p> <p><b>Action:</b> Check the input source file path and retry "lcdrp".</p>
004151	FATAL	lcdMalloc failure	<p><b>Cause:</b> Failed to allocate memory with requested size in shared memory.</p> <p><b>Action:</b> A fatal error will be produced.</p>
004152	ERROR	having \"%s\" depend on \"%s\" would produce a loop	<p><b>Cause:</b> Adding the requested dependency would produce a loop of dependent relationship.</p> <p><b>Action:</b> Correct the requested dependency and retry the dependency creation.</p>
004164	ERROR	Priority mismatch between resources %s and %s. Dependency creation failed.	<p><b>Cause:</b> The priorities for {resource1} and {resource2} do not match.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Resource priorities must match. Change one or both priorities to the same value and retry creating the dependency.</p>
004176	ERROR	%s	<p><b>Cause:</b> The command "doabort" failed to create the {directory} for writing the core file. This is a system error.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004182	ERROR	received signal %d\n	<p><b>Cause:</b> Received signal {signum}.</p>
004186	ERROR	%s: ::receive(%d) protocol error on incoming_mailbox %s	<p><b>Cause:</b> In function {function}, attempting to receive message within timeout {timeout} seconds failed due to a non-idle status of incoming mailbox {mailbox}.</p> <p><b>Action:</b> Check the status of the connections inside the cluster and retry the process.</p>
004190	ERROR	%s: ::receive(%d) did not receive message within %d seconds on incoming_mailbox %s	<p><b>Cause:</b> In function {function}, attempting to receive message within timeout {timeout} seconds failed with incoming mailbox {mailbox}.</p> <p><b>Action:</b> Check the status of the connections inside the cluster and retry the process.</p>
004204	ERROR	attempt to send illegal message	<p><b>Cause:</b> Sending message failed due to a illegal message.</p>
004205	ERROR	destination system \"%s\" is unknown	<p><b>Cause:</b> Sending message failed due to</p>

Code	Severity	Message	Cause/Action
			<p>an unknown destination system name {system}.</p> <p><b>Action:</b> Check the configuration and status of the system and check the logs for related errors. Retry the same process after the system is full initialized.</p>
004206	ERROR	destination mailbox \"%s\" at system \"%s\" is unknown	<p><b>Cause:</b> Sending message failed due to an unknown mailbox {mailbox} on destination system name {system}. This error may be caused by sending a message before the LCD is fully initialized.</p> <p><b>Action:</b> Check the configuration and status of the system and check the logs for related errors. Retry the same process after the system is full initialized.</p>
004208	ERROR	destination system \"%s\" is alive but the \"%s\" mailbox process is not listening.	<p><b>Cause:</b> Sending message failed. The network connection to destination system {system} is alive but the contact to the destination mailbox is lost.</p> <p><b>Action:</b> Check the configuration and status of the system and check the logs for related errors. Retry the same process after the system is full initialized.</p>
004209	ERROR	destination system \"%s\" is dead.	<p><b>Cause:</b> Sending message failed due to losing connection with the destination system {system}.</p> <p><b>Action:</b> Check the configuration and status of the system and check the logs for related errors. Retry the same process after the system is full initialized.</p>
004211	ERROR	can't send to destination \"%s\" error=%d	<p><b>Cause:</b> Sending message to destination</p>

Code	Severity	Message	Cause/Action
			<p>system {system} failed due to internal error {error}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004217	ERROR	destination system \"%s\" is out of service.	<p><b>Cause:</b> Sending message failed due to losing connection with the destination system {system}.</p> <p><b>Action:</b> Check the configuration and status of the system and check the logs for related errors. Retry the same process after the system is full initialized.</p>
004221	ERROR	destination system \"%s\" went out of service.	<p><b>Cause:</b> Sending message failed due to losing connection with the destination system {system}.</p> <p><b>Action:</b> Check the configuration and status of the system and check the logs for related errors. Retry the same process after the system is full initialized.</p>
004228	ERROR	Can't get host name from getaddrinfo()	<p><b>Cause:</b> Creating network object failed due to a failure when getting host name using "getaddrinfo()".</p> <p><b>Action:</b> Check the configuration and status of the system. Do the same process again.</p>
004234	ERROR	IP address pair %s already in use	<p><b>Cause:</b> Creating network object failed due to the IP address pair {pair} is already used for a TCP communication path.</p> <p><b>Action:</b> Check the input IP address pair and do the network creation again.</p>

Code	Severity	Message	Cause/Action
004258	WARN	Communication to %s by %s FAILED	<p><b>Cause:</b> Communication to system {system} by communication path {path} failed.</p> <p><b>Action:</b> Check system configuration and network connection.</p>
004261	WARN	COMMUNICATIONS failover from system \"%s\" will be started.	<p><b>Cause:</b> A failover from system {system} will be started due to all the communications path are down.</p> <p><b>Action:</b> Check system configuration and network connection status. Confirm the system status when failover is done.</p>
004292	ERROR	resource \"%s\" %s	<p><b>Cause:</b> A resource could not be brought in service because its current state is unknown.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004293	ERROR	resource \"%s\" %s	<p><b>Cause:</b> A resource could not be brought into service because its current state disallows it.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004294	ERROR	resource \"%s\" requires a license (for Kit %s/%s) but none is installed	<p><b>Cause:</b> The resource's related recovery kit requires a license.</p> <p><b>Action:</b> Install a license for the recovery kit on the server where the resource was to be brought into service.</p>

Code	Severity	Message	Cause/Action
004297	ERROR	secondary remote resource \"%s\" on machine \"%s\" is already in-service, so resource \"%s\" on machine \"%s\" can't be brought in-service.	<p><b>Cause:</b> A resource {resource} could not be brought into service on machine {machine} because its secondary remote resource {resource} is already in service on machine {machine}.</p> <p><b>Action:</b> Manually change the remote resource out of service and do in service on local resource again.</p>
004298	ERROR	remote resource \"%s\" on machine \"%s\" is still in-service, restore of resource \"%s\" will not be attempted!\n	
004300	ERROR	restore of resource \"%s\" has failed	<p><b>Cause:</b> A resource could not be brought into service.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004311	ERROR	can't perform \"remove\" action on resources in state \"%s\"	<p><b>Cause:</b> A resource could not be put out of service due to the current state being {state}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004313	ERROR	remove of resource \"%s\" has failed	<p><b>Cause:</b> A resource {resource} could not be put out of service.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004318	ERROR	%s,priv_globact(%d,%s): script %s FAILED returning %d	<p><b>Cause:</b> A global action script failed with the specified error code.</p> <p><b>Action:</b> Check the logs for related error</p>

Code	Severity	Message	Cause/Action
			s and try to resolve the reported problem.
004332	ERROR	action \"%s\" has failed on resource \"%s\"	<p><b>Cause:</b> A resource action failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004351	ERROR	a \"%s\" equivalency must have one remote resource	<p><b>Cause:</b> Creating of a {eqvtype} equivalency failed due to the two input tag names exist on the same system.</p> <p><b>Action:</b> Correct the inputs resource tag names and do the same process again.</p>
004356	WARN	Use unsupported option %s for remove. This option may be removed in future upgrades.	
004376	FATAL	wait period of %u seconds for LCM to become available has been exceeded (lock file \"%s\" not removed)	<p><b>Cause:</b> The LCM daemon did not become available within a reasonable time and the LCD cannot operate without the LCM.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004386	ERROR	initlcdMalloc;shmget	<p><b>Cause:</b> A shared memory segment could not be initialized.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Review the product documentation and ensure that the server meets the minimum requirements and that the operating system is configured properly.</p>
004439	WARN	intermachine recovery skipped for %s. Fai	

Code	Severity	Message	Cause/Action
		led to obtain resource_state_change lock.\n	
004444	WARN	License key (for Kit %s/%s) has EXPIRED	<p><b>Cause:</b> Your license has expired.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
004445	WARN	License key (for Kit %s/%s) will expire at midnight in %ld days	<p><b>Cause:</b> Your license is about to expire.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
004466	ERROR	system \"%s\" not defined on machine \"%s\".	<p><b>Cause:</b> The specified system name is not known.</p> <p><b>Action:</b> Verify the system name, and try the operation again.</p>
004467	ERROR	system \"%s\" unknown on machine \"%s\"	<p><b>Cause:</b> The specified system name is not recognized.</p> <p><b>Action:</b> Verify the system name, and try the operation again.</p>
004494	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004495	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
004496	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004497	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004498	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004499	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004500	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004501	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
004502	ERROR	COMMAND OUTPUT: %s	<b>Cause:</b> An action or event script produced unexpected output. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
004503	ERROR	COMMAND OUTPUT: %s	<b>Cause:</b> An action or event script produced unexpected output. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
004504	ERROR	COMMAND OUTPUT: %s	<b>Cause:</b> An action or event script produced unexpected output. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
004505	ERROR	COMMAND OUTPUT: %s	<b>Cause:</b> An action or event script produced unexpected output. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
004506	ERROR	COMMAND OUTPUT: %s	<b>Cause:</b> An action or event script produced unexpected output. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.

Code	Severity	Message	Cause/Action
004507	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004508	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004509	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004510	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004511	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> An action or event script produced unexpected output.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004512	ERROR		<p><b>Cause:</b> An error occurred on the remote machine.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check the logs on the remote machine for additional details.
004565	ERROR	can't set resource state type to ILLSTATE	<p><b>Cause:</b> An attempt was made to put a resource in an illegal state.</p> <p><b>Action:</b> Do not try to put a resource into an illegal state.</p>
004567	ERROR	split brain detected while setting resource \"%s\" to \"%s\" state (SHARED equivalency to resource \"%s\" on machine \"%s\" which is in state \"%s\"). Setting local resource ISP but aborting the operation.	<p><b>Cause:</b> Changing resource {resource} to state {state} failed since its SHARED equivalent resource {resource} on machine {machine} is in state {state}.</p> <p><b>Action:</b> A split brain situation has occurred. The failover operation has been aborted. You should manually put the split brain resources (ISP on multiple systems) in the proper state.</p>
004575	ERROR	COMMAND OUTPUT: %s	
004607	ERROR	no resource instance has tag \"%s\"	<p><b>Cause:</b> No resource with the provided tag exists.</p> <p><b>Action:</b> Provide a valid tag, or check the logs for related errors and try to resolve the reported problem.</p>
004608	ERROR	no resource instance has identifier \"%s\"	<p><b>Cause:</b> No resource exists with the provided identifier.</p> <p><b>Action:</b> Provide a valid identifier, or check the logs for related errors and try to resolve the reported problem.</p>
004619	ERROR	resource with tag \"%s\" already exists with identifier \"%s\"	<p><b>Cause:</b> The provided tag name already exists.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Choose a different tag name.
004620	ERROR	resource with identifier \"%s\" already exists with tag \"%s\"	<b>Cause:</b> The provided identifier already exists. <b>Action:</b> Choose a different identifier to use for this resource.
004643	ERROR	Instance tag name is too long. It must be shorter than %d characters.	<b>Cause:</b> Tag name is too long. <b>Action:</b> Provide a tag name that is less than 256 characters.
004646	ERROR	Tag name contains illegal characters	<b>Cause:</b> Tag name contains an illegal character. <b>Action:</b> Specify a tag name that does not include one of these characters: _-./
004691	ERROR	can't set both tag and identifier at same time	<b>Cause:</b> Both a tag and an identifier were specified. <b>Action:</b> Provide only one of tag or identifier.
004745	ERROR	failed to access lkxterrlog path=%s	<b>Cause:</b> The utility "lkxterrlog" can not be accessed for collecting system information. <b>Action:</b> Check the installation of package "steeleye-lk" and make sure the utility "lkxterrlog" is accessible.
004746	ERROR	lkxterrlog failed runret=%d cmdline=%s	<b>Cause:</b> The execution of utility "lkxterrlog" failed when collecting system information.

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004782	ERROR	Resource \"%s\" was in state \"%s\" before event occurred – recovery will not be attempted	<p><b>Cause:</b> The resource is already in service. Recovery will not be attempted.</p>
004783	ERROR	Resource \"%s\" was already in state \"%s\" before event occurred	<p><b>Cause:</b> A resource was not in an appropriate state to allow recovery.</p> <p><b>Action:</b> Put the resource in the ISP state if recovery is still needed.</p>
004786	ERROR	%s on failing resource \"%s\"	<p><b>Cause:</b> An error occurred why attempting to recover a resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004788	EMERG	failed to remove resource '%s'. SYSTEM HALTED.	<p><b>Cause:</b> A error occur that prevented a resource from being taken out of service during a recovery. The system has been restarted to ensure the resource is not active on two systems.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004793	ERROR	lcdsendremote transfer resource \"%s\" to \"%s\" on machine \"%s\" failed (rt=%d	<p><b>Cause:</b> A failure occurred while transferring a resource and its dependencies to another system.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Check the log on the other system for related errors.</p>

Code	Severity	Message	Cause/Action
004797	ERROR	Restore of SHARED resource \"%s\" has failed	<p><b>Cause:</b> There was an error while restoring a resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
004806	ERROR	Restore in parallel of resource \"%s\" has failed; will re-try serially	<p><b>Cause:</b> Parallel recovery failed. Check the logs for related errors and try to resolve the reported problem.</p> <p><b>Action:</b> No action is required. The system will continue to recover serially. If the recovery fails, check for error messages related to the resources which failed to recover to find out what further actions to take.</p>
004819	ERROR	read_temporal_recovery_log(): failed to open file: %s. fopen() %s.	<p><b>Cause:</b> The opening of the temporal recovery log file {file} in preparation for loading into it into memory failed with the error {error}.</p> <p><b>Action:</b> Check system log files and correct any reported errors before retrying the operation.</p>
004820	ERROR	read_temporal_recovery_log(): failed to malloc initial buf for temporal_recovery_stamp.	<p><b>Cause:</b> Loading the temporal recovery log information into memory failed when attempting to acquire memory to store the log information.</p> <p><b>Action:</b> Check system log files and correct any reported errors before retrying the operation.</p>
004821	ERROR	read_temporal_recovery_log(): failed to re-allocate buffer for temporal_recovery_stamp.	<p><b>Cause:</b> Loading the temporal recovery log information into memory failed when attempting to increase the amount of m</p>

Code	Severity	Message	Cause/Action
			<p>emory required to store the log information.</p> <p><b>Action:</b> Check system log files and correct any reported errors before retrying the operation.</p>
004822	ERROR	write_temporal_recovery_log(): failed to open file: %s.	<p><b>Cause:</b> The update of the temporal recovery log file was terminated when the open of the temporary file {temporary name} failed.</p> <p><b>Action:</b> Check system log files and correct any reported errors before retrying the operation.</p>
004823	ERROR	rename(%s, %s) failed.	<p><b>Cause:</b> The update of the temporal recovery log file was terminated when the rename of the temporary file {temporary name} to the real log file {real name} failed.</p> <p><b>Action:</b> Check system log files and correct any reported errors before retrying the operation.</p>
004827	ERROR	b	
004829	FATAL	err=%s line=%d Semid=%d numops=%zd perror=%s	<p><b>Cause:</b> The modification of semaphore ID {semaphore} failed with error {err} and error message description {perror}.</p> <p><b>Action:</b> Check adjacent log messages for more details. Also, check the system log files and correct any reported errors before retrying the operation.</p>
004860	ERROR	restore ftok failed for resource %s with path %s	<p><b>Cause:</b> The attempt to generate an IPC key for use in semaphore operations for resource {tag} using path {path} failed.</p>

Code	Severity	Message	Cause/Action
			<p>This is a system error.</p> <p><b>Action:</b> Check adjacent log messages for more details. Also check system log files and correct any reported errors before retrying the operation.</p>
004861	ERROR	semget failed with error %d	<p><b>Cause:</b> The attempt to retrieve the semaphore identification associated with the instances files has failed. This is a system error.</p> <p><b>Action:</b> Check adjacent log messages for more details. Also check system log files and correct any reported errors before retrying the operation.</p>
004862	ERROR	semctl SEMSET failed with error %d	<p><b>Cause:</b> The attempt to create and initialize a semaphore used during the recovery process has failed with the error {error number}. This is a system error.</p> <p><b>Action:</b> Check adjacent log messages for more details. Also, check system log files and correct any reported errors before retrying the operation.</p>
004863	ERROR	semop failed with error %d	<p><b>Cause:</b> The attempt to set a semaphore used during the recovery process has failed with the error {error number}. This is a system error.</p> <p><b>Action:</b> Check adjacent log messages for more details. Also, check system log files and correct any reported errors before retrying the operation.</p>
004864	ERROR	semctl SEMSET failed with error %d	<p><b>Cause:</b> The attempt to release a semaphore used during the recovery process</p>

Code	Severity	Message	Cause/Action
			<p>has failed with the error {error number}. This is a system error.</p> <p><b>Action:</b> Check adjacent log messages for more details. Also, check system log files and correct any reported errors before retrying the operation.</p>
004865	ERROR	restore action failed for resource %s (exit: %d)	<p><b>Cause:</b> The attempt to bring resource {tag} In Service has failed.</p> <p><b>Action:</b> Check adjacent log messages for more details. Correct any reported errors and retry the operation.</p>
004872	ERROR	Remote remove of resource \"%s\" on machine \"%s\" failed (rt=%d)	<p><b>Cause:</b> The request to take resource {tag} Out of Service on {server} for transfer to the local system has failed.</p> <p><b>Action:</b> Check adjacent log messages for more details on the local system. Also, check the log messages on {server} for further details on the failure to remove the resource.</p>
004875	ERROR	remote remove of resource \"%s\" on machine \"%s\" failed	<p><b>Cause:</b> The request to take resource {tag} Out of Service on {server} for transfer to the local system has failed.</p> <p><b>Action:</b> Check adjacent log messages for more details on the local system. Also, check the log messages on {server} for further details on the failure to remove the resource.</p>
004876	ERROR	remote remove of resource \"%s\" on machine \"%s\" failed	<p><b>Cause:</b> The request to take resource {tag} Out of Service on {server} for transfer to the local system has failed.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for more details on the local system. Also, check the log messages on {server} for further details on the failure to remove the resource.</p>
005045	ERROR	tli_fdget_j::execute unable to establish a listener port	<p><b>Cause:</b> A network connection could not be properly configured.</p> <p><b>Action:</b> Verify that all network hardware and drivers are properly configured. If this message continues and resources cannot be put into service, contact Support.</p>
005055	ERROR	tli_fdget_o::execute – async connect failure	<p><b>Cause:</b> A network connection could not be properly configured.</p> <p><b>Action:</b> Verify that all network hardware and drivers are properly configured. If this message continues and resources cannot be put into service, contact Support.</p>
005061	ERROR	tli_fdget_o::execute – bind socket	<p><b>Cause:</b> A network connection could not be properly configured.</p> <p><b>Action:</b> Verify that all network hardware and drivers are properly configured. If this message continues and resources cannot be put into service, contact Support.</p>
005090	WARN	system_driver::add_driver: cmd=%s\n	
005108	WARN	system_driver::rm_driver: cmd=%s\n	
005145	ERROR	opening the file	<p><b>Cause:</b> A pipe could not be opened or created.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check adjacent log messages for more details.
005164	ERROR	tli_handler::handle-error:sending/receiving data message	<p><b>Cause:</b> A message failed to be sent or received.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error but if this error continues and servers can't communicate then verify the network configuration on the servers.</p>
005165	WARN	errno %d\n	<p><b>Cause:</b> A message failed to be sent or received.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error but if this error continues and servers can't communicate then verify the network configuration on the servers.</p>
005166	WARN	poll 0x%hx\n	<p><b>Cause:</b> A message failed to be sent or received.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error but if this error continues and servers can't communicate then verify the network configuration on the servers.</p>
005167	WARN	handler for sys %s\n	<p><b>Cause:</b> A message failed to be sent or received.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error but if this error continues and servers can't communicate then verify the network configuration on the servers.</p>

Code	Severity	Message	Cause/Action
005225	WARN	so_driver::handle_error: sending/receiving data message errno %d: %s	<p><b>Cause:</b> A message failed to be sent or received.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error, but if this error continues and servers cannot communicate, verify the network configuration on the servers.</p>
005235	WARN	found tcp connection to iwstp\n	
005236	WARN	didn't find tcp connection to iwstp\n	
005237	WARN	lcm_handler retry send from %s:%s to %s:%s (%d)\n	
005238	WARN	detected duplicate request from %s:%s to %s:%s\n	
005239	WARN	lcm_handler retry timer set to %d based on lcd remote timeout of %d s (%d s	
005240	WARN	lcm_handler retry count/time (%d/%zu) has exceeded the maximum. Giving up...\n	
005241	WARN	dup_list: %s %s %s %s %d %d (%d	
005242	WARN	clean up stale dup_list entry (%d s): %s:%s to %s:%s last: %s	
005243	WARN	add new dup_list entry %s:%s to %s:%s %d %d (%d	
005244	WARN	closing fd %d\n	
005245	WARN	openpoll fd %d\n	
006012	ERROR	quickCheck script '%s' (%d) failed to exit after %lu seconds. Forcibly terminated. Please examine the script or adjust the LKCHECKINTERVAL parameter in %s.	<p><b>Cause:</b> A quickCheck script is probably taking too long or hanging.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
006014	ERROR	LKCHECKINTERVAL parameter is too short. It is currently set to %ld seconds. It should be at least %ld seconds. Please adjust this parameter in %s and execute 'kill %d' to restart the lkcheck daemon.	

Code	Severity	Message	Cause/Action
006102	ERROR	COMMAND OUTPUT: \$LKROOT/bin/sen devent	<b>Cause:</b> This is output from a "sendevent" (event generator) command.  <b>Action:</b> Check adjacent log messages for more details.
006103	ERROR	COMMAND OUTPUT: \$LKROOT/bin/sen devent	<b>Cause:</b> This is output from a "sendevent" (event generator) command.  <b>Action:</b> Check adjacent log messages for more details.
006104	ERROR	COMMAND OUTPUT: \$LKROOT/bin/sen devent	<b>Cause:</b> This is output from a "sendevent" (event generator) command.  <b>Action:</b> Check adjacent log messages for more details.
006502	ERROR	CPU usage has exceeded the threshold (\$threshold%) for \$count check cycles.	
006504	ERROR	Could not open /proc/meminfo	
006505	ERROR	Memory usage has exceeded the threshold (\$threshold%) for \$count check cycles.	
006508	ERROR	[\$SUBJECT event] mail returned \$err	
006509	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
006511	ERROR	snmptrap returned \$err for Trap 190	
006512	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
006514	ERROR	[\$SUBJECT event] mail returned \$err	
006515	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
006517	ERROR	snmptrap returned \$err for Trap 200	
006518	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
006520	ERROR	Failed to update error count in \$cpu_file: \$!	
006521	ERROR	Failed to update error count in \$mem_file: \$!	

Code	Severity	Message	Cause/Action
006523	ERROR	The SNHC_CPUCHECK_THRESHOLD setting is not valid. Please set SNHC_CPU CHECK_THRESHOLD to a value between 10 and 99 in /etc/default/LifeKeeper. If not set, the value will default to 99.	
006524	ERROR	The SNHC_CPUCHECK_TIME setting is not valid. Please set SNHC_CPUCHECK_TIME to a value between 1 and 100 in /etc/default/LifeKeeper. If not set, the value will default to 1.	
006525	ERROR	The SNHC_MEMCHECK_THRESHOLD setting is not valid. Please set SNHC_MEMCHECK_THRESHOLD to a value between 10 and 99 in /etc/default/LifeKeeper. If not set, the value will default to 99.	
006526	ERROR	The SNHC_MEMCHECK_TIME setting is not valid. Please set SNHC_MEMCHECK_TIME to a value between 1 and 100 in /etc/default/LifeKeeper. If not set, the value will default to 1.	
006528	ERROR	Could not open /proc/stat	
006529	ERROR	Could not open /proc/stat	
006530	ERROR	Could not use \$tmp_path	
007053	ERROR	malloc failed. Assume that it is a monitoring target device.	
007058	ERROR	%s: %s failed on '%s', result:%d, Sense Key = %d.	<p><b>Cause:</b> A SCSI device couldn't be reserved or have its status checked. This may be because the storage is malfunctioning or because the disk has been reserved by another server.</p> <p><b>Action:</b> Check adjacent log messages for more details and verify that resources are being handled properly.</p>
007059	ERROR	%s: %s failed on '%s', result:%d.	<p><b>Cause:</b> A SCSI device couldn't be reserved or have its status checked. This may be because the storage is malfunctioning</p>

Code	Severity	Message	Cause/Action
			<p>ing or because the disk has been reserved by another server.</p> <p><b>Action:</b> Check adjacent log messages for more details and verify that resources are being handled properly.</p>
007060	EMERG	%s: failure on device '%s'. SYSTEM HALTED.	<p><b>Cause:</b> A SCSI device couldn't be reserved or have its status checked. This may be because the storage is malfunctioning or because the disk has been reserved by another server. THE SERVER WILL BE REBOOTED/HALTED.</p> <p><b>Action:</b> Verify that the storage is functioning properly and, if so, that resources were handled properly and have been put in service on another server.</p>
007072	ERROR	%s: failed to open SCSI device '%s', initiate recovery. errno=0x%x, retry count=%d.	<p><b>Cause:</b> The protected SCSI device could not be opened. The device may be failing or may have been removed from the system.</p> <p><b>Action:</b> The system will be halted or a failover to the backup node will be initiated. The default action in this case is a failover, but this can be modified with the SCSIERROR tunable.</p>
007073	ERROR	%s: failed to open SCSI device '%s', RETRY. errno=%d, retry count=%d.	<p><b>Cause:</b> The protected SCSI device could not be opened. The device may be failing or may have been removed from the system.</p> <p><b>Action:</b> This error is not critical. The operation will be retried in 5 seconds. If the problem persists, the system will perform a halt or resource failover.</p>

Code	Severity	Message	Cause/Action
007075	ERROR	%s: RESERVATION CONFLICT on SCSI device '%s'. ret=%d, errno=0x%x, retry count=%d.	<p><b>Cause:</b> A SCSI device couldn't be reserved due to a conflict with another server. This may be because the storage is malfunctioning or because the disk has been reserved by another server.</p> <p><b>Action:</b> Check adjacent log messages for more details and verify that resources are handled properly.</p>
007077	ERROR	%s: DEVICE FAILURE on SCSI device '%s', initiate recovery. ret=%d, errno=0x%x, retry count=%d.	<p><b>Cause:</b> A SCSI device couldn't be reserved or have its status checked. This may be because the storage is malfunctioning or because the disk has been reserved by another server.</p> <p><b>Action:</b> Check adjacent log messages for more details and verify that resources are handled properly.</p>
007078	ERROR	%s: DEVICE FAILURE on SCSI device '%s', RETRY. ret=%d, errno=0x%x, retry count=%d.	<p><b>Cause:</b> A SCSI device couldn't be reserved or have its status checked. This may be because the storage is malfunctioning or because the disk has been reserved by another server.</p> <p><b>Action:</b> Check adjacent log messages for more details and verify that resources are handled properly.</p>
010002	WARN	flag \$flag not present, send message again.	<p><b>Cause:</b> This message indicates an incomplete process that will be retried.</p> <p><b>Action:</b> Check adjacent log messages for repeated errors.</p>
010003	ERROR	COMMAND OUTPUT: \$LKBIN/ins_remove	<p><b>Cause:</b> This message is part of the output from an "ins_remove" command.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for more details. This may not be a true error.</p>
010006	WARN	flg_list -d \$i took more than \$pswait seconds to complete...	<p><b>Cause:</b> A flag list operation on a server took much longer than expected. There may be a problem communicating with the other server.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>
010007	ERROR	flag \$flag not present, switchovers may occur.	<p><b>Cause:</b> One of the servers in the cluster could not be told to disallow failover operations from the current server.</p> <p><b>Action:</b> Check adjacent log messages for more details and monitor the cluster for unexpected behavior.</p>
010008	WARN	flag \$flag not present, send message again.	<p><b>Cause:</b> A process is incomplete but will be retried.</p> <p><b>Action:</b> Check adjacent log messages for more details and for repeated warnings/errors.</p>
010023	FATAL	LifeKeeper failed to initialize properly.	<p><b>Cause:</b> There was a fatal error while attempting to start LifeKeeper.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>
010025	ERROR	`printf 'Unable to get a unique tag name on server "%s" for template resource "%s" \$MACH \$DISK`	<p><b>Cause:</b> A suitable tag during the create process for a storage resource could not be automatically generated.</p> <p><b>Action:</b> Check adjacent log messages for</p>

Code	Severity	Message	Cause/Action
			or more details. Retry the operation if there are no other errors.
010034	FATAL	Unable to start lcm.	<p><b>Cause:</b> A core component of the software could not be started.</p> <p><b>Action:</b> Check adjacent log messages for more details and try to resolve the reported problem.</p>
010038	WARN	Waiting for LifeKeeper core components to initialize has exceeded 10 seconds. Continuing anyway, check logs for further details.	<p><b>Cause:</b> Some parts of the software are taking longer than expected to start up.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
010039	WARN	Waiting for LifeKeeper core components to initialize has exceeded 10 seconds. Continuing anyway, check logs for further details.	<p><b>Cause:</b> Some parts of the software are taking longer than expected to start up.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
010046	ERROR	The dependency creation failed on server \$SERVER:" `cat \$TEMP_FILE`	<p><b>Cause:</b> A dependency relationship could not be created on the given server.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem</p>
010063	ERROR	\$REMSH error	<p><b>Cause:</b> A command to request data to be backup from another server failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010085	ERROR	lkswitchback(\$MACH): Automatic switchback of \"\$loctag\" failed	<p><b>Cause:</b> The resource was not switched</p>

Code	Severity	Message	Cause/Action
			<p>back over as expected.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010102	ERROR	admin machine not specified	<p><b>Cause:</b> Invalid parameters were specified for the "getlocks" operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation. If this error happens during normal operation, contact Support.</p>
010107	WARN	Lock for \$m is ignored because system is OOS	<p><b>Cause:</b> A lock was ignored because the system for which the lock was created is not alive.</p> <p><b>Action:</b> Check the logs for related errors. This may be a harmless error.</p>
010108	ERROR	lock acquisition timeout	<p><b>Cause:</b> Acquiring a lock took longer than expected/allowed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010109	ERROR	could not get admin locks." `cat /tmp/ERR\$\$`	<p><b>Cause:</b> The software failed to acquire a lock that is required to manage resources.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010112	ERROR	lcdrpc failed with error no: \$LCDRPCRES	<p><b>Cause:</b> A file could not be copied to another server.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010116	ERROR	unable to set !lkstop flag	<p><b>Cause:</b> A flag could not be set to indicate that the server is being stopped by user request.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010121	ERROR	Extended logs aborted due to a failure in opening \$destination. (\$syserrmsg	<p><b>Cause:</b> The execution of utility "lkexterrlog" failed when opening the extended log file.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010132	ERROR	Unable to retrieve reservation id from "%s". Error: "%s". Attempting to regenerate.	<p><b>Cause:</b> Unable to open the file /opt/LifeKeeper/config/.reservation_id to retrieve the unique id used for SCSI 3 persistent reservations.</p> <p><b>Action:</b> None. An attempt will be made to regenerate the ID and update the file.</p>
010135	ERROR	The current reservation ID of "%s" is not unique within the cluster. A new ID must be generated by running "%s/bin/genresid -g" on "%s".	<p><b>Cause:</b> The reservation id defined for the system is not unique within the cluster and cannot be used.</p> <p><b>Action:</b> Take all resources out of service on this node and then run "/opt/LifeKeeper/bin/genresid -g" to generate a unique reservation id.</p>
010136	ERROR	Unable to store reservation id in "%s". Error: "%s"	<p><b>Cause:</b> Unable to open the file /opt/Life</p>

Code	Severity	Message	Cause/Action
			<p>Keeper/config.reservation_id to store the unique id used for SCSI 3 persistent reservations.</p> <p><b>Action:</b> Correct the error listed as the reason the open failed and then take all resources out of service on this node and then run "/opt/LifeKeeper/bin/genresid -g" to generate a new unique reservation id.</p>
010137	ERROR	Failed to generate a reservation ID that is unique within the cluster.	<p><b>Cause:</b> The generated reservation id is already defined on another node in the cluster and must be unique within the cluster.</p> <p><b>Action:</b> Take all resources out of service on this node and then run "/opt/LifeKeeper/bin/genresid -g" to generate a new unique reservation id.</p>
010138	ERROR	\$message	
010139	WARN	\$message	
010140	ERROR	\$COMMAND_SNMPTRAP returned \$exit code for Trap \$oid:\$result	
010141	ERROR	LK_TRAP_MGR is specified in /etc/default/LifeKeeper, but \$COMMAND_SNMPTRAP command is not in PATH.	
010142	ERROR	\$COMMAND_EMAIL returned \$exitcode:\$result	
010143	ERROR	LK_NOTIFY_ALIAS is specified in /etc/default/LifeKeeper, but \$COMMAND_EMAIL command is not in PATH.	
010144	ERROR	can't opendir \$LICENSE_DIR: \$!	
010145	ERROR	lktest failed	
010146	ERROR	lkcheck failed	
010147	ERROR	ins_list failed: exit code = \$exit_code	
010159	ERROR	Maintenance mode disable currently in pr	

Code	Severity	Message	Cause/Action
		ogress, can't enable maintenance mode. If this problem persists, consider using the <code>—force</code> option.	
010160	ERROR	Maintenance mode enable currently in progress, can't disable maintenance mode. If this problem persists, consider using the <code>—force</code> option.	
010161	ERROR	<code>\$tag</code> is not a valid resource tag on the local machine, aborting. Please check the spelling and try again.	
010163	ERROR	<code>\$cmd</code> script not found or not executable on system <code>\$sys</code> .	
010165	ERROR	An error occurred while running <code>\\$LKROOT/ikadm/subsys/appsuite/sap/bin/\$cmd -m=\${opt_mode}\${tag_cmd}\${force_cmd}</code> on system <code>\$sys</code> (exit code: <code>\$remexc_ret</code> ). Please inspect the logs on that system for more information.	
010168	WARN	Maintenance mode was not fully <code>\${opt_mode}</code> d for at least one resource on system <code>\$sys</code> .	
010172	ERROR	Maintenance mode was not fully <code>\${opt_mode}</code> d for the requested resources on at least one system in the cluster.	
010173	ERROR	LifeKeeper is not running on system <code>\$me</code> . Unable to <code>\${opt_mode}</code> maintenance mode. Aborting.	
010179	ERROR	Maintenance mode action <code>\`\${opt_mode}</code> ' did not complete successfully for resource <code>\$tag</code> on system <code>\$me</code> .	
010181	ERROR	An error occurred while attempting maintenance mode action <code>\`\${opt_mode}</code> ' on system <code>\$me</code> for resources: <code>@{local_hier_tags}</code> .	
010187	ERROR	Resource <code>\$tag</code> has not been extended to system <code>\$sys</code> .	
010222	ERROR	<code>scsifree(%s): LKSCSI_Release(%s) unsuccessful</code>	<b>Cause:</b> A SCSI device that appeared to be reserved was not released as expected

Code	Severity	Message	Cause/Action
			<p>ed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve any reported problems. This error may be benign if the system is functioning properly.</p>
010231	ERROR	scsipllock(%s): reserve failed.	<p><b>Cause:</b> A reservation on a SCSI device could not be acquired.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010250	ERROR	Failed to exec command '%s'	<p><b>Cause:</b> The "lklogmsg" tool failed to execute a sub-command {command}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve any reported problems. Verify that the sub-command exists and is a valid command or program. If this message happens during normal operation, contact Support.</p>
010256	ERROR	scsi_tur(%s): open failed.	
010260	ERROR	scsi_tur(%s): test unit ready failed.	
010402	EMERG	local recovery failure on resource \$opts{'N'}, trigger VMware HA...	<p><b>Cause:</b> When in LifeKeeper Single Server Protection operation, a resource could not be recovered and VMware-HA is about to be triggered to handle the failure (if VMware-HA is enabled).</p> <p><b>Action:</b> No action is required. VMware should handle the failure.</p>
010413	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	<p><b>Cause:</b> This is the output from an "snmptrap" command that may have failed.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010420	EMERG	local recovery failure on resource \$opts{'N'}, trigger reboot...	<p><b>Cause:</b> When in LifeKeeper Single Server Protection operation, a resource could not be recovered and a reboot is about to be triggered to handle the failure.</p> <p><b>Action:</b> No action is required.</p>
010440	ERROR	[\$SUBJECT event] mail returned \$err	<p><b>Cause:</b> This indicates a notification email could not be sent via the "mail" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010443	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	<p><b>Cause:</b> This is the output from a "mail" command that may have failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010445	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	<p><b>Cause:</b> This is the output from a "mail" command that may have failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
010463	ERROR	LifeKeeper: name of machine is not specified, ARGS=\$ARGS	<p><b>Cause:</b> Invalid arguments were specified for the "comm_down" event.</p> <p><b>Action:</b> Check your LifeKeeper configuration and retry the operation.</p>

Code	Severity	Message	Cause/Action
010471	ERROR	COMM_DOWN: Attempt to obtain local comm_down lock flag failed	<p><b>Cause:</b> During the handling of a communication failure with another node, a local lock could not be acquired. This will likely stop a failover from proceeding.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. If failovers are not taking place properly, contact Support.</p>
010482	ERROR	LifeKeeper: name of machine is not specified, ARGS=\$ARGS	<p><b>Cause:</b> Invalid arguments were specified for the "comm_up" event.</p> <p><b>Action:</b> Check your LifeKeeper configuration and retry the operation.</p>
010484	WARN	flg_list -d \$MACH check timed-out (\$delay seconds).	<p><b>Cause:</b> "flg_list" command reached its timeout value {delay} seconds.</p>
010487	WARN	flg_list -d \$MACH check timed-out, unintended switchovers may occur.	<p><b>Cause:</b> "flg_list" command reached its timeout value.</p> <p><b>Action:</b> Switch back the resource tree if unintended switchover occurs.</p>
010492	WARN	\$m	<p><b>Cause:</b> One of other servers looked dead to this server {server}, but witness servers did not agree.</p> <p><b>Action:</b> Ensure other server is dead and switch over the resource manually.</p>
010494	ERROR	LifeKeeper: COMM_UP to machine \$MACH completed with errors.	<p><b>Cause:</b> An unexpected failure occurred during "COMM_UP" event.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>

Code	Severity	Message	Cause/Action
010503	ERROR	lcdrecover hung or returned error, attempting kill of process \$FPID	<b>Cause:</b> "lcdrecover" took too long or errored out.
010506	ERROR	Intelligent Switchback Check Failed	<b>Cause:</b> Failed 5 times to perform "lcdrecover." <b>Action:</b> Switch over the resource tree manually.
010535	ERROR	LifeKeeper: name of machine is not specified, ARGS=\$ARGS	
010600	ERROR	removing hierarchy remnants	
010627	WARN	Equivalency Trim: does not have a full complement of equivalencies. Hierarchy will be unextended from	
010629	WARN	Your hierarchy exists on only one server. Your application has no protection until you extend it to at least one other server.	
010712	ERROR	Unextend hierarchy failed	<b>Cause:</b> A resource hierarchy failed to be unextended from a server. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
010746	ERROR	\$ERRMSG Target machine \"\$TARGET_MACH\" does not have an active LifeKeeper communication path to machine \"\$aMach\" in the hierarchy." >&2	<b>Cause:</b> A hierarchy cannot be unextended because the target server does not have adequate communication with the other servers in the cluster. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Ensure that all servers have communication paths to each other.
010763	ERROR	lock failed	

Code	Severity	Message	Cause/Action
011000	ERROR	appremote: unknown command type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011001	ERROR	depremote: unknown command type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011002	ERROR	eqvremote: unknown command type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011003	ERROR	flgremote: unknown command type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011004	WARN	Illegal creation of resource	<b>Cause:</b> This will not occur under normal circumstances.
011009	ERROR	insremote: unknown change field command type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011010	ERROR	insremote: unknown command type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011011	FATAL	%s	<b>Cause:</b> LifeKeeper could not get IPC key. <b>Action:</b> Check adjacent log messages for more details.
011012	FATAL	semget(%s,%c	<b>Cause:</b> LifeKeeper could not get semap

Code	Severity	Message	Cause/Action
			<p>hore set id.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>
011013	FATAL	shmget(%s,%c	<p><b>Cause:</b> System could not allocate a shared memory segment.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>
011014	FATAL	prefix_lkroot("out"	<p><b>Cause:</b> A system error has occurred while accessing /opt/LifeKeeper/out.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper/out is not accessible.</p>
011015	ERROR	DEMO_UPGRADE_MSG	<p><b>Cause:</b> You are running a demo license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011016	ERROR	lic_single_node_msg	<p><b>Cause:</b> You have a license for LifeKeeper Single Server Protection but you do not have LifeKeeper Single Server Protection installed.</p> <p><b>Action:</b> Either install LifeKeeper Single Server Protection or obtain a license that matches the product you are running.</p>
011018	ERROR	lic_init_fail_msg, lc_errstring(lm_job	<p><b>Cause:</b> License manager initialization failed.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>

Code	Severity	Message	Cause/Action
011020	EMERG	lic_init_fail_msg, lc_errstring(lm_job	<p><b>Cause:</b> License manager initialization failed.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>
011021	EMERG	lic_error_msg, lc_errstring(lm_job	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011022	EMERG	lic_error_msg, lc_errstring(lm_job	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011023	EMERG	lic_no_rest_suite, ""	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011024	EMERG	lic_error_msg, lc_errstring(lm_job	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011025	EMERG	lic_no_license, ""	<p><b>Cause:</b> LifeKeeper could not find valid license keys.</p> <p><b>Action:</b> Ensure license keys are valid for the server and retry the operation.</p>
011026	EMERG	lic_error_msg, lc_errstring(lm_job	<p><b>Cause:</b> There is an unknown problem w</p>

Code	Severity	Message	Cause/Action
			<p>ith your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011027	EMERG	lic_no_license, ""	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011028	ERROR	lang_error_msg	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011029	FATAL	can't set reply system	<p><b>Cause:</b> A message failed to be sent.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error.</p>
011030	FATAL	can't set reply mailbox	<p><b>Cause:</b> A message failed to be sent.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error.</p>
011031	ERROR	Failure reading output of '%s' on behalf of %s	<p><b>Cause:</b> A system error has occurred while accessing temporary file /tmp/OUT.{pid}.</p> <p><b>Action:</b> Determine why /tmp/OUT.{pid} is not accessible.</p>
011032	ERROR	Failure reading output of '%s'	<p><b>Cause:</b> A system error has occurred wh</p>

Code	Severity	Message	Cause/Action
			<p>ile accessing temporary file /tmp/ERR.{pid}.</p> <p><b>Action:</b> Determine why /tmp/ERR.{pid} is not accessible.</p>
011033	ERROR	event \"%s,%s\" already posted for resource with id \"%s\"	<p><b>Cause:</b> This message is for information only.</p>
011034	ERROR	no resource has id of \"%s\"	<p><b>Cause:</b> LifeKeeper could not find the {id} resource.</p> <p><b>Action:</b> Verify the parameters and retry the "sendevent" operation.</p>
011044	ERROR	flagcleanup:fopen(%s	<p><b>Cause:</b> A system error has occurred while reading /opt/LifeKeeper/config/flg.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper/config/flg is not readable.</p>
011045	ERROR	flagcleanup:fopen(%s	<p><b>Cause:</b> A system error has occurred while writing /opt/LifeKeeper/config/flg.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper/config/flg is not writable.</p>
011046	ERROR	flagcleanup:fputs(%s	<p><b>Cause:</b> A system error has occurred while writing /opt/LifeKeeper/config/flg.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper/config/flg is not writable.</p>
011047	ERROR	flagcleanup:rename(%s,%s	<p><b>Cause:</b> A system error has occurred while renaming /opt/LifeKeeper/config/.flg to /opt/LifeKeeper/config/flg.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Determine why /opt/LifeKeeper/config/flg was not able to be renamed.
011048	ERROR	flagcleanup:chmod(%s	<p><b>Cause:</b> A system error has occurred while changing permissions in /opt/LifeKeeper/config/flg.</p> <p><b>Action:</b> Determine why LifeKeeper could not change permissions in /opt/LifeKeeper/config/flg.</p>
011049	ERROR	License check failed with error code %d	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011051	ERROR	lcdinit: clearing Disk Reserve file failed	<p><b>Cause:</b> A system error has occurred while writing /opt/LifeKeeper/subsys/scsi/resources/disk/disk.reserve.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper/subsys/scsi/resources/disk/disk.reserve is not writable.</p>
011052	FATAL	malloc() failed	<p><b>Cause:</b> The system could not allocate memory for LifeKeeper.</p> <p><b>Action:</b> Increase the process limit for the data segment.</p>
011053	FATAL	lcm_is_unavail	<p><b>Cause:</b> A system error has occurred while writing /tmp/LK_IS_UNAVAIL.</p> <p><b>Action:</b> Determine why /tmp/LK_IS_UNAVAIL is not writable.</p>

Code	Severity	Message	Cause/Action
011054	FATAL	lk_is_unavail	<p><b>Cause:</b> A system error has occurred while writing /opt/LifeKeeper/config/LK_IS_ON.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper/config/LK_IS_ON is not writable.</p>
011055	FATAL	usr_alarm_config_LK_IS_ON	<p><b>Cause:</b> A system error has occurred while writing /tmp/LCM_IS_UNAVAI.</p> <p><b>Action:</b> Determine why /tmp/LCM_IS_UNAVAI is not writable.</p>
011056	ERROR	License check failed with error code %d	<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011057	ERROR	lcdremote: unknown command type %d('%c')\n	<p><b>Cause:</b> A message failed to be received.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error, but if this error continues and servers cannot communicate, verify the network configuration on the servers.</p>
011059	FATAL	Could not write to: %s	<p><b>Cause:</b> A system error has occurred while accessing /opt/LifeKeeper/config/LK_START_TIME.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper/config/LK_START_TIME is not accessible.</p>
011060	FATAL	received NULL message	<p><b>Cause:</b> A message failed to be received.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error, but if this error continues and servers cannot communicate, verify the network configuration on the servers.</p>
011061	ERROR	unknown data type %d('%c') on machine \"%s\" \\%s\\n	<p><b>Cause:</b> A message failed to be received.</p> <p><b>Action:</b> Check adjacent log messages for more details. This may be a temporary error, but if this error continues and servers cannot communicate, verify the network configuration on the servers.</p>
011062	WARN	LifeKeeper shutdown in progress. Unable to perform failover recovery processing for %s\\n	<p><b>Cause:</b> LifeKeeper was unable to fail over the given resource during shutdown.</p> <p><b>Action:</b> Switch over the resource tree to other server manually.</p>
011063	WARN	LifeKeeper resource initialization in progress. Unable to perform failover recovery processing for %s\\n	<p><b>Cause:</b> LifeKeeper was unable to fail over the given resource during start up.</p> <p><b>Action:</b> Switch over the resource tree manually after LifeKeeper starts up.</p>
011068	ERROR	ERROR on command %s	<p><b>Cause:</b> An error occurred while running the "rlslocks" command.</p> <p><b>Action:</b> Check adjacent messages for more details.</p>
011070	ERROR	ERROR on command %s	<p><b>Cause:</b> An error occurred while running the "getlocks" command.</p> <p><b>Action:</b> Check adjacent log messages for more details.</p>

Code	Severity	Message	Cause/Action
011080	FATAL	out of memory	<p><b>Cause:</b> Internal error.</p> <p><b>Action:</b> Increase the process limit for the data segment.</p>
011081	FATAL	Failed to ask ksh to run: %s	<p><b>Cause:</b> A system error has occurred while invoking ksh.</p> <p><b>Action:</b> Make sure the pdksh (v8.0 and earlier) or the steeleye-pdksh (v81 and later) package is installed.</p>
011082	ERROR	Failed to remove: %s	<p><b>Cause:</b> A system error has occurred while removing /tmp/LCM_IS_UNAVAIL.</p> <p><b>Action:</b> Determine why /tmp/LCM_IS_UNAVAIL is not removable.</p>
011083	ERROR	Failed to remove: %s	<p><b>Cause:</b> A system error has occurred while trying to unlink /tmp/LK_IS_UNAVAIL.</p> <p><b>Action:</b> Determine why /tmp/LK_IS_UNAVAIL is not removable.</p>
011084	FATAL	Failed to generate an IPC key based on: %s	<p><b>Cause:</b> A system error has occurred while accessing /opt/LifeKeeper.</p> <p><b>Action:</b> Determine why /opt/LifeKeeper is not accessible.</p>
011085	ERROR	semget(%s,%c) failed	<p><b>Cause:</b> A system error has occurred while removing a semaphore.</p> <p><b>Action:</b> Try to remove the semaphore manually.</p>

Code	Severity	Message	Cause/Action
011086	ERROR	shmget(%s,%c) failed	<p><b>Cause:</b> A system error has occurred while removing a shared memory segment.</p> <p><b>Action:</b> Try to remove the shared memory segment manually.</p>
011087	ERROR	semctl(IPC_RMID) failed	<p><b>Cause:</b> A system error has occurred while removing a semaphore.</p> <p><b>Action:</b> Try to remove the semaphore manually.</p>
011088	ERROR	shmctl(IPC_RMID) failed	<p><b>Cause:</b> A system error has occurred while removing a shared memory segment.</p> <p><b>Action:</b> Try to remove the shared memory segment manually.</p>
011089	FATAL	Execution of lcdstatus on remote system <%s> failed\n	<p><b>Cause:</b> The remote {node} is down, inaccessible via the network or some other system problem occurred on the remote node.</p> <p><b>Action:</b> Bring the remote node back online, or check adjacent messages for additional information, or check the logs on the remote node for additional information.</p>
011090	FATAL		<p><b>Cause:</b> Internal error.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
011091	WARN		<p><b>Cause:</b> This will not occur under normal circumstances.</p>

Code	Severity	Message	Cause/Action
011092	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
011093	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
011094	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
011095	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
011096	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
011097	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011098	FATAL		<p><b>Cause:</b> There is a problem with your lic</p>

Code	Severity	Message	Cause/Action
			<p>ense.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011099	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011100	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011101	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011102	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011103	FATAL		<p><b>Cause:</b> There is a problem with your license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011104	FATAL		<p><b>Cause:</b> There is a problem with your license.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Contact Support to obtain a new license.
011105	FATAL		<b>Cause:</b> There is a problem with your license. <b>Action:</b> Perform the steps listed in the message text.
011111	ERROR	action \"%s\" on resource with tag \"%s\" has failed	<b>Cause:</b> The {action} for resource {tag} has failed. <b>Action:</b> See adjacent error messages for further details.
011112	ERROR		<b>Cause:</b> LifeKeeper could not find the network device. <b>Action:</b> Check your LifeKeeper configuration.
011113	ERROR	netremote: unknown subcommand type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011114	ERROR	netremote: unknown command type %d('%c')\n	<b>Cause:</b> Internal error. <b>Action:</b> Try restarting the product.
011117	ERROR	sysremote: system \"%s\" not found on \"%s\"	<b>Cause:</b> An invalid system name was provided. <b>Action:</b> Recheck the system name and rerun the command.

Code	Severity	Message	Cause/Action
011118	ERROR	sysremote: unknown subcommand type %d('%c')\n	<p><b>Cause:</b> Internal error.</p> <p><b>Action:</b> Try restarting the product.</p>
011119	ERROR	sysremote: unknown command type %d('%c')\n	<p><b>Cause:</b> Internal error.</p> <p><b>Action:</b> Try restarting the product.</p>
011120	ERROR	typremote: unknown command type %d('%c')\n	<p><b>Cause:</b> Internal error.</p> <p><b>Action:</b> Try restarting the product.</p>
011129	ERROR	Failure during run of '%s' on behalf of %s	<p><b>Cause:</b> Command execution failed.</p> <p><b>Action:</b> See message details to determine the problem.</p>
011130	ERROR	%s	<p><b>Cause:</b> The command {command} produced unexpected output.</p> <p><b>Action:</b> Action should be determined by the content of adjacent error messages.</p>
011131	EMERG	demo_update_msg	<p><b>Cause:</b> There is a problem with your demo license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011132	EMERG	demo_tamper_msg	<p><b>Cause:</b> You have a demo license and clock tampering has been detected.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011133	EMERG	demo_tamper_msg	<p><b>Cause:</b> You have a demo license and cl</p>

Code	Severity	Message	Cause/Action
			<p>clock tampering has been detected.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011134	EMERG	demo_expire_msg	<p><b>Cause:</b> The demo license for this product has expired.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011135	EMERG	demo_tamper_msg	<p><b>Cause:</b> You have a demo license and clock tampering has been detected.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011136	EMERG	buf	<p><b>Cause:</b> You are running a demo license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011138	EMERG	buf	<p><b>Cause:</b> You are running a demo license.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011142	WARN	LifeKeeper Recovery Kit %s license key NOT FOUND	<p><b>Cause:</b> An Application Recovery Kit license for {kit} was not found.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011150	ERROR	COMMAND OUTPUT: %s	<p><b>Cause:</b> The command "eventslicm" produced unexpected output.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
011151	EMERG	&localebuf <sup>3</sup>	<p><b>Cause:</b> This version of the LifeKeeper core package is restricted to being used within the territories of the People's Republic of China or Japan.</p>
011152	EMERG	Localized license failure	<p><b>Cause:</b> There was a mis-match between your locale and the locale for which the product license was created.</p> <p><b>Action:</b> Contact Support to obtain a new license which matches your locale.</p>
011154	EMERG	Single Node flag check failed.	<p><b>Cause:</b> You have a license for LifeKeeper Single Server Protection but you do not have LifeKeeper Single Server Protection installed.</p> <p><b>Action:</b> Either install LifeKeeper Single Server Protection or obtain a license that matches the product you are running.</p>
011155	EMERG	lic_master_exp_msg, ""	<p><b>Cause:</b> Your license key for this product has expired.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>
011162	EMERG	lic_restricted_exp_msg, ""	<p><b>Cause:</b> Your license key for this product has expired.</p> <p><b>Action:</b> Contact Support to obtain a new license.</p>

Code	Severity	Message	Cause/Action
011163	EMERG	Single Node license check failed	<p><b>Cause:</b> You have a license for LifeKeeper Single Server Protection but you do not have LifeKeeper Single Server Protection installed.</p> <p><b>Action:</b> Either install LifeKeeper Single Server Protection or obtain a license that matches the product you are running.</p>
011164	EMERG	demo_expire_msg, DEMO_UPGRADE_MSG	<p><b>Cause:</b> Your license key for this product has expired.</p> <p><b>Action:</b> Please contact Support to obtain a permanent license key for your product.</p>
011165	ERROR	LifeKeeper initialize timed out in tag \"%s\"	
015000	ERROR	COMMAND OUTPUT: /opt/LifeKeeper/sbin/chpst	<p><b>Cause:</b> An error has occurred with the "steeleye-lighttpd" process. Specific details of the error are included in the actual log message.</p> <p><b>Action:</b> Correct the configuration and "steeleye-lighttpd" will automatically be restarted.</p>
103001	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The db2nodes.cfg does not contain any server names.</p> <p><b>Action:</b> Ensure the db2nodes.cfg is valid.</p>
103002	ERROR	LifeKeeper was unable to get the version for the requested instance \"%s\"	<p><b>Cause:</b> "db2level" command did not return DB2 version.</p> <p><b>Action:</b> Check your DB2 configuration.</p>

Code	Severity	Message	Cause/Action
103003	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103004	ERROR	Unable to get the information for resource "%s"	<p><b>Cause:</b> Failed to get resource information.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103005	ERROR	Unable to get the information for resource "%s"	<p><b>Cause:</b> Failed to get resource information.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103006	ERROR	Unable to get the instance information for resource "%s"	<p><b>Cause:</b> Failed to get the instance information.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103007	ERROR	Unable to get the instance home directory information for resource "%s"	<p><b>Cause:</b> Failed to get the instance home directory path.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103008	ERROR	Unable to get the instance type information for resource "%s"	<p><b>Cause:</b> The DB2 Application Recovery Kit found invalid instance type.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103009	ERROR	LifeKeeper has encountered an error while trying to get the database configuration	<p><b>Cause:</b> There was an unexpected error</p>

Code	Severity	Message	Cause/Action
		parameters for database \"\$DB\"	<p>running "db2 get db cfg for \$DB" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103012	ERROR	LifeKeeper was unable to start the database server for instance "%s"	<p><b>Cause:</b> The requested startup of the DB 2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "restore" operation.</p>
103013	ERROR	LifeKeeper was unable to start the database server for instance "%s"	<p><b>Cause:</b> The requested startup of the DB 2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "restore" operation.</p>
103015	ERROR	An entry for the home directory "%s" of instance "%s" does not exist in "/etc/fstab"	<p><b>Cause:</b> The home directory of instance of Multiple Partition database should exist in "/etc/fstab".</p> <p><b>Action:</b> Ensure the home directory exists in "/etc/fstab".</p>
103016	ERROR	LifeKeeper was unable to mount the home directory for the DB2 instance "%s"	<p><b>Cause:</b> Failed to mount the home directory of instance of Multiple Partition database.</p> <p><b>Action:</b> Ensure the home directory is mounted and retry the operation.</p>
103017	ERROR	Unable to get the instance nodes information for resource "%s"	<p><b>Cause:</b> Failed to get the instance node</p>

Code	Severity	Message	Cause/Action
			<p>s.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103018	ERROR	LifeKeeper was unable to start database partition server "%s" for instance "%s"	<p><b>Cause:</b> The requested startup of the DB 2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before re trying the "restore" operation.</p>
103020	ERROR	LifeKeeper was unable to stop the database server for instance "%s"	<p><b>Cause:</b> The requested shutdown of the DB2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before re trying the "remove" operation.</p>
103021	ERROR	LifeKeeper was unable to stop the database server for instance "%s"	<p><b>Cause:</b> The requested shutdown of the DB2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before re trying the "remove" operation.</p>
103023	ERROR	Unable to get the instance nodes information for resource "%s"	<p><b>Cause:</b> Failed to get the instance nodes.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103024	ERROR	LifeKeeper was unable to stop database partition server "%s" for instance "%s"	<p><b>Cause:</b> The requested shutdown of the DB2 instance failed.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "remove" operation.</p>
103026	ERROR	Unable to get the instance nodes information for resource "%s"	<p><b>Cause:</b> Failed to get the instance nodes.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103027	FATAL	The argument for the DB2 instance is empty	<p><b>Cause:</b> Invalid parameters were specified for the create operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103028	FATAL	Unable to determine the DB2 instance home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>
103029	FATAL	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103030	FATAL	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>

Code	Severity	Message	Cause/Action
103031	ERROR	The path "%s" is not on a shared filesystem	<p><b>Cause:</b> The instance home directory should be on a shared filesystem.</p> <p><b>Action:</b> Ensure the path is on shared filesystem and retry the create operation.</p>
103032	ERROR	LifeKeeper was unable to get the DB tablespace containers for instance "%s" or the log path for one of its databases	<p><b>Cause:</b> LifeKeeper could not determine the location of the database table space containers or verify that they are located in a path which is on a mounted filesystem.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "create" operation.</p>
103033	ERROR	The path "%s" is not on a shared filesystem	<p><b>Cause:</b> The path of database table space container should be on a shared filesystem.</p> <p><b>Action:</b> Ensure database table space container is on a shared filesystem and retry the operation.</p>
103034	ERROR	A DB2 Hierarchy already exists for instance "%s"	<p><b>Cause:</b> An attempt was made to protect the DB2 instance that is already under LifeKeeper protection.</p> <p><b>Action:</b> You must select a different DB2 instance for LifeKeeper protection.</p>
103035	ERROR	The file system resource "%s" is not in-service	<p><b>Cause:</b> The file system which the DB2 resource depends on should be in service.</p> <p><b>Action:</b> Ensure the file system resource is in service and retry the "create" operation.</p>

Code	Severity	Message	Cause/Action
103036	ERROR	Unable to create the hierarchy for raw device "%s"	<p><b>Cause:</b> LifeKeeper was unable to create the resource {raw device} .</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the "create" operation.</p>
103037	ERROR	A RAW hierarchy does not exist for the tag "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the raw resource {tag} .</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103038	ERROR	LifeKeeper was unable to create a dependency between the DB2 hierarchy "%s" and the Raw hierarchy "%s"	<p><b>Cause:</b> The requested dependency creation between the parent DB2 resource and the child Raw resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the "create" operation.</p>
103039	ERROR	LifeKeeper could not disable the automatic startup feature of DB2 instance "%s"	<p><b>Cause:</b> An unexpected error occurred while attempting to update the DB2 setting.</p> <p><b>Action:</b> The DB2AUTOSTART will need to be updated manually to turn off the automatic startup of the instance at system boot.</p>
103040	ERROR	DB2 version "%s" is not installed on server "%s"	<p><b>Cause:</b> LifeKeeper could not find DB2 installed location.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103041	ERROR	The instance owner "%s" does not exist on target server "%s"	<p><b>Cause:</b> An attempt to retrieve the DB2 i</p>

Code	Severity	Message	Cause/Action
			<p>instance owner from template server during a "canextend" or "extend" operation failed.</p> <p><b>Action:</b> Verify the DB2 instance owner exists on the specified server. If the user does not exist, it should be created with the same uid and gid on all servers in the cluster.</p>
103042	ERROR	The instance owner "%s" uids are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The user id on the target server {target server} for the DB2 instance owner {user} does not match the value of the user {user} on the template server {template server}.</p> <p><b>Action:</b> The user ids for the DB2 instance owner {user} must match on all servers in the cluster. The user id mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103043	ERROR	The instance owner "%s" gids are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The group id on the target server {target server} for the DB2 instance owner {user} does not match the value of the user {user} on the template server {template server}.</p> <p><b>Action:</b> The group ids for the DB2 instance owner {user} must match on all servers in the cluster. The group id mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103044	ERROR	The instance owner "%s" home directories are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The home directory location of the user {user} on the target server {target server} does not match the DB2 instance owner's home directory on the temp</p>

Code	Severity	Message	Cause/Action
			<p>late server {template server}.</p> <p><b>Action:</b> The home directory location of the DB2 instance owner {user} must match on all servers in the cluster. The location mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103045	ERROR	LifeKeeper was unable to get the DB2 "SVCENAME" parameter for the DB2 instance	<p><b>Cause:</b> There was an unexpected error running "db2 get dbm cfg" command.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103046	ERROR	Unable to get the value of the DB2 "SVCENAME" parameter for the DB2 instance %s.	<p><b>Cause:</b> The DB2 "SVCENAME" parameter is set to null.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103047	ERROR	LifeKeeper was unable to get the contents of the "/etc/services" file on the server "%s"	<p><b>Cause:</b> "/etc/services" on the template server does not contain the service names for the DB2 instance.</p> <p><b>Action:</b> The service names in "/etc/services" for the DB2 instance must match on all servers in the cluster. The service names mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103048	ERROR	LifeKeeper was unable to get the contents of the "/etc/services" file on the server "%s"	<p><b>Cause:</b> "/etc/services" on the target server does not contain the service names for the DB2 instance.</p> <p><b>Action:</b> The service names in "/etc/services" for the DB2 instance must match on all servers in the cluster. The service names mismatch should be corrected manually on all servers before retrying th</p>

Code	Severity	Message	Cause/Action
			e "canextend" operation.
103049	ERROR	The "/etc/services" entries for the instance "%s" are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The "/etc/services" entries for the instance are mismatched.</p> <p><b>Action:</b> The service names in "/etc/services" for the DB2 instance must match on all servers in the cluster. The service names mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103050	ERROR	The home directory "%s" for instance "%s" is not mounted on server "%s"	<p><b>Cause:</b> LifeKeeper could not find db2nodes.cfg for Multiple Partition instance.</p> <p><b>Action:</b> Ensure the home directory is mounted and retry the operation.</p>
103051	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Failed to get resource information from the template server.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103052	ERROR	LifeKeeper was unable to add instance "%s" and/or its variables to the DB2 registry	<p><b>Cause:</b> There was an unexpected error running "db2iset" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103053	ERROR	Usage: %s instance	
103054	ERROR	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>

Code	Severity	Message	Cause/Action
103055	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103056	ERROR	Usage: %s instance	
103058	ERROR	Usage: %s instance	
103059	ERROR	Usage: %s instance	
103060	ERROR	Unable to determine the DB2 instance home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>
103061	ERROR	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103062	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find the node for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103063	ERROR	Unable to determine the DB2 install path	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find DB2 for the instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103064	ERROR	Usage: nodes -t tag -a add_nodenum   nodes -t tag -d delete_nodenum   nodes -t tag -p	

Code	Severity	Message	Cause/Action
103065	ERROR	Invalid input provided for "%s" utility operation, characters are not allowed.	<p><b>Cause:</b> Invalid parameters were specified for the "nodes" command.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103066	ERROR	Unable to get the information for resource "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag}.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103067	ERROR	The DB2 instance "%s" is not a EEE or Multiple Partition instance	<p><b>Cause:</b> The resource {tag} is single partition instance.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103069	ERROR	Node "%s" is already protected by this hierarchy	<p><b>Cause:</b> Invalid parameters were specified for the "nodes" command.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103070	ERROR	Node number "%s" is the last remaining node protected by resource "%s". Deleting all nodes is not allowed.	<p><b>Cause:</b> Invalid parameters were specified for the "nodes" command.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103071	ERROR	LifeKeeper is unable to get the equivalent instance for resource "%s"	<p><b>Cause:</b> There was an unexpected error running "nodes" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
103072	ERROR	Unable to set NodesInfo for resource "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "nodes" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103073	ERROR	Unable to set NodesInfo for resource "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "nodes" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103074	ERROR	Usage: %s instance	
103075	ERROR	Usage: %s instance	
103076	ERROR	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103077	ERROR	Unable to determine the DB2 instance home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>
103078	ERROR	The database server is not running for instance "%s"	<p><b>Cause:</b> A process check for the DB2 instance did not find any processes running.</p> <p><b>Action:</b> The DB2 instance must be started.</p>

Code	Severity	Message	Cause/Action
103079	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103080	ERROR	One or more of the database partition servers for instance "%s" is down	<p><b>Cause:</b> All database partition servers should be running.</p> <p><b>Action:</b> Ensure all database partition servers are running and retry the operation.</p>
103081	ERROR	DB2 local recovery detected another recovery process in progress for "%s" and will exit.	
103082	ERROR	Failed to create flag "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create a flag for controlling DB2 local recovery processing.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
103083	ERROR	Failed to remove flag "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to remove a flag for controlling DB2 local recovery processing.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
103084	ERROR	Unable to determine the DB2 instance \"\$Instance\" home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>

Code	Severity	Message	Cause/Action
104002	FATAL	\$msg	<p><b>Cause:</b> This message indicates an internal software error.</p> <p><b>Action:</b> The stack trace indicates the source of the error.</p>
104003	FATAL	\$self->Val('Tag') . " is not an SDR resource"	<p><b>Cause:</b> A data replication action was attempted on a non data replication resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
104010	ERROR	\$self->{'md'}: bitmap merge failed, \$action	<p><b>Cause:</b> The bitmap merge operation has failed.</p> <p><b>Action:</b> The target server may have the mirror and/or protected filesystem mounted, or the bitmap file may be missing on the target. Check the target server.</p>
104022	ERROR	\$argv <sup>1</sup> : mdadm failed (\$ret	<p><b>Cause:</b> The "mdadm" command has failed to add a device into the mirror.</p> <p><b>Action:</b> This is usually a temporary condition.</p>
104023	ERROR	\$_	<p><b>Cause:</b> The message contains the output of the "mdadm" command.</p>
104025	ERROR	failed to spawn monitor	<p><b>Cause:</b> The system failed to start the 'mdadm -F' monitor process. This should not happen under normal circumstances.</p> <p><b>Action:</b> Reboot the system to ensure that any potential conflicts are resolved.</p>

Code	Severity	Message	Cause/Action
104026	ERROR	cannot create \$md	<p><b>Cause:</b> The mirror device could not be created.</p> <p><b>Action:</b> Ensure the device is not already in use and that all other parameters for the mirror creation are correct.</p>
104027	ERROR	\$_	<p><b>Cause:</b> This message contains the "mdadm" command output.</p>
104035	ERROR	Too many failures. Aborting resync of \$md	<p><b>Cause:</b> The device was busy for an abnormally long period of time.</p> <p><b>Action:</b> Reboot the system to be sure that the device is no longer busy.</p>
104036	ERROR	Failed to start nbd-server on \$target (error \$port)	<p><b>Cause:</b> The nbd-server process could not be started on the target server.</p> <p><b>Action:</b> Ensure that the target disk device is available and that its Device ID has not changed.</p>
104037	ERROR	Failed to start compression (error \$port)	<p><b>Cause:</b> The system was unable to start the 'balance' tunnel process or there was a network problem.</p> <p><b>Action:</b> Ensure that the network is operating properly and that TCP ports in the range 10000-10512 are opened and unused. Ensure that the software is installed properly on all systems.</p>
104038	ERROR	Failed to start nbd-client on \$source (error \$ret)	<p><b>Cause:</b> The nbd-client process has failed to start on the source server.</p> <p><b>Action:</b> Look up the reported errno value and try to resolve the problem reported.</p>

Code	Severity	Message	Cause/Action
			d. For example, an errno value of 110 means "Connection timed out", which may indicate a network or firewall problem.
104039	ERROR	Failed to add \$nbd to \$md on \$source	<p><b>Cause:</b> This is usually a temporary condition.</p> <p><b>Action:</b> If this error persists, reboot the system to resolve any potential conflicts.</p>
104045	ERROR	failed to stop \$self->{'md'}	<p><b>Cause:</b> The mirror device could not be stopped.</p> <p><b>Action:</b> Ensure that the device is not busy or mounted. Try running "mdadm —stop" manually to stop the device.</p>
104048	WARN	failed to kill \$proc, pid \$pid	<p><b>Cause:</b> The process could not be signaled. This may indicate that the process has already died.</p> <p><b>Action:</b> Ensure that the process in question is no longer running. If it is, then reboot the system to clear up the unkillable process.</p>
104050	ERROR	Setting \$name on \$dest failed: \$ret. Please try again.	<p><b>Cause:</b> The system failed to set a 'mirrorinfo' file setting.</p> <p><b>Action:</b> Check the network and systems and retry the mirror setting operation.</p>
104052	FATAL	Specified existing mount point "%s" is not mounted	<p><b>Cause:</b> The mount point became unmounted.</p> <p><b>Action:</b> Ensure that the mount point is mounted and retry the operation.</p>

Code	Severity	Message	Cause/Action
104055	ERROR	Failed to set up temporary \$type access to data for \$self->{'tag'}. Error: \$ret	<p><b>Cause:</b> The filesystem or device was not available on the target server. The mirrored data will not be available on the target server until the mirror is paused and resumed again.</p> <p><b>Action:</b> Reboot the target server to resolve any potential conflicts.</p>
104057	ERROR	Failed to undo temporary access for \$self->{'tag'} on \$self->{'sys'}. Error: \$ret. Please verify that \$fsid is not mounted on server \$self->{'sys'}.	<p><b>Cause:</b> The filesystem could not be unmounted on the target server.</p> <p><b>Action:</b> Ensure that the filesystem and device are not busy on the target server. Reboot the target server to resolve any potential conflicts.</p>
104062	FATAL	Cannot find a device with unique ID "%s"	<p><b>Cause:</b> The target disk could not be identified.</p> <p><b>Action:</b> Ensure that the appropriate storage recovery kits are installed on the target server. Ensure that the Device ID of the target disk has not changed.</p>
104066	FATAL	Cannot get the hardware ID of device "%s"	<p><b>Cause:</b> A unique ID could not be found for the target disk device.</p> <p><b>Action:</b> Ensure that the appropriate storage recovery kits are installed on the target server. Ensure that the Device ID of the target disk has not changed.</p>
104067	FATAL	Asynchronous writes cannot be enabled without a bitmap file	<p><b>Cause:</b> An attempt was made to create a mirror with invalid parameters.</p> <p><b>Action:</b> A bitmap file parameter must be specified or synchronous writes must be specified.</p>

Code	Severity	Message	Cause/Action
104068	FATAL	Failed to extend dependent resource %s to system %s. Error %s	<p><b>Cause:</b> The hierarchy extend operation failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
104070	FATAL	Unable to extend the mirror "%s" to system "%s"	<p><b>Cause:</b> The hierarchy extend operation failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
104071	ERROR	Failed to restore target device resources on \$target->{'sys'} : \$err	<p><b>Cause:</b> The in-service operation has failed on the target server.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
104074	FATAL	Cannot get the hardware ID of device "%s"	<p><b>Cause:</b> There is no storage recovery kit that recognizes the underlying disk device that you are attempting to use for the mirror.</p> <p><b>Action:</b> Make sure the appropriate storage recovery kits are installed. If necessary, place your device name in the <code>/opt/LifeKeeper/subsys/scsi/resources/DEVNAME/device_pattern</code> file.</p>
104081	FATAL	Cannot make the %s filesystem on "%s" (%d)	<p><b>Cause:</b> The "mkfs" command failed.</p> <p><b>Action:</b> Ensure that the disk device is writable and free of errors and that the filesystem tools for the selected filesystem are installed.</p>

Code	Severity	Message	Cause/Action
104082	FATAL	%s	<b>Cause:</b> This message contains the output of the "mkfs" command.
104083	FATAL	Cannot create filesystem hierarchy "%s"	<b>Cause:</b> The resource creation failed.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104086	ERROR	The "%s_data_corrupt" flag is set in "%s/subsys/scsi/resources/netraid/" on system "%s". To avoid data corruption, LifeKeeper will not restore the resource.	<b>Cause:</b> The data corrupt flag file has been set as a precaution to prevent accidental data corruption. The mirror cannot be restored on this server until the file is removed.  <b>Action:</b> If you are sure that the data is valid on the server in question, you can either: 1) remove the file and restore the mirror, or 2) force the mirror online using the LifeKeeper GUI or 'mirror_action force' command.
104092	ERROR	Mirror target resource movement to system %s : status %s	<b>Cause:</b> The hierarchy switchover operation has failed.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104099	ERROR	Unable to unextend the mirror for resource "%s" from system "%s"	<b>Cause:</b> The hierarchy unextend operation failed.  <b>Action:</b> Reboot the target server to resolve any potential conflicts and retry the operation.
104106	ERROR	remote 'bitmap -m' command failed on \$target->{'sys'}: \$ranges	<b>Cause:</b> The bitmap merge command failed on the target server. This may be ca

Code	Severity	Message	Cause/Action
			<p>used by one of two things: 1) The bitmap file may be missing or corrupted, or 2) the mirror (md) device may be active on the target.</p> <p><b>Action:</b> Make sure that the mirror and protected filesystem are not active on the target. If the target's bitmap file is missing, pause and resume the mirror to recreate the bitmap file.</p>
104107	ERROR	Asynchronous writes cannot be enabled without a bitmap file	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104108	ERROR	Local Partition not available	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104109	ERROR	Cannot get the hardware ID of device "%s"	<p><b>Cause:</b> A unique ID could not be determined for the disk device.</p> <p><b>Action:</b> Ensure that the appropriate storage recovery kits are installed on the server. Ensure that the Device ID of the disk has not changed.</p>
104111	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104112	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104113	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104114	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified

Code	Severity	Message	Cause/Action
			ed for the mirror create operation.
104115	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104117	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104118	FATAL	Cannot unmount existing Mount Point "%s"	<b>Cause:</b> The mount point is busy.  <b>Action:</b> Make sure the filesystem is not busy. Stop any processes or applications that may be accessing the filesystem.
104119	FATAL	Invalid data replication resource type requested ("%s")	<b>Cause:</b> An invalid parameter was specified for the mirror create operation.
104124	EMERG	WARNING: A temporary communication failure has occurred between systems %s and %s. In order to avoid data corruption, data resynchronization will not occur. MANUAL INTERVENTION IS REQUIRED. In order to initiate data resynchronization, you should take one of the following resources out of service: %s on %s or %s on %s. The resource that is taken out of service will become the mirror target.	<b>Cause:</b> A temporary communication failure (split-brain scenario) has occurred between the source and target servers.  <b>Action:</b> Perform the steps listed in the message text.
104125	ERROR	failed to start '\$cmd \$_ <sup>2</sup> \$user_args' on '\$_ <sup>3</sup> '	<b>Cause:</b> The specified command failed.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104126	ERROR	\$_	<b>Cause:</b> This message contains the output of the command that was reported as failing in message 104125.

Code	Severity	Message	Cause/Action
104128	FATAL	comm path/server not specified	<p><b>Cause:</b> The netraid.down script was called without specifying the communication path or the server name. This script is called internally so should always have the proper parameters.</p> <p><b>Action:</b> Report this error to SIOS support.</p>
104129	WARN		<p><b>Cause:</b> The replication connection for the mirror is down.</p> <p><b>Action:</b> Check the network.</p>
104130	ERROR	Mirror resize failed on %s (%s). You must successfully complete this operation before using the mirror. Please try again.	<p><b>Cause:</b> The mirror resize operation has failed to update the mirror metadata on the listed system.</p> <p><b>Action:</b> You must successfully complete the resize before using the mirror. Re-run mirror_resize (possibly using -f to force the operation if necessary).</p>
104132	ERROR	The partition "%s" has an odd number of sectors and system "%s" is running kernel >= 4.12. Mirrors with this configuration will not work correctly with DataKeeper. Please see the SIOS product documentation for information on how to resize the mirror.	<p><b>Cause:</b> The partition or disk chosen for mirror creation has an odd number of disk sectors and will have to be resized to be used with DataKeeper.</p> <p><b>Action:</b> Resize the partition using the 'parted' command or resize the disk (if possible) using platform (VMware, AWS) tools. Caution: data may be lost if this is not done carefully.</p>
104136	ERROR	Extend failed.	<p><b>Cause:</b> The hierarchy extend operation failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
			m.
104143	ERROR	Mirror resume was unsuccessful (\$ret	<p><b>Cause:</b> The mirror could not be established.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problems.</p>
104144	ERROR	Unable to stop the mirror access for \$self->{'md'} on system \$self->{'sys'}. Error: \$ret. Use the \"mdadm —stop \$self->{'md'}\" command to manually stop the mirror.	<p><b>Cause:</b> The mirror device created on the target node when the mirror was paused could not be stopped.</p> <p><b>Action:</b> Ensure that the device is not busy or mounted. Try running \"mdadm —stop\" manually to stop the device.</p>
104145	WARN	Unable to dirty full bitmap. Setting fullsync flag.	<p><b>Cause:</b> A full resync could not be done by dirtying the full bitmap. The fullsync flag will be used instead. This is a non-fatal error as a full synchronization will still be done.</p> <p><b>Action:</b> None</p>
104146	EMERG	WARNING: The target system %s for mirror %s has the target mirror %s currently active. In order to avoid data corruption, data resynchronization will not occur. MANUAL INTERVENTION IS REQUIRED. In order to initiate data resynchronization, you should reboot system %s.	<p><b>Cause:</b> The mirror is configured on the target system.</p> <p><b>Action:</b> The target system should be rebooted. DataKeeper should then be able to resync the mirror.</p>
104147	EMERG	WARNING: The target system %s for mirror %s has the target disk %s currently mounted. In order to avoid data corruption, data resynchronization will not occur. MANUAL INTERVENTION IS REQUIRED. In order to initiate data resynchronization, you should unmount %s on %s. A full resync will occur.	<p><b>Cause:</b> The mirror disk is mounted on the target system.</p> <p><b>Action:</b> The mirror disk should be unmounted on the target system, in order to initiate a full mirror resync. A full resync</p>

Code	Severity	Message	Cause/Action
			<p>is required because untracked changes have occurred on the disk.</p>
104148	EMERG	<p>The storage configuration for mirror "%s (%s)" does not have a unique identifier and may have potential for data corruption in some environments in certain circumstances. Please refer to the SIOS Product Documentation for details on DataKeeper storage configuration options.</p>	<p><b>Cause:</b> The disk chosen for mirroring does not provide a UUID to the operating system. DataKeeper cannot mirror a disk without a UUID.</p> <p><b>Action:</b> You may be able to create a GPT partition table on the disk to provide a UUID for the disk partitions.</p>
104155	EMERG	<p>The mirror %s cannot be forced online at this time. The underlying disk %s is mounted, indicating possible data corruption. MANUAL INTERVENTION IS REQUIRED. You must unmount %s on %s and restore the mirror to the last known mirror source system. A full resync will need to be performed from the source system to %s.</p>	<p><b>Cause:</b> You have mounted the underlying mirrored disk on the target system.</p> <p><b>Action:</b> You must unmount the disk immediately in order to avoid data corruption.</p>
104156	WARN	<p>Resynchronization of "%s" is in PENDING state. Current sync_action is: "%s"</p>	<p><b>Cause:</b> The resynchronization of the md device is detected in PENDING state.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by forcing a resynchronization. Check the logs for related errors. When successful assure that the PENDING state has been cleared in /proc/mdstat and the resynchronization is in progress or has been completed for the datarep resource.</p>
104157	WARN	<p>/etc/sysconfig/raid-check update failed. Please %s \"md%d\" to SKIP_DEVS.</p>	<p><b>Cause:</b> Unable to make changes in /etc/sysconfig/raid-check to add or remove an entry to the list of MD devices to skip (SKIP_DEVS).</p> <p><b>Action:</b> Check system logs for any errors related to raid-check or SKIP_DEVS. Manually add or remove md listed.</p>

Code	Severity	Message	Cause/Action
104158	EMERG	WARNING: The local disk partition \$self->{'part'} for data replication device\n\$self->{'md'} has failed. MANUAL INTERVENTION IS REQUIRED.	<p><b>Cause:</b> The local device for a mirror failed. The recovery action has been set to "nothing" in LKDR_FAILURE requiring manual intervention to recover.</p> <p><b>Action:</b> Check system logs and LifeKeeper logs for errors related to the local disk.</p>
104163	WARN	The "%s_data_corrupt" flag is set in "%s/subsys/scsi/resources/netraid/" on system "%s". The mirror is being forced online.	<p><b>Cause:</b> The mirror is being forced online, overriding the data_corrupt flag. The data on the specified system will be treated as the latest data. If this is not correct then this can lead to data corruption or data loss.</p> <p><b>Action:</b> None</p>
104164	ERROR	The "%s_data_corrupt" flag for related mirror resource "%s" is set in "%s/subsys/scsi/resources/netraid/" on system "%s". To avoid data corruption, LifeKeeper will not restore this mirror or any related mirrors in the hierarchy.	<p><b>Cause:</b> The data_corrupt flag exists for one or more mirrors in the hierarchy. To avoid corrupting additional data none of the mirrors are brought in-service until all of the data_corrupt flags are resolved.</p> <p><b>Action:</b> Check the LifeKeeper logs to determine where each mirror was last in-service, aka where the latest data for each mirror resides. The mirrors should be brought in-service on the "previous source" where the full hierarchy was in-service and allow the mirrors to synchronize with all targets.</p>
104165	ERROR	The "%s_data_corrupt" flag for related mirror resource "%s" is set in "%s/subsys/scsi/resources/netraid/" on system "%s". The mirror resource "%s" is being forced online.	<p><b>Cause:</b> The mirror is being forced online, overriding the data_corrupt flag. The data on the specified system will be treated as the correct data to be synchronized with all targets. This can lead to data corruption or data loss if this is not the latest data.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> None
104170	ERROR	Failed to create \"source\" flag file on shared source %s to track mirror source. This may result in a full resync.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104171	ERROR	Failed to create \"source\" flag file on %s to track mirror source. Target %s will not be added to mirror.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104172	ERROR	The \"source\" flag file on %s does not contain a valid target (%s). Full resync to remaining targets is required.	<p><b>Cause:</b> The 'source' flag file should contain the system name of a previous source but the name listed was not found in the list of systems configured.</p> <p><b>Action:</b> Report this problem to SIOS support.</p>
104173	ERROR	Failed to create \"source\" flag file on %s to track mirror source.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104174	ERROR	Failed to create \"previous_source\" flag file to track time waiting on source. Will not be able to timeout.	<p><b>Cause:</b> The 'previous_source' flag file was not created on the mirror source to track the mirror's previous source.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.
104175	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.  <b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.
104176	ERROR	The \"source\" flag file on %s to track mirror or source does not exist. Full resync is required.	<b>Cause:</b> The 'source' flag file should exist on the system and without it the consistency of the mirror can not be verified. A full resync is required to assure data reliability.  <b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for error or that it is full.
104177	ERROR	Failed to determine amount of time waiting on %s.	<b>Cause:</b> The amount of time waiting for the previous source could not be determined. Targets will be added with a full resync if the previous source is not found.  <b>Action:</b> none
104178	ERROR	Failed to update "source" flag file on target "%s", previous source must be merged first.	<b>Cause:</b> The source flag file on the target is updated when it is in-sync and stopped so the the next in-service does not require the previous source.  <b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.
104180	ERROR	Internal Error: \"previous_source\" has the local system name (%s).	<b>Cause:</b> The local system name should

Code	Severity	Message	Cause/Action
			<p>not be in the previous_source flag file.</p> <p><b>Action:</b> Report this error to SIOS support.</p>
104181	ERROR	Internal Error: There are no targets waiting on %s to be merged.	<p><b>Cause:</b> There are no targets waiting for a previous source to merge.</p> <p><b>Action:</b> Report this error to SIOS support.</p>
104182	ERROR	Failed to create \"source\" flag file on %s to track mirror source. This may result in a full resync.	<p><b>Cause:</b> The 'source' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104186	ERROR	Failed to create \"last_owner\" flag file on %s to track mirror source. This may allow in-service of mirror on old data.	<p><b>Cause:</b> The 'last_owner' flag file was not created on the source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104187	WARN	\$REM_MACH has \${REM_TAG}_last_owner file, create flag \${FLAGTAG}_data_corrupt.	<p><b>Cause:</b> The system listed had the mirror in-service last.</p> <p><b>Action:</b> The system listed has the last_owner file that indicates it has the most recent data and is most likely the best system to in-service the mirror to avoid losing data.</p>
104188	WARN	\$REM_MACH is not alive, create flag \${FLAGTAG}_data_corrupt.	<p><b>Cause:</b> The system listed is not alive.</p> <p><b>Action:</b> Since the system listed is not alive, it cannot be determined whether th</p>

Code	Severity	Message	Cause/Action
			<p>at system was a more recent mirror source than the local system. Therefore the local system should not automatically be allowed to bring the mirror in-service.</p>
104200	EMERG	<p>Continue to wait for %s to merge bitmap and do partial resyncs to all targets, no timeout set.</p>	<p><b>Cause:</b> In a multi-target configuration targets will not be configured until the previous source is available to merge its bitmap so that all targets will be able to partially resynchronize. The LKDR_WAIT_ON_PREVIOUS_SOURCE_TIMEOUT entry in /etc/defaults/LifeKeeper is set to "-1" to wait indefinitely.</p> <p><b>Action:</b> Check on the status of the previous source listed in the message and resolve any issues that are preventing it from rejoining the cluster.</p>
104201	EMERG	<p>To stop waiting for the previous source (forcing a full resync to remaining waiting targets) run: "\"%s/bin/mirror_action %s fullresync %s %s\" on %s.</p>	<p><b>Cause:</b> In a multi-target configuration targets are not being configured, waiting on the previous source to rejoin the cluster.</p> <p><b>Action:</b> Run the command listed in the message to force an immediate full resynchronization to this target and any remaining targets waiting to be resynchronized.</p>
104202	EMERG	<p>Continue to wait for %s to merge bitmap and do partial resyncs to all targets. Continue to wait %d more seconds.</p>	<p><b>Cause:</b> In a multi-target configuration targets will not be configured until the previous source is available to merge its bitmap so that all targets will be able to partially resynchronize. The LKDR_WAIT_ON_PREVIOUS_SOURCE_TIMEOUT entry in /etc/defaults/LifeKeeper is set to the number of seconds to wait. If the previous source does not join the cluster in that time then targets will be added</p>

Code	Severity	Message	Cause/Action
			<p>with a full resynchronization.</p> <p><b>Action:</b> Check on the status of the previous source listed in the message and resolve any issues that are preventing it from rejoining the cluster.</p>
104203	EMERG	To stop waiting for the previous source (forcing a full resync to remaining waiting targets) run: \" <code>%s/bin/mirror_action %s fullresync %s %s\</code> \" on %s.	<p><b>Cause:</b> In a multi-target configuration targets are not being configured, waiting on the previous source to rejoin the cluster.</p> <p><b>Action:</b> Run the command listed in the message to force an immediate full resynchronization to this target and any remaining targets waiting to be resynchronized.</p> <p>Note: Run this command to stop waiting even if the target listed is deleted and never returning.</p>
104207	ERROR	Failed to create "data_corrupt" flag file on "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the source listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104208	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104209	ERROR	Failed to create "data_corrupt" flag file on "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the source listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeK</p>

Code	Severity	Message	Cause/Action
			eeper) file system for errors or that it is full.
104210	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104211	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104212	ERROR	The "\"source\" flag file on %s to track mirror source does not exist. Full resync to remaining targets is required.	<p><b>Cause:</b> The "source" flag file should exist on the system and without it the consistency of the mirror can not be verified. A full resync is required to assure data reliability. All targets not already being mirrored will require a full resync.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for error or that it is full.</p>
104214	ERROR	Failed to create "\"source\" flag file on %s to track mirror source.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104216	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104217	ERROR	Failed to create \"source\" flag file on shared source %s to track mirror source. This may result in a full resync.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104218	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full</p>
104221	ERROR	Failed to create \"last_owner\" flag file on %s to track mirror source. This may allow in-service of mirror on old data.	<p><b>Cause:</b> The 'last_owner' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104222	ERROR	Failed to create "last_owner" flag file to track mirror source". This may allow in-service of mirror on old data.	<p><b>Cause:</b> The 'last_owner' flag file is used to know where the mirror was last in-service.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104223	ERROR	Failed to create "last_owner" flag file to track mirror source". This may allow in-service of mirror on old data.	<p><b>Cause:</b> The 'last_owner' flag file is used to know where the mirror was last in-service.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104224	ERROR	Failed to create \"previous_source\" flag file.	<p><b>Cause:</b> The 'previous_source' flag file was not created. This is needed to merge the previous source bitmap to avoid a full resync.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104227	ERROR	Failed to set %s to %s.	<p><b>Cause:</b> This message indicates a failure to set a sysfs parameter for the nbd driver (/sys/block/nbdX).</p> <p><b>Action:</b> It may be necessary to adjust one or more of:</p> <p>NBD_NR_REQUESTS                      NBD_SCHEDULER                      LKDR_ASYNC_LIMIT</p> <p>in /etc/default/LifeKeeper to avoid this error.</p>
104232	ERROR	Mirror resize failed on %s (%s). Could not set size to %d.	<p><b>Cause:</b> The mirror resize operation failed.</p>
104233	ERROR	Mirror resize failed on %s (%s). Could not set bitmap to %s and bitmap-chunk to %d.	<p><b>Cause:</b> The mirror resize operation failed.</p>
104234	ERROR	The mirror %s failed to resize. You must successfully complete this operation before using the mirror. Please try again.	<p><b>Cause:</b> The mirror resize operation failed.</p>
104235	ERROR	mirror_resize of mirror %s failed due to signal \"%s\".	<p><b>Cause:</b> The mirror resize operation failed.</p>

Code	Severity	Message	Cause/Action
			d.
104236	EMERG	Resource "%s" is "OSF". The mirror "%s" will wait to replicate data until all resources in the hierarchy%s are in-service. This may indicate inconsistent data. Verify the data is correct before replicating data; replication will continue when all resources in the hierarchy%s are in-service. A full resync may be necessary (see "LKDR_WAIT_TO_RESYNC" in /etc/default/LifeKeeper).	<p><b>Cause:</b> The specified resource is OSF. This prevents replication until this resource is repaired. The LKDR_WAIT_TO_RESYNC setting in /etc/default/LifeKeeper determines what resources must be in-service before replication is allowed.</p> <p><b>Action:</b> Determine the cause of the corruption and repair. This may involve bringing the resource in-service on another node. A full resync is most likely required once the data is repaired.</p>
104237	WARN	Resource "%s" is "OSU". The mirror "%s" will wait to replicate data until all resources in the hierarchy%s are in-service. To replicate data immediately run: "%s/bin/mirror_action %s resume" on "%s" (see "LKDR_WAIT_TO_RESYNC" in /etc/default/LifeKeeper).	<p><b>Cause:</b> The specified resource is not in-service and is required before replication is resumed during a restore operation.</p> <p><b>Action:</b> Bring the required resources in-service to resume replication. Replication can also resume using the GUI command to resume or using the mirror_action command.</p>
104238	ERROR	Unable to read \$nbd_taint_file. Assuming that the SIOS 'nbd' kernel module is not loaded.	<p><b>Cause:</b> The /sys/module/nbd/taint file could not be opened for reading.</p> <p><b>Action:</b> Verify that the /sys/module/nbd/taint file exists and is read-enabled.</p>
104239	EMERG	\$failure_msg	<p><b>Cause:</b> At least one kernel module required by SIOS DataKeeper failed to load.</p> <p><b>Action:</b> Verify that the current running kernel is supported by this version of SIOS Protection Suite for Linux. If the kernel was recently updated, re-run the SIOS Protection Suite for Linux setup script lo</p>

Code	Severity	Message	Cause/Action
			cated on the SIOS installation media to install kernel modules that are compatible with the current running kernel.
104242	ERROR	Unable to read /proc/modules. Assuming that the SIOS '\$module' kernel module is not loaded.	<p><b>Cause:</b> Unable to read /proc/modules.</p> <p><b>Action:</b> Verify that the /proc/modules file exists and is read-enabled.</p>
104243	ERROR	Internal script or routine \$caller was called for unsupported kernel module '\$module'. Supported kernel modules for this script or routine are: \$module_list.	<p><b>Cause:</b> The given script or routine was called for an unsupported kernel module.</p> <p><b>Action:</b> This is an internal error. Please contact SIOS Customer Support.</p>
104244	ERROR	Output from 'modprobe \$module' command: \$pretty_modprobe_out	<p><b>Cause:</b> Failed to load the given kernel module using the modprobe command.</p> <p><b>Action:</b> Inspect the output provided in the log message from the failed modprobe attempt and resolve any issues found there.</p>
104251	ERROR	There is no LifeKeeper protected resource with tag \$tag on system \$me.	<p><b>Cause:</b> The given tag does not correspond to a LifeKeeper protected resource on the given system.</p> <p><b>Action:</b> Verify that the resource tag and system name are correct.</p>
104252	ERROR	Resource \$tag is not a \$app/\$typ resource. Please use the \$ins_app/\$ins_typ resource-specific canfailover script instead.	<p><b>Cause:</b> The scsi/netraid-specific canfailover script was called for a non-scsi/netraid resource.</p> <p><b>Action:</b> Use the canfailover script, if it exists, corresponding to the appropriate app and type of the given resource.</p>

Code	Severity	Message	Cause/Action
107015	ERROR	Exported file system \$opt_t cannot be accessed on \$me.	
112000	FATAL	Usage: "%s %s". Specify the correct usage for the requested command.	<p><b>Cause:</b> Invalid parameters passed to the SAP create script.</p> <p><b>Action:</b> Please provide appropriate parameters for the SAP create script.</p>
112004	ERROR	Neither the SID/instance pair nor the tag parameter were specified for the internal "%s" routine on %s. If this was a command line operation, specify the correct parameters. Otherwise, consult the troubleshooting documentation.	<p><b>Cause:</b> Invalid parameters were specified when trying to create an internal SAP object.</p> <p><b>Action:</b> Provide either the SID and instance or the LifeKeeper resource tag name.</p>
112013	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and/or is readable.	<p><b>Cause:</b> The sapservices file either does not exist or is not readable.</p> <p><b>Action:</b> Verify that the sapservices file exists and is readable.</p>
112017	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112019	ERROR	The attempt to update the resource information field for resource %s has failed on %s. View the resource properties manually using "ins_list -t <tag>" to verify that the resource is functional.	<p><b>Cause:</b> Unable to update the LifeKeeper resource information field for the given resource on the given server.</p> <p><b>Action:</b> Verify that the resource exists and that LifeKeeper is running on the given server.</p>

Code	Severity	Message	Cause/Action
112022	ERROR	An error occurred while trying to find the IP address corresponding to "%s" on %s. Verify the IP address or host name exists in DNS or the hosts file.	<p><b>Cause:</b> Unable to find the given IP address or host name on the given server.</p> <p><b>Action:</b> Verify that the IP address or DNS name exists in DNS or in the local hosts file.</p>
112024	FATAL	There was an error verifying the NFS connections for SAP related mount points on %me. One or more NFS servers is not operational and needs to be restarted.	<p><b>Cause:</b> At least one NFS shared file system listed in the SAP_NFS_CHECK_DIRS parameter is currently unavailable.</p> <p><b>Action:</b> Verify that the NFS server is alive, all necessary NFS-related services are running, and that all necessary file systems are being exported.</p>
112028	ERROR	Unable to determine the user name for the SAP administrative user for resource %s on %s.	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112029	ERROR	The canextend script "%s" either does not exist or is not executable on %s.	<p><b>Cause:</b> The SAP canextend script either does not exist or is not executable.</p> <p><b>Action:</b> Verify that the SAP canextend script exists and is executable.</p>
112037	ERROR	Unable to create an internal object for the SAP instance using SID %s, instance %s, and tag %s on server %s. Verify that all necessary SAP file systems are mounted and accessible before reattempting the operation.	<p><b>Cause:</b> Unable to create an internal SAP object to represent the given SAP instance.</p> <p><b>Action:</b> Verify that the SAP instance is properly installed and configured and that all necessary file systems are mounted.</p>

Code	Severity	Message	Cause/Action
112040	ERROR	The SAP Directory "%s" ("%s") does not exist on %s. Verify that the directory exists and that the SAP software is properly installed.	<p><b>Cause:</b> The given SAP installation directory does not exist on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted.</p>
112041	ERROR	The required utility "%s" was not found or was not executable on %s. Verify the SAP installation and location of the required utility.	<p><b>Cause:</b> The saphostexec or saposcol utility either could not be found or is not executable.</p> <p><b>Action:</b> Verify that the SAP Host Agent package is installed correctly and that all necessary file systems are mounted.</p>
112042	ERROR	One or more SAP or LifeKeeper validation checks has failed on %s. Please update the SAP software on this host to include the SAPHOST and SAPCONTROL packages.	<p><b>Cause:</b> The saphostexec or saposcol utility either could not be found or is not executable.</p> <p><b>Action:</b> Verify that the SAP Host Agent package is installed correctly and that all necessary file systems are mounted.</p>
112048	ERROR	The SAP instance %s is already under LifeKeeper protection on server %s. Choose another SAP instance to protect or specify the correct instance.	<p><b>Cause:</b> The given SAP instance is already protected by LifeKeeper on the given server.</p> <p><b>Action:</b> Choose an SAP instance which is not already under LifeKeeper protection.</p>
112049	ERROR	Unable to locate the SAP Mount directory on %s. Verify that all SAP file systems are mounted and accessible before reattempting the operation.	<p><b>Cause:</b> Unable to determine the location of the SAP Mount (sapmnt) directory.</p> <p><b>Action:</b> Verify that all necessary file systems are mounted and that all necessary SAP instance profiles are accessible.</p>

Code	Severity	Message	Cause/Action
112050	ERROR	Detected multiple virtual IP addresses/hostname for instance %s on %s. Verify that the instance is configured correctly.	<p><b>Cause:</b> Multiple virtual IPs or hostnames were detected for the given SAP instance on the given server.</p> <p><b>Action:</b> Verify that the virtual IP or hostname associated to the instance is configured correctly.</p>
112051	ERROR	The "%s" or "%s" value in the default profile "%s" is still set to the physical hostname on %s. The value(s) must be set to a virtual hostname.	<p><b>Cause:</b> The given host name parameter is set to a physical server host name on the given server.</p> <p><b>Action:</b> Set the given host name parameter to a virtual host name.</p>
112053	ERROR	Detected multiple instances under SID %s with the same instance number (%s) on %s. Each instance within a particular SID must have a unique instance number.	<p><b>Cause:</b> Multiple SAP instances with the same instance number were detected under the same SAP SID.</p> <p><b>Action:</b> Reconfigure the SAP environment so that each instance under a given SAP SID has a unique instance number.</p>
112056	ERROR	The NFS export for the path "%s" required by the instance %s for the "%s" directory does not have an NFS hierarchy protecting it on %s. You must create an NFS hierarchy to protect this NFS export before creating the SAP resource hierarchy.	<p><b>Cause:</b> The NFS export for the given file system is not currently protected by LifeKeeper.</p> <p><b>Action:</b> Create a LifeKeeper NFS hierarchy for the given exported file system and reattempt SAP resource creation.</p>
112057	ERROR	Unable to create a file system resource hierarchy for the file system "%s" on %s.	<p><b>Cause:</b> Unable to create a LifeKeeper file system resource hierarchy to protect the given file system on the given server.</p> <p><b>Action:</b> Check the LifeKeeper and system logs for more information.</p>

Code	Severity	Message	Cause/Action
112058	ERROR	Unable to create a dependency between parent tag "%s" and child tag "%s" on "%s".	<p><b>Cause:</b> Unable to create a LifeKeeper dependency between the given resources on the given server.</p> <p><b>Action:</b> Check the LifeKeeper logs for more information.</p>
112060	ERROR	All attempts at local recovery for the SAP resource %s have failed on %s. A failover to the backup server will be attempted.	<p><b>Cause:</b> Unable to recover the given SAP resource on the given server.</p> <p><b>Action:</b> A failover of the SAP resource hierarchy will be attempted automatically. No user intervention is required.</p>
112061	ERROR	The values specified for the target and the template servers are the same. Please specify the correct values for the target and template servers.	<p><b>Cause:</b> The template and target servers provided during SAP resource extension are the same.</p> <p><b>Action:</b> Provide the correct names for the template and target servers and reattempt the extend operation.</p>
112062	ERROR	Unable to find the home directory "%s" for the SAP administrative user "%s" on %s. Verify that the SAP software is installed correctly.	<p><b>Cause:</b> Unable to find the home directory for the given SAP user on the given server.</p> <p><b>Action:</b> Verify that the SAP software is installed correctly and that the appropriate SAP administrative user for the given SID exists on the server.</p>
112063	ERROR	The SAP administrative user "%s" does not exist on %s. Verify that the SAP software is installed correctly or create the required SAP user on %s.	<p><b>Cause:</b> The given SAP administrative user does not exist on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and create the required SAP administrative user if necessary.</p>

Code	Severity	Message	Cause/Action
112064	ERROR	The group ID for user "%s" is not the same on template server "%s" and target server "%s". Please correct the group ID for the user so that it is the same on the template and target servers.	<p><b>Cause:</b> The group ID's for the given SAP administrative user on the template and target servers do not match.</p> <p><b>Action:</b> Modify the group ID of the given SAP administrative user so that it is the same on the template and target servers.</p>
112065	ERROR	The user ID for user "%s" is not the same on template server "%s" and target server "%s". Please correct the user ID so that it is the same on the template and target servers.	<p><b>Cause:</b> The user ID's for the given SAP administrative user on the template and target servers do not match.</p> <p><b>Action:</b> Modify the user ID of the given SAP administrative user so that it is the same on the template and target servers.</p>
112066	ERROR	Required SAP utilities could not be found in "%s" on %s. Verify that the SAP software is installed correctly.	<p><b>Cause:</b> Unable to locate necessary SAP executables or the SAP instance profile.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and configured and that all necessary file systems are mounted.</p>
112069	ERROR	The command "%s" is not found in the "%s" perl module ("%s") on %s. Please check the command specified and retry the operation.	<p><b>Cause:</b> The given command was not found in the sap perl module on the given server.</p> <p><b>Action:</b> If this error resulted from a user-initiated command line action, verify that the correct routine name was provided to the remoteControl script. If this error occurred during normal LifeKeeper operation, please submit an issue report to SIOS customer support.</p>

Code	Severity	Message	Cause/Action
112071	ERROR	The file "%s" exists, but was not read and write enabled on server %s. Enable read and write permissions on the specified file.	<p><b>Cause:</b> The given file does not have read/write permissions enabled on the given server.</p> <p><b>Action:</b> Enable read/write permissions on the given file.</p>
112073	ERROR	Unable to create an internal object for the SAP instance using either SID "%s" and instance "%s" or tag "%s" on server "%s". Either the values specified for the object initialization (SID/instance pair or tag, system) were not valid, or an error occurred while attempting to gather information about the SAP instance. If all specified parameters are correct, verify that all necessary SAP file systems are mounted and accessible before reattempting the operation.	<p><b>Cause:</b> Unable to create an internal SAP object to represent the given SAP instance.</p> <p><b>Action:</b> Verify that the SAP instance is properly installed and configured and that all necessary file systems are mounted.</p>
112074	WARN	WARNING: The profile "%s" for SID %s and instance %s has Autostart enabled on %s. Disable Autostart for the specified instance by setting Autostart=0 in the profile.	<p><b>Cause:</b> The Autostart parameter is enabled in the given instance profile on the given server.</p> <p><b>Action:</b> Disable Autostart for the given SAP instance by setting 'Autostart = 0' in the instance profile.</p>
112076	FATAL	Unable to start the sapstartsrv service for SID \$sid and instance \$Instance on \$me. Verify that the sapservices file is correct and that the process can be started manually.	<p><b>Cause:</b> Unable to start the SAP Start Service (sapstartsrv) process for the given SAP instance.</p> <p><b>Action:</b> Verify that the sapservices file contains the appropriate command to start the sapstartsrv process and that the process can be started manually.</p>
112077	ERROR	Unable to stop the sapstartsrv service for SAP SID %s and SAP instance %s on %s. Verify that the sapservices file is correct and the process can be stopped manually.	<p><b>Cause:</b> Unable to stop the SAP Start Service (sapstartsrv) process for the given SAP instance.</p> <p><b>Action:</b> Verify that the sapservices file</p>

Code	Severity	Message	Cause/Action
			contains the appropriate command to start the sapstartsrv process and that the process can be stopped manually.
112078	ERROR	ERSv1 is only supported in two-node clusters. Resource %s is unable to be extended to system %s. Upgrade to ERSv2 in order to extend the hierarchy to three or more nodes.	<p><b>Cause:</b> Unable to extend an SAP resource representing an ERSv1 instance to three or more nodes.</p> <p><b>Action:</b> In order to extend an SAP resource representing an ERS instance to three or more nodes, upgrade to ERSv2. Upgrade instructions are provided in the online product documentation.</p>
112082	WARN	Instance %s is running a different version of the enqueue server than its corresponding enqueue replication server. This configuration is not supported by SAP and will lead to unexpected resource behavior. See SAP Note 2711036 – "Usage of the Standalone Enqueue Server 2 in an HA Environment" for more details. Please review the online product documentation for instructions on how to modify the instance profiles for the enqueue server and enqueue replication server so that they use the same version.	<p><b>Cause:</b> The versions of the enqueue server and enqueue replication server do not match.</p> <p><b>Action:</b> Consult the online product documentation for instructions on how to modify the instance profiles so that the enqueue server and enqueue replication server are using the same version.</p>
112086	ERROR	The ERS resource corresponding to resource %s is in-service and maintaining backup locks on a remote system. Bringing resource %s in-service on %s would result in a loss of the backup lock table. Please bring resource %s in-service on the system where the corresponding ERS resource is currently in-service in order to maintain consistency of the lock table. In order to force resource %s in-service on %s, either (i) run the command <code>\'/opt/LifeKeeper/bin/flg_create -f sap_cs_force_restore_%s'</code> as root on %s and reattempt the in-service operation or (ii) take the corresponding ERS resource out of service on the remote system. Both of these actions will result in	<p><b>Cause:</b> An in-service operation was attempted for an ASCS/SCS resource while its corresponding ERS instance was running and storing a backup enqueue table on a different server in the cluster.</p> <p><b>Action:</b> Bring the ASCS/SCS resource in-service on the server where its corresponding ERS instance is running in order for it to retrieve the backup enqueue table. If the ASCS/SCS resource must be forced in-service on the given node, either (i) run the command <code>'/opt/LifeKeeper/bin/flg_create -f sap_cs_force_restore</code></p>

Code	Severity	Message	Cause/Action
		a loss of the backup lock table.	e_<ASCS/SCS Tag>' and reattempt the in-service operation, or (ii) take the corresponding ERS resource out of service on the remote server. Both of these actions will result in a loss of the backup enqueue table.
112089	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, or TYPE_NEW_ERS (1, 2, or 5).	<p><b>Cause:</b> The GetEnqVersion routine was called for an unsupported SAP instance type. Only instance types 1 (TYPE_CS), 2 (TYPE_ERS), and 5 (TYPE_NEW_ERS) are supported.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112092	ERROR	The profile "%s" either does not exist or cannot be read on %s. Unable to determine whether enqueue replication is enabled for resource %s. Please verify that the file exists and can be read.	<p><b>Cause:</b> The given instance profile either does not exist or cannot be read on the given server.</p> <p><b>Action:</b> Verify that the file exists and is read-enabled</p>
112095	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The given resource could not be created on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted and accessible on the given server.</p>
112096	ERROR	Resource %s is not currently in-service on server %s. Manually bring the resource in-service and retry the operation.	<p><b>Cause:</b> While attempting to create a dependency, the given resource was not in-service on the given server.</p> <p><b>Action:</b> Bring the resource in-service on the given server and retry the operation.</p>

Code	Severity	Message	Cause/Action
112101	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112102	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> The given resource cannot be extended to the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed on the target server and that any necessary shared file systems are accessible from the target system.</p>
112103	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The given resource could not be created on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted and accessible on the given server.</p>
112104	ERROR	The extend script "%s" either does not exist or is not executable on %s.	<p><b>Cause:</b> The given extend script does not exist or is not executable on the given server.</p> <p><b>Action:</b> Verify that all necessary recovery kits are installed and that the given extend script is executable.</p>
112106	ERROR	Unable to create an internal SAP object for resource "%s" on %s. If the tag is correct, verify that all necessary SAP file systems are mounted and accessible.	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>

Code	Severity	Message	Cause/Action
112112	ERROR	Error getting resource information for resource "%s" on server "%s".	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112115	FATAL	Error getting resource information for resource \$tag on server \$sap::me.	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112119	ERROR	Error getting resource information for resource "%s" on server "%s".	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112123	ERROR	Error getting resource information for resource "%s" on server "%s".	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112124	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server %s for app "%s" and type "%s".	<p><b>Cause:</b> A resource with the given resource tag or ID and the same app and type already exists on the given server.</p> <p><b>Action:</b> Verify that the SAP instance is not already under LifeKeeper protection on the given server. If it is not already protected, choose a different resource tag name.</p>

Code	Severity	Message	Cause/Action
112125	ERROR	Unable to create an SAP object for resource %s on system %s.	<p><b>Cause:</b> Unable to create an internal SAP object to represent the given resource on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted and accessible on the given server.</p>
112126	FATAL	Usage: "%s %s". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP canextend script.</p> <p><b>Action:</b> Specify the correct parameters for the canextend script: canextend &lt;template server&gt; &lt;template tag&gt;</p>
112127	FATAL	Usage: "\\$cmd \$usage\" Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP delete script.</p> <p><b>Action:</b> Specify the correct parameters for the delete script: delete [-U] -t &lt;tag&gt; -i &lt;id&gt;</p>
112128	FATAL	Usage: "%s %s" Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP depstoextend script.</p> <p><b>Action:</b> Specify the correct parameters for the depstoextend script: depstoextend &lt;template server&gt; &lt;template tag&gt;</p>
112129	FATAL	Usage: "%s %s". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP extend script.</p> <p><b>Action:</b> Specify the correct parameters for the extend script: extend &lt;template server&gt; &lt;template tag&gt; &lt;switchback&gt; &lt;target tag&gt;</p>

Code	Severity	Message	Cause/Action
112130	FATAL	Usage: \"\$cmd \$usage\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP quickCheck script.</p> <p><b>Action:</b> Specify the correct parameters for the quickCheck script: quickCheck -t &lt;tag&gt; -i &lt;id&gt;</p>
112131	FATAL	Usage: \"%s %s\" Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP recover script.</p> <p><b>Action:</b> Specify the correct parameters for the recover script: recover -d &lt;tag&gt;</p>
112132	FATAL	Usage: \"\$cmd \$usage\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP remoteControl script.</p> <p><b>Action:</b> Specify the correct parameters for the remoteControl script: remoteControl &lt;tag&gt; &lt;remote instance&gt; &lt;remote cmd&gt; &lt;remote cmd option&gt; &lt;primary system&gt; &lt;primary tag&gt;</p>
112133	FATAL	Usage: \"%s %s\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP remove script.</p> <p><b>Action:</b> Specify the correct parameters for the remove script: remove -t &lt;tag&gt; -i &lt;id&gt;</p>
112134	FATAL	Usage: \"\$cmd \$usage\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP restore script.</p> <p><b>Action:</b> Specify the correct parameters for the restore script: restore -t &lt;tag&gt; -i &lt;id&gt;</p>
112137	ERROR	The required parameter \"parent\" was either not provided or was invalid in the \$func routine on \$me.	<p><b>Cause:</b> Incorrect usage of the CleanUp routine.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.
112138	ERROR	At least one required process for instance %s was not started successfully during "%s" on server %s. Please check the LifeKeeper and system logs for additional information.	<p><b>Cause:</b> At least one required process for the given SAP instance did not start successfully on the given server.</p> <p><b>Action:</b> Correct any issues found in the LifeKeeper or system logs or SAP trace files and retry the operation.</p>
112140	FATAL	The tag parameter was not specified for the internal "\\$func\" routine on \$me. If this was a command line operation, specify the correct parameters. Otherwise, consult the troubleshooting documentation.	<p><b>Cause:</b> The tag parameter was not specified in the GetLK routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112141	ERROR	Either the SID ("%s") or instance ("%s") parameter was not specified for the "%s" routine on %s.	<p><b>Cause:</b> Either the SID or instance parameter was not specified in the StatusSapServer routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112142	ERROR	Either the SID ("%s"), instance ("%s"), or instance number ("%s") parameter was not specified for the "%s" routine on %s.	<p><b>Cause:</b> Either the SID, instance, or instance number parameter was not provided to the StartSapServer routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112143	ERROR	The SID, instance, or instance number parameter was not specified for the "%s" routine on %s.	<b>Cause:</b> Either the SID, instance, or instance number parameter was not provided to the StartSapServer routine on the given server.

Code	Severity	Message	Cause/Action
			<p>ance number parameter was not specified in the StopSapServer routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112173	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and is readable.	<p><b>Cause:</b> The given file does not exist or is not readable on the given server.</p> <p><b>Action:</b> Verify that the file exists and/or modify its permissions so that it is readable.</p>
112174	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and is readable.	<p><b>Cause:</b> The given file does not exist or is not readable on the given server.</p> <p><b>Action:</b> Verify that the file exists and/or modify its permissions so that it is readable.</p>
112175	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and is readable.	<p><b>Cause:</b> The given file does not exist or is not readable on the given server.</p> <p><b>Action:</b> Verify that the file exists and/or modify its permissions so that it is readable.</p>
112194	ERROR	There was an error verifying the NFS connections for SAP related mount points on %s. One or more NFS servers is not operational and needs to be restarted.	<p><b>Cause:</b> At least one critical NFS shared file system whose mount point is listed in the SAP_NFS_CHECK_DIRS entry in /etc/default/LifeKeeper is currently unavailable.</p> <p><b>Action:</b> Verify that all necessary NFS shared file systems are accessible and restart any NFS server which is not currently operational.</p>

Code	Severity	Message	Cause/Action
112195	FATAL	There was an error verifying the NFS connections for SAP related mount points on \$me. One or more NFS servers is not operational and needs to be restarted.	<p><b>Cause:</b> At least one critical NFS shared file system whose mount point is listed in the SAP_NFS_CHECK_DIRS entry in /etc/default/LifeKeeper is currently unavailable.</p> <p><b>Action:</b> Verify that all necessary NFS shared file systems are accessible and restart any NFS server which is not currently operational.</p>
112196	WARN	There was an error verifying the NFS connections for SAP related mount points on %s. One or more NFS servers is not operational and needs to be restarted.	<p><b>Cause:</b> At least one critical NFS shared file system whose mount point is listed in the SAP_NFS_CHECK_DIRS entry in /etc/default/LifeKeeper is currently unavailable.</p> <p><b>Action:</b> Verify that all necessary NFS shared file systems are accessible and restart any NFS server which is not currently operational.</p>
112201	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s.	<p><b>Cause:</b> No resource tag argument was provided to the GetLKEquiv routine.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112203	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s. Additional information available in the LifeKeeper and system logs.	<p><b>Cause:</b> The SAP instance number was not provided to the IsInstanceRunning routine.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112204	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s.	<p><b>Cause:</b> Unable to determine either the appropriate saphostexec or saposcol co</p>

Code	Severity	Message	Cause/Action
			<p>mmand to use.</p> <p><b>Action:</b> If using the SAP_SRVHOST_CMD, SAP_HOSTCTL_CMD, or SAP_OS_COL_CMD LifeKeeper tunable values to provide the appropriate commands for a version of SAP NetWeaver prior to SAP kernel 7.3, ensure that these tunable values are set appropriately.</p>
112205	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s.	<p><b>Cause:</b> The SAP instance was not provided to the SAPRemExec routine on the given system.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112206	ERROR	The required action parameter was not provided. Unable to complete "%s" on %s.	<p><b>Cause:</b> No action was provided to the SAPRemExec routine on the given system.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112208	ERROR	Error getting the value of "%s" or "%s" from the default profile "%s" on %s. Verify that the specified value exists.	<p><b>Cause:</b> Unable to obtain the virtual IP or host name from the given profile on the given server.</p> <p><b>Action:</b> Verify that the appropriate entry exists in the profile.</p>
112209	FATAL	Unable to gather required information from the SAP default profile for SID \$sid (\$DPFL) on \$me. Verify that the default profile exists and is accessible.	<p><b>Cause:</b> Unable to obtain information about the SAP instance from the given profile on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that the given pro</p>

Code	Severity	Message	Cause/Action
			file exists and is read-enabled.
112214	ERROR	Unable to determine the status of the path "%s" ("%s") on %s. The path on %s may require the execution of the command: "mount <ip>:<export> %s". Verify that the SAP software is correctly installed and that all SAP file systems are mounted and accessible.	<p><b>Cause:</b> The status of the file system on the given path could not be determined.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted.</p>
112219	ERROR	[HACONNECTOR:%s] Unable to write to file "%s" on %s. If the file already exists, manually enable write permissions on it.	<p><b>Cause:</b> The given file does not have read/write permissions enabled on the given server.</p> <p><b>Action:</b> Enable read/write permissions on the given file.</p>
112220	ERROR	Unable to start the sapstartsrv service for SID %s and SAP instance %s on %s. Verify that the sapservices file is correct and the process can be started manually.	<p><b>Cause:</b> Unable to start the SAP Start Service (sapstartsrv) process for the given SAP instance.</p> <p><b>Action:</b> Verify that the sapservices file contains the appropriate command to start the sapstartsrv process and that the process can be started manually.</p>
112221	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, or TYPE_NEW_ERS (1, 2, or 5).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112222	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, or TYPE_NEW_ERS (1, 2, or 5).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please</p>

Code	Severity	Message	Cause/Action
			submit an issue report to SIOS customer support.
112223	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_CS (1).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112224	ERROR	The internal "%s\" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_CS (1).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112225	ERROR	The internal "%s\" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_NEW_ERS (5).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112226	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, and TYPE_NEW_ERS (1, 2, and 5).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112227	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This routine only supports SAP instance type TYPE_CS (1).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.
112228	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_NEW_ERS (5).	<b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.  <b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.
112229	ERROR	The profile \"%s\" either does not exist or cannot be read on %s. Unable to determine whether enqueue replication is enabled for resource %s. Please verify that the file exists and can be read.	<b>Cause:</b> The given instance profile either does not exist or cannot be read on the given server.  <b>Action:</b> Verify that the file exists and is read-enabled
112325	ERROR	The \$SAP_CONTROL utility cannot be located. Verify that all necessary file systems are mounted and that the \$SAP_CONTROL utility can be located in the \$sapadm in user's PATH.	<b>Cause:</b> Unable to locate the sapcontrol utility required for SAP instance administration.  <b>Action:</b> Verify that all necessary file systems are mounted at that the sapcontrol utility can be located in the SAP administrative user's PATH.
112326	ERROR	The \"which \$SAP_CONTROL\" command for user \$sapadmin returned \$sapcmd as the location of the \$SAP_CONTROL utility, but the utility could not be found or was not executable in this location. Verify that all necessary file systems are mounted and that the \$SAP_CONTROL utility can be located in the \$sapadmin user's PATH.	<b>Cause:</b> Unable to locate the sapcontrol utility required for SAP instance administration.  <b>Action:</b> Verify that all necessary file systems are mounted at that the sapcontrol utility can be located in the SAP administrative user's PATH.
112433	ERROR	Unsupported SAPENQ_VERSION (\$enqversion) for resource \$tag on \$me. Unable to obtain enqueue replication status.	<b>Cause:</b> An unsupported value was detected.

Code	Severity	Message	Cause/Action
			<p>cted for the SAPENQ_VERSION parameter for the given resource on the given server.</p> <p><b>Action:</b> Verify that SAPENQ_VERSION is set to a valid value (1 or 2, representing the version of the enqueue server currently in use) in the info file for the given resource.</p>
112437	ERROR	Profile \"\$srvpf\" not found on \$me. Unable to obtain enqueue replication status for instance \$inst.	<p><b>Cause:</b> The given instance profile either does not exist or cannot be read on the given server.</p> <p><b>Action:</b> Verify that the file exists and is read-enabled.</p>
112438	ERROR	Unsupported SAPENQ_VERSION (\$enqverson) for resource \$tag on \$me. Unable to obtain enqueue replication status for instance \$inst.	<p><b>Cause:</b> An unsupported value was detected for the SAPENQ_VERSION parameter for the given resource on the given server.</p> <p><b>Action:</b> Verify that SAPENQ_VERSION is set to a valid value (1 or 2, representing the version of the enqueue server currently in use) in the info file for the given resource.</p>
112470	ERROR	The instance profile %s could not be found on %s. Verify that the SAP software is installed correctly and that all necessary file systems are mounted and accessible.	<p><b>Cause:</b> The given instance profile either does not exist or cannot be read on the given server.</p> <p><b>Action:</b> Verify that the file exists and is read-enabled.</p>
112490	ERROR	[HACONNECTOR] Unable to determine the corresponding tag for resource with ID \"\$res\" on \$me.	<p><b>Cause:</b> Unable to find a LifeKeeper SAP resource with the given resource ID on the given server.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Verify that the resource ID provided as the <code>—res</code> argument of the <code>fra</code> command corresponds to a valid LifeKeeper SAP resource.</p>
112507	ERROR	[HACONNECTOR] At least one required process for the instance was not killed successfully during the <code>fra migrate</code> action on <code>\$me</code> . Aborting resource migration.	<p><b>Cause:</b> The SAP instance was not successfully stopped on the given server while attempting a <code>fra migrate</code> action.</p> <p><b>Action:</b> Manually kill any processes still running for the SAP instance and reattempt the migrate action.</p>
112539	ERROR	[HACONNECTOR] Unable to find rpm information for one or more packages on <code>\$me</code> .	<p><b>Cause:</b> The HA Connector <code>gvi</code> ("Get Version Information") routine was unable to determine the current version number of LifeKeeper and/or the SAP Recovery Kit.</p> <p><b>Action:</b> Verify that the LifeKeeper Core and SAP Recovery Kit rpm information can be obtained with the <code>rpm -q</code> command.</p>
112976	ERROR	There is no LifeKeeper protected resource with tag <code>\$tag</code> on system <code>\$me</code> .	<p><b>Cause:</b> The resource tag provided to the SAP canfailover script does not correspond to any existing LifeKeeper resource.</p> <p><b>Action:</b> Verify that the resource tag name is correct and execute the command again.</p>
112977	ERROR	Resource <code>\$tag</code> is not a <code>\$app/\$typ</code> resource. Please use the <code>\$ins_app/\$ins_typ</code> resource-specific canfailover script instead.	<p><b>Cause:</b> The resource provided to the SAP canfailover script is not an <code>appsuite/sap</code> resource.</p> <p><b>Action:</b> Use the appropriate type-specific canfailover script for the given resource.</p>

Code	Severity	Message	Cause/Action
			ce.
122005	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "getlocks".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122007	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "rlslocks".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122009	ERROR	The path %s is not a valid file.	<p><b>Cause:</b> There is no listener.ora file.</p> <p><b>Action:</b> Ensure the file exists and retry the operation.</p>
122010	ERROR	The listener user does not exist on the server %s.	<p><b>Cause:</b> "Stat" command could not get user id.</p> <p><b>Action:</b> Retry the operation.</p>
122011	ERROR	The listener user does not exist on the server %s.	<p><b>Cause:</b> UID is not in passwd file.</p> <p><b>Action:</b> Ensure the UID exists in passwd file and retry the operation.</p>
122012	ERROR	The listener user does not exist on the server %s.	<p><b>Cause:</b> User name is not in passwd file.</p> <p><b>Action:</b> Ensure the user name exists in passwd file; retry the operation.</p>

Code	Severity	Message	Cause/Action
122023	ERROR	The %s command failed (%d	<p><b>Cause:</b> This message contains the return code of the "lsnrctl" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122024	ERROR	\$line	<p><b>Cause:</b> The message contains the output of the "lsnrctl" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122039	ERROR	Usage error	<p><b>Cause:</b> Invalid parameters were specified for the restore operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122040	ERROR	Script \$cmd has hung on the restore of \"\$opt_t\". Forcibly terminating.	<p><b>Cause:</b> The listener restore script reached its timeout value.</p> <p><b>Action:</b> Ensure listener.ora is valid and that LSNR_START_TIME (default 35 seconds) in /etc/default/LifeKeeper is set to a value greater than or equal to the time needed to start the listener.</p>
122041	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to restore the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122045	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Failed to get resource information.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check your LifeKeeper configuration.
122046	ERROR	Usage error	<b>Cause:</b> Invalid parameters were specified for the restore operation.  <b>Action:</b> Verify the parameters and retry the operation.
122049	ERROR	The script \$cmd has hung on remove of \"\$opt_t\". Forcibly terminating.	<b>Cause:</b> The listener remove script reached its timeout value.  <b>Action:</b> Ensure listener.ora is valid and that LSNR_STOP_TIME (default 35 seconds) in /etc/default/LifeKeeper is set to a value greater than or equal to the time needed to stop the listener.
122051	ERROR	Error getting resource information for resource \"%s\" on server \"%s\"	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.  <b>Action:</b> Check your LifeKeeper configuration.
122055	ERROR	END failed %s of \"%s\" on server \"%s\" due to a \"%s\" signal	<b>Cause:</b> LifeKeeper was unable to quick Check the resource {resource} on {server}.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122057	ERROR	Error getting resource information for resource \"%s\" on server \"%s\"	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.  <b>Action:</b> Check your LifeKeeper configuration.

Code	Severity	Message	Cause/Action
122064	WARN	The %s level is set to %s a %s will not occur.	<p><b>Cause:</b> The minimal Listener protection level is Start and Monitor.</p> <p><b>Action:</b> Start the listener manually.</p>
122066	ERROR	Script has hung checking \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The listener quickCheck script reached its timeout value.</p> <p><b>Action:</b> Ensure listener.ora is valid and that LSNR_STATUS_TIME (default 15 seconds) in /etc/default/LifeKeeper is set to a value greater than or equal to the time needed to check the listener.</p>
122067	ERROR	Usage error	<p><b>Cause:</b> Invalid parameters were specified for the quickCheck operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122069	ERROR	Usage error	<p><b>Cause:</b> Invalid parameters were specified for the delete operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122072	ERROR	%s: resource "%s" not found on local server	<p><b>Cause:</b> Invalid parameters were specified for the recover operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122074	WARN	The local recovery attempt has failed but %s level is set to %s preventing a failover to another node in the cluster. With %s recovery set all local recovery failures will exit successfully to prevent resource failovers.	<p><b>Cause:</b> The optional listener recovery level is set to local recovery only.</p> <p><b>Action:</b> Switch over the resource tree manually.</p>

Code	Severity	Message	Cause/Action
122078	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to recover the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122082	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122083	ERROR	\$cmd has hung checking \"\$tag\". Forcibly terminating	<p><b>Cause:</b> The recover script was stopped by signal.</p> <p><b>Action:</b> Ensure listener.ora is valid.</p>
122084	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122085	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the canextend operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122086	ERROR	The values specified for the target and the template servers are the same. Please specify the correct values for the target and template servers.	<p><b>Cause:</b> The values specified for the target and the template servers are the same.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>

Code	Severity	Message	Cause/Action
122087	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122088	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Failed to get listener user name from resource information.</p> <p><b>Action:</b> Ensure the resource info field is valid then retry the operation.</p>
122089	ERROR	The listener user %s does not exist on the server %s.	<p><b>Cause:</b> User name is not in passwd file.</p> <p><b>Action:</b> Ensure the user name exists in passwd file and retry the operation.</p>
122090	ERROR	The id for user %s is not the same on template server %s and target server %s.	<p><b>Cause:</b> User ID should be same on both servers.</p> <p><b>Action:</b> Trim user ID to the same.</p>
122091	ERROR	The group id for user %s is not the same on template server %s and target server %s.	<p><b>Cause:</b> Group ID should be same on both servers.</p> <p><b>Action:</b> Trim group ID to the same.</p>
122092	ERROR	Cannot access canextend script "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to run pre-extend checks because it was unable to find the "canextend" script on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122097	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "configActions" operation.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Verify the arguments and retry the operation.
122098	ERROR	Error getting resource information for resource "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}. <b>Action:</b> Check your LifeKeeper configuration.
122099	ERROR	Unable to update the resource %s to change the %s to %s on %s.	<b>Cause:</b> LifeKeeper failed to put information into the info field. <b>Action:</b> Restart LifeKeeper and retry the operation.
122100	ERROR	Error getting resource information for resource "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}. <b>Action:</b> Check your LifeKeeper configuration.
122101	ERROR	Unable to update the resource %s to change the %s to %s on %s.	<b>Cause:</b> LifeKeeper failed to put information to info field on {server}. <b>Action:</b> Restart LifeKeeper on {server} and retry the operation.
122103	ERROR	Usage: %s %s	<b>Cause:</b> Invalid parameters were specified for the create operation. <b>Action:</b> Verify the parameters and retry the operation.
122124	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	<b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}. <b>Action:</b> Check the logs for related error

Code	Severity	Message	Cause/Action
			s and try to resolve the reported problem.
122126	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "rlslocks".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122127	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122129	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "getlocks."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122131	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122133	ERROR	Unable to create a file system resource hierarchy for the file system %s.	<p><b>Cause:</b> There was an unexpected error running "filesyshier."</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
122135	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error</p>

Code	Severity	Message	Cause/Action
			<p>running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122140	ERROR	Resource "%s" is not ISP on server "%s" Manually bring the resource in service and retry the operation	<p><b>Cause:</b> IP resource {tag} which the listener resource depends on should be ISP.</p> <p><b>Action:</b> Perform the steps listed in the message text.</p>
122141	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122144	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the "create_ins" operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122145	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for details and retry the operation.	<p><b>Cause:</b> There was an unexpected error running "app_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122146	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for details and retry the operation.	<p><b>Cause:</b> There was an unexpected error running "typ_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
			m.
122147	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for details and retry the operation.	<p><b>Cause:</b> There was an unexpected error running "newtag."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122148	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122149	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for details and retry the operation.	<p><b>Cause:</b> There was an unexpected error running "ins_setstate."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122150	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122151	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "depstoextend" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122152	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the "extend" operation.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Verify the parameters and retry the operation.
122153	ERROR	Error getting resource information for resource "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}. <b>Action:</b> Check your LifeKeeper configuration.
122154	ERROR	Cannot extend resource "%s" to server "%s"	<b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.
122155	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server "%s" for App "%s" and Type "%s"	<b>Cause:</b> During the Listener resource extension, a resource instance was found using the same {tag} and/or {id} but with a different resource application and type. <b>Action:</b> Resource IDs must be unique. The resource instance with the ID matching the Oracle Listener resource instance must be removed.
122156	ERROR	Cannot access extend script "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to extend the resource hierarchy because it was unable to find the script EXTEND on {server}. <b>Action:</b> Check your LifeKeeper configuration.
122157	ERROR	Usage: %s %s	<b>Cause:</b> Invalid arguments were specified for the "getConfigIps" operation. <b>Action:</b> Verify the arguments and retry the operation.

Code	Severity	Message	Cause/Action
122158	ERROR	The file %s is not a valid listener file. The file does not contain any listener definitions.	<p><b>Cause:</b> Failed to find any listener definitions.</p> <p><b>Action:</b> Ensure listener definition is in the listener.ora and retry the operation.</p>
122159	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "getSidListeners" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122160	ERROR	The file %s is not a valid listener file. The file does not contain any listener definitions.	<p><b>Cause:</b> Failed to find any listener definitions.</p> <p><b>Action:</b> Ensure listener definition is in the listener.ora and retry the operation.</p>
122161	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "lsn-display" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122162	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122163	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the updateHelper operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122164	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of	<p><b>Cause:</b> LifeKeeper was unable to update</p>

Code	Severity	Message	Cause/Action
		%d	<p>e the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122166	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the "updateHelper" operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122170	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122171	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122172	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "updIPDeps" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122173	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	<p><b>Cause:</b> LifeKeeper was unable to update the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
122175	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "rlslocks."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122177	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "getlocks."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122180	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122181	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122183	ERROR	The path %s is not a valid file.	<p><b>Cause:</b> There is no listener.ora file.</p> <p><b>Action:</b> Ensure the file exists and retry the operation.</p>
122185	ERROR	The file %s is not a valid listener file. The file does not contain any listener definitions.	<p><b>Cause:</b> LifeKeeper failed to find any valid listener definitions.</p> <p><b>Action:</b> Ensure there are valid listener definitions in the listener.ora and retry the operation.</p>

Code	Severity	Message	Cause/Action
			he operation.
122186	ERROR	The value specified for %s cannot be empty. Please specify a value for this field.	<p><b>Cause:</b> The config and/or executable {path} field is empty.</p> <p><b>Action:</b> Input a non-empty value for {path} and retry the operation.</p>
122187	ERROR	The path %s is not a valid file or directory.	<p><b>Cause:</b> The defined {path} is invalid.</p> <p><b>Action:</b> Ensure the {path} exists and retry the operation.</p>
122188	ERROR	The path %s is not a valid file or directory.	<p><b>Cause:</b> There is no {path}.</p> <p><b>Action:</b> Ensure the {path} exists and retry the operation.</p>
122189	ERROR	The value specified for %s cannot be empty. Please specify a value for this field.	<p><b>Cause:</b> The config and/or executable Path field is empty.</p> <p><b>Action:</b> Input path for the field.</p>
122190	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "valid_rpath" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122191	ERROR	The values specified for the target and the template servers are the same.	<p><b>Cause:</b> Invalid argument of valid_rpath.</p> <p><b>Action:</b> Ensure arguments and retry the operation.</p>
122192	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<p><b>Cause:</b> There is no oratab file in /etc/or</p>

Code	Severity	Message	Cause/Action
			<p>atab or {path}.</p> <p><b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.</p>
122193	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<p><b>Cause:</b> There is no oratab file in /etc/oratab or {path}.</p> <p><b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.</p>
122194	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<p><b>Cause:</b> There is no oratab file in /etc/oratab or {path}.</p> <p><b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.</p>
122195	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<p><b>Cause:</b> There is no oratab file in /etc/oratab or {path}.</p> <p><b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.</p>
122196	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to remove the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122197	ERROR	Unable to find the configuration file \"oratab\" in its default locations, /etc/oratab or \$!istener::oraTab on \"\$me\"	<p><b>Cause:</b> There is no oratab file in /etc/oratab or {path}.</p> <p><b>Action:</b> Ensure oratab file exists in {pat</p>

Code	Severity	Message	Cause/Action
			h} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.
122198	ERROR	remove for \$okListener failed.	
122251	ERROR	Update of pluggable database info field for "%s" on "%s" failed (%s).	
122252	ERROR	Initial connect with query buffer to database "%s" on "%s" failed, testing output.	
122253	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	
122261	ERROR	The Oracle resource (%s) and dependency are not set on %s.	
122262	ERROR	Usage: %s %s	
122263	ERROR	The restore of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_RESTORE_TIMEOUT in /etc/default/LifeKeeper.	
122264	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122268	ERROR	Failed to create object instance for Oracle on "%s".	
122269	ERROR	no dependency for Oracle on "%s".	
122270	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122271	ERROR	Usage: %s %s	
122272	ERROR	The remove of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_REMOVE_TIMEOUT in /etc/default/LifeKeeper.	
122273	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122277	ERROR	Usage: %s %s	
122278	ERROR	Failed to create object instance for Oracle	

Code	Severity	Message	Cause/Action
		on "%s".	
122279	ERROR	no dependency for Oracle on "%s".	
122280	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122281	ERROR	The quickCheck of %s has timed out on server %s. The default TIMEOUT is 45 seconds. To increase the TIMEOUT, set ORACLE_QUICKCHECK_TIMEOUT in /etc/default/LifeKeeper.	
122282	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122284	ERROR	Failed to create object instance for Oracle on "%s".	
122285	ERROR	no dependency for Oracle on "%s".	
122287	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122288	ERROR	Usage: %s %s	
122291	ERROR	Cannot extend resource "%s" to server "%s"	
122292	ERROR	The values specified for the target and the template servers are the same: "%s".	
122294	ERROR	Cannot access canextend script "%s" on server "%s"	
122295	ERROR	Usage: %s %s	
122296	ERROR	DB instance "%s" is not protected on "%s".	
122297	ERROR	Failed to create object instance for Oracle PDB on "%s".	
122298	ERROR	Unable to locate the oratab file "%s" on "%s".	
122299	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	
122301	ERROR	Unable to "%s" on "%s" during resource create.	
122302	ERROR	END failed %s of "%s" on server "%s" du	

Code	Severity	Message	Cause/Action
		e to a "%s" signal	
122304	ERROR	Unable to "%s" on "%s" during resource create.	
122305	ERROR	Unable to determine Oracle user for "%s" on "%s".	
122306	ERROR	Error creating resource "%s" on server "%s"	
122308	ERROR	Dependency creation between Oracle pluggable database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	
122309	ERROR	%s	
122311	ERROR	In-service attempted failed for tag "%s" on "%s".	
122312	ERROR	Usage: %s %s	
122313	ERROR	Create of app "%s" on "%s" failed with return code of "%d".	
122314	ERROR	Create of typ "%s" for app "%s" on "%s" failed with return code of "%d".	
122316	ERROR	Create of resource tag via "newtag" on "%s" failed.	
122318	ERROR	Error creating resource "%s" on server "%s"	
122320	ERROR	Setting "resstate" for resource "%s" on "%s" failed with return code of "%d".	
122321	ERROR	Error creating resource "%s" on server "%s"	
122322	ERROR	Usage: %s %s	
122323	ERROR	Usage: %s %s	
122324	ERROR	Usage: %s %s	
122325	ERROR	Cannot extend resource "%s" to server "%s"	
122326	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server "%s" for App "%s" and Type "%s"	
122327	ERROR	Error creating resource "%s" on server	

Code	Severity	Message	Cause/Action
		"%s"	
122328	ERROR	Cannot access extend script "%s" on server "%s"	
122329	ERROR	Cannot extend resource "%s" to server "%s"	
122330	ERROR	Usage: %s %s	
122331	ERROR	Failed to create object instance for Oracle PDB on "%s".	
122332	ERROR	Usage: %s %s	
122334	ERROR	Backup node %s is unreachable; abort protection PDB changes.	
122336	ERROR	Update of protection PDB failed for "%s" on "%s".	
122339	ERROR	Usage: %s %s	
122340	ERROR	Usage: %s %s	
122341	ERROR	Usage: %s %s	
122342	ERROR	Failed to create object instance for Oracle PDB on "%s".	
122343	ERROR	Usage: %s %s	
122344	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122348	ERROR	Failed to create flag "%s" on "%s".	
122350	ERROR	Failed to create object instance for Oracle on "%s".	
122351	ERROR	no dependency for Oracle on "%s".	
122353	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122356	ERROR	Usage: %s %s	
122357	ERROR	Failed to create object instance for Oracle PDB on "%s".	
122358	ERROR	The selected oracle SID "%s" is not a CDB.	
122359	ERROR	No protectable PDB found for the selected SID "%s".	

Code	Severity	Message	Cause/Action
122360	ERROR	No protected Oracle database found on "%s".	
122500	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the create operation.</p> <p><b>Action:</b> Verify the parameters are correct and retry the operation.</p>
122501	ERROR	DB instance "%s" is already protected on "%s".	<p><b>Cause:</b> An attempt was made to protect an Oracle database instance {sid} that is already under LifeKeeper protection on {server}.</p> <p><b>Action:</b> You must select a different database instance {sid} for LifeKeeper protection.</p>
122502	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122503	ERROR	Unable to locate the oratab file "%s" on "%s".	<p><b>Cause:</b> The oratab file was not found at the default or alternate locations on {server}.</p> <p><b>Action:</b> Verify the oratab file exists and has proper permissions for the Oracle user. A valid oratab file is required to complete the "create" operation.</p>
122504	ERROR	Unable to determine Oracle user for "%s" on "%s".	<p><b>Cause:</b> The Oracle Application Recovery Kit was unable to determine the ownership of the Oracle database installation binaries.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> The owner of the Oracle binaries must be a valid non-root user on {server}. Correct the permissions and ownership of the Oracle database installation and retry the operation.</p>
122505	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	<p><b>Cause:</b> The database instance {sid} was not running or connections to the database were not available via the credentials provided.</p> <p><b>Action:</b> The database instance {sid} must be started on {server} and the proper credentials must be provided for the completion of the "create" operation.</p>
122506	ERROR	Unable to determine Oracle tablespaces and logfiles for "%s" on "%s".	<p><b>Cause:</b> A query to determine the location of required tablespaces, logfiles and related database files failed. This may have been caused by an internal database error.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122507	ERROR	Unknown chunk type found for "%s" on "%s".	<p><b>Cause:</b> The specified tablespace, logfile or other required database file is not one of the LifeKeeper supported file or character device types.</p> <p><b>Action:</b> The specified file {database_file} must reference an existing character device or file. Consult the Oracle installation documentation to recreate the specified file {database_file} as a supported file or character device type.</p>

Code	Severity	Message	Cause/Action
122508	ERROR	DB Chunk "%s" for "%s" on "%s" does not reside on a shared file system.	<p><b>Cause:</b> The specified tablespace, logfile or other required database file {database_file} does not reside on a file system that is shared with other systems in the cluster.</p> <p><b>Action:</b> Use the LifeKeeper UI or "lkdstatus (1M)" to verify that communication paths have been properly created. Use "rpm" to verify that the necessary Application Recovery Kits for storage protection have been installed. Verify that the file is, in fact, not on shared storage, and if not, move it to a shared storage device.</p>
122510	ERROR	File system create failed for "%s" on "%s". Reason	<p><b>Cause:</b> LifeKeeper was unable to create the resource {filesystem} on the specified server {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the "create" operation.</p>
122511	ERROR	%s	<p><b>Cause:</b> The message contains the output of the "filesyshier" command.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
122513	ERROR	Dependency creation between Oracle database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	<p><b>Cause:</b> LifeKeeper was unable to create a dependency between the database resource {tag} and the necessary child resource {childtag}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Once any problems have been corrected,</p>

Code	Severity	Message	Cause/Action
			<p>d, it may be possible to create the dependency between {tag} and {childtag} manually.</p>
122514	ERROR	Unable to "%s" on "%s" during resource create.	<p><b>Cause:</b> The Oracle Application Recovery Kit was unable to release the administrative lock using the "rlslocks" command.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
122516	ERROR	Raw device resource created failed for "%s" on "%s". Reason	<p><b>Cause:</b> LifeKeeper was unable to create the resource {raw device} on the specified server {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the "create" operation.</p>
122519	ERROR	In-service attempted failed for tag "%s" on "%s".	<p><b>Cause:</b> The "perform_action" command for {tag} on {server} failed to start the database {sid}. The in-service operation has failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the "create" operation.</p>
122521	ERROR	Create of app "%s" on "%s" failed with return code of "%d".	<p><b>Cause:</b> There was an error running the command "app_create" to create the internal application type.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>

Code	Severity	Message	Cause/Action
122522	ERROR	Create of typ "%s" for app "%s" on "%s" failed with return code of "%d".	<p><b>Cause:</b> There was an error running the command "typ_create" to create the internal resource type.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122524	ERROR	Setting "resstate" for resource "%s" on "%s" failed with return code of "%d".	<p><b>Cause:</b> There was an error running the command "ins_setstate" to set the resource state to {state}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122525	ERROR	The values specified for the target and the template servers are the same: "%s".	<p><b>Cause:</b> The value specified for the target and template servers for the "extend" operation were the same.</p> <p><b>Action:</b> You must specify the correct parameter for the {target server} and {template server}. The {target server} is the server where the {tag} will be extended.</p>
122526	ERROR	Unable to locate the oratab file in "/etc" or in "%s" on "%s".	<p><b>Cause:</b> The oratab file was not found at the default or alternate locations on {server}.</p> <p><b>Action:</b> Verify the oratab file exists and has proper permissions for the Oracle user. A valid oratab file is required to complete the "extend" operation.</p>
122527	ERROR	Unable to retrieve the Oracle user on "%s".	<p><b>Cause:</b> An attempt to retrieve the Oracle user from {template server} during a "canextend" or "extend" operation failed.</p>

Code	Severity	Message	Cause/Action
			<p>d.</p> <p><b>Action:</b> The owner of the Oracle binaries must be a valid user on {target server} and {template server}. Correct the permissions and ownership of the Oracle database installation and retry the operation.</p>
122528	ERROR	The Oracle user and/or group information for user "%s" does not exist on the server "%s".	<p><b>Cause:</b> LifeKeeper is unable to find the Oracle user and/or group information for the Oracle user {user} on the server {server}.</p> <p><b>Action:</b> Verify the Oracle user {user} exists on the specified {server}. If the user {user} does not exist, it should be created with the same uid and gid on all servers in the cluster.</p>
122529	ERROR	The id for user "%s" is not the same on template server "%s" and target server "%s".	<p><b>Cause:</b> The user id on the target server {target server} for the Oracle user {user} does not match the value of the user {user} on the template server {template server}.</p> <p><b>Action:</b> The user ids for the Oracle user {user} must match on all servers in the cluster. The user id mismatch should be corrected manually on all servers before retrying the "extend" operation.</p>
122530	ERROR	The group id for user "%s" is not the same on template server "%s" and target server "%s".	<p><b>Cause:</b> The group id on the target server {target server} for the Oracle user {user} does not match the value of the user {user} on the template server {template server}.</p> <p><b>Action:</b> The group ids for the Oracle user {user} must match on all servers in th</p>

Code	Severity	Message	Cause/Action
			e cluster. The group id mismatch should be corrected manually on all servers before retrying the "extend" operation.
122532	ERROR	No file system or raw devices found to extend for "%s" on "%s".	<p><b>Cause:</b> There were no dependent file system or raw device resources found for the Oracle resource {tag} on server {template server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
122533	WARN	A RAMDISK (%s) was detected in the ORACLE Database configuration for "%s" on "%s". LifeKeeper cannot protect RAMDISK. This RAMDISK resource will not be protected by LifeKeeper! ORACLE hierarchy creation will continue.	<p><b>Cause:</b> The specified tablespace, logfile or other database file {database_file} was detected as a ramdisk. No protection is available for this type of resource in the current LifeKeeper product.</p> <p><b>Action:</b> The ramdisk will not be protected. You must manually ensure that the required database file {database_file} will be available during all Oracle database operations.</p>
122534	ERROR	Failed to initialize object instance for Oracle sid "%s" on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122537	ERROR	Update of instance info field for "%s" on "%s" failed (%s).	<p><b>Cause:</b> There was an error while running the command "ins_setinfo" to update the internal resource information field.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>

Code	Severity	Message	Cause/Action
			<p>You must correct any reported errors before retrying the operation.</p>
122538	ERROR	Initial connect with query buffer to database "%s" on "%s" failed, testing output.	<p><b>Cause:</b> A connection attempt to the Oracle database {sid} to determine the database status has failed.</p> <p><b>Action:</b> The connection attempt failed with the specified credentials. Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122542	ERROR	The "%s [%s]" attempt of the database "%s" appears to have failed on "%s".	<p><b>Cause:</b> The attempted Oracle action {action} using method {action_method} for the database instance {sid} failed on the server {server}.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122543	ERROR	All attempts to "%s" database "%s" on "%s" failed	<p><b>Cause:</b> All efforts to perform the action {action} on the Oracle database {sid} on server {server} have failed.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122544	ERROR	Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".	<p><b>Cause:</b> An unexpected error occurred with</p>

Code	Severity	Message	Cause/Action
			<p>hile attempting to update the oratab entry for the database {sid}. The error occurred while attempting to open the oratab file.</p> <p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122545	ERROR	Unable to locate the oratab file in "/etc" or in "%s" on "%s".	<p><b>Cause:</b> The oratab file was not found at the default or alternate locations on {server}.</p> <p><b>Action:</b> Verify the oratab file exists and has proper permissions for the oracle user. A valid oratab file is required to complete the "extend" operation.</p>
122546	ERROR	Unable to open file "%s" on "%s" (%s).	<p><b>Cause:</b> The specified file {file} could not be opened or accessed on the server {server} due to the error {error}.</p> <p><b>Action:</b> Verify the existence and permissions on the specified file {file}. Check adjacent log messages for further details and related errors. You must correct any reported errors before retrying the operation.</p>
122547	ERROR	(cleanUpPids):Forcefully killing hung pid(s):pid(s)="%s"	<p><b>Cause:</b> The process {pid} failed to respond to the request to terminate gracefully. The process {pid} will be forcefully terminated.</p> <p><b>Action:</b> Use the command line to verify that the process {pid} has been terminated. Check the adjacent log messages for further details and related messages.</p>

Code	Severity	Message	Cause/Action
122548	ERROR	Unable to locate the DB utility (%s/%s) on this host.	<p><b>Cause:</b> The Oracle binaries and required database utility {utility} located at {path/utility} were not found on this server {server}.</p> <p><b>Action:</b> Verify that the Oracle binaries and required software utilities are installed and properly configured on the server {server}. The Oracle binaries must be installed locally on each node or located on shared storage available to all nodes in the cluster.</p>
122549	ERROR	Oracle internal error or non-standard Oracle configuration detected. Oracle User and/or Group set to "root".	<p><b>Cause:</b> The detected ownership of the Oracle database installation resolves to the root user and/or root group. Ownership of the Oracle installation by root is a non-standard configuration.</p> <p><b>Action:</b> The owner of the Oracle binaries must be a valid non-root user on {server}. Correct the permissions and ownership of the Oracle database installation and retry the operation.</p>
122550	ERROR	Initial inspection of "%s" failed, verifying failure or success of received output.	<p><b>Cause:</b> The previous Oracle query {query} or command {cmd} failed to return success.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122551	ERROR	Logon failed with "%s" for "%s" on "%s". Please check username/password and privileges.	<p><b>Cause:</b> The logon with the credentials {credentials} for the database instance {sid} on server {server} failed. An invalid user {user} or password was specified.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Verify that the Oracle database user {user} and password {password} are indeed valid. In addition, the Oracle database user {user} must have sufficient privileges for the attempted action.</p>
122552	ERROR	%s	<p><b>Cause:</b> The message contains the output of the "sqlplus" command.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
122553	ERROR	Unable to open file "%s" on "%s" (%s).	<p><b>Cause:</b> The specified file {file} could not be opened or accessed on the server {server} due to the error {error}.</p> <p><b>Action:</b> Verify the existence and permissions on the specified file {file}. Check adjacent log messages for further details and related errors. You must correct any reported errors before retrying the operation.</p>
122554	ERROR	The tag "%s" on "%s" is not an Oracle instance or it does not exist.	<p><b>Cause:</b> The specified tag {tag} on server {server} does not refer to an existing and valid Oracle resource instance.</p> <p><b>Action:</b> Use the UI or "lcdstatus (1M)" to verify the existence of the resource tag {tag}. The resource tag {tag} must be an Oracle resource instance to use the command "ora-display."</p>
122555	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to update the authorized user, password and database role for the Oracle resource instance.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122557	ERROR	Update of user and password failed for "%s" on "%s".	<p><b>Cause:</b> A request to update the user and password for the resource tag {tag} failed. The specified credentials failed the initial validation/connection attempt on server {server}.</p> <p><b>Action:</b> Verify the correct credentials {user/password} were specified for the attempted operation. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122559	ERROR	Update of user and password failed for "%s" on "%s".	<p><b>Cause:</b> The update of the user and password information for the resource tag {tag} on server {server} failed.</p> <p><b>Action:</b> Verify the correct credentials {user/password} were specified for the attempted operation. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122562	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The required Oracle executable {exe} was not found on this server {server}.</p> <p><b>Action:</b> Verify that the Oracle binaries and required software utilities are installed and properly configured on the server {server}. The Oracle binaries must be installed locally on each node or located on shared storage available to all nodes in the cluster.</p>

Code	Severity	Message	Cause/Action
122566	ERROR	Unable to find Oracle home for "%s" on "%s".	<p><b>Cause:</b> The Oracle home directory {Oracle home} does not appear to contain files necessary for the proper operation of the Oracle instance {sid}.</p> <p><b>Action:</b> Verify using the command line that the Oracle home directory {Oracle home} contains the Oracle binaries, a valid spfile{sid}.ora or init{sid}.ora file.</p>
122567	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected. The specified internal ID {id} does not match the expected SID {sid}.</p> <p><b>Action:</b> Verify the parameters are correct. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122568	ERROR	DB Processes are not running on "%s".	<p><b>Cause:</b> A process check for the Oracle instance did not find any processes running on server {server}.</p> <p><b>Action:</b> If local recovery is enabled, the Oracle instance will be restarted locally. Check adjacent log messages for further details and related messages.</p>
122572	ERROR	Failed to create flag "%s" on "%s".	<p><b>Cause:</b> An unexpected error occurred attempting to create a flag for controlling Oracle local recovery processing causing a failover to the standby node.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>

Code	Severity	Message	Cause/Action
122574	ERROR	all attempts to shutdown the database %s failed on "%s".	<p><b>Cause:</b> The shutdown of the Oracle database failed during a local recovery process most likely caused because the maximum number of database connections has been reached.</p> <p><b>Action:</b> Check the Oracle logs for connection failures caused by the maximum number of available connections being reached, and if found, consider increasing the value. Additionally, set the tunable LK_ORA_NICE to 1 to prevent connection failures from causing a quickCheck failure followed by a local recovery attempt.</p>
122597	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected during pre-extend checking.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the pre-extend.</p>
122598	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being created while attempting to determine the validity of the Oracle home directory.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the "create."</p>
122599	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while</p>

Code	Severity	Message	Cause/Action
			<p>attempting to look up the Oracle user on the template system.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the extend.</p>
122600	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to display the resource properties.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the display of the resource properties.</p>
122601	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to check for valid database authorization.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the command.</p>
122603	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to perform health checks on the Oracle resource instance.</p> <p><b>Action:</b> Check that correct arguments were passed to the quickCheck command and also check the adjacent log messages.</p>

Code	Severity	Message	Cause/Action
			<p>es for further details and related messages. Correct any reported errors before retrying the restore.</p>
122604	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to perform a local recovery on the Oracle resource instance.</p> <p><b>Action:</b> Check that correct arguments were passed to the "recover" command, and also check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the recover.</p>
122606	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	<p><b>Cause:</b> The database instance {sid} was not running or connections to the database are not available via the credentials provided.</p> <p><b>Action:</b> The database instance {sid} must be started on {server} and the proper credentials must be provided for the completion of the selected operation.</p>
122607	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	<p><b>Cause:</b> The database instance {sid} was not running or connections to the database are not available via the credentials provided.</p> <p><b>Action:</b> The database instance {sid} must be started on {server} and the proper credentials must be provided for the completion of the selected operation.</p>
122608	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> The "remove" operation failed to create the resource object instance re</p>

Code	Severity	Message	Cause/Action
			<p>quired to take the Oracle resource Out of Service.</p> <p><b>Action:</b> Check that correct arguments were passed to the "remove" command and also check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the restore.</p>
122609	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> The "restore" operation failed to create the resource object instance required to put the Oracle resource In Service.</p> <p><b>Action:</b> Check that correct arguments were passed to the "restore" command and also check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the "restore."</p>
122610	ERROR	Unable to "%s" on "%s" during resource create.	<p><b>Cause:</b> The Oracle Application Recovery Kit was unable to create the administrative lock using the "getlocks" command during resource creation.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create.</p>
122611	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle resource and the child File System resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create.</p>

Code	Severity	Message	Cause/Action
			ng the create operation.
122612	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle resource and the child Raw resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create operation.</p>
122613	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle resource and the child Raw resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create operation.</p>
122614	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle resource and the child Listener resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create operation.</p>
122616	ERROR	%s	<p><b>Cause:</b> The requested start up or shutdown of the Oracle database failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the "restore" or "remove" operation.</p>

Code	Severity	Message	Cause/Action
122618	ERROR	Dependency creation between Oracle database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	<p><b>Cause:</b> LifeKeeper was unable to create a dependency between the database resource {tag} and the necessary child resource {childtag}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Once any problems have been corrected, it may be possible to create the dependency between {tag} and {childtag} manually.</p>
122619	ERROR	Dependency creation between Oracle database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	<p><b>Cause:</b> LifeKeeper was unable to create a dependency between the database resource {tag} and the necessary child resource {childtag}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Once any problems have been corrected, it may be possible to create the dependency between {tag} and {childtag} manually.</p>
122625	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The quickCheck process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages. Correct any reported problems.</p>
122626	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The remove process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages.</p>

Code	Severity	Message	Cause/Action
			Correct any reported problems.
122627	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The restore process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages. Correct any reported problems.</p>
122628	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The recover process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages. Correct any reported problems.</p>
122632	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> During a remove, the resource instance {sid} passed to the remove process does not match internal resource instance information for the {sid}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122633	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> During a restore, the resource instance {sid} passed to restore does not match internal resource instance information for the {sid}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122634	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> During resource recovery, the r</p>

Code	Severity	Message	Cause/Action
			<p>resource instance {sid} passed to recovery does not match internal resource instance information for the {sid}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122636	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	<p><b>Cause:</b> The create of the Oracle resource hierarchy {tag} failed on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122638	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The create action for the Oracle database resource {tag} on server {server} failed. The signal {sig} was received by the create process.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122640	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122641	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>

Code	Severity	Message	Cause/Action
			Correct any reported errors.
122642	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122643	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
122644	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to retrieve resource instance information for {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors and retry the extend.</p>
122645	ERROR	Cannot access canextend script "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to run pre-extend checks because it was unable to find the "canextend" script on {server} for a dependent child resource.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122646	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to retrieve resource instance information for {tag} on {server}.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors and retry the extend.</p>
122647	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122648	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server "%s" for App "%s" and Type "%s"	<p><b>Cause:</b> During the database resource extension, a resource instance was found using the same {tag} and/or {id} but with a different resource application and type.</p> <p><b>Action:</b> Resource IDs must be unique. The resource instance with the ID matching the Oracle resource instance must be removed.</p>
122649	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122650	ERROR	Cannot access extend script "%s" on server "%s"	<p><b>Cause:</b> The request to extend the database resource {resource} to {server} failed because it was unable to find the script {extend} on {server} for a dependent child resource.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>

Code	Severity	Message	Cause/Action
122651	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> The request to extend the database resource {resource} to {server} failed because of an error attempting to extend a dependent child resource.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
122654	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The health check for the database {sid} was terminated because the quickCheck process received a signal. This is most likely caused by the quickCheck process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The health check time for an Oracle resource is controlled by the tunable value ORACLE_QUICKCHECK_TIMEOUT. Set it to a value greater than 45 seconds to allow more time for the health check process to complete.</p>
122655	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The request to take database {sid} "Out of Service" was terminated because the remove process received a signal. This is most likely caused by the remove process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The remove time for an Oracle resource is controlled by the tunable value ORACLE_REMOVE_TIMEOUT. Set the tunable to a value greater than 240 seconds to allow more time for the remove process to complete.</p>
122659	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The request to place database {sid} "In Service" was terminated because the restore process received a signal.</p>

Code	Severity	Message	Cause/Action
			<p>I. This is most likely caused by the restore process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The restore time for an Oracle resource is controlled by the tunable value ORACLE_RESTORE_TIMEOUT. Set the tunable to a value greater than 240 seconds to allow more time for the restore process to complete.</p>
122663	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The recovery of the failed database was terminated because the recovery process received a signal. This is most likely caused by the recovery process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The recovery time for an Oracle resource is controlled by the tunable values ORACLE_RESTORE_TIMEOUT and ORACLE_REMOVE_TIMEOUT. Set one or both of these to a value greater than 240 seconds to allow more time for a recovery to complete.</p>
122670	ERROR	Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".	<p><b>Cause:</b> An unexpected error occurred while attempting to update the oratab entry for the database {sid}. The error occurred while attempting to open the temporary file used in the update process.</p> <p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122671	ERROR	Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".	<p><b>Cause:</b> An unexpected error occurred while attempting to update the oratab entry for the database {sid}. The error occurred</p>

Code	Severity	Message	Cause/Action
			<p>red while attempting to close the temporary file used in the update process.</p> <p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122672	ERROR	Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".	<p><b>Cause:</b> An unexpected error occurred while attempting to update the oratab entry for the database {sid}. The error occurred while attempting to rename the temporary file back to oratab.</p> <p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122673	ERROR	Unable to log messages queued while running as oracle user %s on %s. Reason: \$!	<p><b>Cause:</b> An unexpected error {reason} occurred while attempting to add messages to the log file. These messages were generated while running as the Oracle user.</p> <p><b>Action:</b> Review the reason for the failure and take corrective action.</p>
122674	ERROR	Unable to open %s Reason: %s.	<p><b>Cause:</b> An unexpected error occurred while attempting to open a connection to the Oracle database and run the database {cmd}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Additionally, check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct any reported problems.</p>

Code	Severity	Message	Cause/Action
122680	ERROR	Unable to find Oracle home for "%s" on "%s".	<p><b>Cause:</b> The Oracle home directory {Oracle home} does not appear to contain files necessary for the proper operation of the Oracle instance {sid}.</p> <p><b>Action:</b> Verify using the command line that the Oracle home directory {Oracle home} contains the Oracle binaries, a valid spfile{sid}.ora or init{sid}.ora file.</p>
122681	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122682	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The required Oracle executable {exe} was not found on this server {server}.</p> <p><b>Action:</b> Verify that the Oracle binaries and required software utilities are installed and properly configured on the server {server}. The Oracle binaries must be installed locally on each node or located on shared storage available to all nodes in the cluster.</p>
122683	ERROR	Backup node %s is unreachable; abort username/password changes.	<p><b>Cause:</b> Backup node {node} is currently unreachable. Pending changes to username/password have been aborted.</p> <p><b>Action:</b> Ensure that SIOS LifeKeeper is running on the given backup node {node} and that all communication paths are up, then retry the operation.</p>

Code	Severity	Message	Cause/Action
122684	ERROR	The restore of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_RESTORE_TIMEOUT in /etc/default/LifeKeeper.	<p><b>Cause:</b> The restore for Oracle resource {resource} has timed out on server {server}. This occurs when the in-service operation has taken longer than the time allotted.</p> <p><b>Action:</b> The restore time for an Oracle resource is controlled by the tunable value ORACLE_RESTORE_TIMEOUT. Set the tunable to a value greater than 240 seconds to allow more time for the restore process to complete.</p>
122685	ERROR	The remove of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_REMOVE_TIMEOUT in /etc/default/LifeKeeper.	<p><b>Cause:</b> The remove for Oracle resource {resource} has timed out on server {server}. This occurs when the out-of-service operation has taken longer than the time allotted.</p> <p><b>Action:</b> The remove time for an Oracle resource is controlled by the tunable value ORACLE_REMOVE_TIMEOUT. Set the tunable to a value greater than 240 seconds to allow more time for the remove process to complete.</p>
122686	ERROR	The quickCheck of %s has timed out on server %s. The default TIMEOUT is 45 seconds. To increase the TIMEOUT, set ORACLE_QUICKCHECK_TIMEOUT in /etc/default/LifeKeeper.	<p><b>Cause:</b> The quickCheck for Oracle resource {resource} has timed out on server {server}. This occurs when the health check has taken longer than the time allotted.</p> <p><b>Action:</b> The health check time for an Oracle resource is controlled by the tunable value ORACLE_QUICKCHECK_TIMEOUT. Set it to a value greater than 45 seconds to allow more time for the health check process to complete.</p>
122687	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified.</p>

Code	Severity	Message	Cause/Action
			<p>ed for the getalertlog script.</p> <p><b>Action:</b> The tag name of the Oracle resource must be provided to the script as the first command line parameter. Verify the parameters are correct and retry the operation.</p>
122702	ERROR	Failed start OHAS.	
122704	ERROR	Failed start ASM.	
122706	ERROR	Failed start Diskgroup.	
122708	ERROR	Failed stop Diskgroup.	
122710	ERROR	Failed stop ASM.	
122712	ERROR	Failed stop OHAS.	
122713	ERROR	OHAS is not running.	
122714	ERROR	ASM is not running.	
122715	ERROR	Diskgroup is not running.	
122716	ERROR	\$usage	
122717	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" does not exist	
122718	ERROR	Cannot extend resource \"\$template_tag\" to server \"\$me\"	
122719	ERROR	\$usage	
122720	ERROR	END failed create of \"\$Tag\" due to a \"\$sig\" signal	
122722	ERROR	Unable to getlocks on \$me during resource create.	
122725	ERROR	Error creating resource \$Tag. Error (\$rc	
122726	ERROR	END failed hierarchy create of resource \$Tag with return value of \$code.	
122728	ERROR	Unable to rlslocks on \$me during resource create.	
122729	ERROR	\$usage	
122730	ERROR	END failed extend of \"\$Tag\" due to a	

Code	Severity	Message	Cause/Action
		\"\$sig\" signal	
122732	ERROR	Template resource \"\$TemplateTag\" on server \"\$TemplateSys\" does not exist	
122733	ERROR	Error creating resource \"\$Tag\" on server \"\$me\"	
122735	ERROR	END failed hierarchy extend of resource \$Tag with return value of \$ecode.	
122736	ERROR	\$usage	
122737	ERROR	Failed to create object instance for Oracle ASM on \$oracle::me	
122738	ERROR	Usage : \$cmd \$usage	
122739	ERROR	Failed to create object instance for Oracle ASM on \"\$me\".	
122740	ERROR	\$usage	
122742	ERROR	backup node \$back_dead_name is unreachable; abort protection Diskgroup changes.	
122744	ERROR	Update of protection Diskgroup failed for \"\$tag\" on \"\$oracleasm::me\"	
122746	ERROR	Update of protection Diskgroup failed for \"\$tag\" on \"\$remote\"	
122750	ERROR	Failed to init the object.	
122751	ERROR	Usage: %s %s	
123006	FATAL	Unknown version %s of IP address	<p><b>Cause:</b> The IP address does not appear to be valid for either IPv4 or IPv6.</p> <p><b>Action:</b> Provide a valid IP address.</p>
123008	ERROR	No pinglist found for %s.	<p><b>Cause:</b> Problem while opening the pinglist for this IP address.</p> <p><b>Action:</b> Make sure you have provided a pinglist for this IP address.</p>

Code	Severity	Message	Cause/Action
123009	ERROR	List ping test failed for virtual IP %s	<p><b>Cause:</b> No response was received from any of the addresses in the ping list.</p> <p><b>Action:</b> Check network connectivity of this node and the systems on which the IPs in the ping list reside.</p>
123013	ERROR	Link check failed for virtual IP %s on interface %s.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present on the physical layer connection.</p> <p><b>Action:</b> Check the physical connections for the interface and bring the physical layer link up.</p>
123015	ERROR	Link check failed for virtual IP %s on interface %s.	<p><b>Cause:</b> The requested interface is a bonded interface, and one of the slaves is showing 'NO-CARRIER' indicating that no link is present on the physical layer connection.</p> <p><b>Action:</b> Check the physical connections for the slave interface and bring the physical layer link up.</p>
123024	ERROR	IP address seems to still exist somewhere else.	<p><b>Cause:</b> The IP address appears to be in use elsewhere on the network.</p> <p><b>Action:</b> Either select a different IP address to use or locate and disable the current use of this IP address.</p>
123037	ERROR	must specify machine name containing primary hierarchy	<p><b>Cause:</b> Not enough arguments were provided to crelPhier.</p> <p><b>Action:</b> Supply all of the needed arguments to crelPhier.</p>

Code	Severity	Message	Cause/Action
123038	ERROR	must specify IP resource name	<p><b>Cause:</b> Not enough arguments were passed to crelPhier.</p> <p><b>Action:</b> Supply all of the needed arguments to crelPhier.</p>
123039	ERROR	must specify primary IP Resource tag	<p><b>Cause:</b> The argument specifying the primary IP Resource tag was missing from the "crelPhier" command.</p> <p><b>Action:</b> Supply all of the needed arguments.</p>
123042	ERROR	An unknown error has occurred in utility validmask on machine %s.	<p><b>Cause:</b> There was an unexpected error running the "validmask" utility.</p> <p><b>Action:</b> Check adjacent log messages for additional details.</p>
123045	ERROR	An unknown error has occurred in utility getlocks.	<p><b>Cause:</b> There was an unexpected error running the "getlocks" utility.</p> <p><b>Action:</b> Check adjacent log messages for additional details.</p>
123053	ERROR	Cannot resolve hostname %s	<p><b>Cause:</b> A hostname was provided for the IP address, but the system was unable to resolve the name to an IP address.</p> <p><b>Action:</b> Check the correctness of the hostname and verify that name resolution (DNS or /etc/hosts) is working correctly and returns the IP for the hostname.</p>
123055	ERROR	An unknown error has occurred in utility %s on machine %s.	<p><b>Cause:</b> There was a failure while creating the IP resource.</p> <p><b>Action:</b> Check the logs for related error</p>

Code	Severity	Message	Cause/Action
			s and try to resolve the reported problem.
123056	ERROR	create ip hierarchy failure: perform_action failed	<p><b>Cause:</b> Unexpected error trying to restore the IP address during creation.</p> <p><b>Action:</b> Check adjacent log messages for additional details.</p>
123059	ERROR	Resource already exists on machine %s	<p><b>Cause:</b> Attempted to create an IP address that already exists.</p> <p><b>Action:</b> Reuse the existing resource or manually remove the IP address if it exists or use a different IP address.</p>
123060	ERROR	ins_create failed on machine %s	<p><b>Cause:</b> An unexpected failure occurred while creating an IP resource.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
123064	ERROR	An unknown error has occurred in utility %s on machine %s.	<p><b>Cause:</b> There was a failure while creating a dependency for the IP resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
123066	ERROR	An error occurred during creation of LifeKeeper application=comm on %s.	<p><b>Cause:</b> A failure occurred while calling "app_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
123068	ERROR	An error occurred during creation of LifeKeeper resource type=ip on %s.	<p><b>Cause:</b> A failure occurred while calling "t</p>

Code	Severity	Message	Cause/Action
			<p>yp_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
123089	ERROR	Link check failed for virtual IP %s on interface %s.	
123091	ERROR	the link for interface %s is down	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present on the physical layer connection.</p> <p><b>Action:</b> Check the physical connections for the interface and bring the physical layer link up.</p>
123093	ERROR	the ping list check failed	<p><b>Cause:</b> No response was received from any of the addresses in the ping list.</p> <p><b>Action:</b> Check network connectivity of this node and the systems on which the IPs in the ping list reside.</p>
123094	ERROR	IP address is not assigned to interface %s.	
123095	ERROR	broadcast ping failed	<p><b>Cause:</b> No replies were received from a broadcast ping.</p> <p><b>Action:</b> Verify that at least one host on the subnet will respond to broadcast pings. Verify that virtual IP is on the correct network interface. Consider using a ping list instead of a broadcast ping.</p>
123096	ERROR	\$msg	<p><b>Cause:</b> The broadcast ping used to determine the viability of the virtual IP failed.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Please ensure that the ping list for this resource is properly configured in the properties panel or that broadcast ping checking is disabled by adding NOBCASTPING=1 to the /etc/default/LifeKeeper configuration file.</p>
123097	ERROR	exec_list_ping(): broadcast ping failed.	<p><b>Cause:</b> The broadcast ping used to determine the viability of the virtual IP failed.</p> <p><b>Action:</b> Ensure that the ping list for this resource is properly configured in the properties panel or that broadcast ping checking is disabled by adding NOBCASTPING=1 to the /etc/default/LifeKeeper configuration file.</p>
123299	ERROR	Unable to open %s. Reason %s	
123410	ERROR	Usage error OSUquickCheck	
123411	ERROR	OSUquickCheck: both tag and id name not specified	
123412	ERROR	resource \$Tag not found on local server	
123414	ERROR	The link for network interface \$IPObj->{'device'} is down	
123415	ERROR	No pinglist found for \$IPObj->{'ipaddr'}	
123416	ERROR	List ping test failed for virtual IP \$IPObj->{'ipaddr'}	
124004	FATAL	resource tag name not specified	<p><b>Cause:</b> Invalid arguments were specified for the "quickCheck" operation.</p> <p><b>Action:</b> Ensure that the correct arguments are passed.</p>
124005	FATAL	resource id not specified	<p><b>Cause:</b> Invalid arguments were specified for the "quickCheck" operation.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Ensure that the correct arguments are passed.
124007	FATAL	Failed to get resource information	<p><b>Cause:</b> The filesystem resource's info field does not contain the correct information.</p> <p><b>Action:</b> Put the correct information in the resource's info field or restore the system from a recent "lkbackup" to restore the original info field.</p>
124008	ERROR	getld failed	<p><b>Cause:</b> The filesystem resource could not find the underlying disk device.</p> <p><b>Action:</b> Check adjacent log messages for further details. Verify that the resource hierarchy is valid and that all required storage kits are installed.</p>
124009	ERROR	LifeKeeper protected filesystem is in service but quickCheck detects the following error	<p><b>Cause:</b> The filesystem kit has found something wrong with the resource.</p> <p><b>Action:</b> Check the messages immediately following this one for more details.</p>
124010	ERROR	\"\$id\" is not mounted	<p><b>Cause:</b> The filesystem resource is no longer mounted.</p> <p><b>Action:</b> No action is required. Allow local recovery to remount the resource.</p>
124011	ERROR	\"\$id\" is mounted but with the incorrect mount options (current mount option list: \$mntopts, expected mount option list: \$infopts	<p><b>Cause:</b> The filesystem resource is mounted incorrectly.</p> <p><b>Action:</b> No action is required. Allow local recovery to remount the resource.</p>

Code	Severity	Message	Cause/Action
124012	ERROR	"\$id\" is mounted but on the wrong device (current mount device: \$tmpdev, expected mount device: \$dev	<p><b>Cause:</b> The filesystem resource has the wrong device mounted.</p> <p><b>Action:</b> No action is required. Allow local recovery to remount the resource.</p>
124015	ERROR	LifeKeeper protected filesystem \"\$tag\" (\$id) is \$percent% full (\$blocksfree free blocks).	<p><b>Cause:</b> The filesystem is getting full.</p> <p><b>Action:</b> Remove or migrate data from the filesystem.</p>
124016	WARN	LifeKeeper protected filesystem \"\$tag\" (\$id) is \$percent% full (\$blocksfree free blocks).	<p><b>Cause:</b> The filesystem is getting full.</p> <p><b>Action:</b> Remove or migrate data from the filesystem.</p>
124020	FATAL	cannot find device information for filesystem \$id	<p><b>Cause:</b> The filesystem resource could not find the underlying disk device.</p> <p><b>Action:</b> Check adjacent log messages for further details. Verify that the resource hierarchy is valid and that all required storage kits are installed.</p>
124029	ERROR	Failed to find child resource.	<p><b>Cause:</b> The filesystem resource could not determine its underlying disk resource.</p> <p><b>Action:</b> Ensure that the resource hierarchy is correct.</p>
124032	FATAL	Script has hung. Exiting.	<p><b>Cause:</b> Processes had files open on a mounted filesystem that needed to be unmounted. Killing those processes has taken too long.</p> <p><b>Action:</b> If this error continues, try to temporarily stop all software that may be</p>

Code	Severity	Message	Cause/Action
			using the mount point to allow it to be unmounted. If the filesystem still cannot be unmounted, contact Support.
124042	ERROR	file system \$fs failed unmount; will try again	<p><b>Cause:</b> Processes had files open on a mounted filesystem that needed to be unmounted. It can take multiple attempts to clear those processes.</p> <p><b>Action:</b> No action is required. Allow the process to continue.</p>
124046	ERROR	file system \$fsname failed unmount	<p><b>Cause:</b> A filesystem could not be unmounted.</p> <p><b>Action:</b> If this error continues, try to temporarily stop all software that may be using the mount point to allow it to be unmounted. If the filesystem still cannot be unmounted, contact Support.</p>
124049	ERROR	Local recovery of resource has failed (error=\$err)	<p><b>Cause:</b> A filesystem resource has a problem that cannot be repaired locally.</p> <p><b>Action:</b> No action is required. Allow the resource to be failed over to another system.</p>
124051	WARN	getld failed, try count : \$cnt/\$try	
124052	ERROR	"\$id" is mounted but filesystem is shutdown state.	
124054	ERROR	Failed to change mount option from "\$old_opts" to "\$new_opts" for migration	
124103	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Check adjacent log messages for</p>

Code	Severity	Message	Cause/Action
			or further details.
124104	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Invalid arguments were specified for the canextend operation.</p> <p><b>Action:</b> Ensure that the arguments are correct. If this error happens during normal operation, please contact Support.</p>
124105	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Invalid arguments were specified for the canextend operation.</p> <p><b>Action:</b> Ensure that the arguments are correct. If this error happens during normal operation, please contact Support.</p>
124106	ERROR	\$ERRMSG Unable to access template resource \"\$TemplateName\	<p><b>Cause:</b> The resource's underlying disk information cannot be determined.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template system before extending.</p>
124107	ERROR	\$ERRMSG Resource \"\$TemplateName\" must have one and only one device resource dependency	<p><b>Cause:</b> The resource has too many underlying devices in the hierarchy.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template system before extending.</p>
124108	ERROR	\$ERRMSG Unable to access template resource \"\$TemplateName\	<p><b>Cause:</b> The resource cannot be found on the template system.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template system before extending.</p>

Code	Severity	Message	Cause/Action
124109	ERROR	\$ERRMSG Can not access canextend for scsi/\$DeviceResType resources on machine \"\${TargetSysName}	<p><b>Cause:</b> The target system is missing some required components.</p> <p><b>Action:</b> Ensure that the target system has all the correct kits installed and licensed.</p>
124110	ERROR	\$ERRMSG Either filesystem \"\${TemplateLKId}\" is not mounted on \"\${TemplateSysName}\" or filesystem is not shareable with \"\${TargetSysName}	<p><b>Cause:</b> The filesystem isn't in service on the template system or doesn't meet the requirements for extending to the target system.</p> <p><b>Action:</b> Make sure the resource is in service on the template system and review the product documentation regarding the requirements for extending filesystems.</p>
124111	ERROR	\$ERRMSG File system type \"\${FSType}\" is not supported by the kernel currently running on \"\${TargetSysName}	<p><b>Cause:</b> The filesystem's type cannot be mounted on the target system due to lack of kernel support.</p> <p><b>Action:</b> Ensure that the target system has all its kernel modules configured correctly before extending the resource.</p>
124112	ERROR	must specify machine name containing primary hierarchy	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>
124113	ERROR	must specify primary ROOT tag	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>

Code	Severity	Message	Cause/Action
124114	ERROR	must specify primary mount point	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>
124115	ERROR	must specify primary switchback type	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>
124118	ERROR	dep_remove failure on machine \"'\$PRIMACH'\" for parent \"'\$PRITAG'\" and child \"'\$DEVTAG.\\	<p><b>Cause:</b> Cleanup after a dependency creation failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124119	ERROR	ins_remove failure on machine \"'\$PRIMACH'\" for \"'\$PRITAG.\\	<p><b>Cause:</b> Cleanup after an instance creation failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124121	ERROR	ins_remove failure on machine \"'\$PRIMACH'\"	<p><b>Cause:</b> Cleanup after a resource creation failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124122	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124123	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Invalid arguments were specified</p>

Code	Severity	Message	Cause/Action
			<p>d for the depstoextend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124124	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Invalid arguments were specified for the depstoextend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124125	ERROR	\$ERRMSG Unable to access template resource \"\$TemplateName\	<p><b>Cause:</b> The resource was unable to locate its underlying disk resource.</p> <p><b>Action:</b> Ensure the hierarchy and all dependencies are correct before extending.</p>
124126	ERROR	unextmgr failure on machine \"\$PRIMACH\"	<p><b>Cause:</b> The cleanup, after a failed resource extend operation, failed.</p> <p><b>Action:</b> Manually clean up any remaining resources and check adjacent log messages for further details.</p>
124128	ERROR	unextmgr failure on machine \"\$PRIMACH\" for \"\$PRITAG.\	<p><b>Cause:</b> The cleanup, after a failed resource extend operation, failed.</p> <p><b>Action:</b> Manually clean up any remaining resources and check adjacent log messages for further details.</p>
124129	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Look for additional log message</p>

Code	Severity	Message	Cause/Action
			s for more details.
124130	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Invalid arguments were specified for the extend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124131	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Invalid arguments were specified for the extend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124132	ERROR	\$ERRMSG Required target mount point is null	<p><b>Cause:</b> Invalid arguments were specified for the extend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124133	ERROR	\$ERRMSG Unable to access template resource \"\$TemplateName\	<p><b>Cause:</b> The tag being extended doesn't exist on the template system.</p> <p><b>Action:</b> Ensure that the hierarchy is correct on the template system before extending.</p>
124134	ERROR	\$ERRMSG Detected conflict in expected tag name \"\$TargetTagName\" on target machine.	<p><b>Cause:</b> A resource already exists on the target system with the same tag as the resource being extended.</p> <p><b>Action:</b> Recreate one of the conflicting resources with a different tag.</p>

Code	Severity	Message	Cause/Action
124135	ERROR	\$ERRMSG Resource \"TemplateTagName\" does not have required device resource dependency or unable to access this resource on template machine.	<p><b>Cause:</b> The resource or its underlying disk resource cannot be found on the template system.</p> <p><b>Action:</b> Ensure that the hierarchy is correct on the template system before extending.</p>
124136	ERROR	\$ERRMSG Resource \"TemplateTagName\" must have one and only one device resource dependency	<p><b>Cause:</b> The resource has multiple underlying devices in the hierarchy on the template system.</p> <p><b>Action:</b> Ensure the hierarchy is correct before extending and that the filesystem resource only depends on a single disk resource.</p>
124137	ERROR	\$ERRMSG Can not access extend for scsi/DeviceResType resources on machine \"TargetSysName\"	<p><b>Cause:</b> The files required to support the given storage type aren't available on the target system.</p> <p><b>Action:</b> Ensure that the required kits are installed on the target system and licensed.</p>
124138	ERROR	\$ERRMSG Unable to access target device resource \"DeviceTagName\" on machine \"TargetSysName\"	<p><b>Cause:</b> The required underlying disk resource doesn't exist on the target system.</p> <p><b>Action:</b> Check adjacent log messages for further details and ensure that the target system is properly configured for hosting the resources being extended.</p>
124141	ERROR	\$ERRMSG Unable to find mount point \"TemplateLKId\" mode on template machine	<p><b>Cause:</b> The details of the mount point on the template system cannot be determined.</p> <p><b>Action:</b> Ensure that the resource is in s</p>

Code	Severity	Message	Cause/Action
			ervice and accessible on the template system before extending.
124142	ERROR	\$ERRMSG Unable to create or access mount point \"\${TargetLKId}\" on target machine	<p><b>Cause:</b> The mount point could not be created on the target system.</p> <p><b>Action:</b> Ensure that the mount point's parent directory exists and is accessible on the target system.</p>
124143	ERROR	\$ERRMSG Two or more conflicting entries found in /etc/fstab on \"\${TargetSysName}\	<p><b>Cause:</b> The device or mount point appears to be mounted more than once on the target system.</p> <p><b>Action:</b> Ensure that the mount point is not mounted on the target system before extending.</p>
124144	ERROR	\$ERRMSG Failed to create resource instance on \"\${TargetSysName}\	<p><b>Cause:</b> The resource creation on the target system failed.</p> <p><b>Action:</b> Check adjacent log messages for further details. Make sure to check the logs on the target server.</p>
124145	ERROR	\$ERRMSG Failed to set resource instance state for \"\${TargetTagName}\" on \"\${TargetSysName}\	<p><b>Cause:</b> The source state could not be changed to OSU on the target system.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124146	ERROR	must specify machine name containing primary hierarchy	<p><b>Cause:</b> Invalid arguments were specified for the filesyshier operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>

Code	Severity	Message	Cause/Action
124147	ERROR	must specify primary mount point	<p><b>Cause:</b> Invalid arguments were specified for the filesyshier operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124149	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The process of finding the resource instance failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124150	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The system failed to read the /etc/mtab file.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124152	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The underlying disk resource could not be found.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124153	ERROR	create file system hierarchy failure	<p><b>Cause:</b> Creating the filesystem resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124154	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The info field for the resource could not be updated.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>

Code	Severity	Message	Cause/Action
124155	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The switchback strategy could not be set on the resource.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124157	ERROR	create file system hierarchy failure \(\conflicting entries in /etc/fstab\	<p><b>Cause:</b> The mount point could not be removed from the /etc/fstab file.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124160	ERROR	Unknown error in script filesysins, err=\$err	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124161	ERROR	create filesys instance – existid – failure	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124163	ERROR	create filesys instance – ins_list – failure	<p><b>Cause:</b> Checking for an existing resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124164	ERROR	create filesys instance – newtag – failure	<p><b>Cause:</b> The system failed to generate a suggested tag for the resource.</p> <p><b>Action:</b> If this error happens during normal operation, contact Support.</p>
124168	ERROR	create filesys instance – ins_create – failure	<p><b>Cause:</b> The filesystem resource could not</p>

Code	Severity	Message	Cause/Action
			<p>ot be created.</p> <p><b>Action:</b> Check adjacent log messages f or further details.</p>
124169	ERROR	filesys instance – ins_setstate – failure	<p><b>Cause:</b> The new filesystem resource's s tate could not be initialized.</p> <p><b>Action:</b> Check adjacent log messages f or further details.</p>
124173	ERROR	create filesys instance – dep_create – fail ure	<p><b>Cause:</b> The resource's dependency rela tionship with its underlying disk could n ot be created.</p> <p><b>Action:</b> Check adjacent log messages f or further details.</p>
124174	ERROR	machine not specified	<p><b>Cause:</b> Invalid arguments were specifie d for the rmenu_mp operation.</p> <p><b>Action:</b> Ensure the script is called corre ctly. If this error happens during normal operation, please contact Support.</p>
124175	ERROR	mount point not specified	<p><b>Cause:</b> Invalid arguments were specifie d for the rmenu_mp operation.</p> <p><b>Action:</b> Ensure the script is called corre ctly. If this error happens during normal operation, please contact Support.</p>
124177	ERROR	unexpected multiple matches found	<p><b>Cause:</b> One or more systems show a fil esystem or mount point used more than once.</p> <p><b>Action:</b> Verify filesystem devices and m ount points and ensure that filesystems</p>

Code	Severity	Message	Cause/Action
			are only mounted once. Look for additional log messages for more details.
124178	ERROR	machine name not specified	<p><b>Cause:</b> Invalid arguments were specified for the rmenump operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124180	ERROR	must specify filesystem type	<p><b>Cause:</b> Invalid arguments were specified for the validfstype operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124181	ERROR	mount point not specified	<p><b>Cause:</b> Invalid arguments were specified for the validmp operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124182	ERROR	The mount point \$MP is not an absolute path	<p><b>Cause:</b> A mount point was specified that isn't an absolute path (doesn't start with a '/').</p> <p><b>Action:</b> Specify a mount point as an absolute path starting with a '/'.</p>
124183	ERROR	\$MP is already mounted on \$MACH	<p><b>Cause:</b> The requested mount point is already in use on the system.</p> <p><b>Action:</b> Specify a mount point that isn't in use or unmount it before retrying the operation.</p>

Code	Severity	Message	Cause/Action
124184	ERROR	The mount point \$MP is already protected by LifeKeeper on \$MACH	<p><b>Cause:</b> The system is already protecting the specified mount point.</p> <p><b>Action:</b> Choose a different mount point that isn't already being protected.</p>
124185	ERROR	The mount point \$MP is not a directory on \$MACH	<p><b>Cause:</b> The mount point refers to a non-directory such as a regular file.</p> <p><b>Action:</b> Choose a mount point that refers to a directory.</p>
124186	ERROR	The mount point directory \$MP is not empty on \$MACH	<p><b>Cause:</b> The specified mount point refers to a directory that isn't empty.</p> <p><b>Action:</b> Choose a mount point that is empty or remove the contents of the specified directory before retrying the operation.</p>
124187	ERROR	server name not specified	<p><b>Cause:</b> Invalid arguments were specified for the valuepmp operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124188	ERROR	There are no mount points on server \$MACH	<p><b>Cause:</b> There are no possible mount points for filesystem resource on the server.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124194	WARN	Please correct conflicting \"/etc/fstab\" entries on server \$UNAME for: \$FSDEV, \$FSNAME	<p><b>Cause:</b> After deleting a filesystem resource, some entries in /etc/fstab need to be manually cleaned up.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Manually clean up the /etc/fstab file.
124195	ERROR	getchildinfo found no \$OKAPP child for \$PTAG	<b>Cause:</b> The system could not find a child resource in the hierarchy.  <b>Action:</b> Check adjacent log messages for further details and ensure that the hierarchy is correct before retrying the operation.
124196	ERROR	enablequotas – quotacheck may have failed for \$FS_NAME	<b>Cause:</b> The quota operation failed.  <b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and in /var/log/messages.
124198	ERROR	enablequotas – quotaon failed to turn on quotas for \$FS_NAME, reason	<b>Cause:</b> The quota operation failed.  <b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and /var/log/messages.
124200	ERROR	The device node \$dev was not found or did not appear in the udev create time limit of \$delay seconds	<b>Cause:</b> A device node (/dev/...) was not created by udev. This may indicate an issue with the storage or the server's connection to the storage.  <b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and in /var/log/messages.
124201	WARN	Device \$device not found. Will retry wait to see if it appears.	<b>Cause:</b> This can happen under normal conditions while udev creates device node entries for storage. This message should not happen repeatedly.  <b>Action:</b> Check adjacent log messages for

Code	Severity	Message	Cause/Action
			<p>or further details in both the lifekeeper log and in /var/log/messages.</p>
124202	ERROR	<p>Command \"\$commandwithargs\" failed. Retrying ....</p>	<p><b>Cause:</b> The given command failed but may have failed temporarily. This failure may happen during normal operations but should not keep failing.</p> <p><b>Action:</b> Check adjacent log messages for further details if this message continues.</p>
124204	WARN	<p>cannot make file system \$FSNAME mount point</p>	<p><b>Cause:</b> The mount point directory could not be created.</p> <p><b>Action:</b> Ensure that the mount point can be created. This may be due to filesystem permissions, mount options, etc.</p>
124207	ERROR	<p>\"fsck\"ing file system \$FSNAME failed, trying alternative superblock</p>	<p><b>Cause:</b> This message indicates that the typical filesystem check failed. This message may be ok for ext2 filesystems or other filesystems where an alternative superblock location is used.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124209	ERROR	<p>\"fsck\"ing file system \$FSNAME with alternative superblock failed</p>	<p><b>Cause:</b> This indicates that an ext2 filesystem (or other filesystem where an alternative superblock location is used) check failed with the alternative superblock location.</p> <p><b>Action:</b> Check adjacent log messages for further details and instructions on how to proceed.</p>

Code	Severity	Message	Cause/Action
124210	WARN	POSSIBLE FILESYSTEM CORRUPTION ON \$FSNAME (\$FPNAME)	<p><b>Cause:</b> A filesystem was put in service or failed over when it was out of sync with its mirror source.</p> <p><b>Action:</b> Check adjacent log messages for further details and review the product documentation for information on how to bring the filesystem in service safely.</p>
124211	ERROR	Reason for fsck failure (\$retval): \$ret	<p><b>Cause:</b> This log message is part of a series of messages and gives the actual exit code from the fsck process.</p> <p><b>Action:</b> Check adjacent log messages for further details and instructions on how to proceed.</p>
124212	ERROR	"fsck" of file system \$FSNAME failed	<p><b>Cause:</b> The check of the filesystem failed. This is usually due to the filesystem having corruption.</p> <p><b>Action:</b> Check adjacent log messages for further details. Review the product documentation for instructions on how to handle possible filesystem corruption.</p>
124213	WARN	POSSIBLE FILESYSTEM CORRUPTION ON \$FSNAME (\$FPNAME)	<p><b>Cause:</b> The system or user tried to bring into service a filesystem that may be corrupted. This can happen if a filesystem is switched or failed over when it was out of sync with its mirror source.</p> <p><b>Action:</b> Check adjacent log messages for further details and review the product documentation for instructions on how to bring the resource into service safely.</p>
124214	ERROR	Reason for fsck failure (\$retval)	<p><b>Cause:</b> This message should follow a previous log message about a filesystem</p>

Code	Severity	Message	Cause/Action
			<p>check failure and gives the process exit code of the fsck process.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124218	ERROR	File system \$FSNAME was found to be already	<p><b>Cause:</b> This message is part of a series of messages.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124219	ERROR	mounted after initial mount attempt failed.	<p><b>Cause:</b> This message is part of a series of messages. This should not happen under normal circumstances but may not be fatal if the resource can be put in service.</p> <p><b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and in /var/log/messages.</p>
124220	ERROR	File system \$FSNAME failed to mount.	<p><b>Cause:</b> The filesystem could not be mounted.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124221	WARN	Protected Filesystem \$ID is full	<p><b>Cause:</b> The filesystem is full.</p> <p><b>Action:</b> Remove unused data from the filesystem or migrate to a larger filesystem.</p>
124222	WARN	Dependent Applications may be affected <>	<p><b>Cause:</b> This indicates that an operation on a resource is likely to cause operation on other resources based on the resource hierarchy.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Make sure it's acceptable for the indicated resources to be affected before continuing.
124223	ERROR	Put \"\${t}\" Out-Of-Service Failed By Signal	<b>Cause:</b> This message should not occur under normal circumstances. <b>Action:</b> Check adjacent log messages for further details.
124227	ERROR	Put \"\${i}\" Out-Of-Service Failed	<b>Cause:</b> The operation failed. <b>Action:</b> Check adjacent log messages for further details.
124230	ERROR	Put \"\${t}\" In-Service Failed By Signal	<b>Cause:</b> This message should not occur under normal circumstances. <b>Action:</b> Check adjacent log messages for further details.
124231	ERROR	Put \"\${t}\" In-Service Failed	<b>Cause:</b> The operation failed. <b>Action:</b> Check adjacent log messages for further details.
124234	ERROR	Put \"\${t}\" In-Service Failed	<b>Cause:</b> The operation failed. <b>Action:</b> Check adjacent log messages for further details.
125102	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125103	ERROR	`printf '%s is not shareable with any machine.' \$DEV`	<b>Cause:</b> The device does not appear to

Code	Severity	Message	Cause/Action
			<p>be shared with any other systems.</p> <p><b>Action:</b> Verify that the device is accessible from all servers in the cluster. Ensure that all relevant storage drivers and software are installed and configured properly.</p>
125104	ERROR	`printf 'Failed to create disk hierarchy for "%s" on "%s" \$PRIMACH \$DEV`	<p><b>Cause:</b> The creation of a resource to protect a physical disk failed.</p> <p><b>Action:</b> Check adjacent log messages for more details and try to resolve the reported problem.</p>
125107	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p>
125114	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p>
125120	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p>
125123	ERROR	`printf 'Cannot access depstoextend script "%s" on server "%s" \$depstoextend \$TargetSysName`	<p><b>Cause:</b> LifeKeeper was unable to run pre-extend checks on the resource hierarchy because it was unable to find the script "DEPSTOEXTEND" on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
125126	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$ChildTag \$TemplateSysName`	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p>

Code	Severity	Message	Cause/Action
125129	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateTagName \$TemplateSysName`	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125155	ERROR	SCSI \$DEV failed to lock.	<b>Cause:</b> There was a problem locking a SCSI device. <b>Action:</b> Check adjacent log messages for more details and try to resolve the reported problem.
125164	ERROR	SCSI \$INFO failed to unlock.	<b>Cause:</b> There was a problem unlocking a SCSI device. <b>Action:</b> Check adjacent log messages for more details and try to resolve the reported problem.
125181	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateTag \$TemplateSysName`	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125194	ERROR	Failed to check disk.(tag="\\$opt_t")	
126105	ERROR	script not specified – \$PTH is a directory	<b>Cause:</b> The specified script path is a directory. <b>Action:</b> Correct the path of the script.
126110	ERROR	script \$PTH does not exist	<b>Cause:</b> The specified script path does not exist. <b>Action:</b> Correct the path of the script.
126115	ERROR	script \$PTH is a zero length file	<b>Cause:</b> The specified script is an empty file. <b>Action:</b> Correct the script's file path and check the contents inside the script.

Code	Severity	Message	Cause/Action
126117	ERROR	script \$PTH is not executable	<p><b>Cause:</b> The specified script is not executable.</p> <p><b>Action:</b> Correct the script's file path, check the contents inside the script file and make sure it has the proper execute permissions.</p>
126125	ERROR	required template machine name is null	<p><b>Cause:</b> The input template machine name is null.</p> <p><b>Action:</b> Correct the input template machine name.</p>
126130	ERROR	required template resource tag name is null	<p><b>Cause:</b> The input template resource {tag} is null.</p> <p><b>Action:</b> Correct the input template resource tag name.</p>
126135	ERROR	Unable to generate a new tag	<p><b>Cause:</b> Failed to generate a new tag as the same as the template tag name on the target node using the "newtag" script during the extension. The tag name is already existing.</p> <p><b>Action:</b> Avoid using duplicate tag name on the same node and check the log for detail.</p>
126140	ERROR	Unable to generate a new tag	<p><b>Cause:</b> Failed to generate a new tag as input target tag name on the target node using the "newtag" script during the extension. The tag name is already existing.</p> <p><b>Action:</b> Avoid using duplicate tag name on the same node and check log for detail.</p>

Code	Severity	Message	Cause/Action
126150	ERROR	unable to remote copy template \"\$_lscript\" script file	<p><b>Cause:</b> Failed to remote copy template script file. The cause may be due to the non-existence/availability of template script file on template node or any transaction failure during "lcdrp" process.</p> <p><b>Action:</b> Check the existence/availability of template script and the connection to template node. Also check the logs for related errors and try to resolve the reported problem.</p>
126155	ERROR	failed to create resource instance on \"TargetSysName\"	<p><b>Cause:</b> Failed to create resource instance using "ins_create".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126160	ERROR	failed to set resource instance state for \"TargetTagName\" on \"TargetSysName\"	<p><b>Cause:</b> Failed to set resource instance state using "ins_setstate" during GenApp resource extension.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126170	ERROR	getlocks failure	<p><b>Cause:</b> Failed to get the administrative lock when creating a resource hierarchy.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126175	ERROR	instance create failed	<p><b>Cause:</b> Failed to create a GenApp instance using "appins".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
			m.
126180	ERROR	unable to set state to OSU	<p><b>Cause:</b> Failed to set resource instance state using "ins_setstate" during GenApp resource creation.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126190	ERROR	resource restore has failed	<p><b>Cause:</b> Failed to restore GenApp resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126200	ERROR	create application hierarchy rlslocks failure	<p><b>Cause:</b> Failed to release lock after GenApp resource created.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126210	ERROR	copy \$ltype script \$lscript failure	<p><b>Cause:</b> Failed to copy user provided script to appropriate GenApp directory during resource creation.</p> <p><b>Action:</b> Check the existence/availability of user provided script and the GenApp directory as well. Also check the logs for related errors and try to resolve the reported problem.</p>
126215	ERROR	no \$ltype script specified	<p><b>Cause:</b> Missing user defined script during GenApp resource creation.</p> <p><b>Action:</b> Check the input action script an</p>

Code	Severity	Message	Cause/Action
			d run resource creation again.
126220	ERROR	no machine name specified	<p><b>Cause:</b> Missing specified machine name during GenApp resource creation. Failed to copy specified user script due to missing the input for machine name.</p> <p><b>Action:</b> Check the input for machine name and run resource creation again.</p>
126225	ERROR	no resource tag specified	<p><b>Cause:</b> Missing specified tag name during resource creation.</p> <p><b>Action:</b> Check the input for source tag name and run resource creation again.</p>
126230	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> Failed to extend GenApp resource due to unknown reason.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126235	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Missing the input for template machine name during GenApp resource extension.</p> <p><b>Action:</b> Check the input for template machine name and do the resource extension again.</p>
126240	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Missing the input for template resource tag name during GenApp resource extension.</p> <p><b>Action:</b> Check the input for template resource tag name and do the resource extension again.</p>

Code	Severity	Message	Cause/Action
126245	ERROR	\$ERRMSG Can not access extend for \$AppType/\$ResType resources on machine \"\$TargetSysName\	<p><b>Cause:</b> Failed to locate "extend" script during GenApp resource extension on target node.</p> <p><b>Action:</b> Check the existence/availability of "extend" script and do GenApp resource extension again.</p>
126250	ERROR	create application failure – ins_list failed	<p><b>Cause:</b> Failed when calling "ins_list" during GenApp resource creation.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126255	ERROR	create application failure – unable to generate a new tag	<p><b>Cause:</b> Failed to generate a new tag during the GenApp resource creation.</p> <p><b>Action:</b> Avoid using duplicate tag name on the same node. Also check the logs for related errors and try to resolve the reported problem.</p>
126270	ERROR	create application failure – ins_create failed	<p><b>Cause:</b> Failed using "ins_create" to create GenApp instance.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126275	ERROR	create application failure – copy_actions failed	<p><b>Cause:</b> Failed using "copy_actions" to copy user specified template script file.</p> <p><b>Action:</b> Check the existence/availability of template script. Also check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
126290	ERROR	Unable to obtain tag for resource with id \$ID	<p><b>Cause:</b> Failed to fetch GenApp resource tag name by input ID during recovery.</p> <p><b>Action:</b> Check the correctness of input ID and existence/availability of GenApp resource in LCD. Also check the logs for related errors and try to resolve the reported problem.</p>
126300	ERROR	generic application recover script for \$TAG was not found or was not executable	<p><b>Cause:</b> Failed to locate the user defined script for GenApp resource during recovery.</p> <p><b>Action:</b> Check the existence/availability of the user defined script and do the GenApp recovery process again.</p>
126310	ERROR	-t flag not specified	<p><b>Cause:</b> Missing the input for resource tag name during GenApp resource restore.</p> <p><b>Action:</b> Check the input for resource tag name and do resource restore again.</p>
126315	ERROR	-i flag not specified	<p><b>Cause:</b> Missing the input for resource internal id during GenApp resource restore.</p> <p><b>Action:</b> Check the input for resource internal id and do resource restore again.</p>
126327	ERROR	END timeout restore of \"\$TAG\" (forcibly terminating)	
126335	ERROR	restore script \"\$LCDAS/\$APP_RESTOREDIR/\$TAG\" was not found or is not executable	<p><b>Cause:</b> Failed to locate the user defined script for GenApp resource during restore.</p> <p><b>Action:</b> Check the existence/availability of the user defined script and do the GenApp recovery process again.</p>

Code	Severity	Message	Cause/Action
			nApp restore process again.
126340	ERROR	-t flag not specified	<p><b>Cause:</b> Missing the input for resource tag name during GenApp resource remove.</p> <p><b>Action:</b> Check the input for resource tag name and do resource remove again.</p>
126345	ERROR	-i flag not specified	<p><b>Cause:</b> Missing the input for resource internal id during GenApp resource remove.</p> <p><b>Action:</b> Check the input for resource internal id and do resource remove again.</p>
126357	ERROR	END timeout remove of \"\$TAG\" (forcibly terminating)	
126365	ERROR	remove script \"\$LCDAS/\$APP_REMOVE DIR/\$TAG\" was not found or was not executable	<p><b>Cause:</b> Failed to locate the user defined script for GenApp resource during remove.</p> <p><b>Action:</b> Check the existence/availability of the user defined script and do the GenApp remove process again.</p>
126375	ERROR	Script has hung checking \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "quickCheck" Script will be forcibly terminated for GenApp resource with tag name {tag} due to a waiting time over the user defined timeout.</p> <p><b>Action:</b> Check the GenApp resource performance and restart quickChecking. Also check the logs for related errors and try to resolve the reported problem.</p>
126380	ERROR	Usage error: no tag specified	<p><b>Cause:</b> Missing the input for resource t</p>

Code	Severity	Message	Cause/Action
			<p>ag name during GenApp resource quick Check.</p> <p><b>Action:</b> Check the input for resource tag name and retry resource quickCheck.</p>
126385	ERROR	Internal error: ins_list failed on \$tag.	<p><b>Cause:</b> Failed using "ins_list" to fetch the GenApp resource information by input tag name during the quickCheck process.</p> <p><b>Action:</b> Correct the input tag name and do the quickCheck process again. Also check the logs for related errors and try to resolve the reported problem.</p>
126390	FATAL	Failed to fork process to execute \$userscript: \$!	<p><b>Cause:</b> Failed to fork process to execute user defined "quickCheck" script during the GenApp resource "quickCheck" process.</p> <p><b>Action:</b> Check the existency/availability of the user defined "quickCheck" script and do the "quickCheck" process again.</p>
126391	ERROR	quickCheck has failed for \"\$tag\". Starting recovery.	<p><b>Cause:</b> The GenApp resource with tag name {tag} is determined to be failed by using the user defined health monitoring script – "quickCheck" and the recovery process will be initiated.</p> <p><b>Action:</b> Check the performance of the GenApp resource when local recovery finished. Also check the logs for related errors and try to resolve the reported problem.</p>
126392	WARN	\${convtag}_TIMEOUT: This parameter is old. This parameter will not be supported soon.	

Code	Severity	Message	Cause/Action
126400	ERROR	-t flag not specified	<p><b>Cause:</b> Missing the input for resource tag name during GenApp resource deletion process.</p> <p><b>Action:</b> Check the input for resource tag name and do resource deletion process again.</p>
126405	ERROR	-i flag not specified	<p><b>Cause:</b> Missing the input for resource internal id during GenApp resource deletion process.</p> <p><b>Action:</b> Check the input for resource internal id and do resource deletion process again.</p>
126478	ERROR	Failed to create tag '\\$new_leaf'.	<p><b>Cause:</b> The 'creapphier' utility failed to create the specified tag.</p> <p><b>Action:</b> Check /var/log/lifekeeper.log for additional messages from 'creapphier'.</p>
126479	ERROR	Failed to extend tag '\\$new_leaf'.	<p><b>Action:</b> Check /var/log/lifekeeper.log for additional errors from extend manager.</p>
126481	ERROR	Failed to create dependency on '\\$sys' for '\\$new_leaf' to '\\$hier{\\$leaf}{\\$sys}'Tag'.	<p><b>Action:</b> Check /var/log/lifekeeper.log for errors from the 'dep_create' function</p>
126484	ERROR	Tag '\\$root_tag' is not in-service.	<p><b>Cause:</b> The specified tag is not in-service on any node in the cluster.</p> <p><b>Action:</b> Bring the specified tag in-service on any node in the cluster and re-run 'create_terminal_leaf'.</p>
126485	ERROR	Tag '\\$root_tag_1' was not found, select the root tag for a hierarchy to add a terminal leaf resource.	<p><b>Cause:</b> The first tag passed to 'create_terminal_leaf' was not found on the system.</p>

Code	Severity	Message	Cause/Action
			<p>em where the utility was run.</p> <p><b>Action:</b> Verify each resource is in-service on a node in the cluster and fully extended to all nodes.</p>
126486	ERROR	Unable to create leaf tag from '\$tag'.	<p><b>Cause:</b> A unique terminal leaf tag could not be created. A unique terminal leaf tag could not be created after 100 tries.</p> <p><b>Action:</b> Check for multiple leaf tags and for errors in /var/log/lifekeeper.log that may indicate the problem.</p>
126487	ERROR	Tag '\$root_tag_2' was not found, select the root tag for a hierarchy to add a terminal leaf resource.	<p><b>Cause:</b> The second tag passed to 'create_terminal_leaf' was not found on the system where the utility was run.</p> <p><b>Action:</b> Verify the resource is in-service on a node in the cluster and fully extended to all nodes.</p>
126488	ERROR	Tag '\$root_tag_1' is not extended to 3 or more systems.	<p><b>Cause:</b> The specified tag is not extended to 3 or more nodes.</p> <p><b>Action:</b> Extend the specified tag to at least 3 nodes and retry 'create_terminal_leaf'.</p>
126489	ERROR	\$cmd does not support SDRS resources.	<p><b>Cause:</b> A multi-site configuration was detected.</p> <p><b>Action:</b> none</p>
126492	ERROR	Remove resource \$tag failed.	
126494	ERROR	Delete dependency failed on '\$sys' for '\$tag' to '\$parent'.	
126495	ERROR	New tag '\$new_tag' was modified during create, expected '\$new_leaf'.	

Code	Severity	Message	Cause/Action
126496	ERROR	Tag '\\$root_tag_1\' is not in-service.	
126497	ERROR	Tag '\\$root_tag_2\' is not in-service.	
126498	ERROR	Tag '\\$root_tag_1\' and '\\$root_tag_2\' are not extended to same systems.	
128005	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The quickCheck of {resource} on {server} failed due to an operating system signal {signal}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128008	ERROR	Usage: quickCheck -t <tag name> -i <id>	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device quickCheck command preventing it from running.</p> <p><b>Action:</b> Make sure all software components are properly installed and at the correct version. Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource to be quickChecked.</p>
128010	ERROR	quickCheck for "%s" failed checks of underlying paths, initiate recovery. retry count=%s.	<p><b>Cause:</b> The dmmp kit failed to quickCheck a device after {count} times of retries. A recovery of the protected dmmp resource will be executed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128021	ERROR	unable to find device for uuid "%s".	<p><b>Cause:</b> The device could not be found by unique id during a restore operation.</p> <p><b>Action:</b> Verify that the resource is configured properly. Rerun the command and supply the correct device id that identifies the resource.</p>

Code	Severity	Message	Cause/Action
			ifies the dmmp device resource to be restored.
128025	ERROR	Device "%s" failed to unlock.	<p><b>Cause:</b> A non working {device} was detected and could not be unlocked during the restore operation.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128026	ERROR	Device "%s" failed to lock.	<p><b>Cause:</b> The {device} could not be locked during the restore.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128031	ERROR	unable to find device for uuid "%s".	<p><b>Cause:</b> The device could not be found by unique id during the remove operation.</p> <p><b>Action:</b> Verify that the resource is configured properly. Rerun the command and supply the correct device id that identifies the dmmp device resource to be removed.</p>
128034	ERROR	Device "%s" failed to unlock.	<p><b>Cause:</b> The {device} could not be unlocked during the remove.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>

Code	Severity	Message	Cause/Action
128036	ERROR	unable to load existing information for device with uuid "%s".	<p><b>Cause:</b> The device information could not be loaded by unique id.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct device id that identifies the dmmp device resource.</p>
128037	ERROR	unable to load existing information for device "%s".	<p><b>Cause:</b> The device information could not be loaded by name.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct device name that identifies the dmmp device resource.</p>
128038	ERROR	unable to load existing information for device, no dev or uuid defined.	<p><b>Cause:</b> The device information could not be loaded since neither a unique device id nor name of the device were defined.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct device id or name that identifies the dmmp device resource.</p>
128041	ERROR	unable to load existing information for device with uuid "%s".	<p><b>Cause:</b> The device information could not be loaded by unique id.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct device id that identifies the dmmp device resource.</p>
128057	ERROR	All paths are failed on "%s".	<p><b>Cause:</b> LifeKeeper detected all paths listed to the protected dmmp device are in the failed state.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check adjacent log messages for further details and related messages.
128058	ERROR	could not determine registrations for "%s"! All paths failed.	<p><b>Cause:</b> LifeKeeper could not determine registrations for protected dmmp {device}. All paths to the dmmp {device} are in the failed state.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128059	WARN	path "%s" no longer configured for "%s", remove from path list.	<p><b>Cause:</b> LifeKeeper detected listed {path} to protected {device} is not valid anymore and will remove it from the path list.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128060	WARN	registration failed on path "%s" for "%s".	<p><b>Cause:</b> LifeKeeper failed the registration on {path} for protected dmmp {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128062	ERROR	all paths failed for "%s".	<p><b>Cause:</b> LifeKeeper failed to verify a valid path to protected dmmp {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128072	ERROR	The daemon "%s" does not appear to be running and could not be restarted. Path failures may not be correctly handled without this daemon.	<p><b>Cause:</b> LifeKeeper failed to verify dmmp daemon is running and could not restart it.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>

Code	Severity	Message	Cause/Action
128078	ERROR	"%s" resource type is not installed on "%s".	<p><b>Cause:</b> The Device Mapper Multipath Recovery Kit for dmmp device support is not installed on the system.</p> <p><b>Action:</b> Install the steeleye-ikDMMP Device Mapper Multipath Recovery Kit rpm on the system.</p>
128083	ERROR	This script must be executed on "%s".	<p><b>Cause:</b> An incorrect system name has been supplied as an argument to the devicehier script used to create the dmmp device resource.</p> <p><b>Action:</b> Make sure the cluster nodes and comm-paths are properly configured. Supply the correct system name to the devicehier script. The name must match the name of the system on which the command is run.</p>
128084	ERROR	The device %s is not active.	<p><b>Cause:</b> LifeKeeper failed to find the specified {device} as a valid device on the system during resource creation.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource to be created.</p>
128086	ERROR	Failed to create "%s" hierarchy.	<p><b>Cause:</b> LifeKeeper failed to create resource hierarchy for {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>

Code	Severity	Message	Cause/Action
128088	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper failed to create the resource with {tagname} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128090	ERROR	Failed to create dependency "%s"-"%s" on system "%s".	<p><b>Cause:</b> LifeKeeper failed to create dependency {resource tag name} – {resource tag name} on {system} during creation.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128091	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper failed to create {resource} on {system}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128101	ERROR	"%s" constructor requires a valid argument.	<p><b>Cause:</b> LifeKeeper failed to create an object for the dmmp resource during construction.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource.</p>
128102	ERROR	Invalid tag "%s".	<p><b>Cause:</b> A resource instance could not be found for the given tag name.</p> <p><b>Action:</b> Make sure the resource is confi</p>

Code	Severity	Message	Cause/Action
			<p>gured properly. Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource.</p>
128111	ERROR	<p>Failed to get registrations for "%s": %s. Verify the storage supports persistent reservations.</p>	<p><b>Cause:</b> LifeKeeper failed to get the registrations of {device} with the message, "bad field in Persistent reservation in cdb".</p> <p><b>Action:</b> Verify if the storage supports persistent reservations. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128112	ERROR	<p>Failed to get registrations for "%s": %s. Verify the storage supports persistent reservations.</p>	<p><b>Cause:</b> LifeKeeper failed to get the registrations of {device} with the message, "illegal request".</p> <p><b>Action:</b> Verify if the storage supports persistent reservations. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128136	ERROR	<p>A previous quickCheck with PID "%s" running for device "%s" has been terminated.</p>	<p><b>Cause:</b> LifeKeeper detected that a previous quickCheck operation is still running during the dmmp resource restore operation. It has been terminated by LifeKeeper.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128137	ERROR	<p>SCSI reservation conflict on %s during LifeKeeper resource initialization. Manual intervention required.</p>	<p><b>Cause:</b> LifeKeeper detected a SCSI reservation conflict on {device} during dmm</p>

Code	Severity	Message	Cause/Action
			<p>p resource restore.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Manual intervention and fix of the reservation conflict on {device} is required.</p>
128138	ERROR	unable to clear registrations on %s.	<p><b>Cause:</b> LifeKeeper failed to clear all the registrations on {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128140	WARN	registration failed on path %s for %s.	<p><b>Cause:</b> LifeKeeper failed to make the registration on {path} for {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128143	ERROR	reserve failed (%d) on %s.	<p><b>Cause:</b> LifeKeeper failed to make reservation for {resource} on {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128145	ERROR	The server ID "%s" returned by "%s" is not valid.	<p><b>Cause:</b> LifeKeeper failed to generate a valid host {ID}.</p> <p><b>Action:</b> The ID used to register a device is made up of 1 to 12 Hex digits that uniquely identifies the server in the cluster. Check adjacent log messages for further details and related messages. You</p>

Code	Severity	Message	Cause/Action
			must correct any reported errors before retrying the operation.
128146	ERROR	device failure on %s. SYSTEM HALTED.	<p><b>Cause:</b> LifeKeeper detected failure on {device} and will reboot the server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128148	ERROR	device failure on %s. SYSTEM HALTED DISABLED.	<p><b>Cause:</b> LifeKeeper detected a failure on {device}. The reboot was skipped due to LifeKeeper configuration.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Turn on the configuration "SCSIHALT" to make the reboot available for any detected device failure.</p>
128149	ERROR	device failure or SCSI Error on %s. SEND EVENT DISABLED.	<p><b>Cause:</b> LifeKeeper detected a failure on {device}. The event generation was skipped due to LifeKeeper configuration.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Turn on the configuration "SCSIEVENT" to make the sendevent available for any detected device failure.</p>
128150	ERROR	%s does not have EXCLUSIVE access to %s, halt system.	<p><b>Cause:</b> LifeKeeper detected a reservation conflict for {device} on {server} and will reboot the server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128151	ERROR	%s does not have EXCLUSIVE access to %s, halt system DISABLED.	<p><b>Cause:</b> LifeKeeper detected a reservation conflict for {device} on {server}. The</p>

Code	Severity	Message	Cause/Action
			<p>reboot was skipped due to LifeKeeper configuration.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Turn on the configuration "RESERVATIONCONFLICT" to make the reboot available for any detected reservation conflicts.</p>
128154	WARN	unable to flush buffers on %s.	<p><b>Cause:</b> LifeKeeper failed to flush the buffers for {device} during dmmp resource remove.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128157	WARN	%s utility not found, limited healthcheck for %s.	<p><b>Cause:</b> LifeKeeper failed to find "dd" utility for the health check of {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128160	ERROR	%s failed to read %s.	<p><b>Cause:</b> LifeKeeper failed a disk I/O test for {device} when using {utility}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128163	ERROR	Registration ID "%s" for "%s" is not valid.	<p><b>Cause:</b> LifeKeeper failed to generate a valid registration {ID} for {device}.</p> <p><b>Action:</b> The ID used to register a device</p>

Code	Severity	Message	Cause/Action
			e is made up of 4 Hex digits derived from the path to the device. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.
128170	ERROR	Usage: canextend <Template system name> <Template tag name>	
128500	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device restore command preventing it from running.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource to be restored.</p>
128504	ERROR	"%s" resource type is not installed on "%s".	<p><b>Cause:</b> The Device Mapper Multipath Recovery Kit for dmmp device support is not installed on the system.</p> <p><b>Action:</b> Install the steeleye-ikDMMP Device Mapper Multipath Recovery Kit rpm on the system.</p>
128506	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device devShared command preventing it from running.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: &lt;Template Resource System Name&gt; and &lt;Template Resource Tag&gt; that identifies the dmmp device resource to be created.</p>
128507	FATAL	This script must be executed on "%s".	<p><b>Cause:</b> An incorrect system name has</p>

Code	Severity	Message	Cause/Action
			<p>been supplied as an argument to the devicehier script used to create the dmmp device resource.</p> <p><b>Action:</b> Supply the correct system name to the devicehier script. The name must match the name of the system on which the command is run.</p>
128511	ERROR	Failed to get the ID for the device "%s". Hierarchy create failed.	<p><b>Cause:</b> The devicehier script used to create the dmmp device resource was unable to determine the SCSI ID for the supplied device.</p> <p><b>Action:</b> Check that the supplied device path exists and that is for a supported SCSI storage array.</p>
128512	ERROR	Failed to get the disk ID for the device "%s". Hierarchy create failed.	<p><b>Cause:</b> The devicehier script used to create the dmmp disk resource was unable to determine the SCSI ID for the supplied disk.</p> <p><b>Action:</b> Check that the supplied device path exists and that is for a supported SCSI storage array.</p>
128513	ERROR	Failed to create the underlying resource for device "%s". Hierarchy create failed.	<p><b>Cause:</b> The creation of the underlying dmmp disk resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128515	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The creation of the dmmp device resource failed.</p> <p><b>Action:</b> Check adjacent log messages for</p>

Code	Severity	Message	Cause/Action
			<p>or further details and related messages. You must correct any reported errors before retrying the operation.</p>
128517	ERROR	Failed to create dependency "%s"-"%s" on system "%s".	<p><b>Cause:</b> The parent child dependency creation between the dmmp device and dmmp disk resources failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128519	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The attempt to bring the newly created dmmp device resource in service has failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128521	ERROR	Either TEMPLATESYS or TEMPLATETAG argument missing	<p><b>Cause:</b> Incorrect arguments have been supplied to the extend command for the dmmp device resource.</p> <p><b>Action:</b> Rerun the dmmp device resource extend and supply the correct template system and tag names.</p>
128540	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device getId command used to retrieve the SCSI ID.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -i &lt;device path&gt; or -b &lt;device ID&gt;.</p>

Code	Severity	Message	Cause/Action
128541	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the command used to delete the dmmp device resource.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;dmmp device resource tag&gt;.</p>
128543	ERROR	device node \"\$dev\" does not exist.	<p><b>Cause:</b> The device node required for restoring the dmmp device resource does not exist. The allocated wait time in restore for udev device creation has been exceeded.</p> <p><b>Action:</b> Rerun the dmmp device resource restore once udev has created the device.</p>
128544	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the remove command used to take the dmmp device resource out of service.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;dmmp device resource tag&gt;.</p>
128550	ERROR	Failed to get path information.(tag=\"\$opt_t\")	
128551	ERROR	Failed to get the status of path.(tag=\"\$opt_t\")	
128552	ERROR	Failed to check the status of path.(tag=\"\$opt_t\")	
128553	ERROR	Failed to get the resource information.(tag=\"\$opt_t\")	
129100	FATAL	Failed to load instance from LifeKeeper.	<p><b>Cause:</b> An invalid resource tag or ID was specified.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check that the tag or ID is valid and re-run the command.
129103	FATAL	No resource matches tag \"\${self->{tag}}\".	<b>Cause:</b> An invalid resource tag was specified. <b>Action:</b> Check the tag and re-run the command.
129104	FATAL	An error occurred setting LifeKeeper resource information	<b>Cause:</b> An internal error has occurred in LifeKeeper.
129110	ERROR	Could not get the Elastic Network Interface ID for \$dev	<b>Cause:</b> The EC2 API call failed, possibly due to a network issue. <b>Action:</b> Check the network and the Amazon console and retry the operation.
129111	ERROR	Failed to get Allocation ID of Elastic IP \"\${elasticIp}\".	<b>Cause:</b> The EC2 API call failed, possibly due to a network issue. <b>Action:</b> Check the network and the Amazon console and retry the operation.
129113	ERROR	Failed to get my instance ID.	<b>Cause:</b> The EC2 instance metadata access failed. <b>Action:</b> Check the Amazon console and retry the operation.
129114	ERROR	Failed to get ENI ID.	<b>Cause:</b> The EC2 API call failed, possibly due to a network issue. <b>Action:</b> Check the network and the Amazon console and retry the operation.

Code	Severity	Message	Cause/Action
129116	ERROR	Failed to associate Elastic IP \"\$self->{'EIP'}\" on \"\$self->{'DEV'}\".	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129118	WARN	\$self->{'EIP'} is not associated with any instance.	<p><b>Cause:</b> The Elastic IP is not associated with any instance.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check adjacent log messages for more details.</p>
129119	WARN	\$self->{'EIP'} is associated with another instance.	<p><b>Cause:</b> The Elastic IP is associated with another instance.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check adjacent log messages for more details.</p>
129120	ERROR	Failed to recover Elastic IP.	<p><b>Cause:</b> The EC2 API call failed to associate the Elastic IP.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129121	ERROR	Recovery process ended but Elastic IP is not associated with this instance. Please check AWS console.	<p><b>Cause:</b> The EC2 API call failed to associate the Elastic IP.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129122	ERROR	Error creating resource \"\$target_tag\" with return code of \"\$err\".	<p><b>Cause:</b> LifeKeeper was unable to create the resource instance on the server.</p> <p><b>Action:</b> Check adjacent log messages f</p>

Code	Severity	Message	Cause/Action
			or further details and related messages. Correct any reported errors.
129123	ERROR	Failed to get ENI ID.	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129124	WARN	\$self->{'EIP'} is associated with another network interface.	<p><b>Cause:</b> The Elastic IP is associated with the proper instance, but the wrong ENI.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check adjacent log messages for more details.</p>
129125	ERROR	Link check failed for interface '\$dev'.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present.</p> <p><b>Action:</b> Check the network interface and bring the link up.</p>
129126	ERROR	Link check failed for interface '\$dev'. Reason: down slave.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present.</p> <p><b>Action:</b> Check the network interface and bring the link up.</p>
129129	WARN	The link for network interface '\$self->{'DEV'}' is down. Attempting to bring the link up.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by bringing the link up and associatin</p>

Code	Severity	Message	Cause/Action
			g the Elastic IP with the interface. Check adjacent log messages for more details.
129130	ERROR	Failed to modify \"\$opt_t\" to endpoint URL \"\$endpoint\".	
129137	ERROR	The link for network interface \"\$self->{'DEV'}\" is still down.	<p><b>Cause:</b> LifeKeeper could not bring the link up.</p> <p><b>Action:</b> Ensure the interface is enabled and up. Check adjacent log messages for more details.</p>
129139	WARN	The link for network interface \"\$self->{'DEV'}\" is down.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present.</p> <p><b>Action:</b> Check the network interface and bring the link up.</p>
129140	ERROR	Could not get ENI ID for \$self->{IP}.	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129142	ERROR	Failed to update route table	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129143	ERROR	You must have exactly one IP address resource as the parent of the RouteTable EC2 resource. Please reconfigure your resource hierarchy.	<p><b>Cause:</b> The Route Table EC2 resource must have one and only one IP resource as a parent.</p> <p><b>Action:</b> Repair the resource hierarchy a</p>

Code	Severity	Message	Cause/Action
			s necessary.
129144	ERROR	\$func called with invalid timeout: \$timeout	<p><b>Cause:</b> An invalid timeout value was specified in the /etc/default/LifeKeeper file.</p> <p><b>Action:</b> Verify all EC2_*_TIMEOUT settings are valid in /etc/default/LifeKeeper.</p>
129145	ERROR	\$func action timed out after \$timeout seconds	<p><b>Cause:</b> The action did not complete within the timeout period.</p> <p><b>Action:</b> Consider increasing the EC2_*_TIMEOUT value for the given action (in /etc/default/LifeKeeper).</p>
129146	ERROR	failed to run \$func with timeout: \$@	<p><b>Cause:</b> This is an internal error.</p>
129148	ERROR	Amazon describe-route-tables call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129150	ERROR	Elastic IP \"\$elasticIp\" is associated with another instance.	<p><b>Cause:</b> The Elastic IP is not associated with the proper instance.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check adjacent log messages for more details.</p>
129151	ERROR	Could not get the Association ID for Elastic IP \"\$elasticIp\".	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>

Code	Severity	Message	Cause/Action
129152	ERROR	Failed to disassociate Elastic IP \"\$self->{'EIP'}\" on \"\$self->{'DEV'}\".	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129153	ERROR	Failed to disassociate Elastic IP \"\$elasticIp\", (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129154	ERROR	Amazon describe-addresses call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129155	ERROR	Amazon describe-address call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129157	ERROR	curl call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 instance metadata access failed.</p> <p><b>Action:</b> Check the Amazon console and retry the operation.</p>
129159	ERROR	Amazon associate-address call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129160	ERROR	Amazon describe-addresses call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly</p>

Code	Severity	Message	Cause/Action
			<p>y due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129161	ERROR	Error deleting resource \"\$otherTag\" on \"\$otherSys\" with return code of \"\$err\".	<p><b>Cause:</b> LifeKeeper was unable to delete the resource instance on the server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
129162	ERROR	Could not getRouteTablesByIP	
129163	ERROR	Could not getRouteTablesByIP	
129164	ERROR	[\$SUBJECT event] mail returned \$err	
129165	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
129167	ERROR	snmptrap returned \$err for Trap 180	
129168	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
129170	ERROR	This resource is in the old format. Please update.	
129171	ERROR	Error disabling sourceDest checks with return code of \"\$ret\".	
129172	ERROR	Failed to lookup IP from resource \$self->{tag}.	
129173	ERROR	Failed to lookup device GUID for IP resource \$iptag.	
129174	ERROR	Could not find eni_id for device GUID \$dev.	
129175	ERROR	Failed to disable sourceDestChecks. (err=%s)(output=%s	
129176	ERROR	Failed to disable sourceDestChecks for \$eni_id. (err=%s)(output=%s	
129177	ERROR	Error disabling sourceDest checks with return code of \"\$ret\".	
129178	ERROR	Failed to disable sourceDestChecks for \$eni_id. (err=%s)(output=%s	

Code	Severity	Message	Cause/Action
129180	ERROR	Elastic Network Interface \$eni sourceDest Checks are enabled.	
129181	WARN	WARNING Failed to find IP resources since EC2 resource is extended first. Skipping sourceDestChecks.	
129182	WARN	WARNING SourceDestCheck was skipped because Amazon describe-network-interface-attribute call failed. Please check if ec2:DescribeNetworkInterfaceAttribute and ec2:ModifyNetworkInterfaceAttribute are granted.	
129403	ERROR	END failed create of \$TAG due to a \$sig signal	<b>Cause:</b> The create process was interrupted by a signal.
129409	ERROR	The IP resource \$IP_RES is not \"ISP\".	<b>Cause:</b> The IP resource is not in service. <b>Action:</b> Bring the resource in service and retry the operation.
129410	ERROR	Could not find IP resource \$IP_RES	<b>Cause:</b> Ensure that the IP resource exists and retry the operation.
129412	ERROR	EC2 resource \$ID is already protected	<b>Cause:</b> A resource with the specified ID already exists. <b>Action:</b> Make sure to clean up any remnants of an old resource before re-creating a new resource.
129416	ERROR	Error creating resource \"\$TAG\" with return code of \"\$!cderror\".	<b>Cause:</b> LifeKeeper was unable to create the resource instance. <b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.

Code	Severity	Message	Cause/Action
129418	ERROR	Dependency creation between \"\$IP_RE S\" and \"\$TAG\" failed with return code of \"\$!cderror\".	<p><b>Cause:</b> LifeKeeper was unable to create the resource dependency.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
129420	ERROR	In-service failed for tag \"\$TAG\".	<p><b>Cause:</b> LifeKeeper could not bring the resource instance into service.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
129423	ERROR	Could not get ENI ID for \$dev.	
129425	ERROR	Failed to update route table	
129426	ERROR	In-service (dummy) failed for tag \"\$TAG\".	
129800	ERROR	canextend checks failed for \"\$self->{tag}\" (err=\$ret	<p><b>Cause:</b> The pre-extend checks failed for the target server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
129801	ERROR	canextend checks failed for \"\$self->{tag}\". EC2_HOME \"\$self->{EC2_HOME}\" does not exist on \$me.	
133106	ERROR	You must have exactly one IP address resource as the child of the route53 resource. Please reconfigure your resource hierarchy.	
133111	ERROR	Failed to init the object.	
133114	ERROR	Failed start Route53 resource \$self->{Tag}	
133115	ERROR	Failed start Route53 resource \$self->{Tag}	
133116	WARN	Could not Get A Record Value Address :	

Code	Severity	Message	Cause/Action
		\$aAWSresult	
133117	WARN	aws call failed : \$aAWSresult	
133118	WARN	aws call failed : \$upsertResp	
133119	ERROR	Could not Update and Create A Record Set : \$upsertResp	
133120	WARN	Could not get Route53 API batch request ID form UPSERT response XML data : \$upsertResp	
133121	WARN	Could not Get A Record Value Address : \$aAWSresult	
133122	WARN	aws call failed : \$statusResult	
133123	ERROR	Failed check change batch request status : \$statusResult	
133126	WARN	Failed to Route53 API response : \$rc	
133129	ERROR	Route53 resource \$self->{Tag} is stopped	
133130	ERROR	\$self->{Tag} is in the old format. Please update.	
133601	ERROR	\$usage	
133602	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" does not exist	
133608	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" does not exist	
133614	ERROR	Error creating resource \"\$Tag\" on server \"\$me\"	
133617	ERROR	END failed hierarchy extend of resource \$Tag with return value of \$ecode.	
133619	ERROR	\$usage	
133620	ERROR	END failed extend of \"\$Tag\" due to a \"\$sig\" signal	
133700	ERROR	\$usage	
133701	ERROR	END failed create of \"\$Route53Tag\" due to a \"\$sig\" signal	
133703	ERROR	Unable to getlocks on \$me during resource create.	

Code	Severity	Message	Cause/Action
133704	ERROR	The route53 on \$Route53HostName is already protected by LifeKeeper.	
133706	ERROR	Error creating resource \$Route53Tag. Error (\$rc	
133708	ERROR	Dependency create failed between \$Route53Tag and \$Route53IPResTag. Error (\$rc).	
133710	ERROR	In-service attempted failed for tag \$Route53Tag.	
133813	ERROR	Usage: \$usage	
133815	ERROR	The host name \"\$hostname\" is too long.	
133816	ERROR	The host name \"\$hostname\" is too short.	
133817	ERROR	The host name \"\$hostname\" contains invalid character.	
133818	ERROR	The first character must be an alpha character.	
133819	ERROR	The last character must not be a minus sign or period.	
133826	ERROR	Host Name cannot be blank.	
134003	ERROR	catch a \"\$sig\" signal	<p><b>Cause:</b> The "create" process was interrupted by a signal.</p> <p><b>Action:</b> Check adjacent log messages.</p>
134004	ERROR	Unable to getlocks on \$server during resource create. Error (\$rc	<p><b>Cause:</b> Failed to get the administrative lock when creating a resource hierarchy.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
134005	ERROR	The service \"\$serviceName\" is not supported on \$server. Error (\$rc	<p><b>Cause:</b> The service does not exist or cannot be protected.</p> <p><b>Action:</b> Input the appropriate service name.</p>

Code	Severity	Message	Cause/Action
134006	ERROR	The service \"\$serviceName\" is already protected on \$server.	<p><b>Cause:</b> This service is already protected.</p> <p><b>Action:</b> Can not create a resource for the protection of this service.</p>
134007	ERROR	Error creating resource \$tag. Error (\$rc	<p><b>Cause:</b> LifeKeeper was unable to create the resource instance.</p> <p><b>Action:</b> Check adjacent log messages. Correct the cause of the error.</p>
134011	ERROR	In-service attempted failed for tag \$tag.	<p><b>Cause:</b> Failed to restore QSP resource.</p> <p><b>Action:</b> Check the log related to the service that you want to protect. Resolve the problem.</p>
134015	ERROR	Unable to rlslocks on \$server during resource create. Error (\$rc	<p><b>Cause:</b> Failed to release lock after QSP resource created.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
134103	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" does not exist	<p><b>Cause:</b> The resource cannot be found on the template server.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template server before extending.</p>
134104	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" is not QSP resource (app=\$ins <sup>1</sup> , res=\$ins <sup>2</sup>	<p><b>Cause:</b> The template resource is not QSP resource.</p> <p><b>Action:</b> Expand to the same type of resources as a template resource.</p>

Code	Severity	Message	Cause/Action
134105	ERROR	The service \"\$service\" is not supported on \$me. Error (\$check	<p><b>Cause:</b> The service does not exist on target server.</p> <p><b>Action:</b> Install the service on target server before extending.</p>
134106	ERROR	The service \"\$service\" is already protected on \$me.	<p><b>Cause:</b> There is already a resource of the same ID on target server.</p> <p><b>Action:</b> Cannot create the resource of the same service.</p>
134203	ERROR	catch a \"\$sig\" signal	<p><b>Cause:</b> The "extend" process resource was interrupted by a signal.</p> <p><b>Action:</b> Check adjacent log messages.</p>
134204	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" does not exist	<p><b>Cause:</b> The resource cannot be found on the template server.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template server before extending.</p>
134208	ERROR	Error creating resource \"\$tag\" on server \"\$me\"	<p><b>Cause:</b> LifeKeeper was unable to create the resource instance on target server.</p> <p><b>Action:</b> Check adjacent log messages. Correct the cause of the error.</p>
134401	ERROR	timeout \$cmd for \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "restore" process of the service does not terminate within the specified time.</p> <p><b>Action:</b> Check about the protected service and retry the "restore" operation. Also check the logs for related errors and t</p>

Code	Severity	Message	Cause/Action
			ry to resolve the reported problem.
134405	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p> <p><b>Action:</b> Determine why fork fails.</p>
134407	ERROR	service command has failed for \"\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p> <p><b>Action:</b> It is an error to manually run the service command with "start" option. Correct the cause of the error by reference to error message.</p>
134501	ERROR	timeout \$cmd for \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "remove" process of the service does not terminate within the specified time.</p> <p><b>Action:</b> Check about the protected service and retry the "remove" operation. Also check the logs for related errors and try to resolve the reported problem.</p>
134505	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p> <p><b>Action:</b> Determine why fork fails.</p>
134507	ERROR	service command has failed for \"\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p> <p><b>Action:</b> It is an error to manually run the service command with "stop" option. Correct the cause of the error by reference to error message.</p>

Code	Severity	Message	Cause/Action
134601	ERROR	timeout \$cmd for \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "quickCheck" process will be forcibly terminated due to a waiting time over the user defined timeout.</p> <p><b>Action:</b> Check about the protected service. Also check the logs for related errors and try to resolve the reported problem.</p>
134605	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p> <p><b>Action:</b> Determine why fork fails.</p>
134607	ERROR	service command has failed for \"\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p> <p><b>Action:</b> It is an error to manually run the service command with "status" option. Correct the cause of the error by reference to error message.</p>
134701	ERROR	timeout \$cmd for \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "recover" process of the service does not terminate within the specified time.</p> <p><b>Action:</b> Check about the protected service. Also check the logs for related errors and try to resolve the reported problem.</p>
134706	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p> <p><b>Action:</b> Determine why fork fails.</p>
134708	ERROR	service command has failed for \"\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> It is an error to manually run the service command with "start" option. Correct the cause of the error by reference to error message.
134803	ERROR	tag \"\$tag\" does not exist on server \"\$mel\"	<b>Cause:</b> The specified tag does not exist. This is an internal error.
134804	ERROR	app type \"\$ins <sup>1</sup> \" is not \$app	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134805	ERROR	res type \"\$ins <sup>2</sup> \" is not \$res	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134823	ERROR	tag \"\$tag\" does not exist on server \"\$mel\"	<b>Cause:</b> The specified tag does not exist. This is an internal error.
134824	ERROR	app type \"\$ins <sup>1</sup> \" is not \$app	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134825	ERROR	res type \"\$ins <sup>2</sup> \" is not \$res	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134843	ERROR	tag \"\$tag\" does not exist	<b>Cause:</b> The specified tag does not exist. This is an internal error.
134844	ERROR	app type \"\$ins <sup>1</sup> \" is not \$app	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134845	ERROR	res type \"\$ins <sup>2</sup> \" is not \$res	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
135802	ERROR	object[%d] not iterated for %lld ms.	

Code	Severity	Message	Cause/Action
135806	ERROR	qwk_config() failed.	
135807	ERROR	thread_initialize() failed.	
135808	ERROR	state_monitor_initialize() failed.	
135809	ERROR	state_monitor() failed.	
135810	ERROR	start_server() failed.	
135812	ERROR	'fopen' for %s failed: %m	
135813	ERROR	'qwk_object_path=' line cannot be found f or %s.	
135814	ERROR	'%s' is unknown config.	
135815	ERROR	configuration of hbeattime '%d' is incorrec t.	
135816	ERROR	configuration of numhbeats '%d' is incorre ct.	
135817	ERROR	configuration of timeout_multiplier '%d' is i ncorrect.	
135818	ERROR	configuration of lcmhbeattime '%d' is incor rect.	
135819	ERROR	configuration of lcmnumhbeats '%d' is inc orrect.	
135820	ERROR	configuration of qwk_object_type is incorr ect.	
135821	ERROR	configuration of my_node is incorrect.	
135822	ERROR	configuration of number_of_node '%d' is i ncorrect.	
135823	ERROR	configuration of number_of_object '%d' is incorrect.	
135824	ERROR	configuration of object node is incorrect.	
135825	ERROR	configuration of object path is incorrect.	
135826	ERROR	my_node is not include in qwk_objects.	
135827	ERROR	'open' for %s failed: %m	
135828	ERROR	'read' for %s failed: %m	
135829	ERROR	'popen' for %s failed: %m	
135830	ERROR	'open' for %s failed: %m	

Code	Severity	Message	Cause/Action
135831	ERROR	'write' for %s failed: %m	
135832	ERROR	'popen' for %s failed: %m	
135833	ERROR	(bug) buffer overflow	
135834	ERROR	(bug) data is corrupted	
135835	ERROR	'signature=' line cannot be found.	
135836	ERROR	signature '%s' does not match.	
135837	ERROR	'local_node=' line cannot be found.	
135838	ERROR	local_node '%s' does not match.	
135839	ERROR	'time=' line cannot be found.	
135840	ERROR	'sequence=' line cannot be found.	
135841	ERROR	sequence '%s' scan failed.	
135842	WARN	'node=' line cannot be found. index=%d	
135843	ERROR	'commstat=' line cannot be found. index=%d	
135844	ERROR	'checksum=' line cannot be found.	
135845	ERROR	checksum '%s' scan failed.	
135846	ERROR	checksum does not match.	
135849	ERROR	qwk object was not found.	
135851	ERROR	failed to read qwk object.	
135852	ERROR	failed to decode node_info.	
135854	WARN	sequence backed down from %llu to %llu.	
135855	ERROR	'malloc' for %zu failed: %m	
135856	ERROR	thread_create() failed. index=%d	
135870	ERROR	(bug) data is corrupted	
135871	ERROR	format error in request.	
135872	ERROR	format error in quorum_verify request. lke vent cannot be found.	
135873	ERROR	format error in witness_verify request. lke vent cannot be found.	
135874	ERROR	format error in witness_verify request. target_node cannot be found.	

Code	Severity	Message	Cause/Action
135875	ERROR	'%s' is unknown command.	
135877	ERROR	qwk_receive() did not receive full header. Close socket.	
135878	ERROR	Request is too long. Close socket.	
135879	ERROR	qwk_receive() did not receive full request. Close socket.	
135880	ERROR	do_request() failed.	
135881	ERROR	qwk_send() did not send full header. Close socket.	
135882	ERROR	qwk_send() did not send full response. Close socket.	
135884	ERROR	qwk_accept() failed.	
135885	ERROR	thread_create() failed.	
135886	ERROR	cannot create socket.	
135887	ERROR	'bind' failed: %m	
135888	ERROR	'listen' failed: %m	
135889	ERROR	create_sockets() failed.	
135890	ERROR	thread_create() failed.	
135891	ERROR	'pthread_attr_init' failed: %m	
135892	ERROR	'pthread_attr_setstacksize' failed: %m	
135893	ERROR	'pthread_create' failed: %m	
135894	ERROR	request is too long. header.size=%zu	
135895	ERROR	qwk_send() failed for header.	
135896	ERROR	qwk_send() failed for request.	
135897	ERROR	qwk_send() failed for termination.	
135898	ERROR	qwk_receive() did not receive full header.	
135899	ERROR	Response buffer is not enough large. Server sent %zu bytes.	
135900	ERROR	qwk_receive() did not receive full response.	
135901	ERROR	'%s' is unknown command.	
135903	ERROR	Cannot create socket.	

Code	Severity	Message	Cause/Action
135904	ERROR	'connect' failed: %m	
135905	ERROR	request_send() failed.	
135906	ERROR	request_receive() failed.	
135907	ERROR	'%s' is unknown lkevent.	
135908	ERROR	'%s' is unknown qwktype.	
135909	ERROR	'%s' is unknown node state.	
135910	ERROR	'%s' is unknown quorum state.	
135911	ERROR	'accept' failed: %m	
135912	ERROR	You must install the LifeKeeper license key for Storage Quorum Witness Kit.	
135913	EMERG	WARNING: Cannot find the configuration for node '%s' in the configuration file. If the configuration of the cluster is changed after the initialization, it needs to be initialized again. Please see the SIOS product documentation for information on how to reinitialize.	
135999	ERROR	-c \$cmd	
136002	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA create script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: &lt;Resource Tag&gt; &lt;SAP SID&gt; &lt;HDB Instance&gt; [Switchback Type] [Virtual IP Resource Tag]</p>
136003	ERROR	END failed create of resource \$tag on server \$me with return value of \$errcode.	<p><b>Cause:</b> Failure during SAP HANA resource creation.</p> <p><b>Action:</b> Verify that SAP HANA System Replication is fully configured and enabled on both the primary and secondary replication sites and reattempt the resource creation operation.</p>
136005	ERROR	An unknown error has occurred in utility rlocks on server \$me. View the LifeKeeper	<b>Cause:</b> Failure of the LifeKeeper rlocks

Code	Severity	Message	Cause/Action
		r logs for details and retry the operation.	s utility during SAP HANA resource creation.  <b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource creation operation.
136006	ERROR	END failed create of resource \$tag on server \$me with signal \$sig.	<b>Cause:</b> Failure during SAP HANA resource creation due to a signal.  <b>Action:</b> Review the LifeKeeper log file for details.
136008	ERROR	An unknown error has occurred in utility getlocks on server \$me. View the LifeKeeper logs for details and retry the operation.	<b>Cause:</b> Failure of the LifeKeeper getlocks utility during SAP HANA resource creation.  <b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource creation operation.
136009	ERROR	The SAP HANA product was not found in the directory \$obj->{'hana_util_path'} on server \$me. Verify that SAP HANA is correctly installed and that all necessary file systems are mounted.	<b>Cause:</b> Required SAP HANA binaries could not be located during resource creation.  <b>Action:</b> Verify that SAP HANA is correctly installed and configured on all servers in the cluster.
136010	ERROR	Failed to create resource as id \$id already exists on system \$me.	<b>Cause:</b> A LifeKeeper resource with the given ID already exists on the system.  <b>Action:</b> Check whether the SAP HANA database is already protected by LifeKeeper.
136011	ERROR	Failed to create new tag \$tag for SAP HANA resource on \$me.	<b>Cause:</b> The provided SAP HANA resource tag is already in use by another Life

Code	Severity	Message	Cause/Action
			<p>Keeper resource and the LifeKeeper ne wtag utility failed to create a new tag.</p> <p><b>Action:</b> Choose a different tag name for the SAP HANA resource.</p>
136012	ERROR	Failed creation of resource with \$tag on system \$me.	<p><b>Cause:</b> Failed to create the given SAP HANA resource in LifeKeeper.</p> <p><b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource creation operation.</p>
136014	ERROR	Failed to create resource dependency for parent \$tag and child \$virtual_ip_tag.	<p><b>Cause:</b> Failed to create a dependency between the SAP HANA resource and its dependent virtual IP resource.</p> <p><b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource creation operation.</p>
136015	ERROR	The info field for resource \$tag could not be successfully generated using values [SID: \$info_sid, Instance: \$info_instance, Replication Mode: \$info_repl_mode, Site Name: \$info_site_name, Operation Mode: \$info_oper_mode]. If using SAP HANA System Replication, please verify that it is fully configured and enabled on both the primary and secondary systems before creating the SAP HANA resource.	<p><b>Cause:</b> An invalid value was found when creating the SAP HANA resource info field.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured and that SAP HANA System Replication is fully configured and enabled on all servers in the cluster.</p>
136016	ERROR	The selected server \$me is not the primary/source system for SAP HANA System Replication for the selected SID \$sid and HDB instance \$instance. Please select 'Cancel' and start this action on the primary/source HANA System Replication system.	<p><b>Cause:</b> Resource creation is initiated on a secondary system in HANA System Replication.</p> <p><b>Action:</b> Initiate the create action on a primary system in HANA System Replication.</p>

Code	Severity	Message	Cause/Action
136031	ERROR	END failed extend of resource \$target_tag on server \$me with return value of \$err_code.	<p><b>Cause:</b> Failure during SAP HANA resource extension.</p> <p><b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource extension operation.</p>
136033	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA extend script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: &lt;Template System&gt; &lt;Template Tag&gt; &lt;Switchback Type&gt; &lt;Target Tag&gt;</p>
136035	ERROR	Template resource \$template_tag on server \$template_sys does not exist.	<p><b>Cause:</b> The template SAP HANA resource to be extended does not exist on the template server.</p> <p><b>Action:</b> Verify that the SAP HANA resource that is being extended exists on the template server.</p>
136036	ERROR	Resource with matching id \$target_id already exists on server \$me for App \$app_type and Type \$res_type.	<p><b>Cause:</b> An SAP HANA resource with the same LifeKeeper ID already exists on the target system.</p> <p><b>Action:</b> Check whether the SAP HANA database is already protected by LifeKeeper on the target server.</p>
136037	ERROR	Resource with matching tag \$target_tag already exists on server \$me for App \$app_type and Type \$res_type	<p><b>Cause:</b> An SAP HANA resource with the same LifeKeeper resource tag already exists on the target server.</p> <p><b>Action:</b> Check whether the SAP HANA database is already protected by LifeKeeper on the target server.</p>

Code	Severity	Message	Cause/Action
136039	ERROR	Error creating resource \$target_tag on system \$me.	<p><b>Cause:</b> Failed to create an equivalent SAP HANA resource on the target server.</p> <p><b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource extension operation.</p>
136040	ERROR	The target tag (\$target_tag) and template tag (\$template_tag) must be the same.	<p><b>Cause:</b> Resource tag name used on primary system is different than resource tag name used on secondary system. Both must be same.</p> <p><b>Action:</b> While resource creation use same name as primary system tag and secondary system tag.</p>
136045	ERROR	Cannot extend resource \$template_tag to server \$me.	<p><b>Cause:</b> The SAP HANA canextend script indicates that the resource cannot be extended to the given target system.</p> <p><b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource extension operation.</p>
136047	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA canextend script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: &lt;Template System&gt; &lt;Template Tag&gt;</p>
136048	ERROR	Resource \$template_tag does not exist on server \$template_sys.	<p><b>Cause:</b> The template SAP HANA resource to be extended does not exist on the template server.</p> <p><b>Action:</b> Verify that the SAP HANA resource that is being extended exists on the template server.</p>

Code	Severity	Message	Cause/Action
136049	ERROR	The system user \$hana_user does not exist on server \$me.	<p><b>Cause:</b> The SAP Administrative User for the SAP HANA database does not exist on the given server.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured.</p>
136050	ERROR	The user id for user \$hana_user (\$template_uid) on template server \$template_sys is not the same as user id (\$uid) on target server \$me.	<p><b>Cause:</b> The user ID for the SAP Administrative User differs between the template and target servers.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>
136051	ERROR	The group id for user \$hana_user (\$template_gid) on template server \$template_sys is not the same as group id (\$gid) on target server \$me.	<p><b>Cause:</b> The group ID for the SAP Administrative User differs between the template and target servers.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>
136052	ERROR	The home directory for user \$hana_user (\$template_home) on template server \$template_sys is not the same as home directory (\$user_home) on target server \$me.	<p><b>Cause:</b> The home directory for the SAP Administrative User differs between the template and target servers.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>
136053	ERROR	The SAP HANA instance \$instance does not exist for \$sid on server \$me.	<p><b>Cause:</b> Installation directories for the given SAP HANA database instance could not be located.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>

Code	Severity	Message	Cause/Action
136055	ERROR	The SAP HANA site name \$target_obj->{'site_name'} on server \$me must be different from site name \$template_obj->{'site_name'} on \$template_obj->{'sys'}.	<p><b>Cause:</b> The SAP HANA System Replication site name is the same on both the primary and secondary servers.</p> <p><b>Action:</b> Stop the SAP HANA database on the secondary server and use the hdhnsutil utility to re-register the secondary replication site using a different site name.</p>
136056	ERROR	Unable to obtain SAP HANA System Replication parameters for database \$instance on server \$me. Please verify that SAP HANA System Replication is enabled and properly configured and that the database instance is running on all servers in the cluster.	<p><b>Cause:</b> SAP HANA System Replication parameters could not be determined for the database on the given system.</p> <p><b>Action:</b> Verify that SAP HANA System Replication is enabled and properly configured and that the database is running on all servers in the cluster.</p>
136057	ERROR	Unable to create a HANA object for database \$instance on server \$me. Please verify that database instance \$instance is properly installed.	<p><b>Cause:</b> Unable to create an internal hana object representing the given instance on the given server.</p> <p><b>Action:</b> Verify that the database instance is properly installed and that all necessary file systems are mounted.</p>
136083	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the SAP HANA delete script.</p> <p><b>Action:</b> Provide both a valid HANA Life Keeper resource tag name and resource ID. Usage: delete -t &lt;tag&gt; -i &lt;id&gt; [-U]</p>
136096	ERROR	One of the required parameters (server, tag, or action) was missing. Unable to set the local recovery policy.	<p><b>Cause:</b> One of the required parameters (server, tag, or action) was not provided.</p> <p><b>Action:</b> Provide the required parameter</p>

Code	Severity	Message	Cause/Action
			s and reattempt the operation.
136097	ERROR	Failed to \$flg_action local recovery for SAP HANA resource \$tag on server \$node.	<p><b>Cause:</b> Failed to enable or disable local recovery for the given SAP HANA resource on the given server.</p> <p><b>Action:</b> Verify that LifeKeeper is running and fully initialized, then reattempt the operation.</p>
136099	ERROR	Failed start of SAP Host Agent on server \$sys.	<p><b>Cause:</b> Failed to start the SAP Host Agent processes (i.e., saphostexec, sapocol) on the given server.</p> <p><b>Action:</b> Check SAP Host Agent process trace logs and correct any issues found, then reattempt the operation.</p>
136100	ERROR	Failed start of SAP Host Agent on server \$sys.	<p><b>Cause:</b> Failed to start SAP Host Agent processes (i.e., saphostexec, sapocol) on the given server.</p> <p><b>Action:</b> Inspect the SAP Host Agent process trace logs and correct any issues found, then reattempt the operation.</p>
136160	ERROR	Unable to create SAP HANA object. The SID and instance for the SAP HANA database must be provided.	<p><b>Cause:</b> Either the SAP SID or the SAP HANA instance name were missing while trying to create an SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136161	ERROR	Unable to create SAP HANA object. Resource system name and tag name must be provided.	<p><b>Cause:</b> Either the system name or resource tag name were missing while trying to create an SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for</p>

Code	Severity	Message	Cause/Action
			or more details.
136162	ERROR	Could not find any information regarding resource \$tag on \$sys.	<p><b>Cause:</b> Failed to obtain information about the SAP HANA resource with the given tag.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136168	ERROR	Unable to check status of SAP Host Agent on server \$self->{'sys'}. Command \"\$curr_cmd\" returned exit code \$ret.	<p><b>Cause:</b> Failed to determine the status of the SAP Host Agent processes on the given server.</p> <p><b>Action:</b> Inspect the SAP Host Agent trace files (e.g., dev_saphostexec) for more details.</p>
136174	ERROR	Unable to create SAP HANA object. The SID or instance value is missing.	<p><b>Cause:</b> Either the SAP SID, the SAP HANA instance name, or one of the SAP HANA System Replication values were missing while trying to create an SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136182	ERROR	Failed to \$flg_action flag \"\$HANA_FLAG_DATA_OUT_OF_SYNC}_\${seqv_tag}\$sys}\" on server \$sys.	<p><b>Cause:</b> Failed to create or remove the !HANA_DATA_OUT_OF_SYNC_&lt;tag&gt; flag on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136190	ERROR	Failed to start SAP Host Agent processes on server \$self->{'sys'}. Command \"\$host_agent_cmd\" returned \$ret.	<p><b>Cause:</b> Failed to start the SAP Host Agent processes on the given server.</p> <p><b>Action:</b> Inspect the SAP Host Agent tra</p>

Code	Severity	Message	Cause/Action
			ce files (e.g., dev_saphostexec) for more details.
136191	ERROR	Failed to start SAP OS Collector process on server \$self->{'sys'}. Command \"\$oscol_cmd\" returned \$ret.	<p><b>Cause:</b> Failed to start the SAP OS Collector process on the given server.</p> <p><b>Action:</b> Inspect the SAP OS Collector trace files (e.g., dev_coll) for more details.</p>
136193	ERROR	\$takeover_text of SAP HANA System Replication for SAP HANA database \$self->{'instance'} failed on server \$node with exit code \$ret.	<p><b>Cause:</b> Failed to register the given server as primary master for the given database in SAP HANA System Replication.</p> <p><b>Action:</b> Inspect the SAP HANA trace files (e.g., nameserver_&lt;hostname&gt;.xxxxx.xxx.trc) for more details.</p>
136197	ERROR	Update of resource info field for \$seqv_tag{\$sys} on \$sys failed with exit code \$setinfo_ret. Current info: [\$info]. Attempted new info: [\$new_info].	<p><b>Cause:</b> Failed to update the info field for the given resource on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136202	EMERG	Failed to disable Autostart for SAP HANA instance \$self->{'instance'} on server \$sys with exit code \$remexec_ret. Please manually set \"Autostart = 0\" in the instance profile \$profile on \$sys.	<p><b>Cause:</b> The value of the Autostart parameter could not be modified in the given HDB instance profile on the given server.</p> <p><b>Action:</b> Edit the HDB instance profile manually and set \"Autostart = 0\".</p>
136205	ERROR	Failed start of SAP Start Service for SAP HANA database \$instance on server \$sys.	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database.</p> <p><b>Action:</b> Inspect the SAP Start Service trace files (e.g., sapstartsvr.log) for more</p>

Code	Severity	Message	Cause/Action
			details.
136208	ERROR	LifeKeeper was unable to determine the SAP HANA System Replication mode for database \$rem_obj->{'instance'} on server \$rem_obj->{'sys'} while attempting to identify the previous primary replication site. Please resolve the issue and bring the SAP HANA resource in-service on the system where the database should be registered as primary master.	<p><b>Cause:</b> Failed to determine the SAP HANA System Replication mode on the given server. As a result, the previous primary replication site could not be identified.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136210	ERROR	Failed start of SAP Start Service for SAP HANA database \$rem_obj->{'instance'} on server \$rem_obj->{'sys'}. Unable to stop the database on the server where it is currently registered as primary master.	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database on the given server. As a result, the database could not be stopped on the current primary SAP HANA System Replication site.</p> <p><b>Action:</b> Inspect the SAP Start Service trace files (e.g., sapstartsrv.log) for more details.</p>
136212	ERROR	Failed stop of SAP HANA database \$rem_obj->{'instance'} on server \$rem_obj->{'sys'} where it is currently registered as primary master.	<p><b>Cause:</b> Failed to stop the given SAP HANA database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA trace files and LifeKeeper log file for more details.</p>
136217	ERROR	Failed start of SAP HANA database \$instance on server \$sys.	<p><b>Cause:</b> Failed to start the given SAP HANA database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA trace files and LifeKeeper log file for more details.</p>
136220	ERROR	Unable to register \$sys as a secondary SAP HANA System Replication site for database \$instance. The host name of the current primary replication site was not provided.	<p><b>Cause:</b> The host name of the current primary SAP HANA System Replication site was not provided.</p>

Code	Severity	Message	Cause/Action
		ded.	<p>e was not provided when attempting to register a secondary replication site.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136233	EMERG	<p>WARNING: A temporary communication failure has occurred between servers \$self-&gt;{'sys'} and \$sys. Manual intervention is required in order to minimize the risk of data loss. To resolve this situation, please take one of the following resource hierarchies out of service: \$self-&gt;{'tag'} on \$self-&gt;{'sys'} or \$eqv{\$sys} on \$sys. The server that the resource hierarchy is taken out of service on will become the standby server for SAP HANA database \$self-&gt;{'instance'}.</p>	<p><b>Cause:</b> A temporary communication failure has caused the given equivalent HANA resources to both be brought in-service at the same time on their respective host servers.</p> <p><b>Action:</b> Take the entire HANA resource hierarchy out of service on the server which should become the secondary replication site. Once the database has been stopped on that server, LifeKeeper will automatically register it as a secondary replication site during the next quickCheck cycle.</p>
136234	EMERG	<p>WARNING: SAP HANA database \$self-&gt;{'instance'} is running and registered as primary master on the following servers: \$primary_node_list. Manual intervention is required in order to minimize the risk of data loss. To resolve this situation, please stop database \$self-&gt;{'instance'} on the standby server by running the command 'su - \$self-&gt;{'sid_admin'} -c \"\$SAP_CONTROLLER -nr \$self-&gt;{'instance_number'} -function StopWait \$HANA_STOP_WAIT 5\" on that server, allow LifeKeeper to register the standby server as a secondary replication site, then use LifeKeeper to bring resource \$self-&gt;{'tag'} in-service on the intended primary replication site.</p>	<p><b>Cause:</b> The given SAP HANA database is running and registered as primary master on two cluster servers concurrently.</p> <p><b>Action:</b> Use the command provided in the message to stop the database on the standby server. Once the database is stopped, LifeKeeper will automatically register the standby server as a secondary replication site.</p>
136236	ERROR	<p>Failed to remove \${HANA_FLAG_DATA_OUT_OF_SYNC}_\$tag flag on server \$sys.</p>	<p><b>Cause:</b> Failed to remove the !HANA_DATA_OUT_OF_SYNC_&lt;HANA Tag&gt; LifeKeeper flag on the given server. This will cause the given SAP HANA resource</p>

Code	Severity	Message	Cause/Action
			<p>e to fail to come in-service on the given server until it is removed.</p> <p><b>Action:</b> Verify that LifeKeeper is running and fully initialized. If it can be verified that SAP HANA System Replication is in-sync, the out-of-sync flag can be removed manually with the command <code>"/opt/LifeKeeper/bin/flg_remove -f '!HANA_DATA_OUT_OF_SYNC_&lt;HANA Tag&gt;'"</code>.</p>
136238	ERROR	Unable to create SAP HANA object. Resource system name and tag name must be provided.	<p><b>Cause:</b> Either the server name or the resource tag name were missing while trying to create an SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136239	ERROR	Failed start of SAP Start Service for SAP HANA database \$obj->'instance' on server \$obj->'sys'. Unable to determine status of the database on \$obj->'sys'.	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database on the given server. As a result, the status of the database could not be determined.</p> <p><b>Action:</b> Inspect the SAP Start Service trace files (e.g., sapstartsrv.log) for more details.</p>
136242	ERROR	Unable to create SAP HANA object. At least one of the SAP HANA System Replication values is missing.	<p><b>Cause:</b> SAP HANA System Replication parameters could not be determined for the database on the given system.</p> <p><b>Action:</b> Verify that SAP HANA System Replication is enabled and properly configured and that the database is running on all servers in the cluster.</p>
136243	ERROR	Unable to locate the pingnfs utility (\$pingnfs) on server \$me. Please verify that this utility exists and is executable.	<p><b>Cause:</b> Unable to locate the pingnfs utility used for testing available of exported</p>

Code	Severity	Message	Cause/Action
			<p>NFS shares.</p> <p><b>Action:</b> Verify that the pingnfs utility exists in the given location and is executable.</p>
136245	ERROR	Critical NFS shares being exported by server \$sys (\$export_list) are not currently available. Please verify that the NFS server is alive and that all NFS-related services are running.	<p><b>Cause:</b> The given critical NFS shared file systems are not currently available.</p> <p><b>Action:</b> Verify that the NFS server exporting the file systems is alive and that all necessary NFS-related services are running.</p>
136248	ERROR	Unable to open file \$crit_mount_file on server \$me. Please verify that this file exists and is read-enabled.	<p><b>Cause:</b> Unable to open the file containing mount information for critical NFS shares on the given server.</p> <p><b>Action:</b> Verify that the file exists in the specified location and is read-enabled.</p>
136253	ERROR	The SAP HANA "takeover with handshake" feature is available only for SAP HANA versions 2.0 SPS04 and greater. Database \$self->{'instance'} cannot be resumed.	<p><b>Cause:</b> Resuming a suspended database is not supported in SAP HANA versions earlier than 2.0 SPS04.</p> <p><b>Action:</b> Upgrade to SAP HANA 2.0 SPS04 or later in order to use features related to "Takeover with Handshake".</p>
136254	ERROR	Attempt to resume suspended primary database \$self->{'instance'} on server \$self->{'sys'} failed with exit code \$ret.	<p><b>Cause:</b> The attempt to resume the suspended database instance on the given server failed.</p> <p><b>Action:</b> Inspect the LifeKeeper and SAP log files to determine the cause of the failure. Either reattempt the operation or bring the corresponding LifeKeeper resource in-service on a different server.</p>

Code	Severity	Message	Cause/Action
136258	ERROR	Attempt to register server \$node as the secondary SAP HANA System Replication site for database \$self->{"instance"} failed with exit code \$ret.	<p><b>Cause:</b> The attempt to register the given server as a secondary SAP HANA System Replication site for the given database instance failed.</p> <p><b>Action:</b> Inspect the LifeKeeper and SAP log files to determine the cause of the failure and manually correct any issues found. While the SAP HANA resource is in-service, LifeKeeper will automatically continue attempting to register the backup server in a secondary HSR role.</p>
136263	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the SAP HANA resource restore script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: -t &lt;Resource Tag&gt; -i &lt;Resource ID&gt;</p>
136265	ERROR	Error getting resource information for \$tag on server \$me.	<p><b>Cause:</b> Failed to obtain information about the given SAP HANA resource on the given server.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136266	ERROR	The resource \$tag protecting SAP HANA database \$instance is not in sync. To protect the data LifeKeeper will not restore the resource on \$me. Please restore the resource on the previous source server to allow the resync to complete.	<p><b>Cause:</b> SAP HANA System Replication was not in sync before attempting to bring the database resource in-service on the backup server.</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the previous primary server and allow the resynchronization to complete.</p>
136275	ERROR	Failed to determine SAP HANA System Replication mode for database \$instance on server \$me.	<p><b>Cause:</b> Failed to determine the SAP HA</p>

Code	Severity	Message	Cause/Action
			<p>NA System Replication mode for the given database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA trace files (e.g., nameserver_&lt;hostname&gt;.xxxxx.xxx.trc) and the LifeKeeper log file for more details.</p>
136351	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the SAP HANA resource quickCheck script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: -t &lt;Resource Tag&gt; -i &lt;Resource ID&gt;</p>
136353	ERROR	Error getting resource information for \$tag.	<p><b>Cause:</b> Failed to obtain information about the given SAP HANA resource.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136354	EMERG	The replication mode of the SAP HANA database \$instance corresponding to resource \$tag was modified outside of LifeKeeper and is no longer registered as primary master on server \$me. Please bring the SAP HANA resource in-service on the server where the database should be registered as primary master. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System Replication mode for the given database was modified outside of LifeKeeper.</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the server where it should be registered as primary master.</p>
136363	EMERG	LifeKeeper was unable to determine the SAP HANA System Replication mode for database \$instance corresponding to resource \$tag on server \$me. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> Failed to determine the SAP HANA System Replication mode on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>

Code	Severity	Message	Cause/Action
136375	EMERG	An NFS server exporting a critical shared file system for resource \$tag is currently unavailable. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> A critical NFS shared file system is currently unavailable.</p> <p><b>Action:</b> Verify that the NFS server is alive and that all necessary NFS-related services are running.</p>
136376	EMERG	WARNING: LifeKeeper resource \$tag is designed for use in situations where SAP HANA System Replication (HSR) is disabled, but HSR was found to be enabled on server \$me. Please ensure that the correct LifeKeeper resource type is being used for your current SAP HANA configuration.	<p><b>Cause:</b> The given LifeKeeper resource is designed to protect a SAP HANA database environment where HANA System Replication (HSR) is disabled, but HSR is currently enabled on the given server.</p> <p><b>Action:</b> Verify that the correct LifeKeeper resource type is being used based for your SAP HANA configuration. If you have migrated from a SAP HANA configuration where HSR was disabled to one where it is enabled, the corresponding SAP HANA resource must be recreated in LifeKeeper.</p>
136377	EMERG	SAP HANA database \$instance corresponding to resource \$tag is currently suspended on server \$me due to actions performed outside of LifeKeeper. Please take the SAP HANA resource out of service on server \$me and bring it in-service on the server where the database should be registered as primary master. Bringing resource \$tag back in-service on \$me will resume the suspended database. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The given SAP HANA database instance has been suspended on the given server due to actions performed outside of LifeKeeper.</p> <p><b>Action:</b> If you would like to resume the suspended database on the server where the corresponding LifeKeeper resource is currently in-service (Active), execute the command given in the message. Otherwise, bring the LifeKeeper SAP HANA resource in-service on the intended primary replication site.</p>
136450	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA remove script.</p> <p><b>Action:</b> Please provide appropriate arguments.</p>

Code	Severity	Message	Cause/Action
			<p>uments in the form: &lt;Template Tag&gt; &lt;Template Id&gt;</p>
136454	ERROR	<p>Error getting resource information for \$tag.</p>	<p><b>Cause:</b> Failed to obtain information about the given SAP HANA resource.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136456	ERROR	<p>Failed start of SAP Start Service for SAP HANA database \$instance on server \$me.</p>	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database.</p> <p><b>Action:</b> Inspect the SAP Start Service trace files (e.g., sapstartsrv.log) for more details.</p>
136459	EMERG	<p>WARNING: LifeKeeper resource \$tag is designed for use in situations where SAP HANA System Replication (HSR) is disabled, but HSR was found to be enabled on server \$me. Please ensure that the correct LifeKeeper resource type is being used for your current SAP HANA configuration.</p>	<p><b>Cause:</b> The given LifeKeeper resource is designed to protect a SAP HANA database environment where HANA System Replication (HSR) is disabled, but HSR is currently enabled on the given server.</p> <p><b>Action:</b> Verify that the correct LifeKeeper resource type is being used based for your SAP HANA configuration. If you have migrated from a SAP HANA configuration where HSR was disabled to one where it is enabled, the corresponding SAP HANA resource must be recreated in LifeKeeper.</p>
136462	ERROR	<p>Failed to remove flag "\${hana::HANA_FLAG_LEAVE_DB_RUNNING}_\${tag}" on server \$me. This may cause subsequent remove actions for resource \$tag on server \$me to unintentionally fail to stop the database.</p>	<p><b>Cause:</b> Failed to remove the LifeKeeper !volatile!hana_leave_db_running_&lt;tag&gt; flag on the given server.</p> <p><b>Action:</b> Manually remove the flag with the command "/opt/LifeKeeper/bin/flg_remove -f !volatile!hana_leave_db_runnin</p>

Code	Severity	Message	Cause/Action
			<p>g_&lt;tag&gt;" on the given server. While the flag exists, out-of-service operations for the SAP HANA resource on the given server will leave the protected database instance running.</p>
136550	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the HANA resource recover script.</p> <p><b>Action:</b> Provide both a valid HANA LifeKeeper resource tag name and resource ID. Usage: recover -d &lt;tag&gt; -n &lt;id&gt;</p>
136555	ERROR	Error getting resource information for \$tag on server \$me.	<p><b>Cause:</b> Failed to obtain information about the given HANA resource on the given server.</p> <p><b>Action:</b> Verify that the server is online, LifeKeeper is running, and the HANA resource exists.</p>
136556	EMERG	The replication mode of the SAP HANA database \$instance corresponding to resource \$tag was modified outside of LifeKeeper and is no longer registered as primary master on server \$me. Please bring the SAP HANA resource in-service on the server where the database should be registered as primary master. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System Replication mode for the given database was modified outside of LifeKeeper.</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the server where it should be registered as primary master.</p>
136558	EMERG	LifeKeeper was unable to determine the SAP HANA System Replication mode for database \$instance corresponding to resource \$tag on server \$me. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System Replication mode could not be determined for the given database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA trace files and LifeKeeper log file for more details.</p>

Code	Severity	Message	Cause/Action
136559	ERROR	Resource \$tag is no longer ISP on server \$me. Exiting \$cmd for \$tag.	<p><b>Cause:</b> The given SAP HANA resource is no longer ISP on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136650	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA hana_stop_all_dbs script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: hana_stop_all_dbs -t &lt;tag&gt;</p>
136654	ERROR	Error getting resource information for \$tag.	<p><b>Cause:</b> Failed to obtain information about the given SAP HANA resource.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136658	ERROR	Failed start of SAP Start Service for SAP HANA database \$x->{'instance'} on server \$x->{'sys'}. Could not determine status of SAP HANA DB on \$x->{'sys'}.	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database.</p> <p><b>Action:</b> Inspect the SAP Start Service trace files (e.g., sapstartsvr.log) for more details.</p>
136661	ERROR	Failed stop of SAP HANA database \$x->{'instance'} on server \$x->{'sys'}.	<p><b>Cause:</b> Failed to stop the given SAP HANA database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA trace files and LifeKeeper log file for more details.</p>
136673	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA hana_takeover_with_handshake script.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Please provide appropriate arguments in the form: -t &lt;tag&gt; [-s &lt;target server&gt;] [-b]</p>
136674	ERROR	<p>Unable to remove the "\\${hana::HANA_FLAG_LEAVE_DB_RUNNING}_\$tag\" flag on server \$isp_node. This may cause subsequent remove actions for resource \$tag on server \$isp_node to unexpectedly fail to stop the database.</p>	<p><b>Cause:</b> Failed to remove the LifeKeeper !volatile!hana_leave_db_running_&lt;tag&gt; flag on the given server.</p> <p><b>Action:</b> Manually remove the flag with the command "/opt/LifeKeeper/bin/flg_remove -f !volatile!hana_leave_db_running_&lt;tag&gt;" on the given server. While the flag exists, out-of-service operations for the SAP HANA resource on the given server will leave the protected database instance running.</p>
136675	ERROR	<p>Script \$cmd exited unexpectedly due to signal "\$sig\" on server \$me. This may leave the \$tag resource hierarchy as well as SAP HANA System Replication in an unexpected state. Please verify that the cluster resources are in the expected state.</p>	<p><b>Cause:</b> The hana_takeover_with_hands hake script exited unexpectedly due to the given signal on the given server.</p> <p><b>Action:</b> Verify that the SAP HANA cluster resources are in the expected state. If not, fix any issues that are found and bring the SAP HANA resource hierarchy in-service on the intended primary server.</p>
136677	ERROR	<p>Unable to find equivalent SAP HANA resource corresponding to \$tag on server \$target_node.</p>	<p><b>Cause:</b> Unable to find an equivalent SAP HANA resource corresponding to the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource tag is correct, the resource has been extended to the given target server, and that LifeKeeper is running and fully initialized on the target server.</p>
136678	ERROR	<p>Unable to obtain information about equivalent SAP HANA resource \$tag on server \$target_node. Verify that LifeKeeper is ru</p>	<p><b>Cause:</b> Unable to find an equivalent SA</p>

Code	Severity	Message	Cause/Action
		<p>nnng and fully initialized.</p>	<p>P HANA resource corresponding to the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource tag is correct, the resource has been extended to the given target server, and that LifeKeeper is running and fully initialized on the target server.</p>
136679	ERROR	<p>Resource \$tag is not a SAP HANA resource.</p>	<p><b>Cause:</b> The given resource is not a SAP HANA (database/hana) resource.</p> <p><b>Action:</b> Verify that the resource tag is correct.</p>
136680	ERROR	<p>Resource \$tag is designed for use in an environment where SAP HANA System Replication is disabled. Takeover with handshake cannot be performed for this resource type. Please use the standard "\In Service...\\" command instead.</p>	<p><b>Cause:</b> Features related to "Takeover with Handshake" may only be used in environments where SAP HANA System Replication is enabled.</p> <p><b>Action:</b> Use the standard "In Service..." option to bring the SAP HANA resource in-service.</p>
136681	ERROR	<p>SAP HANA resource \$tag is not currently in-service on any server in the cluster. The resource must be in-service and SAP HANA System Replication must be in-sync before performing a takeover with handshake.</p>	<p><b>Cause:</b> The given SAP HANA resource is not in-service on any server in the cluster while attempting "Takeover with Handshake".</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the intended primary server.</p>
136682	ERROR	<p>SAP HANA resource \$tag is currently in-service on multiple servers: \$isp_node_list. The resource must be in-service on only one server and SAP HANA System Replication must be in-sync before performing a takeover with handshake.</p>	<p><b>Cause:</b> The given SAP HANA resource is in-service on multiple servers in the cluster while attempting a "Takeover with Handshake".</p> <p><b>Action:</b> Take the resource out of service.</p>

Code	Severity	Message	Cause/Action
			<p>e on every server except the one where it is intended to be registered as primary master.</p>
136684	ERROR	<p>Unable to create internal SAP HANA object for resource \$tag on server \$target_node. Verify that all necessary file systems are mounted and that LifeKeeper is running and fully initialized on \$target_node.</p>	<p><b>Cause:</b> Unable to create an internal hana object representing the given instance on the given server.</p> <p><b>Action:</b> Verify that the database instance is properly installed and that all necessary file systems are mounted.</p>
136685	ERROR	<p>Takeover with handshake is only supported in SAP HANA versions 2.0 SPS04 and greater. The SAP HANA software must be upgraded in order to use this feature. Please use the standard "In Service...\\" command instead.</p>	<p><b>Cause:</b> The "Takeover with Handshake" feature cannot be used when the underlying SAP HANA database version is less than 2.0 SPS04.</p> <p><b>Action:</b> Upgrade to SAP HANA 2.0 SPS04 or later in order to use the "Takeover with Handshake" feature.</p>
136686	ERROR	<p>Takeover with handshake cannot be performed for database \$target_obj-&gt;{'instance'} on server \$target_node because the database is not currently running and registered as primary on any other server in the cluster.</p>	<p><b>Cause:</b> The given SAP HANA database instance is not running and registered as primary master on any server in the cluster during an attempted "Takeover with Handshake".</p> <p><b>Action:</b> Bring the corresponding SAP HANA resource in-service on the intended primary server.</p>
136687	ERROR	<p>SAP HANA database \$target_obj-&gt;{'instance'} is running and registered as primary on more than one server in the cluster. Please resolve this situation and reattempt the takeover.</p>	<p><b>Cause:</b> The given SAP HANA database instance is running and registered as primary on multiple servers during an attempted "Takeover with Handshake".</p> <p><b>Action:</b> If the database instance is already running and registered as primary on the intended primary server, bring the</p>

Code	Severity	Message	Cause/Action
			<p>corresponding LifeKeeper resource in-service on that server. If not, stop the database instance on every server except the one where it is currently in-service on LifeKeeper, allow LifeKeeper to resume system replication, then reattempt the takeover.</p>
136689	ERROR	<p>Unable to create internal SAP HANA object for resource \$tag on server \$isp_node. Verify that all necessary file systems are mounted and that LifeKeeper is running and fully initialized on \$isp_node.</p>	<p><b>Cause:</b> Unable to create an internal hana object representing the given instance on the given server.</p> <p><b>Action:</b> Verify that the database instance is properly installed and that all necessary file systems are mounted.</p>
136690	ERROR	<p>Unable to set the \${hana::HANA_FLAG_LEAVE_DB_RUNNING}_\$tag flag on server \$isp_node. Aborting takeover with handshake attempt for database \$target_obj-&gt;{'instance'} on server \$target_node.</p>	<p><b>Cause:</b> Failed to set the !volatile!hana_leave_db_running_&lt;tag&gt; LifeKeeper flag on the given server during an attempted "Takeover with Handshake".</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more information. Correct any issues found and reattempt the takeover.</p>
136692	ERROR	<p>Takeover with handshake for database \$target_obj-&gt;{'instance'} failed on server \$target_node.</p>	<p><b>Cause:</b> The "Takeover with Handshake" attempt failed for the given SAP HANA database instance on the given target server.</p> <p><b>Action:</b> If HANA_HANDSHAKE_TAKEOVER_FAILBACK=true in /etc/default/LifeKeeper, the SAP HANA resource hierarchy will automatically be brought back in-service on the previous database host server. Otherwise, the SAP HANA resource hierarchy must be manually brought in-service on the intended primary server.</p>

Code	Severity	Message	Cause/Action
136695	ERROR	Resource \$tag does not exist on server \$me.	<p><b>Cause:</b> The given resource does not exist on the given server.</p> <p><b>Action:</b> Verify that the resource tag is correct.</p>
136696	ERROR	Unable to verify the status of resource \$tag on server \$target_node. Assuming that it is not in-service.	<p><b>Cause:</b> Failed to determine the status of the given resource on the given server while checking whether the "Takeover with Handshake" was successful.</p> <p><b>Action:</b> Verify that LifeKeeper is running and fully initialized on the given server and that the communication path between the local and target servers is active. If HANA_HANDSHAKE_TAKEOVER_FAILBACK=true in /etc/default/LifeKeeper, the SAP HANA resource hierarchy will automatically be brought back in-service on the previous database host server. Otherwise, the SAP HANA resource hierarchy must be manually brought in-service on the intended primary server.</p>
136697	ERROR	Resource \$res was not successfully brought in-service on server \$target_node.	<p><b>Cause:</b> The given resource failed to come in-service on the given server during an attempted "Takeover with Handshake".</p> <p><b>Action:</b> If HANA_HANDSHAKE_TAKEOVER_FAILBACK=true in /etc/default/LifeKeeper, the SAP HANA resource hierarchy will automatically be brought back in-service on the previous database host server. Otherwise, the SAP HANA resource hierarchy must be manually brought in-service on the intended primary server.</p>

Code	Severity	Message	Cause/Action
136698	ERROR	LifeKeeper is not running or is not fully initialized on server \$me.	<p><b>Cause:</b> LifeKeeper is either not running or not fully initialized on the given server during an attempted "Takeover with Handshake".</p> <p><b>Action:</b> Either start LifeKeeper with /opt/LifeKeeper/bin/lkstart or allow it to fully initialize, then reattempt the takeover.</p>
136699	ERROR	Unknown server \$target_node.	<p><b>Cause:</b> The given server host name is not recognized.</p> <p><b>Action:</b> Verify that the server host name is correct and that communication paths have been created between the local server and the target server.</p>
136700	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA remoteregisterdb script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: remoteregisterdb -d &lt;tag&gt; -n &lt;id&gt;</p>
136705	ERROR	Unable to obtain information about resource \$tag on server \$me. Exiting \$cmd for \$tag.	<p><b>Cause:</b> Failed to determine information about the given SAP HANA resource on the given server.</p> <p><b>Action:</b> Verify that the given resource tag is correct, the resource exists on the given server, all necessary file systems are mounted, and that LifeKeeper is running and fully initialized on the given server.</p>
136706	ERROR	Resource \$tag is no longer ISP on server \$me. Exiting \$cmd for \$tag.	<p><b>Cause:</b> The SAP HANA remoteregisterdb script is exiting because the SAP HANA resource is no longer in service (ISP) on the given server.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> No action is required.</p>
136707	ERROR	<p>The \$cmd event is intended for use only in environments in which SAP HANA System Replication is enabled. Exiting \$cmd for \$tag.</p>	<p><b>Cause:</b> The SAP HANA remoteregisterdb event detected that SAP HANA System Replication is disabled for the database protected by the given SAP HANA resource.</p> <p><b>Action:</b> Verify that SAP HANA System Replication is enabled on the server where the remoteregisterdb script was running. If necessary, System Replication may be enabled by executing the command 'hdbnsutil -sr_enable --name=&lt;Site Name&gt;' as the SAP HANA administrative user.</p>
136708	ERROR	<p>Error getting resource information for \$tag on server \$me. Exiting \$cmd for \$tag.</p>	<p><b>Cause:</b> Failed to obtain information about the given SAP HANA resource on the given server.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136709	EMERG	<p>LifeKeeper was unable to determine the SAP HANA System Replication mode for database \$instance corresponding to resource \$tag on server \$me. Resource monitoring for \$tag will be suspended until the issue is resolved.</p>	<p><b>Cause:</b> The SAP HANA System Replication mode could not be determined for the given database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA trace files and LifeKeeper log file for more details.</p>
136710	EMERG	<p>The replication mode of the SAP HANA database \$instance corresponding to resource \$tag was modified outside of LifeKeeper and is no longer registered as primary master on server \$me. Please bring the SAP HANA resource in-service on the server where the database should be registered as primary master. Resource monitoring</p>	<p><b>Cause:</b> The SAP HANA System Replication mode for the given database was modified outside of LifeKeeper.</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the server where it should</p>

Code	Severity	Message	Cause/Action
		g for \$tag will be suspended until the issue is resolved.	be registered as primary master.
136711	EMERG	SAP HANA database \$instance corresponding to resource \$tag is currently suspended on server \$me due to actions performed outside of LifeKeeper. Please take the SAP HANA resource out of service on server \$me and bring it in-service on the server where the database should be registered as primary master. Bringing resource \$tag back in-service on \$me will resume the suspended database. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The given SAP HANA database instance has been suspended on the given server due to actions performed outside of LifeKeeper.</p> <p><b>Action:</b> If you would like to resume the suspended database on the server where the corresponding LifeKeeper resource is currently in-service (Active), execute the command given in the message. Otherwise, bring the LifeKeeper SAP HANA resource in-service on the intended primary replication site.</p>
137000	ERROR	PowerShell is not installed.	
137001	ERROR	PowerCLI is not installed.	
137002	ERROR	A valid network interface was not found.	
137005	ERROR	Failed to attach VMDK.	
137010	ERROR	Failed to detach VMDK.	
137020	ERROR	Failed to execute VMDK status checker daemon.	
137030	ERROR	Disk not specified.	
137031	ERROR	Cannot get disk uuid for \$Disk. Please check your ESXi settings.	
137032	ERROR	PowerCLI failed. %s	
137034	ERROR	Cannot bring VMDK resource \"%s\" in service on server \"%s\".	
137037	ERROR	This system is not a VMware guest.	
137050	ERROR	Failed to connect to ESXi server \$addr.	
137051	ERROR	There is no ESXi server connected.	
137055	ERROR	Cannot determine ESXi VM ID because multiple network interfaces were found with the MAC address \$MAC_ADDR.	
137057	ERROR	Usable SCSI controller not found.	

Code	Severity	Message	Cause/Action
137058	ERROR	Cannot find VMDK with ID \$UUID.	
137059	ERROR	This guest has snapshots present.	
137060	ERROR	The VMDK with ID \$UUID cannot be attached to this guest.	
137061	ERROR	The virtual storage controller has an incompatible sharing mode configured.	
137068	ERROR	The VMDK detection failed. Retry count exceeded.	
137070	ERROR	Connect failed.	
137071	ERROR	Get-LocalVM failed.	
137072	ERROR	The VMDK has been detached remotely. This server has lost ownership.	
137075	ERROR	Cannot find virtual SCSI controller \$CONTROLLER.	
137076	ERROR	The virtual storage controller has an incompatible sharing mode configured.	
137077	ERROR	Cannot find VM with MAC address \$MAC_ADDR.	
137078	ERROR	Cannot find VM with MAC address \$MAC_ADDR.	
137101	WARN	Partition information not defined for %s on %s. Retry.	
137102	ERROR	Partition information not defined for %s on %s.	
137105	ERROR	Device not specified.	
137106	ERROR	Cannot get device uuid for \$Device. Please check your ESXi settings.	
137107	ERROR	%s is not shareable with any machine.	
137111	ERROR	Failed to create dependency \"%s\"-\"%s\" on machine \"%s\".	
137112	ERROR	Cannot bring VMDKP resource \"%s\" in service on server \"%s\".	
139000	ERROR	Cannot find the "oci" command in directories of the PATH. Please confirm that it is installed and the PATH is set correctly.	<b>Cause:</b> Unable to get PATH for "oci" oci command.

Code	Severity	Message	Cause/Action
			<b>Action:</b> Make sure that the <code>oci</code> command is installed and that the installation directory for the <code>oci</code> command is added to your PATH.
139012	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<b>Cause:</b> Failed to execute the <code>oci</code> command.  <b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.
139013	ERROR	Failed to assign \$ip, error:\"\$result\"	<b>Cause:</b> Failed to allocate \$ip.  <b>Action:</b> Resolve the cause of the error based on the \$result details.
139022	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<b>Cause:</b> Failed to execute the <code>oci</code> command.  <b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.
139023	ERROR	Failed to unassign \$ip, error:\"\$result\"	<b>Cause:</b> Failed to allocate \$ip.  <b>Action:</b> Resolve the cause of the error based on the \$result details. If \$ip was unallocated before the remove process was performed, make sure that the resource state is OSU and that \$ip is not allocated to the target VNIC on the OCI console.
139032	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<b>Cause:</b> Failed to execute the <code>oci</code> command.

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139033	ERROR	Failed to \"\$ocicmdstr\", unknown error: \"\$ip_list\"	<p><b>Cause:</b> The <code>oci</code> command terminated abnormally due to a cause other than code 139032.</p> <p><b>Action:</b> Unknown error: Please take action based on the \"\$ip_list\".</p>
139034	ERROR	There is no secondary IPs on \$device.	<p><b>Cause:</b> No secondary IP address has been assigned to \$device.</p> <p><b>Action:</b> If local recovery is enabled, make sure that local recovery was performed and therefore it recovered from the failure.</p>
139035	ERROR	\$ip is not assigned to \$device.	<p><b>Cause:</b> \$ip is not assigned to \$device.</p> <p><b>Action:</b> If local recovery is enabled, make sure that local recovery was performed and therefore it recovered from the failure.</p>
139042	ERROR	\"oci\" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<p><b>Cause:</b> Failed to execute the <code>oci</code> command.</p> <p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139043	ERROR	Failed to assign \$ip, error: \"\$result\"	<p><b>Cause:</b> Failed to allocate \$ip.</p> <p><b>Action:</b> Resolve the cause of the error based on the \$result details.</p>

Code	Severity	Message	Cause/Action
139060	ERROR	\$cmd is invalid.	<p><b>Cause:</b> \$cmd is invalid.</p> <p><b>Action:</b> Please contact us.</p>
139061	ERROR	OCIVIP does not support IPv6.	<p><b>Cause:</b> OCIVIP resources do not support use with IPv6.</p> <p><b>Action:</b> Please use IPv4.</p>
139062	ERROR	\$cmd is invalid.	<p><b>Cause:</b> \$cmd is invalid.</p> <p><b>Action:</b> Please contact us.</p>
139063	ERROR	IPv\$ipversion is unknown version.	<p><b>Cause:</b> The IP address version is incorrect.</p> <p><b>Action:</b> Please use IPv4.</p>
139070	ERROR	Failed to access the \$IMDS_URL with the curl command, status code is \$status_code.	<p><b>Cause:</b> Command curl to \$IMDS_URL failed.</p> <p><b>Action:</b> Make sure that curl to \$IMDS_URL completes successfully.</p>
139071	ERROR	Failed to decode JSON, error: \"\$e\".	<p><b>Cause:</b> Failed to decode JSON.</p> <p><b>Action:</b> Make sure that the result obtained from Instance Meta Data Service is in JSON format.</p>
139073	ERROR	Cannot find the vnicId of \$device.	<p><b>Cause:</b> The vnicId corresponding to \$device could not be obtained.</p> <p><b>Action:</b> Check if vnicId exists in the JSON record of Instance Meta Data Service.</p>

Code	Severity	Message	Cause/Action
139074	ERROR	Cannot find the subnetcidr of \$device.	<p><b>Cause:</b> The subnetCidrBlock corresponding to \$device could not be obtained.</p> <p><b>Action:</b> Check if subnetcidr exists in the JSON record of Instance Meta Data Service.</p>
139075	ERROR	Cannot find the vRouterIp of \$device.	<p><b>Cause:</b> The virtualRouterIp corresponding to \$device could not be obtained.</p> <p><b>Action:</b> Check if virtualRouterIp exists in the JSON record of Instance Meta Data Service.</p>
139080	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<p><b>Cause:</b> Failed to execute the <code>oci</code> command.</p> <p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139081	ERROR	Failed to get subnet id, error: \"\$subnetid\"	<p><b>Cause:</b> Failed to obtain the subnet id.</p> <p><b>Action:</b> Check if the information can be obtained correctly with the following command. The \$vnicid corresponds to the VNIC-ID (OCID) of each VNIC assigned to the node.</p> <pre>oci network vnic get --vnic-id \$vnicid --raw-output --query 'data.\"subnet-id\"'</pre>
139100	ERROR	IP address is not specified.	<p><b>Cause:</b> IP address is not specified.</p> <p><b>Action:</b> Make sure the IP address is specified.</p>
139101	ERROR	Device is not specified.	<p><b>Cause:</b> Device is not specified.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Make sure the device is specified.</p>
139105	ERROR	Cannot bring OCIVIP resource \$Tag in service on server \$SysName.	<p><b>Cause:</b> Failed to restore the OCIVIP resource \$Tag on \$SysName.</p> <p><b>Action:</b> Please refer to the log related to the restore process.</p>
139111	ERROR	IP-\$ip already exists, and device and net mask are inconsistent with those of \$ocvipTag.	<p><b>Cause:</b> An IP resource corresponding to the IP address of the OCIVIP resource exists, but the device information and netmask information are inconsistent.</p> <p><b>Action:</b> Change the configuration of the existing IP resource to have the same settings as the OCIVIP resource and create a dependency with the OCIVIP resource.</p>
139112	ERROR	Cannot bring IP resource \$iptag in service on server \$SysName.	<p><b>Cause:</b> Failed to restore the IP resource \$iptag corresponding to the OCIVIP resource on \$SysName.</p> <p><b>Action:</b> Review the configuration of the IP resource and make sure that restoring of the IP resource is successful. Then create a dependency with the OCIVIP resource.</p>
139200	ERROR	A resource with the specified tag name "%s" already exists on the target machine "%s".	<p><b>Cause:</b> A resource with the specified tag name already exists on the extended node.</p> <p><b>Action:</b> Please use a different tag name or delete the resource with the target tag name of the extended node.</p>

Code	Severity	Message	Cause/Action
139201	ERROR	Template resource "%s" does not exist on the server "%s".	<p><b>Cause:</b> The resource specified as the extension source does not exist on the node.</p> <p><b>Action:</b> Please specify the tag name correctly.</p>
139250	ERROR	Failed to valuenetOnOCI, error: \"\$stderr_log\".	<p><b>Cause:</b> valuentOnOCI failed.</p> <p><b>Action:</b> The full error message is in \$stderr_log. Please take action based on it.</p>

## 8.1.1. DataKeeper Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
104002	FATAL	\$msg	<p><b>Cause:</b> This message indicates an internal software error.</p> <p><b>Action:</b> The stack trace indicates the source of the error.</p>
104003	FATAL	\$self->Val('Tag') . " is not an SDR resource"	<p><b>Cause:</b> A data replication action was attempted on a non data replication resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
104010	ERROR	\$self->{'md'}: bitmap merge failed, \$action	<p><b>Cause:</b> The bitmap merge operation has failed.</p> <p><b>Action:</b> The target server may have the mirror and/or protected filesystem mounted, or the bitmap file may be missing on the target. Check the target server.</p>
104022	ERROR	\$argv <sup>1</sup> : mdadm failed (\$ret	<p><b>Cause:</b> The "mdadm" command has failed to add a device into the mirror.</p> <p><b>Action:</b> This is usually a temporary condition.</p>
104023	ERROR	\$_	<p><b>Cause:</b> The message contains the output of the "mdadm" command.</p>
104025	ERROR	failed to spawn monitor	<p><b>Cause:</b> The system failed to start the 'mdadm -F' monitor process. This should not happen under normal circumstances.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Reboot the system to ensure that any potential conflicts are resolved.
104026	ERROR	cannot create \$md	<b>Cause:</b> The mirror device could not be created.  <b>Action:</b> Ensure the device is not already in use and that all other parameters for the mirror creation are correct.
104027	ERROR	\$_	<b>Cause:</b> This message contains the "mdadm" command output.
104035	ERROR	Too many failures. Aborting resync of \$md	<b>Cause:</b> The device was busy for an abnormally long period of time.  <b>Action:</b> Reboot the system to be sure that the device is no longer busy.
104036	ERROR	Failed to start nbd-server on \$target (error \$port	<b>Cause:</b> The nbd-server process could not be started on the target server.  <b>Action:</b> Ensure that the target disk device is available and that its Device ID has not changed.
104037	ERROR	Failed to start compression (error \$port	<b>Cause:</b> The system was unable to start the 'balance' tunnel process or there was a network problem.  <b>Action:</b> Ensure that the network is operating properly and that TCP ports in the range 10000-10512 are opened and unused. Ensure that the software is installed properly on all systems.
104038	ERROR	Failed to start nbd-client on \$source (error \$ret	<b>Cause:</b> The nbd-client process has failed to start on the source server.  <b>Action:</b> Look up the reported errno value and try to

Code	Severity	Message	Cause/Action
			resolve the problem reported. For example, an errno value of 110 means "Connection timed out", which may indicate a network or firewall problem.
104039	ERROR	Failed to add \$nbd to \$md on \$source	<p><b>Cause:</b> This is usually a temporary condition.</p> <p><b>Action:</b> If this error persists, reboot the system to resolve any potential conflicts.</p>
104045	ERROR	failed to stop \$self->{'md'}	<p><b>Cause:</b> The mirror device could not be stopped.</p> <p><b>Action:</b> Ensure that the device is not busy or mounted. Try running "mdadm —stop" manually to stop the device.</p>
104048	WARN	failed to kill \$proc, pid \$pid	<p><b>Cause:</b> The process could not be signalled. This may indicate that the process has already died.</p> <p><b>Action:</b> Ensure that the process in question is no longer running. If it is, then reboot the system to clear up the unkillable process.</p>
104050	ERROR	Setting \$name on \$dest failed: \$ret. Please try again.	<p><b>Cause:</b> The system failed to set a 'mirrorinfo' file setting.</p> <p><b>Action:</b> Check the network and systems and retry the mirror setting operation.</p>
104052	FATAL	Specified existing mount point "%s" is not mounted	<p><b>Cause:</b> The mount point became unmounted.</p> <p><b>Action:</b> Ensure that the mount point is mounted and retry the operation.</p>
104055	ERROR	Failed to set up temporary \$type access to data for \$self->{'tag'}. Error: \$ret	<p><b>Cause:</b> The filesystem or device was not available on the target server. The mirrored data will not be available on the target server until the mirror is paused and resumed again.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Reboot the target server to resolve any potential conflicts.
104057	ERROR	Failed to undo temporary access for \$self->{'tag'} on \$self->{'sys'}. Error: \$ret. Please verify that \$fsid is not mounted on server \$self->{'sys'}.	<b>Cause:</b> The filesystem could not be unmounted on the target server. <b>Action:</b> Ensure that the filesystem and device are not busy on the target server. Reboot the target server to resolve any potential conflicts.
104062	FATAL	Cannot find a device with unique ID "%s"	<b>Cause:</b> The target disk could not be identified. <b>Action:</b> Ensure that the appropriate storage recovery kits are installed on the target server. Ensure that the Device ID of the target disk has not changed.
104066	FATAL	Cannot get the hardware ID of device "%s"	<b>Cause:</b> A unique ID could not be found for the target disk device. <b>Action:</b> Ensure that the appropriate storage recovery kits are installed on the target server. Ensure that the Device ID of the target disk has not changed.
104067	FATAL	Asynchronous writes cannot be enabled without a bitmap file	<b>Cause:</b> An attempt was made to create a mirror with invalid parameters. <b>Action:</b> A bitmap file parameter must be specified or synchronous writes must be specified.
104068	FATAL	Failed to extend dependent resource %s to system %s. Error %s	<b>Cause:</b> The hierarchy extend operation failed. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104070	FATAL	Unable to extend the mirror "%s" to system "%s"	<b>Cause:</b> The hierarchy extend operation failed.

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104071	ERROR	Failed to restore target device resources on \$target->{'sys'} : \$err	<b>Cause:</b> The in-service operation has failed on the target server.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104074	FATAL	Cannot get the hardware ID of device "%s"	<b>Cause:</b> There is no storage recovery kit that recognizes the underlying disk device that you are attempting to use for the mirror.  <b>Action:</b> Make sure the appropriate storage recovery kits are installed. If necessary, place your device name in the /opt/LifeKeeper/subsys/scsi/resources/DEVNAME/device_pattern file.
104081	FATAL	Cannot make the %s filesystem on "%s" (%d)	<b>Cause:</b> The "mkfs" command failed.  <b>Action:</b> Ensure that the disk device is writable and free of errors and that the filesystem tools for the selected filesystem are installed.
104082	FATAL	%s	<b>Cause:</b> This message contains the output of the "mkfs" command.
104083	FATAL	Cannot create filesys hierarchy "%s"	<b>Cause:</b> The resource creation failed.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104086	ERROR	The "%s_data_corrupt" flag is set in "%s/subsys/scsi/resources/netraid/" on system "%s". To avoid data corruption, LifeKeeper will not restore the resource.	<b>Cause:</b> The data corrupt flag file has been set as a precaution to prevent accidental data corruption. The mirror cannot be restored on this server until the file is removed.

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> If you are sure that the data is valid on the server in question, you can either: 1) remove the file and restore the mirror, or 2) force the mirror online using the LifeKeeper GUI or 'mirror_action force' command.</p>
104092	ERROR	Mirror target resource movement to system %s : status %s	<p><b>Cause:</b> The hierarchy switchover operation has failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
104099	ERROR	Unable to unextend the mirror for resource "%s" from system "%s"	<p><b>Cause:</b> The hierarchy unextend operation failed.</p> <p><b>Action:</b> Reboot the target server to resolve any potential conflicts and retry the operation.</p>
104106	ERROR	remote 'bitmap -m' command failed on \$target->{'sys'}: \$ranges	<p><b>Cause:</b> The bitmap merge command failed on the target server. This may be caused by one of two things: 1) The bitmap file may be missing or corrupted, or 2) the mirror (md) device may be active on the target.</p> <p><b>Action:</b> Make sure that the mirror and protected filesystem are not active on the target. If the target's bitmap file is missing, pause and resume the mirror to recreate the bitmap file.</p>
104107	ERROR	Asynchronous writes cannot be enabled without a bitmap file	<p><b>Cause:</b> Invalid parameters were specified for the mirror create operation.</p>
104108	ERROR	Local Partition not available	<p><b>Cause:</b> Invalid parameters were specified for the mirror create operation.</p>
104109	ERROR	Cannot get the hardware ID of device "%s"	<p><b>Cause:</b> A unique ID could not be determined for the disk device.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Ensure that the appropriate storage recovery kits are installed on the server. Ensure that the Device ID of the disk has not changed.
104111	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104112	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104113	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104114	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104115	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104117	FATAL	Insufficient input parameters for "%s" creation	<b>Cause:</b> Invalid parameters were specified for the mirror create operation.
104118	FATAL	Cannot unmount existing Mount Point "%s"	<b>Cause:</b> The mount point is busy.  <b>Action:</b> Make sure the filesystem is not busy. Stop any processes or applications that may be accessing the filesystem.
104119	FATAL	Invalid data replication resource type requested ("%s")	<b>Cause:</b> An invalid parameter was specified for the mirror create operation.
104124	EMERG	WARNING: A temporary communication failure has occurred between systems %s and %s. In order to avoid data	<b>Cause:</b> A temporary communication failure (split-brain scenario) has occurred between the source

Code	Severity	Message	Cause/Action
		corruption, data resynchronization will not occur. MANUAL INTERVENTION IS REQUIRED. In order to initiate data resynchronization, you should take one of the following resources out of service: %s on %s or %s on %s. The resource that is taken out of service will become the mirror target.	and target servers.  <b>Action:</b> Perform the steps listed in the message text.
104125	ERROR	failed to start '\$cmd \$_ <sup>2</sup> \$user_args' on '\$_ <sup>3</sup> '	<b>Cause:</b> The specified command failed.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
104126	ERROR	\$_	<b>Cause:</b> This message contains the output of the command that was reported as failing in message 104125.
104128	FATAL	comm path/server not specified	<b>Cause:</b> The netraid.down script was called without specifying the communication path or the server name. This script is called internally so should always have the proper parameters.  <b>Action:</b> Report this error to SIOS support.
104129	WARN		<b>Cause:</b> The replication connection for the mirror is down.  <b>Action:</b> Check the network.
104130	ERROR	Mirror resize failed on %s (%s). You must successfully complete this operation before using the mirror. Please try again.	<b>Cause:</b> The mirror resize operation has failed to update the mirror metadata on the listed system.  <b>Action:</b> You must successfully complete the resize before using the mirror. Re-run mirror_resize (possibly using -f to force the operation if

Code	Severity	Message	Cause/Action
			necessary).
104132	ERROR	The partition "%s" has an odd number of sectors and system "%s" is running kernel >= 4.12. Mirrors with this configuration will not work correctly with DataKeeper. Please see the SIOS product documentation for information on how to resize the mirror.	<p><b>Cause:</b> The partition or disk chosen for mirror creation has an odd number of disk sectors and will have to be resized to be used with DataKeeper.</p> <p><b>Action:</b> Resize the partition using the 'parted' command or resize the disk (is possible) using platform (VMware, AWS) tools. Caution: data may be lost if this is not done carefully.</p>
104136	ERROR	Extend failed.	<p><b>Cause:</b> The hierarchy extend operation failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
104143	ERROR	Mirror resume was unsuccessful (\$ret	<p><b>Cause:</b> The mirror could not be established.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problems.</p>
104144	ERROR	Unable to stop the mirror access for \$self->{'md'} on system \$self->{'sys'}. Error: \$ret. Use the "\"mdadm —stop \$self->{'md'}\" command to manually stop the mirror.	<p><b>Cause:</b> The mirror device created on the target node when the mirror was paused could not be stopped.</p> <p><b>Action:</b> Ensure that the device is not busy or mounted. Try running "mdadm —stop" manually to stop the device.</p>
104145	WARN	Unable to dirty full bitmap. Setting fullsync flag.	<p><b>Cause:</b> A full resync could not be done by dirtying the full bitmap. The fullsync flag will be used instead. This is a non-fatal error as a full synchronization will still be done.</p> <p><b>Action:</b> None</p>
104146	EMERG	WARNING: The target system %s for mirror %s has the target	<p><b>Cause:</b> The mirror is configured on the target</p>

Code	Severity	Message	Cause/Action
		mirror %s currently active. In order to avoid data corruption, data resynchronization will not occur. MANUAL INTERVENTION IS REQUIRED. In order to initiate data resynchronization, you should reboot system %s.	system.  <b>Action:</b> The target system should be rebooted. DataKeeper should then be able to resync the mirror.
104147	EMERG	WARNING: The target system %s for mirror %s has the target disk %s currently mounted. In order to avoid data corruption, data resynchronization will not occur. MANUAL INTERVENTION IS REQUIRED. In order to initiate data resynchronization, you should unmount %s on %s. A full resync will occur.	<b>Cause:</b> The mirror disk is mounted on the target system.  <b>Action:</b> The mirror disk should be unmounted on the target system, in order to initiate a full mirror resync. A full resync is required because untracked changes have occurred on the disk.
104148	EMERG	The storage configuration for mirror "%s (%s)" does not have a unique identifier and may have potential for data corruption in some environments in certain circumstances. Please refer to the SIOS Product Documentation for details on DataKeeper storage configuration options.	<b>Cause:</b> The disk chosen for mirroring does not provide a UUID to the operating system. DataKeeper cannot mirror a disk without a UUID.  <b>Action:</b> You may be able to create a GPT partition table on the disk to provide a UUID for the disk partitions.
104155	EMERG	The mirror %s cannot be forced online at this time. The underlying disk %s is mounted, indicating possible data corruption. MANUAL INTERVENTION IS REQUIRED. You must unmount %s on %s and restore the mirror to the last known mirror source system. A full resync will need to be performed from the source system to %s.	<b>Cause:</b> You have mounted the underlying mirrored disk on the target system.  <b>Action:</b> You must unmount the disk immediately in order to avoid data corruption.

Code	Severity	Message	Cause/Action
104156	WARN	Resynchronization of "%s" is in PENDING state. Current sync_action is: "%s"	<p><b>Cause:</b> The resynchronization of the md device is detected in PENDING state.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by forcing a resynchronization. Check the logs for related errors. When successful assure that the PENDING state has been cleared in /proc/mdstat and the resynchronization is in progress or has been completed for the datarep resource.</p>
104157	WARN	/etc/sysconfig/raid-check update failed. Please %s \"md%d\" to SKIP_DEVS.	<p><b>Cause:</b> Unable to make changes in /etc/sysconfig/raid-check to add or remove an entry to the list of MD devices to skip (SKIP_DEVS).</p> <p><b>Action:</b> Check system logs for any errors related to raid-check or SKIP_DEVS. Manually add or remove md listed.</p>
104158	EMERG	WARNING: The local disk partition \$self->{'part'} for data replication device\n\$self->{'md'} has failed. MANUAL INTERVENTION IS REQUIRED.	<p><b>Cause:</b> The local device for a mirror failed. The recovery action has been set to "nothing" in LKDR_FAILURE requiring manual intervention to recover.</p> <p><b>Action:</b> Check system logs and LifeKeeper logs for errors related to the local disk.</p>
104163	WARN	The "%s_data_corrupt" flag is set in "%s/subsys/scsi/resources/netraid/" on system "%s". The mirror is being forced online.	<p><b>Cause:</b> The mirror is being forced online, overriding the data_corrupt flag. The data on the specified system will be treated as the latest data. If this is not correct then this can lead to data corruption or data loss.</p> <p><b>Action:</b> None</p>
104164	ERROR	The "%s_data_corrupt" flag for related mirror resource "%s" is set in "%s/subsys/scsi/resources/netraid/" on system "%s". To avoid data corruption, LifeKeeper will not restore this mirror or any related mirrors in	<p><b>Cause:</b> The data_corrupt flag exists for one or more mirrors in the hierarchy. To avoid corrupting additional data none of the mirrors are brought in-service until all of the data_corrupt flags are resolved.</p>

Code	Severity	Message	Cause/Action
		the hierarchy.	<p><b>Action:</b> Check the LifeKeeper logs to determine where each mirror was last in-service, aka where the latest data for each mirror resides. The mirrors should be brought in-service on the "previous source" where the full hierarchy was in-service and allow the mirrors to synchronize with all targets.</p>
104165	ERROR	The "%s_data_corrupt" flag for related mirror resource "%s" is set in "%s/subsys/scsi/resources/netraid/" on system "%s". The mirror resource "%s" is being forced online.	<p><b>Cause:</b> The mirror is being forced online, overriding the data_corrupt flag. The data on the specified system will be treated as the correct data to be synchronized with all targets. This can lead to data corruption or data loss if this is not the latest data.</p> <p><b>Action:</b> None</p>
104170	ERROR	Failed to create "\"source\" flag file on shared source %s to track mirror source. This may result in a full resync.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104171	ERROR	Failed to create "\"source\" flag file on %s to track mirror source. Target %s will not be added to mirror.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104172	ERROR	The "\"source\" flag file on %s does not contain a valid target (%s). Full resync to remaining targets is required.	<p><b>Cause:</b> The 'source' flag file should contain the system name of a previous source but the name listed was not found in the list of systems configured.</p> <p><b>Action:</b> Report this problem to SIOS support.</p>
104173	ERROR	Failed to create "\"source\" flag file on %s to track mirror source.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file</p>

Code	Severity	Message	Cause/Action
			system for errors or that it is full.
104174	ERROR	Failed to create \"previous_source\" flag file to track time waiting on source. Will not be able to timeout.	<p><b>Cause:</b> The 'previous_source' flag file was not created on the mirror source to track the mirror's previous source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104175	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104176	ERROR	The \"source\" flag file on %s to track mirror source does not exist. Full resync is required.	<p><b>Cause:</b> The 'source' flag file should exist on the system and without it the consistency of the mirror can not be verified. A full resync is required to assure data reliability.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for error or that it is full.</p>
104177	ERROR	Failed to determine amount of time waiting on %s.	<p><b>Cause:</b> The amount of time waiting for the previous source could not be determined. Targets will be added with a full resync if the previous source is not found.</p> <p><b>Action:</b> none</p>
104178	ERROR	Failed to update "source" flag file on target "%s", previous source must be merged first.	<p><b>Cause:</b> The source flag file on the target is updated when it is in-sync and stopped so the the next in-service does not require the previous source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>

Code	Severity	Message	Cause/Action
104180	ERROR	Internal Error: \"previous_source\" has the local system name (%s).	<b>Cause:</b> The local system name should not be in the previous_source flag file.  <b>Action:</b> Report this error to SIOS support.
104181	ERROR	Internal Error: There are no targets waiting on %s to be merged.	<b>Cause:</b> There are no targets waiting for a previous source to merge.  <b>Action:</b> Report this error to SIOS support.
104182	ERROR	Failed to create \"source\" flag file on %s to track mirror source. This may result in a full resync.	<b>Cause:</b> The 'source' flag file was not created on the target listed.  <b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.
104186	ERROR	Failed to create \"last_owner\" flag file on %s to track mirror source. This may allow in-service of mirror on old data.	<b>Cause:</b> The 'last_owner' flag file was not created on the source.  <b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.
104187	WARN	\$REM_MACH has \${REM_TAG}_last_owner file, create flag \${FLAGTAG}_data_corrupt.	<b>Cause:</b> The system listed had the mirror in-service last.  <b>Action:</b> The system listed has the last_owner file that indicates it has the most recent data and is most likely the best system to in-service the mirror to avoid losing data.
104188	WARN	\$REM_MACH is not alive, create flag \${FLAGTAG}_data_corrupt.	<b>Cause:</b> The system listed is not alive.  <b>Action:</b> Since the system listed is not alive, it cannot be determined whether that system was a more recent mirror source than the local system. Therefore the local system should not automatically be allowed to bring the mirror in-service.

Code	Severity	Message	Cause/Action
104200	EMERG	Continue to wait for %s to merge bitmap and do partial resyncs to all targets, no timeout set.	<p><b>Cause:</b> In a multi-target configuration targets will not be configured until the previous source is available to merge its bitmap so that all targets will be able to partially resynchronize. The LKDR_WAIT_ON_PREVIOUS_SOURCE_TIMEOUT entry in /etc/defaults/LifeKeeper is set to "-1" to wait indefinitely.</p> <p><b>Action:</b> Check on the status of the previous source listed in the message and resolve any issues that are preventing it from rejoining the cluster.</p>
104201	EMERG	To stop waiting for the previous source (forcing a full resync to remaining waiting targets) run: <code>\"%s/bin/mirror_action %s fullresync %s %s\"</code> on %s.	<p><b>Cause:</b> In a multi-target configuration targets are not being configured, waiting on the previous source to rejoin the cluster.</p> <p><b>Action:</b> Run the command listed in the message to force an immediate full resynchronization to this target and any remaining targets waiting to be resynchronized.</p>
104202	EMERG	Continue to wait for %s to merge bitmap and do partial resyncs to all targets. Continue to wait %d more seconds.	<p><b>Cause:</b> In a multi-target configuration targets will not be configured until the previous source is available to merge its bitmap so that all targets will be able to partially resynchronize. The LKDR_WAIT_ON_PREVIOUS_SOURCE_TIMEOUT entry in /etc/defaults/LifeKeeper is set to the number of seconds to wait. If the previous source does not join the cluster in that time then targets will be added with a full resynchronization.</p> <p><b>Action:</b> Check on the status of the previous source listed in the message and resolve any issues that are preventing it from rejoining the cluster.</p>
104203	EMERG	To stop waiting for the previous source (forcing a full resync to remaining waiting targets) run: <code>\"%s/bin/mirror_action %s fullresync %s %s\"</code> on %s.	<p><b>Cause:</b> In a multi-target configuration targets are not being configured, waiting on the previous source to rejoin the cluster.</p> <p><b>Action:</b> Run the command listed in the message to</p>

Code	Severity	Message	Cause/Action
			<p>force an immediate full resynchronization to this target and any remaining targets waiting to be resynchronized.</p> <p>Note: Run this command to stop waiting even if the target listed is deleted and never returning.</p>
104207	ERROR	Failed to create "data_corrupt" flag file on "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the source listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104208	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104209	ERROR	Failed to create "data_corrupt" flag file on "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the source listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104210	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104211	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104212	ERROR	The "\"source\" flag file on %s to track mirror source does not	<p><b>Cause:</b> The "source" flag file should exist on the</p>

Code	Severity	Message	Cause/Action
		exist. Full resync to remaining targets is required.	<p>system and without it the consistency of the mirror can not be verified. A full resync is required to assure data reliability. All targets not already being mirrored will require a full resync.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for error or that it is full.</p>
104214	ERROR	Failed to create \"source\" flag file on %s to track mirror source.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104216	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104217	ERROR	Failed to create \"source\" flag file on shared source %s to track mirror source. This may result in a full resync.	<p><b>Cause:</b> The 'source' flag file was not created on the specified system to track the mirror source.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104218	ERROR	Failed to create "data_corrupt" flag file on target "%s".	<p><b>Cause:</b> The 'data_corrupt' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full</p>
104221	ERROR	Failed to create \"last_owner\" flag file on %s to track mirror source. This may allow in-service of mirror on old data.	<p><b>Cause:</b> The 'last_owner' flag file was not created on the target listed.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>

Code	Severity	Message	Cause/Action
104222	ERROR	Failed to create "last_owner" flag file to track mirror source". This may allow in-service of mirror on old data.	<p><b>Cause:</b> The 'last_owner' flag file is used to know where the mirror was last in-service.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104223	ERROR	Failed to create "last_owner" flag file to track mirror source". This may allow in-service of mirror on old data.	<p><b>Cause:</b> The 'last_owner' flag file is used to know where the mirror was last in-service.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104224	ERROR	Failed to create \"previous_source\" flag file.	<p><b>Cause:</b> The 'previous_source' flag file was not created. This is needed to merge the previous source bitmap to avoid a full resync.</p> <p><b>Action:</b> Check the LKROOT (/opt/LifeKeeper) file system for errors or that it is full.</p>
104227	ERROR	Failed to set %s to %s.	<p><b>Cause:</b> This message indicates a failure to set a sysfs parameter for the nbd driver (/sys/block/nbdX).</p> <p><b>Action:</b> It may be necessary to adjust one or more of:</p> <p>NBD_NR_REQUESTS NBD_SCHEDULER LKDR_ASYNC_LIMIT</p> <p>in /etc/default/LifeKeeper to avoid this error.</p>
104232	ERROR	Mirror resize failed on %s (%s). Could not set size to %d.	<p><b>Cause:</b> The mirror resize operation failed.</p>
104233	ERROR	Mirror resize failed on %s (%s). Could not set bitmap to %s and bitmap-chunk to %d.	<p><b>Cause:</b> The mirror resize operation failed.</p>

Code	Severity	Message	Cause/Action
104234	ERROR	The mirror %s failed to resize. You must successfully complete this operation before using the mirror. Please try again.	<b>Cause:</b> The mirror resize operation failed.
104235	ERROR	mirror_resize of mirror %s failed due to signal "%s".	<b>Cause:</b> The mirror resize operation failed.
104236	EMERG	Resource "%s" is "OSF". The mirror "%s" will wait to replicate data until all resources in the hierarchy%s are in-service. This may indicate inconsistent data. Verify the data is correct before replicating data; replication will continue when all resources in the hierarchy%s are in-service. A full resync may be necessary (see "LKDR_WAIT_TO_RESYNC" in /etc/default/LifeKeeper).	<b>Cause:</b> The specified resource is OSF. This prevents replication until this resource is repaired. The LKDR_WAIT_TO_RESYNC setting in /etc/default/LifeKeeper determines what resources must be in-service before replication is allowed.  <b>Action:</b> Determine the cause of the corruption and repair. This may involve bringing the resource in-service on another node. A full resync is most likely required once the data is repaired.
104237	WARN	Resource "%s" is "OSU". The mirror "%s" will wait to replicate data until all resources in the hierarchy%s are in-service. To replicate data immediately run: "%s/bin/mirror_action %s resume" on "%s" (see "LKDR_WAIT_TO_RESYNC" in /etc/default/LifeKeeper).	<b>Cause:</b> The specified resource is not in-service and is required before replication is resumed during a restore operation.  <b>Action:</b> Bring the required resources in-service to resume replication. Replication can also resume using the GUI command to resume or using the mirror_action command.
104238	ERROR	Unable to read \$nbd_taint_file. Assuming that the SIOS 'nbd' kernel module is not loaded.	<b>Cause:</b> The /sys/module/nbd/taint file could not be opened for reading.  <b>Action:</b> Verify that the /sys/module/nbd/taint file exists and is read-enabled.
104239	EMERG	\$failure_msg	<b>Cause:</b> At least one kernel module required by SIOS DataKeeper failed to load.

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Verify that the current running kernel is supported by this version of LifeKeeper for Linux. If the kernel was recently updated, re-run the LifeKeeper for Linux setup script located on the SIOS installation media to install kernel modules that are compatible with the current running kernel.</p>
104242	ERROR	Unable to read /proc/modules. Assuming that the SIOS '\$module' kernel module is not loaded.	<p><b>Cause:</b> Unable to read /proc/modules.</p> <p><b>Action:</b> Verify that the /proc/modules file exists and is read-enabled.</p>
104243	ERROR	Internal script or routine \$caller was called for unsupported kernel module '\$module'. Supported kernel modules for this script or routine are: \$module_list.	<p><b>Cause:</b> The given script or routine was called for an unsupported kernel module.</p> <p><b>Action:</b> This is an internal error. Please contact SIOS Customer Support.</p>
104244	ERROR	Output from 'modprobe \$module' command: \$pretty_modprobe_out	<p><b>Cause:</b> Failed to load the given kernel module using the modprobe command.</p> <p><b>Action:</b> Inspect the output provided in the log message from the failed modprobe attempt and resolve any issues found there.</p>
104251	ERROR	There is no LifeKeeper protected resource with tag \$tag on system \$me.	<p><b>Cause:</b> The given tag does not correspond to a LifeKeeper protected resource on the given system.</p> <p><b>Action:</b> Verify that the resource tag and system name are correct.</p>
104252	ERROR	Resource \$tag is not a \$app/\$typ resource. Please use the \$ins_app/\$ins_typ resource-specific canfailover script instead.	<p><b>Cause:</b> The scsi/netraid-specific canfailover script was called for a non-scsi/netraid resource.</p> <p><b>Action:</b> Use the canfailover script, if it exists, corresponding to the appropriate app and type of the given resource.</p>

## 8.1.2. DB2 Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
103001	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The db2nodes.cfg does not contain any server names.</p> <p><b>Action:</b> Ensure the db2nodes.cfg is valid.</p>
103002	ERROR	LifeKeeper was unable to get the version for the requested instance "%s"	<p><b>Cause:</b> "db2level" command did not return DB2 version.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103003	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103004	ERROR	Unable to get the information for resource "%s"	<p><b>Cause:</b> Failed to get resource information.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103005	ERROR	Unable to get the information for resource "%s"	<p><b>Cause:</b> Failed to get resource information.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103006	ERROR	Unable to get the instance information for resource "%s"	<p><b>Cause:</b> Failed to get the instance information.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check your LifeKeeper configuration.
103007	ERROR	Unable to get the instance home directory information for resource "%s"	<b>Cause:</b> Failed to get the instance home directory path. <b>Action:</b> Check your LifeKeeper configuration.
103008	ERROR	Unable to get the instance type information for resource "%s"	<b>Cause:</b> The DB2 Application Recovery Kit found invalid instance type. <b>Action:</b> Check your LifeKeeper configuration.
103009	ERROR	LifeKeeper has encountered an error while trying to get the database configuration parameters for database \"\$DB\"	<b>Cause:</b> There was an unexpected error running "db2 get db cfg for \$DB" command. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
103012	ERROR	LifeKeeper was unable to start the database server for instance "%s"	<b>Cause:</b> The requested startup of the DB2 instance failed. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "restore" operation.
103013	ERROR	LifeKeeper was unable to start the database server for instance "%s"	<b>Cause:</b> The requested startup of the DB2 instance failed. <b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "restore" operation.

Code	Severity	Message	Cause/Action
103015	ERROR	An entry for the home directory "%s" of instance "%s" does not exist in "/etc/fstab"	<p><b>Cause:</b> The home directory of instance of Multiple Partition database should exist in "/etc/fstab".</p> <p><b>Action:</b> Ensure the home directory exists in "/etc/fstab".</p>
103016	ERROR	LifeKeeper was unable to mount the home directory for the DB2 instance "%s"	<p><b>Cause:</b> Failed to mount the home directory of instance of Multiple Partition database.</p> <p><b>Action:</b> Ensure the home directory is mounted and retry the operation.</p>
103017	ERROR	Unable to get the instance nodes information for resource "%s"	<p><b>Cause:</b> Failed to get the instance nodes.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103018	ERROR	LifeKeeper was unable to start database partition server "%s" for instance "%s"	<p><b>Cause:</b> The requested startup of the DB2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "restore" operation.</p>
103020	ERROR	LifeKeeper was unable to stop the database server for instance "%s"	<p><b>Cause:</b> The requested shutdown of the DB2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "remove" operation.</p>
103021	ERROR	LifeKeeper was unable to stop the database server for instance "%s"	<p><b>Cause:</b> The requested shutdown of the DB2 instance failed.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "remove" operation.</p>
103023	ERROR	Unable to get the instance nodes information for resource "%s"	<p><b>Cause:</b> Failed to get the instance nodes.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103024	ERROR	LifeKeeper was unable to stop database partition server "%s" for instance "%s"	<p><b>Cause:</b> The requested shutdown of the DB2 instance failed.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "remove" operation.</p>
103026	ERROR	Unable to get the instance nodes information for resource "%s"	<p><b>Cause:</b> Failed to get the instance nodes.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103027	FATAL	The argument for the DB2 instance is empty	<p><b>Cause:</b> Invalid parameters were specified for the create operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103028	FATAL	Unable to determine the DB2 instance home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>

Code	Severity	Message	Cause/Action
103029	FATAL	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103030	FATAL	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103031	ERROR	The path "%s" is not on a shared filesystem	<p><b>Cause:</b> The instance home directory should be on a shared filesystem.</p> <p><b>Action:</b> Ensure the path is on shared filesystem and retry the create operation.</p>
103032	ERROR	LifeKeeper was unable to get the DB tablespace containers for instance "%s" or the log path for one of its databases	<p><b>Cause:</b> LifeKeeper could not determine the location of the database table space containers or verify that they are located in a path which is on a mounted filesystem.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem. Correct any reported errors before retrying the "create" operation.</p>
103033	ERROR	The path "%s" is not on a shared filesystem	<p><b>Cause:</b> The path of database table space container should be on a shared filesystem.</p> <p><b>Action:</b> Ensure database table space container is on a shared filesystem and retry the operation.</p>

Code	Severity	Message	Cause/Action
103034	ERROR	A DB2 Hierarchy already exists for instance "%s"	<p><b>Cause:</b> An attempt was made to protect the DB2 instance that is already under LifeKeeper protection.</p> <p><b>Action:</b> You must select a different DB2 instance for LifeKeeper protection.</p>
103035	ERROR	The file system resource "%s" is not in-service	<p><b>Cause:</b> The file system which the DB2 resource depends on should be in service.</p> <p><b>Action:</b> Ensure the file system resource is in service and retry the "create" operation.</p>
103036	ERROR	Unable to create the hierarchy for raw device "%s"	<p><b>Cause:</b> LifeKeeper was unable to create the resource {raw device} .</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the "create" operation.</p>
103037	ERROR	A RAW hierarchy does not exist for the tag "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the raw resource {tag} .</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103038	ERROR	LifeKeeper was unable to create a dependency between the DB2 hierarchy "%s" and the Raw hierarchy "%s"	<p><b>Cause:</b> The requested dependency creation between the parent DB2 resource and the child Raw resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the "create" operation.</p>

Code	Severity	Message	Cause/Action
103039	ERROR	LifeKeeper could not disable the automatic startup feature of DB2 instance "%s"	<p><b>Cause:</b> An unexpected error occurred while attempting to update the DB2 setting.</p> <p><b>Action:</b> The DB2AUTOSTART will need to be updated manually to turn off the automatic startup of the instance at system boot.</p>
103040	ERROR	DB2 version "%s" is not installed on server "%s"	<p><b>Cause:</b> LifeKeeper could not find DB2 installed location.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103041	ERROR	The instance owner "%s" does not exist on target server "%s"	<p><b>Cause:</b> An attempt to retrieve the DB2 instance owner from template server during a "canextend" or "extend" operation failed.</p> <p><b>Action:</b> Verify the DB2 instance owner exists on the specified server. If the user does not exist, it should be created with the same uid and gid on all servers in the cluster.</p>
103042	ERROR	The instance owner "%s" uids are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The user id on the target server {target server} for the DB2 instance owner {user} does not match the value of the user {user} on the template server {template server}.</p> <p><b>Action:</b> The user ids for the DB2 instance owner {user} must match on all servers in the cluster. The user id mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103043	ERROR	The instance owner "%s" gids are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The group id on the target</p>

Code	Severity	Message	Cause/Action
			<p>server {target server} for the DB2 instance owner {user} does not match the value of the user {user} on the template server {template server}.</p> <p><b>Action:</b> The group ids for the DB2 instance owner {user} must match on all servers in the cluster. The group id mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103044	ERROR	The instance owner "%s" home directories are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The home directory location of the user {user} on the target server {target server} does not match the DB2 instance owner's home directory on the template server {template server}.</p> <p><b>Action:</b> The home directory location of the DB2 instance owner {user} must match on all servers in the cluster. The location mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103045	ERROR	LifeKeeper was unable to get the DB2 "SVCENAME" parameter for the DB2 instance	<p><b>Cause:</b> There was an unexpected error running "db2 get dbm cfg" command.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103046	ERROR	Unable to get the value of the DB2 "SVCENAME" parameter for the DB2 instance %s.	<p><b>Cause:</b> The DB2 "SVCENAME" parameter is set to null.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103047	ERROR	LifeKeeper was unable to get the contents of the "/etc/services" file on the server "%s"	<p><b>Cause:</b> "/etc/services" on the template server does not contain the service names for the DB2 instance.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> The service names in "/etc/services" for the DB2 instance must match on all servers in the cluster. The service names mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103048	ERROR	LifeKeeper was unable to get the contents of the "/etc/services" file on the server "%s"	<p><b>Cause:</b> "/etc/services" on the target server does not contain the service names for the DB2 instance.</p> <p><b>Action:</b> The service names in "/etc/services" for the DB2 instance must match on all servers in the cluster. The service names mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103049	ERROR	The "/etc/services" entries for the instance "%s" are different on target server "%s" and template server "%s"	<p><b>Cause:</b> The "/etc/services" entries for the instance are mismatched.</p> <p><b>Action:</b> The service names in "/etc/services" for the DB2 instance must match on all servers in the cluster. The service names mismatch should be corrected manually on all servers before retrying the "canextend" operation.</p>
103050	ERROR	The home directory "%s" for instance "%s" is not mounted on server "%s"	<p><b>Cause:</b> LifeKeeper could not find db2nodes.cfg for Multiple Partition instance.</p> <p><b>Action:</b> Ensure the home directory is mounted and retry the operation.</p>
103051	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Failed to get resource</p>

Code	Severity	Message	Cause/Action
			<p>information from the template server.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
103052	ERROR	LifeKeeper was unable to add instance "%s" and/or its variables to the DB2 registry	<p><b>Cause:</b> There was an unexpected error running "db2iset" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103053	ERROR	Usage: %s instance	
103054	ERROR	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103055	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103056	ERROR	Usage: %s instance	
103058	ERROR	Usage: %s instance	
103059	ERROR	Usage: %s instance	
103060	ERROR	Unable to determine the DB2 instance home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>

Code	Severity	Message	Cause/Action
103061	ERROR	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103062	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find the node for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103063	ERROR	Unable to determine the DB2 install path	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find DB2 for the instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103064	ERROR	Usage: nodes -t tag -a add_nodenum   nodes -t tag -d delete_nodenum   nodes -t tag -p	
103065	ERROR	Invalid input provided for "%s" utility operation, characters are not allowed.	<p><b>Cause:</b> Invalid parameters were specified for the "nodes" command.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103066	ERROR	Unable to get the information for resource "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag}.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103067	ERROR	The DB2 instance "%s" is not a EEE or Multiple Partition instance	<p><b>Cause:</b> The resource {tag} is single partition instance.</p> <p><b>Action:</b> Verify the parameters and retry</p>

Code	Severity	Message	Cause/Action
			the operation.
103069	ERROR	Node "%s" is already protected by this hierarchy	<p><b>Cause:</b> Invalid parameters were specified for the "nodes" command.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103070	ERROR	Node number "%s" is the last remaining node protected by resource "%s". Deleting all nodes is not allowed.	<p><b>Cause:</b> Invalid parameters were specified for the "nodes" command.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
103071	ERROR	LifeKeeper is unable to get the equivalent instance for resource "%s"	<p><b>Cause:</b> There was an unexpected error running "nodes" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103072	ERROR	Unable to set NodesInfo for resource "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "nodes" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103073	ERROR	Unable to set NodesInfo for resource "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "nodes" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
103074	ERROR	Usage: %s instance	
103075	ERROR	Usage: %s instance	

Code	Severity	Message	Cause/Action
103076	ERROR	Unable to determine the DB2 instance type	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance type.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103077	ERROR	Unable to determine the DB2 instance home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>
103078	ERROR	The database server is not running for instance "%s"	<p><b>Cause:</b> A process check for the DB2 instance did not find any processes running.</p> <p><b>Action:</b> The DB2 instance must be started.</p>
103079	ERROR	LifeKeeper has detected an error while trying to determine the node number(s) of the DB partition server(s) for the instance	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to find any nodes for the DB2 instance.</p> <p><b>Action:</b> Check your DB2 configuration.</p>
103080	ERROR	One or more of the database partition servers for instance "%s" is down	<p><b>Cause:</b> All database partition servers should be running.</p> <p><b>Action:</b> Ensure all database partition servers are running and retry the operation.</p>
103081	ERROR	DB2 local recovery detected another recovery process in progress for "%s" and will exit.	

Code	Severity	Message	Cause/Action
103082	ERROR	Failed to create flag "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create a flag for controlling DB2 local recovery processing.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
103083	ERROR	Failed to remove flag "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to remove a flag for controlling DB2 local recovery processing.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
103084	ERROR	Unable to determine the DB2 instance \"\${Instance}\" home directory	<p><b>Cause:</b> The DB2 Application Recovery Kit was unable to determine the DB2 instance home directory.</p> <p><b>Action:</b> Ensure the instance owner name is same as the instance name and retry the operation.</p>

## 8.1.3. DMMP Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
128005	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The quickCheck of {resource} on {server} failed due to an operating system signal {signal}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128008	ERROR	Usage: quickCheck -t <tag name> -i <id>	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device quickCheck command preventing it from running.</p> <p><b>Action:</b> Make sure all software components are properly installed and at the correct version. Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource to be quickChecked.</p>
128010	ERROR	quickCheck for "%s" failed checks of underlying paths, initiate recovery. retry count=%s.	<p><b>Cause:</b> The dmmp kit failed to quickCheck a device after {count} times of retries. A recovery of the protected dmmp resource will be executed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128021	ERROR	unable to find device for uuid "%s".	<p><b>Cause:</b> The device could not be found</p>

Code	Severity	Message	Cause/Action
			<p>by unique id during a restore operation.</p> <p><b>Action:</b> Verify that the resource is configured properly. Rerun the command and supply the correct device id that identifies the dmmp device resource to be restored.</p>
128025	ERROR	Device "%s" failed to unlock.	<p><b>Cause:</b> A non working {device} was detected and could not be unlocked during the restore operation.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128026	ERROR	Device "%s" failed to lock.	<p><b>Cause:</b> The {device} could not be locked during the restore.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128031	ERROR	unable to find device for uuid "%s".	<p><b>Cause:</b> The device could not be found by unique id during the remove operation.</p> <p><b>Action:</b> Verify that the resource is configured properly. Rerun the command and supply the correct device id that identifies the dmmp device resource to be removed.</p>
128034	ERROR	Device "%s" failed to unlock.	<p><b>Cause:</b> The {device} could not be unlocked during the remove.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128036	ERROR	unable to load existing information for device with uuid "%s".	<p><b>Cause:</b> The device information could not be loaded by unique id.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct device id that identifies the dmmp device resource.</p>
128037	ERROR	unable to load existing information for device "%s".	<p><b>Cause:</b> The device information could not be loaded by name.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct device name that identifies the dmmp device resource.</p>
128038	ERROR	unable to load existing information for device, no dev or uuid defined.	<p><b>Cause:</b> The device information could not be loaded since neither a unique device id nor name of the device were defined.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct device id or name that identifies the dmmp device resource.</p>
128041	ERROR	unable to load existing information for device with uuid "%s".	<p><b>Cause:</b> The device information could not be loaded by unique id.</p> <p><b>Action:</b> Make sure the resource is</p>

Code	Severity	Message	Cause/Action
			<p>configured properly. Rerun the command and supply the correct device id that identifies the dmmp device resource.</p>
128057	ERROR	All paths are failed on "%s".	<p><b>Cause:</b> LifeKeeper detected all paths listed to the protected dmmp device are in the failed state.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128058	ERROR	could not determine registrations for "%s"! All paths failed.	<p><b>Cause:</b> LifeKeeper could not determine registrations for protected dmmp {device}. All paths to the dmmp {device} are in the failed state.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128059	WARN	path "%s" no longer configured for "%s", remove from path list.	<p><b>Cause:</b> LifeKeeper detected listed {path} to protected {device} is not valid anymore and will remove it from the path list.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128060	WARN	registration failed on path "%s" for "%s".	<p><b>Cause:</b> LifeKeeper failed the registration on {path} for protected dmmp {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>

Code	Severity	Message	Cause/Action
128062	ERROR	all paths failed for "%s".	<p><b>Cause:</b> LifeKeeper failed to verify a valid path to protected dmp {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128072	ERROR	The daemon "%s" does not appear to be running and could not be restarted. Path failures may not be correctly handled without this daemon.	<p><b>Cause:</b> LifeKeeper failed to verify dmp daemon is running and could not restart it.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128078	ERROR	"%s" resource type is not installed on "%s".	<p><b>Cause:</b> The Device Mapper Multipath Recovery Kit for dmp device support is not installed on the system.</p> <p><b>Action:</b> Install the steeleye-ikDMMP Device Mapper Multipath Recovery Kit rpm on the system.</p>
128083	ERROR	This script must be executed on "%s".	<p><b>Cause:</b> An incorrect system name has been supplied as an argument to the devicehier script used to create the dmp device resource.</p> <p><b>Action:</b> Make sure the cluster nodes and comm-paths are properly configured. Supply the correct system name to the devicehier script. The name must match the name of the system on which the command is run.</p>
128084	ERROR	The device %s is not active.	<p><b>Cause:</b> LifeKeeper failed to find the specified {device} as a valid device on the system during resource creation.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and related messages. Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource to be created.</p>
128086	ERROR	Failed to create "%s" hierarchy.	<p><b>Cause:</b> LifeKeeper failed to create resource hierarchy for {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128088	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper failed to create the resource with {tagname} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128090	ERROR	Failed to create dependency "%s"-"%s" on system "%s".	<p><b>Cause:</b> LifeKeeper failed to create dependency {resource tag name} – {resource tag name} on {system} during creation.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128091	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper failed to create {resource} on {system}.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128101	ERROR	"%s" constructor requires a valid argument.	<p><b>Cause:</b> LifeKeeper failed to create an object for the dmmp resource during construction.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource.</p>
128102	ERROR	Invalid tag "%s".	<p><b>Cause:</b> A resource instance could not be found for the given tag name.</p> <p><b>Action:</b> Make sure the resource is configured properly. Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource.</p>
128111	ERROR	Failed to get registrations for "%s": %s. Verify the storage supports persistent reservations.	<p><b>Cause:</b> LifeKeeper failed to get the registrations of {device} with the message, "bad field in Persistent reservation in cdb".</p> <p><b>Action:</b> Verify if the storage supports persistent reservations. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128112	ERROR	Failed to get registrations for "%s": %s. Verify the storage supports persistent reservations.	<p><b>Cause:</b> LifeKeeper failed to get the registrations of {device} with the</p>

Code	Severity	Message	Cause/Action
			<p>message, "illegal request".</p> <p><b>Action:</b> Verify if the storage supports persistent reservations. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128136	ERROR	A previous quickCheck with PID "%s" running for device "%s" has been terminated.	<p><b>Cause:</b> LifeKeeper detected that a previous quickCheck operation is still running during the dmmp resource restore operation. It has been terminated by LifeKeeper.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128137	ERROR	SCSI reservation conflict on %s during LifeKeeper resource initialization. Manual intervention required.	<p><b>Cause:</b> LifeKeeper detected a SCSI reservation conflict on {device} during dmmp resource restore.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Manual intervention and fix of the reservation conflict on {device} is required.</p>
128138	ERROR	unable to clear registrations on %s.	<p><b>Cause:</b> LifeKeeper failed to clear all the registrations on {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128140	WARN	registration failed on path %s for %s.	<p><b>Cause:</b> LifeKeeper failed to make the</p>

Code	Severity	Message	Cause/Action
			<p>registratioin on {path} for {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128143	ERROR	reserve failed (%d) on %s.	<p><b>Cause:</b> LifeKeeper failed to make reservation for {resource} on {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128145	ERROR	The server ID "%s" returned by "%s" is not valid.	<p><b>Cause:</b> LifeKeeper failed to generate a valid host {id}.</p> <p><b>Action:</b> The ID used to register a device is made up of 1 to 12 Hex digits that uniquely identifies the server in the cluster. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128146	ERROR	device failure on %s. SYSTEM HALTED.	<p><b>Cause:</b> LifeKeeper detected failure on {device} and will reboot the server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128148	ERROR	device failure on %s. SYSTEM HALTED DISABLED.	<p><b>Cause:</b> LifeKeeper detected a failure on {device}. The reboot was skipped due to LifeKeeper configuration.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and related messages. Turn on the configuration "SCSIHALT" to make the reboot available for any detected device failure.</p>
128149	ERROR	device failure or SCSI Error on %s. SENDEVENT DISABLED.	<p><b>Cause:</b> LifeKeeper detected a failure on {device}. The event generation was skipped due to LifeKeeper configuration.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Turn on the configuration "SCSIEVENT" to make the sendevent available for any detected device failure.</p>
128150	ERROR	%s does not have EXCLUSIVE access to %s, halt system.	<p><b>Cause:</b> LifeKeeper detected a reservation conflict for {device} on {server} and will reboot the server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
128151	ERROR	%s does not have EXCLUSIVE access to %s, halt system DISABLED.	<p><b>Cause:</b> LifeKeeper detected a reservation conflict for {device} on {server}. The reboot was skipped due to LifeKeeper configuration.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Turn on the configuration "RESERVATIONCONFLICT" to make the reboot available for any detected reservation conflicts.</p>

Code	Severity	Message	Cause/Action
128154	WARN	unable to flush buffers on %s.	<p><b>Cause:</b> LifeKeeper failed to flush the buffers for {device} during dmmp resource remove.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128157	WARN	%s utility not found, limited healthcheck for %s.	<p><b>Cause:</b> LifeKeeper failed to find "dd" utility for the health check of {device}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128160	ERROR	%s failed to read %s.	<p><b>Cause:</b> LifeKeeper failed a disk I/O test for {device} when using {utility}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128163	ERROR	Registration ID "%s" for "%s" is not valid.	<p><b>Cause:</b> LifeKeeper failed to generate a valid registration {id} for {device}.</p> <p><b>Action:</b> The ID used to register a device is made up of 4 Hex digits derived from the path to the device. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128170	ERROR	Usage: canextend <Template system name> <Template tag name>	

Code	Severity	Message	Cause/Action
128500	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device restore command preventing it from running.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;Resource Tag&gt; and -i &lt;Resource ID&gt; that identifies the dmmp device resource to be restored.</p>
128504	ERROR	"%s" resource type is not installed on "%s".	<p><b>Cause:</b> The Device Mapper Multipath Recovery Kit for dmmp device support is not installed on the system.</p> <p><b>Action:</b> Install the steeleye-ikDMMP Device Mapper Multipath Recovery Kit rpm on the system.</p>
128506	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device devShared command preventing it from running.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: &lt;Template Resource System Name&gt; and &lt;Template Resource Tag&gt; that identifies the dmmp device resource to be created.</p>
128507	FATAL	This script must be executed on "%s".	<p><b>Cause:</b> An incorrect system name has been supplied as an argument to the devicehier script used to create the dmmp device resource.</p> <p><b>Action:</b> Supply the correct system name to the devicehier script. The name must match the name of the system on which the command is run.</p>

Code	Severity	Message	Cause/Action
128511	ERROR	Failed to get the ID for the device "%s". Hierarchy create failed.	<p><b>Cause:</b> The devicehier script used to create the dmmp device resource was unable to determine the SCSI ID for the supplied device.</p> <p><b>Action:</b> Check that the supplied device path exists and that is for a supported SCSI storage array.</p>
128512	ERROR	Failed to get the disk ID for the device "%s". Hierarchy create failed.	<p><b>Cause:</b> The devicehier script used to create the dmmp disk resource was unable to determine the SCSI ID for the supplied disk.</p> <p><b>Action:</b> Check that the supplied device path exists and that is for a supported SCSI storage array.</p>
128513	ERROR	Failed to create the underlying resource for device "%s". Hierarchy create failed.	<p><b>Cause:</b> The creation of the underlying dmmp disk resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128515	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The creation of the dmmp device resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128517	ERROR	Failed to create dependency "%s"-"%s" on system "%s".	<p><b>Cause:</b> The parent child dependency creation between the dmmp device and dmmp disk resources failed.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128519	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The attempt to bring the newly created dmmp device resource in service has failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
128521	ERROR	Either TEMPLATESYS or TEMPLATETAG argument missing	<p><b>Cause:</b> Incorrect arguments have been supplied to the extend command for the dmmp device resource.</p> <p><b>Action:</b> Rerun the dmmp device resource extend and supply the correct template system and tag names.</p>
128540	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the dmmp device getid command used to retrieve the SCSI ID.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -i &lt;device path&gt; or -b &lt;device ID&gt;.</p>
128541	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the command used to delete the dmmp device resource.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;dmmp device resource tag&gt;.</p>

Code	Severity	Message	Cause/Action
128543	ERROR	device node \"\$dev\" does not exist.	<p><b>Cause:</b> The device node required for restoring the dmmp device resource does not exist. The allocated wait time in restore for udev device creation has been exceed.</p> <p><b>Action:</b> Rerun the dmmp device resource restore once udev has created the device.</p>
128544	ERROR	Usage error	<p><b>Cause:</b> Incorrect arguments have been supplied to the remove command used to take the dmmp device resource out of service.</p> <p><b>Action:</b> Rerun the command and supply the correct argument list: -t &lt;dmmp device resource tag&gt;.</p>
128550	ERROR	Failed to get path information.(tag=\"\$opt_t\")	
128551	ERROR	Failed to get the status of path.(tag=\"\$opt_t\")	
128552	ERROR	Failed to check the status of path.(tag=\"\$opt_t\")	
128553	ERROR	Failed to get the resource information.(tag=\"\$opt_t\")	

## 8.1.4. Recovery Kit for EC2 Message Catalog

The Recovery Kit for EC2 Message Catalog below contains listings of all messages that may be encountered while utilizing the Recovery Kit for EC2.

The [Combined Message Catalog](#) provides a listing of all messages that may be encountered while using LifeKeeper for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received,

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
129100	FATAL	Failed to load instance from LifeKeeper.	<p><b>Cause:</b> An invalid resource tag or ID was specified.</p> <p><b>Action:</b> Check that the tag or ID is valid and re-run the command.</p>
129103	FATAL	No resource matches tag \"\${self->{tag}}\".	<p><b>Cause:</b> An invalid resource tag was specified.</p> <p><b>Action:</b> Check the tag and re-run the command.</p>
129104	FATAL	An error occurred setting LifeKeeper resource information	<p><b>Cause:</b> An internal error has occurred in LifeKeeper.</p>
129110	ERROR	Could not get the Elastic Network Interface ID for \$dev	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129111	ERROR	Failed to get Allocation ID of Elastic IP \"\${elasticIp}\".	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129113	ERROR	Failed to get my instance ID.	<p><b>Cause:</b> The EC2 instance metadata access failed.</p> <p><b>Action:</b> Check the Amazon console and retry the operation.</p>
129114	ERROR	Failed to get ENI ID.	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129116	ERROR	Failed to associate Elastic IP \"\$self->{'EIP'}\" on \"\$self->{'DEV'}\".	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129118	WARN	\$self->{'EIP'} is not associated with any instance.	<p><b>Cause:</b> The Elastic IP is not associated with any instance.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check adjacent log messages for more details.</p>
129119	WARN	\$self->{'EIP'} is associated with another instance.	<p><b>Cause:</b> The Elastic IP is associated with another instance.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check</p>

Code	Severity	Message	Cause/Action
			adjacent log messages for more details.
129120	ERROR	Failed to recover Elastic IP.	<p><b>Cause:</b> The EC2 API call failed to associate the Elastic IP.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129121	ERROR	Recovery process ended but Elastic IP is not associated with this instance. Please check AWS console.	<p><b>Cause:</b> The EC2 API call failed to associate the Elastic IP.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129122	ERROR	Error creating resource \"\$target_tag\" with return code of \"\$err\".	<p><b>Cause:</b> LifeKeeper was unable to create the resource instance on the server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
129123	ERROR	Failed to get ENI ID.	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129124	WARN	\$self->{'EIP'} is associated with another network interface.	<p><b>Cause:</b> The Elastic IP is associated with the proper instance, but the wrong ENI.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check</p>

Code	Severity	Message	Cause/Action
			adjacent log messages for more details.
129125	ERROR	Link check failed for interface '\\$dev'.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present.</p> <p><b>Action:</b> Check the network interface and bring the link up.</p>
129126	ERROR	Link check failed for interface '\\$dev'. Reason: down slave.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present.</p> <p><b>Action:</b> Check the network interface and bring the link up.</p>
129129	WARN	The link for network interface '\\$self->{'DEV'}' is down. Attempting to bring the link up.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by bringing the link up and associating the Elastic IP with the interface. Check adjacent log messages for more details.</p>
129130	ERROR	Failed to modify '\\$opt_t' to end pint URL '\\$endpoint'.	
129137	ERROR	The link for network interface '\\$self->{'DEV'}' is still down.	<p><b>Cause:</b> LifeKeeper could not bring the link up.</p> <p><b>Action:</b> Ensure the interface is enabled and up. Check adjacent log messages for more details.</p>
129139	WARN	The link for network interface '\\$self->{'DEV'}' is down.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that</p>

Code	Severity	Message	Cause/Action
			<p>no link is present.</p> <p><b>Action:</b> Check the network interface and bring the link up.</p>
129140	ERROR	Could not get ENI ID for \$self->{IP}.	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129142	ERROR	Failed to update route table	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129143	ERROR	You must have exactly one IP address resource as the parent of the RouteTable EC2 resource. Please reconfigure your resource hierarchy.	<p><b>Cause:</b> The Route Table EC2 resource must have one and only one IP resource as a parent.</p> <p><b>Action:</b> Repair the resource hierarchy as necessary.</p>
129144	ERROR	\$func called with invalid timeout: \$timeout	<p><b>Cause:</b> An invalid timeout value was specified in the /etc/default/LifeKeeper file.</p> <p><b>Action:</b> Verify all EC2_*_TIMEOUT settings are valid in /etc/default/LifeKeeper.</p>
129145	ERROR	\$func action timed out after \$timeout seconds	<p><b>Cause:</b> The action did not complete within the timeout period.</p> <p><b>Action:</b> Consider increasing the</p>

Code	Severity	Message	Cause/Action
			<p>EC2_*_TIMEOUT value for the given action (in /etc/default/LifeKeeper).</p>
129146	ERROR	<p>failed to run \$func with timeout: \$@</p>	<p><b>Cause:</b> This is an internal error.</p>
129148	ERROR	<p>Amazon describe-route-tables call failed (err=%s)(output=%s                      It is recommended to confirm the following.</p> <p>_____</p> <p>Verify that you're running the most recent AWS CLI version                      -&gt; Make sure that you're using the most recent version of the AWS CLI.</p> <p>Unable to locate credentials                      -&gt; Verify that the AWS CLI is installed and configured correctly.</p> <p>An error occurred (UnauthorizedOperation) and (AuthFailure)                      -&gt; Make sure that the AWS IAM role or IAM user has the correct permissions to run the relevant commands.                      -&gt; Make sure that the time on your Linux or Windows instance is correct.                      The instance's UTC time should match to AWS UTC time. (The time zone setting is irrelevant.)                      -&gt; Make sure that you're using the correct Amazon Simple Token Service (AWS STS) token format.                      -&gt; Make sure that you're using the correct credentials to make the API call.                      If there are multiple sets of credentials on the instance, credential precedence might affect which credentials the instance uses to make the API call.                      Verify which set of credentials you're using by running the aws sts get-caller-identity command.</p>	<p><b>Cause:</b> Failed to call the AWS CLI command.</p> <p><b>Action:</b> Check your environment by referring to the displayed troubleshooting information.</p>

Code	Severity	Message	Cause/Action
		<p>An error occurred (ExpiredToken)                      -&gt; Temporary credentials expire at the time interval specified during creation.</p> <p>If the credentials for your IAM role are expired, obtain a new STS token by assuming a new IAM role.</p> <p>If your error does not match the above errors, you can get more information by using the following on the command line. (It is the same with the command line and the current error, add the '– –debug' option.)                      # [Execute the AWS command + option]                      – –debug</p>	
129150	ERROR	Elastic IP \"\$elasticIp\" is associated with another instance.	<p><b>Cause:</b> The Elastic IP is not associated with the proper instance.</p> <p><b>Action:</b> LifeKeeper will try to fix the issue by calling the EC2 API to associate the Elastic IP. Check adjacent log messages for more details.</p>
129151	ERROR	Could not get the Association ID for Elastic IP \"\$elasticIp\".	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129152	ERROR	Failed to disassociate Elastic IP \"\$self->{'EIP'}\" on \"\$self->{'DEV'}\".	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129153	ERROR	Failed to disassociate Elastic IP \"\$elasticIp\", (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129154	ERROR	Amazon describe-addresses call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129155	ERROR	Amazon describe-address call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129157	ERROR	curl call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 instance metadata access failed.</p> <p><b>Action:</b> Check the Amazon console and retry the operation.</p>
129159	ERROR	Amazon associate-address call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>
129160	ERROR	Amazon describe-addresses call failed (err=%s)(output=%s	<p><b>Cause:</b> The EC2 API call failed, possibly due to a network issue.</p> <p><b>Action:</b> Check the network and the Amazon console and retry the operation.</p>

Code	Severity	Message	Cause/Action
129161	ERROR	Error deleting resource \"\$otherTag\" on \"\$otherSys\" with return code of \"\$err\".	<p><b>Cause:</b> LifeKeeper was unable to delete the resource instance on the server.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
129162	ERROR	Could not getRouteTablesByIP	
129163	ERROR	Could not getRouteTablesByIP	
129164	ERROR	[\$SUBJECT event] mail returned \$err	
129165	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
129167	ERROR	snmptrap returned \$err for Trap 180	
129168	ERROR	COMMAND OUTPUT: cat /tmp/err\$\$	
129170	ERROR	This resource is in the old format. Please update.	
129171	ERROR	Error disabling sourceDest checks with return code of \"\$ret\".	
129172	ERROR	Failed to lookup IP from resource \$self->{'tag'}.	
129173	ERROR	Failed to lookup device GUID for IP resource \$iptag.	
129174	ERROR	Could not find eni_id for device GUID \$dev.	
129175	ERROR	Failed to disable sourceDestChecks. (err=%s)(output=%s	
129176	ERROR	Failed to disable sourceDestChecks for \$eni_id. (err=%s)(output=%s	
129177	ERROR	Error disabling sourceDest checks with return code of \"\$ret\".	
129178	ERROR	Failed to disable sourceDestChecks for \$eni_id. (err=%s)(output=%s	
129180	ERROR	Elastic Network Interface \$eni sourceDestChecks are enabled.	
129181	WARN	WARNING Failed to find IP resources since EC2 resource is extended first.	

Code	Severity	Message	Cause/Action
		Skipping sourceDestChecks.	
129182	WARN	WARNING SourceDestCheck was skipped because Amazon describe-network-interface-attribute call failed. Please check if ec2:DescribeNetworkInterfaceAttribute and ec2:ModifyNetworkInterfaceAttribute are granted.	
129403	ERROR	END failed create of \$TAG due to a \$sig signal	<b>Cause:</b> The create process was interrupted by a signal.
129409	ERROR	The IP resource \$IP_RES is not \"ISP\".	<b>Cause:</b> The IP resource is not in service. <b>Action:</b> Bring the resource in service and retry the operation.
129410	ERROR	Could not find IP resource \$IP_RES	<b>Cause:</b> Ensure that the IP resource exists and retry the operation.
129412	ERROR	EC2 resource \$ID is already protected	<b>Cause:</b> A resource with the specified ID already exists. <b>Action:</b> Make sure to clean up any remnants of an old resource before re-creating a new resource.
129416	ERROR	Error creating resource \"\$TAG\" with return code of \"\$!cderror\".	<b>Cause:</b> LifeKeeper was unable to create the resource instance. <b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.
129418	ERROR	Dependency creation between \"\$IP_RES\" and \"\$TAG\" failed with return code of \"\$!cderror\".	<b>Cause:</b> LifeKeeper was unable to create the resource dependency.

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.
129420	ERROR	In-service failed for tag \" <code>\$TAG</code> \".	<b>Cause:</b> LifeKeeper could not bring the resource instance into service.  <b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.
129423	ERROR	Could not get ENI ID for <code>\$dev</code> .	
129425	ERROR	Failed to update route table	
129426	ERROR	In-service (dummy) failed for tag \" <code>\$TAG</code> \".	
129800	ERROR	canextend checks failed for \" <code>\$self-&gt;{tag}</code> \" (err= <code>\$ret</code> )	<b>Cause:</b> The pre-extend checks failed for the target server.  <b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.
129801	ERROR	canextend checks failed for \" <code>\$self-&gt;{tag}</code> \". <code>EC2_HOME</code> \" <code>\$self-&gt;{EC2_HOME}</code> \" does not exist on <code>\$me</code> .	

## 8.1.5. File System Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
124004	FATAL	resource tag name not specified	<p><b>Cause:</b> Invalid arguments were specified for the "quickCheck" operation.</p> <p><b>Action:</b> Ensure that the correct arguments are passed.</p>
124005	FATAL	resource id not specified	<p><b>Cause:</b> Invalid arguments were specified for the "quickCheck" operation.</p> <p><b>Action:</b> Ensure that the correct arguments are passed.</p>
124007	FATAL	Failed to get resource information	<p><b>Cause:</b> The filesystem resource's info field does not contain the correct information.</p> <p><b>Action:</b> Put the correct information in the resource's info field or restore the system from a recent "lkbackup" to restore the original info field.</p>
124008	ERROR	getld failed	<p><b>Cause:</b> The filesystem resource could not find the underlying disk device.</p> <p><b>Action:</b> Check adjacent log messages for further details. Verify that the resource hierarchy is valid and that all required storage kits are installed.</p>
124009	ERROR	LifeKeeper protected filesystem is in service but quickCheck detects the following error	<p><b>Cause:</b> The filesystem kit has found</p>

Code	Severity	Message	Cause/Action
			<p>something wrong with the resource.</p> <p><b>Action:</b> Check the messages immediately following this one for more details.</p>
124010	ERROR	\"\$id\" is not mounted	<p><b>Cause:</b> The filesystem resource is no longer mounted.</p> <p><b>Action:</b> No action is required. Allow local recovery to remount the resource.</p>
124011	ERROR	\"\$id\" is mounted but with the incorrect mount options (current mount option list: \$mntopts, expected mount option list: \$infopts	<p><b>Cause:</b> The filesystem resource is mounted incorrectly.</p> <p><b>Action:</b> No action is required. Allow local recovery to remount the resource.</p>
124012	ERROR	\"\$id\" is mounted but on the wrong device (current mount device: \$tmpdev, expected mount device: \$dev	<p><b>Cause:</b> The filesystem resource has the wrong device mounted.</p> <p><b>Action:</b> No action is required. Allow local recovery to remount the resource.</p>
124015	ERROR	LifeKeeper protected filesystem \"\$tag\" (\$id) is \$percent% full (\$blocksfree free blocks).	<p><b>Cause:</b> The filesystem is getting full.</p> <p><b>Action:</b> Remove or migrate data from the filesystem.</p>
124016	WARN	LifeKeeper protected filesystem \"\$tag\" (\$id) is \$percent% full (\$blocksfree free blocks).	<p><b>Cause:</b> The filesystem is getting full.</p> <p><b>Action:</b> Remove or migrate data from the filesystem.</p>
124020	FATAL	cannot find device information for filesystem \$id	<p><b>Cause:</b> The filesystem resource could not find the underlying disk device.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details. Verify that the resource hierarchy is valid and that all required storage kits are installed.</p>
124029	ERROR	Failed to find child resource.	<p><b>Cause:</b> The filesystem resource could not determine its underlying disk resource.</p> <p><b>Action:</b> Ensure that the resource hierarchy is correct.</p>
124032	FATAL	Script has hung. Exiting.	<p><b>Cause:</b> Processes had files open on a mounted filesystem that needed to be unmounted. Killing those processes has taken too long.</p> <p><b>Action:</b> If this error continues, try to temporarily stop all software that may be using the mount point to allow it to be unmounted. If the filesystem still cannot be unmounted, contact Support.</p>
124042	ERROR	file system \$fs failed unmount; will try again	<p><b>Cause:</b> Processes had files open on a mounted filesystem that needed to be unmounted. It can take multiple attempts to clear those processes.</p> <p><b>Action:</b> No action is required. Allow the process to continue.</p>
124046	ERROR	file system \$fsname failed unmount	<p><b>Cause:</b> A filesystem could not be unmounted.</p> <p><b>Action:</b> If this error continues, try to temporarily stop all software that may be using the mount point to allow it to be unmounted. If the filesystem still cannot be unmounted, contact Support.</p>

Code	Severity	Message	Cause/Action
124049	ERROR	Local recovery of resource has failed (err=\$err)	<p><b>Cause:</b> A filesystem resource has a problem that cannot be repaired locally.</p> <p><b>Action:</b> No action is required. Allow the resource to be failed over to another system.</p>
124051	WARN	getld failed, try count : \$cnt/\$try	
124052	ERROR	\"\$id\" is mounted but filesystem is shutdown state.	
124054	ERROR	Failed to change mount option from \"\$old_opts\" to \"\$new_opts\" for migration	
124103	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124104	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Invalid arguments were specified for the canextend operation.</p> <p><b>Action:</b> Ensure that the arguments are correct. If this error happens during normal operation, please contact Support.</p>
124105	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Invalid arguments were specified for the canextend operation.</p> <p><b>Action:</b> Ensure that the arguments are correct. If this error happens during normal operation, please contact Support.</p>
124106	ERROR	\$ERRMSG Unable to access template resource \"\$TemplateName\"	<p><b>Cause:</b> The resource's underlying disk information cannot be determined.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Ensure the hierarchy is correct on the template system before extending.</p>
124107	ERROR	\$ERRMSG Resource \"\${TemplateTagName}\" must have one and only one device resource dependency	<p><b>Cause:</b> The resource has too many underlying devices in the hierarchy.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template system before extending.</p>
124108	ERROR	\$ERRMSG Unable to access template resource \"\${TemplateTagName}\"	<p><b>Cause:</b> The resource cannot be found on the template system.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template system before extending.</p>
124109	ERROR	\$ERRMSG Can not access canextend for scsi/\${DeviceResType} resources on machine \"\${TargetSysName}\"	<p><b>Cause:</b> The target system is missing some required components.</p> <p><b>Action:</b> Ensure that the target system has all the correct kits installed and licensed.</p>
124110	ERROR	\$ERRMSG Either filesystem \"\${TemplateLkId}\" is not mounted on \"\${TemplateSysName}\" or filesystem is not shareable with \"\${TargetSysName}\"	<p><b>Cause:</b> The filesystem isn't in service on the template system or doesn't meet the requirements for extending to the target system.</p> <p><b>Action:</b> Make sure the resource is in service on the template system and review the product documentation regarding the requirements for extending filesystems.</p>
124111	ERROR	\$ERRMSG File system type \"\${FSType}\" is not supported by the kernel currently running on \"\${TargetSysName}\"	<p><b>Cause:</b> The filesystem's type cannot be mounted on the target system due to</p>

Code	Severity	Message	Cause/Action
			<p>lack of kernel support.</p> <p><b>Action:</b> Ensure that the target system has all its kernel modules configured correctly before extending the resource.</p>
124112	ERROR	must specify machine name containing primary hierarchy	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>
124113	ERROR	must specify primary ROOT tag	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>
124114	ERROR	must specify primary mount point	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>
124115	ERROR	must specify primary switchback type	<p><b>Cause:</b> Invalid arguments were specified for the creFShier operation.</p> <p><b>Action:</b> If this error happens during normal operation, please contact Support.</p>
124118	ERROR	dep_remove failure on machine \"'\$PRIMACH'\" for parent \"'\$PRITAG'\" and child \"'\$DEVTAG'.\	<p><b>Cause:</b> Cleanup after a dependency creation failed.</p> <p><b>Action:</b> Check adjacent log messages</p>

Code	Severity	Message	Cause/Action
			for further details.
124119	ERROR	ins_remove failure on machine \\\"\$PRIMACH\" for \\\"\$PRITAG.\	<p><b>Cause:</b> Cleanup after an instance creation failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124121	ERROR	ins_remove failure on machine \\\"\$PRIMACH\"	<p><b>Cause:</b> Cleanup after a resource creation failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124122	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124123	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Invalid arguments were specified for the depstoextend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124124	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Invalid arguments were specified for the depstoextend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>

Code	Severity	Message	Cause/Action
124125	ERROR	\$ERRMSG Unable to access template resource \"\${TemplateTagName}\"	<p><b>Cause:</b> The resource was unable to locate its underlying disk resource.</p> <p><b>Action:</b> Ensure the hierarchy and all dependencies are correct before extending.</p>
124126	ERROR	unextmgr failure on machine \"\${PRIMACH}\"	<p><b>Cause:</b> The cleanup, after a failed resource extend operation, failed.</p> <p><b>Action:</b> Manually clean up any remaining resources and check adjacent log messages for further details.</p>
124128	ERROR	unextmgr failure on machine \"\${PRIMACH}\" for \"\${PRITAG}.\"	<p><b>Cause:</b> The cleanup, after a failed resource extend operation, failed.</p> <p><b>Action:</b> Manually clean up any remaining resources and check adjacent log messages for further details.</p>
124129	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> This message should not occur under normal circumstances.</p> <p><b>Action:</b> Look for additional log messages for more details.</p>
124130	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Invalid arguments were specified for the extend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124131	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Invalid arguments were</p>

Code	Severity	Message	Cause/Action
			<p>specified for the extend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124132	ERROR	\$ERRMSG Required target mount point is null	<p><b>Cause:</b> Invalid arguments were specified for the extend operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124133	ERROR	\$ERRMSG Unable to access template resource \"\${TemplateName}\"	<p><b>Cause:</b> The tag being extended doesn't exist on the template system.</p> <p><b>Action:</b> Ensure that the hierarchy is correct on the template system before extending.</p>
124134	ERROR	\$ERRMSG Detected conflict in expected tag name \"\${TargetTagName}\" on target machine.	<p><b>Cause:</b> A resource already exists on the target system with the same tag as the resource being extended.</p> <p><b>Action:</b> Recreate one of the conflicting resources with a different tag.</p>
124135	ERROR	\$ERRMSG Resource \"\${TemplateName}\" does not have required device resource dependency or unable to access this resource on template machine.	<p><b>Cause:</b> The resource or its underlying disk resource cannot be found on the template system.</p> <p><b>Action:</b> Ensure that the hierarchy is correct on the template system before extending.</p>
124136	ERROR	\$ERRMSG Resource \"\${TemplateName}\" must have one and only one device resource	<p><b>Cause:</b> The resource has multiple</p>

Code	Severity	Message	Cause/Action
		dependency	<p>underlying devices in the hierarchy on the template system.</p> <p><b>Action:</b> Ensure the hierarchy is correct before extending and that the filesystem resource only depends on a single disk resource.</p>
124137	ERROR	\$ERRMSG Can not access extend for scsi/\$DeviceResType resources on machine \"\$TargetSysName\	<p><b>Cause:</b> The files required to support the given storage type aren't available on the target system.</p> <p><b>Action:</b> Ensure that the required kits are installed on the target system and licensed.</p>
124138	ERROR	\$ERRMSG Unable to access target device resource \"\$DeviceTagName\" on machine \"\$TargetSysName\	<p><b>Cause:</b> The required underlying disk resource doesn't exist on the target system.</p> <p><b>Action:</b> Check adjacent log messages for further details and ensure that the target system is properly configured for hosting the resources being extended.</p>
124141	ERROR	\$ERRMSG Unable to find mount point \"\$TemplateLKId\" mode on template machine	<p><b>Cause:</b> The details of the mount point on the template system cannot be determined.</p> <p><b>Action:</b> Ensure that the resource is in service and accessible on the template system before extending.</p>
124142	ERROR	\$ERRMSG Unable to create or access mount point \"\$TargetLKId\" on target machine	<p><b>Cause:</b> The mount point could not be created on the target system.</p> <p><b>Action:</b> Ensure that the mount point's parent directory exists and is accessible on the target system.</p>

Code	Severity	Message	Cause/Action
124143	ERROR	\$ERRMSG Two or more conflicting entries found in /etc/fstab on \"\$TargetSysName\"	<p><b>Cause:</b> The device or mount point appears to be mounted more than once on the target system.</p> <p><b>Action:</b> Ensure that the mount point is not mounted on the target system before extending.</p>
124144	ERROR	\$ERRMSG Failed to create resource instance on \"\$TargetSysName\"	<p><b>Cause:</b> The resource creation on the target system failed.</p> <p><b>Action:</b> Check adjacent log messages for further details. Make sure to check the logs on the target server.</p>
124145	ERROR	\$ERRMSG Failed to set resource instance state for \"\$TargetTagName\" on \"\$TargetSysName\"	<p><b>Cause:</b> The source state could not be changed to OSU on the target system.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124146	ERROR	must specify machine name containing primary hierarchy	<p><b>Cause:</b> Invalid arguments were specified for the filesyshier operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124147	ERROR	must specify primary mount point	<p><b>Cause:</b> Invalid arguments were specified for the filesyshier operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124149	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The process of finding the</p>

Code	Severity	Message	Cause/Action
			<p>resource instance failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124150	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The system failed to read the /etc/mstab file.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124152	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The underlying disk resource could not be found.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124153	ERROR	create file system hierarchy failure	<p><b>Cause:</b> Creating the filesystem resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124154	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The info field for the resource could not be updated.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124155	ERROR	create file system hierarchy failure	<p><b>Cause:</b> The switchback strategy could not be set on the resource.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124157	ERROR	create file system hierarchy failure (conflicting entries in /etc/fstab)	<p><b>Cause:</b> The mount point could not be removed from the /etc/fstab file.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check adjacent log messages for further details.
124160	ERROR	Unknown error in script filesysins, err=\$err	<b>Cause:</b> This message should not occur under normal circumstances.  <b>Action:</b> Check adjacent log messages for further details.
124161	ERROR	create filesys instance – existid – failure	<b>Cause:</b> This message should not occur under normal circumstances.  <b>Action:</b> Check adjacent log messages for further details.
124163	ERROR	create filesys instance – ins_list – failure	<b>Cause:</b> Checking for an existing resource failed.  <b>Action:</b> Check adjacent log messages for further details.
124164	ERROR	create filesys instance – newtag – failure	<b>Cause:</b> The system failed to generate a suggested tag for the resource.  <b>Action:</b> If this error happens during normal operation, contact Support.
124168	ERROR	create filesys instance – ins_create – failure	<b>Cause:</b> The filesystem resource could not be created.  <b>Action:</b> Check adjacent log messages for further details.
124169	ERROR	filesys instance – ins_setstate – failure	<b>Cause:</b> The new filesystem resource's state could not be initialized.  <b>Action:</b> Check adjacent log messages

Code	Severity	Message	Cause/Action
			for further details.
124173	ERROR	create fileys instance – dep_create – failure	<p><b>Cause:</b> The resource's dependency relationship with its underlying disk could not be created.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124174	ERROR	machine not specified	<p><b>Cause:</b> Invalid arguments were specified for the rmenu_mp operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124175	ERROR	mount point not specified	<p><b>Cause:</b> Invalid arguments were specified for the rmenu_mp operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124177	ERROR	unexpected multiple matches found	<p><b>Cause:</b> One or more systems show a filesystem or mount point used more than once.</p> <p><b>Action:</b> Verify filesystem devices and mount points and ensure that filesystems are only mounted once. Look for additional log messages for more details.</p>
124178	ERROR	machine name not specified	<p><b>Cause:</b> Invalid arguments were specified for the rmenump operation.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124180	ERROR	must specify filesystem type	<p><b>Cause:</b> Invalid arguments were specified for the validfstype operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124181	ERROR	mount point not specified	<p><b>Cause:</b> Invalid arguments were specified for the validmp operation.</p> <p><b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.</p>
124182	ERROR	The mount point \$MP is not an absolute path	<p><b>Cause:</b> A mount point was specified that isn't an absolute path (doesn't start with a '/').</p> <p><b>Action:</b> Specify a mount point as an absolute path starting with a '/'.</p>
124183	ERROR	\$MP is already mounted on \$MACH	<p><b>Cause:</b> The requested mount point is already in use on the system.</p> <p><b>Action:</b> Specify a mount point that isn't in use or unmount it before retrying the operation.</p>
124184	ERROR	The mount point \$MP is already protected by LifeKeeper on \$MACH	<p><b>Cause:</b> The system is already protecting the specified mount point.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Choose a different mount point that isn't already being protected.
124185	ERROR	The mount point \$MP is not a directory on \$MACH	<b>Cause:</b> The mount point refers to a non-directory such as a regular file. <b>Action:</b> Choose a mount point that refers to a directory.
124186	ERROR	The mount point directory \$MP is not empty on \$MACH	<b>Cause:</b> The specified mount point refers to a directory that isn't empty. <b>Action:</b> Choose a mount point that is empty or remove the contents of the specified directory before retrying the operation.
124187	ERROR	server name not specified	<b>Cause:</b> Invalid arguments were specified for the valuepmp operation. <b>Action:</b> Ensure the script is called correctly. If this error happens during normal operation, please contact Support.
124188	ERROR	There are no mount points on server \$MACH	<b>Cause:</b> There are no possible mount points for filesystem resource on the server. <b>Action:</b> Check adjacent log messages for further details.
124194	WARN	Please correct conflicting \"/etc/fstab\" entries on server \$UNAME for: \$FSDEV, \$FSNAME	<b>Cause:</b> After deleting a filesystem resource, some entries in /etc/fstab need to be manually cleaned up. <b>Action:</b> Manually clean up the /etc/fstab file.

Code	Severity	Message	Cause/Action
124195	ERROR	getchildinfo found no \$OKAPP child for \$PTAG	<p><b>Cause:</b> The system could not find a child resource in the hierarchy.</p> <p><b>Action:</b> Check adjacent log messages for further details and ensure that the hierarchy is correct before retrying the operation.</p>
124196	ERROR	enablequotas – quotacheck may have failed for \$FS_NAME	<p><b>Cause:</b> The quota operation failed.</p> <p><b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and in /var/log/messages.</p>
124198	ERROR	enablequotas – quotaon failed to turn on quotas for \$FS_NAME, reason	<p><b>Cause:</b> The quota operation failed.</p> <p><b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and /var/log/messages.</p>
124200	ERROR	The device node \$dev was not found or did not appear in the udev create time limit of \$delay seconds	<p><b>Cause:</b> A device node (/dev/...) was not created by udev. This may indicate an issue with the storage or the server's connection to the storage.</p> <p><b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and in /var/log/messages.</p>
124201	WARN	Device \$device not found. Will retry wait to see if it appears.	<p><b>Cause:</b> This can happen under normal conditions while udev creates device node entries for storage. This message should not happen repeatedly.</p> <p><b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and in /var/log/messages.</p>

Code	Severity	Message	Cause/Action
124202	ERROR	Command \" <code>\$commandwithargs</code> \" failed. Retrying ....	<p><b>Cause:</b> The given command failed but may have failed temporarily. This failure may happen during normal operations but should not keep failing.</p> <p><b>Action:</b> Check adjacent log messages for further details if this message continues.</p>
124204	WARN	cannot make file system <code>\$FSNAME</code> mount point	<p><b>Cause:</b> The mount point directory could not be created.</p> <p><b>Action:</b> Ensure that the mount point can be created. This may be due to filesystem permissions, mount options, etc.</p>
124207	ERROR	\" <code>fsck</code> \"ing file system <code>\$FSNAME</code> failed, trying alternative superblock	<p><b>Cause:</b> This message indicates that the typical filesystem check failed. This message may be ok for ext2 filesystems or other filesystems where an alternative superblock location is used.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124209	ERROR	\" <code>fsck</code> \"ing file system <code>\$FSNAME</code> with alternative superblock failed	<p><b>Cause:</b> This indicates that an ext2 filesystem (or other filesystem where an alternative superblock location is used) check failed with the alternative superblock location.</p> <p><b>Action:</b> Check adjacent log messages for further details and instructions on how to proceed.</p>
124210	WARN	POSSIBLE FILESYSTEM CORRUPTION ON <code>\$FSNAME</code> ( <code>\$FPNAME</code> )	<p><b>Cause:</b> A filesystem was put in service or failed over when it was out of sync with its mirror source.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check adjacent log messages for further details and review the product documentation for information on how to bring the filesystem in service safely.</p>
124211	ERROR	Reason for fsck failure (\$retval): \$ret	<p><b>Cause:</b> This log message is part of a series of messages and gives the actual exit code from the fsck process.</p> <p><b>Action:</b> Check adjacent log messages for further details and instructions on how to proceed.</p>
124212	ERROR	\"fsck\" of file system \$FSNAME failed	<p><b>Cause:</b> The check of the filesystem failed. This is usually due to the filesystem having corruption.</p> <p><b>Action:</b> Check adjacent log messages for further details. Review the product documentation for instructions on how to handle possible filesystem corruption.</p>
124213	WARN	POSSIBLE FILESYSTEM CORRUPTION ON \$FSNAME (\$FPNAME	<p><b>Cause:</b> The system or user tried to bring into service a filesystem that may be corrupted. This can happen if a filesystem is switched or failed over when it was out of sync with its mirror source.</p> <p><b>Action:</b> Check adjacent log messages for further details and review the product documentation for instructions on how to bring the resource into service safely.</p>
124214	ERROR	Reason for fsck failure (\$retval)	<p><b>Cause:</b> This message should follow a previous log message about a</p>

Code	Severity	Message	Cause/Action
			<p>filesystem check failure and gives the process exit code of the fsck process.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124218	ERROR	File system \$FSNAME was found to be already	<p><b>Cause:</b> This message is part of a series of messages.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124219	ERROR	mounted after initial mount attempt failed.	<p><b>Cause:</b> This message is part of a series of messages. This should not happen under normal circumstances but may not be fatal if the resource can be put in service.</p> <p><b>Action:</b> Check adjacent log messages for further details in both the lifekeeper log and in /var/log/messages.</p>
124220	ERROR	File system \$FSNAME failed to mount.	<p><b>Cause:</b> The filesystem could not be mounted.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
124221	WARN	Protected Filesystem \$ID is full	<p><b>Cause:</b> The filesystem is full.</p> <p><b>Action:</b> Remove unused data from the filesystem or migrate to a larger filesystem.</p>
124222	WARN	Dependent Applications may be affected <>	<p><b>Cause:</b> This indicates that an operation on a resource is likely to cause operation on other resources based on the resource hierarchy.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Make sure it's acceptable for the indicated resources to be affected before continuing.
124223	ERROR	Put \" <i>\$t</i> \" Out-Of-Service Failed By Signal	<b>Cause:</b> This message should not occur under normal circumstances. <b>Action:</b> Check adjacent log messages for further details.
124227	ERROR	Put \" <i>\$i</i> \" Out-Of-Service Failed	<b>Cause:</b> The operation failed. <b>Action:</b> Check adjacent log messages for further details.
124230	ERROR	Put \" <i>\$t</i> \" In-Service Failed By Signal	<b>Cause:</b> This message should not occur under normal circumstances. <b>Action:</b> Check adjacent log messages for further details.
124231	ERROR	Put \" <i>\$t</i> \" In-Service Failed	<b>Cause:</b> The operation failed. <b>Action:</b> Check adjacent log messages for further details.
124234	ERROR	Put \" <i>\$t</i> \" In-Service Failed	<b>Cause:</b> The operation failed. <b>Action:</b> Check adjacent log messages for further details.

## 8.1.6. Gen/App Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
126105	ERROR	script not specified – \$PTH is a directory	<p><b>Cause:</b> The specified script path is a directory.</p> <p><b>Action:</b> Correct the path of the script.</p>
126110	ERROR	script \$PTH does not exist	<p><b>Cause:</b> The specified script path does not exist.</p> <p><b>Action:</b> Correct the path of the script.</p>
126115	ERROR	script \$PTH is a zero length file	<p><b>Cause:</b> The specified script is an empty file.</p> <p><b>Action:</b> Correct the script's file path and check the contents inside the script.</p>
126117	ERROR	script \$PTH is not executable	<p><b>Cause:</b> The specified script is not executable.</p> <p><b>Action:</b> Correct the script's file path, check the contents inside the script file and make sure it has the proper execute permissions.</p>
126125	ERROR	required template machine name is null	<p><b>Cause:</b> The input template machine name is null.</p> <p><b>Action:</b> Correct the input template machine name.</p>
126130	ERROR	required template resource tag name is null	<p><b>Cause:</b> The input template resource</p>

Code	Severity	Message	Cause/Action
			<p>{tag} is null.</p> <p><b>Action:</b> Correct the input template resource tag name.</p>
126135	ERROR	Unable to generate a new tag	<p><b>Cause:</b> Failed to generate a new tag as the same as the template tag name on the target node using the "newtag" script during the extension. The tag name is already existing.</p> <p><b>Action:</b> Avoid using duplicate tag name on the same node and check the log for detail.</p>
126140	ERROR	Unable to generate a new tag	<p><b>Cause:</b> Failed to generate a new tag as input target tag name on the target node using the "newtag" script during the extension. The tag name is already existing.</p> <p><b>Action:</b> Avoid using duplicate tag name on the same node and check log for detail.</p>
126150	ERROR	unable to remote copy template \"\$_lscript\" script file	<p><b>Cause:</b> Failed to remote copy template script file. The cause may be due to the non-existence/availability of template script file on template node or any transaction failure during "lcdrp" process.</p> <p><b>Action:</b> Check the existence/availability of template script and the connection to template node. Also check the logs for related errors and try to resolve the reported problem.</p>
126155	ERROR	failed to create resource instance on \"\$TargetSysName\"	<p><b>Cause:</b> Failed to create resource</p>

Code	Severity	Message	Cause/Action
			<p>instance using "ins_create".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126160	ERROR	failed to set resource instance state for \" <code>\$TargetTagName</code> \" on \" <code>\$TargetSysName</code> \"	<p><b>Cause:</b> Failed to set resource instance state using "ins_setstate" during GenApp resource extension.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126170	ERROR	getlocks failure	<p><b>Cause:</b> Failed to get the administrative lock when creating a resource hierarchy.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126175	ERROR	instance create failed	<p><b>Cause:</b> Failed to create a GenApp instance using "appins".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126180	ERROR	unable to set state to OSU	<p><b>Cause:</b> Failed to set resource instance state using "ins_setstate" during GenApp resource creation.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126190	ERROR	resource restore has failed	<p><b>Cause:</b> Failed to restore GenApp</p>

Code	Severity	Message	Cause/Action
			<p>resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126200	ERROR	create application hierarchy rlslocks failure	<p><b>Cause:</b> Failed to release lock after GenApp resource created.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126210	ERROR	copy \$ltype script \$lscript failure	<p><b>Cause:</b> Failed to copy user provided script to appropriate GenApp directory during resource creation.</p> <p><b>Action:</b> Check the existence/availability of user provided script and the GenApp directory as well. Also check the logs for related errors and try to resolve the reported problem.</p>
126215	ERROR	no \$ltype script specified	<p><b>Cause:</b> Missing user defined script during GenApp resource creation.</p> <p><b>Action:</b> Check the input action script and run resource creation again.</p>
126220	ERROR	no machine name specified	<p><b>Cause:</b> Missing specified machine name during GenApp resource creation. Failed to copy specified user script due to missing the input for machine name.</p> <p><b>Action:</b> Check the input for machine name and run resource creation again.</p>
126225	ERROR	no resource tag specified	<p><b>Cause:</b> Missing specified tag name</p>

Code	Severity	Message	Cause/Action
			<p>during resource creation.</p> <p><b>Action:</b> Check the input for source tag name and run resource creation again.</p>
126230	ERROR	\$ERRMSG Script was terminated for unknown reason	<p><b>Cause:</b> Failed to extend GenApp resource due to unknown reason.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126235	ERROR	\$ERRMSG Required template machine name is null	<p><b>Cause:</b> Missing the input for template machine name during GenApp resource extension.</p> <p><b>Action:</b> Check the input for template machine name and do the resource extension again.</p>
126240	ERROR	\$ERRMSG Required template resource tag name is null	<p><b>Cause:</b> Missing the input for template resource tag name during GenApp resource extension.</p> <p><b>Action:</b> Check the input for template resource tag name and do the resource extension again.</p>
126245	ERROR	\$ERRMSG Can not access extend for \$AppType/\$ResType resources on machine \"\$TargetSysName\	<p><b>Cause:</b> Failed to locate "extend" script during GenApp resource extension on target node.</p> <p><b>Action:</b> Check the existence/availability of "extend" script and do GenApp resource extension again.</p>
126250	ERROR	create application failure – ins_list failed	<p><b>Cause:</b> Failed when calling "ins_list" during GenApp resource creation.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126255	ERROR	create application failure – unable to generate a new tag	<p><b>Cause:</b> Failed to generate a new tag during the GenApp resource creation.</p> <p><b>Action:</b> Avoid using duplicate tag name on the same node. Also check the logs for related errors and try to resolve the reported problem.</p>
126270	ERROR	create application failure – ins_create failed	<p><b>Cause:</b> Failed using "ins_create" to create GenApp instance.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
126275	ERROR	create application failure – copy_actions failed	<p><b>Cause:</b> Failed using "copy_actions" to copy user specified template script file.</p> <p><b>Action:</b> Check the existence/availability of template script. Also check the logs for related errors and try to resolve the reported problem.</p>
126290	ERROR	Unable to obtain tag for resource with id \$ID	<p><b>Cause:</b> Failed to fetch GenApp resource tag name by input ID during recovery.</p> <p><b>Action:</b> Check the correctness of input ID and existence/availability of GenApp resource in LCD. Also check the logs for related errors and try to resolve the reported problem.</p>
126300	ERROR	generic application recover script for \$TAG was not found or was not executable	<p><b>Cause:</b> Failed to locate the user</p>

Code	Severity	Message	Cause/Action
			<p>defined script for GenApp resource during recovery.</p> <p><b>Action:</b> Check the existence/availability of the user defined script and do the GenApp recovery process again.</p>
126310	ERROR	-t flag not specified	<p><b>Cause:</b> Missing the input for resource tag name during GenApp resource restore.</p> <p><b>Action:</b> Check the input for resource tag name and do resource restore again.</p>
126315	ERROR	-i flag not specified	<p><b>Cause:</b> Missing the input for resource internal id during GenApp resource restore.</p> <p><b>Action:</b> Check the input for resource internal id and do resource restore again.</p>
126327	ERROR	END timeout restore of \"\$TAG\" (forcibly terminating)	
126335	ERROR	restore script \"\$LCDAS/\$APP_RESTOREDIR/\$TAG\" was not found or is not executable	<p><b>Cause:</b> Failed to locate the user defined script for GenApp resource during restore.</p> <p><b>Action:</b> Check the existence/availability of the user defined script and do the GenApp restore process again.</p>
126340	ERROR	-t flag not specified	<p><b>Cause:</b> Missing the input for resource tag name during GenApp resource remove.</p> <p><b>Action:</b> Check the input for resource tag name and do resource remove</p>

Code	Severity	Message	Cause/Action
			again.
126345	ERROR	-i flag not specified	<p><b>Cause:</b> Missing the input for resource internal id during GenApp resource remove.</p> <p><b>Action:</b> Check the input for resource internal id and do resource remove again.</p>
126357	ERROR	END timeout remove of \"\$TAG\" (forcibly terminating)	
126365	ERROR	remove script \"\$LCDAS/\$APP_REMOVEDIR/\$TAG\" was not found or was not executable	<p><b>Cause:</b> Failed to locate the user defined script for GenApp resource during remove.</p> <p><b>Action:</b> Check the existence/availability of the user defined script and do the GenApp remove process again.</p>
126375	ERROR	Script has hung checking \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "quickCheck" Script will be forcibly terminated for GenApp resource with tag name {tag} due to a waiting time over the user defined timeout.</p> <p><b>Action:</b> Check the GenApp resource performance and restart quickChecking. Also check the logs for related errors and try to resolve the reported problem.</p>
126380	ERROR	Usage error: no tag specified	<p><b>Cause:</b> Missing the input for resource tag name during GenApp resource quickCheck.</p> <p><b>Action:</b> Check the input for resource tag name and retry resource quickCheck.</p>

Code	Severity	Message	Cause/Action
126385	ERROR	Internal error: ins_list failed on \$tag.	<p><b>Cause:</b> Failed using "ins_list" to fetch the GenApp resource information by input tag name during the quickCheck process.</p> <p><b>Action:</b> Correct the input tag name and do the quickCheck process again. Also check the logs for related errors and try to resolve the reported problem.</p>
126390	FATAL	Failed to fork process to execute \$userscript: \$!	<p><b>Cause:</b> Failed to fork process to execute user defined "quickCheck" script during the GenApp resource "quickCheck" process.</p> <p><b>Action:</b> Check the existency/availability of the user defined "quickCheck" script and do the "quickCheck" process again.</p>
126391	ERROR	quickCheck has failed for \"\$tag\". Starting recovery.	<p><b>Cause:</b> The GenApp resource with tag name {tag} is determined to be failed by using the user defined health monitoring script – "quickCheck" and the recovery process will be initiated.</p> <p><b>Action:</b> Check the performance of the GenApp resource when local recovery finished. Also check the logs for related errors and try to resolve the reported problem.</p>
126392	WARN	\${convtag}_TIMEOUT: This parameter is old. This parameter will not be supported soon.	
126400	ERROR	-t flag not specified	<p><b>Cause:</b> Missing the input for resource tag name during GenApp resource deletion process.</p> <p><b>Action:</b> Check the input for resource tag name and do resource deletion</p>

Code	Severity	Message	Cause/Action
			process again.
126405	ERROR	-i flag not specified	<p><b>Cause:</b> Missing the input for resource internal id during GenApp resource deletion process.</p> <p><b>Action:</b> Check the input for resource internal id and do resource deletion process again.</p>
126478	ERROR	Failed to create tag '\\$new_leaf'.	<p><b>Cause:</b> The 'creapphier' utility failed to create the specified tag.</p> <p><b>Action:</b> Check /var/log/lifekeeper.log for additional messages from 'creapphier'.</p>
126479	ERROR	Failed to extend tag '\\$new_leaf'.	<p><b>Action:</b> Check /var/log/lifekeeper.log for additional errors from extend manager.</p>
126481	ERROR	Failed to create dependency on '\\$sys' for '\\$new_leaf' to '\\$hier{\\$leaf}{\\$sys}'Tag'.	<p><b>Action:</b> Check /var/log/lifekeeper.log for errors from the 'dep_create' function</p>
126484	ERROR	Tag '\\$root_tag' is not in-service.	<p><b>Cause:</b> The specified tag is not in-service on any node in the cluster.</p> <p><b>Action:</b> Bring the specified tag in-service on any node in the cluster and re-run 'create_terminal_leaf'.</p>
126485	ERROR	Tag '\\$root_tag_1' was not found, select the root tag for a hierarchy to add a terminal leaf resource.	<p><b>Cause:</b> The first tag passed to 'create_terminal_leaf' was not found on the system where the utility was run.</p> <p><b>Action:</b> Verify each resource is in-service on a node in the cluster and fully extended to all nodes.</p>

Code	Severity	Message	Cause/Action
126486	ERROR	Unable to create leaf tag from '\\$tag'.	<p><b>Cause:</b> A unique terminal leaf tag could not be created. A unique terminal leaf tag could not be created after 100 tries.</p> <p><b>Action:</b> Check for multiple leaf tags and for errors in /var/log/lifekeeper.log that may indicate the problem.</p>
126487	ERROR	Tag '\\$root_tag_2' was not found, select the root tag for a hierarchy to add a terminal leaf resource.	<p><b>Cause:</b> The second tag passed to 'create_terminal_leaf' was not found on the system where the utility was run.</p> <p><b>Action:</b> Verify the resource is in-service on a node in the cluster and fully extended to all nodes.</p>
126488	ERROR	Tag '\\$root_tag_1' is not extended to 3 or more systems.	<p><b>Cause:</b> The specified tag is not extended to 3 or more nodes.</p> <p><b>Action:</b> Extend the specified tag to at least 3 nodes and retry 'create_terminal_leaf'.</p>
126489	ERROR	\$cmd does not support SDRS resources.	<p><b>Cause:</b> A multi-site configuration was detected.</p> <p><b>Action:</b> none</p>
126492	ERROR	Remove resource \$tag failed.	
126494	ERROR	Delete dependency failed on '\\$sys' for '\\$tag' to '\\$parent'.	
126495	ERROR	New tag '\\$new_tag' was modified during create, expected '\\$new_leaf'.	
126496	ERROR	Tag '\\$root_tag_1' is not in-service.	
126497	ERROR	Tag '\\$root_tag_2' is not in-service.	
126498	ERROR	Tag '\\$root_tag_1' and '\\$root_tag_2' are not extended to same systems.	

## 8.1.7. IP Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
123006	FATAL	Unknown version %s of IP address	<p><b>Cause:</b> The IP address does not appear to be valid for either IPv4 or IPv6.</p> <p><b>Action:</b> Provide a valid IP address.</p>
123008	ERROR	No pinglist found for %s.	<p><b>Cause:</b> Problem while opening the pinglist for this IP address.</p> <p><b>Action:</b> Make sure you have provided a pinglist for this IP address.</p>
123009	ERROR	List ping test failed for virtual IP %s	<p><b>Cause:</b> No response was received from any of the addresses in the ping list.</p> <p><b>Action:</b> Check network connectivity of this node and the systems on which the IPs in the ping list reside.</p>
123013	ERROR	Link check failed for virtual IP %s on interface %s.	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present on the physical layer connection.</p> <p><b>Action:</b> Check the physical connections for the interface and bring the physical layer link up.</p>
123015	ERROR	Link check failed for virtual IP %s on interface %s.	<p><b>Cause:</b> The requested interface is a bonded interface, and one of the slaves is showing 'NO-CARRIER' indicating that no link is present on the physical</p>

Code	Severity	Message	Cause/Action
			<p>layer connection.</p> <p><b>Action:</b> Check the physical connections for the slave interface and bring the physical layer link up.</p>
123024	ERROR	IP address seems to still exist somewhere else.	<p><b>Cause:</b> The IP address appears to be in use elsewhere on the network.</p> <p><b>Action:</b> Either select a different IP address to use or locate and disable the current use of this IP address.</p>
123037	ERROR	must specify machine name containing primary hierarchy	<p><b>Cause:</b> Not enough arguments were provided to crelPhier.</p> <p><b>Action:</b> Supply all of the needed arguments to crelPhier.</p>
123038	ERROR	must specify IP resource name	<p><b>Cause:</b> Not enough arguments were passed to crelPhier.</p> <p><b>Action:</b> Supply all of the needed arguments to crelPhier.</p>
123039	ERROR	must specify primary IP Resource tag	<p><b>Cause:</b> The argument specifying the primary IP Resource tag was missing from the "crelPhier" command.</p> <p><b>Action:</b> Supply all of the needed arguments.</p>
123042	ERROR	An unknown error has occurred in utility validmask on machine %s.	<p><b>Cause:</b> There was an unexpected error running the "validmask" utility.</p> <p><b>Action:</b> Check adjacent log messages for additional details.</p>

Code	Severity	Message	Cause/Action
123045	ERROR	An unknown error has occurred in utility getlocks.	<p><b>Cause:</b> There was an unexpected error running the "getlocks" utility.</p> <p><b>Action:</b> Check adjacent log messages for additional details.</p>
123053	ERROR	Cannot resolve hostname %s	<p><b>Cause:</b> A hostname was provided for the IP address, but the system was unable to resolve the name to an IP address.</p> <p><b>Action:</b> Check the correctness of the hostname and verify that name resolution (DNS or /etc/hosts) is working correctly and returns the IP for the hostname.</p>
123055	ERROR	An unknown error has occurred in utility %s on machine %s.	<p><b>Cause:</b> There was a failure while creating the IP resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
123056	ERROR	create ip hierarchy failure: perform_action failed	<p><b>Cause:</b> Unexpected error trying to restore the IP address during creation.</p> <p><b>Action:</b> Check adjacent log messages for additional details.</p>
123059	ERROR	Resource already exists on machine %s	<p><b>Cause:</b> Attempted to create an IP address that already exists.</p> <p><b>Action:</b> Reuse the existing resource or manually remove the IP address if it exists or use a different IP address.</p>
123060	ERROR	ins_create failed on machine %s	<p><b>Cause:</b> An unexpected failure occurred</p>

Code	Severity	Message	Cause/Action
			<p>while creating an IP resource.</p> <p><b>Action:</b> Check adjacent log messages for further details.</p>
123064	ERROR	An unknown error has occurred in utility %s on machine %s.	<p><b>Cause:</b> There was a failure while creating a dependency for the IP resource.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
123066	ERROR	An error occurred during creation of LifeKeeper application=comm on %s.	<p><b>Cause:</b> A failure occurred while calling "app_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
123068	ERROR	An error occurred during creation of LifeKeeper resource type=ip on %s.	<p><b>Cause:</b> A failure occurred while calling "typ_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
123089	ERROR	Link check failed for virtual IP %s on interface %s.	
123091	ERROR	the link for interface %s is down	<p><b>Cause:</b> The requested interface is showing 'NO-CARRIER' indicating that no link is present on the physical layer connection.</p> <p><b>Action:</b> Check the physical connections for the interface and bring the physical layer link up.</p>

Code	Severity	Message	Cause/Action
123093	ERROR	the ping list check failed	<p><b>Cause:</b> No response was received from any of the addresses in the ping list.</p> <p><b>Action:</b> Check network connectivity of this node and the systems on which the IPs in the ping list reside.</p>
123095	ERROR	broadcast ping failed	<p><b>Cause:</b> No replies were received from a broadcast ping.</p> <p><b>Action:</b> Verify that at least one host on the subnet will respond to broadcast pings. Verify that virtual IP is on the correct network interface. Consider using a pinglist instead of a broadcast ping.</p>
123096	ERROR	\$msg	<p><b>Cause:</b> The broadcast ping used to determine the viability of the virtual IP failed.</p> <p><b>Action:</b> Please ensure that the ping list for this resource is properly configured in the properties panel or that broadcast ping checking is disabled by adding NOBCASTPING=1 to the /etc/default/LifeKeeper configuration file.</p>
123097	ERROR	exec_list_ping(): broadcast ping failed.	<p><b>Cause:</b> The broadcast ping used to determine the viability of the virtual IP failed.</p> <p><b>Action:</b> Ensure that the ping list for this resource is properly configured in the properties panel or that broadcast ping checking is disabled by adding NOBCASTPING=1 to the /etc/default/LifeKeeper configuration file.</p>
123299	ERROR	Unable to open %s. Reason %s	

Code	Severity	Message	Cause/Action
123410	ERROR	Usage error OSUquickCheck	
123411	ERROR	OSUquickCheck: both tag and id name not specified	
123412	ERROR	resource \$Tag not found on local server	
123414	ERROR	The link for network interface \$IObj->{'device'} is down	
123415	ERROR	No pinglist found for \$IObj->{'ipaddr'}	
123416	ERROR	List ping test failed for virtual IP \$IObj->{'ipaddr'}	

## 8.1.8. Recovery Kit for Oracle Cloud Infrastructure Message Catalog

Use Control F to search for a specific error code in each catalog. To search for any error code, select the Search button at the top right of the screen.

For a complete list of all the messages you may encounter while using LifeKeeper for Linux, please refer to the [Combined Message Catalog](#).

Code	Severity	Message	Cause/Action
139000	ERROR	Cannot find the "oci" command in directories of the PATH. Please confirm that it is installed and the PATH is set correctly.	<p><b>Cause:</b> Unable to get PATH for "oci" <code>oci</code> command.</p> <p><b>Action:</b> Make sure that the <code>oci</code> command is installed and that the installation directory for the <code>oci</code> command is added to your PATH.</p>
139012	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<p><b>Cause:</b> Failed to execute the <code>oci</code> command.</p> <p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139013	ERROR	Failed to assign \$ip, error:\"\$result\"	<p><b>Cause:</b> Failed to allocate \$ip.</p> <p><b>Action:</b> Resolve the cause of the error based on the \$result details.</p>
139022	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<p><b>Cause:</b> Failed to execute the <code>oci</code> command.</p> <p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139023	ERROR	Failed to unassign \$ip, error:\"\$result\"	<p><b>Cause:</b> Failed to allocate \$ip.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Resolve the cause of the error based on the \$result details. If \$ip was unallocated before the remove process was performed, make sure that the resource state is OSU and that \$ip is not allocated to the target VNIC on the OCI console.</p>
139032	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<p><b>Cause:</b> Failed to execute the <code>oci</code> command.</p> <p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139033	ERROR	Failed to \"%s\", unknown error: \"%s\"	<p><b>Cause:</b> The <code>oci</code> command terminated abnormally due to a cause other than code 139032.</p> <p><b>Action:</b> Unknown error: Please take action based on the \"%s\".</p>
139034	ERROR	There is no secondary IPs on \$device.	<p><b>Cause:</b> No secondary IP address has been assigned to \$device.</p> <p><b>Action:</b> If local recovery is enabled, make sure that local recovery was performed and therefore it recovered from the failure.</p>
139035	ERROR	\$ip is not assigned to \$device.	<p><b>Cause:</b> \$ip is not assigned to \$device.</p> <p><b>Action:</b> If local recovery is enabled, make sure that local recovery was performed and therefore it recovered from the failure.</p>
139042	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<p><b>Cause:</b> Failed to execute the <code>oci</code></p>

Code	Severity	Message	Cause/Action
		%s	<p>command.</p> <p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139043	ERROR	Failed to assign \$ip, error:\ "\$result"	<p><b>Cause:</b> Failed to allocate \$ip.</p> <p><b>Action:</b> Resolve the cause of the error based on the \$result details.</p>
139060	ERROR	\$cmd is invalid.	<p><b>Cause:</b> \$cmd is invalid.</p> <p><b>Action:</b> Please contact us.</p>
139061	ERROR	OCIVIP does not support IPv6.	<p><b>Cause:</b> OCIVIP resources do not support use with IPv6.</p> <p><b>Action:</b> Please use IPv4.</p>
139062	ERROR	\$cmd is invalid.	<p><b>Cause:</b> \$cmd is invalid.</p> <p><b>Action:</b> Please contact us.</p>
139063	ERROR	IPv\$ipversion is unknown version.	<p><b>Cause:</b> The IP address version is incorrect.</p> <p><b>Action:</b> Please use IPv4.</p>
139070	ERROR	Failed to access the \$IMDS_URL with the curl command, status code is \$status_code.	<p><b>Cause:</b> Command curl to \$IMDS_URL failed.</p> <p><b>Action:</b> Make sure that curl to \$IMDS_URL completes successfully.</p>

Code	Severity	Message	Cause/Action
139071	ERROR	Failed to decode JSON, error:\ "\$e\".	<p><b>Cause:</b> Failed to decode JSON.</p> <p><b>Action:</b> Make sure that the result obtained from Instance Meta Data Service is in JSON format.</p>
139073	ERROR	Cannot find the vnicId of \$device.	<p><b>Cause:</b> The vnicId corresponding to \$device could not be obtained.</p> <p><b>Action:</b> Check if vnicId exists in the JSON record of Instance Meta Data Service.</p>
139074	ERROR	Cannot find the subnetCidr of \$device.	<p><b>Cause:</b> The subnetCidrBlock corresponding to \$device could not be obtained.</p> <p><b>Action:</b> Check if subnetCidr exists in the JSON record of Instance Meta Data Service.</p>
139075	ERROR	Cannot find the vRouterIp of \$device.	<p><b>Cause:</b> The virtualRouterIp corresponding to \$device could not be obtained.</p> <p><b>Action:</b> Check if virtualRouterIp exists in the JSON record of Instance Meta Data Service.</p>
139080	ERROR	"oci" command(request-id:%s) failed with code %s, message \"%s\" and status %s	<p><b>Cause:</b> Failed to execute the <code>oci</code> command.</p> <p><b>Action:</b> Configure the <code>oci</code> command to end successfully based on the status code and message.</p>
139081	ERROR	Failed to get subnet id, error:\ "\$subnetid\"	<p><b>Cause:</b> Failed to obtain the subnet id.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check if the information can be obtained correctly with the following command. The \$vnicid corresponds to the VNIC-ID (OCID) of each VNIC assigned to the node.</p> <pre>oci network vnic get --vnic-id \$vnicid --raw-output --query 'data."subnet-id"'</pre>
139100	ERROR	IP address is not specified.	<p><b>Cause:</b> IP address is not specified.</p> <p><b>Action:</b> Make sure the IP address is specified.</p>
139101	ERROR	Device is not specified.	<p><b>Cause:</b> Device is not specified.</p> <p><b>Action:</b> Make sure the device is specified.</p>
139105	ERROR	Cannot bring OCIVIP resource \$Tag in service on server \$SysName.	<p><b>Cause:</b> Failed to restore the OCIVIP resource \$Tag on \$SysName.</p> <p><b>Action:</b> Please refer to the log related to the restore process.</p>
139111	ERROR	IP-\$ip already exists, and device and netmask are inconsistent with those of \$ocivipTag.	<p><b>Cause:</b> An IP resource corresponding to the IP address of the OCIVIP resource exists, but the device information and netmask information are inconsistent.</p> <p><b>Action:</b> Change the configuration of the existing IP resource to have the same settings as the OCIVIP resource and create a dependency with the OCIVIP resource.</p>
139112	ERROR	Cannot bring IP resource \$iptag in service on server \$SysName.	<p><b>Cause:</b> Failed to restore the IP</p>

Code	Severity	Message	Cause/Action
			<p>resource \$iptag corresponding to the OCIVIP resource on \$SysName.</p> <p><b>Action:</b> Review the configuration of the IP resource and make sure that restoring of the IP resource is successful. Then create a dependency with the OCIVIP resource.</p>
139200	ERROR	A resource with the specified tag name "%s" already exists on the target machine "%s".	<p><b>Cause:</b> A resource with the specified tag name already exists on the extended node.</p> <p><b>Action:</b> Please use a different tag name or delete the resource with the target tag name of the extended node.</p>
139201	ERROR	Template resource "%s" does not exist on the server "%s".	<p><b>Cause:</b> The resource specified as the extension source does not exist on the node.</p> <p><b>Action:</b> Please specify the tag name correctly.</p>
139250	ERROR	Failed to valuenetOnOCI, error: \"\$stderr_log\".	<p><b>Cause:</b> valuentOnOCI failed.</p> <p><b>Action:</b> The full error message is in \$stderr_log. Please take action based on it.</p>

## 8.1.9. Oracle Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
122500	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the create operation.</p> <p><b>Action:</b> Verify the parameters are correct and retry the operation.</p>
122501	ERROR	DB instance "%s" is already protected on "%s".	<p><b>Cause:</b> An attempt was made to protect an Oracle database instance {sid} that is already under LifeKeeper protection on {server}.</p> <p><b>Action:</b> You must select a different database instance {sid} for LifeKeeper protection.</p>
122502	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122503	ERROR	Unable to locate the oratab file "%s" on "%s".	<p><b>Cause:</b> The oratab file was not found at the default or alternate locations on {server}.</p> <p><b>Action:</b> Verify the oratab file exists and has proper permissions for the Oracle user. A valid oratab file is required to complete the "create" operation.</p>

Code	Severity	Message	Cause/Action
122504	ERROR	Unable to determine Oracle user for "%s" on "%s".	<p><b>Cause:</b> The Oracle Application Recovery Kit was unable to determine the ownership of the Oracle database installation binaries.</p> <p><b>Action:</b> The owner of the Oracle binaries must be a valid non-root user on {server}. Correct the permissions and ownership of the Oracle database installation and retry the operation.</p>
122505	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	<p><b>Cause:</b> The database instance {sid} was not running or connections to the database were not available via the credentials provided.</p> <p><b>Action:</b> The database instance {sid} must be started on {server} and the proper credentials must be provided for the completion of the "create" operation.</p>
122506	ERROR	Unable to determine Oracle tablespaces and logfiles for "%s" on "%s".	<p><b>Cause:</b> A query to determine the location of required tablespaces, logfiles and related database files failed. This may have been caused by an internal database error.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122507	ERROR	Unknown chunk type found for "%s" on "%s".	<p><b>Cause:</b> The specified tablespace, logfile or other required database file is not one of the LifeKeeper supported file or character device types.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> The specified file {database_file} must reference an existing character device or file. Consult the Oracle installation documentation to recreate the specified file {database_file} as a supported file or character device type.</p>
122508	ERROR	DB Chunk "%s" for "%s" on "%s" does not reside on a shared file system.	<p><b>Cause:</b> The specified tablespace, logfile or other required database file {database_file} does not reside on a file system that is shared with other systems in the cluster.</p> <p><b>Action:</b> Use the LifeKeeper UI or "lcdstatus (1M)" to verify that communication paths have been properly created. Use "rpm" to verify that the necessary Application Recovery Kits for storage protection have been installed. Verify that the file is, in fact, not on shared storage, and if not, move it to a shared storage device.</p>
122510	ERROR	File system create failed for "%s" on "%s". Reason	<p><b>Cause:</b> LifeKeeper was unable to create the resource {filesystem} on the specified server {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the "create" operation.</p>
122511	ERROR	%s	<p><b>Cause:</b> The message contains the output of the "filesyshier" command.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>

Code	Severity	Message	Cause/Action
122513	ERROR	Dependency creation between Oracle database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	<p><b>Cause:</b> LifeKeeper was unable to create a dependency between the database resource {tag} and the necessary child resource {childtag}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Once any problems have been corrected, it may be possible to create the dependency between {tag} and {childtag} manually.</p>
122514	ERROR	Unable to "%s" on "%s" during resource create.	<p><b>Cause:</b> The Oracle Application Recovery Kit was unable to release the administrative lock using the "rlslocks" command.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
122516	ERROR	Raw device resource created failed for "%s" on "%s". Reason	<p><b>Cause:</b> LifeKeeper was unable to create the resource {raw device} on the specified server {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the "create" operation.</p>
122519	ERROR	In-service attempted failed for tag "%s" on "%s".	<p><b>Cause:</b> The "perform_action" command for {tag} on {server} failed to start the database {sid}. The in-service operation has failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the</p>

Code	Severity	Message	Cause/Action
			"create" operation.
122521	ERROR	Create of app "%s" on "%s" failed with return code of "%d".	<p><b>Cause:</b> There was an error running the command "app_create" to create the internal application type.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122522	ERROR	Create of typ "%s" for app "%s" on "%s" failed with return code of "%d".	<p><b>Cause:</b> There was an error running the command "typ_create" to create the internal resource type.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122524	ERROR	Setting "resstate" for resource "%s" on "%s" failed with return code of "%d".	<p><b>Cause:</b> There was an error running the command "ins_setstate" to set the resource state to {state}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122525	ERROR	The values specified for the target and the template servers are the same: "%s".	<p><b>Cause:</b> The value specified for the target and template servers for the "extend" operation were the same.</p> <p><b>Action:</b> You must specify the correct parameter for the {target server} and {template server}. The {target server} is</p>

Code	Severity	Message	Cause/Action
			<p>the server where the {tag} will be extended.</p>
122526	ERROR	<p>Unable to locate the oratab file in "/etc" or in "%s" on "%s".</p>	<p><b>Cause:</b> The oratab file was not found at the default or alternate locations on {server}.</p> <p><b>Action:</b> Verify the oratab file exists and has proper permissions for the Oracle user. A valid oratab file is required to complete the "extend" operation.</p>
122527	ERROR	<p>Unable to retrieve the Oracle user on "%s".</p>	<p><b>Cause:</b> An attempt to retrieve the Oracle user from {template server} during a "canextend" or "extend" operation failed.</p> <p><b>Action:</b> The owner of the Oracle binaries must be a valid user on {target server} and {template server}. Correct the permissions and ownership of the Oracle database installation and retry the operation.</p>
122528	ERROR	<p>The Oracle user and/or group information for user "%s" does not exist on the server "%s".</p>	<p><b>Cause:</b> LifeKeeper is unable to find the Oracle user and/or group information for the Oracle user {user} on the server {server}.</p> <p><b>Action:</b> Verify the Oracle user {user} exists on the specified {server}. If the user {user} does not exist, it should be created with the same uid and gid on all servers in the cluster.</p>
122529	ERROR	<p>The id for user "%s" is not the same on template server "%s" and target server "%s".</p>	<p><b>Cause:</b> The user id on the target server {target server} for the Oracle user {user} does not match the value of the user {user} on the template server</p>

Code	Severity	Message	Cause/Action
			<p>{template server}.</p> <p><b>Action:</b> The user ids for the Oracle user {user} must match on all servers in the cluster. The user id mismatch should be corrected manually on all servers before retrying the "extend" operation.</p>
122530	ERROR	The group id for user "%s" is not the same on template server "%s" and target server "%s".	<p><b>Cause:</b> The group id on the target server {target server} for the Oracle user {user} does not match the value of the user {user} on the template server {template server}.</p> <p><b>Action:</b> The group ids for the Oracle user {user} must match on all servers in the cluster. The group id mismatch should be corrected manually on all servers before retrying the "extend" operation.</p>
122532	ERROR	No file system or raw devices found to extend for "%s" on "%s".	<p><b>Cause:</b> There were no dependent file system or raw device resources found for the Oracle resource {tag} on server {template server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
122533	WARN	A RAMDISK (%s) was detected in the ORACLE Database configuration for "%s" on "%s". LifeKeeper cannot protect RAMDISK. This RAMDISK resource will not be protected by LifeKeeper! ORACLE hierarchy creation will continue.	<p><b>Cause:</b> The specified tablespace, logfile or other database file {database_file} was detected as a ramdisk. No protection is available for this type of resource in the current LifeKeeper product.</p> <p><b>Action:</b> The ramdisk will not be protected. You must manually ensure that the required database file</p>

Code	Severity	Message	Cause/Action
			<p>{database_file} will be available during all Oracle database operations.</p>
122534	ERROR	Failed to initialize object instance for Oracle sid "%s" on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122537	ERROR	Update of instance info field for "%s" on "%s" failed (%s).	<p><b>Cause:</b> There was an error while running the command "ins_setinfo" to update the internal resource information field.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122538	ERROR	Initial connect with query buffer to database "%s" on "%s" failed, testing output.	<p><b>Cause:</b> A connection attempt to the Oracle database {sid} to determine the database status has failed.</p> <p><b>Action:</b> The connection attempt failed with the specified credentials. Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122542	ERROR	The "%s [ %s ]" attempt of the database "%s" appears to have failed on "%s".	<p><b>Cause:</b> The attempted Oracle action {action} using method {action_method} for the database instance {sid} failed on</p>

Code	Severity	Message	Cause/Action
			<p>the server {server}.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122543	ERROR	All attempts to "%s" database "%s" on "%s" failed	<p><b>Cause:</b> All efforts to perform the action {action} on the Oracle database {sid} on server {server} have failed.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122544	ERROR	Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".	<p><b>Cause:</b> An unexpected error occurred while attempting to update the oratab entry for the database {sid}. The error occurred while attempting to open the oratab file.</p> <p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122545	ERROR	Unable to locate the oratab file in "/etc" or in "%s" on "%s".	<p><b>Cause:</b> The oratab file was not found at the default or alternate locations on {server}.</p> <p><b>Action:</b> Verify the oratab file exists and has proper permissions for the oracle user. A valid oratab file is required to complete the "extend" operation.</p>

Code	Severity	Message	Cause/Action
122546	ERROR	Unable to open file "%s" on "%s" (%s).	<p><b>Cause:</b> The specified file {file} could not be opened or accessed on the server {server} due to the error {error}.</p> <p><b>Action:</b> Verify the existence and permissions on the specified file {file}. Check adjacent log messages for further details and related errors. You must correct any reported errors before retrying the operation.</p>
122547	ERROR	(cleanUpPids):Forcefully killing hung pid(s):pid(s)="%s"	<p><b>Cause:</b> The process {pid} failed to respond to the request to terminate gracefully. The process {pid} will be forcefully terminated.</p> <p><b>Action:</b> Use the command line to verify that the process {pid} has been terminated. Check the adjacent log messages for further details and related messages.</p>
122548	ERROR	Unable to locate the DB utility (%s/%s) on this host.	<p><b>Cause:</b> The Oracle binaries and required database utility {utility} located at {path/utility} were not found on this server {server}.</p> <p><b>Action:</b> Verify that the Oracle binaries and required software utilities are installed and properly configured on the server {server}. The Oracle binaries must be installed locally on each node or located on shared storage available to all nodes in the cluster.</p>
122549	ERROR	Oracle internal error or non-standard Oracle configuration detected. Oracle User and/or Group set to "root".	<p><b>Cause:</b> The detected ownership of the Oracle database installation resolves to the root user and/or root group. Ownership of the Oracle installation by root is a non-standard configuration.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> The owner of the Oracle binaries must be a valid non-root user on {server}. Correct the permissions and ownership of the Oracle database installation and retry the operation.</p>
122550	ERROR	Initial inspection of "%s" failed, verifying failure or success of received output.	<p><b>Cause:</b> The previous Oracle query {query} or command {cmd} failed to return success.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related errors. Check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct the reported problem(s).</p>
122551	ERROR	Logon failed with "%s" for "%s" on "%s". Please check username/password and privileges.	<p><b>Cause:</b> The logon with the credentials {credentials} for the database instance {sid} on server {server} failed. An invalid user {user} or password was specified.</p> <p><b>Action:</b> Verify that the Oracle database user {user} and password {password} are indeed valid. In addition, the Oracle database user {user} must have sufficient privileges for the attempted action.</p>
122552	ERROR	%s	<p><b>Cause:</b> The message contains the output of the "sqlplus" command.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages.</p>
122553	ERROR	Unable to open file "%s" on "%s" (%s).	<p><b>Cause:</b> The specified file {file} could not be opened or accessed on the server {server} due to the error {error}.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Verify the existence and permissions on the specified file {file}. Check adjacent log messages for further details and related errors. You must correct any reported errors before retrying the operation.</p>
122554	ERROR	The tag "%s" on "%s" is not an Oracle instance or it does not exist.	<p><b>Cause:</b> The specified tag {tag} on server {server} does not refer to an existing and valid Oracle resource instance.</p> <p><b>Action:</b> Use the UI or "lcdstatus (1M)" to verify the existence of the resource tag {tag}. The resource tag {tag} must be an Oracle resource instance to use the command "ora-display."</p>
122555	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to update the authorized user, password and database role for the Oracle resource instance.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122557	ERROR	Update of user and password failed for "%s" on "%s".	<p><b>Cause:</b> A request to update the user and password for the resource tag {tag} failed. The specified credentials failed the initial validation/connection attempt on server {server}.</p> <p><b>Action:</b> Verify the correct credentials {user/password} were specified for the</p>

Code	Severity	Message	Cause/Action
			attempted operation. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.
122559	ERROR	Update of user and password failed for "%s" on "%s".	<p><b>Cause:</b> The update of the user and password information for the resource tag {tag} on server {server} failed.</p> <p><b>Action:</b> Verify the correct credentials {user/password} were specified for the attempted operation. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122562	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The required Oracle executable {exe} was not found on this server {server}.</p> <p><b>Action:</b> Verify that the Oracle binaries and required software utilities are installed and properly configured on the server {server}. The Oracle binaries must be installed locally on each node or located on shared storage available to all nodes in the cluster.</p>
122566	ERROR	Unable to find Oracle home for "%s" on "%s".	<p><b>Cause:</b> The Oracle home directory {Oracle home} does not appear to contain files necessary for the proper operation of the Oracle instance {sid}.</p> <p><b>Action:</b> Verify using the command line that the Oracle home directory {Oracle home} contains the Oracle binaries, a valid spfile{sid}.ora or init{sid}.ora file.</p>

Code	Severity	Message	Cause/Action
122567	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected. The specified internal ID {id} does not match the expected SID {sid}.</p> <p><b>Action:</b> Verify the parameters are correct. Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122568	ERROR	DB Processes are not running on "%s".	<p><b>Cause:</b> A process check for the Oracle instance did not find any processes running on server {server}.</p> <p><b>Action:</b> If local recovery is enabled, the Oracle instance will be restarted locally. Check adjacent log messages for further details and related messages.</p>
122572	ERROR	Failed to create flag "%s" on "%s".	<p><b>Cause:</b> An unexpected error occurred attempting to create a flag for controlling Oracle local recovery processing causing a failover to the standby node.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
122574	ERROR	all attempts to shutdown the database %s failed on "%s".	<p><b>Cause:</b> The shutdown of the Oracle database failed during a local recovery process most likely caused because the maximum number of database connections has been reached.</p> <p><b>Action:</b> Check the Oracle logs for connection failures caused by the</p>

Code	Severity	Message	Cause/Action
			<p>maximum number of available connections being reached, and if found, consider increasing the value. Additionally, set the tunable LK_ORA_NICE to 1 to prevent connection failures from causing a quickCheck failure followed by a local recovery attempt.</p>
122597	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected during pre-extend checking.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the pre-extend.</p>
122598	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being created while attempting to determine the validity of the Oracle home directory.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the "create."</p>
122599	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to look up the Oracle user on the template system.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the extend.</p>

Code	Severity	Message	Cause/Action
122600	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to display the resource properties.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the display of the resource properties.</p>
122601	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to check for valid database authorization.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the command.</p>
122603	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected while attempting to perform health checks on the Oracle resource instance.</p> <p><b>Action:</b> Check that correct arguments were passed to the quickCheck command and also check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the restore.</p>
122604	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected</p>

Code	Severity	Message	Cause/Action
			<p>while attempting to perform a local recovery on the Oracle resource instance.</p> <p><b>Action:</b> Check that correct arguments were passed to the "recover" command, and also check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the recover.</p>
122606	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	<p><b>Cause:</b> The database instance {sid} was not running or connections to the database are not available via the credentials provided.</p> <p><b>Action:</b> The database instance {sid} must be started on {server} and the proper credentials must be provided for the completion of the selected operation.</p>
122607	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	<p><b>Cause:</b> The database instance {sid} was not running or connections to the database are not available via the credentials provided.</p> <p><b>Action:</b> The database instance {sid} must be started on {server} and the proper credentials must be provided for the completion of the selected operation.</p>
122608	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> The "remove" operation failed to create the resource object instance required to take the Oracle resource Out of Service.</p> <p><b>Action:</b> Check that correct arguments were passed to the "remove" command</p>

Code	Severity	Message	Cause/Action
			<p>and also check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the restore.</p>
122609	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> The "restore" operation failed to create the resource object instance required to put the Oracle resource In Service.</p> <p><b>Action:</b> Check that correct arguments were passed to the "restore" command and also check the adjacent log messages for further details and related messages. Correct any reported errors before retrying the "restore."</p>
122610	ERROR	Unable to "%s" on "%s" during resource create.	<p><b>Cause:</b> The Oracle Application Recovery Kit was unable to create the administrative lock using the "getlocks" command during resource creation.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create.</p>
122611	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle resource and the child File System resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create operation.</p>
122612	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle</p>

Code	Severity	Message	Cause/Action
			<p>resource and the child Raw resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create operation.</p>
122613	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle resource and the child Raw resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create operation.</p>
122614	ERROR	%s	<p><b>Cause:</b> The requested dependency creation between the parent Oracle resource and the child Listener resource failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the create operation.</p>
122616	ERROR	%s	<p><b>Cause:</b> The requested start up or shutdown of the Oracle database failed.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors before retrying the "restore" or "remove" operation.</p>
122618	ERROR	Dependency creation between Oracle database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	<p><b>Cause:</b> LifeKeeper was unable to create a dependency between the</p>

Code	Severity	Message	Cause/Action
			<p>database resource {tag} and the necessary child resource {childtag}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Once any problems have been corrected, it may be possible to create the dependency between {tag} and {childtag} manually.</p>
122619	ERROR	Dependency creation between Oracle database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	<p><b>Cause:</b> LifeKeeper was unable to create a dependency between the database resource {tag} and the necessary child resource {childtag}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Once any problems have been corrected, it may be possible to create the dependency between {tag} and {childtag} manually.</p>
122625	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The quickCheck process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages. Correct any reported problems.</p>
122626	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The remove process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages. Correct any reported</p>

Code	Severity	Message	Cause/Action
			problems.
122627	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The restore process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages. Correct any reported problems.</p>
122628	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The recover process was unable to find the Oracle executable "sqlplus."</p> <p><b>Action:</b> Check the Oracle configuration and also check adjacent log messages for further details and related messages. Correct any reported problems.</p>
122632	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> During a remove, the resource instance {sid} passed to the remove process does not match internal resource instance information for the {sid}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122633	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> During a restore, the resource instance {sid} passed to restore does not match internal resource instance information for the {sid}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>

Code	Severity	Message	Cause/Action
122634	ERROR	Oracle SID mismatch. The instance SID "%s" does not match the SID "%s" specified for the command.	<p><b>Cause:</b> During resource recovery, the resource instance {sid} passed to recovery does not match internal resource instance information for the {sid}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122636	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	<p><b>Cause:</b> The create of the Oracle resource hierarchy {tag} failed on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122638	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The create action for the Oracle database resource {tag} on server {server} failed. The signal {sig} was received by the create process.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122640	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122641	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle</p>

Code	Severity	Message	Cause/Action
			<p>resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122642	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122643	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
122644	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to retrieve resource instance information for {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors and retry the extend.</p>
122645	ERROR	Cannot access canextend script "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to run pre-extend checks because it was unable to find the "canextend" script on {server} for a dependent child resource.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>

Code	Severity	Message	Cause/Action
122646	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to retrieve resource instance information for {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors and retry the extend.</p>
122647	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122648	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server "%s" for App "%s" and Type "%s"	<p><b>Cause:</b> During the database resource extension, a resource instance was found using the same {tag} and/or {id} but with a different resource application and type.</p> <p><b>Action:</b> Resource IDs must be unique. The resource instance with the ID matching the Oracle resource instance must be removed.</p>
122649	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> An unexpected error occurred attempting to create the Oracle resource instance {tag} on {server}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Correct any reported errors.</p>
122650	ERROR	Cannot access extend script "%s" on server "%s"	<p><b>Cause:</b> The request to extend the database resource {resource} to {server} failed because it was unable to</p>

Code	Severity	Message	Cause/Action
			<p>the find the script {extend} on {server} for a dependent child resource.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122651	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> The request to extend the database resource {resource} to {server} failed because of an error attempting to extend a dependent child resource.</p> <p><b>Action:</b> Check the adjacent log messages for further details and related messages. Correct any reported errors.</p>
122654	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The health check for the database {sid} was terminated because the quickCheck process received a signal. This is most likely caused by the quickCheck process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The health check time for an Oracle resource is controlled by the tunable value ORACLE_QUICKCHECK_TIMEOUT. Set it to a value greater than 45 seconds to allow more time for the health check process to complete.</p>
122655	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> The request to take database {sid} "Out of Service" was terminated because the remove process received a signal. This is most likely caused by the remove process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The remove time for an Oracle resource is controlled by the tunable value ORACLE_REMOVE_TIMEOUT.</p>

Code	Severity	Message	Cause/Action
			<p>Set the tunable to a value greater than 240 seconds to allow more time for the remove process to complete.</p>
122659	ERROR	<p>END failed %s of "%s" on server "%s" due to a "%s" signal</p>	<p><b>Cause:</b> The request to place database {sid} "In Service" was terminated because the restore process received a signal. This is most likely caused by the restore process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The restore time for an Oracle resource is controlled by the tunable value ORACLE_RESTORE_TIMEOUT. Set the tunable to a value greater than 240 seconds to allow more time for the restore process to complete.</p>
122663	ERROR	<p>END failed %s of "%s" on server "%s" due to a "%s" signal</p>	<p><b>Cause:</b> The recovery of the failed database was terminated because the recovery process received a signal. This is most likely caused by the recovery process requiring more time to complete than was allotted.</p> <p><b>Action:</b> The recovery time for an Oracle resource is controlled by the tunable values ORACLE_RESTORE_TIMEOUT and ORACLE_REMOVE_TIMEOUT. Set one or both of these to a value greater than 240 seconds to allow more time for a recovery to complete.</p>
122670	ERROR	<p>Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".</p>	<p><b>Cause:</b> An unexpected error occurred while attempting to update the oratab entry for the database {sid}. The error occurred while attempting to open the temporary file used in the update process.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122671	ERROR	Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".	<p><b>Cause:</b> An unexpected error occurred while attempting to update the oratab entry for the database {sid}. The error occurred while attempting to close the temporary file used in the update process.</p> <p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122672	ERROR	Update of "%s" sid "%s" on "%s" failed. Reason: "%s" "%s" failed: "%s".	<p><b>Cause:</b> An unexpected error occurred while attempting to update the oratab entry for the database {sid}. The error occurred while attempting to rename the temporary file back to oratab.</p> <p><b>Action:</b> The oratab file entry for {sid} will need to be updated manually to turn off the automatic start up of the database at system boot.</p>
122673	ERROR	Unable to log messages queued while running as oracle user %s on %s. Reason: \$!	<p><b>Cause:</b> An unexpected error {reason} occurred while attempting to add messages to the log file. These messages were generated while running as the Oracle user.</p> <p><b>Action:</b> Review the reason for the failure and take corrective action.</p>
122674	ERROR	Unable to open %s Reason: %s.	<p><b>Cause:</b> An unexpected error occurred while attempting to open a connection</p>

Code	Severity	Message	Cause/Action
			<p>to the Oracle database and run the database {cmd}.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. Additionally, check the Oracle log (alert.log) and related trace logs (*.trc) for additional information and correct any reported problems.</p>
122680	ERROR	Unable to find Oracle home for "%s" on "%s".	<p><b>Cause:</b> The Oracle home directory {Oracle home} does not appear to contain files necessary for the proper operation of the Oracle instance {sid}.</p> <p><b>Action:</b> Verify using the command line that the Oracle home directory {Oracle home} contains the Oracle binaries, a valid spfile{sid}.ora or init{sid}.ora file.</p>
122681	ERROR	Failed to create object instance for Oracle on "%s".	<p><b>Cause:</b> There was an unexpected error creating an internal representation of the Oracle instance being protected.</p> <p><b>Action:</b> Check adjacent log messages for further details and related messages. You must correct any reported errors before retrying the operation.</p>
122682	ERROR	Unable to find the Oracle executable "%s" on "%s".	<p><b>Cause:</b> The required Oracle executable {exe} was not found on this server {server}.</p> <p><b>Action:</b> Verify that the Oracle binaries and required software utilities are installed and properly configured on the server {server}. The Oracle binaries must be installed locally on each node or located on shared storage available</p>

Code	Severity	Message	Cause/Action
			to all nodes in the cluster.
122683	ERROR	Backup node %s is unreachable; abort username/password changes.	<p><b>Cause:</b> Backup node {node} is currently unreachable. Pending changes to username/password have been aborted.</p> <p><b>Action:</b> Ensure that SIOS LifeKeeper is running on the given backup node {node} and that all communication paths are up, then retry the operation.</p>
122684	ERROR	The restore of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_RESTORE_TIMEOUT in /etc/default/LifeKeeper.	<p><b>Cause:</b> The restore for Oracle resource {resource} has timed out on server {server}. This occurs when the in-service operation has taken longer than the time allotted.</p> <p><b>Action:</b> The restore time for an Oracle resource is controlled by the tunable value ORACLE_RESTORE_TIMEOUT. Set the tunable to a value greater than 240 seconds to allow more time for the restore process to complete.</p>
122685	ERROR	The remove of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_REMOVE_TIMEOUT in /etc/default/LifeKeeper.	<p><b>Cause:</b> The remove for Oracle resource {resource} has timed out on server {server}. This occurs when the out-of-service operation has taken longer than the time allotted.</p> <p><b>Action:</b> The remove time for an Oracle resource is controlled by the tunable value ORACLE_REMOVE_TIMEOUT. Set the tunable to a value greater than 240 seconds to allow more time for the remove process to complete.</p>
122686	ERROR	The quickCheck of %s has timed out on server %s. The default TIMEOUT is 45 seconds. To increase the TIMEOUT, set	<p><b>Cause:</b> The quickCheck for Oracle</p>

Code	Severity	Message	Cause/Action
		ORACLE_QUICKCHECK_TIMEOUT in /etc/default/LifeKeeper.	<p>resource {resource} has timed out on server {server}. This occurs when the health check has taken longer than the time allotted.</p> <p><b>Action:</b> The health check time for an Oracle resource is controlled by the tunable value ORACLE_QUICKCHECK_TIMEOUT. Set it to a value greater than 45 seconds to allow more time for the health check process to complete.</p>
122687	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the getalertlog script.</p> <p><b>Action:</b> The tag name of the Oracle resource must be provided to the script as the first command line parameter. Verify the parameters are correct and retry the operation.</p>
122702	ERROR	Failed start OHAS.	
122704	ERROR	Failed start ASM.	
122706	ERROR	Failed start Diskgroup.	
122708	ERROR	Failed stop Diskgroup.	
122710	ERROR	Failed stop ASM.	
122712	ERROR	Failed stop OHAS.	
122713	ERROR	OHAS is not running.	
122714	ERROR	ASM is not running.	
122715	ERROR	Diskgroup is not running.	
122716	ERROR	\$usage	
122717	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" does not exist	
122718	ERROR	Cannot extend resource \"\$template_tag\" to server \"\$me\"	
122719	ERROR	\$usage	

Code	Severity	Message	Cause/Action
122720	ERROR	END failed create of \"\$Tag\" due to a \"\$sig\" signal	
122722	ERROR	Unable to getlocks on \$me during resouce create.	
122725	ERROR	Error creating resource \$Tag. Error (\$rc	
122726	ERROR	END failed hierarchy create of resource \$Tag with return value of \$ecode.	
122728	ERROR	Unable to rlslocks on \$me during resource create.	
122729	ERROR	\$usage	
122730	ERROR	END failed extend of \"\$Tag\" due to a \"\$sig\" signal	
122732	ERROR	Template resource \"\$TemplateTag\" on server \"\$TemplateSys\" does not exist	
122733	ERROR	Error creating resouce \"\$Tag\"on server \"\$me\"	
122735	ERROR	END failed hierarchy extend of resource \$Tag with return value of \$ecode.	
122736	ERROR	\$usage	
122737	ERROR	Failed to create object instance for Oracle ASM on \$oracle::me	
122738	ERROR	Usage : \$cmd \$usage	
122739	ERROR	Failed to create object instance for Oracle ASM on \"\$me\".	
122740	ERROR	\$usage	
122742	ERROR	backup node \$back_dead_name is unreachable; abort protection Diskgroup changes.	
122744	ERROR	Update of protection Diskgroup failed for \"\$tag\" on \"\$oracleasm::me\"	
122746	ERROR	Update of protection Diskgroup failed foor \"\$tag\" on \"\$remote\"	
122750	ERROR	Failed to init the object.	
122751	ERROR	Usage: %s %s	

## 8.1.10. Oracle Listener Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
122005	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "getlocks".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122007	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "rlslocks".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122009	ERROR	The path %s is not a valid file.	<p><b>Cause:</b> There is no listener.ora file.</p> <p><b>Action:</b> Ensure the file exists and retry the operation.</p>
122010	ERROR	The listener user does not exist on the server %s.	<p><b>Cause:</b> "Stat" command could not get user id.</p> <p><b>Action:</b> Retry the operation.</p>
122011	ERROR	The listener user does not exist on the server %s.	<p><b>Cause:</b> UID is not in passwd file.</p> <p><b>Action:</b> Ensure the UID exists in passwd file and retry the operation.</p>
122012	ERROR	The listener user does not exist on the server %s.	<p><b>Cause:</b> User name is not in passwd file.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Ensure the user name exists in passwd file; retry the operation.
122023	ERROR	The %s command failed (%d	<p><b>Cause:</b> This message contains the return code of the "lsnrctl" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122024	ERROR	\$line	<p><b>Cause:</b> The message contains the output of the "lsnrctl" command.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122039	ERROR	Usage error	<p><b>Cause:</b> Invalid parameters were specified for the restore operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122040	ERROR	Script \$cmd has hung on the restore of \"\$opt_t\". Forcibly terminating.	<p><b>Cause:</b> The listener restore script reached its timeout value.</p> <p><b>Action:</b> Ensure listener.ora is valid and that LSNR_START_TIME (default 35 seconds) in /etc/default/LifeKeeper is set to a value greater than or equal to the time needed to start the listener.</p>
122041	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to restore the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported</p>

Code	Severity	Message	Cause/Action
			problem.
122045	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Failed to get resource information.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122046	ERROR	Usage error	<p><b>Cause:</b> Invalid parameters were specified for the restore operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122049	ERROR	The script \$cmd has hung on remove of \"\$opt_t\". Forcibly terminating.	<p><b>Cause:</b> The listener remove script reached its timeout value.</p> <p><b>Action:</b> Ensure listener.ora is valid and that LSNR_STOP_TIME (default 35 seconds) in /etc/default/LifeKeeper is set to a value greater than or equal to the time needed to stop the listener.</p>
122051	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122055	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to quickCheck the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
122057	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122064	WARN	The %s level is set to %s a %s will not occur.	<p><b>Cause:</b> The minimal Listener protection level is Start and Monitor.</p> <p><b>Action:</b> Start the listener manually.</p>
122066	ERROR	Script has hung checking \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The listener quickCheck script reached its timeout value.</p> <p><b>Action:</b> Ensure listener.ora is valid and that LSNR_STATUS_TIME (default 15 seconds) in /etc/default/LifeKeeper is set to a value greater than or equal to the time needed to check the listener.</p>
122067	ERROR	Usage error	<p><b>Cause:</b> Invalid parameters were specified for the quickCheck operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122069	ERROR	Usage error	<p><b>Cause:</b> Invalid parameters were specified for the delete operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122072	ERROR	%s: resource "%s" not found on local server	<p><b>Cause:</b> Invalid parameters were specified for the recover operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>

Code	Severity	Message	Cause/Action
122074	WARN	The local recovery attempt has failed but %s level is set to %s preventing a failover to another node in the cluster. With %s recovery set all local recovery failures will exit successfully to prevent resource failovers.	<p><b>Cause:</b> The optional listener recovery level is set to local recovery only.</p> <p><b>Action:</b> Switch over the resource tree manually.</p>
122078	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to recover the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122082	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122083	ERROR	\$cmd has hung checking \"\$tag\". Forcibly terminating	<p><b>Cause:</b> The recover script was stopped by signal.</p> <p><b>Action:</b> Ensure listener.ora is valid.</p>
122084	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122085	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the canextend operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>

Code	Severity	Message	Cause/Action
122086	ERROR	The values specified for the target and the template servers are the same. Please specify the correct values for the target and template servers.	<b>Cause:</b> The values specified for the target and the template servers are the same.  <b>Action:</b> Perform the steps listed in the message text.
122087	ERROR	Error getting resource information for resource "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.  <b>Action:</b> Check your LifeKeeper configuration.
122088	ERROR	Error getting resource information for resource "%s" on server "%s"	<b>Cause:</b> Failed to get listener user name from resource information.  <b>Action:</b> Ensure the resource info field is valid then retry the operation.
122089	ERROR	The listener user %s does not exist on the server %s.	<b>Cause:</b> User name is not in passwd file.  <b>Action:</b> Ensure the user name exists in passwd file and retry the operation.
122090	ERROR	The id for user %s is not the same on template server %s and target server %s.	<b>Cause:</b> User ID should be same on both servers.  <b>Action:</b> Trim user ID to the same.
122091	ERROR	The group id for user %s is not the same on template server %s and target server %s.	<b>Cause:</b> Group ID should be same on both servers.  <b>Action:</b> Trim group ID to the same.
122092	ERROR	Cannot access canextend script "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to run pre-extend checks because it was unable to find the "canextend" script on

Code	Severity	Message	Cause/Action
			<p>{server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122097	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "configActions" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122098	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122099	ERROR	Unable to update the resource %s to change the %s to %s on %s.	<p><b>Cause:</b> LifeKeeper failed to put information into the info field.</p> <p><b>Action:</b> Restart LifeKeeper and retry the operation.</p>
122100	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122101	ERROR	Unable to update the resource %s to change the %s to %s on %s.	<p><b>Cause:</b> LifeKeeper failed to put information to info field on {server}.</p> <p><b>Action:</b> Restart LifeKeeper on {server} and retry the operation.</p>
122103	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were</p>

Code	Severity	Message	Cause/Action
			<p>specified for the create operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122124	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	<p><b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122126	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "rlslocks".</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122127	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<p><b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122129	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "getlocks."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122131	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check your LifeKeeper configuration.
122133	ERROR	Unable to create a file system resource hierarchy for the file system %s.	<b>Cause:</b> There was an unexpected error running "filesyshier."  <b>Action:</b> Check adjacent log messages for further details.
122135	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<b>Cause:</b> There was an unexpected error running "dep_create."  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122140	ERROR	Resource "%s" is not ISP on server "%s" Manually bring the resource in service and retry the operation	<b>Cause:</b> IP resource {tag} which the listener resource depends on should be ISP.  <b>Action:</b> Perform the steps listed in the message text.
122141	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<b>Cause:</b> There was an unexpected error running "dep_create."  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122144	ERROR	Usage: %s %s	<b>Cause:</b> Invalid parameters were specified for the "create_ins" operation.  <b>Action:</b> Verify the parameters and retry the operation.
122145	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for	<b>Cause:</b> There was an unexpected error

Code	Severity	Message	Cause/Action
		details and retry the operation.	running "app_create."  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122146	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for details and retry the operation.	<b>Cause:</b> There was an unexpected error running "typ_create."  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122147	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for details and retry the operation.	<b>Cause:</b> There was an unexpected error running "newtag."  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122148	ERROR	Error creating resource "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.  <b>Action:</b> Check your LifeKeeper configuration.
122149	ERROR	An error has occurred in utility %s on server %s. View the LifeKeeper logs for details and retry the operation.	<b>Cause:</b> There was an unexpected error running "ins_setstate."  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122150	ERROR	Error creating resource "%s" on server "%s"	<b>Cause:</b> LifeKeeper was unable to create the resource {resource} on {server}.

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Check your LifeKeeper configuration.</p>
122151	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "depstoextend" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122152	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the "extend" operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122153	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122154	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource {resource} on {server}.</p>
122155	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server "%s" for App "%s" and Type "%s"	<p><b>Cause:</b> During the Listener resource extension, a resource instance was found using the same {tag} and/or {id} but with a different resource application and type.</p> <p><b>Action:</b> Resource IDs must be unique. The resource instance with the ID matching the Oracle Listener resource instance must be removed.</p>

Code	Severity	Message	Cause/Action
122156	ERROR	Cannot access extend script "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to extend the resource hierarchy because it was unable to find the script EXTEND on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122157	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "getConfigIps" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122158	ERROR	The file %s is not a valid listener file. The file does not contain any listener definitions.	<p><b>Cause:</b> Failed to find any listener definitions.</p> <p><b>Action:</b> Ensure listener definition is in the listener.ora and retry the operation.</p>
122159	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "getSidListeners" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122160	ERROR	The file %s is not a valid listener file. The file does not contain any listener definitions.	<p><b>Cause:</b> Failed to find any listener definitions.</p> <p><b>Action:</b> Ensure listener definition is in the listener.ora and retry the operation.</p>
122161	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "Isn-display" operation.</p> <p><b>Action:</b> Verify the arguments and retry</p>

Code	Severity	Message	Cause/Action
			the operation.
122162	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
122163	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the updateHelper operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122164	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	<p><b>Cause:</b> LifeKeeper was unable to update the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122166	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid parameters were specified for the "updateHelper" operation.</p> <p><b>Action:</b> Verify the parameters and retry the operation.</p>
122170	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>

Code	Severity	Message	Cause/Action
122171	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122172	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "updIPDeps" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122173	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	<p><b>Cause:</b> LifeKeeper was unable to update the resource {resource} on {server}.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122175	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "rlslocks."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122177	ERROR	Unable to "%s" on "%s"	<p><b>Cause:</b> There was an unexpected error running "getlocks."</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
122180	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<p><b>Cause:</b> There was an unexpected error running "dep_create."</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122181	ERROR	Unable to create a dependency between parent tag %s and child tag %s.	<b>Cause:</b> There was an unexpected error running "dep_create."  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122183	ERROR	The path %s is not a valid file.	<b>Cause:</b> There is no listener.ora file.  <b>Action:</b> Ensure the file exists and retry the operation.
122185	ERROR	The file %s is not a valid listener file. The file does not contain any listener definitions.	<b>Cause:</b> LifeKeeper failed to find any valid listener definitions.  <b>Action:</b> Ensure there are valid listener definitions in the listener.ora and retry the operation.
122186	ERROR	The value specified for %s cannot be empty. Please specify a value for this field.	<b>Cause:</b> The config and/or executable {path} field is empty.  <b>Action:</b> Input a non-empty value for {path} and retry the operation.
122187	ERROR	The path %s is not a valid file or directory.	<b>Cause:</b> The defined {path} is invalid.  <b>Action:</b> Ensure the {path} exists and retry the operation.
122188	ERROR	The path %s is not a valid file or directory.	<b>Cause:</b> There is no {path}.  <b>Action:</b> Ensure the {path} exists and

Code	Severity	Message	Cause/Action
			retry the operation.
122189	ERROR	The value specified for %s cannot be empty. Please specify a value for this field.	<p><b>Cause:</b> The config and/or executable Path field is empty.</p> <p><b>Action:</b> Input path for the field.</p>
122190	ERROR	Usage: %s %s	<p><b>Cause:</b> Invalid arguments were specified for the "valid_rpath" operation.</p> <p><b>Action:</b> Verify the arguments and retry the operation.</p>
122191	ERROR	The values specified for the target and the template servers are the same.	<p><b>Cause:</b> Invalid argument of valid_rpath.</p> <p><b>Action:</b> Ensure arguments and retry the operation.</p>
122192	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<p><b>Cause:</b> There is no oratab file in /etc/oratab or {path}.</p> <p><b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.</p>
122193	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<p><b>Cause:</b> There is no oratab file in /etc/oratab or {path}.</p> <p><b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.</p>
122194	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<p><b>Cause:</b> There is no oratab file in /etc/oratab or {path}.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.
122195	ERROR	Unable to find the configuration file "oratab" in its default locations, /etc/oratab or %s on "%s"	<b>Cause:</b> There is no oratab file in /etc/oratab or {path}.  <b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.
122196	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	<b>Cause:</b> LifeKeeper was unable to remove the resource {resource} on {server}.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.
122197	ERROR	Unable to find the configuration file \"oratab\" in its default locations, /etc/oratab or \$listener::oraTab on \"\$me\"	<b>Cause:</b> There is no oratab file in /etc/oratab or {path}.  <b>Action:</b> Ensure oratab file exists in {path} or ORACLE_ORATABLOC in /etc/default/Lifekeeper is set to a valid path.
122198	ERROR	remove for \$okListener failed.	

## 8.1.11. Oracle PDB Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
122251	ERROR	Update of pluggable database info field for "%s" on "%s" failed (%s).	
122252	ERROR	Initial connect with query buffer to database "%s" on "%s" failed, testing output.	
122253	ERROR	The Oracle database "%s" is not running or no open connections are available on "%s".	
122261	ERROR	The Oracle resource (%s) and dependency are not set on %s.	
122262	ERROR	Usage: %s %s	
122263	ERROR	The restore of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_RESTORE_TIMEOUT in /etc/default/LifeKeeper.	
122264	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122268	ERROR	Failed to create object instance for Oracle on "%s".	
122269	ERROR	no dependency for Oracle on "%s".	
122270	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122271	ERROR	Usage: %s %s	
122272	ERROR	The remove of %s has timed out on server %s. The default TIMEOUT is 300 seconds. To increase the TIMEOUT, set ORACLE_REMOVE_TIMEOUT in /etc/default/LifeKeeper.	
122273	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122277	ERROR	Usage: %s %s	

Code	Severity	Message	Cause/Action
122278	ERROR	Failed to create object instance for Oracle on "%s".	
122279	ERROR	no dependency for Oracle on "%s".	
122280	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122281	ERROR	The quickCheck of %s has timed out on server %s. The default TIMEOUT is 45 seconds. To increase the TIMEOUT, set ORACLE_QUICKCHECK_TIMEOUT in /etc/default/LifeKeeper.	
122282	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122284	ERROR	Failed to create object instance for Oracle on "%s".	
122285	ERROR	no dependency for Oracle on "%s".	
122287	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122288	ERROR	Usage: %s %s	
122291	ERROR	Cannot extend resource "%s" to server "%s"	
122292	ERROR	The values specified for the target and the template servers are the same: "%s".	
122294	ERROR	Cannot access canextend script "%s" on server "%s"	
122295	ERROR	Usage: %s %s	
122296	ERROR	DB instance "%s" is not protected on "%s".	
122297	ERROR	Failed to create object instance for OraclePDB on "%s".	
122298	ERROR	Unable to locate the oratab file "%s" on "%s".	
122299	ERROR	END failed hierarchy "%s" of resource "%s" on server "%s" with return value of %d	
122301	ERROR	Unable to "%s" on "%s" during resource create.	

Code	Severity	Message	Cause/Action
122302	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122304	ERROR	Unable to "%s" on "%s" during resource create.	
122305	ERROR	Unable to determine Oracle user for "%s" on "%s".	
122306	ERROR	Error creating resource "%s" on server "%s"	
122308	ERROR	Dependency creation between Oracle pluggable database "%s (%s)" and the dependent resource "%s" on "%s" failed. Reason	
122309	ERROR	%s	
122311	ERROR	In-service attempted failed for tag "%s" on "%s".	
122312	ERROR	Usage: %s %s	
122313	ERROR	Create of app "%s" on "%s" failed with return code of "%d".	
122314	ERROR	Create of typ "%s" for app "%s" on "%s" failed with return code of "%d".	
122316	ERROR	Create of resource tag via "newtag" on "%s" failed.	
122318	ERROR	Error creating resource "%s" on server "%s"	
122320	ERROR	Setting "resstate" for resource "%s" on "%s" failed with return code of "%d".	
122321	ERROR	Error creating resource "%s" on server "%s"	
122322	ERROR	Usage: %s %s	
122323	ERROR	Usage: %s %s	
122324	ERROR	Usage: %s %s	
122325	ERROR	Cannot extend resource "%s" to server "%s"	
122326	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server "%s" for App "%s" and Type "%s"	

Code	Severity	Message	Cause/Action
122327	ERROR	Error creating resource "%s" on server "%s"	
122328	ERROR	Cannot access extend script "%s" on server "%s"	
122329	ERROR	Cannot extend resource "%s" to server "%s"	
122330	ERROR	Usage: %s %s	
122331	ERROR	Failed to create object instance for OraclePDB on "%s".	
122332	ERROR	Usage: %s %s	
122334	ERROR	Backup node %s is unreachable; abort protection PDB changes.	
122336	ERROR	Update of protection PDB failed for "%s" on "%s".	
122339	ERROR	Usage: %s %s	
122340	ERROR	Usage: %s %s	
122341	ERROR	Usage: %s %s	
122342	ERROR	Failed to create object instance for OraclePDB on "%s".	
122343	ERROR	Usage: %s %s	
122344	ERROR	END failed %s of "%s" on server "%s" due to a "%s" signal	
122348	ERROR	Failed to create flag "%s" on "%s".	
122350	ERROR	Failed to create object instance for Oracle on "%s".	
122351	ERROR	no dependency for Oracle on "%s".	
122353	ERROR	Unable to find the Oracle executable "%s" on "%s".	
122356	ERROR	Usage: %s %s	
122357	ERROR	Failed to create object instance for OraclePDB on "%s".	
122358	ERROR	The selected oracle SID "%s" is not a CDB.	
122359	ERROR	No protectable PDB found for the	

Code	Severity	Message	Cause/Action
		selected SID "%s".	
122360	ERROR	No protected Oracle database found on "%s".	

## 8.1.12. SCSI Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
125102	ERROR	<code>`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`</code>	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125103	ERROR	<code>`printf '%s is not shareable with any machine.' \$DEV`</code>	<b>Cause:</b> The device does not appear to be shared with any other systems.  <b>Action:</b> Verify that the device is accessible from all servers in the cluster. Ensure that all relevant storage drivers and software are installed and configured properly.
125104	ERROR	<code>`printf 'Failed to create disk hierarchy for "%s" on "%s"' \$PRIMACH \$DEV`</code>	<b>Cause:</b> The creation of a resource to protect a physical disk failed.  <b>Action:</b> Check adjacent log messages for more details and try to resolve the reported problem.
125107	ERROR	<code>`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`</code>	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125114	ERROR	<code>`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`</code>	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125120	ERROR	<code>`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateName \$TemplateSysName`</code>	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125123	ERROR	<code>`printf 'Cannot access depstoextend script "%s" on server "%s"' \$depstoextend \$TargetSysName`</code>	<b>Cause:</b> LifeKeeper was unable to run pre-extend checks on the resource

Code	Severity	Message	Cause/Action
			<p>hierarchy because it was unable to find the script "DEPSTOEXTEND" on {server}.</p> <p><b>Action:</b> Check your LifeKeeper configuration.</p>
125126	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$ChildTag \$TemplateSysName`	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125129	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateTagName \$TemplateSysName`	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125155	ERROR	SCSI \$DEV failed to lock.	<p><b>Cause:</b> There was a problem locking a SCSI device.</p> <p><b>Action:</b> Check adjacent log messages for more details and try to resolve the reported problem.</p>
125164	ERROR	SCSI \$INFO failed to unlock.	<p><b>Cause:</b> There was a problem unlocking a SCSI device.</p> <p><b>Action:</b> Check adjacent log messages for more details and try to resolve the reported problem.</p>
125181	ERROR	`printf 'Template resource "%s" on server "%s" does not exist' \$TemplateTag \$TemplateSysName`	<b>Cause:</b> LifeKeeper was unable to find the resource {tag} on {server}.
125194	ERROR	Failed to check disk.(tag="\\$opt_t\	

## 8.1.13. Quick Service Protection Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
134003	ERROR	catch a \" <code>\$sig</code> \" signal	<p><b>Cause:</b> The "create" process was interrupted by a signal.</p> <p><b>Action:</b> Check adjacent log messages.</p>
134004	ERROR	Unable to getlocks on \$server during resource create. Error ( <code>\$rc</code> )	<p><b>Cause:</b> Failed to get the administrative lock when creating a resource hierarchy.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
134005	ERROR	The service \" <code>\$serviceName</code> \" is not supported on \$server. Error ( <code>\$rc</code> )	<p><b>Cause:</b> The service does not exist or cannot be protected.</p> <p><b>Action:</b> Input the appropriate service name.</p>
134006	ERROR	The service \" <code>\$serviceName</code> \" is already protected on \$server.	<p><b>Cause:</b> This service is already protected.</p> <p><b>Action:</b> Can not create a resource for the protection of this service.</p>
134007	ERROR	Error creating resource \$tag. Error ( <code>\$rc</code> )	<p><b>Cause:</b> LifeKeeper was unable to create the resource instance.</p> <p><b>Action:</b> Check adjacent log messages. Correct the cause of the error.</p>

Code	Severity	Message	Cause/Action
134011	ERROR	In-service attempted failed for tag \$tag.	<p><b>Cause:</b> Failed to restore QSP resource.</p> <p><b>Action:</b> Check the log related to the service that you want to protect. Resolve the problem.</p>
134015	ERROR	Unable to rlslocks on \$server during resource create. Error (\$rc	<p><b>Cause:</b> Failed to release lock after QSP resource created.</p> <p><b>Action:</b> Check the logs for related errors and try to resolve the reported problem.</p>
134103	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" does not exist	<p><b>Cause:</b> The resource cannot be found on the template server.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template server before extending.</p>
134104	ERROR	Template resource \"\$template_tag\" on server \"\$template_sys\" is not QSP resource (app=\$ins <sup>1</sup> , res=\$ins <sup>2</sup>	<p><b>Cause:</b> The template resource is not QSP resource.</p> <p><b>Action:</b> Expand to the same type of resources as a template resource.</p>
134105	ERROR	The service \"\$service\" is not supported on \$me. Error (\$check	<p><b>Cause:</b> The service does not exist on target server.</p> <p><b>Action:</b> Install the service on target server before extending.</p>
134106	ERROR	The service \"\$service\" is already protected on \$me.	<p><b>Cause:</b> There is already a resource of the same ID on target server.</p> <p><b>Action:</b> Cannot create the resource of the same service.</p>

Code	Severity	Message	Cause/Action
134203	ERROR	catch a \"\$\$sig\" signal	<p><b>Cause:</b> The "extend" process resource was interrupted by a signal.</p> <p><b>Action:</b> Check adjacent log messages.</p>
134204	ERROR	Template resource \"\$\$template_tag\" on server \"\$\$template_sys\" does not exist	<p><b>Cause:</b> The resource cannot be found on the template server.</p> <p><b>Action:</b> Ensure the hierarchy is correct on the template server before extending.</p>
134208	ERROR	Error creating resource \"\$\$tag\" on server \"\$\$me\"	<p><b>Cause:</b> LifeKeeper was unable to create the resource instance on target server.</p> <p><b>Action:</b> Check adjacent log messages. Correct the cause of the error.</p>
134401	ERROR	timeout \$cmd for \"\$\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "restore" process of the service does not terminate within the specified time.</p> <p><b>Action:</b> Check about the protected service and retry the "restore" operation. Also check the logs for related errors and try to resolve the reported problem.</p>
134405	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p> <p><b>Action:</b> Determine why fork fails.</p>
134407	ERROR	service command has failed for \"\$\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p> <p><b>Action:</b> It is an error to manually run</p>

Code	Severity	Message	Cause/Action
			the service command with "start" option. Correct the cause of the error by reference to error message.
134501	ERROR	timeout \$cmd for \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "remove" process of the service does not terminate within the specified time.</p> <p><b>Action:</b> Check about the protected service and retry the "remove" operation. Also check the logs for related errors and try to resolve the reported problem.</p>
134505	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p> <p><b>Action:</b> Determine why fork fails.</p>
134507	ERROR	service command has failed for \"\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p> <p><b>Action:</b> It is an error to manually run the service command with "stop" option. Correct the cause of the error by reference to error message.</p>
134601	ERROR	timeout \$cmd for \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "quickCheck" process will be forcibly terminated due to a waiting time over the user defined timeout.</p> <p><b>Action:</b> Check about the protected service. Also check the logs for related errors and try to resolve the reported problem.</p>
134605	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Determine why fork fails.
134607	ERROR	service command has failed for \"\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p> <p><b>Action:</b> It is an error to manually run the service command with "status" option. Correct the cause of the error by reference to error message.</p>
134701	ERROR	timeout \$cmd for \"\$tag\". Forcibly terminating.	<p><b>Cause:</b> The "recover" process of the service does not terminate within the specified time.</p> <p><b>Action:</b> Check about the protected service. Also check the logs for related errors and try to resolve the reported problem.</p>
134706	FATAL	Failed to fork process to execute service command: \$!	<p><b>Cause:</b> Failed to fork. This is a system error.</p> <p><b>Action:</b> Determine why fork fails.</p>
134708	ERROR	service command has failed for \"\$tag\"	<p><b>Cause:</b> Failed to execute service command.</p> <p><b>Action:</b> It is an error to manually run the service command with "start" option. Correct the cause of the error by reference to error message.</p>
134803	ERROR	tag \"\$tag\" does not exist on server \"\$me\"	<b>Cause:</b> The specified tag does not exist. This is an internal error.
134804	ERROR	app type \"\$ins <sup>1</sup> \" is not \$app	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.

Code	Severity	Message	Cause/Action
134805	ERROR	res type \"\$ins <sup>2</sup> \" is not \$res	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134823	ERROR	tag \"\$tag\" does not exist on server \"\$me\"	<b>Cause:</b> The specified tag does not exist. This is an internal error.
134824	ERROR	app type \"\$ins <sup>1</sup> \" is not \$app	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134825	ERROR	res type \"\$ins <sup>2</sup> \" is not \$res	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134843	ERROR	tag \"\$tag\" does not exist	<b>Cause:</b> The specified tag does not exist. This is an internal error.
134844	ERROR	app type \"\$ins <sup>1</sup> \" is not \$app	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.
134845	ERROR	res type \"\$ins <sup>2</sup> \" is not \$res	<b>Cause:</b> The specified tag is not QSP resource. This is an internal error.

## 8.1.14. GUI Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
000200	ERROR	pam_start() failed	
000201	ERROR	pam_authenticate failed (user %s, retval %d	
000202	ERROR	pam_end() failed?!?!	
000203	ERROR	Did not find expected group 'lkguest'	
000204	ERROR	Did not find expected group 'lkoper'	
000205	ERROR	Did not find expected group 'lkadmin'	
000208	ERROR	pam_setcred establish credentials failed (user %s, retval %d	<p><b>Cause:</b> Unable to establish valid login credentials for user {user}. The pam_setcred call returned: {retval}.</p> <p><b>Action:</b> Check /var/log/security and /var/log/messages for more information.</p>
000209	ERROR	pam_setcred delete credentials failed (user %s, retval %d	<p><b>Cause:</b> Unable to clear login credentials for user {user}. The pam_setcred call returned: {retval}.</p> <p><b>Action:</b> Check /var/log/security and /var/log/messages for more information.</p>
000902	ERROR	Error removing system name from loopback address line in /etc/hosts file. You must do this manually before starting the GUI server.	<p><b>Cause:</b> System name did not get removed from /etc/hosts file.</p> <p><b>Action:</b> Remove system name manually then restart the GUI server, then enter the following: run &lt;action name&gt;</p>
000918	ERROR	LifeKeeper GUI Server error during Startup	<p><b>Cause:</b> The GUI server terminated due</p>

Code	Severity	Message	Cause/Action
			to an abnormal condition.  <b>Action:</b> Check the logs for related errors and try to resolve the reported problem.

## 8.1.15. SAP Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
112000	FATAL	Usage: "%s %s". Specify the correct usage for the requested command.	<p><b>Cause:</b> Invalid parameters passed to the SAP create script.</p> <p><b>Action:</b> Please provide appropriate parameters for the SAP create script.</p>
112004	ERROR	Neither the SID/instance pair nor the tag parameter were specified for the internal "%s" routine on %s. If this was a command line operation, specify the correct parameters. Otherwise, consult the troubleshooting documentation.	<p><b>Cause:</b> Invalid parameters were specified when trying to create an internal SAP object.</p> <p><b>Action:</b> Provide either the SID and instance or the LifeKeeper resource tag name.</p>
112013	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and/or is readable.	<p><b>Cause:</b> The sapservices file either does not exist or is not readable.</p> <p><b>Action:</b> Verify that the sapservices file exists and is readable.</p>
112017	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112019	ERROR	The attempt to update the resource information field for resource %s has failed on %s. View the resource properties manually using "ins_list -t <tag>" to verify that the resource is functional.	<p><b>Cause:</b> Unable to update the LifeKeeper resource information field for the given resource on the given server.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Verify that the resource exists and that LifeKeeper is running on the given server.</p>
112022	ERROR	<p>An error occurred while trying to find the IP address corresponding to "%s" on %s. Verify the IP address or host name exists in DNS or the hosts file.</p>	<p><b>Cause:</b> Unable to find the given IP address or host name on the given server.</p> <p><b>Action:</b> Verify that the IP address or DNS name exists in DNS or in the local hosts file.</p>
112024	FATAL	<p>There was an error verifying the NFS connections for SAP related mount points on \$me. One or more NFS servers is not operational and needs to be restarted.</p>	<p><b>Cause:</b> At least one NFS shared file system listed in the SAP_NFS_CHECK_DIRS parameter is currently unavailable.</p> <p><b>Action:</b> Verify that the NFS server is alive, all necessary NFS-related services are running, and that all necessary file systems are being exported.</p>
112028	ERROR	<p>Unable to determine the user name for the SAP administrative user for resource %s on %s.</p>	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112029	ERROR	<p>The canextend script "%s" either does not exist or is not executable on %s.</p>	<p><b>Cause:</b> The SAP canextend script either does not exist or is not executable.</p> <p><b>Action:</b> Verify that the SAP canextend script exists and is executable.</p>

Code	Severity	Message	Cause/Action
112037	ERROR	Unable to create an internal object for the SAP instance using SID %s, instance %s, and tag %s on server %s. Verify that all necessary SAP file systems are mounted and accessible before reattempting the operation.	<p><b>Cause:</b> Unable to create an internal SAP object to represent the given SAP instance.</p> <p><b>Action:</b> Verify that the SAP instance is properly installed and configured and that all necessary file systems are mounted.</p>
112040	ERROR	The SAP Directory "%s" ("%s") does not exist on %s. Verify that the directory exists and that the SAP software is properly installed.	<p><b>Cause:</b> The given SAP installation directory does not exist on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted.</p>
112041	ERROR	The required utility "%s" was not found or was not executable on %s. Verify the SAP installation and location of the required utility.	<p><b>Cause:</b> The saphostexec or saposcol utility either could not be found or is not executable.</p> <p><b>Action:</b> Verify that the SAP Host Agent package is installed correctly and that all necessary file systems are mounted.</p>
112042	ERROR	One or more SAP or LifeKeeper validation checks has failed on %s. Please update the SAP software on this host to include the SAPHOST and SAPCONTROL packages.	<p><b>Cause:</b> The saphostexec or saposcol utility either could not be found or is not executable.</p> <p><b>Action:</b> Verify that the SAP Host Agent package is installed correctly and that all necessary file systems are mounted.</p>
112048	ERROR	The SAP instance %s is already under LifeKeeper protection on server %s. Choose another SAP instance to protect or specify the correct instance.	<p><b>Cause:</b> The given SAP instance is already protected by LifeKeeper on the given server.</p> <p><b>Action:</b> Choose an SAP instance which is not already under LifeKeeper</p>

Code	Severity	Message	Cause/Action
			protection.
112049	ERROR	Unable to locate the SAP Mount directory on %s. Verify that all SAP file systems are mounted and accessible before reattempting the operation.	<p><b>Cause:</b> Unable to determine the location of the SAP Mount (sapmnt) directory.</p> <p><b>Action:</b> Verify that all necessary file systems are mounted and that the all necessary SAP instance profiles are accessible.</p>
112050	ERROR	Detected multiple virtual IP addresses/ host names for instance %s on %s. Verify that the instance is configured correctly.	<p><b>Cause:</b> Multiple virtual IPs or hostnames were detected for the given SAP instance on the given server.</p> <p><b>Action:</b> Verify that the virtual IP or hostname associated to the instance is configured correctly.</p>
112051	ERROR	The "%s" or "%s" value in the default profile "%s" is still set to the physical host name on %s. The value(s) must be set to a virtual host name.	<p><b>Cause:</b> The given host name parameter is set to a physical server host name on the given server.</p> <p><b>Action:</b> Set the given host name parameter to a virtual host name.</p>
112053	ERROR	Detected multiple instances under SID %s with the same instance number (%s) on %s. Each instance within a particular SID must have a unique instance number.	<p><b>Cause:</b> Multiple SAP instances with the same instance number were detected under the same SAP SID.</p> <p><b>Action:</b> Reconfigure the SAP environment so that each instance under a given SAP SID has a unique instance number.</p>
112056	ERROR	The NFS export for the path "%s" required by the instance %s for the "%s" directory does not have an NFS hierarchy protecting it on %s. You must create an NFS hierarchy to protect this NFS export	<p><b>Cause:</b> The NFS export for the given file system is not currently protected by LifeKeeper.</p>

Code	Severity	Message	Cause/Action
		before creating the SAP resource hierarchy.	<b>Action:</b> Create a LifeKeeper NFS hierarchy for the given exported file system and reattempt SAP resource creation.
112057	ERROR	Unable to create a file system resource hierarchy for the file system "%s" on %s.	<b>Cause:</b> Unable to create a LifeKeeper file system resource hierarchy to protect the given file system on the given server.  <b>Action:</b> Check the LifeKeeper and system logs for more information.
112058	ERROR	Unable to create a dependency between parent tag "%s" and child tag "%s" on "%s".	<b>Cause:</b> Unable to create a LifeKeeper dependency between the given resources on the given server.  <b>Action:</b> Check the LifeKeeper logs for more information.
112060	ERROR	All attempts at local recovery for the SAP resource %s have failed on %s. A failover to the backup server will be attempted.	<b>Cause:</b> Unable to recover the given SAP resource on the given server.  <b>Action:</b> A failover of the SAP resource hierarchy will be attempted automatically. No user intervention is required.
112061	ERROR	The values specified for the target and the template servers are the same. Please specify the correct values for the target and template servers.	<b>Cause:</b> The template and target servers provided during SAP resource extension are the same.  <b>Action:</b> Provide the correct names for the template and target servers and reattempt the extend operation.
112062	ERROR	Unable to find the home directory "%s" for the SAP administrative user "%s" on %s. Verify that the SAP software is installed correctly.	<b>Cause:</b> Unable to find the home directory for the given SAP user on the

Code	Severity	Message	Cause/Action
			<p>given server.</p> <p><b>Action:</b> Verify that the SAP software is installed correctly and that the appropriate SAP administrative user for the given SID exists on the server.</p>
112063	ERROR	The SAP administrative user "%s" does not exist on %s. Verify that the SAP software is installed correctly or create the required SAP user on %s.	<p><b>Cause:</b> The given SAP administrative user does not exist on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and create the required SAP administrative user if necessary.</p>
112064	ERROR	The group ID for user "%s" is not the same on template server "%s" and target server "%s". Please correct the group ID for the user so that it is the same on the template and target servers.	<p><b>Cause:</b> The group ID's for the given SAP administrative user on the template and target servers do not match.</p> <p><b>Action:</b> Modify the group ID of the given SAP administrative user so that it is the same on the template and target servers.</p>
112065	ERROR	The user ID for user "%s" is not the same on template server "%s" and target server "%s". Please correct the user ID so that it is the same on the template and target servers.	<p><b>Cause:</b> The user ID's for the given SAP administrative user on the template and target servers do not match.</p> <p><b>Action:</b> Modify the user ID of the given SAP administrative user so that it is the same on the template and target servers.</p>
112066	ERROR	Required SAP utilities could not be found in "%s" on %s. Verify that the SAP software is installed correctly.	<p><b>Cause:</b> Unable to locate necessary SAP executables or the SAP instance profile.</p> <p><b>Action:</b> Verify that the SAP software is</p>

Code	Severity	Message	Cause/Action
			properly installed and configured and that all necessary file systems are mounted.
112069	ERROR	The command "%s" is not found in the "%s" perl module ("%s") on %s. Please check the command specified and retry the operation.	<p><b>Cause:</b> The given command was not found in the sap perl module on the given server.</p> <p><b>Action:</b> If this error resulted from a user-initiated command line action, verify that the correct routine name was provided to the remoteControl script. If this error occurred during normal LifeKeeper operation, please submit an issue report to SIOS customer support.</p>
112071	ERROR	The file "%s" exists, but was not read and write enabled on server %s. Enable read and write permissions on the specified file.	<p><b>Cause:</b> The given file does not have read/write permissions enabled on the given server.</p> <p><b>Action:</b> Enable read/write permissions on the given file.</p>
112073	ERROR	Unable to create an internal object for the SAP instance using either SID "%s" and instance "%s" or tag "%s" on server "%s". Either the values specified for the object initialization (SID/instance pair or tag, system) were not valid, or an error occurred while attempting to gather information about the SAP instance. If all specified parameters are correct, verify that all necessary SAP file systems are mounted and accessible before reattempting the operation.	<p><b>Cause:</b> Unable to create an internal SAP object to represent the given SAP instance.</p> <p><b>Action:</b> Verify that the SAP instance is properly installed and configured and that all necessary file systems are mounted.</p>
112074	WARN	WARNING: The profile "%s" for SID %s and instance %s has Autostart enabled on %s. Disable Autostart for the specified instance by setting Autostart=0 in the profile.	<p><b>Cause:</b> The Autostart parameter is enabled in the given instance profile on the given server.</p> <p><b>Action:</b> Disable Autostart for the given</p>

Code	Severity	Message	Cause/Action
			SAP instance by setting 'Autostart = 0' in the instance profile.
112076	FATAL	Unable to start the sapstartsrv service for SID \$sid and instance \$Instance on \$me. Verify that the sapservices file is correct and that the process can be started manually.	<p><b>Cause:</b> Unable to start the SAP Start Service (sapstartsrv) process for the given SAP instance.</p> <p><b>Action:</b> Verify that the sapservices file contains the appropriate command to start the sapstartsrv process and that the process can be started manually.</p>
112077	ERROR	Unable to stop the sapstartsrv service for SAP SID %s and SAP instance %s on %s. Verify that the sapservices file is correct and the process can be stopped manually.	<p><b>Cause:</b> Unable to stop the SAP Start Service (sapstartsrv) process for the given SAP instance.</p> <p><b>Action:</b> Verify that the sapservices file contains the appropriate command to start the sapstartsrv process and that the process can be stopped manually.</p>
112078	ERROR	ERSv1 is only supported in two-node clusters. Resource %s is unable to be extended to system %s. Upgrade to ERSv2 in order to extend the hierarchy to three or more nodes.	<p><b>Cause:</b> Unable to extend an SAP resource representing an ERSv1 instance to three or more nodes.</p> <p><b>Action:</b> In order to extend an SAP resource representing an ERS instance to three or more nodes, upgrade to ERSv2. Upgrade instructions are provided in the online product documentation.</p>
112082	WARN	Instance %s is running a different version of the enqueue server than its corresponding enqueue replication server. This configuration is not supported by SAP and will lead to unexpected resource behavior. See SAP Note 2711036 – "Usage of the Standalone Enqueue Server 2 in an HA Environment" for more details. Please review the online	<p><b>Cause:</b> The versions of the enqueue server and enqueue replication server do not match.</p> <p><b>Action:</b> Consult the online product documentation for instructions on how to modify the instance profiles so that</p>

Code	Severity	Message	Cause/Action
		product documentation for instructions on how to modify the instance profiles for the enqueue server and enqueue replication server so that they use the same version.	the enqueue server and enqueue replication server are using the same version.
112086	ERROR	The ERS resource corresponding to resource %s is in-service and maintaining backup locks on a remote system. Bringing resource %s in-service on %s would result in a loss of the backup lock table. Please bring resource %s in-service on the system where the corresponding ERS resource is currently in-service in order to maintain consistency of the lock table. In order to force resource %s in-service on %s, either (i) run the command <code>\opt/LifeKeeper/bin/flg_create -f sap_cs_force_restore_%s'</code> as root on %s and reattempt the in-service operation or (ii) take the corresponding ERS resource out of service on the remote system. Both of these actions will result in a loss of the backup lock table.	<p><b>Cause:</b> An in-service operation was attempted for an ASCS/SCS resource while its corresponding ERS instance was running and storing a backup enqueue table on a different server in the cluster.</p> <p><b>Action:</b> Bring the ASCS/SCS resource in-service on the server where its corresponding ERS instance is running in order for it to retrieve the backup enqueue table. If the ASCS/SCS resource must be forced in-service on the given node, either (i) run the command <code>'/opt/LifeKeeper/bin/flg_create -f sap_cs_force_restore_&lt;ASCS/SCS Tag&gt;'</code> and reattempt the in-service operation, or (ii) take the corresponding ERS resource out of service on the remote server. Both of these actions will result in a loss of the backup enqueue table.</p>
112089	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, or TYPE_NEW_ERS (1, 2, or 5).	<p><b>Cause:</b> The GetEnqVersion routine was called for an unsupported SAP instance type. Only instance types 1 (TYPE_CS), 2 (TYPE_ERS), and 5 (TYPE_NEW_ERS) are supported.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112092	ERROR	The profile "%s" either does not exist or cannot be read on %s. Unable to determine whether enqueue replication is enabled for resource %s. Please verify	<p><b>Cause:</b> The given instance profile either does not exist or cannot be read on the</p>

Code	Severity	Message	Cause/Action
		that the file exists and can be read.	<p>given server.</p> <p><b>Action:</b> Verify that the file exists and is read-enabled</p>
112095	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The given resource could not be created on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted and accessible on the given server.</p>
112096	ERROR	Resource %s is not currently in-service on server %s. Manually bring the resource in-service and retry the operation.	<p><b>Cause:</b> While attempting to create a dependency, the given resource was not in-service on the given server.</p> <p><b>Action:</b> Bring the resource in-service on the given server and retry the operation.</p>
112101	ERROR	Error getting resource information for resource "%s" on server "%s"	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112102	ERROR	Cannot extend resource "%s" to server "%s"	<p><b>Cause:</b> The given resource cannot be extended to the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed on the target server and that any necessary shared file systems are accessible from the target system.</p>

Code	Severity	Message	Cause/Action
112103	ERROR	Error creating resource "%s" on server "%s"	<p><b>Cause:</b> The given resource could not be created on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted and accessible on the given server.</p>
112104	ERROR	The extend script "%s" either does not exist or is not executable on %s.	<p><b>Cause:</b> The given extend script does not exist or is not executable on the given server.</p> <p><b>Action:</b> Verify that all necessary recovery kits are installed and that the given extend script is executable.</p>
112106	ERROR	Unable to create an internal SAP object for resource "%s" on %s. If the tag is correct, verify that all necessary SAP file systems are mounted and accessible.	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112112	ERROR	Error getting resource information for resource "%s" on server "%s".	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112115	FATAL	Error getting resource information for resource \$tag on server \$sap::me.	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>

Code	Severity	Message	Cause/Action
112119	ERROR	Error getting resource information for resource "%s" on server "%s".	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112123	ERROR	Error getting resource information for resource "%s" on server "%s".	<p><b>Cause:</b> Unable to find information about the given resource on the given server.</p> <p><b>Action:</b> Verify that the given resource exists and that all necessary file systems are mounted on the given server.</p>
112124	ERROR	Resource with either matching tag "%s" or id "%s" already exists on server %s for app "%s" and type "%s".	<p><b>Cause:</b> A resource with the given resource tag or ID and the same app and type already exists on the given server.</p> <p><b>Action:</b> Verify that the SAP instance is not already under LifeKeeper protection on the given server. If it is not already protected, choose a different resource tag name.</p>
112125	ERROR	Unable to create an SAP object for resource %s on system %s.	<p><b>Cause:</b> Unable to create an internal SAP object to represent the given resource on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted and accessible on the given server.</p>
112126	FATAL	Usage: "%s %s". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP canextend script.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Specify the correct parameters for the canextend script: canextend &lt;template server&gt; &lt;template tag&gt;</p>
112127	FATAL	Usage: \"\$cmd \$usage\" Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP delete script.</p> <p><b>Action:</b> Specify the correct parameters for the delete script: delete [-U] -t &lt;tag&gt; -i &lt;id&gt;</p>
112128	FATAL	Usage: \"%s %s\" Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP depstoextend script.</p> <p><b>Action:</b> Specify the correct parameters for the depstoextend script: depstoextend &lt;template server&gt; &lt;template tag&gt;</p>
112129	FATAL	Usage: \"%s %s\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP extend script.</p> <p><b>Action:</b> Specify the correct parameters for the extend script: extend &lt;template server&gt; &lt;template tag&gt; &lt;switchback&gt; &lt;target tag&gt;</p>
112130	FATAL	Usage: \" \$cmd \$usage\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP quickCheck script.</p> <p><b>Action:</b> Specify the correct parameters for the quickCheck script: quickCheck -t &lt;tag&gt; -i &lt;id&gt;</p>
112131	FATAL	Usage: \"%s %s\" Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP recover script.</p> <p><b>Action:</b> Specify the correct parameters</p>

Code	Severity	Message	Cause/Action
			for the recover script: recover -d <tag>
112132	FATAL	Usage: \"\$cmd \$usage\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP remoteControl script.</p> <p><b>Action:</b> Specify the correct parameters for the remoteControl script: remoteControl &lt;tag&gt; &lt;remote instance&gt; &lt;remote cmd&gt; &lt;remote cmd option&gt; &lt;primary system&gt; &lt;primary tag&gt;</p>
112133	FATAL	Usage: \"%s %s\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP remove script.</p> <p><b>Action:</b> Specify the correct parameters for the remove script: remove -t &lt;tag&gt; -i &lt;id&gt;</p>
112134	FATAL	Usage: \"\$cmd \$usage\". Specify the correct usage for the requested command.	<p><b>Cause:</b> Incorrect usage of the SAP restore script.</p> <p><b>Action:</b> Specify the correct parameters for the restore script: restore -t &lt;tag&gt; -i &lt;id&gt;</p>
112137	ERROR	The required parameter \"parent\" was either not provided or was invalid in the \$func routine on \$me.	<p><b>Cause:</b> Incorrect usage of the CleanUp routine.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112138	ERROR	At least one required process for instance %s was not started successfully during \"%s\" on server %s. Please check the LifeKeeper and system logs for additional information.	<p><b>Cause:</b> At least one required process for the given SAP instance did not start successfully on the given server.</p> <p><b>Action:</b> Correct any issues found in the LifeKeeper or system logs or SAP trace</p>

Code	Severity	Message	Cause/Action
			files and retry the operation.
112140	FATAL	The tag parameter was not specified for the internal \"\$func\" routine on \$me. If this was a command line operation, specify the correct parameters. Otherwise, consult the troubleshooting documentation.	<p><b>Cause:</b> The tag parameter was not specified in the GetLK routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112141	ERROR	Either the SID ("%s") or instance ("%s") parameter was not specified for the "%s" routine on %s.	<p><b>Cause:</b> Either the SID or instance parameter was not specified in the StatusSapServer routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112142	ERROR	Either the SID ("%s"), instance ("%s"), or instance number ("%s") parameter was not specified for the "%s" routine on %s.	<p><b>Cause:</b> Either the SID, instance, or instance number parameter was not provided to the StartSapServer routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112143	ERROR	The SID, instance, or instance number parameter was not specified for the "%s" routine on %s.	<p><b>Cause:</b> Either the SID, instance, or instance number parameter was not specified in the StopSapServer routine on the given server.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112173	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and is readable.	<p><b>Cause:</b> The given file does not exist or</p>

Code	Severity	Message	Cause/Action
			<p>is not readable on the given server.</p> <p><b>Action:</b> Verify that the file exists and/or modify its permissions so that it is readable.</p>
112174	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and is readable.	<p><b>Cause:</b> The given file does not exist or is not readable on the given server.</p> <p><b>Action:</b> Verify that the file exists and/or modify its permissions so that it is readable.</p>
112175	ERROR	The file "%s" does not exist or was not readable on %s. Verify that the specified file exists and is readable.	<p><b>Cause:</b> The given file does not exist or is not readable on the given server.</p> <p><b>Action:</b> Verify that the file exists and/or modify its permissions so that it is readable.</p>
112194	ERROR	There was an error verifying the NFS connections for SAP related mount points on %s. One or more NFS servers is not operational and needs to be restarted.	<p><b>Cause:</b> At least one critical NFS shared file system whose mount point is listed in the SAP_NFS_CHECK_DIRS entry in /etc/default/LifeKeeper is currently unavailable.</p> <p><b>Action:</b> Verify that all necessary NFS shared file systems are accessible and restart any NFS server which is not currently operational.</p>
112195	FATAL	There was an error verifying the NFS connections for SAP related mount points on \$me. One or more NFS servers is not operational and needs to be restarted.	<p><b>Cause:</b> At least one critical NFS shared file system whose mount point is listed in the SAP_NFS_CHECK_DIRS entry in /etc/default/LifeKeeper is currently unavailable.</p> <p><b>Action:</b> Verify that all necessary NFS shared file systems are accessible and</p>

Code	Severity	Message	Cause/Action
			restart any NFS server which is not currently operational.
112196	WARN	There was an error verifying the NFS connections for SAP related mount points on %s. One or more NFS servers is not operational and needs to be restarted.	<p><b>Cause:</b> At least one critical NFS shared file system whose mount point is listed in the SAP_NFS_CHECK_DIRS entry in /etc/default/LifeKeeper is currently unavailable.</p> <p><b>Action:</b> Verify that all necessary NFS shared file systems are accessible and restart any NFS server which is not currently operational.</p>
112201	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s.	<p><b>Cause:</b> No resource tag argument was provided to the GetLKEquiv routine.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112203	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s. Additional information available in the LifeKeeper and system logs.	<p><b>Cause:</b> The SAP instance number was not provided to the IsInstanceRunning routine.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112204	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s.	<p><b>Cause:</b> Unable to determine either the appropriate saphostexec or saposcol command to use.</p> <p><b>Action:</b> If using the SAP_SRVHOST_CMD, SAP_HOSTCTL_CMD, or SAP_OSCOL_CMD LifeKeeper tunable values to provide the appropriate commands for a version of SAP</p>

Code	Severity	Message	Cause/Action
			NetWeaver prior to SAP kernel 7.3, ensure that these tunable values are set appropriately.
112205	ERROR	The internal object value "%s" was empty. Unable to complete "%s" on %s.	<p><b>Cause:</b> The SAP instance was not provided to the SAPRemExec routine on the given system.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112206	ERROR	The required action parameter was not provided. Unable to complete "%s" on %s.	<p><b>Cause:</b> No action was provided to the SAPRemExec routine on the given system.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112208	ERROR	Error getting the value of "%s" or "%s" from the default profile "%s" on %s. Verify that the specified value exists.	<p><b>Cause:</b> Unable to obtain the virtual IP or host name from the given profile on the given server.</p> <p><b>Action:</b> Verify that the appropriate entry exists in the profile.</p>
112209	FATAL	Unable to gather required information from the SAP default profile for SID \$sid (\$DPFL) on \$me. Verify that the default profile exists and is accessible.	<p><b>Cause:</b> Unable to obtain information about the SAP instance from the given profile on the given server.</p> <p><b>Action:</b> Verify that the SAP software is properly installed and that the given profile exists and is read-enabled.</p>
112214	ERROR	Unable to determine the status of the path "%s" ("%s") on %s. The path on %s may require the execution of the command: "mount <ip>:<export> %s". Verify that the SAP software is correctly installed and	<p><b>Cause:</b> The status of the file system on the given path could not be determined.</p>

Code	Severity	Message	Cause/Action
		that all SAP file systems are mounted and accessible.	<b>Action:</b> Verify that the SAP software is properly installed and that all necessary file systems are mounted.
112219	ERROR	[HACONNECTOR:%s] Unable to write to file "%s" on %s. If the file already exists, manually enable write permissions on it.	<b>Cause:</b> The given file does not have read/write permissions enabled on the given server.  <b>Action:</b> Enable read/write permissions on the given file.
112220	ERROR	Unable to start the sapstartsrv service for SID %s and SAP instance %s on %s. Verify that the sapservices file is correct and the process can be started manually.	<b>Cause:</b> Unable to start the SAP Start Service (sapstartsrv) process for the given SAP instance.  <b>Action:</b> Verify that the sapservices file contains the appropriate command to start the sapstartsrv process and that the process can be started manually.
112221	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, or TYPE_NEW_ERS (1, 2, or 5).	<b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.  <b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.
112222	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, or TYPE_NEW_ERS (1, 2, or 5).	<b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.  <b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.
112223	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_CS (1).	<b>Cause:</b> The given routine was called for an internal SAP object with an

Code	Severity	Message	Cause/Action
			<p>unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112224	ERROR	<p>The internal \"%s\" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_CS (1).</p>	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112225	ERROR	<p>The internal \"%s\" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_NEW_ERS (5).</p>	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112226	ERROR	<p>The internal \"%s\" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance types TYPE_CS, TYPE_ERS, and TYPE_NEW_ERS (1, 2, and 5).</p>	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112227	ERROR	<p>The internal \"%s\" routine was called for a resource with unsupported instance type %s. This routine only supports SAP instance type TYPE_CS (1).</p>	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>

Code	Severity	Message	Cause/Action
112228	ERROR	The internal "%s" routine was called for a resource with unsupported instance type %s. This method supports only SAP instance type TYPE_NEW_ERS (5).	<p><b>Cause:</b> The given routine was called for an internal SAP object with an unsupported instance type.</p> <p><b>Action:</b> This is an internal error. Please submit an issue report to SIOS customer support.</p>
112229	ERROR	The profile \"%s\" either does not exist or cannot be read on %s. Unable to determine whether enqueue replication is enabled for resource %s. Please verify that the file exists and can be read.	<p><b>Cause:</b> The given instance profile either does not exist or cannot be read on the given server.</p> <p><b>Action:</b> Verify that the file exists and is read-enabled</p>
112325	ERROR	The \$SAP_CONTROL utility cannot be located. Verify that all necessary file systems are mounted and that the \$SAP_CONTROL utility can be located in the \$sapadmin user's PATH.	<p><b>Cause:</b> Unable to locate the sapcontrol utility required for SAP instance administration.</p> <p><b>Action:</b> Verify that all necessary file systems are mounted and that the sapcontrol utility can be located in the SAP administrative user's PATH.</p>
112326	ERROR	The \"which \$SAP_CONTROL\" command for user \$sapadmin returned \$sapcmd as the location of the \$SAP_CONTROL utility, but the utility could not be found or was not executable in this location. Verify that all necessary file systems are mounted and that the \$SAP_CONTROL utility can be located in the \$sapadmin user's PATH.	<p><b>Cause:</b> Unable to locate the sapcontrol utility required for SAP instance administration.</p> <p><b>Action:</b> Verify that all necessary file systems are mounted at that the sapcontrol utility can be located in the SAP administrative user's PATH.</p>
112433	ERROR	Unsupported SAPENQ_VERSION (\$enqversion) for resource \$tag on \$me. Unable to obtain enqueue replication status.	<p><b>Cause:</b> An unsupported value was detected for the SAPENQ_VERSION parameter for the given resource on the given server.</p> <p><b>Action:</b> Verify that SAPENQ_VERSION</p>

Code	Severity	Message	Cause/Action
			is set to a valid value (1 or 2, representing the version of the enqueue server currently in use) in the info file for the given resource.
112437	ERROR	Profile \"\$srvpf\" not found on \$me. Unable to obtain enqueue replication status for instance \$inst.	<p><b>Cause:</b> The given instance profile either does not exist or cannot be read on the given server.</p> <p><b>Action:</b> Verify that the file exists and is read-enabled.</p>
112438	ERROR	Unsupported SAPENQ_VERSION (\$enqversion) for resource \$tag on \$me. Unable to obtain enqueue replication status for instance \$inst.	<p><b>Cause:</b> An unsupported value was detected for the SAPENQ_VERSION parameter for the given resource on the given server.</p> <p><b>Action:</b> Verify that SAPENQ_VERSION is set to a valid value (1 or 2, representing the version of the enqueue server currently in use) in the info file for the given resource.</p>
112470	ERROR	The instance profile %s could not be found on %s. Verify that the SAP software is installed correctly and that all necessary file systems are mounted and accessible.	<p><b>Cause:</b> The given instance profile either does not exist or cannot be read on the given server.</p> <p><b>Action:</b> Verify that the file exists and is read-enabled.</p>
112490	ERROR	[HACONNECTOR] Unable to determine the corresponding tag for resource with ID \"\$res\" on \$me.	<p><b>Cause:</b> Unable to find a LifeKeeper SAP resource with the given resource ID on the given server.</p> <p><b>Action:</b> Verify that the resource ID provided as the —res argument of the fra command corresponds to a valid LifeKeeper SAP resource.</p>

Code	Severity	Message	Cause/Action
112507	ERROR	[HACONNECTOR] At least one required process for the instance was not killed successfully during the fra migrate action on \$me. Aborting resource migration.	<p><b>Cause:</b> The SAP instance was not successfully stopped on the given server while attempting a fra migrate action.</p> <p><b>Action:</b> Manually kill any processes still running for the SAP instance and reattempt the migrate action.</p>
112539	ERROR	[HACONNECTOR] Unable to find rpm information for one or more packages on \$me.	<p><b>Cause:</b> The HA Connector gvi ("Get Version Information") routine was unable to determine the current version number of LifeKeeper and/or the SAP Recovery Kit.</p> <p><b>Action:</b> Verify that the LifeKeeper Core and SAP Recovery Kit rpm information can be obtained with the rpm -q command.</p>
112976	ERROR	There is no LifeKeeper protected resource with tag \$tag on system \$me.	<p><b>Cause:</b> The resource tag provided to the SAP canfailover script does not correspond to any existing LifeKeeper resource.</p> <p><b>Action:</b> Verify that the resource tag name is correct and execute the command again.</p>
112977	ERROR	Resource \$tag is not a \$app/\$typ resource. Please use the \$ins_app/\$ins_typ resource-specific canfailover script instead.	<p><b>Cause:</b> The resource provided to the SAP canfailover script is not an appsuite/sap resource.</p> <p><b>Action:</b> Use the appropriate type-specific canfailover script for the given resource.</p>

## 8.1.16. SAP HANA Recovery Kit Message Catalog

Use **Control F** to search for a specific error code in each catalog. To search for any error code, select the **Search** button at the top right of the screen.

Code	Severity	Message	Cause/Action
136002	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided for the HANA create script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: &lt;Resource Tag&gt; &lt;SAP S/4HANA Instance&gt; [Switchback Type] [Virtual Instance Tag]</p>
136003	ERROR	END failed create of resource \$tag on server \$me with return value of \$errcode.	<p><b>Cause:</b> Failure during SAP HANA resource creation.</p> <p><b>Action:</b> Verify that SAP HANA System Replication is fully configured and enabled on the primary and secondary replication sites and retry the resource creation operation.</p>
136005	ERROR	An unknown error has occurred in utility rlslocks on server \$me. View the LifeKeeper logs for details and retry the operation.	<p><b>Cause:</b> Failure of the LifeKeeper resource manager during SAP HANA resource creation.</p> <p><b>Action:</b> Resolve any issues found in the LifeKeeper log file and reattempt the resource creation operation.</p>
136006	ERROR	END failed create of resource \$tag on server \$me with signal \$sig.	<p><b>Cause:</b> Failure during SAP HANA resource creation due to a signal.</p> <p><b>Action:</b> Review the LifeKeeper logs for details and retry the operation.</p>
136008	ERROR	An unknown error has occurred in utility getlocks on server \$me. View the LifeKeeper logs for details and retry the operation.	<p><b>Cause:</b> Failure of the LifeKeeper resource manager during SAP HANA resource creation.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Resolve any issues found in the log file and reattempt the resource creation operation.</p>
136009	ERROR	The SAP HANA product was not found in the directory \$obj->{'hana_util_path'} on server \$me. Verify that SAP HANA is correctly installed and that all necessary file systems are mounted.	<p><b>Cause:</b> Required SAP HANA binaries are not located during resource creation.</p> <p><b>Action:</b> Verify that SAP HANA is correctly installed and configured on all servers in the cluster.</p>
136010	ERROR	Failed to create resource as id \$id already exists on system \$me.	<p><b>Cause:</b> A LifeKeeper resource with the same ID already exists on the system.</p> <p><b>Action:</b> Check whether the SAP HANA resource is already protected by LifeKeeper.</p>
136011	ERROR	Failed to create new tag \$tag for SAP HANA resource on \$me.	<p><b>Cause:</b> The provided SAP HANA resource is already in use by another LifeKeeper resource or the LifeKeeper newtag utility failed to create the tag.</p> <p><b>Action:</b> Choose a different tag name for the SAP HANA resource.</p>
136012	ERROR	Failed creation of resource with \$tag on system \$me.	<p><b>Cause:</b> Failed to create the given resource in LifeKeeper.</p> <p><b>Action:</b> Resolve any issues found in the log file and reattempt the resource creation operation.</p>
136014	ERROR	Failed to create resource dependency for parent \$tag and child \$virtual_ip_tag.	<p><b>Cause:</b> Failed to create a dependency between the SAP HANA resource and its dependent resource.</p> <p><b>Action:</b> Resolve any issues found in the log file and reattempt the resource creation operation.</p>

Code	Severity	Message	Cause/Action
136015	ERROR	The info field for resource \$tag could not be successfully generated using values [SID: \$info_sid, Instance: \$info_instance, Replication Mode: \$info_rep_mode, Site Name: \$info_site_name, Operation Mode: \$info_oper_mode]. If using SAP HANA System Replication, please verify that it is fully configured and enabled on both the primary and secondary systems before creating the SAP HANA resource.	<p><b>Cause:</b> An invalid value was found in the SAP HANA resource info field.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured and that SAP HANA System Replication is fully configured and enabled on both servers in the cluster.</p>
136016	ERROR	The selected server \$me is not the primary/source system for SAP HANA System Replication for the selected SID \$sid and HDB instance \$instance. Please select 'Cancel' and start this action on the primary/source HANA System Replication system.	<p><b>Cause:</b> Resource creation is initiated on a secondary system in HANA System Replication.</p> <p><b>Action:</b> Initiate the create action on the primary system in HANA System Replication.</p>
136031	ERROR	END failed extend of resource \$target_tag on server \$me with return value of \$err_code.	<p><b>Cause:</b> Failure during SAP HANA resource extension.</p> <p><b>Action:</b> Resolve any issues found in the log file and reattempt the resource creation operation.</p>
136033	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the SAP HANA extend script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: &lt;Template System&gt; &lt;Template Instance&gt; &lt;Switchback Type&gt; &lt;Target Tag&gt;</p>
136035	ERROR	Template resource \$template_tag on server \$template_sys does not exist.	<p><b>Cause:</b> The template SAP HANA resource being extended does not exist on the target server.</p> <p><b>Action:</b> Verify that the SAP HANA resource being extended exists on the target server.</p>
136036	ERROR	Resource with matching id \$target_id already exists on server \$me for App \$app_type and Type \$res_type.	<p><b>Cause:</b> An SAP HANA resource with the same LifeKeeper ID already exists on the target server.</p> <p><b>Action:</b> Check whether the SAP HANA resource with the same LifeKeeper ID already exists on the target server.</p>

Code	Severity	Message	Cause/Action
			already protected by LifeKeeper on server.
136037	ERROR	Resource with matching tag \$target_tag already exists on server \$me for App \$app_type and Type \$res_type	<p><b>Cause:</b> An SAP HANA resource with LifeKeeper resource tag already exists on target server.</p> <p><b>Action:</b> Check whether the SAP HANA resource already protected by LifeKeeper on server.</p>
136039	ERROR	Error creating resource \$target_tag on system \$me.	<p><b>Cause:</b> Failed to create an equivalent resource on the target server.</p> <p><b>Action:</b> Resolve any issues found in log file and reattempt the resource creation operation.</p>
136040	ERROR	The target tag (\$target_tag) and template tag (\$template_tag) must be the same.	<p><b>Cause:</b> Resource tag name used on primary system is different than resource tag name on secondary system. Both must be same.</p> <p><b>Action:</b> While resource creation on primary system tag and secondary system tag must be same.</p>
136045	ERROR	Cannot extend resource \$template_tag to server \$me.	<p><b>Cause:</b> The SAP HANA canextend script on target system that the resource cannot be extended to target system.</p> <p><b>Action:</b> Resolve any issues found in log file and reattempt the resource creation operation.</p>
136047	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to HANA canextend script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: &lt;Template System&gt; &lt;Target System&gt;</p>

Code	Severity	Message	Cause/Action
136048	ERROR	Resource \$template_tag does not exist on server \$template_sys.	<p><b>Cause:</b> The template SAP HANA extended does not exist on the template server.</p> <p><b>Action:</b> Verify that the SAP HANA instance being extended exists on the template server.</p>
136049	ERROR	The system user \$hana_user does not exist on server \$me.	<p><b>Cause:</b> The SAP Administrative User on the target server HANA database does not exist on the target server.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured.</p>
136050	ERROR	The user id for user \$hana_user (\$template_uid) on template server \$template_sys is not the same as user id (\$uid) on target server \$me.	<p><b>Cause:</b> The user ID for the SAP Administrative User differs between the template and target servers.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>
136051	ERROR	The group id for user \$hana_user (\$template_gid) on template server \$template_sys is not the same as group id (\$gid) on target server \$me.	<p><b>Cause:</b> The group ID for the SAP Administrative User differs between the template and target servers.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>
136052	ERROR	The home directory for user \$hana_user (\$template_home) on template server \$template_sys is not the same as home directory (\$user_home) on target server \$me.	<p><b>Cause:</b> The home directory for the SAP Administrative User differs between the template and target servers.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>
136053	ERROR	The SAP HANA instance \$instance does not exist for \$sid on server \$me.	<p><b>Cause:</b> Installation directories for the SAP HANA database instance could not be found on the target server.</p> <p><b>Action:</b> Verify that SAP HANA is properly installed and configured on all servers in the cluster.</p>

Code	Severity	Message	Cause/Action
136055	ERROR	The SAP HANA site name \$target_obj->{'site_name'} on server \$me must be different from site name \$template_obj->{'site_name'} on \$template_obj->{'sys'}.	<p><b>Cause:</b> The SAP HANA System Replication site name is the same on both the primary and secondary servers.</p> <p><b>Action:</b> Stop the SAP HANA data replication on the secondary server and use the hdbcli to register the secondary replication with a different site name.</p>
136056	ERROR	Unable to obtain SAP HANA System Replication parameters for database \$instance on server \$me. Please verify that SAP HANA System Replication is enabled and properly configured and that the database instance is running on all servers in the cluster.	<p><b>Cause:</b> SAP HANA System Replication parameters could not be determined for the database instance on the given system.</p> <p><b>Action:</b> Verify that SAP HANA System Replication is enabled and properly configured and that the database instance is running on all servers in the cluster.</p>
136057	ERROR	Unable to create a HANA object for database \$instance on server \$me. Please verify that database instance \$instance is properly installed.	<p><b>Cause:</b> Unable to create an internal object representing the given instance on the server.</p> <p><b>Action:</b> Verify that the database instance is properly installed and that all necessary file systems are mounted.</p>
136083	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the SIOS script.</p> <p><b>Action:</b> Provide both a valid HANA resource tag name and resource ID. -t &lt;tag&gt; -i &lt;id&gt; [-U]</p>
136096	ERROR	One of the required parameters (server, tag, or action) was missing. Unable to set the local recovery policy.	<p><b>Cause:</b> One of the required parameters (server, tag, or action) was not provided.</p> <p><b>Action:</b> Provide the required parameters and reattempt the operation.</p>

Code	Severity	Message	Cause/Action
136097	ERROR	Failed to \$flg_action local recovery for SAP HANA resource \$tag on server \$node.	<p><b>Cause:</b> Failed to enable or disable for the given SAP HANA resource server.</p> <p><b>Action:</b> Verify that LifeKeeper is initialized, then reattempt the operation.</p>
136099	ERROR	Failed start of SAP Host Agent on server \$sys.	<p><b>Cause:</b> Failed to start the SAP Host processes (i.e., saphostexec, sap on given server.</p> <p><b>Action:</b> Check SAP Host Agent logs and correct any issues found, then reattempt the operation.</p>
136100	ERROR	Failed start of SAP Host Agent on server \$sys.	<p><b>Cause:</b> Failed to start SAP Host Agent processes (i.e., saphostexec, saposcol) on the given server.</p> <p><b>Action:</b> Inspect the SAP Host Agent logs and correct any issues found, then reattempt the operation.</p>
136160	ERROR	Unable to create SAP HANA object. The SID and instance for the SAP HANA database must be provided.	<p><b>Cause:</b> Either the SAP SID or the instance name were missing while trying to create an SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper logs for details.</p>
136161	ERROR	Unable to create SAP HANA object. Resource system name and tag name must be provided.	<p><b>Cause:</b> Either the system name or tag name were missing while trying to create an SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper logs for details.</p>
136162	ERROR	Could not find any information regarding resource \$tag on \$sys.	<p><b>Cause:</b> Failed to obtain information regarding the SAP HANA resource with the given tag name.</p>

Code	Severity	Message	Cause/Action
			<b>Action:</b> Verify that a SAP HANA <code>node</code> tag exists on the given server.
136168	ERROR	Unable to check status of SAP Host Agent on server <code>\$self-&gt;{'sys'}</code> . Command <code>"\$curr_cmd"</code> returned exit code <code>\$ret</code> .	<p><b>Cause:</b> Failed to determine the status of SAP Host Agent processes on the given server.</p> <p><b>Action:</b> Inspect the SAP Host Agent logs (e.g., <code>dev_saphostexec</code>) for more details.</p>
136174	ERROR	Unable to create SAP HANA object. The SID or instance name value is missing.	<p><b>Cause:</b> Either the SAP SID, the SAP instance name, or one of the SAP System Replication values were missing when trying to create an SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper logs for more details.</p>
136182	ERROR	Failed to <code>\$flg_action</code> flag <code>"\${HANA_FLAG_DATA_OUT_OF_SYNC}_\${seqv_tag}\$sys"</code> on server <code>\$sys</code> .	<p><b>Cause:</b> Failed to create or remove the <code>!HANA_DATA_OUT_OF_SYNC_&lt;</code> flag on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper logs for more details.</p>
136190	ERROR	Failed to start SAP Host Agent processes on server <code>\$self-&gt;{'sys'}</code> . Command <code>"\$hostagent_cmd"</code> returned <code>\$ret</code> .	<p><b>Cause:</b> Failed to start the SAP Host Agent processes on the given server.</p> <p><b>Action:</b> Inspect the SAP Host Agent logs (e.g., <code>dev_saphostexec</code>) for more details.</p>
136191	ERROR	Failed to start SAP OS Collector process on server <code>\$self-&gt;{'sys'}</code> . Command <code>"\$socol_cmd"</code> returned <code>\$ret</code> .	<p><b>Cause:</b> Failed to start the SAP OS Collector process on the given server.</p> <p><b>Action:</b> Inspect the SAP OS Collector logs (e.g., <code>dev_coll</code>) for more details.</p>
136193	ERROR	<code>\$takeover_text</code> of SAP HANA System Replication for SAP HANA database <code>\$self-&gt;{'instance'}</code> failed on server <code>\$node</code>	<b>Cause:</b> Failed to register the given

Code	Severity	Message	Cause/Action
		with exit code \$ret.	<p>primary master for the given database on the given server.                      HANA System Replication.</p> <p><b>Action:</b> Inspect the SAP HANA trace file <code>nameserver_&lt;hostname&gt;.xxxxx.xxxxxx</code> for more details.</p>
136197	ERROR	Update of resource info field for <code>\$seqv_tag{\$sys}</code> on <code>\$sys</code> failed with exit code <code>\$setinfo_ret</code> . Current info: <code>[\$info]</code> . Attempted new info: <code>[\$new_info]</code> .	<p><b>Cause:</b> Failed to update the info field for the given resource on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136202	EMERG	Failed to disable Autostart for SAP HANA instance <code>\$self-&gt;{'instance'}</code> on server <code>\$sys</code> with exit code <code>\$remexec_ret</code> . Please manually set <code>"Autostart = 0"</code> in the instance profile <code>\$profile</code> on <code>\$sys</code> .	<p><b>Cause:</b> The value of the Autostart parameter for the given HDB instance cannot be modified in the given HDB instance profile on the given server.</p> <p><b>Action:</b> Edit the HDB instance profile and set <code>"Autostart = 0"</code>.</p>
136205	ERROR	Failed start of SAP Start Service for SAP HANA database <code>\$instance</code> on server <code>\$sys</code> .	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database.</p> <p><b>Action:</b> Inspect the SAP Start Service log file (e.g., <code>sapstartsrv.log</code>) for more details.</p>
136208	ERROR	LifeKeeper was unable to determine the SAP HANA System Replication mode for database <code>\$rem_obj-&gt;{'instance'}</code> on server <code>\$rem_obj-&gt;{'sys'}</code> while attempting to identify the previous primary replication site. Please resolve the issue and bring the SAP HANA resource in-service on the system where the database should be registered as primary master.	<p><b>Cause:</b> Failed to determine the SAP HANA System Replication mode on the given server while attempting to identify the previous primary replication site.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136210	ERROR	Failed start of SAP Start Service for SAP HANA database <code>\$rem_obj-&gt;{'instance'}</code> on server <code>\$rem_obj-&gt;{'sys'}</code> . Unable to stop the database on the server where it is currently registered as primary master.	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database on the given server. As a result, the database could not be brought in-service on the current primary SAP HANA System Replication site.</p>

Code	Severity	Message	Cause/Action
			<p>site.</p> <p><b>Action:</b> Inspect the SAP Start Server log file (e.g., sapstartsrv.log) for more details.</p>
136212	ERROR	Failed stop of SAP HANA database \$rem_obj->{'instance'} on server \$rem_obj->{'sys'} where it is currently registered as primary master.	<p><b>Cause:</b> Failed to stop the given SAP HANA database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA transaction log file and LifeKeeper log file for more details.</p>
136217	ERROR	Failed start of SAP HANA database \$instance on server \$sys.	<p><b>Cause:</b> Failed to start the given SAP HANA database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA transaction log file and LifeKeeper log file for more details.</p>
136220	ERROR	Unable to register \$sys as a secondary SAP HANA System Replication site for database \$instance. The host name of the current primary replication site was not provided.	<p><b>Cause:</b> The host name of the current primary SAP HANA System Replication site was not provided when attempting to register a secondary replication site.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136233	EMERG	WARNING: A temporary communication failure has occurred between servers \$self->{'sys'} and \$sys. Manual intervention is required in order to minimize the risk of data loss. To resolve this situation, please take one of the following resource hierarchies out of service: \$self->{'tag'} on \$self->{'sys'} or \$seqv{\$sys} on \$sys. The server that the resource hierarchy is taken out of service on will become the standby server for SAP HANA database \$self->{'instance'}.	<p><b>Cause:</b> A temporary communication failure between the given equivalent HANA System Replication sites caused the given equivalent HANA System Replication sites to be brought in-service at the same time on both respective host servers.</p> <p><b>Action:</b> Take the entire HANA System Replication site out of service on the server which is currently the secondary replication site. Once the primary replication site has been stopped on that server, the secondary replication site will automatically register it as a secondary replication site during the next quickCheck cycle.</p>
136234	EMERG	WARNING: SAP HANA database \$self->{'instance'} is running and registered as primary master on the following servers: \$primary_node_list. Manual intervention is required	<p><b>Cause:</b> The given SAP HANA database is running and registered as primary master on the following servers: \$primary_node_list.</p>

Code	Severity	Message	Cause/Action
		<p>in order to minimize the risk of data loss. To resolve this situation, please stop database \$self-&gt;{'instance'} on the standby server by running the command <code>\su – \$self-&gt;{'sid_admin'} -c \"\\$SAP_CONTROL -nr \$self-&gt;{'instance_number'} -function StopWait \$HANA_STOP_WAIT 5\"</code> on that server, allow LifeKeeper to register the standby server as a secondary replication site, then use LifeKeeper to bring resource \$self-&gt;{'tag'} in-service on the intended primary replication site.</p>	<p>and registered as primary master servers concurrently.</p> <p><b>Action:</b> Use the command provided to stop the database on the standby server. Once the database is stopped, LifeKeeper will automatically register the standby server as a secondary replication site.</p>
136236	ERROR	<p>Failed to remove <code>!HANA_FLAG_DATA_OUT_OF_SYNC</code> flag on server \$sys.</p>	<p><b>Cause:</b> Failed to remove the <code>!HANA_FLAG_DATA_OUT_OF_SYNC</code> LifeKeeper flag on the given server. To bring the given SAP HANA resource to service on the given server until it is in-service.</p> <p><b>Action:</b> Verify that LifeKeeper is properly initialized. If it can be verified that LifeKeeper System Replication is in-sync, the flag can be removed manually with the command <code>LifeKeeper/bin/flg_remove -f 'HANA_FLAG_DATA_OUT_OF_SYNC'</code>.</p>
136238	ERROR	<p>Unable to create SAP HANA object. Resource system name and tag name must be provided.</p>	<p><b>Cause:</b> Either the server name or resource name were missing while trying to create the SAP HANA object.</p> <p><b>Action:</b> Inspect the LifeKeeper log for more details.</p>
136239	ERROR	<p>Failed start of SAP Start Service for SAP HANA database \$obj-&gt;{'instance'} on server \$obj-&gt;{'sys'}. Unable to determine status of the database on \$obj-&gt;{'sys'}.</p>	<p><b>Cause:</b> Failed to start SAP Start Service for given SAP HANA database on the server. As a result, the status of the database could not be determined.</p> <p><b>Action:</b> Inspect the SAP Start Service log (e.g., sapstartsrv.log) for more details.</p>
136242	ERROR	<p>Unable to create SAP HANA object. At least one of the SAP HANA System Replication values is missing.</p>	<p><b>Cause:</b> SAP HANA System Replication values could not be determined for the database.</p>

Code	Severity	Message	Cause/Action
			<p>given system.</p> <p><b>Action:</b> Verify that SAP HANA Sy is enabled and properly configured database is running on all servers</p>
136243	ERROR	Unable to locate the pingnfs utility (\$pingnfs) on server \$me. Please verify that this utility exists and is executable.	<p><b>Cause:</b> Unable to locate the pingnfs utility on the given server, testing available of exported NFS shares.</p> <p><b>Action:</b> Verify that the pingnfs utility exists at the given location and is executable.</p>
136245	ERROR	Critical NFS shares being exported by server \$sys (\$export_list) are not currently available. Please verify that the NFS server is alive and that all NFS-related services are running.	<p><b>Cause:</b> The given critical NFS shares are not currently available.</p> <p><b>Action:</b> Verify that the NFS server on the given systems is alive and that all necessary NFS-related services are running.</p>
136248	ERROR	Unable to open file \$crit_mount_file on server \$me. Please verify that this file exists and is read-enabled.	<p><b>Cause:</b> Unable to open the file containing the export information for critical NFS shares on the given server.</p> <p><b>Action:</b> Verify that the file exists in the given location and is read-enabled.</p>
136253	ERROR	The SAP HANA \"takeover with handshake\" feature is available only for SAP HANA versions 2.0 SPS04 and greater. Database \$self->{'instance'} cannot be resumed.	<p><b>Cause:</b> Resuming a suspended database instance is not supported in SAP HANA versions prior to 2.0 SPS04.</p> <p><b>Action:</b> Upgrade to SAP HANA 2.0 SPS04 or later in order to use features related to \"Handshake\".</p>
136254	ERROR	Attempt to resume suspended primary database \$self->{'instance'} on server \$self->{'sys'} failed with exit code \$ret.	<p><b>Cause:</b> The attempt to resume the database instance on the given server failed.</p> <p><b>Action:</b> Inspect the LifeKeeper logs for more details.</p>

Code	Severity	Message	Cause/Action
			determine the cause of the failure the operation or bring the corresponding resource in-service on a different
136258	ERROR	Attempt to register server \$node as the secondary SAP HANA System Replication site for database \$self->{'instance'} failed with exit code \$ret.	<p><b>Cause:</b> The attempt to register the a secondary SAP HANA System F the given database instance failed</p> <p><b>Action:</b> Inspect the LifeKeeper and determine the cause of the failure correct any issues found. While the resource is in-service, LifeKeeper continue attempting to register the a secondary HSR role.</p>
136263	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the S resource restore script.</p> <p><b>Action:</b> Please provide appropriate the form: -t &lt;Resource Tag&gt; -i &lt;R</p>
136265	ERROR	Error getting resource information for \$tag on server \$me.	<p><b>Cause:</b> Failed to obtain information SAP HANA resource on the given</p> <p><b>Action:</b> Verify that a SAP HANA r given tag exists on the given serv</p>
136266	ERROR	The resource \$tag protecting SAP HANA database \$instance is not in sync. To protect the data LifeKeeper will not restore the resource on \$me. Please restore the resource on the previous source server to allow the resync to complete.	<p><b>Cause:</b> SAP HANA System Replic sync before attempting to bring the resource in-service on the backup</p> <p><b>Action:</b> Bring the SAP HANA reso on the previous primary server and resynchronization to complete.</p>
136275	ERROR	Failed to determine SAP HANA System Replication mode for database \$instance on server \$me.	<p><b>Cause:</b> Failed to determine the S Replication mode for the given da given server.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Inspect the SAP HANA trace file <code>nameserver_&lt;hostname&gt;.xxxxx.xxxxxx.log</code> and the LifeKeeper log file for more details.</p>
136351	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the <code>SAP_HANA_quickCheck</code> resource quickCheck script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: <code>-t &lt;Resource Tag&gt; -i &lt;Resource Instance&gt;</code>.</p>
136353	ERROR	Error getting resource information for \$tag.	<p><b>Cause:</b> Failed to obtain information for the given SAP HANA resource.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136354	EMERG	The replication mode of the SAP HANA database \$instance corresponding to resource \$tag was modified outside of LifeKeeper and is no longer registered as primary master on server \$me. Please bring the SAP HANA resource in-service on the server where the database should be registered as primary master. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System Replication mode for the given database was modified outside of LifeKeeper.</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the server where it should be registered as primary master.</p>
136363	EMERG	LifeKeeper was unable to determine the SAP HANA System Replication mode for database \$instance corresponding to resource \$tag on server \$me. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> Failed to determine the SAP HANA System Replication mode on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136375	EMERG	An NFS server exporting a critical shared file system for resource \$tag is currently unavailable. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> A critical NFS shared file system is currently unavailable.</p> <p><b>Action:</b> Verify that the NFS server is in-service and all necessary NFS-related services are running.</p>
136376	EMERG	WARNING: LifeKeeper resource \$tag is designed for use in situations where SAP HANA System Replication (HSR) is	<p><b>Cause:</b> The given LifeKeeper resource is designed for use in situations where SAP HANA System Replication (HSR) is</p>

Code	Severity	Message	Cause/Action
		disabled, but HSR was found to be enabled on server \$me. Please ensure that the correct LifeKeeper resource type is being used for your current SAP HANA configuration.	<p>to protect a SAP HANA database where HANA System Replication is disabled but HSR is currently enabled on the</p> <p><b>Action:</b> Verify that the correct LifeKeeper resource type is being used based for your current configuration. If you have migrated to a new SAP HANA configuration where HSR was disabled but HSR is currently enabled, the correct LifeKeeper resource type must be recreated.</p>
136377	EMERG	SAP HANA database \$instance corresponding to resource \$tag is currently suspended on server \$me due to actions performed outside of LifeKeeper. Please take the SAP HANA resource out of service on server \$me and bring it in-service on the server where the database should be registered as primary master. Bringing resource \$tag back in-service on \$me will resume the suspended database. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The given SAP HANA database has been suspended on the given server due to actions performed outside of LifeKeeper.</p> <p><b>Action:</b> If you would like to resume the SAP HANA database on the server where the database should be registered as primary master, bring the LifeKeeper resource in-service on the server where the LifeKeeper resource is currently in-service. To resume the database on the server where the database should be registered as primary master, execute the command given in the LifeKeeper SIOS documentation. Otherwise, bring the LifeKeeper resource in-service on the intended replication site.</p>
136450	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the LifeKeeper SIOS SAP HANA remove script.</p> <p><b>Action:</b> Please provide appropriate arguments to the script in the form: &lt;Template Tag&gt; &lt;Template Tag&gt;</p>
136454	ERROR	Error getting resource information for \$tag.	<p><b>Cause:</b> Failed to obtain information for the given SAP HANA resource.</p> <p><b>Action:</b> Verify that a SAP HANA resource exists on the given server with the given tag.</p>
136456	ERROR	Failed start of SAP Start Service for SAP HANA database \$instance on server \$me.	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database.</p> <p><b>Action:</b> Inspect the SAP Start Service logs for the given SAP HANA database.</p>

Code	Severity	Message	Cause/Action
			(e.g., sapstartsrv.log) for more details.
136459	EMERG	WARNING: LifeKeeper resource \$tag is designed for use in situations where SAP HANA System Replication (HSR) is disabled, but HSR was found to be enabled on server \$me. Please ensure that the correct LifeKeeper resource type is being used for your current SAP HANA configuration.	<p><b>Cause:</b> The given LifeKeeper resource is designed to protect a SAP HANA database where HANA System Replication is disabled, but HSR is currently enabled on the given server.</p> <p><b>Action:</b> Verify that the correct LifeKeeper resource type is being used based for your current SAP HANA configuration. If you have migrated to a new SAP HANA configuration where HSR was disabled, but one where it is enabled, the correct LifeKeeper resource must be recreated.</p>
136462	ERROR	Failed to remove flag "\\${hana::HANA_FLAG_LEAVE_DB_RUNNING}_\$tag" on server \$me. This may cause subsequent remove actions for resource \$tag on server \$me to unintentionally fail to stop the database.	<p><b>Cause:</b> Failed to remove the LifeKeeper flag !volatile!hana_leave_db_running_ on the given server.</p> <p><b>Action:</b> Manually remove the flag using the command "/opt/LifeKeeper/bin/flag_remove !volatile!hana_leave_db_running_ on the given server. While the flag exists, subsequent operations for the SAP HANA resource on the given server will leave the protected instance running.</p>
136550	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments in the HANA recover script.</p> <p><b>Action:</b> Provide both a valid HANA resource tag name and resource ID. recover -d &lt;tag&gt; -n &lt;id&gt;</p>
136555	ERROR	Error getting resource information for \$tag on server \$me.	<p><b>Cause:</b> Failed to obtain information for the given HANA resource on the given server.</p> <p><b>Action:</b> Verify that the server is online and is running, and the HANA resource exists.</p>

Code	Severity	Message	Cause/Action
136556	EMERG	The replication mode of the SAP HANA database \$instance corresponding to resource \$tag was modified outside of LifeKeeper and is no longer registered as primary master on server \$me. Please bring the SAP HANA resource in-service on the server where the database should be registered as primary master. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System Replication mode for the given database was modified outside of LifeKeeper.</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the server where it should be registered as primary master.</p>
136558	EMERG	LifeKeeper was unable to determine the SAP HANA System Replication mode for database \$instance corresponding to resource \$tag on server \$me. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System Replication mode could not be determined for the given database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA transaction log and LifeKeeper log file for more details.</p>
136559	ERROR	Resource \$tag is no longer ISP on server \$me. Exiting \$cmd for \$tag.	<p><b>Cause:</b> The given SAP HANA resource is no longer ISP on the given server.</p> <p><b>Action:</b> Inspect the LifeKeeper log file for more details.</p>
136650	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the HANA hana_stop_all_dbs script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: hana_stop_all_dbs -t &lt;tag&gt;</p>
136654	ERROR	Error getting resource information for \$tag.	<p><b>Cause:</b> Failed to obtain information for the given SAP HANA resource.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>
136658	ERROR	Failed start of SAP Start Service for SAP HANA database \$x->{'instance'} on server \$x->{'sys'}. Could not determine status of SAP HANA DB on \$x->{'sys'}.	<p><b>Cause:</b> Failed to start SAP Start Service for the given SAP HANA database.</p> <p><b>Action:</b> Inspect the SAP Start Service log file (e.g., sapstartsrv.log) for more details.</p>

Code	Severity	Message	Cause/Action
136661	ERROR	Failed stop of SAP HANA database \$x->{'instance'} on server \$x->{'sys'}.	<p><b>Cause:</b> Failed to stop the given SAP HANA database on the given server.</p> <p><b>Action:</b> Inspect the SAP HANA transaction log file for more details.</p>
136673	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the hana_hana_takeover_with_hands_off command.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: -t &lt;tag&gt; [-s &lt;target server&gt;]</p>
136674	ERROR	Unable to remove the <code>"\${hana::HANA_FLAG_LEAVE_DB_RUNNING}_\$tag"</code> flag on server \$isp_node. This may cause subsequent remove actions for resource \$tag on server \$isp_node to unexpectedly fail to stop the database.	<p><b>Cause:</b> Failed to remove the LifeKeeper flag <code>!volatile!hana_leave_db_running_\$tag</code> on the given server.</p> <p><b>Action:</b> Manually remove the flag using the command <code>"/opt/LifeKeeper/bin/flag_remove !volatile!hana_leave_db_running_\$tag \$isp_node"</code> on the given server. While the flag exists, any subsequent operations for the SAP HANA resource on the given server will leave the protected instance running.</p>
136675	ERROR	Script \$cmd exited unexpectedly due to signal "\$sig" on server \$me. This may leave the \$tag resource hierarchy as well as SAP HANA System Replication in an unexpected state. Please verify that the cluster resources are in the expected state.	<p><b>Cause:</b> The hana_takeover_with_hands_off command exited unexpectedly due to the given signal on the given server.</p> <p><b>Action:</b> Verify that the SAP HANA resources are in the expected state and resolve any issues that are found and bring the resource hierarchy in-service on the primary server.</p>
136677	ERROR	Unable to find equivalent SAP HANA resource corresponding to \$tag on server \$target_node.	<p><b>Cause:</b> Unable to find an equivalent SAP HANA resource corresponding to the given tag on the given server.</p> <p><b>Action:</b> Verify that the given resource exists on the target server.</p>

Code	Severity	Message	Cause/Action
			<p>correct, the resource has been ex given target server, and that LifeK and fully initialized on the target s</p>
136678	ERROR	<p>Unable to obtain information about equivalent SAP HANA resource \$tag on server \$target_node. Verify that LifeKeeper is running and fully initialized.</p>	<p><b>Cause:</b> Unable to find an equivalent resource corresponding to the given the given server.</p> <p><b>Action:</b> Verify that the given resource correct, the resource has been ex given target server, and that LifeK and fully initialized on the target s</p>
136679	ERROR	<p>Resource \$tag is not a SAP HANA resource.</p>	<p><b>Cause:</b> The given resource is not (database/hana) resource.</p> <p><b>Action:</b> Verify that the resource ta</p>
136680	ERROR	<p>Resource \$tag is designed for use in an environment where SAP HANA System Replication is disabled. Takeover with handshake cannot be performed for this resource type. Please use the standard "\"In Service...\" command instead.</p>	<p><b>Cause:</b> Features related to "Takeover with Handshake" may only be used in where SAP HANA System Replication</p> <p><b>Action:</b> Use the standard "In Service...\" command to bring the SAP HANA resource in-s</p>
136681	ERROR	<p>SAP HANA resource \$tag is not currently in-service on any server in the cluster. The resource must be in-service and SAP HANA System Replication must be in-sync before performing a takeover with handshake.</p>	<p><b>Cause:</b> The given SAP HANA resource is not in-service on any server in the cluster. Attempting "Takeover with Handshake" will fail.</p> <p><b>Action:</b> Bring the SAP HANA resource in-service on the intended primary server.</p>
136682	ERROR	<p>SAP HANA resource \$tag is currently in-service on multiple servers: \$isp_node_list. The resource must be in-service on only one server and SAP HANA System Replication must be in-sync before performing a takeover with handshake.</p>	<p><b>Cause:</b> The given SAP HANA resource is in-service on multiple servers in the cluster. Attempting "Takeover with Handshake" will fail.</p> <p><b>Action:</b> Take the resource out of service on all servers except the one where it is intended to be in-service.</p>

Code	Severity	Message	Cause/Action
			registered as primary master.
136684	ERROR	Unable to create internal SAP HANA object for resource \$tag on server \$target_node. Verify that all necessary file systems are mounted and that LifeKeeper is running and fully initialized on \$target_node.	<p><b>Cause:</b> Unable to create an internal object representing the given instance on the target server.</p> <p><b>Action:</b> Verify that the database instance is properly installed and that all necessary file systems are mounted.</p>
136685	ERROR	Takeover with handshake is only supported in SAP HANA versions 2.0 SPS04 and greater. The SAP HANA software must be upgraded in order to use this feature. Please use the standard "\"In Service...\"" command instead.	<p><b>Cause:</b> The "Takeover with Handshake" feature cannot be used when the underlying SAP HANA database version is less than 2.0 SPS04.</p> <p><b>Action:</b> Upgrade to SAP HANA 2.0 SPS04 or greater in order to use the "Takeover with Handshake" feature.</p>
136686	ERROR	Takeover with handshake cannot be performed for database \$target_obj->{'instance'} on server \$target_node because the database is not currently running and registered as primary on any other server in the cluster.	<p><b>Cause:</b> The given SAP HANA database instance is not running and registered as primary on any server in the cluster during an attempted "Takeover with Handshake".</p> <p><b>Action:</b> Bring the corresponding SAP HANA database resource in-service on the intended target server.</p>
136687	ERROR	SAP HANA database \$target_obj->{'instance'} is running and registered as primary on more than one server in the cluster. Please resolve this situation and reattempt the takeover.	<p><b>Cause:</b> The given SAP HANA database instance is running and registered as primary on more than one server during an attempted "Takeover with Handshake".</p> <p><b>Action:</b> If the database instance is running and registered as primary on the intended target server, bring the corresponding LifeKeeper resource in-service on that server. If the database instance is running and registered as primary on every server in the cluster, bring the database instance on every server in the cluster where it is currently in-service on-line. Once LifeKeeper resumes system replication, reattempt the takeover.</p>

Code	Severity	Message	Cause/Action
136689	ERROR	Unable to create internal SAP HANA object for resource \$tag on server \$isp_node. Verify that all necessary file systems are mounted and that LifeKeeper is running and fully initialized on \$isp_node.	<p><b>Cause:</b> Unable to create an internal object representing the given instance on the given server.</p> <p><b>Action:</b> Verify that the database is properly installed and that all necessary file systems are mounted.</p>
136690	ERROR	Unable to set the \${hana::HANA_FLAG_LEAVE_DB_RUNNING}_\$tag flag on server \$isp_node. Aborting takeover with handshake attempt for database \$target_obj->{'instance'} on server \$target_node.	<p><b>Cause:</b> Failed to set the !volatile!hana_leave_db_running_\$tag flag on the given server during an "Takeover with Handshake".</p> <p><b>Action:</b> Inspect the LifeKeeper log for error information. Correct any issues found and then reattempt the takeover.</p>
136692	ERROR	Takeover with handshake for database \$target_obj->{'instance'} failed on server \$target_node.	<p><b>Cause:</b> The "Takeover with Handshake" failed for the given SAP HANA database on the given target server.</p> <p><b>Action:</b> If HANA_HANDSHAKE_TAKEOVER is set to true in /etc/default/LifeKeeper, the SAP HANA service hierarchy will automatically be brought in-service on the previous database server. Otherwise, the SAP HANA resource must be manually brought in-service on the primary server.</p>
136695	ERROR	Resource \$tag does not exist on server \$me.	<p><b>Cause:</b> The given resource does not exist on the given server.</p> <p><b>Action:</b> Verify that the resource tag is correct.</p>
136696	ERROR	Unable to verify the status of resource \$tag on server \$target_node. Assuming that it is not in-service.	<p><b>Cause:</b> Failed to determine the status of the resource on the given server while performing the "Takeover with Handshake".</p> <p><b>Action:</b> Verify that the takeover was successful.</p>

Code	Severity	Message	Cause/Action
			<p><b>Action:</b> Verify that LifeKeeper is not initialized on the given server and that the communication path between the servers is active. If HANA_HANDSHAKE_TAKEOVER is set to true in /etc/default/LifeKeeper, the SAP HANA hierarchy will automatically be brought in-service on the previous database server. Otherwise, the SAP HANA resource must be manually brought in-service on the primary server.</p>
136697	ERROR	Resource \$res was not successfully brought in-service on server \$target_node.	<p><b>Cause:</b> The given resource failed to be brought in-service on the given server during "Takeover with Handshake".</p> <p><b>Action:</b> If HANA_HANDSHAKE_TAKEOVER is set to true in /etc/default/LifeKeeper, the SAP HANA hierarchy will automatically be brought in-service on the previous database server. Otherwise, the SAP HANA resource must be manually brought in-service on the primary server.</p>
136698	ERROR	LifeKeeper is not running or is not fully initialized on server \$me.	<p><b>Cause:</b> LifeKeeper is either not running or not fully initialized on the given server during "Takeover with Handshake".</p> <p><b>Action:</b> Either start LifeKeeper with /usr/sbin/LifeKeeper/bin/lkstart or allow it to start and then reattempt the takeover.</p>
136699	ERROR	Unknown server \$target_node.	<p><b>Cause:</b> The given server host name is not recognized.</p> <p><b>Action:</b> Verify that the server host name is correct and that communication paths have been established between the local server and the target server.</p>

Code	Severity	Message	Cause/Action
136700	ERROR	Usage: \$usage	<p><b>Cause:</b> Invalid arguments provided to the HANA remoteregisterdb script.</p> <p><b>Action:</b> Please provide appropriate arguments in the form: remoteregisterdb -d &lt;tag&gt;</p>
136705	ERROR	Unable to obtain information about resource \$tag on server \$me. Exiting \$cmd for \$tag.	<p><b>Cause:</b> Failed to determine information about the given SAP HANA resource on the given server.</p> <p><b>Action:</b> Verify that the given resource tag is correct, the resource exists on the given server, the necessary file systems are mounted, and the LifeKeeper is running and fully initialized on the given server.</p>
136706	ERROR	Resource \$tag is no longer ISP on server \$me. Exiting \$cmd for \$tag.	<p><b>Cause:</b> The SAP HANA remoteregisterdb script is exiting because the SAP HANA resource is no longer in service (ISP) on the given server.</p> <p><b>Action:</b> No action is required.</p>
136707	ERROR	The \$cmd event is intended for use only in environments in which SAP HANA System Replication is enabled. Exiting \$cmd for \$tag.	<p><b>Cause:</b> The SAP HANA remoteregisterdb script detected that SAP HANA System Replication is disabled for the database protected by the given SAP HANA resource.</p> <p><b>Action:</b> Verify that SAP HANA System Replication is enabled on the server where the remoteregisterdb script was running. System Replication may be enabled by running the command 'hdbnsutil -sr_enable sap://Name&gt;' as the SAP HANA administrator.</p>
136708	ERROR	Error getting resource information for \$tag on server \$me. Exiting \$cmd for \$tag.	<p><b>Cause:</b> Failed to obtain information about the given SAP HANA resource on the given server.</p> <p><b>Action:</b> Verify that a SAP HANA resource with the given tag exists on the given server.</p>

Code	Severity	Message	Cause/Action
136709	EMERG	LifeKeeper was unable to determine the SAP HANA System Replication mode for database \$instance corresponding to resource \$tag on server \$me. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System R could not be determined for the given server.</p> <p><b>Action:</b> Inspect the SAP HANA tra LifeKeeper log file for more details.</p>
136710	EMERG	The replication mode of the SAP HANA database \$instance corresponding to resource \$tag was modified outside of LifeKeeper and is no longer registered as primary master on server \$me. Please bring the SAP HANA resource in-service on the server where the database should be registered as primary master. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The SAP HANA System R for the given database was modified LifeKeeper.</p> <p><b>Action:</b> Bring the SAP HANA resource on the server where it should be registered as primary master.</p>
136711	EMERG	SAP HANA database \$instance corresponding to resource \$tag is currently suspended on server \$me due to actions performed outside of LifeKeeper. Please take the SAP HANA resource out of service on server \$me and bring it in-service on the server where the database should be registered as primary master. Bringing resource \$tag back in-service on \$me will resume the suspended database. Resource monitoring for \$tag will be suspended until the issue is resolved.	<p><b>Cause:</b> The given SAP HANA database has been suspended on the given server due to actions performed outside of LifeKeeper.</p> <p><b>Action:</b> If you would like to resume the database on the server where the LifeKeeper resource is currently in-service, execute the command given in the log file. Otherwise, bring the LifeKeeper S resource in-service on the intended replication site.</p>

# 9. LifeKeeper for Linux Support Matrix

Supported Operating Systems

[Supported Applications](#)

[Supported Virtualization](#)

## Supported Operating Systems

Product	Supported Operating System	LifeKeeper for Linux				
		v9.4.0	v9.4.1	v9.5.0	v9.5.1	v9.5.2
LifeKeeper for Linux	Red Hat Enterprise Linux 8  <b>Note:</b> <a href="#">A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</a>	8.0 64-Bit (Upgrading from one kernel version to another major version such as from RHEL7 to RHEL8 is NOT supported.) (Disks with an odd sector size cannot be used for DataKeeper.)	8.0 64-Bit (Upgrading from one kernel version to another major version such as from RHEL7 to RHEL8 is NOT supported.) (Disks with an odd sector size cannot be used for DataKeeper.)	8.0, 8.1 64-Bit (Upgrading from one kernel version to another major version such as from RHEL7 to RHEL8 is NOT supported.) (Disks with an odd sector size cannot be used for DataKeeper.)	8.0 to 8.3 64-Bit  <b>IMPORTANT NOTICE:</b> As our valued customer, we want to proactively notify you of an issue that we discovered in the md/raid1 kernel module of several Linux releases. Due to this issue, a partial resync of data within supported versions (v9.3.2 – v9.5.1) of LifeKeeper for Linux with DataKeeper cluster nodes may NOT resync all blocks. Refer to <a href="#">this link</a> for more	8.0 to 8.4 64-Bit (Upgrading from one kernel version to another major version such as from RHEL7 to RHEL8 is NOT supported.) (Disks with an odd sector size cannot be used for DataKeeper.)

					<p>information.</p> <p>(Upgrading from one kernel version to another major version such as from RHEL7 to RHEL8 is NOT supported.)</p> <p>(<a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a>)</p> <p><b>Note:</b> If you are using DataKeeper with Red Hat Enterprise Linux 8.3, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.</p>	
Red Hat Enterprise Linux 7	<p>7.0 to 7.7 64-Bit</p> <p><b>Note:</b> If you are using RHEL 7.7 with DataKeeper login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper. (Some kernel versions do not support asynchronous mode.)</p>	<p>7.0 to 7.7 64-Bit</p> <p>(<a href="#">Some kernel versions do not support asynchronous mode.</a>)</p>	<p>7.0 to 7.8 64-Bit</p> <p><b>Note:</b> If you are using RHEL 7.8 with DataKeeper, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper. (Some kernel versions do not support asynchronous mode.)</p>	<p>7.0 to 7.9 64-Bit</p> <p><b>Note:</b> If you are using RHEL 7.9 with DataKeeper, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper. (Some kernel versions do not support asynchronous mode.)</p>	<p>7.0 to 7.9 64-Bit</p> <p>(<a href="#">Some kernel versions do not support asynchronous mode.</a>)</p>	
Red Hat	6.0 to 6.10	6.0 to 6.10	6.0 to 6.10	6.0 to 6.10		

	Enterprise Linux 6	64-Bit ( <a href="#">6.0 NOT Recommended</a> )	64-Bit ( <a href="#">6.0 NOT Recommended</a> )	64-Bit ( <a href="#">6.0 NOT Recommended</a> )	64-Bit ( <a href="#">6.0 NOT Recommended</a> )	
	<p>SUSE Linux Enterprise Server (SLES) 15(*4)</p> <p><b>Note:</b> <a href="#">A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</a></p>	<p>15.0 64-Bit (<a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a>) (Upgrading from one kernel version to another major version such as from SLES 12 to SLES 15 is NOT supported.)</p>	<p>15.0 SP1 64-Bit (<a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a>) (Upgrading from one kernel version to another major version such as from SLES 12 to SLES 15 is NOT supported.)</p>	<p>15.0 SP1 64-Bit (<a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a>) (Upgrading from one kernel version to another major version such as from SLES 12 to SLES 15 is NOT supported.)</p>	<p>15.0 SP1 to SP2 64-Bit</p> <p><b>IMPORTANT NOTICE:</b> As our valued customer, we want to proactively notify you of an issue that we discovered in the md/raid1 kernel module of several Linux releases. Due to this issue, a partial resync of data within supported versions (v9.3.2 – v9.5.1) of LifeKeeper for Linux with DataKeeper cluster nodes may NOT resync all blocks. Refer to <a href="#">this link</a> for more information. (<a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a>) (Upgrading from one</p>	<p>15.0 SP1 to SP2 64-Bit (<a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a>) (Upgrading from one kernel version to another major version such as from SLES 12 to SLES 15 is NOT supported.)</p>

					kernel version to another major version such as from SLES 12 to SLES 15 is NOT supported.)	
	<p>SUSE Linux Enterprise Server (SLES) 12(*2)</p> <p><b>Note:</b> <a href="#">A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</a></p>	12SP1 to SP4 64-Bit	12SP1 to SP4 64-Bit	12SP1 to SP5 64-Bit	<p>12SP1, SP2, SP3, SP4*, SP5* 64-Bit</p> <p><b>IMPORTANT NOTICE:</b> As our valued customer, we want to proactively notify you of an issue that we discovered in the md/raid1 kernel module of several Linux releases. Due to this issue, a partial resync of data within supported versions (v9.3.2 – v9.5.1) of LifeKeeper for Linux with DataKeeper cluster nodes may NOT resync all blocks. Refer to <a href="#">this link</a> for more</p>	12SP1 to SP5 64-Bit

					information.	
	SUSE Linux Enterprise Server (SLES) 11(*1)	11.0 to SP4 64-Bit	11.0 to SP4 64-Bit			
	Oracle Linux 8  <b>Note:</b> <a href="#">A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</a>			8.0, 8.1  (Upgrading the Kernel from OEL 7 to OEL 8 is not supported.)	8.0 to 8.2 (UEK6 is not included)  <b>IMPORTANT NOTICE:</b> As our valued customer, we want to proactively notify you of an issue that we discovered in the md/raid1 kernel module of several Linux releases. Due to this issue, a partial resync of data within supported versions (v9.3.2 – v9.5.1) of LifeKeeper for Linux with DataKeeper cluster nodes may NOT resync all blocks. Refer to <a href="#">this link</a> for more information. (Upgrading the Kernel from OEL 7 to OEL	8.0 to 8.3 (including UEK R6)  The kernel should be updated to 5.4.17-2102.202 or higher for UEK R6.  (Upgrading the Kernel from OEL to OEL 8 is not supported.)  ( <a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a> )

					8 is not supported.)  ( <a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a> )	
	Oracle Linux 7  <b>Note:</b> <a href="#">A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</a>	7.0 to 7.6 (including UEK R3, R4 and R5) 64-Bit ( <a href="#">Some kernel versions do not support asynchronous mode.</a> )	7.0 to 7.7 (including UEK R3, R4 and R5) 64-Bit ( <a href="#">Some kernel versions do not support asynchronous mode.</a> )	7.0 to 7.8 (including UEK R3, R4 and R5) 64-Bit ( <a href="#">Some kernel versions do not support asynchronous mode.</a> )	7.0 to 7.9 (including UEK R3, R4 and R5) 64-Bit  <b>IMPORTANT NOTICE:</b> As our valued customer, we want to proactively notify you of an issue that we discovered in the md/raid1 kernel module of several Linux releases. Due to this issue, a partial resync of data within supported versions (v9.3.2 – v9.5.1) of LifeKeeper for Linux with DataKeeper cluster nodes may NOT resync all blocks. Refer to <a href="#">this link</a> for more	7.0 to 7.9 (including UEK R3, R4, R5 and R6) 64-Bit ( <a href="#">Some kernel versions do not support asynchronous mode.</a> )  The kernel should be updated to 5.4.17-2102.202 or higher for UEK R6.

					information. ( <a href="#">Some kernel versions do not support asynchronous mode.</a> )	
	Oracle Linux 6	6.3 to 6.10 (including UEK R3, R4) 64-Bit	6.3 to 6.10 (including UEK R3, R4) 64-Bit	6.3 to 6.10 (including UEK R3, R4) 64-Bit	6.3 to 6.10 (including UEK R3, R4) 64-Bit	
	CentOS 8  <b>Note:</b> <a href="#">A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</a>			8.0, 8.1  Upgrading from one kernel version to another major version such as CentOS 7 to CentOS 8 is not supported.	8.0 to 8.2  <b>IMPORTANT NOTICE:</b> As our valued customer, we want to proactively notify you of an issue that we discovered in the md/raid1 kernel module of several Linux releases. Due to this issue, a partial resync of data within supported versions (v9.3.2 – v9.5.1) of LifeKeeper for Linux with DataKeeper cluster nodes may NOT resync all blocks. Refer to <a href="#">this link</a> for more	8.0 to 8.3  Upgrading from one kernel version to another major version such as CentOS 7 to CentOS 8 is not supported.  ( <a href="#">Disks with an odd sector size cannot be used for DataKeeper.</a> )

					<p>information. Upgrading from one kernel version to another major version such as CentOS 7 to CentOS 8 is not supported.</p> <p><a href="#">(Disks with an odd sector size cannot be used for DataKeeper.)</a></p>	
CentOS 7	<p><b>Note:</b> <a href="#">A DataKeeper resource configuration where the resource is created with asynchronous mode and extended with synchronous mode is not supported.</a></p>	<p>7.0 to 7.6 64-Bit <a href="#">(Some kernel versions do not support asynchronous mode.)</a></p>	<p>7.0 to 7.7 64-Bit <a href="#">(Some kernel versions do not support asynchronous mode.)</a></p>	<p>7.0 to 7.8 64-Bit <a href="#">(Some kernel versions do not support asynchronous mode.)</a> <b>Note:</b> If you are using DataKeeper with 7.8, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.</p>	<p>7.0 to 7.9 64-Bit <a href="#">(Some kernel versions do not support asynchronous mode.)</a> <b>Note:</b> If you are using DataKeeper with 7.9, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.</p>	<p>7.0 to 7.9 64-Bit <a href="#">(Some kernel versions do not support asynchronous mode.)</a></p>
CentOS 6(*3)	<p>6.0 to 6.10 64-Bit (6.0 DataKeeper Configuration NOT Supported)</p>	<p>6.0 to 6.10 64-Bit (6.0 DataKeeper Configuration NOT Supported)</p>	<p>6.0 to 6.10 64-Bit (6.0 DataKeeper Configuration NOT Supported)</p>	<p>6.0 to 6.10 64-Bit (6.0 DataKeeper Configuration NOT Supported)</p>	<p>6.0 to 6.10 64-Bit (6.0 DataKeeper Configuration NOT Supported)</p>	
Rocky Linux 8						

\*1 The kernel should be updated to 3.0.42-0.7.3 for SLES11SP2.

\*2 The kernel should be updated to 4.4.82-6.9.1 for SLES12SP3.

\*3 CentOS 6.0 – DataKeeper configuration is NOT Supported. This limitation was missing in the past release notes of each version.

\*4 The kernel should be updated to 5.3.18-59.5 for SLES15SP3.

\*5 Rocky Linux is not supported in the Cloud. (i.e. AWS/Azure/GCP/OCI) .

\*6 LifeKeeper Single Server Protection is not supported.

\*7 SAP Recovery Kit and SAP HANA Recovery Kit are not supported.

## Supported Applications

Product	Supported Application	LifeKeeper for Linux					
		v9.4.0	v9.4.1	v9.5.0	v9.5.1	v9.5.2	
Apache ARK	Apache Web Server	2.4	2.4	2.4	2.4	2.4	2.4
SAP ARK (For SAP supported OS information see <a href="#">LifeKeeper for SAP Solution Page</a> )	SAP NetWeaver	7.0 including Enhancement Package 1-3, 7.3 including Enhancement Package 1, 7.4, 7.5	7.0 including Enhancement Package 1-3, 7.3 including Enhancement Package 1, 7.4, 7.5	7.0 including Enhancement Package 1-3, 7.3 including Enhancement Package 1, 7.4, 7.5	7.0 including Enhancement Package 1-3, 7.3 including Enhancement Package 1, 7.4, 7.5	7.0 including Enhancement Package 1-3, 7.3 including Enhancement Package 1, 7.4, 7.5	7.0 including Enhancement Package 1-3, 7.3 including Enhancement Package 1, 7.4, 7.5
	SAP NetWeaver AS for ABAP	7.51 innovation package, 7.52 innovation package					
	SAP S/4HANA Platform	1809 Platform	1809 Platform	1809, 1909 Platforms	1809, 1909, 2020 Platforms	1809, 1909, 2020 Platforms	1809, 1909, 2020 Platforms
SAP MaxDB ARK	SAP MaxDB	7.9	7.9	7.9	7.9	7.9	7.9
SAP HANA ARK  (For SAP HANA supported OS information see <a href="#">SAP HANA Supported Configurations</a> )	SAP HANA			2.0 SPS 04  Versions prior to SAP HANA 2.0 SPS04 are not supported	2.0 SPS 04, 05  Versions prior to SAP HANA 2.0 SPS04 are not supported	2.0 SPS 04, 05  Versions prior to SAP HANA 2.0 SPS04 are not supported	2.0 SPS 04, 05  Versions prior to SAP HANA 2.0 SPS04 are not supported
Postfix ARK	Postfix	provided with					

		the supported Linux distributions					
<b>DB2 ARK</b>	IBM Db2 Advanced, Standard and Community Editions  <b>Note:</b> <a href="#">Advanced Enterprise Server Edition and Enterprise Server Edition</a> are now referred to as <a href="#">Db2 Advanced Edition</a>	11.5	11.5	11.5	11.5	11.5	11.5
	IBM Db2 Enterprise Server Edition (ESE) and Workgroup Server Edition (WSE)	10.5, 11.1	10.5, 11.1	10.5, 11.1	10.5, 11.1	10.5, 11.1	10.5, 11.1
	IBM Db2 Express Edition	10.5, 11.1	10.5, 11.1	10.5, 11.1	10.5, 11.1	10.5, 11.1	10.5, 11.1
	IBM Db2 Advanced Workgroup Server Edition				10.5, 11.1	10.5, 11.1	10.5, 11.1
	IBM Db2 Advanced Enterprise Server Edition				10.5, 11.1	10.5, 11.1	10.5, 11.1

<b>PostgreSQL ARK</b>	PostgreSQL	9.4, 9.5, 9.6,10,11	9.4, 9.5, 9.6,10,11	9.5, 9.6,10,11,12	9.5, 9.6,10,11,12,13	9.6,10,11,12,13	9.6, 14
	EnterpriseDB Postgres Plus Advanced Server/EDB Postgres Advanced Server	9.4, 9.5, 9.6,10,11	9.4, 9.5, 9.6,10,11	9.5, 9.6,10,11,12	9.5, 9.6,10,11,12,13	9.6,10,11,12,13	9.6, 14
	Symfoware Server Enterprise Edition	12.2, 12.3	12.2, 12.3, 12.4	12.2, 12.3, 12.4	12.2, 12.3, 12.4	12.2, 12.3, 12.4	12.2, 12.3, 12.4
	Symfoware Server Standard Edition	12.2, 12.3	12.2, 12.3, 12.4	12.2, 12.3, 12.4	12.2, 12.3, 12.4	12.2, 12.3, 12.4	12.2, 12.3, 12.4
	Symfoware Server Lite Edition	12.3	12.3	12.3	12.3	12.3	12.3
	FUJITSU Software Enterprise Postgres Advanced Edition	9.5, 10, 11	9.5, 10, 11	9.5, 10, 11	9.5, 10, 11, 12	10, 11, 12	10, 11, 12
	FUJITSU Software Enterprise Postgres Standard Edition	9.5, 9.6, 10, 11	9.5, 9.6, 10, 11	9.5, 9.6, 10, 11	9.5, 9.6, 10, 11, 12	9.6, 10, 11, 12	9.6, 10, 11, 12
	FUJITSU Software Enterprise Postgres Community Edition	10, 11	10, 11	10, 11	10, 11, 12	10, 11, 12	10, 11, 12
	PowerGres on Linux		11	11	11	11	11
	PowerGres		10	10	10	10	10

	Plus						
<b>Oracle ARK</b>	Oracle Database Enterprise Edition	11g R2, 12c, 12c R2, 18c, 19c	11g R2, 12c, 12c R2, 18c, 19c	12c, 12c R2, 18c, 19c	12c, 12c R2, 18c, 19c	12c, 12c R2, 18c, 19c	12c, 12c R2, 18c, 19c
	Oracle Database Standard Edition 2	12c, 12c R2, 18c, 19c					
	Oracle Database Standard Edition	11g R2	11g R2				
	Oracle Database Standard Edition One	11g R2	11g R2				
<b>Sybase ASE ARK</b>	SAP Adaptive Server Enterprise	15.7, 16.0	15.7, 16.0	15.7, 16.0	15.7, 16.0	15.7, 16.0	15.7, 16.0
<b>Samba ARK</b>	Standard Samba file shares	provided with the supported Linux distributions					
<b>NFS Server ARK</b>	Linux kernel version	Kernel 2.6 or later (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7 or later.) (NFS over UDP is not supported on RHEL 8 or later.)	Kernel 2.6 or later (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7 or later.) (NFS over UDP is not supported on RHEL 8 or later.)	Kernel 2.6 or later (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7 or later.) (NFS over UDP is not supported on RHEL 8 or later.)	Kernel 2.6 or later (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7 or later.) (NFS over UDP is not supported on RHEL 8 or later.)	Kernel 2.6 or later (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7 or later.) (NFS over UDP is not supported on RHEL 8 or later.)	Kernel 2.6 or later (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7 or later.) (NFS over UDP is not supported on RHEL 8 or later.)
<b>NAS ARK</b>	NFS version of Mounted NFS file systems from an NFS	NFSv2 (NFSv2 is not supported on RHEL 7/ CentOS 7/OL	NFSv2 (NFSv2 is not supported on RHEL 7/ CentOS 7/OL	NFSv2 (NFSv2 is not supported on RHEL 7/ CentOS 7/OL	NFSv2 (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7	NFSv2 (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7	NFSv2 (NFSv2 is not supported on RHEL 7/ CentOS 7/OL 7

	server or Network Attached Storage (NAS) device	7 or later.) NFSv3 NFSv4 (NFS over UDP is not supported on RHEL 8 or later.)	7 or later.) NFSv3 NFSv4 (NFS over UDP is not supported on RHEL 8 or later.)	7 or later.) NFSv3 NFSv4 (NFS over UDP is not supported on RHEL 8 or later.)	or later.) NFSv3 NFSv4 (NFS over UDP is not supported on RHEL 8 or later.)	or later.) NFSv3 NFSv4 (NFS over UDP is not supported on RHEL 8 or later.)	or NF NF (N is on lat
<b>MySQL ARK</b>	MySQL Community Edition, MySQL Enterprise Edition	5.7, 8.0	5.7, 8.0	5.7, 8.0	5.7, 8.0	5.7, 8.0	5.
	MariaDB	5.5, 10.0, 10.3 (Excluding 10.1 and 10.2)	5.5, 10.0, 10.3 (Excluding 10.1 and 10.2)	10.3, 10.4 (Excluding 10.1 and 10.2)	10.3, 10.4 (Excluding 10.1 and 10.2)	10.3, 10.4, 10.5 (Excluding 10.1 and 10.2)	10 (E ar
<b>WebSphere MQ ARK</b>	WebSphere MQ/IBM MQ	8.0, 9.0, 9.1	8.0, 9.0, 9.1	8.0, 9.0, 9.1	8.0, 9.0, 9.1, 9.2	8.0, 9.0, 9.1, 9.2	8.
<b>Software RAID (md) ARK</b>	RHEL6, OL6, CentOS6	6.0 to 6.7	6.0 to 6.7	6.0 to 6.7	6.0 to 6.7		
	SLES11	11.0 to SP1	11.0 to SP1				
<b>VMDK as Shared Storage ARK</b>	VMware vSphere		6.5, 6.7, (RHEL 6.x / CentOS 6.x / OEL 6.x and SLES11.x are not supported.) (Not supported by SSP.)  (An <a href="#">additional step</a> is required before creating resources with the VMDK ARK.)	6.5, 6.7, 7.0 (RHEL 6.x / CentOS 6.x / OEL 6.x and SLES11.x are not supported.) (Not supported by SSP.)  (An <a href="#">additional step</a> is required before creating resources with the VMDK ARK.)	6.5, 6.7, 7.0 (RHEL 6.x / CentOS 6.x / OEL 6.x and SLES11.x are not supported.) (Not supported by SSP.)  (Stopping LifeKeeper without stopping resources (lkstop -f) may halt the entire system. Refer to the following for more	6.5, 6.7, 7.0 (RHEL 6.x / CentOS 6.x / OEL 6.x and SLES11.x are not supported.) (Not supported by SSP.)  (Stopping LifeKeeper without stopping resources (lkstop -f) may halt the entire system. Refer to the following for more	6. (R Ce O SL no (N by  (S Lif wi re (lk ha sy to fo int

			(Stopping LifeKeeper without stopping resources (lkstop -f) may halt the entire system. Refer to the following for more information.)	(Stopping LifeKeeper without stopping resources (lkstop -f) may halt the entire system. Refer to the following for more information.)	information.)	information.)
--	--	--	---------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------	---------------	---------------

## Supported Virtualization

Product	Supported Virtualization	LifeKeeper for Linux					
		v9.4.0	v9.4.1	v9.5.0	v9.5.1	v9.5.2	v9.6.0
LifeKeeper for Linux	VMware vSphere	5.5, 6.0, 6.5, 6.7	5.5, 6.0, 6.5, 6.7	5.5, 6.0, 6.5, 6.7, 7.0	5.5, 6.0, 6.5, 6.7, 7.0	6.5, 6.7, 7.0	6.5, 6.7, 7.0
	KVM	provided with the supported Linux distributions					
	Oracle VM Server for x86	3 (It does not support shared storage.)					
	Nutanix Acropolis Hypervisor (AHV)	20160925.57 20160925.90	20160925.57 20160925.90	20160925.57 20160925.90	20160925.57 20160925.90	20160925.57 20160925.90	20160925.57 20160925.90

For a list of the disk array storage models and adapters currently supported by LifeKeeper in shared storage configurations as well as their type of certification, see the [Supported Storage List](#).

For version requirements for SAP, see the [LifeKeeper for SAP Solution Page](#).

## Supported Cloud Platforms

- Amazon EC2
- Microsoft Azure
- Google Cloud

- [Oracle Cloud Infrastructure](#)
- Virtustream – Dell Technologies

# 10. Supported Storage

## Supported Storage List for LifeKeeper for Linux v9

Last updated: October 27, 2021

The table below is a list of LifeKeeper for Linux v9 supported storage and should be considered when configuring your environment.

### About Supported Storage

Some types of storage used as shared storage in LifeKeeper require certification.

The following storage must be certified. Make sure the storage you plan to use is listed in the table below. Supported storage should be used for shared storage such as SCSI/FC/iSCSI/SAS that refer to the same data from multiple nodes for which LifeKeeper's IO fencing with SCSI-2/3 Reservation is required.

### Configurations that do not require certification

- NAS storage (requires the NAS Recovery Kit)
- All disk devices that make up data replication by DataKeeper (both built-in and external)
- Virtual disks on shared storage protected by the VMDK Recovery Kit on vSphere
- Storage which is used in the environment where all of the requirements below are satisfied:
  - OS, hardware and platform has supported the storage
  - The LifeKeeper SCSI reservation feature is disabled
  - The fencing mechanism of LifeKeeper Quorum/Witness is configured

**Note:** If your shared storage meets the requirements above, it does not require SIOS certification.

### Unsupported Hardware

Consumer storage that is connected via USB or IEEE1394 is not supported.

## DELL

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)	Notes
DELL	Dell EqualLogic: PS4100 PS4110	iSCSI	Multi Path (DMMP)	YES	DMMP ARK	v9.0 – v9.6	With a large number of luns (over 20), change the REMOTETIMEOUT setting in
			Single	YES	DMMP		

PS4210 PS6100 PS6110 PS6210 PS6610		Path (DMMP)		ARK		/etc/default/LifeKeeper to REMOTETIMEOUT=600.
MD3800f MD3820f MD3860f	FC	vSphere (RDM)	No	–	v9.0 – v9.6	
		Multi Path (DMMP)	YES	DMMP ARK		
		Single Path	No	–		
MD3400 MD3420 MD3460	SAS	vSphere (RDM)	No	–	v9.0 – v9.6	
		Multi Path (DMMP)	YES	DMMP ARK		
		Single Path	No	–		
SCv3000 Series: SCv3000 SCv3020	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
		Multi Path (DMMP)	YES	DMMP ARK		
		Single Path	No	–		
	ISCSI	Multi Path (DMMP)	YES	DMMP ARK	v9.0 – v9.6	
		Single Path	No	–		
	SAS	vSphere	No	–	v9.0 – v9.6	
		Multi Path (DMMP)	YES	DMMP ARK		
		Single Path	No	–		
	SCv2000	FC	vSphere	YES	No	v9.0 – v9.6

	Series: SCv2000 SCv2020 SCv2080		(RDM)		Multipath ARK	v9.0 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
		ISCSI	Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
			SAS	vSphere	No		
	Multi Path (DMMP)	YES		DMMP ARK			
	Single Path	No		–			
	Dell Storage: SC7020 SC5020 SC4020	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
		ISCSI	Multi Path (DMMP)	YES	DMMP ARK		
Single Path			No	–			
FC			vSphere (RDM)	No	–		
	Multi Path (DMMP)	No	–				
	Single Path	No	–				
Dell Compellent: SC9000 SC8000	FC	vSphere (RDM)	No	–	–		
		Multi Path (DMMP)	No	–			
		Single Path	No	–			
	ISCSI	Multi	YES	DMMP			v9.0 – v9.6

			Path (DMMP)		ARK		
			Single Path	No	-		

## Fujitsu

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)	Notes
Fujitsu	ETERNUS: DX60 S2 DX80 S2 DX90 S2 DX410 S2 DX440 S2 DX8100 S2 DX8700 S2 DX60 S3 DX100 S3 DX200 S3 DX500 S3 DX600 S3 DX8700 S3 DX8900 S3 DX200F DX60 S4 DX100 S4 DX200 S4 DX500 S4 DX600 S4 DX60 S5 DX100 S5 DX200 S5 DX500 S5 DX600 S5 DX900 S5 AF250 AF650	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	If ETERNUS is used as a boot device, a file system error may occur when the boot device is not accessible from a primary node and switching to a secondary node may fail. In order to avoid this, set a panic when a file system error occurs as shown below.  Example: for xfs, configure the following and reboot the server.  <pre>echo 'fs.xfs.panic_mask=127' &gt; /etc/sysctl.d/01-xfs.conf</pre>
	Multi Path (DMMP)		YES	DMMP ARK			
	Multi Path (EMPD)		YES	No Multipath ARK			
	Single Path		YES	No Multipath ARK			
	ISCSI	Multi Path (DMMP)	YES	DMMP ARK	v9.0 – v9.6		
		Multi Path (EMPD)	YES	No Multipath ARK			
		Single Path	YES	No Multipath ARK			

AF250 S2						
AF650 S2						
AF250 S3						
AF650 S3						
DX8900						
S4						

## Hitachi / Hitachi Vantara

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)
Hitachi / Hitachi Vantara	Hitachi Virtual Storage Platform G1000 Hitachi Virtual Storage Platform G1500 Hitachi Virtual Storage Platform F1500	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6
			Multi Path (HDLM)	YES	HDLM ARK	
			Multi Path (DMMP)	YES	DMMP ARK	
			Single Path	YES	–	
	ISCSI	Multi Path (HDLM)	No	–	–	
		Single Path	No	–		
	Hitachi Virtual Storage Platform	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6
			Multi Path (HDLM)	YES	HDLM ARK	
			Single Path	YES	–	
		ISCSI	Multi Path (HDLM)	No	–	–
Single Path			No	–		
Hitachi Virtual Storage Platform Mid Range Family:	FC	vSphere (RDM)	YES	No Multipath	v9.0 – v9.6	

	Hitachi Virtual Storage Platform F900 (VSP F900) Hitachi Virtual Storage Platform F700 (VSP F700) Hitachi Virtual Storage Platform F370 (VSP F370) Hitachi Virtual Storage Platform F350 (VSP F350) Hitachi Virtual Storage Platform G900 (VSP G900)				ARK			
		Multi Path (HDLM)	YES	HDLM	ARK			
		Multi Path (DMMP)	YES	DMMP	ARK			
		Single Path	YES	No Multipath	ARK			
	Hitachi Virtual Storage Platform G700 (VSP G700) Hitachi Virtual Storage Platform G370 (VSP G370) Hitachi Virtual Storage Platform G350 (VSP G350) Hitachi Virtual Storage Platform G150 (VSP G150) Hitachi Virtual Storage Platform G130 (VSP G130)	ISCSI	Multi Path (DMMP)	YES	DMMP	ARK	v9.0 – v9.6	
			Single Path	YES	No Multipath	ARK		
	Hitachi Virtual Storage Platform Mid Range Family: Hitachi Virtual Storage Platform F800 (VSP F800) Hitachi Virtual Storage Platform F600 (VSP F600) Hitachi Virtual Storage Platform F400 (VSP F400) Hitachi Virtual Storage Platform G800 (VSP G800) Hitachi Virtual Storage Platform G600 (VSP G600)	FC	vSphere (RDM)	YES	No Multipath	ARK	v9.0 – v9.6	
			Multi Path (HDLM)	YES	HDLM	ARK		
			Multi Path (DMMP)	YES	DMMP	ARK		
			Single Path	YES	No Multipath	ARK		
		Hitachi Virtual Storage Platform G400 (VSP G400) Hitachi Virtual Storage Platform G200 (VSP G200) Hitachi Virtual Storage Platform G100 (VSP G100)	ISCSI	Multi Path (DMMP)	YES	DMMP	ARK	v9.0 – v9.6
				Single Path	YES	No Multipath	ARK	
Hitachi Unified Storage VM (HUS VM)	FC	vSphere (RDM)	YES	No Multipath	ARK	v9.0 – v9.6		
		Multi Path (HDLM)	YES	HDLM	ARK			
		Single Path	YES	No Multipath				

		ISCSI	Multi Path (HDLM)	No	–	–
			Single Path	No	–	
	Hitachi Unified Storage 100 Series: Hitachi Unified Storage 150 (HUS 150) Hitachi Unified Storage 130 (HUS 130) Hitachi Unified Storage 110 (HUS 110)	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6
			Multi Path (HDLM)	YES	HDLM ARK	
			Single Path	YES	No Multipath ARK	
		ISCSI	Multi Path (HDLM)	No	–	–
			Single Path	No	–	

## Hitachi

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)	Notes
Hitachi	BR1200	ISCSI	Multi Path (HDLM)	No	–	–	Both the single path and multipath configurations require the RDAC driver. Only the BR1200 configuration using the RDAC driver is supported.  The BR1200 configuration using the HDLM (HDLM ARK)
			Multi Path (DMMP)	No	–		
			Single Path	No	–		
		SAS	Multi Path (RDAC)	YES	No Multipath ARK	v9.0 – v9.6	
			Single Path (RDAC)	YES	No Multipath ARK		

							is not supported.
	BR1250	ISCSI	Multi Path (HDLM)	No	–	–	
			Multi Path (DMMP)	No	–		
			Single Path	No	–		
	SAS	Multi Path (NECSPS)	YES	NEC SPS ARK	v9.0 – v9.6	Both the single path and multipath configurations require the NEC SPS ARK. The BR1250 configuration using HDLM (HDLM ARK) is not supported.  LifeKeeper for Linux v9.1 or later can be used for NEC iStorage StoragePathSavior (NECSPS) Recovery Kit on RHEL7 environment	
		Single Path (NECSPS)	YES	NEC SPS ARK			
	BR1650E BR 1650S	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
Multi Path (HDLM)			YES	HDLM ARK			
Single Path			YES	–			
ISCSI		Multi Path (HDLM)	No	–	–		
		Single Path	No	–			

## HPE

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)	Notes
HPE	Primera A630,	FC	vSphere (RDM)	YES	No Multipath	v9.0 – v9.6	

A650, A670, C630, C650, C670				ARK		
		Multi Path (DMMP)	YES	DMMP ARK		
		Single Path	No	–		
Nimble Storage: AF20 AF20Q AF40 AF60 AF80 AF1000 AF3000 AF5000 AF7000 AF9000 CS1000H CS1000 CS3000 CS5000 CS7000 HF20C HF20H HF20 HF40 HF60	ISCSI	Multi Path (DMMP)	YES	DMMP ARK	v9.2.2 – v9.6	
		Single Path	No	–		
3PAR: 3PAR StoreServ 8200 3PAR StoreServ 8400 3PAR StoreServ 8440 3PAR StoreServ 8450	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
		Multi Path (DMMP)	YES	DMMP ARK		
		Single Path	No	–		
3PAR: 3PAR StoreServ	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	<b>Note:</b> 3PAR StoreServ 7400 returns a reservation conflict with the default

	7400 3PAR StoreServ 7400c 3PAR StoreServ 7440 3PAR StoreServ 7440c 3PAR StoreServ		Multi Path (DMMP)	YES	DMMP ARK		path checker setting. To avoid this conflict, set the following parameter in “/etc/default/LifeKeeper”: DMMP_REGISTRATION_TYPE=hba User friendly device mapping is not supported. Set the following parameter in “multipath.conf” “user_friendly_names no”  * When using the vSphere(RDM) for multipath configurations, no configuration parameters need to be set.
			Single Path	No	–		
	7450 3PAR StoreServ 7450c	ISCSI	Multi Path (DMMP)	YES	DMMP ARK	v9.0 – v9.6	<b>Note:</b> 3PAR StoreServ 7400 iSCSI returns a reservation conflict. To avoid this conflict, set the following parameter in “/etc/default/LifeKeeper”: DMMP_REGISTER_IGNORE=TRUE
			Single Path	No	–		
	3PAR: 3PAR StoreServ 7200 3PAR StoreServ 7200c	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	<b>Note:</b> 3PAR StoreServ 7200 returns a reservation conflict with the default path checker. To avoid this conflict, set the following parameter in “/etc/default/LifeKeeper”: DMMP_REGISTRATION_TYPE=hba
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
		ISCSI	Multi Path (DMMP)	No	–	–	
			Single Path	No	–		
	3PAR: 3PAR StoreServ 9450 3PAR StoreServ 20800 R2 3PAR StoreServ 20840 R2	FC	vSphere (RDM)	YES	No Multipath ARK	v9.2.1 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
Single Path			No	–			
ISCSI		Multi	No	–	–		

	3PAR StoreServ 20850 R2		Path (DMMP)				
			Single Path	No	–		
	3PAR: 3PAR StoreServ 20450 3PAR StoreServ 20800	FC	vSphere (RDM)	YES	No Multipath ARK	v9.1 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
	3PAR StoreServ 20840 3PAR StoreServ 20850	ISCSI	Multi Path (DMMP)	No	–	–	
			Single Path	No	–		
	3PAR StoreServ 10400	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
		ISCSI	Multi Path (DMMP)	No	–	–	
			Single Path	No	–		
	3PAR StoreServ 10800	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	<p><b>Note:</b> 3PAR StoreServ 10800 FC returns a reservation conflict with the default path checker setting. To avoid this conflict, set the following parameter in “/etc/default/LifeKeeper”: DMMP_REGISTRATION_TYPE=hba</p> <p>User friendly device mapping is not supported. Set the following</p>
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		

							parameter in “multipath.conf” “user_friendly_names no”  * When using the vSphere(RDM) for multipath configurations, no configuration parameters need to be set.
		ISCSI	Multi Path (DMMP)	No	–	–	
			Single Path	No	–		
MSA: MSA2060	FC		vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6.0	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
	ISCSI		Multi Path (DMMP)	No	–	–	
			Single Path	No	–		
	SAS		vSphere (RDM)	No	–	–	
			Multi Path (DMMP)	No	–		
			Single Path	No	–		
	MSA: MSA2050 MSA2052	FC		vSphere (RDM)	YES	No Multipath ARK	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		

		ISCSI	Multi Path (DMMP)	YES	DMMP ARK	v9.2 – v9.6	
			Single Path	No	–		
		SAS	vSphere (RDM)	No	–	–	
			Multi Path (DMMP)	No	–		
			Single Path	No	–		
		MSA: MSA2042	FC	vSphere (RDM)	YES	No Multipath ARK	
	Multi Path (DMMP)			YES	DMMP ARK		
	Single Path			No	–		
	ISCSI		Multi Path (DMMP)	YES	DMMP ARK	v9.1 – v9.6	
Single Path			No	–			
SAS	vSphere (RDM)		YES	No Multipath ARK	v9.1 – v9.6		
	Multi Path (DMMP)		YES	DMMP ARK			
	Single Path		No	–			
MSA: MSA1040 MSA2040	FC		vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
		Multi Path (DMMP)	YES	DMMP ARK			

			Single Path	No	–	v9.0 – v9.6	
		ISCSI	Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
		SAS	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
	StoreVirtual (LeftHand) Series LeftHand OS version 12.6	FC	vSphere (RDM)	No	–	–	
			Multi Path (DMMP)	No	–		
			Single Path	No	–		
		ISCSI	Multi Path (DMMP)	YES	DMMP ARK	v9.1 – v9.6	
			Single Path	No	–		
	HP P9500	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
Multi Path (DMMP)			YES	DMMP ARK			
Single Path			YES	No Multipath ARK			
XP7 Storage System	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	LeftHand OS (SAN/iQ) version 12.6 is supported in HPE StoreVirtual (LeftHand) storage.All StoreVirtual series are supported, including HC250 / HC380 and StoreVirtual VSA as the virtual storage appliance.See the release note of each version of LeftHand OS (SAN/iQ) to know the virtual storage appliance in detail.	

			Multi Path (HDLM)	YES	HDLM ARK		
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	YES	No Multipath ARK		

## IBM

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)	Notes
IBM	IBM SAN Volume Controller* *IBM TotalStorage Proven	FC	vSphere (RDM)	No	-	v9.0 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	YES	No Multipath ARK		
		ISCSI	Multi Path (DMMP)	No	-	-	
			Single Path	No	-		
	IBM Storwize V7000	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	-		
ISCSI		Multi Path (DMMP)	YES	DMMP ARK	v9.0 – v9.6	The IBM Storwize V7000 (Firmware Version 6.3.0.1)	

					Quorum/ Witness Server Kit		<p>has been certified by partner testing using iSCSI (iscsi-initiator-utils6.2.0.872-34.el6.x86_64) with DMMP (device-mapper-1.02.66-6.el6, device-mapper-multipath-0.4.9-46.el6). The test was performed with LifeKeeper for Linux v7.5 using RHEL 6.2.</p> <p><b>Restriction:</b> IBM Storwize V7000 must be used in combination with the Quorum/Witness Server Kit and STONITH. See the online documentation for the LifeKeeper version that you are using for more information. <a href="http://docs.us.sios.com/">http://docs.us.sios.com/</a></p>
			Single Path	No	–		
	IBM Storwize V3700	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path	No	–		
		ISCSI	Multi Path (DMMP)	No	–	–	
			Single Path	No	–		
		SAS	vSphere (RDM)	No	–	–	
			Multi Path (DMMP)	No	–		
			Single Path	No	–		

	IBM XIV Storage System	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	If you are creating more than 32 LUNs on an IBM XIV Storage System using LifeKeeper, please contact IBM for further details.
			Multi Path (DMMP)	YES	DMMP ARK		
			Single Path (DMMP)	YES	DMMP ARK		
		ISCSI	Multi Path (DMMP)	No	–	–	
			Single Path	No	–		

## Lenovo

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)
Lenovo	Lenovo Storage V3700 V2 Lenovo Storage V3700 V2 XP	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6
			Multi Path (DMMP)	YES	DMMP ARK	
			Single Path	No	–	
		ISCSI	Multi Path (DMMP)	No	–	–
			Single Path	No	–	
			SAS	vSphere (RDM)	No	
		Multi Path (DMMP)		No	–	
		Single Path		No	–	
		IBM Storwize V3700	FC	vSphere (RDM)	YES	No Multipath ARK
	Multi Path (DMMP)			YES	DMMP ARK	

		ISCSI	Single Path	No	–	–
			Multi Path (DMMP)	No	–	
			Single Path	No	–	
		SAS	vSphere (RDM)	No	–	–
			Multi Path (DMMP)	No	–	
			Single Path	No	–	

## NEC

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)	Notes
NEC	NEC iStorage: M10e M11e M100 M110 M300 M12e M120 M320	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6	LifeKeeper for Linux v9.1 or later can be used for NEC iStorage StoragePathSavior (NECSPS) Recovery Kit on RHEL7 environment.
			Multi Path (NECSPS)	YES	NECSPS ARK		
			Single Path (NECSPS)	YES	NECSPS ARK		
		ISCSI	Multi Path (NECSPS)	YES	NECSPS ARK	v9.0 – v9.6	
			Single Path (NECSPS)	No	–		
		SAS	vSphere (RDM)	No	–	v9.0 – v9.6	
			Multi Path (NECSPS)	YES	NECSPS ARK		
			Single Path (NECSPS)	YES	NECSPS ARK		
NEC iStorage:	FC	vSphere (RDM)	YES	No Multipath	v9.0 – v9.6		

	M310 M310F M320F M500 M510 M700 M710				ARK	
			Multi Path (NECSPS)	YES	NECSPS ARK	
			Single Path (NECSPS)	YES	NECSPS ARK	
			Multi Path (NECSPS)	YES	NECSPS ARK	
	M710F M520 M720 M720F	ISCSI	Multi Path (NECSPS)	YES	NECSPS ARK	v9.0 – v9.6
			Single Path (NECSPS)	No	–	

## Pure Storage

Vendor (*1)	Storage Model	Bus Type	Connect Type (*2)	Support Status	Required ARK (i.e. Multipath ARK)	Supported LK Version (*4)
Pure Storage	FA-400 Series	FC	vSphere (RDM)	YES	No Multipath ARK	v9.0 – v9.6
			Multi Path (DMMP)	YES	DMMP ARK	
			Single Path	No	–	
		ISCSI	Multi Path (DMMP)	No	–	–
			Single Path	No	–	
			FlashArray//m Series //m10 //m20 //m50 //m70	FC	vSphere (RDM)	
	Multi Path (DMMP)	YES	DMMP ARK			
	Single Path	No	–			
	FlashArray//x Series //x10 //x20 //x50 //x70	FC	vSphere (RDM)	No	–	v9.0 – v9.6
			Multi Path (DMMP)	YES	DMMP ARK	
			Single Path	No	–	
	FlashArray//x Series //X10R2	FC	vSphere (RDM)	No	–	v9.3.2 – v9.6

	//X20R2 //X50R2 //X70R2 //X90R2		Multi Path (DMMP)	YES	DMMP ARK	v9.0 – v9.6
			Single Path	No	–	
		iSCSI	vSphere (RDM)	No	–	
			Multi Path (DMMP)	YES	DMMP ARK	
		Single Path	No	–		

**(\*1) EMC storage for use**

LifeKeeper acquired the E-Lab certification. See the matrix below for the EMC storage support.

<http://www.emc.com/interoperability>

**(\*2) Connecting Configurations**

vSphere (RDM)	Connected a shared disk using Raw Device Mapping (RDM) on the VMWare ESX server
Multi Path (DMMP)	Multipath configuration of the shared disk using the Device Mapper Multipath driver
Multi Path (EMPD)	Multipath configuration of the shared disk using the ETERNUS Multipath Driver
Multi Path (HDLM)	Multipath configuration of the shared disk using the Hitachi Dynamic Link Manager driver
Multi Path (RDAC)	Multipath configuration of the shared disk using the IBM Redundant Disk Array Controller driver
Multi Path (NECSPS)	Multipath configuration of the shared disk using the iStorage StoragePathSavior driver
Single Path (DMMP)	Single path configuration of the shared disk using the Device Mapper Multipath driver
Single Path (RDAC)	Single path configuration of the shared disk using the IBM Redundant Disk Array Controller
Single Path	Single path connection not using a specific multipath driver

**(\*3) Multipath Software**

The support information by maker must be referred about the connectivity information of each Multipath software.

Please refer to the release notes for the details about Multipath software support by LifeKeeper.

## (\*4) Supported versions of LifeKeeper

This document is written for LifeKeeper for Linux v9.

If you are using a previous version of LifeKeeper, you must refer to the Release Notes for that version for the supported storage list.

# 11. Evaluation Guides

---

[DataKeeper for Linux Evaluation Guide](#)

[LifeKeeper for Linux Evaluation Guide for Cloud Environments](#)

# 11.1. DataKeeper for Linux Evaluation Guide

---

## Objective

This document is intended to aid you in installing, configuring and using the LifeKeeper for Linux evaluation product with DataKeeper to enable real time, host based, block-level data replication

There are five phases in this process:

- Phase 1 – [Prepare to Install](#)
- Phase 2 – [Configure Storage](#)
- Phase 3 – [Install LifeKeeper for Linux](#)
- Phase 4 – [Configure your LifeKeeper Cluster](#)
- Phase 5 – [Test Your Environment](#)

## 11.1.1. DataKeeper for Linux Terms to Know

The following terms are used throughout this document and, while some may be familiar to you, it may be helpful to review how SIOS defines and uses these terms.

### Network Communication Terms

**Crossover cable** – A cable used to directly connect computing devices together, instead of being connected to a network switch, hub or router. This cable creates an isolated, private network to allow cluster-related and data replication traffic to flow between systems.

### Types of LifeKeeper Servers

**Server** – A computer system dedicated to running software application programs.

**Active Server** – This is the server where the resource hierarchy is currently running (IN SERVICE).

**Standby Server** – This is the server where the resource hierarchy is defined, but is not currently running. This server is available to bring the resource hierarchy into service should something happen to the resource hierarchy on the Active Server.

**Primary Server** – This is the server in a LifeKeeper configuration with the highest priority for a given resource hierarchy. It is the server that provides services for the resource hierarchy under normal circumstances.

**Secondary Server** – This is the server in a LifeKeeper configuration with the 2nd highest priority for a given resource hierarchy.

**Source Server** – In a LifeKeeper cluster, using data replication, this is the Active Server. It is where the resource hierarchy is currently running and the replicated partition (Source Partition) is accessible for writes.

**Target Server** – In a LifeKeeper cluster, using data replication, this is the Standby Server. The replicated partition (Target Partition) is updated with writes from the Source Partition by the SIOS Data Replication system. This partition should not be accessed/modified manually.

### SIOS DataKeeper Terms

**Replication** – Transferring data from one partition to another via a sector-by-sector copy. During replication, the target partition should not be accessed or modified assuring your data integrity.

**Asynchronous** – A replication scheme in which the data is released for writing on the source immediately and is sent to the target(s) simultaneously for writing as fast as the data can get there and can be written on them.

**Rate of Change** – A measure of the amount of data which is changing over a set period of time.

**Compression** – An algorithm which is optionally implemented to reduce the amount of traffic between source and target nodes. Nine levels of compression are offered. Compression is turned off by default.

**Throttling** – An optionally implemented mechanism to limit the bandwidth used for replication.

## LifeKeeper Product Terms

**Communications Path** – A mechanism supporting communication between nodes in a LifeKeeper cluster. SIOS highly recommends implementing multiple communication paths between all servers in the cluster to eliminate a single point of failure.

**Heartbeat** – A periodic message exchanged between nodes in a LifeKeeper cluster that provides server health monitoring. A heartbeat message is one type of inter-node cluster communication sent over a communications path.

**Split Brain** – A situation in which all communications paths between cluster members fail, but all servers remain up and running. In this situation, both systems believe the other has failed and both believe they should keep or bring resources into service.

**Failover** – The unplanned migration of a resource hierarchy to the Standby Server because of a system or resource failure on the Active Server.

**Switchover** – The planned migration of a resource hierarchy from the Active Server to the Standby Server.

**Switchback** – The setting that governs the recovery behavior of the server where the resource was in service when it failed. If the setting is “intelligent”, the server acts as a possible backup for the given resource. If the setting is “Automatic”, the server actively attempts to re-acquire the resource without further notice.

**Resource** – A system asset that can be protected by LifeKeeper. Resources can be used to represent disk partitions, virtual IP addresses, applications, etc.

**Extend a Resource** – Create or define an already configured LifeKeeper resource onto another server in the cluster and build an equivalency relationship that prevents the resource from coming in service on both systems simultaneously

**Resource Hierarchy** – A grouping of resources, in a predetermined order, from high to low. This may also be referred to as simply a Hierarchy.

**Shared Storage** – One or more logical disk partitions that are physically attached to all nodes in a cluster. LifeKeeper ensures that the volume is only accessible by one server at a time. This is formally called I/O fencing.

**Data Replication (Disk Mirroring)** – The replication of logical disk partitions to separate physical hard disks in real time to ensure continuous availability, currency and accuracy of data.

**Source** – The partition on the source server used for replication. The “gold” copy of the data.

**Target** – The partition on the target server used for replication.

**Switchable IP Address** – A unique IP address that may be moved between systems in the cluster. Client systems connect to this address and the system where the virtual IP resource is active will respond to requests.

## 11.1.2. The Evaluation Process

---

SIOS strongly recommends performing your evaluation of LifeKeeper for Linux within a test lab environment. SIOS is not responsible and cannot provide support for evaluation software installed in a production environment.

All questions during the evaluation period should be directed to [evalsupport@us.sios.com](mailto:evalsupport@us.sios.com) or your local Pre-Sales Engineering contact. Pre-sales support will contact you by the next business day to answer questions. Once you are a licensed customer on software maintenance, you will have access to 24 × 7 post-sales technical support.

 **Important:** Your evaluation license is valid for a limited period of time from the day you receive the SIOS product evaluation package and licenses via email from the SIOS sales team.

## 11.1.3. Prepare to Install DK for Linux

---

### Hardware Requirements

#### Primary and Secondary Servers

- Systems must meet the minimum requirements for the Linux distribution to be used during the evaluation.
- 512MB RAM minimum; 1GB RAM recommended.
- 2GB of available hard disk space recommended.
- Multiple Network Interface Cards (NIC's) are recommended.
- Configure one or more additional disk partitions to be used for data replication. On the primary server, these will become the source partitions. On the secondary server(s), these will become the target partitions. In this evaluation example we will be replicating MySQL data (which will be a partition mounted at `/var/lib/mysql`)
- For replicated partitions, a target partition's size must equal to or larger than the size of its source partition.
- The system (`/`) and boot (`/boot`) partitions are not eligible for replication.

### Software Requirements

#### Primary Server and Secondary Server

- Linux Distribution x86\_64, AMD 64:
  - RedHat Enterprise Linux 6.x, 7.x, or 8.x
  - CentOS Linux 6.x, 7.x, or 8.x
  - Oracle Enterprise Linux 6.x, 7.x, or 8.x
  - SuSE Linux Enterprise Server 11, 12, or 15
  - See [Linux Release Notes](#) for a full list of supported Operating Systems
- Current patches / security updates are recommended.
- Satisfied dependencies; especially if the Linux installation package selection was base/minimal you will need to refer to the dependencies documentation at [Linux Dependencies](#)
- It is recommended the firewall be disabled
  - `# service iptables stop (systemctl stop firewalld)`
  - `# chkconfig iptables off (systemctl disable firewalld)`
  - See [here](#) for information regarding the ports LifeKeeper for Linux uses.
- Disable SELinux :

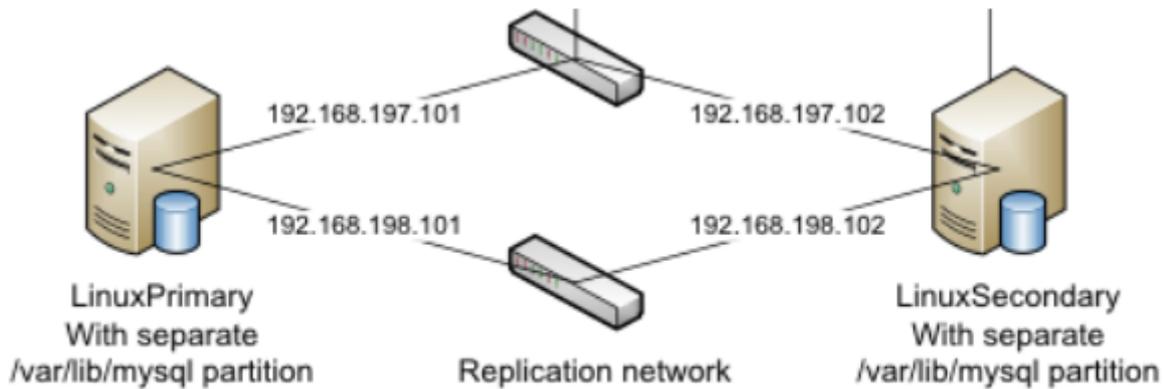
- Edit `/etc/selinux/config`
- Set `SELINUX=disabled` (note: permissive mode is also acceptable)
- Check the configuration of your `/etc/hosts` file
  - `localhost.localdomain` and `localhost` are the only entries that can be on `127.0.0.1`
  - Create a separate entry for your hostname with a static address
- GUI Authentication with PAM
  - LifeKeeper for Linux now leverages the Pluggable Authentication Module (PAM) provided in the Linux Standard Base (LSB).
  - Users are identified and authenticated against the system's PAM configuration. Privilege levels are determined from group membership as provided through PAM.
  - In order to access the GUI, a user must be a member in one of the three LifeKeeper groups: `lkadmin`, `lkoper` or `lkguest`.
  - See [Configuring GUI Users](#) for more information.

## Network Requirements

For your evaluation, we recommend configuring your machines similarly to the following example.

LinuxPrimary and LinuxSecondary are multi-homed, between two LAN segments (the second NIC in each server could even be connected via a cross-over cable if a second physical network is not available). The second NIC is optional in this configuration, but highly recommended in production environments to avoid a single point of failure.

In this example evaluation scenario we will be replicating direct attached storage



**Network Configuration Example**

## IMPORTANT: Rate of Change Analysis

When replicating data in real time, it's critical to ensure that you have sufficient bandwidth to keep the replication in a mirroring state at all times. To perform a Rate of Change analysis on your server, which will collect and analyze Write activity over time vs. bandwidth, refer to [Measuring Rate of Change on a Linux System](#).

### Primary Server and Secondary Servers

- Configure the Host file with entries for all LifeKeeper protected servers. This is typically /etc/hosts.

Example:

```
192.168.197.101 LinuxPrimary
```

```
192.168.197.102 LinuxSecondary
```

- See your Network Administrator to obtain an unused IP Address to be used as the switchable IP Address. This switchable IP Address will be created later in the configuration process.

- Public Network connection(s) configured with:

◦ Static IP address

◦ Correct subnet mask

- Correct gateway address

- Private Network connection(s) configured with:

- Static IP address (on a different subnet from the public network)

- Correct network mask

- No gateway IP address

- No DNS server addresses

## 11.1.4. Configure Storage for DK for Linux

---

### Before You Begin

Ensure the following:

- You have an extra disk/partition on both servers that can be used for data replication. A target volume's size must equal to or larger than the size of its source disk/partition.

## Partition local storage for use with SIOS DataKeeper for Linux

### Primary Server

On your Primary server, perform the following actions:

1. Identify an existing free, unused disk partition to use as the source of the mirror. Alternatively, create a new partition. Use the "gdisk" utility to partition your disk appropriately. In this example /dev/sdb is an unused disk where we will create a single partition
  - a. `gdisk /dev/sdb`
  - b. Press "n" to create a new partition
  - c. This example uses a new disk, so we will use all default values (Partition 1, entire disk and Linux filesystem partition type) Hit Enter four times to confirm these parameters.
  - d. Press "w" to write the partition table
  - e. Press "Y" to confirm to overwrite existing partitions

### Example

```
[root@LinuxPrimary ~]# gdisk /dev/sdb
```

```
GPT fdisk (gdisk) version 0.8.10
```

```
Partition table scan:
```

```
MBR: protective
```

```
BSD: not present
```

```
APM: not present
```

```
GPT: present
```

```
Found valid GPT with protective MBR; using GPT.
```

```
Command (? for help): n
```

```
Partition number (1-128, default 1): <enter>
First sector (34-2047, default = 34) or {+-.}size{KMGTP}: <enter>
Last sector (34-2047, default = 2047) or {+-.}size{KMGTP}: <enter>
Current type is 'Linux filesystem'
Hex code or GUID (L to show codes, Enter = 8300): <enter>
Changed type of partition to 'Linux filesystem'
```

```
Command (? for help): w
```

```
Final checks complete. About to write GPT data. THIS WILL OVERWRITE EXISTING
PARTITIONS!!
```

```
Do you want to proceed? (Y/N): Y
OK; writing new GUID partition table (GPT) to /dev/sdb.
Warning: The kernel is still using the old partition table.
The new table will be used at the next reboot.
The operation has completed successfully.
```

```
[root@LinuxPrimary ~]#
```

2. Format the newly created disk partition

```
# mkfs.ext3 /dev/sdb1
```

3. Mount the partition at the desired location, for example /var/lib/mysql

```
# mount /dev/sdb1 /var/lib/mysql
```

4. Note: there is no need to add an entry to /etc/fstab. Lifekeeper will take care of mounting this automatically.

## Result

```
[root@LinuxPrimary ~]# df /var/lib/mysql
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/sdb1 253855 11083 229666 5% /var/lib/mysql
```

## Secondary Server

5. On your Secondary server, only perform Step #1 above, where you partition the disk. The size of the Target disk/partition needs to be the same size, or greater, than our Source disk/partition. There is no need to format or mount the filesystem.

## 11.1.5. Install LifeKeeper for Linux

For ease of installation, SIOS has provided the LifeKeeper for Linux with an installation script. Towards the end of the script, the desired Application Recovery Kits (ARKs) should be selected for installation. The software will be installed to the following locations:

SPS for Linux Component	Install Location
LifeKeeper Software	/opt/LifeKeeper
LifeKeeper Config File	/etc/default/LifeKeeper

Perform the following actions on **both** Primary and Secondary server.

### Download Software

1. Open the LifeKeeper evaluation email you received from SIOS.
2. Download the LifeKeeper Software from the link provided in your email. It is generally easiest to use “wget” to recursively download all files. Example:

a. # cd /root

b. # wget -r <URL>

c. After successful download you will have downloaded contents similar to the follow directory listing:

```
[root@LinuxPrimary ~]# ls -l <directory> total 63680
-rw-r--r-- 1 root root 23163 May 30 14:03 EULA.pdf
-rw-r--r-- 1 root root 536 May 30 14:03 readme.txt
-rw-r--r-- 1 root root 65179648 May 30 14:03 sps.img
```

3. Download your Evaluation license key from the link specified in your evaluation email. Save the license file to an easy to remember location on both servers.

### Run the LifeKeeper Installer Script

1. Loopback mount the sps.img file previously downloaded, which is an ISO9660 image file. Run the “setup” script inside:

```
# mount -o loop sps.img /mnt
```

```
# cd /mnt
```

```
# ./setup
```

2. During this procedure, you will hit Enter in most cases to accept default values and continue to the next screen. Note the following exceptions:
  - a. On the screen titled “High AvailabilityNFS” you may select “n” as in this particular eval guide we will not be creating a highly available NFS server cluster configuration.
  - b. If you have plans to create a highly available NFS service, adjust your response accordingly.
3. Towards the end of the setup script, you can choose to install a trial license key now, or later. We will install the license key in the next step, so you can safely select “n” at this point
4. In the final screen of the “setup” select the DataKeeper from the list displayed on the screen.
5. The following RPMs should be installed:
  - a. steeleye-lkDR-<version>.noarch.rpm
6. Un-mount the Distribution Enabling disk image:

```
# cd /root
```

```
# umount /mnt
```

## Install the Evaluation License Keys

The last phase of the setup process installs the licensing keys. You must install the evaluation license key file (“.lic”) that you downloaded with your evaluation software before starting the LifeKeeper for Linux.

1. To install your trial license key, run the “lkkeyins” command on both Primary and Secondary Server. This command is located at /opt/LifeKeeper/bin/lkkeyins . Example:

```
# /opt/LifeKeeper/bin/lkkeyins <path_to_license/<filename>.lic
```

2. Validate your license keys were installed via the /opt/LifeKeeper/bin/lklicmgr command

```
# /opt/LifeKeeper/bin/lklicmgr
```

```
License File: 20101230.lic
```

Product	Type	Expiry
LifeKeeper for Linux	Eval	27 Mar 2013 (87 days)
SIOS Data Replication ARK	Eval	27 Mar 2013 (87 days)

## Start the LifeKeeper for Linux

1. Start:

```
# /opt/LifeKeeper/bin/lkstart
```

## 11.1.6. Configure the Cluster – DK for Linux

### Primary Server

Complete the following steps on the primary server to configure the cluster:

- Create TCP Communication (Comm) Path(s)
- Verify the Communication (Comm) Path(s)

Before you begin, SIOS recommends at least two TCP communications paths between each server within the cluster to each remote server for heartbeat redundancy.

 **Important** Supported configurations require that you define redundant comm. paths, so that the failure of a single communication line will not cause a split brain where resource hierarchies may come in-service on multiple servers simultaneously

### Access the LifeKeeper GUI

The LifeKeeper Graphical User Interface (GUI) is a Java based application that can be run as a native Linux application, or as an applet within your Java-Enabled Web Browser.

The LifeKeeper GUI is based on Java RMI with callbacks. Hostnames must be resolvable or you may receive a Java 115 or 116 errors.

1. Verify that both short and fully qualified hostnames of all cluster nodes resolve to the proper locations

```
# ping LinuxPrimary
```

```
# ping LinuxPrimary.domain.com
```

```
# ping LinuxSecondary
```

```
# ping LinuxSecondary.domain.com
```

2. To start the LifeKeeper Linux GUI Application:

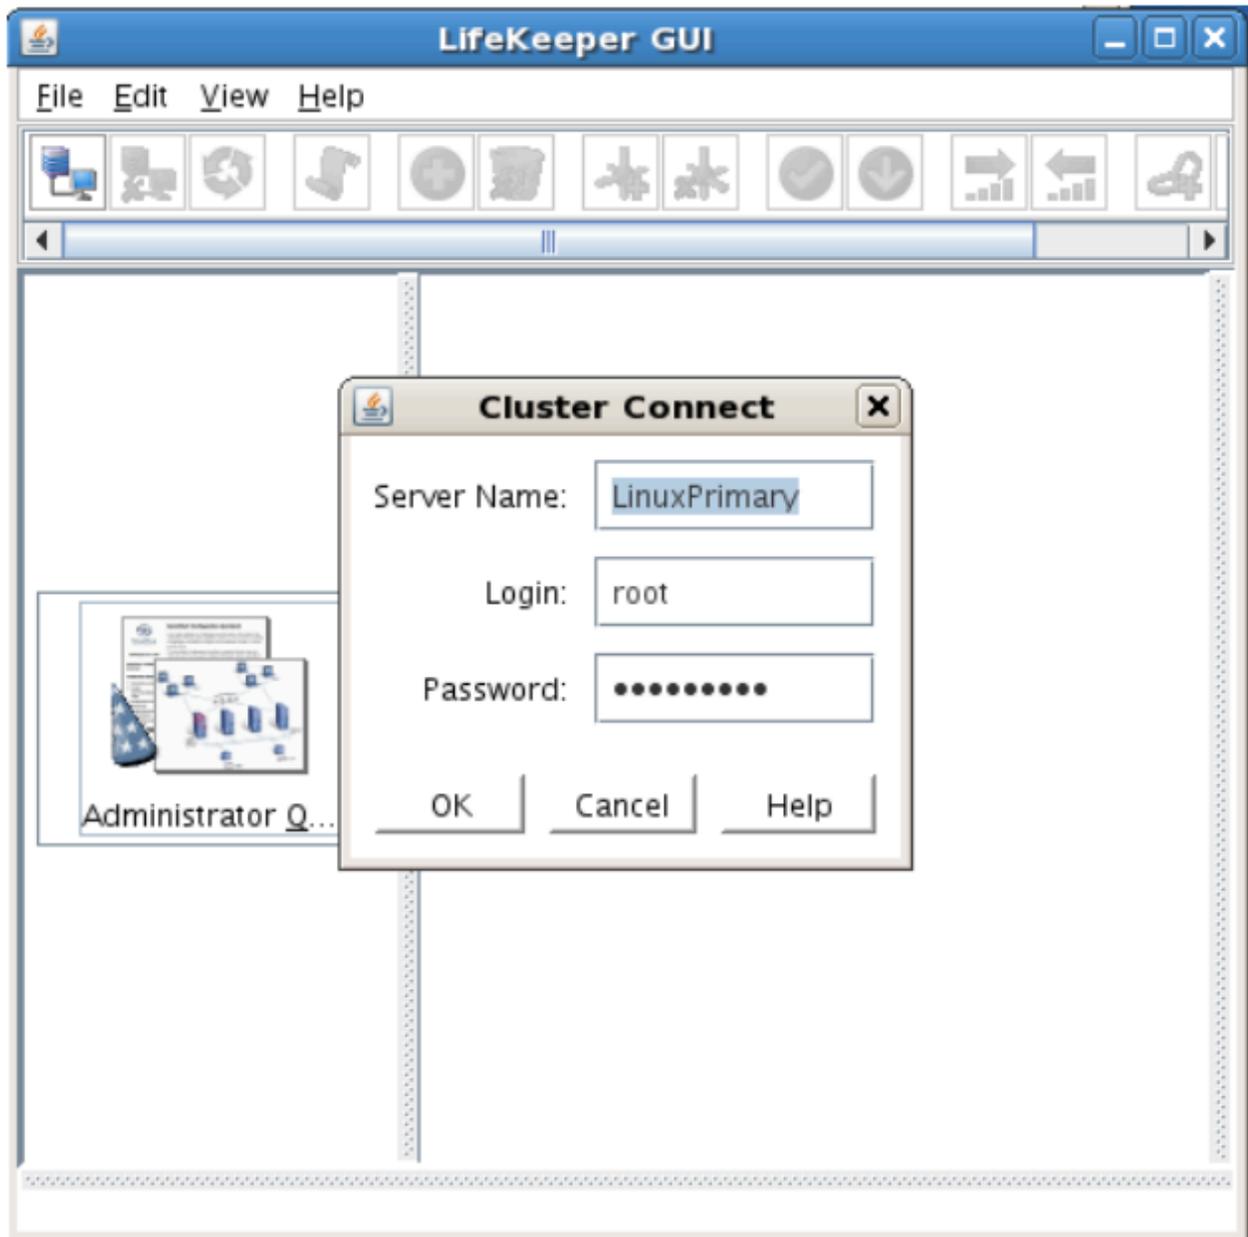
```
a. /opt/LifeKeeper/bin/lkGUIapp &
```

3. To Connect to the LifeKeeper GUI Applet from a Web Browser, go to:

```
a. http://<hostname>:81
```

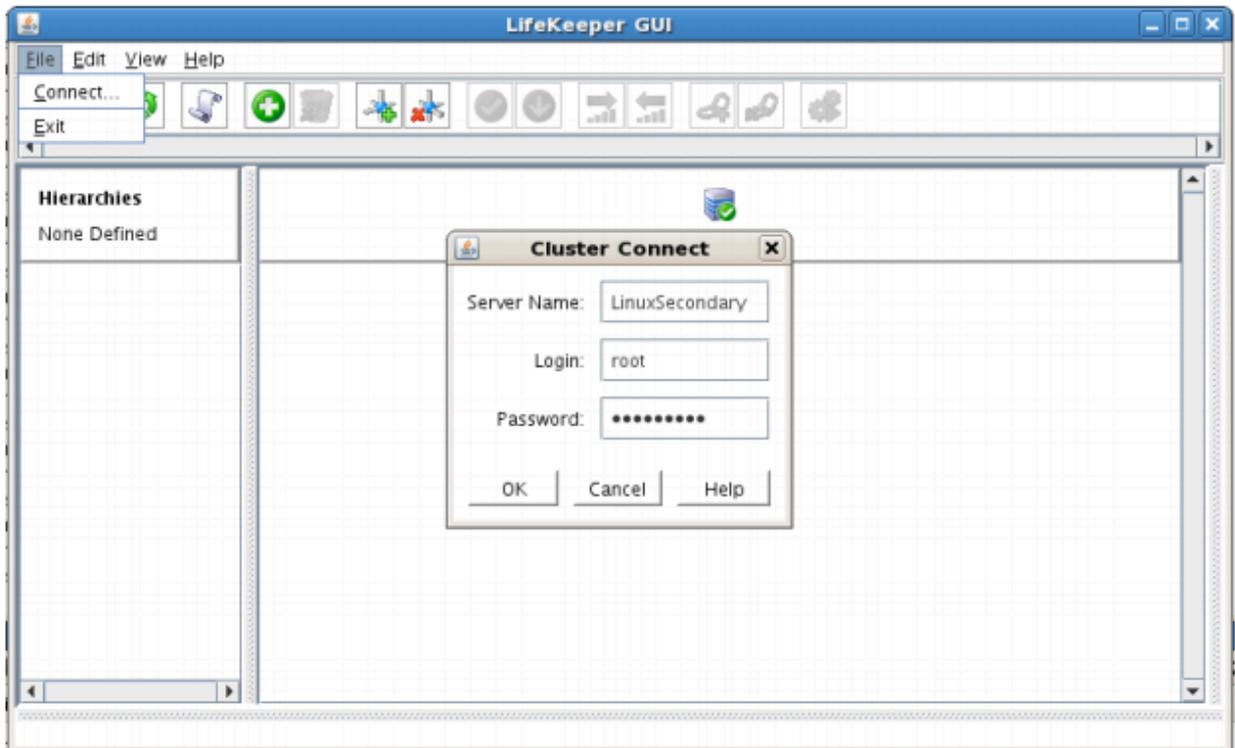
4. Enter the name of the server you wish to connect to (this field will be populated with the name of the server you are on, if you are running the GUI from a server with LifeKeeper installed) along

with your root credentials and click OK.



## Create Communication (Comm) Paths

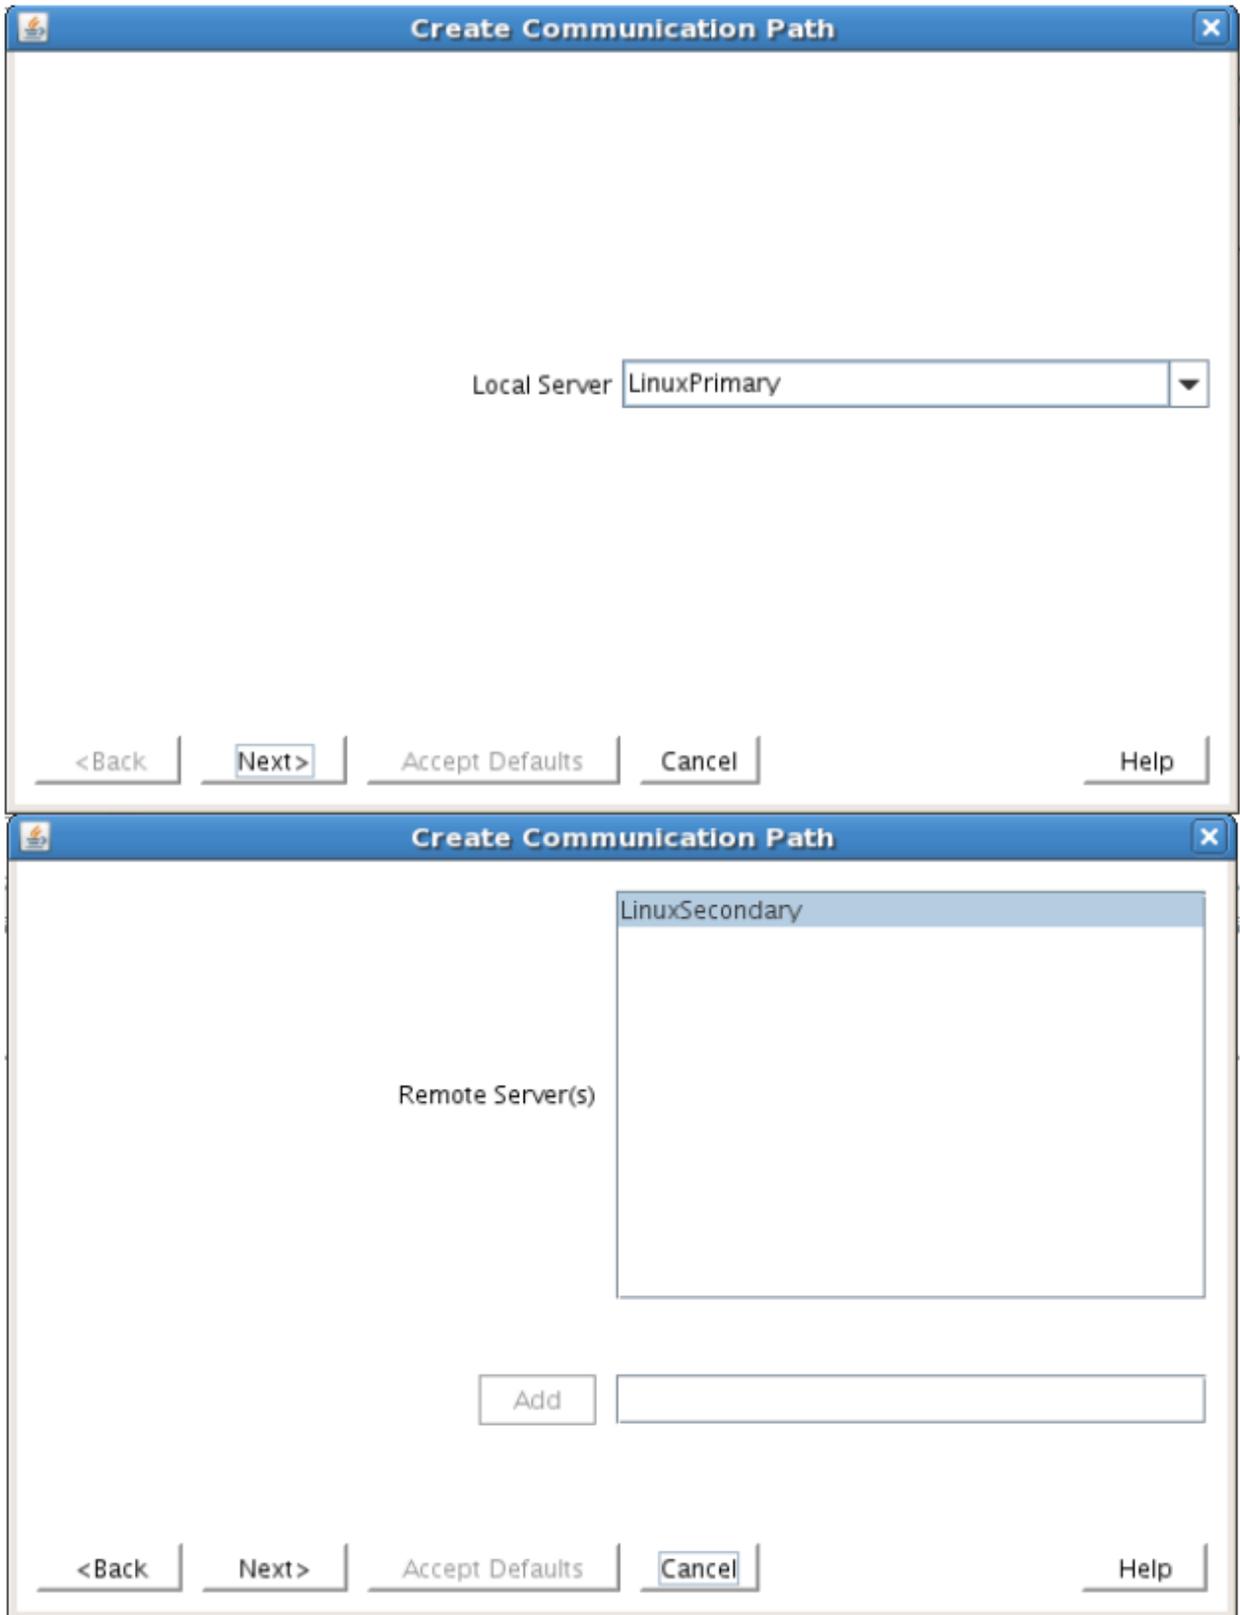
5. Within the LifeKeeper GUI, from the File menu, select Connect. Enter the name of your Secondary server, login and password when the Cluster Connect window displays.



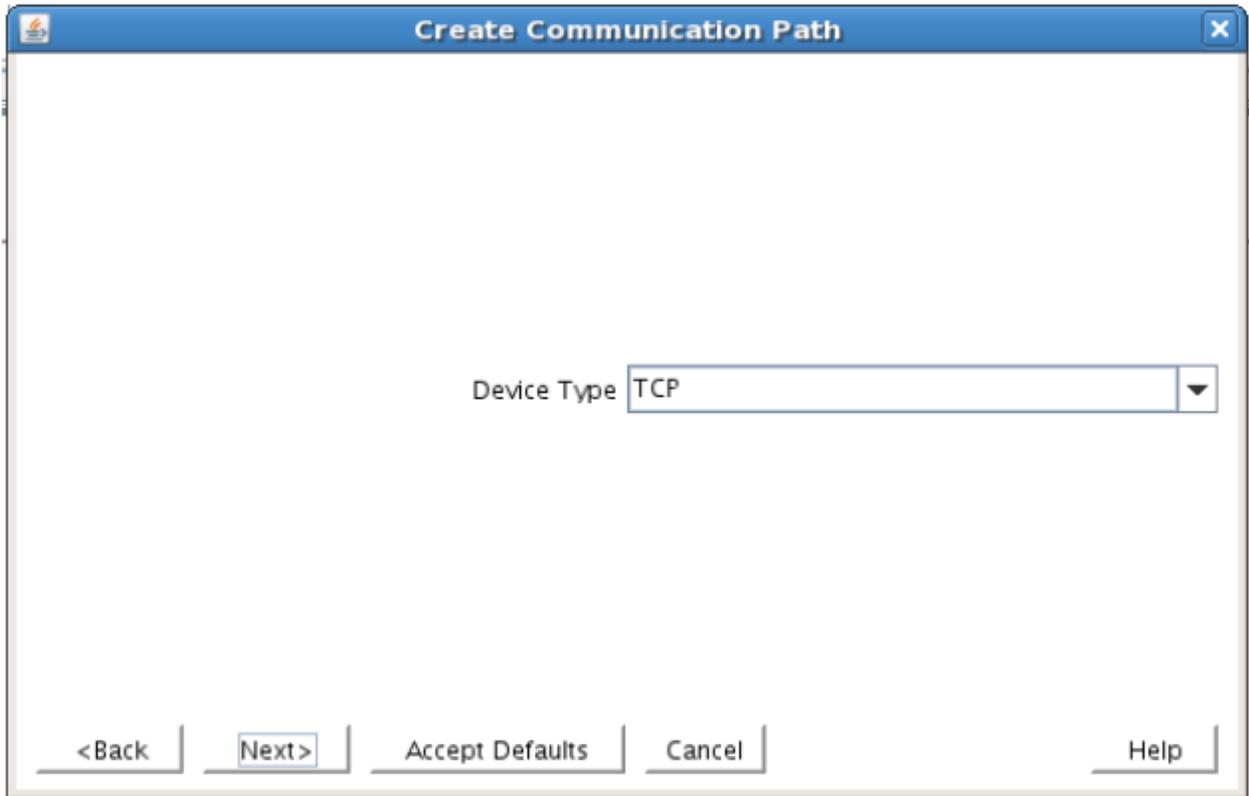
6. Within the LifeKeeper GUI, click the Create Comm Path button on the toolbar. You can also right click one of the servers and click Create Comm Path from the pop-up menu as well.



7. Select your Local and Remote Server(s) from the list box. If a server does not appear in the list box, you may enter it by typing its name and clicking the Add Server button. When using the Add Server procedure, you must make sure that the computer names for both network interfaces on the servers respond correctly when you ping them (from all of the partner server(s)) using the **ping -a IP ADDRESS** syntax. If they do not, this must be corrected prior to continuing. Click Next.

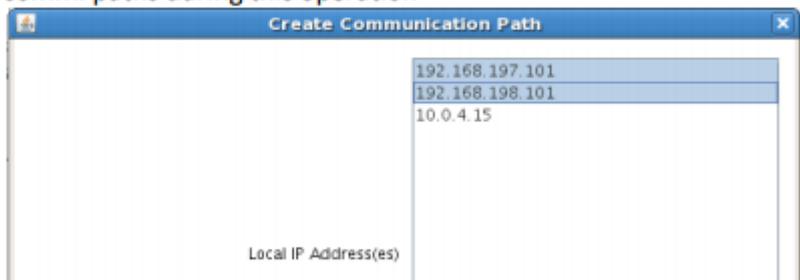


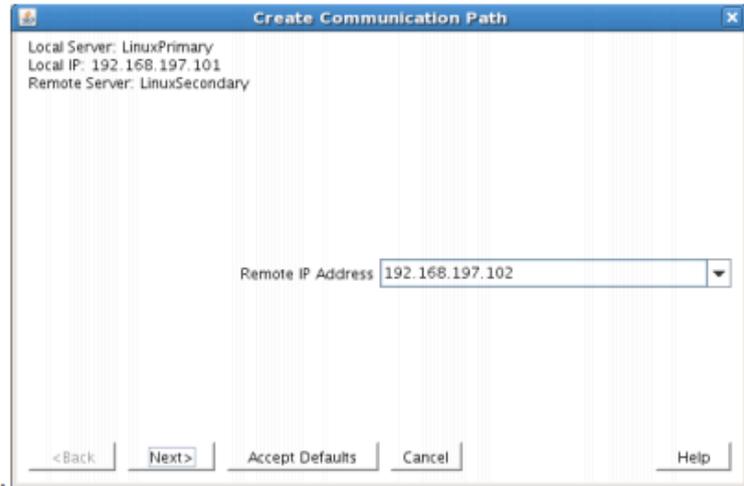
8. Select TCP for Device Type and Click Next.



- 9. Provide all the required information and click Next for the following series of dialog boxes. For each field in the dialog box you can click Help for further information or refer to the table below for an explanation or recommendation

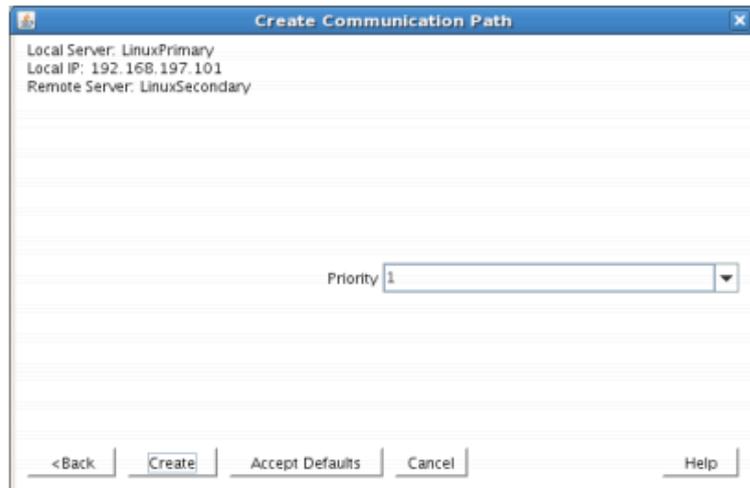
Field	Tips
<b>For TCP/IP Comm Path...</b>	
Local IP Address	Choose the IP address to be used by the local server for this comm path. Select both interfaces so that the wizard creates multiple comm. paths during this operation
Remote IP Address	Choose the IP address to be used by the remote server for this comm path



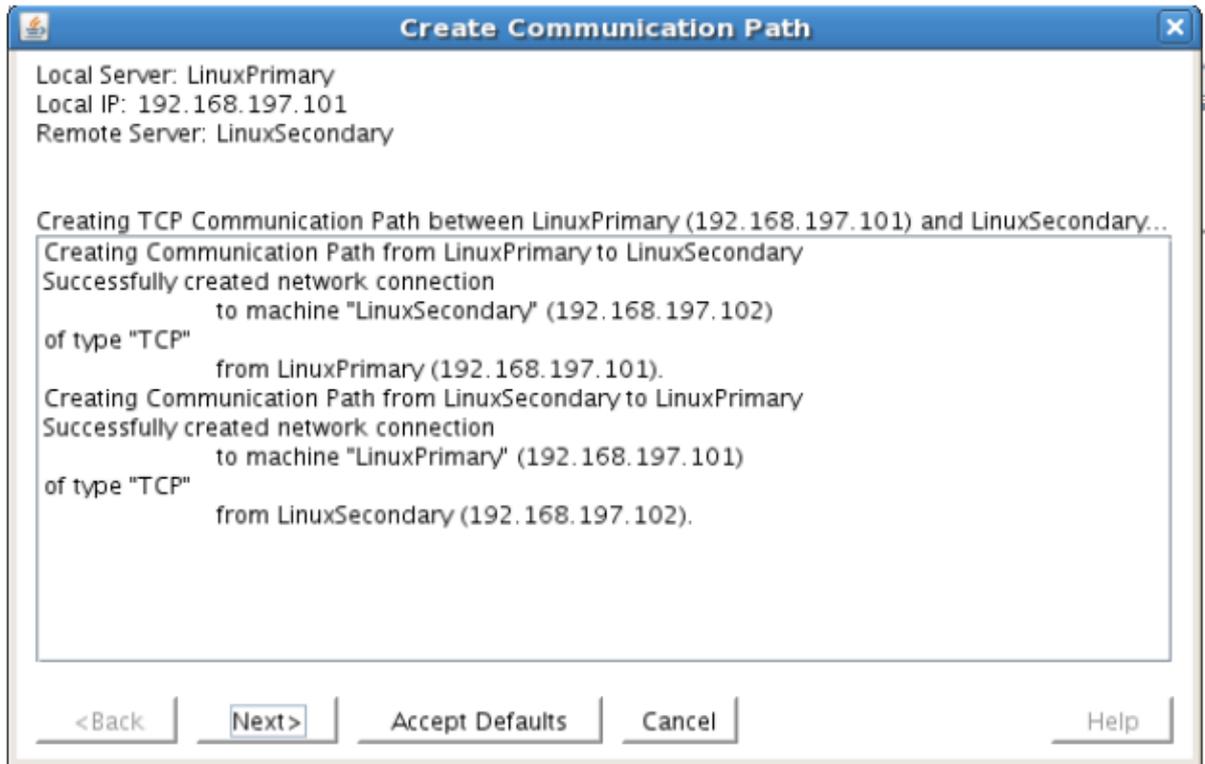


Enter the priority for the comm path on the local server. The priority will be used to determine the order that the comm paths between two servers will be used. Priority 1 is the highest; 99 is the lowest.

Priority



10. After entering data in all the required fields, select Create. A message will display indicating the network communication path is successfully created. Click Next.



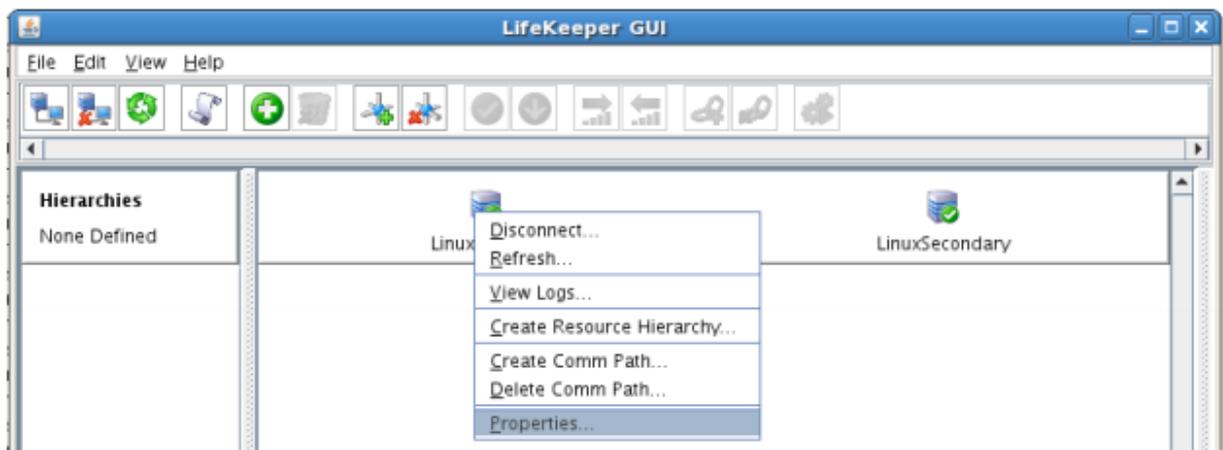
If you selected multiple Local IP Addresses or multiple Remote Servers and the Device Type was set to TCP, then the procedure will return you to the setup wizard the next Comm Path.

11. Select Done in the last dialog box.

Repeat this process until you have defined all the communication paths you plan to use. SIOS strongly recommends that you define at least two communication paths for redundancy.

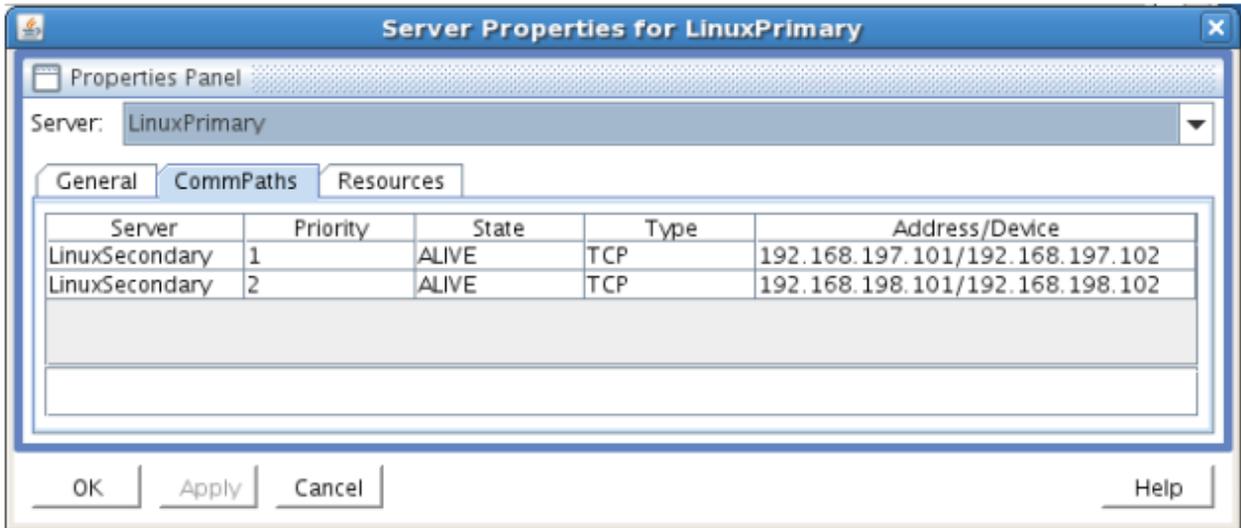
## Verify the Communications Paths

1. Verify that the communications paths are configured properly by viewing the Server Properties dialog box. From the LK GUI, select Edit, Server, Properties and then the Comm Paths tab.



2. Note the State displayed is ALIVE. You can also check the server icon in the right, main pane of

the GUI. If only one comm path has been created, the server icon shows a yellow warning icon on the server icon, indicating that one comm. path is ALIVE, but there is no redundant comm path. The server icon will display a green heartbeat checkmark when there are at least two comm paths configured and ALIVE.



## Create the LifeKeeper Hierarchy

### Create a Mirror and Begin Data Replication

In this section we will setup and configure the Data Replication resource, which be used to synchronize our MySQL's data between cluster nodes. The data we will replicate resides in the /var/lib/mysql partition on our Primary cluster node

Please note:

- The source volume to be replicated must be mounted on the Primary server
- The target volume, which will received replicated data, must NOT be mounted on the Secondary server.
- The target volume's size must equal to or larger than the size of its source volume.

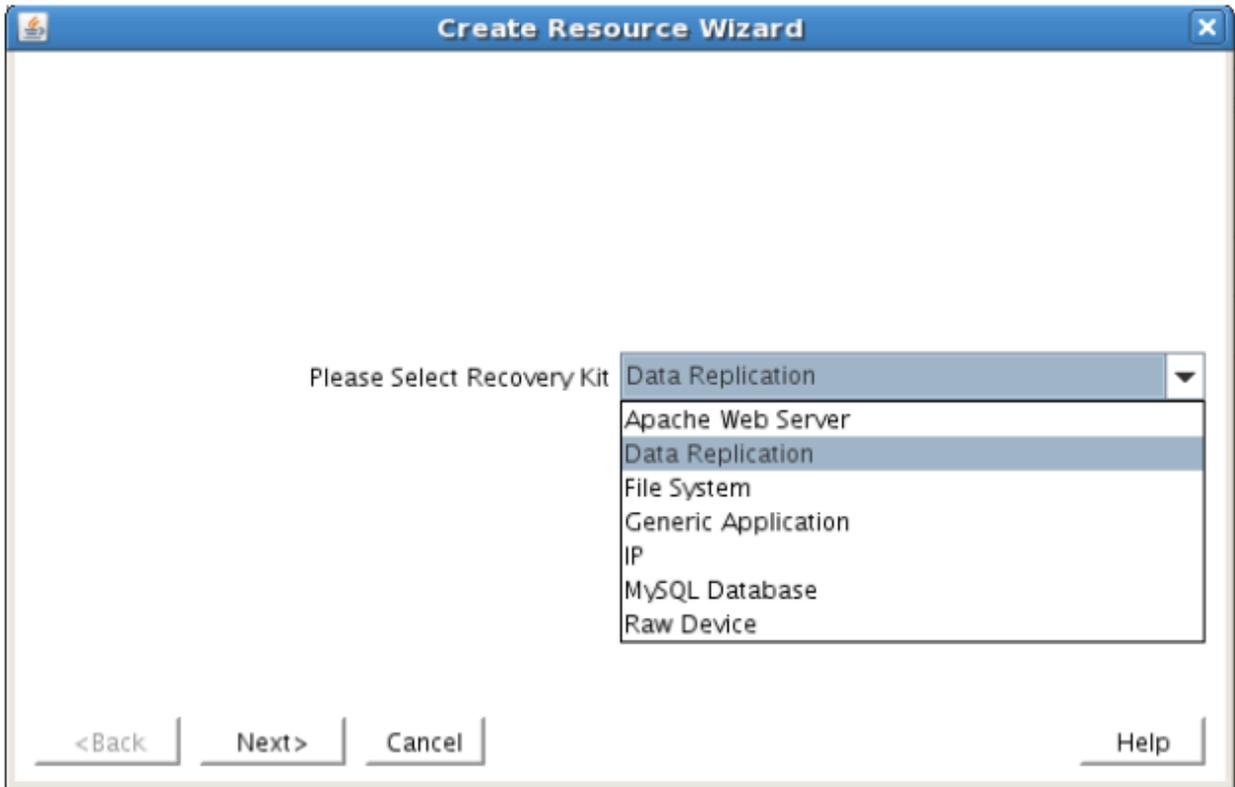
1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all

recognized Recovery Kits installed within the cluster.

2. Select Data Replication and click Next.

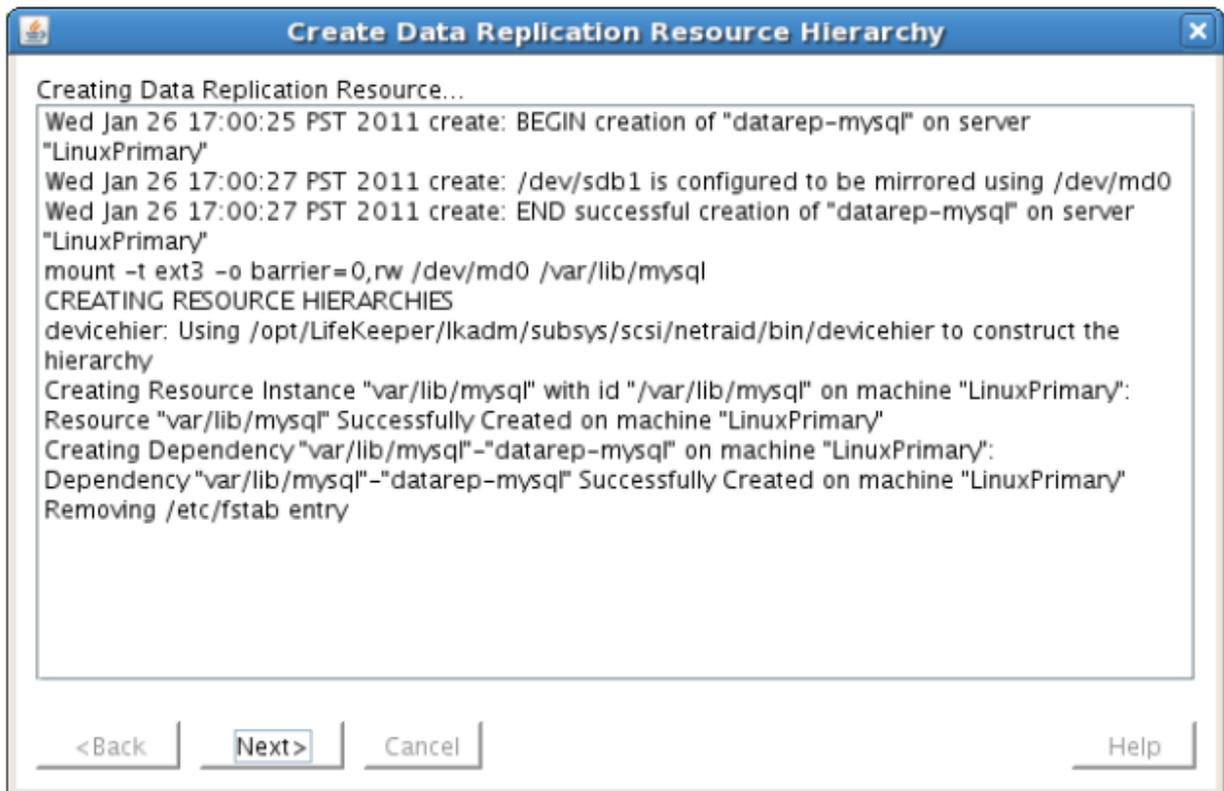


3. Follow the Data Replication wizard, and enter the following values:

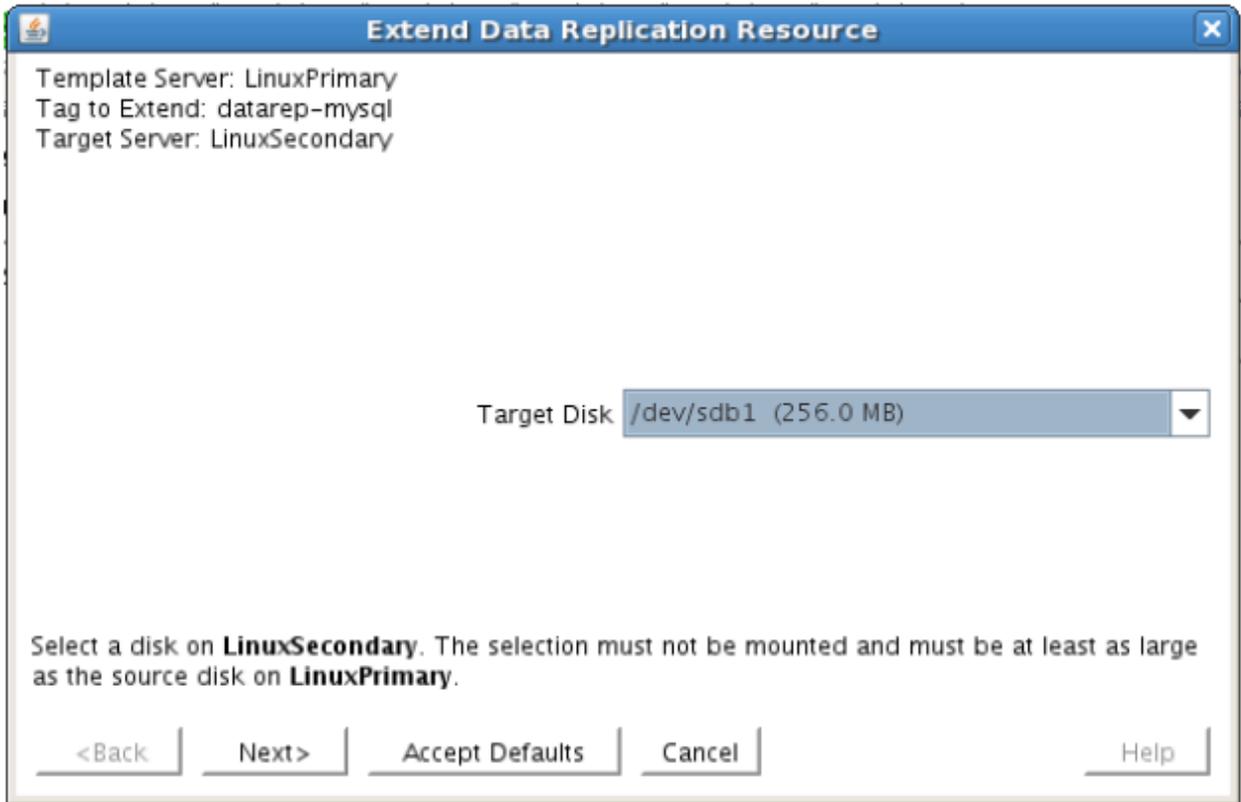
Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node, i.e. Mirror Source)
Hierarchy Type	Select: <b>“Replicate Existing Filesystem”</b>
Existing Mount Point	At this step you will select the mounted partition to replicate. In our example, select <code>“/var/lib/mysql”</code>
Data Replication Resource Tag	Leave as default
File System Resource Tag	Leave as default (Note: if using high speed SSD storage you will want to create a small partition and use it for bitmap placement, i.e. <code>/bitmaps</code> )

Bitmap File	
Enable Asynchronous Replication	Leave as default (Yes)

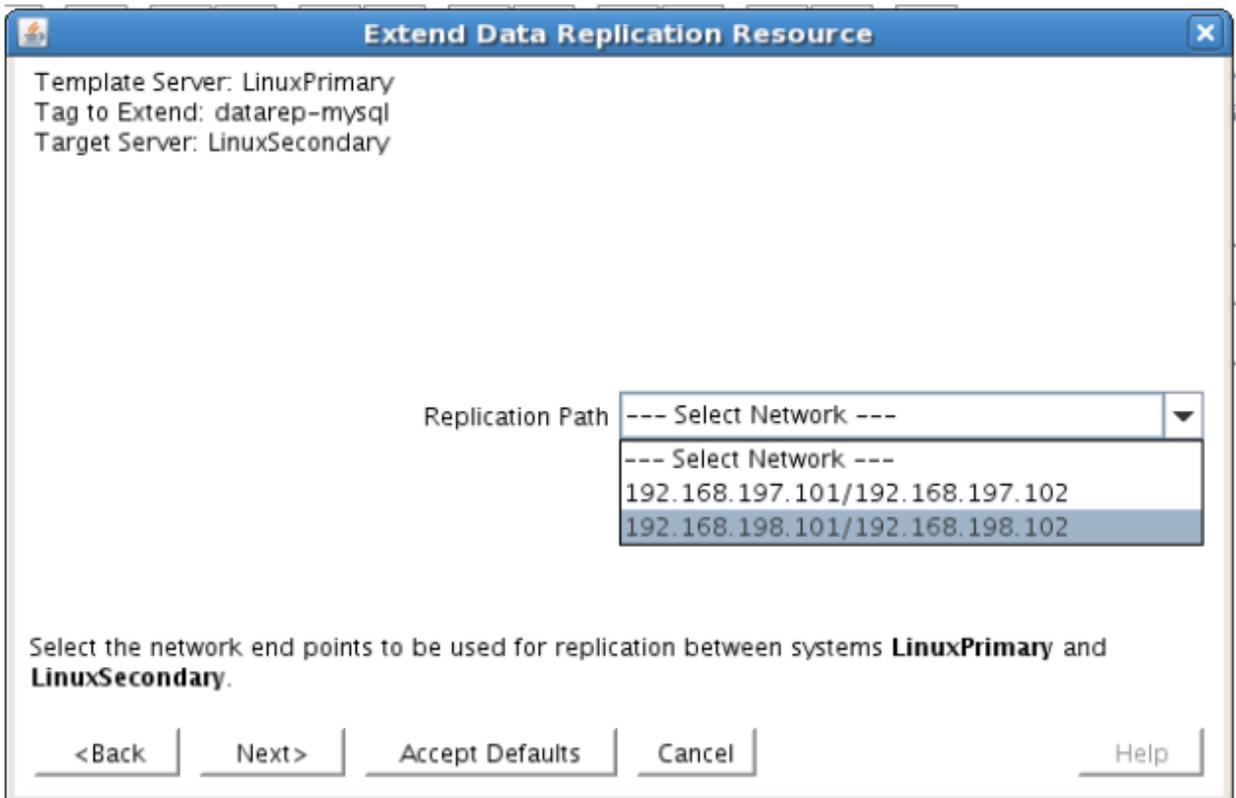
4. Click Next to begin creation of the Data Replication resource hierarchy. Status will be displayed in the GUI as follows:



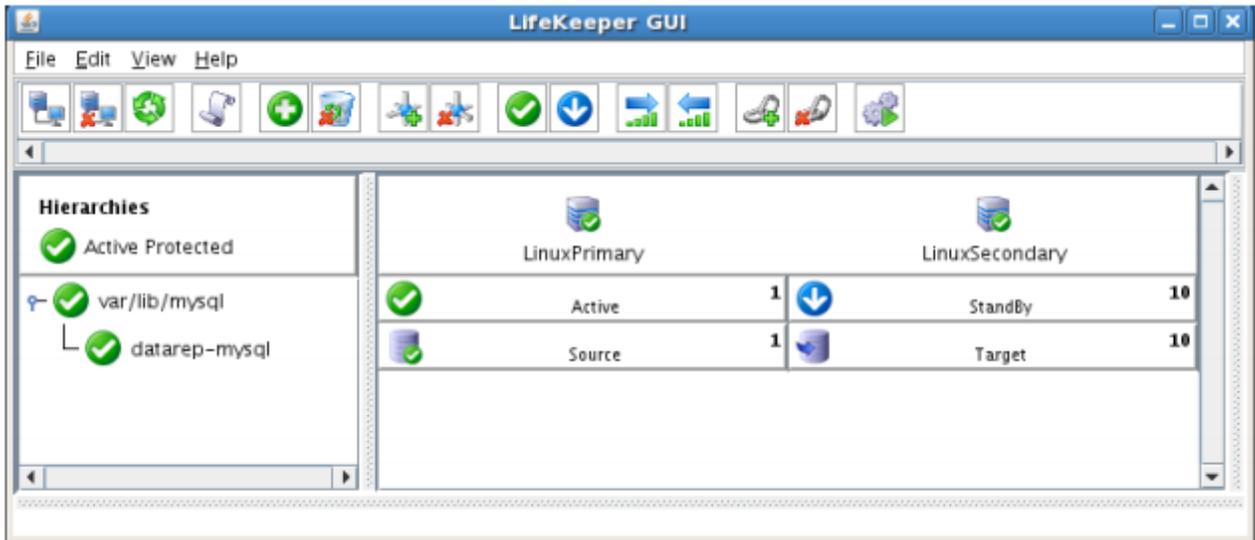
5. Click Next to begin the process to Extend the Data Replication Resource. Select all default settings. When it asks for the target disk, select a free partition on your Target server which is the same size (or greater) than the Source Volume we are replicating. This partition should NOT be mounted on the Target system.



- Continue through the wizard, and you will be prompted to select the network you would like replication to take place over. In general, it's a best practice to separate your user/application and your replication traffic. In our example setup we will replicate over our backend network, 192.168.198.X



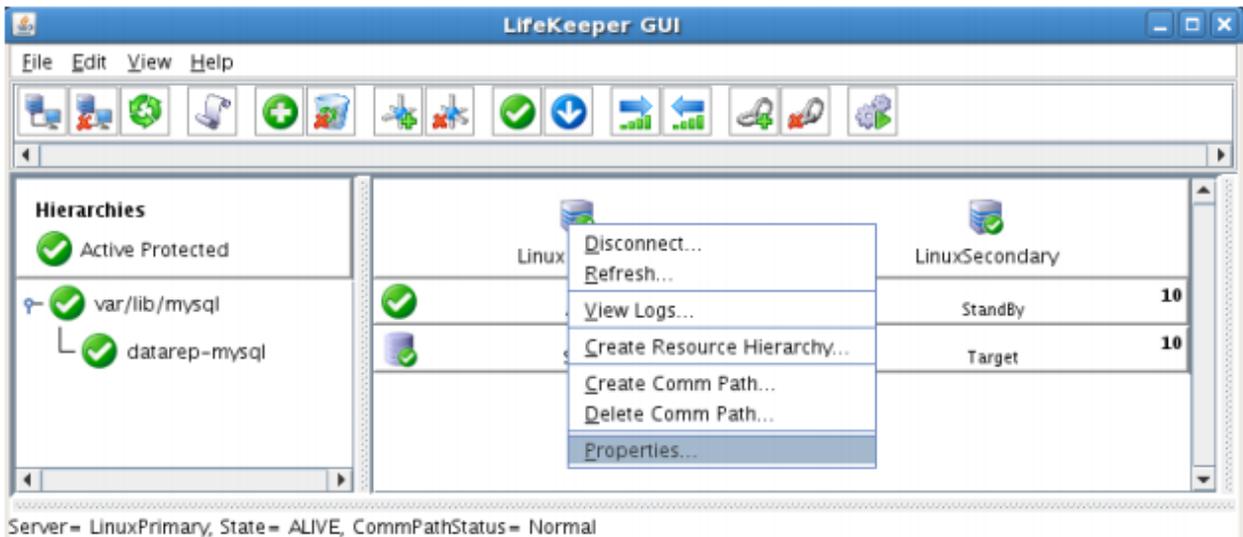
- Click Next and continue through the wizard. Once completed, your resource hierarchy will look as follows



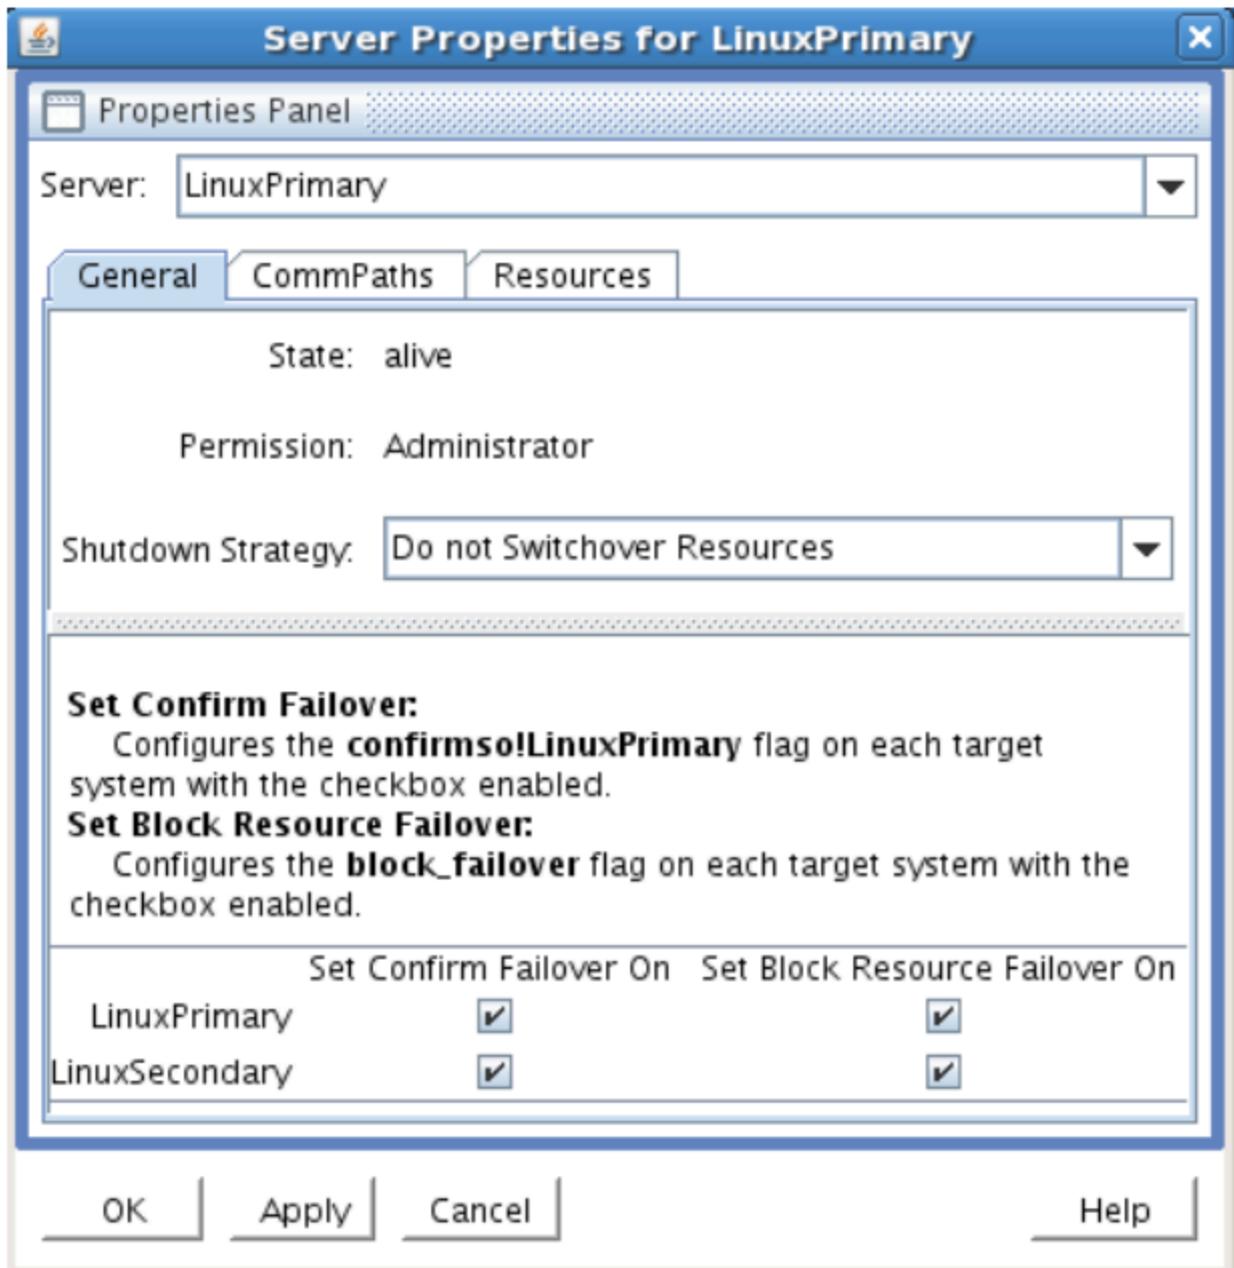
### Disable Automatic Failover

In this section we will review the procedure for disabling automatic failover to the standby server.

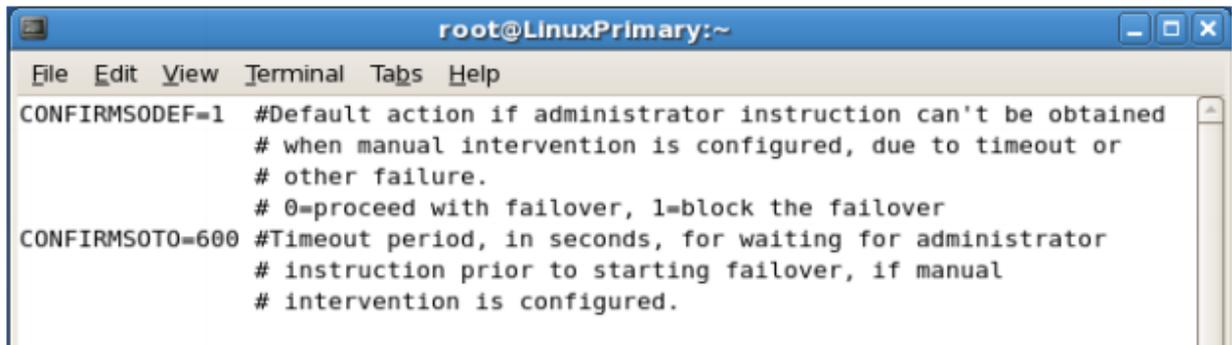
- From the LifeKeeper GUI, right click on one of the cluster nodes and select Properties.



- Select the Source server from the “Server:” drop down at the top of the window
- Once the Server Properties window loads, check all boxes at the bottom of the page. This will prevent any automatic failovers from happening.



4. Click Apply
5. Repeat steps 2-4, this time selecting the Target server from the "Server:" drop down
6. Next, edit /etc/default/LifeKeeper on both nodes
  - a. Set CONFIRMSODEF=1 (change from 0 to 1)

A terminal window titled "root@LinuxPrimary:~" with a menu bar containing "File", "Edit", "View", "Terminal", "Tabs", and "Help". The terminal displays two configuration parameters: "CONFIRMSODEF=1" and "CONFIRMSOTO=600", each followed by a multi-line comment explaining their function in failover scenarios.

```
CONFIRMSODEF=1 #Default action if administrator instruction can't be obtained
                # when manual intervention is configured, due to timeout or
                # other failure.
                # 0=proceed with failover, 1=block the failover
CONFIRMSOTO=600 #Timeout period, in seconds, for waiting for administrator
                # instruction prior to starting failover, if manual
                # intervention is configured.
```

## 11.1.7. Test Your DK for Linux Environment

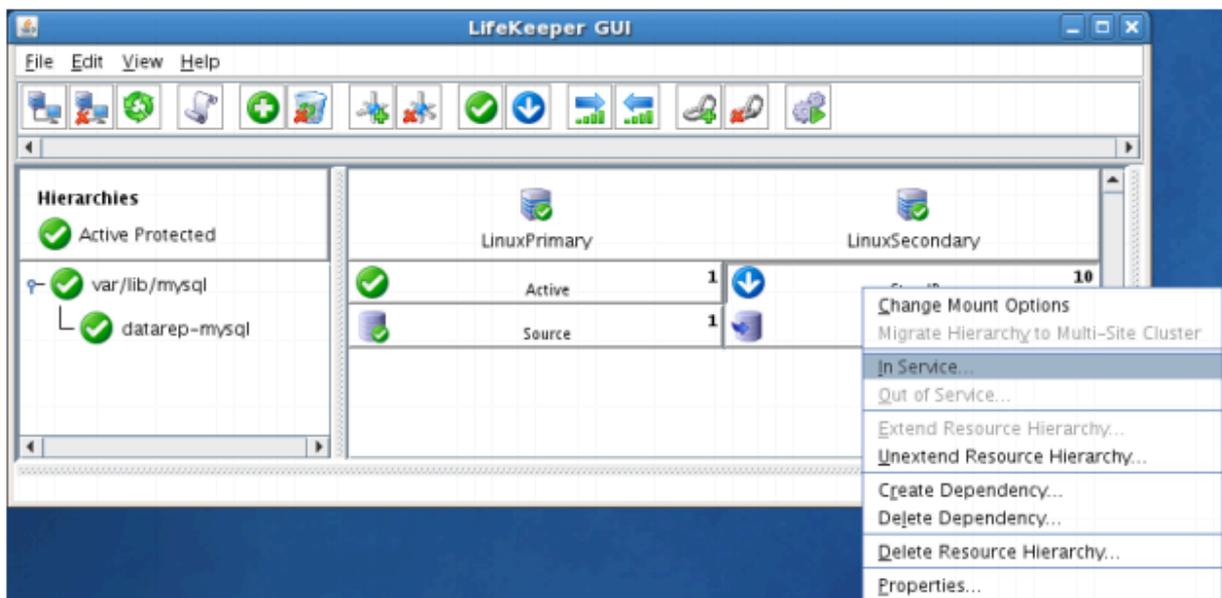
The following test scenarios have been included to guide you as you get started evaluating LifeKeeper for Linux. **Before beginning these tests, make sure the data replication resources are in the mirroring state.**

\* **Note:** For these test examples, the Primary Server is referred to as LINUXPRIMARY. The Backup or Secondary Server is referred to as LINUXSECONDARY.

### 1. Manual Switchover of the Mirror to Secondary Server

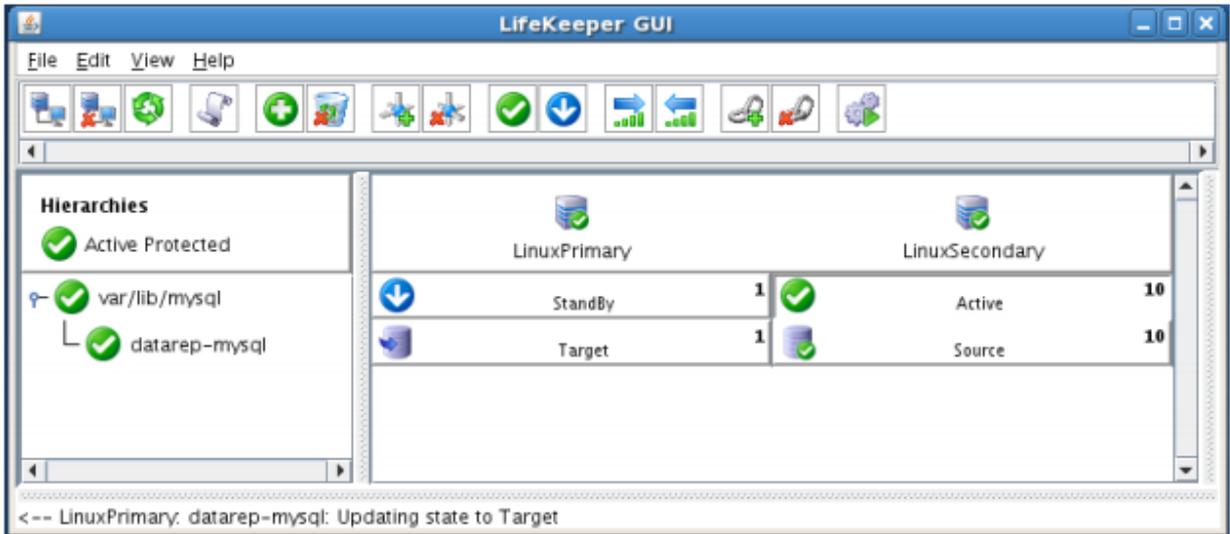
#### Procedure:

- From the LifeKeeper GUI, right click on the top of the resource hierarchy on the Secondary Server (LINUXSECONDARY) and choose IN SERVICE.
- Click "In Service" in the window that pops up



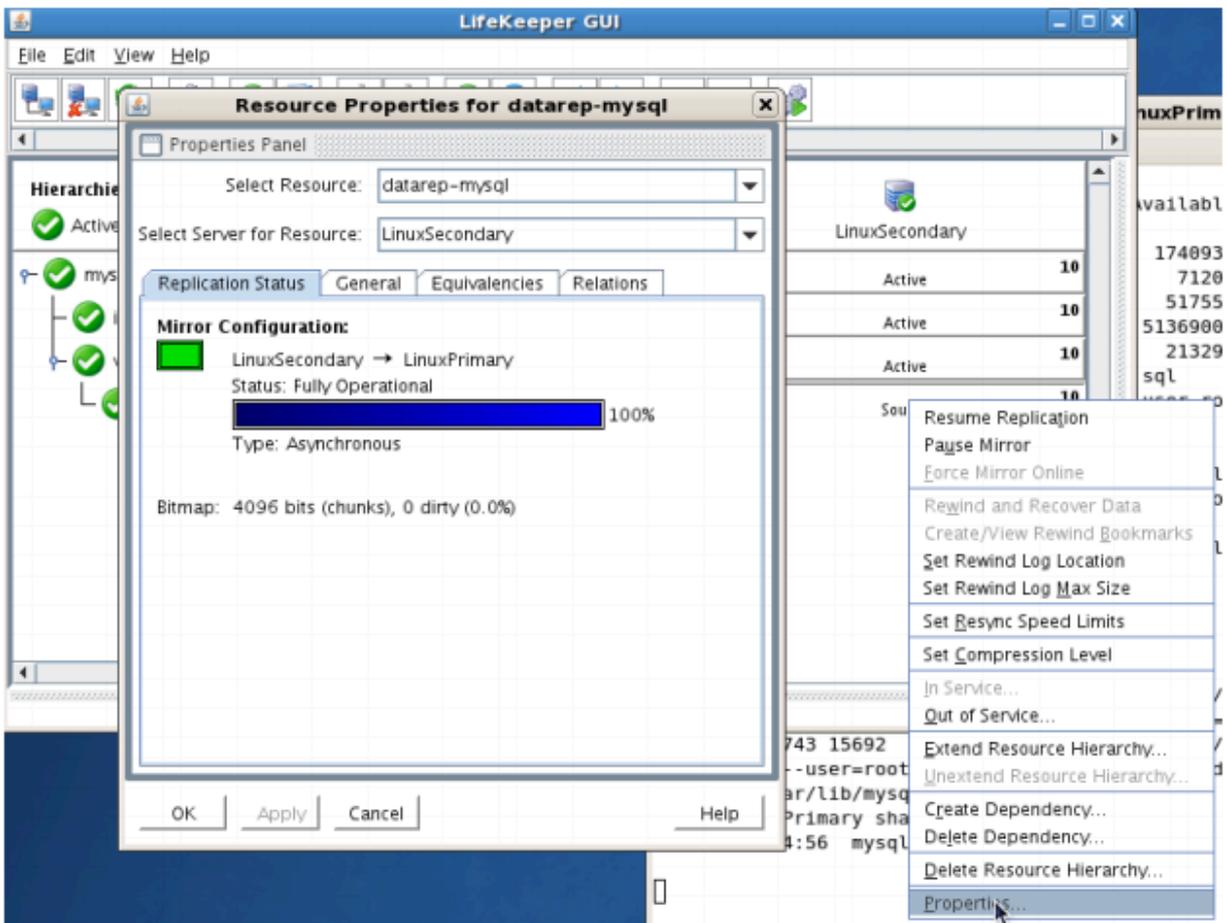
#### Expected Result:

- All resources will be removed from service on the Active Server (LINUXPRIMARY).
- Beginning with the dependent resources, all resources will be brought in service on LINUXSECONDARY.
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXSECONDARY -> LINUXPRIMARY
- At this point, all resources are now active on LINUXSECONDARY.



**Tests/Verification:**

- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the “datarep-mysql” resource and select Properties



- Run “df -h” to verify that the /var/lib/mysql replicated filesystem is mounted as an “md” device (example: /dev/md0”) on LINUXSECONDARY

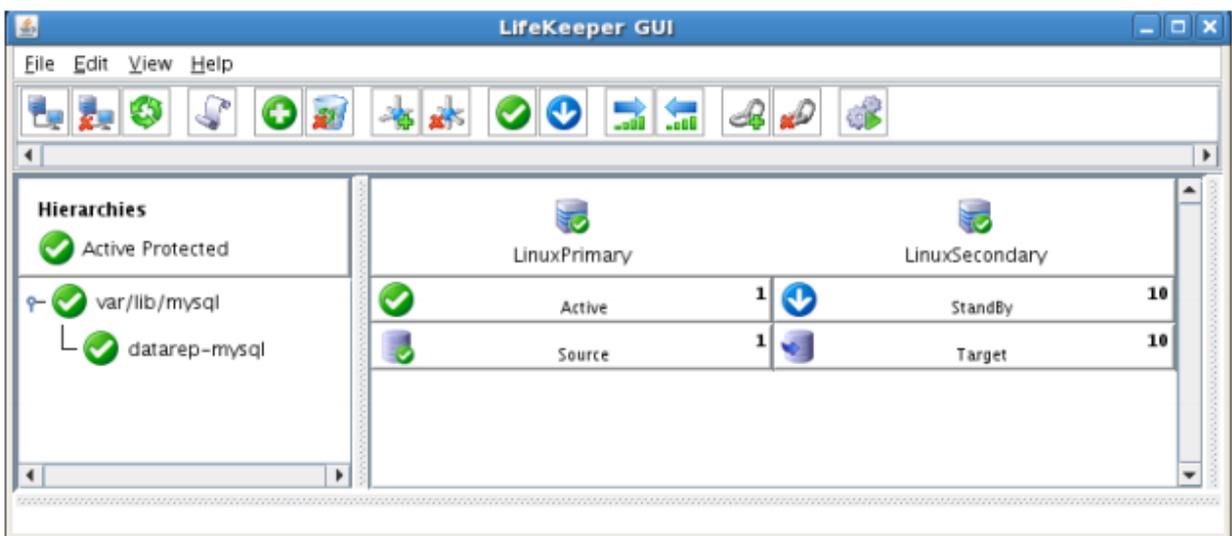
2. Manual Switchover of the Mirror back to Primary Server

**Procedure:**

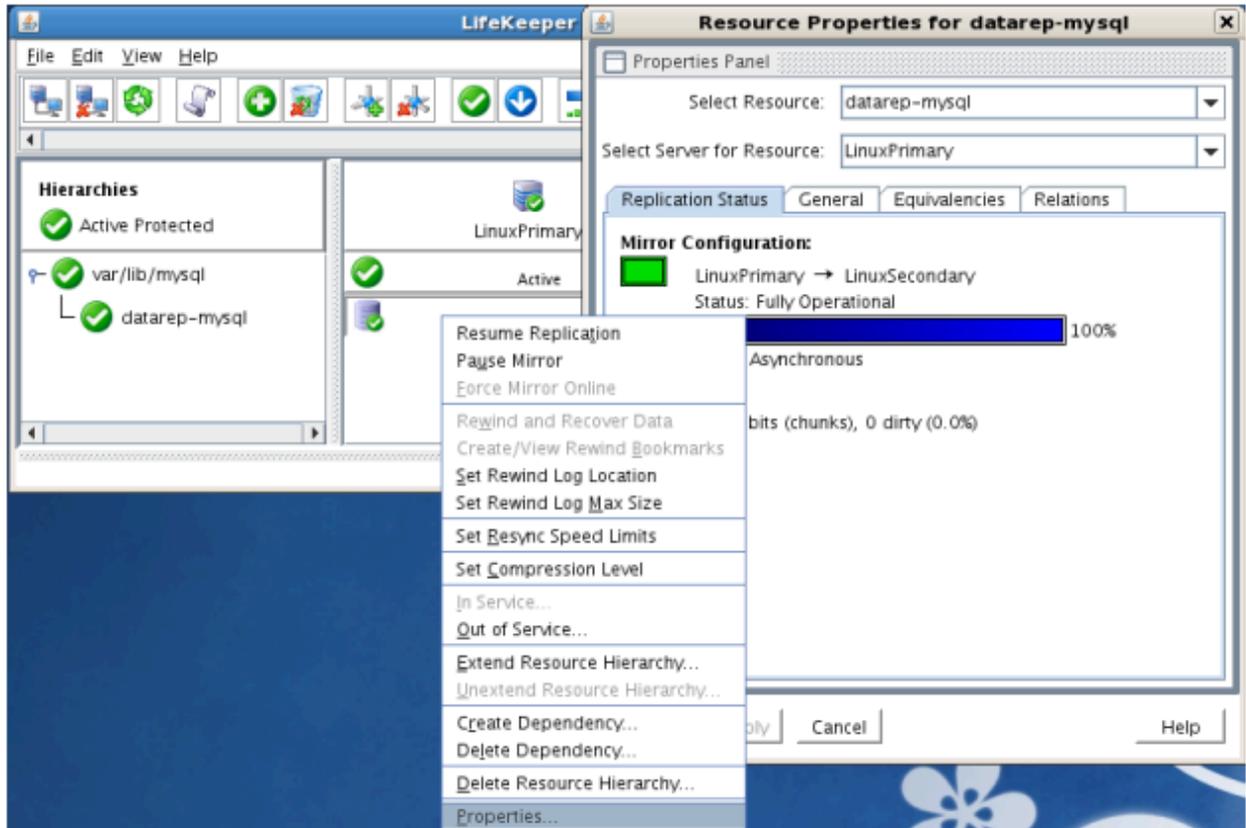
- From the LifeKeeper GUI, right click on the top level of the resource hierarchy on the Primary Server (LINUXPRIMARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up

**Expected Result:**

- All resources will be removed from service on the Active Server (LINUXSECONDARY).
- Beginning with the dependent resources, all resources will be brought in service on LINUXPRIMARY
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXPRIMARY -> LINUXSECONDARY

**Tests/Verification:**

- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the “datarep-mysql” resource and select Properties



- Run “df -h” to verify that the /var/lib/mysql replicated filesystem is mounted as an “md” device (example: /dev/md0) on LINUXPRIMARY

## 11.2. LifeKeeper Evaluation Guide for Cloud Environments

---

This Evaluation Guide explains the basic concepts of the LifeKeeper and then walks through core user interface elements to show how various resources may be protected. The guide discusses the cloud environment from several perspectives and may also be used to evaluate the LifeKeeper in an on-premise environment.

## 11.2.1. Before Starting an Evaluation of LifeKeeper for Linux

---

This section discusses general topics to consider when evaluating software bringing High Availability to mission-critical workloads.

- [High Availability, RTO, and RPO](#)
- [LifeKeeper for Linux – Integrated Components](#)
- [Benefits of LifeKeeper for Linux](#)
- [How Workloads Should be Distributed when Migrating to a Cloud Environment](#)
- [Public Cloud Platforms and their Network Structure Differences](#)
- [How a Client Connects to the Active Node](#)
- [How does Data Replication between Nodes Work?](#)
- [What is “Split Brain” and How to Avoid It](#)

## 11.2.1.1. High Availability, RTO, and RPO

---

High availability (HA) is an information technology term that refers to a computer software or component that is operational and available for more than 99.99% of the time. End users of an application, or system, experience less than 52.5 minutes per year of service interruption. This level of availability is typically achieved through the use of high availability clustering, a configuration that reduces application downtime by eliminating single points-of-failure through the use of redundant servers, networks, storage, and software.

### What are recovery time objectives (RTO) and recovery point objectives (RPO)?

In addition to 99.99% availability time, high availability environments also meet stringent recovery time and recovery point objectives. Recovery time objective (RTO) is a measure of the time elapsed from application failure to restoration of application operation and availability. It is a measure of how long a company can afford to have that application down. Recovery point objectives (RPO) are a measure of how up-to-date the data is when application availability has been restored after a downtime issue. It is often described as the maximum amount of data loss that can be tolerated when a failure happens. SIOS high availability clusters deliver an RPO of zero and an RTO of minutes.

### What is a high availability cluster?

In a high availability cluster, important applications are run on a primary server node, which is connected to one or more secondary nodes for redundancy. Clustering software, such as SIOS LifeKeeper, monitors clustered applications and dependent resources to ensure they are operational on the active node. System level monitoring is accomplished via intervallic heartbeats between cluster nodes. If the primary server fails, the secondary server initiates recovery after the heartbeat timeout interval is exceeded. For application level failures, the clustering software detects that an application is not available on the active node. It then moves the application and dependent resources to the secondary node(s) in a process called a failover, where operation continues and meets stringent RTOs.

In a traditional failover cluster, all nodes in the cluster are connected to the same shared storage, typically a storage area network (SAN). After a failover, the secondary node is granted access to the shared storage, enabling it to meet stringent RPOs.

## 11.2.1.2. LifeKeeper for Linux – Integrated Components

---

LifeKeeper includes the following software components to protect an organization's mission-critical systems.

### SIOS LifeKeeper

SIOS LifeKeeper provides a complete fault-resilient software solution that enables high availability for servers, file systems, applications, and processes. LifeKeeper does not require any customized, fault-tolerant hardware. LifeKeeper simply requires two or more systems to be grouped in a network and site-specific configuration data is then created to provide automatic fault detection and recovery.

In the case of a failure, LifeKeeper migrates protected resources from the failed server to a designated standby server. Users experience a brief interruption during the actual switchover; however, LifeKeeper restores operations on the standby server without operator intervention.

### SIOS DataKeeper

SIOS DataKeeper provides an integrated data replication capability for LifeKeeper environments. This feature enables LifeKeeper resources to operate in shared and non-shared storage environments.

### Application Recovery Kits (ARKs)

Application Recovery Kits (ARKs) include tools and utilities that allow LifeKeeper to manage and control a specific application or service. When an ARK is installed for a specific application, LifeKeeper is able to monitor the health of the application and automatically recover the application if it fails. These Recovery Kits are non-intrusive and require no changes within the application in order for it to be protected by LifeKeeper.

There is a comprehensive library of 'off-the-shelf' Application Recovery Kits available as part of the LifeKeeper portfolio. The types and quantity of ARKs supplied vary based on the edition of LifeKeeper purchased.

## 11.2.1.3. Benefits of LifeKeeper for Linux

---

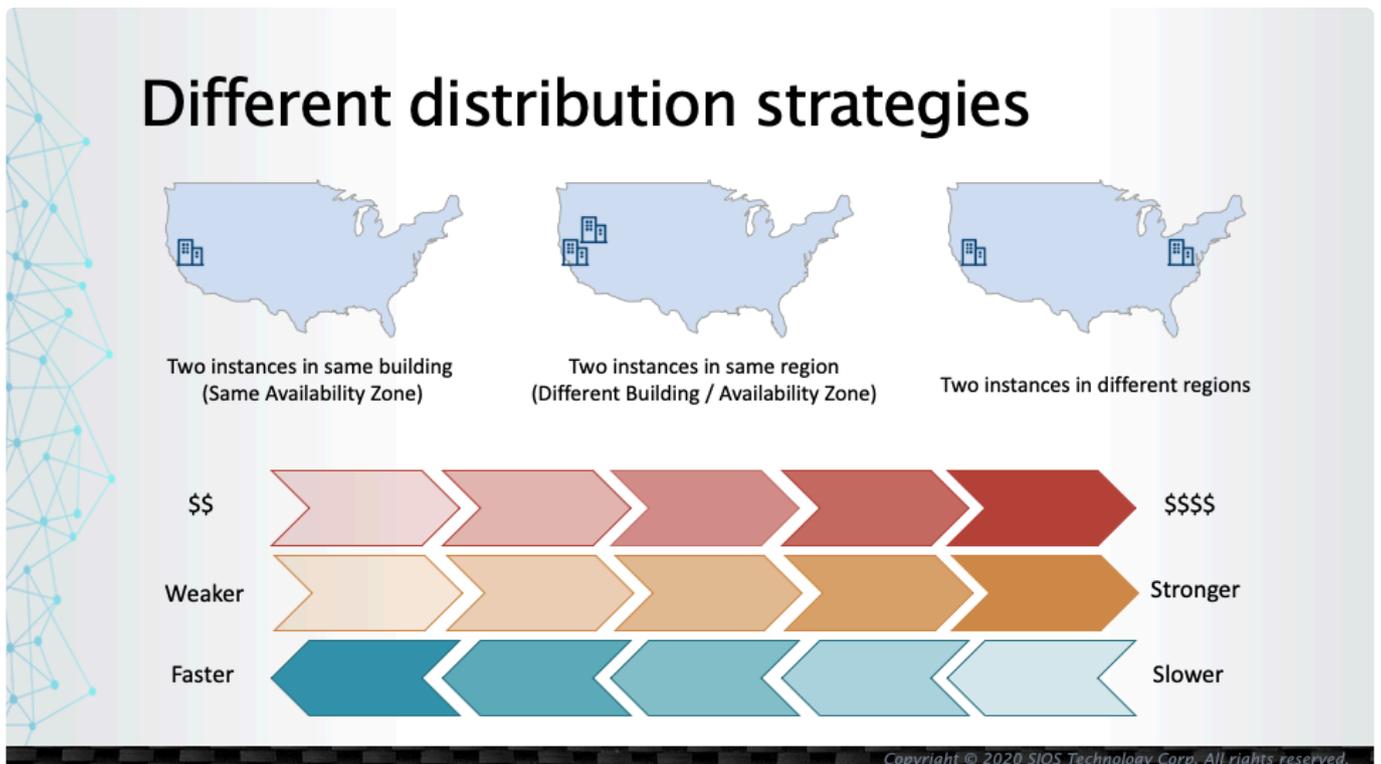
- SIOS software supports multiple operating system versions and flavors (both Linux and Windows).
  - Consistent user experience to protect mission-critical resources regardless of the operating system.
- SIOS delivers high availability solutions for multiple environments including on-premises, virtualized (VMware) and cloud environments including public platforms from AWS, Azure, Google Cloud as well as privately hosted cloud environments. The same tools can be used across different environments.
  - Hybrid environments are also supported. It is possible to configure an on-premise node as the primary location with the secondary node being hosted in a private cloud and a third node located on a public cloud platform providing an additional DR (disaster recovery) option.
- SIOS provides application-aware protection mechanisms including Application Recovery Kits (ARKs), available for the world's leading providers of enterprise applications and databases.
  - Our comprehensive library of ARKs protect the broadest range of applications 'off-the-shelf'.
  - Other 'non-standard' or legacy applications can be protected with the built-in 'GenApp' ARK or by developing a custom ARK either in-house or in collaboration with SIOS engineers.
- Unlike open-source tools (which force the user to manually establish environmental parameters in advance and then type a complex set of command-line parameters), a series of wizard based installation and configuration screens enables intuitive selection of resources requiring protection and helps in selecting the type and extent of protection to be provisioned.
  - Wizards scan the system and environment to identify the resources to protect. In most cases the operator only needs to confirm default selections and enter unique environment-specific host IDs etc. to complete the installation.
  - This reduces the likelihood of misconfiguration of the HA solution and the inevitable unexpected downtime as the result of a system or application failure.
- SIOS provides a comprehensive range of technical support options including 24x7 critical support, offering support options tailored to the available budget, systems complexity or criticality of the application requiring high availability.

# 11.2.1.4. How Workloads Should be Distributed when Migrating to a Cloud Environment

Determining how Workloads (nodes) should be distributed is a common topic of discussion when migrating to the public cloud with High Availability in mind. If workloads are located within an on-premise environment, more often than not the locations of these workloads are defined by the location(s) of established datacenters. In many cases choosing another location in which to host a workload is not an available option. With a public cloud offering there are a wide range of geographical regions as well as availability zones to choose from.

An Availability Zone is generally analogous to one or more datacenters (physical locations) being located in the same physical region (e.g., in California). These datacenters may be located in different areas but are connected using high-speed networks to minimize connection latency between them. (Note that hosting services across several datacenters within an availability region should be transparent to the user).

As a general rule, the greater the physical distance between workloads, the more resilient the environment becomes. It's a reasonable assumption that natural disasters such as earthquakes won't affect different regions at the same time (e.g., both U.S. west coast and east coast at the same time). However, there is still a chance of experiencing service outages across different regions simultaneously due to system-wide failures (some cloud providers have previously reported simultaneous cross-region outages such as in the US & Australia). It may be appropriate to consider creating a DR (disaster recovery) plan defined across different cloud providers.



Another perspective worthy of consideration is the cost to protect the resources. Generally the greater

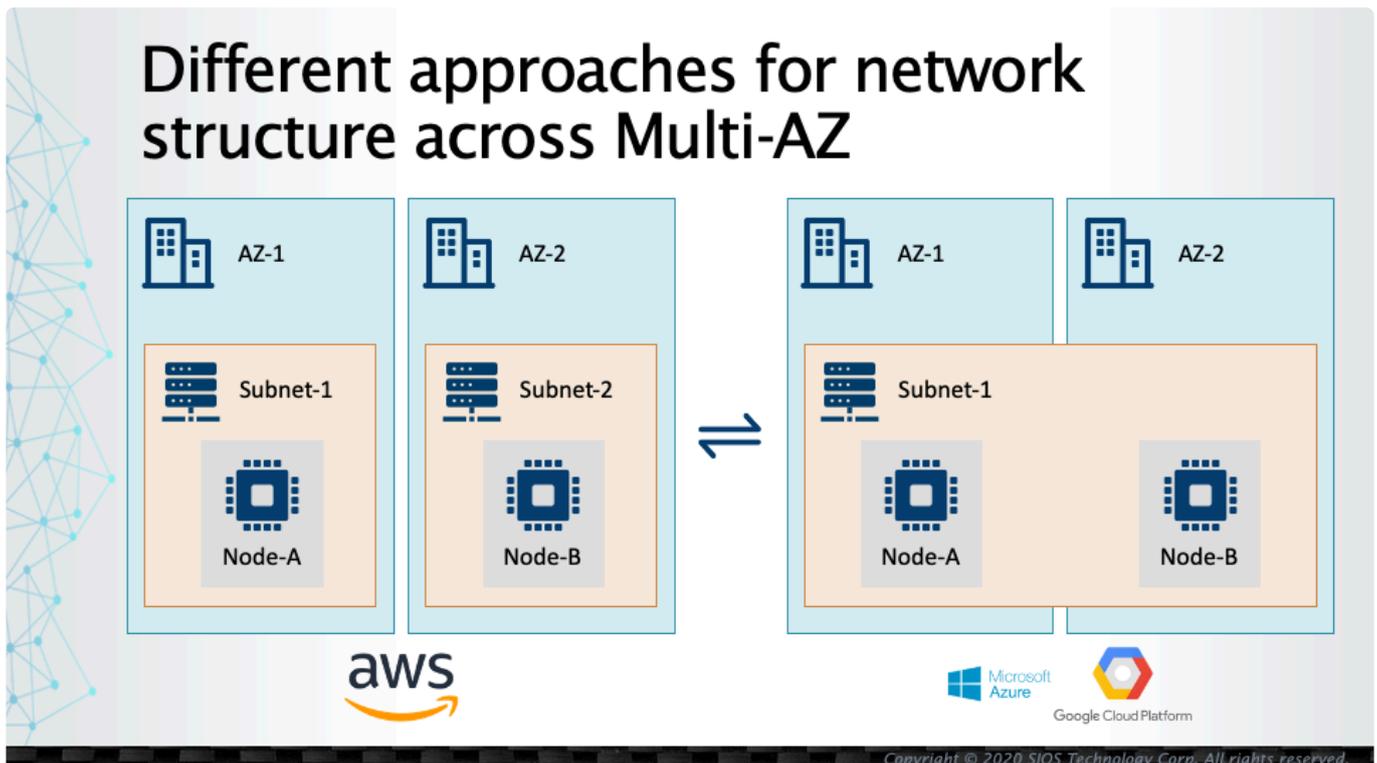
the distance between workloads, the more costs are incurred for data transfer. In many cases, data transfer between nodes within the same datacenter (Availability Zone) is free while it might cost \$0.01/GB or more to transfer data across Availability Zones. This additional cost might be doubled (or more) when data is transferred across regions (i.e. \$0.02 / GB). In addition, due to the increased physical distance between workloads, greater data latency between nodes should be anticipated between locations. Through consideration of these factors, generally speaking, it is recommended to distribute workloads across Availability Zones within the same Region.

## 11.2.1.5. Public Cloud Platforms and their Network Structure Differences

There are several public cloud platforms including Amazon Web Services (AWS), Microsoft Azure and Google Cloud. While there are many similarities in their infrastructures, there are some differences. In many cases a VPC (Virtual Private Cloud) or a VNET (Virtual Network) that is tied to a region is created. One or more VPCs may be defined for a logical group of applications. By so doing, different systems are divided into separate unconnected networks unless different VPCs are specifically connected.

Under a VPC many different subnets can be defined. Based on the purpose, some subnets are configured as “public” subnets which are accessible to the internet and some are configured as “private” subnets which are not accessible to the internet.

Some cloud providers (such as Azure and Google Cloud) allow subnets to be defined across Availability Zones (different datacenters), while some (such as AWS) do not allow subnets to be defined across Availability Zones. In the latter case, a subnet will need to be defined for each Availability Zone.



In this guide, we’ll use different Availability Zones for each node. Once the basic functionality of the SIOS product is understood, it might be appropriate to explore different scenarios (similar to those in use in your own network infrastructure) that involve distributing workloads across different subnets, modifying the IP ranges for these subnets, changing the manner in which the network is connected to the Internet, etc.

## 11.2.1.6. How a Client Connects to the Active Node

As discussed earlier, once a High Availability cluster has been configured, two or more nodes run simultaneously and users connect to the “active” node. When an issue occurs on the active node, a “failover” condition occurs and the “standby” node becomes the new “active” node. When a failover occurs there must be a mechanism that either allows a client to detect the failover condition and to reconnect, or a seamless transfer of the user’s active client session to the active node.

### A Virtual IP Address

Usually a “virtual” IP address is created when a cluster is configured and the client communicates with the active node using a virtual IP address. When a failover occurs, the virtual IP address is reassigned to the new active node and the client reconnects to the same virtual IP address.

As an example, let us assume that there are two nodes, A and B, with IP addresses of 10.20.1.10 and 10.20.2.10. In this example, we will define a virtual IP address of 10.20.0.10 which should be considered to be assigned to the current active node.

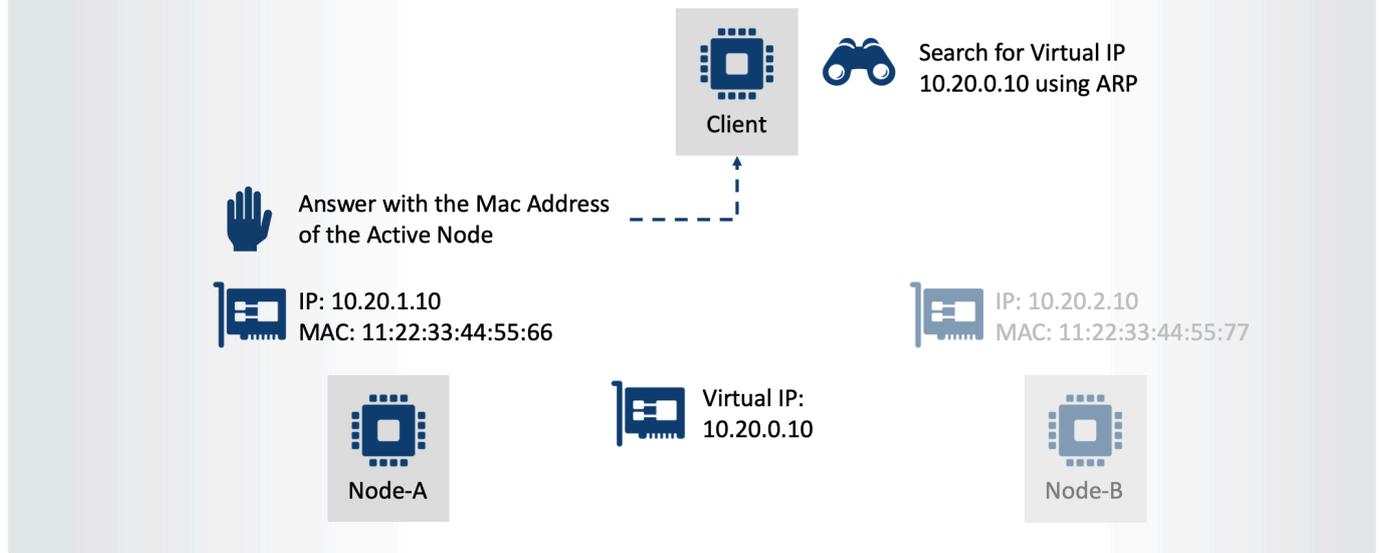
This is similar to assigning a second IP address to one network interface card on one node. If the command `ip a` is entered on the active node, both IP addresses will appear (as on lines 10 and 12 in this Linux example):

```
1 [root@node-a default]# ip a
2 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
3     link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
4     inet 127.0.0.1/8 scope host lo
5         valid_lft forever preferred_lft forever
6     inet6 ::1/128 scope host
7         valid_lft forever preferred_lft forever
8 2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
9     link/ether 00:0d:3a:6e:1f:74 brd ff:ff:ff:ff:ff:ff
10    inet 10.20.1.10/22 brd 10.20.3.255 scope global noprefixroute eth0
11        valid_lft forever preferred_lft forever
12    inet 10.20.0.10/22 scope global secondary eth0
13        valid_lft forever preferred_lft forever
14    inet6 fe80::20d:3aff:fe6e:1f74/64 scope link
15        valid_lft forever preferred_lft forever
```

### The ARP Protocol

When a client attempts to find a server using an IP address, the client typically uses ARP (Address Resolution Protocol) to find the MAC (Media Access Control) address of the target machine.

# How does a client find the active Node?



Once a client broadcasts a message to find the target IP address, the active node answers with its MAC address and the client resolves the request and connects to it.

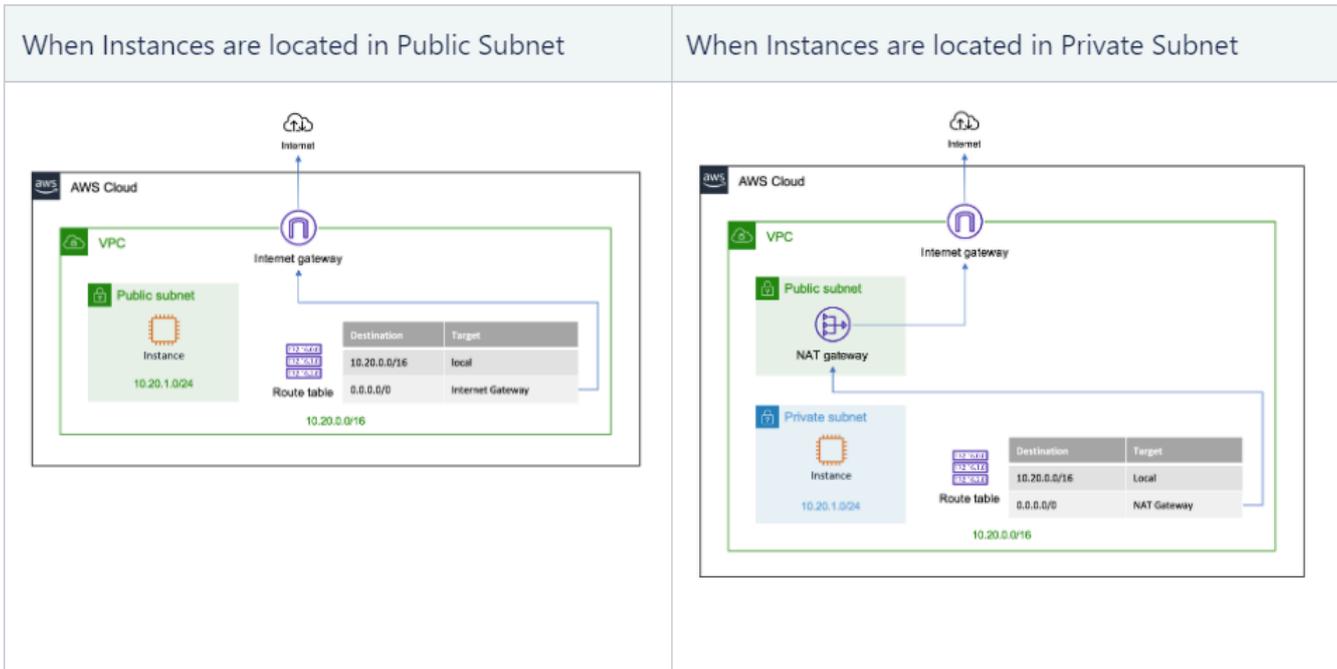
## ARP Alternatives for the Cloud Environment

In the cloud environment, however, it is not possible to identify the active node using ARP as many layers are abstracted in the virtual environment. An alternative method based on the network infrastructure in use in the specific cloud environment may be required. There are normally several options, and a selection should be made from the following list.

- [AWS Route Table Scenario](#)
- [AWS Elastic IP Scenario](#)
- [AWS Route53 Scenario](#)
- [Azure Internal Load Balancer Scenario](#)
- [Google Cloud Internal Load Balancer Scenario](#)

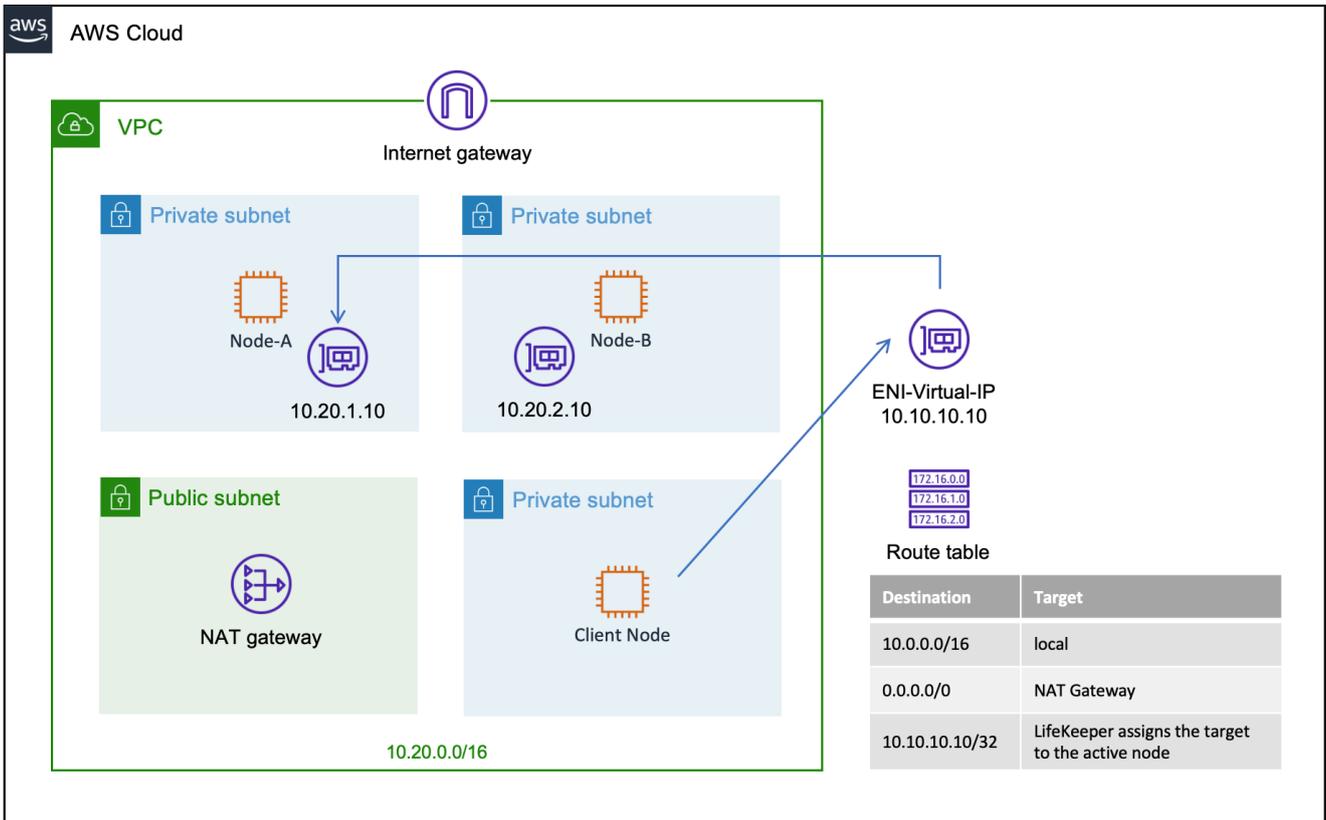
# 11.2.1.6.1. AWS Route Table Scenario

When an instance is created on AWS, one of the first steps is to create a Route Table. The Route Table defines the route for each Destination. In the following diagram the Instance has two routes, one for 10.20.0.0/16, which is the VPC boundary (local network). Any traffic within 10.20.0.0/16 is considered a local connection. The other route is for 0.0.0.0/0 which will allow IP traffic to route outside the local VPC boundary (to other VPCs or the Internet). To connect to the internet either an Internet Gateway is needed (in the case of a public subnet) or a NAT Gateway should be specified (in the case of a private subnet).



## Finding the Active Node within the VPC

When it comes to finding an active node in a clustering environment, we must define a virtual IP address **outside of the VPC CIDR** so that a **route** can be defined. In fact, defining a Virtual IP means **creating a new Destination** in the Route table.

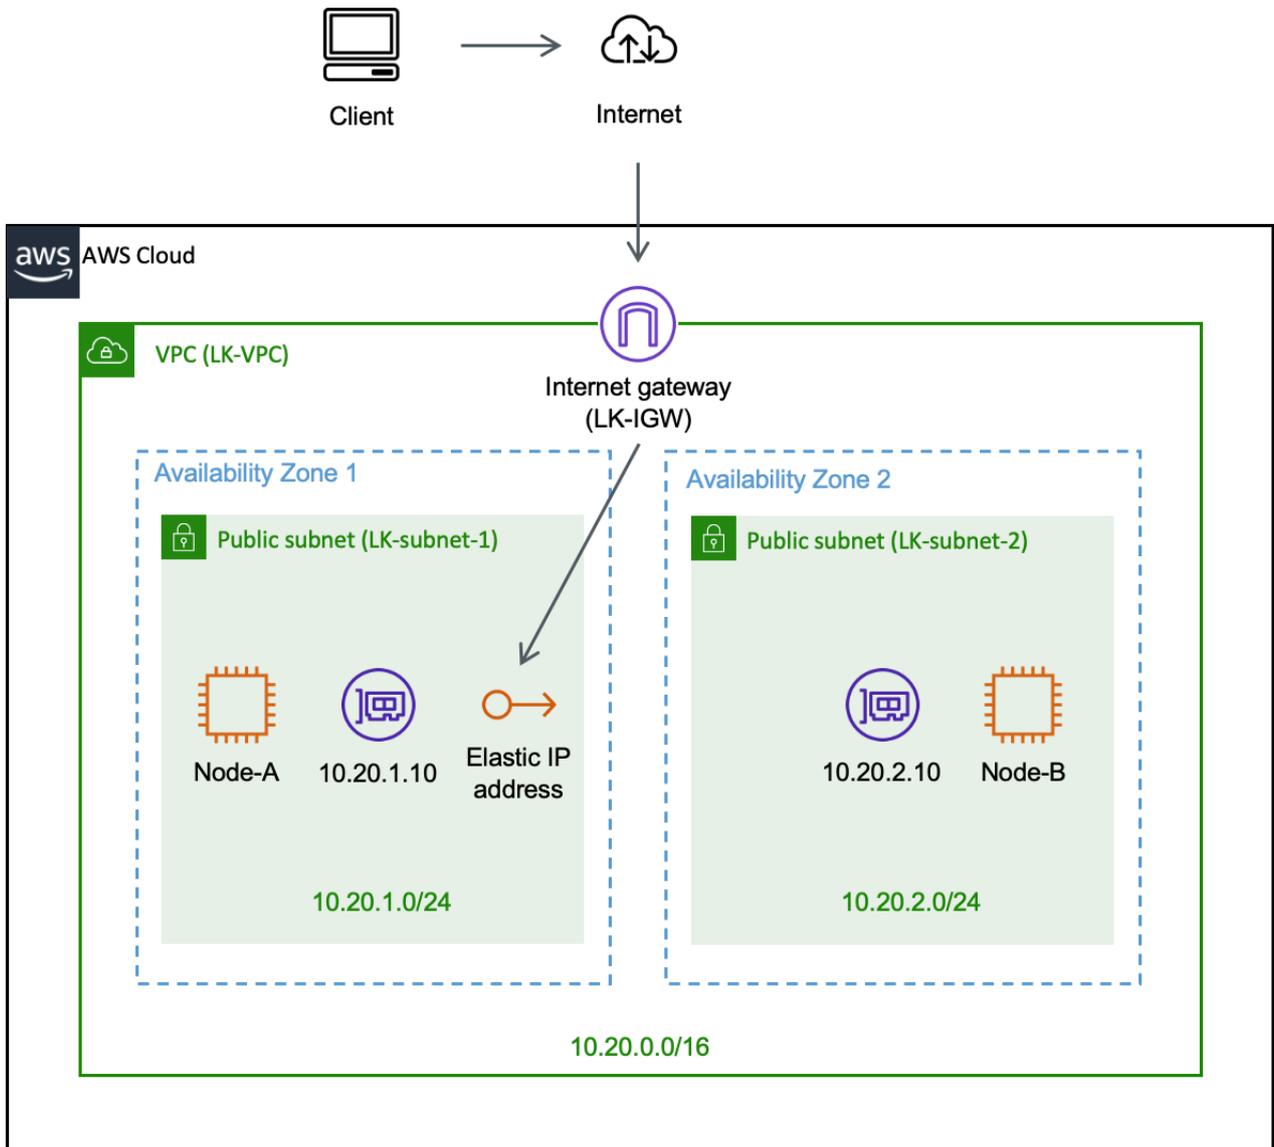


Once a new destination outside of VPC CIDR has been defined, LifeKeeper can change the target of this route to the active node (either 10.20.1.10 or 10.20.2.10 in the example configuration shown above). A client node should simply look for 10.10.10.10 as the destination and the routing table guides the traffic to the active node. The routing table is updated dynamically by SIOS's Recovery Kit for EC2.

\* The technical details of how this scenario works are discussed later in [Creating an AWS EC2 Resource \(RouteTable Scenario\)](#).

# 11.2.1.6.2. AWS Elastic IP Scenario

If clients are located on the internet and require public access to the active cluster node, the Elastic IP scenario should be selected. In this example the SIOS Recovery Kit for EC2 switches the Elastic IP when switching between nodes and attaches the Elastic IP to the ENI (Elastic Network Interface) of the new active node.



\* The technical details of how this scenario works are discussed later in [Creating AWS EC2 Resource \(Elastic IP Scenario\)](#).

# 11.2.1.6.3. Azure Internal Load Balancer Scenario

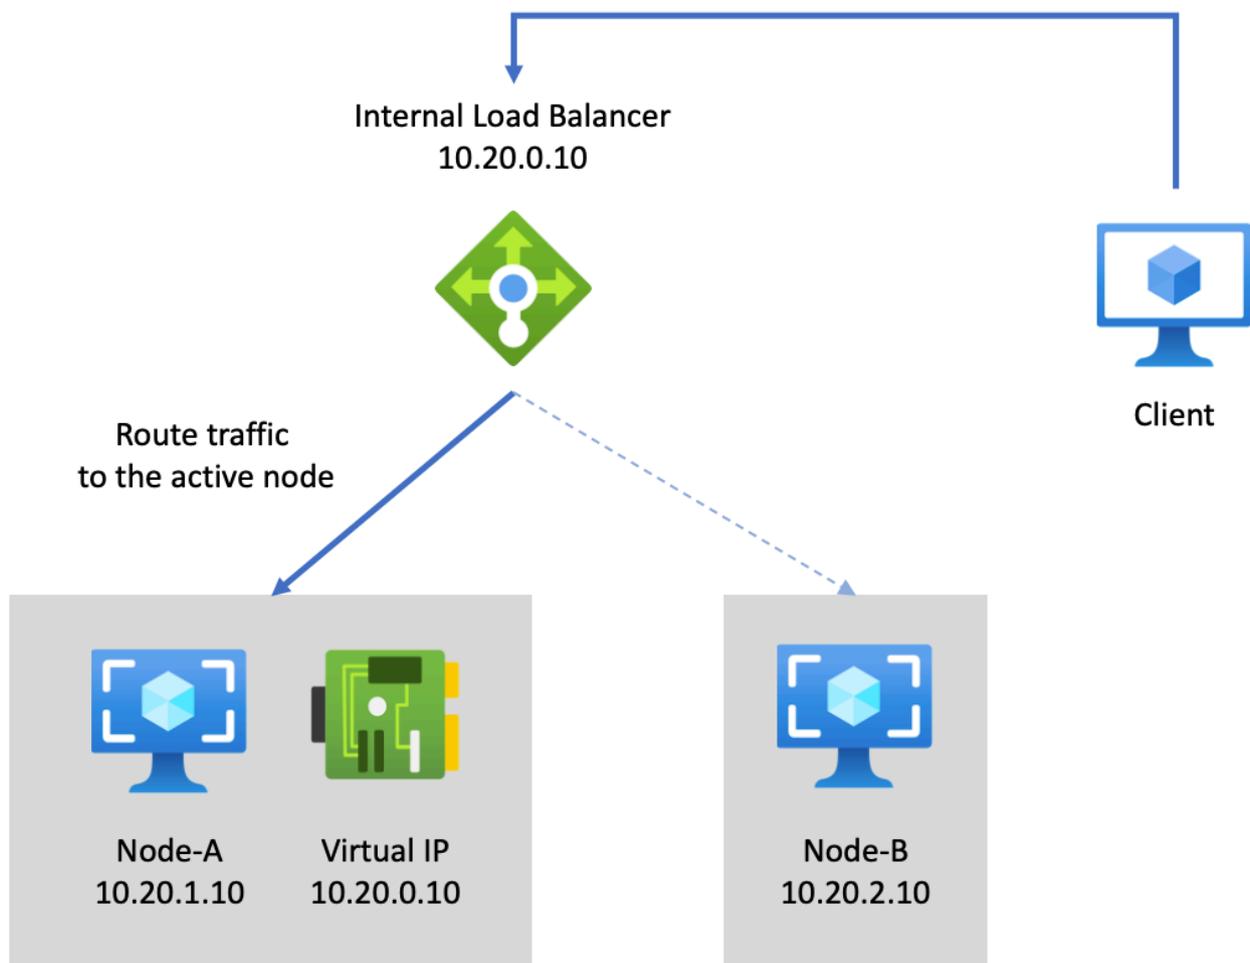
On Azure, using a Virtual IP may not enable traffic to be resolved to the active node as many infrastructure layers are abstracted in the cloud environment. To resolve this issue, an Internal Load Balancer should be used to forward traffic to the active node.

As an example, let us assume there are two nodes (10.20.1.10 and 10.20.2.10), working as a cluster and a virtual IP 10.20.0.10 is created that points to the active node.

In this case, an Azure Load Balancer is created with a Frontend IP address of 10.20.0.10 and a Virtual IP address with the same IP address 10.20.0.10 is created.

This ensures the client always connects to the load balancer first via the Virtual IP address 10.20.0.10.

The load balancer uses a health probe to determine which node is currently hosting the application, and all client connections are forwarded to that active node.



\* The technical details of how this scenario works are discussed later in [Azure – Using Internal Load Balancer](#).

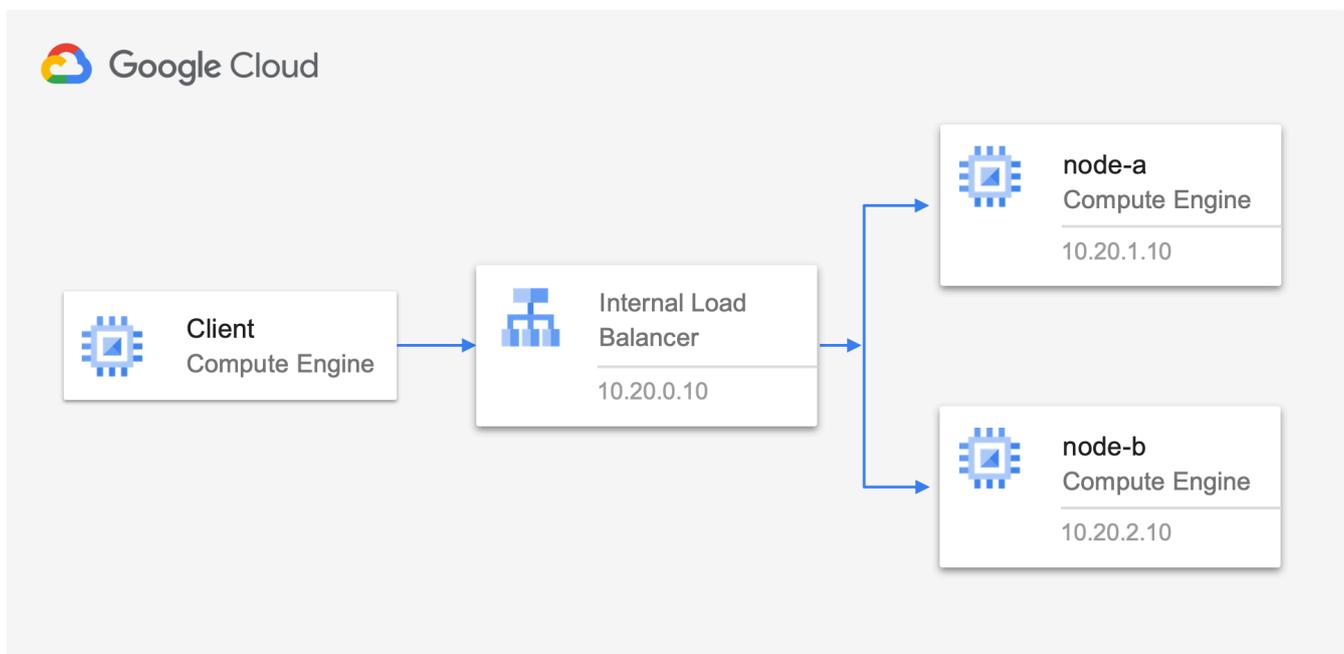
## 11.2.1.6.4. Google Cloud Internal Load Balancer Scenario

On Google Cloud, using a Virtual IP may not enable traffic to be resolved to the active node as many infrastructure layers are abstracted in the cloud environment. To resolve this issue, an Internal Load Balancer should be used to forward traffic to the active node.

As an example, let us assume that there are two nodes (10.20.1.10 and 10.20.2.10), working as a cluster and a virtual IP 10.20.0.10 is created that points to the active node.

In this case, an Internal Load Balancer is created with a Frontend IP address of 10.20.0.10.

The load balancer uses a health probe to determine which node is currently hosting the application, and all client connections are forwarded to that active node.



\* The technical details of how this scenario works are discussed later in [Google Cloud – Using an Internal Load Balancer](#).

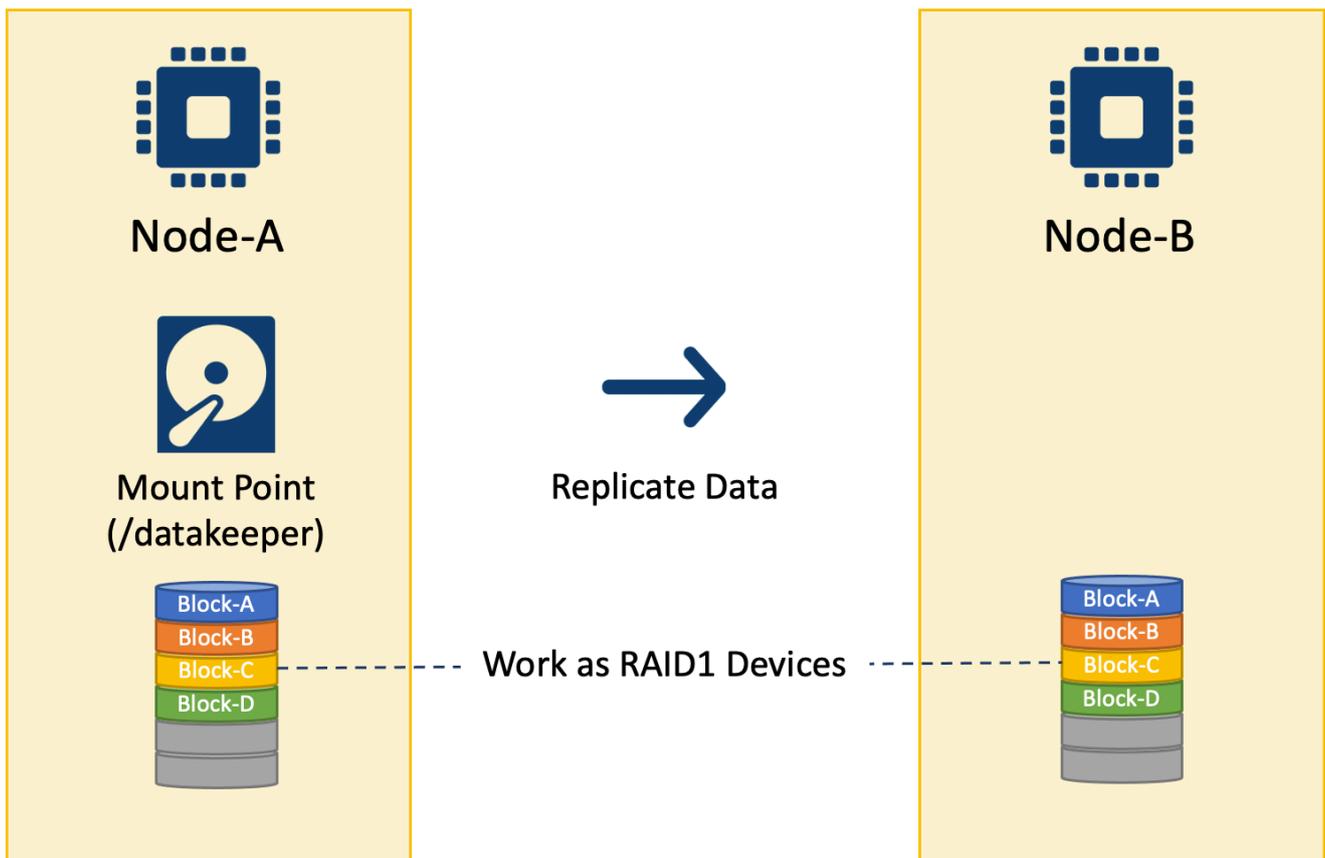
## 11.2.1.7. How does Data Replication between Nodes Work?

In the traditional datacenter scenario, data is commonly stored on a storage area network (SAN). The cloud environment doesn't typically support shared storage.

SIOS DataKeeper presents 'shared' storage using replication technology to create a copy of the currently active data. It creates a NetRAID device that works as a RAID1 device (data mirrored across devices).

Data changes are replicated from the Mirror Source (disk device on the active node – Node A in the diagram below) to the Mirror Target (disk device on the standby node – Node B in the diagram below).

In order to guarantee consistency of data across both devices, only the active node has write access to the replicated device (/datakeeper mount point in the example below). Access to the replicated device (the /datakeeper mount point) is not allowed while it is a Mirror Target (i.e., on the standby node).

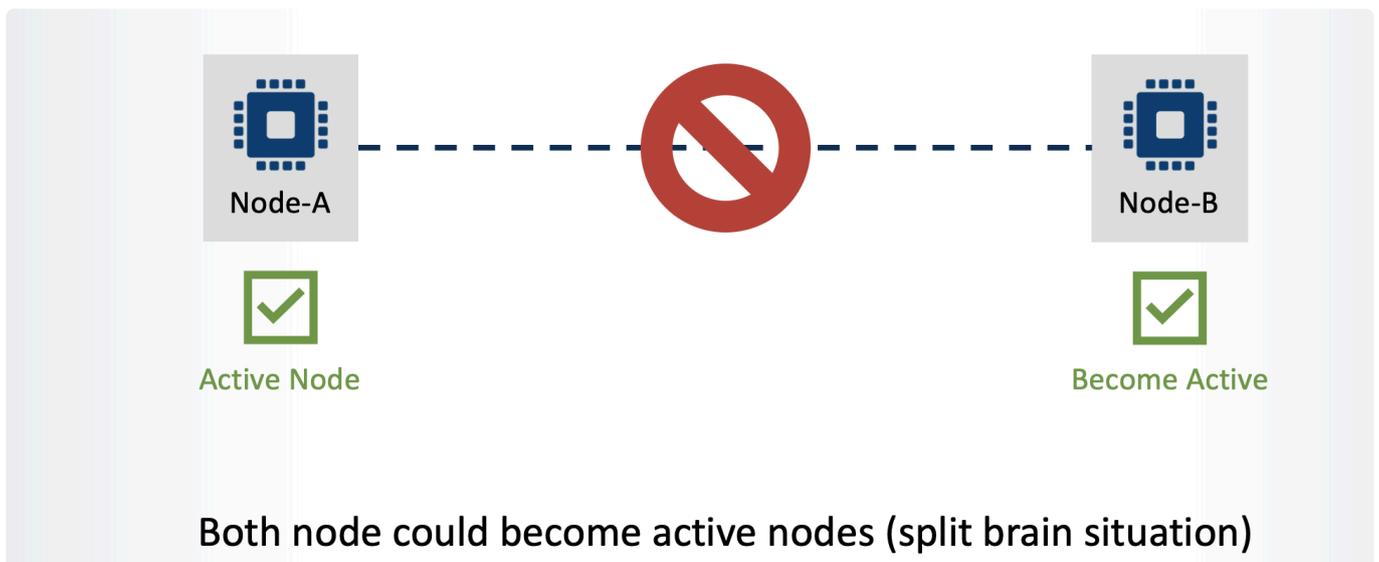


## 11.2.1.8. What is “Split Brain” and How to Avoid It

As we have discussed, in a High Availability cluster environment there is one active node and one or more standby node(s) that will take over service when the active node either fails or stops responding.



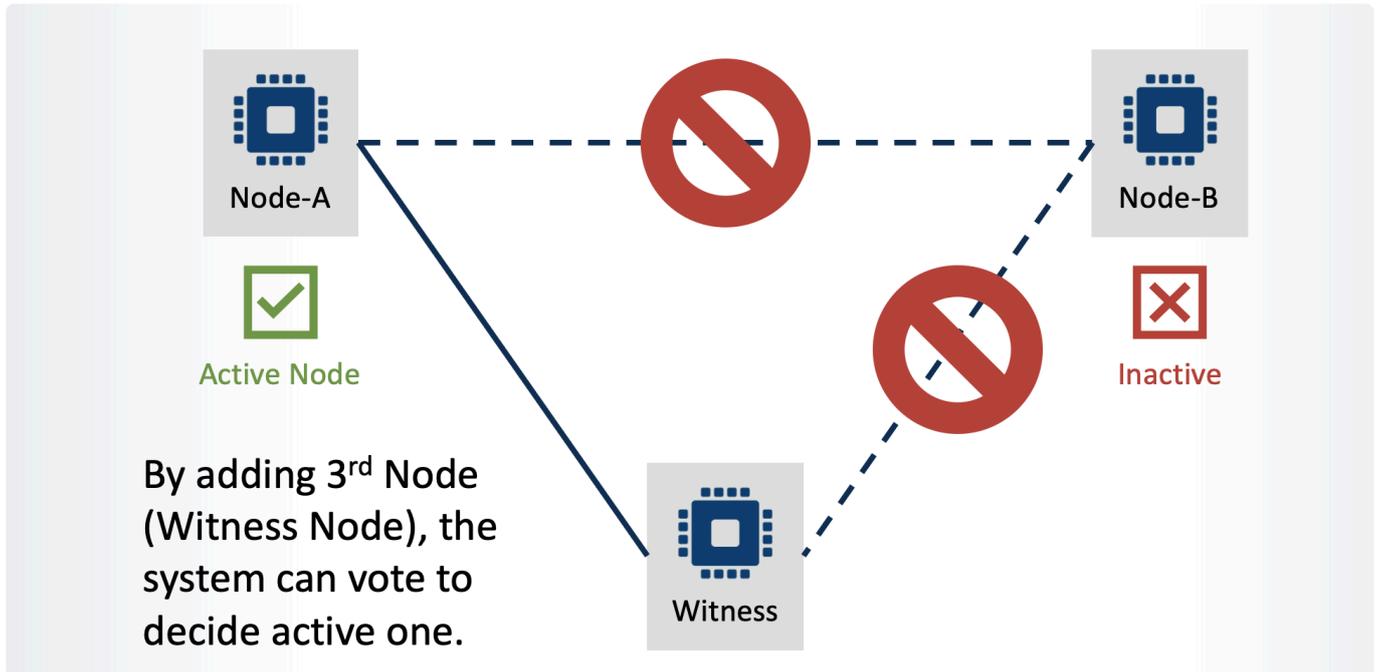
This sounds like a reasonable assumption until the network layer between the nodes is considered. What if the network path between the nodes goes down?



Neither node can now communicate with the other and in this situation the standby server may promote itself to become the active server on the basis that it believes the active node has failed. This results in both nodes becoming ‘active’ as each would see the other as being dead. As a result, data integrity and consistency is compromised as data on both nodes would be changing. This is referred to as “**Split Brain**”.

To avoid a split brain scenario, a Quorum node (also referred to as a ‘Witness’) should be installed within the cluster. Adding the quorum node (to a cluster consisting of an even number of nodes) creates an odd number of nodes (3, 5, 7, etc.), with nodes voting to decide which should act as the active node within the cluster.

In the example below, the server rack containing Node B has lost LAN connectivity. In this scenario, through the addition of a 3rd node to the cluster environment, the system can still determine which node should be the active node.



Quorum/Witness functionality is included in the LifeKeeper. At installation, Quorum / Witness is selected on all nodes (not only the quorum node) and a communication path is defined between all nodes (including the quorum node).

The quorum node doesn't host any active services. Its only role is to participate in node communication in order to determine which are active and to provide a 'tie-break vote' in case of a communication outage.

SIOS also supports [IO Fencing and Storage](#) as quorum devices, and in these configurations an additional quorum node is not required.

## 11.2.2. Documentation Style Used in this Guide

---

Most of the steps documented in this tutorial describe an environment consisting of two nodes (node-a and node-b). Examples of user command execution are shown within the guide. The sections below show the styles used to indicate the different nodes or users used to execute commands.

### Nodes to Execute Actions

When a particular set of command needs to be executed on a specific node, the prompt contains the user-name & host-name as follows:

```
1 [root@node-a]# <command name and its parameters are listed here>
```

An in-line comment similar to the below will also be shown:

 You need to execute the following steps on the primary node (node-a).

If the name of the node is not specified in the prompt, those commands need to be executed on both nodes (see the following example):

```
1 # <command name and its parameters are listed here>
```

An in-line comment similar to the below will also be shown:

 You need to execute the following steps for each node.

### Users to execute commands

Within the guide, the user who executes the command is identified by the prompt used.

If the command is executed as a regular user (e.g., `ec2-user` for most of the EC2 Instances), the prompt is shown as:

```
1 $ date
2 Thu Nov 26 22:58:22 UTC 2020
```

Once you become the `root` user, the prompt is described as `#`

```
1 $ sudo su -  
2 # shutdown -h now
```

If a specific user needs to execute a command, the name of the user is included in the prompt as follows:

```
1 [oracle]$ passwd
```

## Option to Select while Working with the SIOS GUI (Wizards)

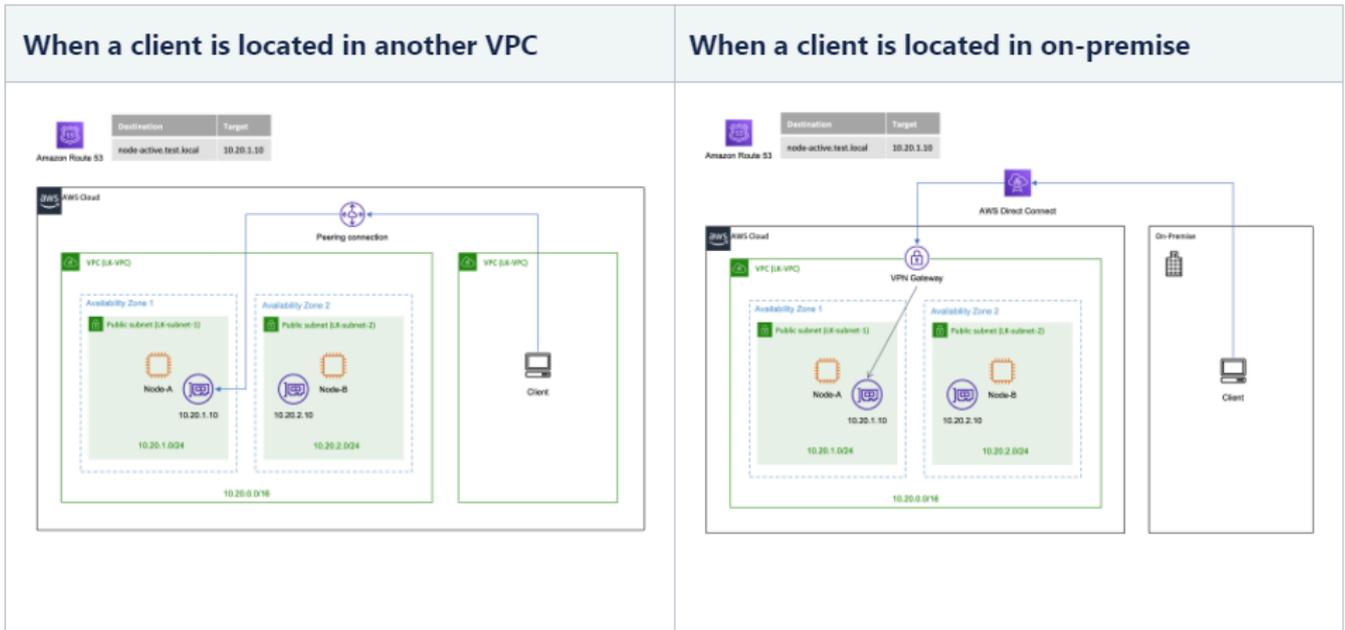
One of the benefits of using LifeKeeper is that wizards scan the system and then provide suggested values within the installation screens. In most cases, it is only necessary to confirm and select the

default values. Where default values may be chosen, the guide indicates them with a checkmark . If no checkmark is shown, a specific value (such as an environment-specific IP address) may need to be entered.

When modifying any parameters given in the guide (e.g., choosing a different resource name, different IP address, etc.), care must be taken to use the updated values for each occurrence of the parameters after they are introduced.

# 11.2.2.1. AWS Route53 Scenario

The previous sections discussed the Route Table Scenario when the client is located within the same VPC as the target node. However, if the client node is located within another VPC or located in an on-premise environment (connected to AWS via Direct Connect or VPN), the route table scenario cannot be used to specify the active target node. In this case Route53 may be used to route client traffic to the active target node.



The Route53 Recovery Kit (provided as part of LifeKeeper) can update the DNS entry allowing the client to connect to the active node as long as it can connect to the node via either VPC peering or Direct Connect.

 The technical details of how this scenario works are discussed later in [Creating AWS Route53 Resource](#).

## 11.2.3. Configuring Network Components and Creating Instances

---

This section outlines the computing resource required for each node, the network structure and the process required to configure these components.

- [Network Structure Used in This Tutorial](#)
- [Computing Resources Used in This Tutorial](#)

In the following sections we will configure these network components and create instances on the cloud platforms.

- [Creating an Instance on AWS from Scratch](#)
- [Creating an Instance on Azure from Scratch](#)
- [Creating an Instance on Google Cloud from Scratch](#)

## 11.2.3.1. Network Structure Used in this Tutorial

---

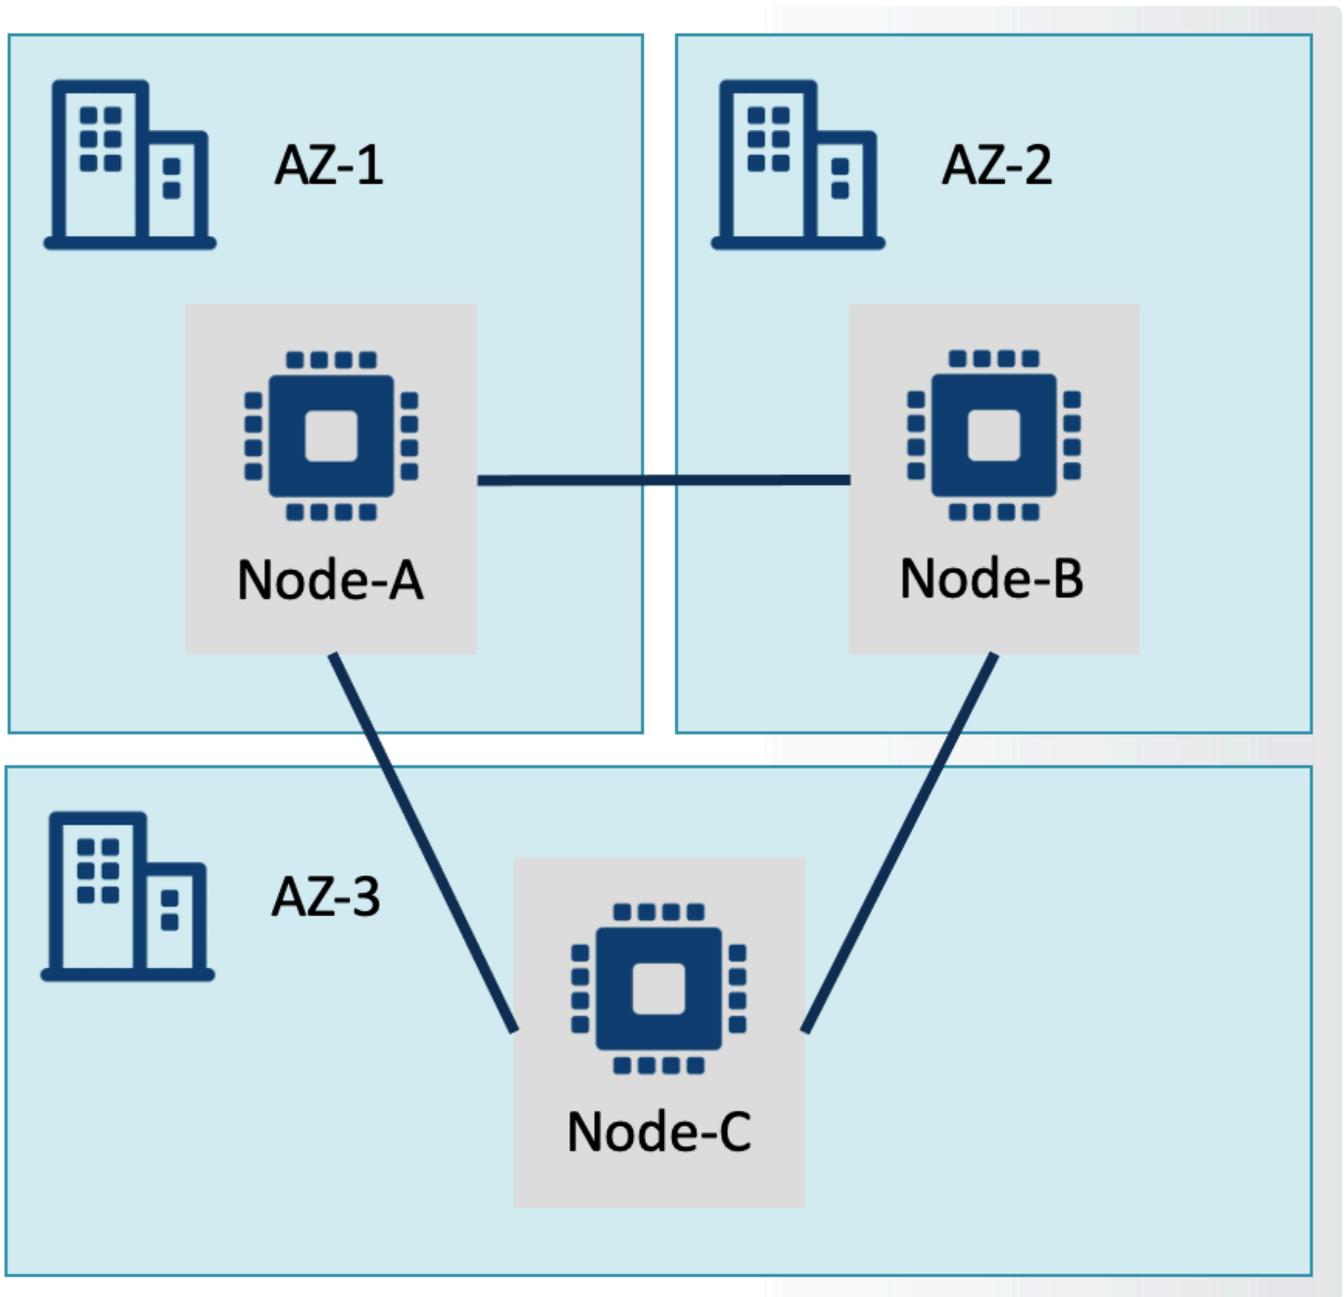
In this tutorial three nodes will be created. Node-A and Node-B will be used as the Active node and Standby node, respectively. In addition to these two nodes, Node-C (the Quorum/Witness node), will also be created. (See [What is](#) for an explanation of the need for a Quorum/Witness Node).

As discussed in [How Workloads Should be Distributed when Migrating to a Cloud Environment](#), SIOS strongly recommends the use of split workloads within different Availability Zones ('datacenters').

\* Although this Evaluation Guide discusses the Quorum / Witness Node (node-c), this third node may be skipped to simplify the learning experience. However, for a production environment (or a formal PoC project), SIOS strongly recommends the use of the Quorum/Witness node.

In the following topics, nodes in different Availability Zones will be created.

- [Creating an Instance in AWS from Scratch](#)
- [Creating an Instance in Azure from Scratch](#)
- [Creating an Instance in Google Cloud from Scratch](#)



## 11.2.3.2. Computing Resources Used in this Tutorial

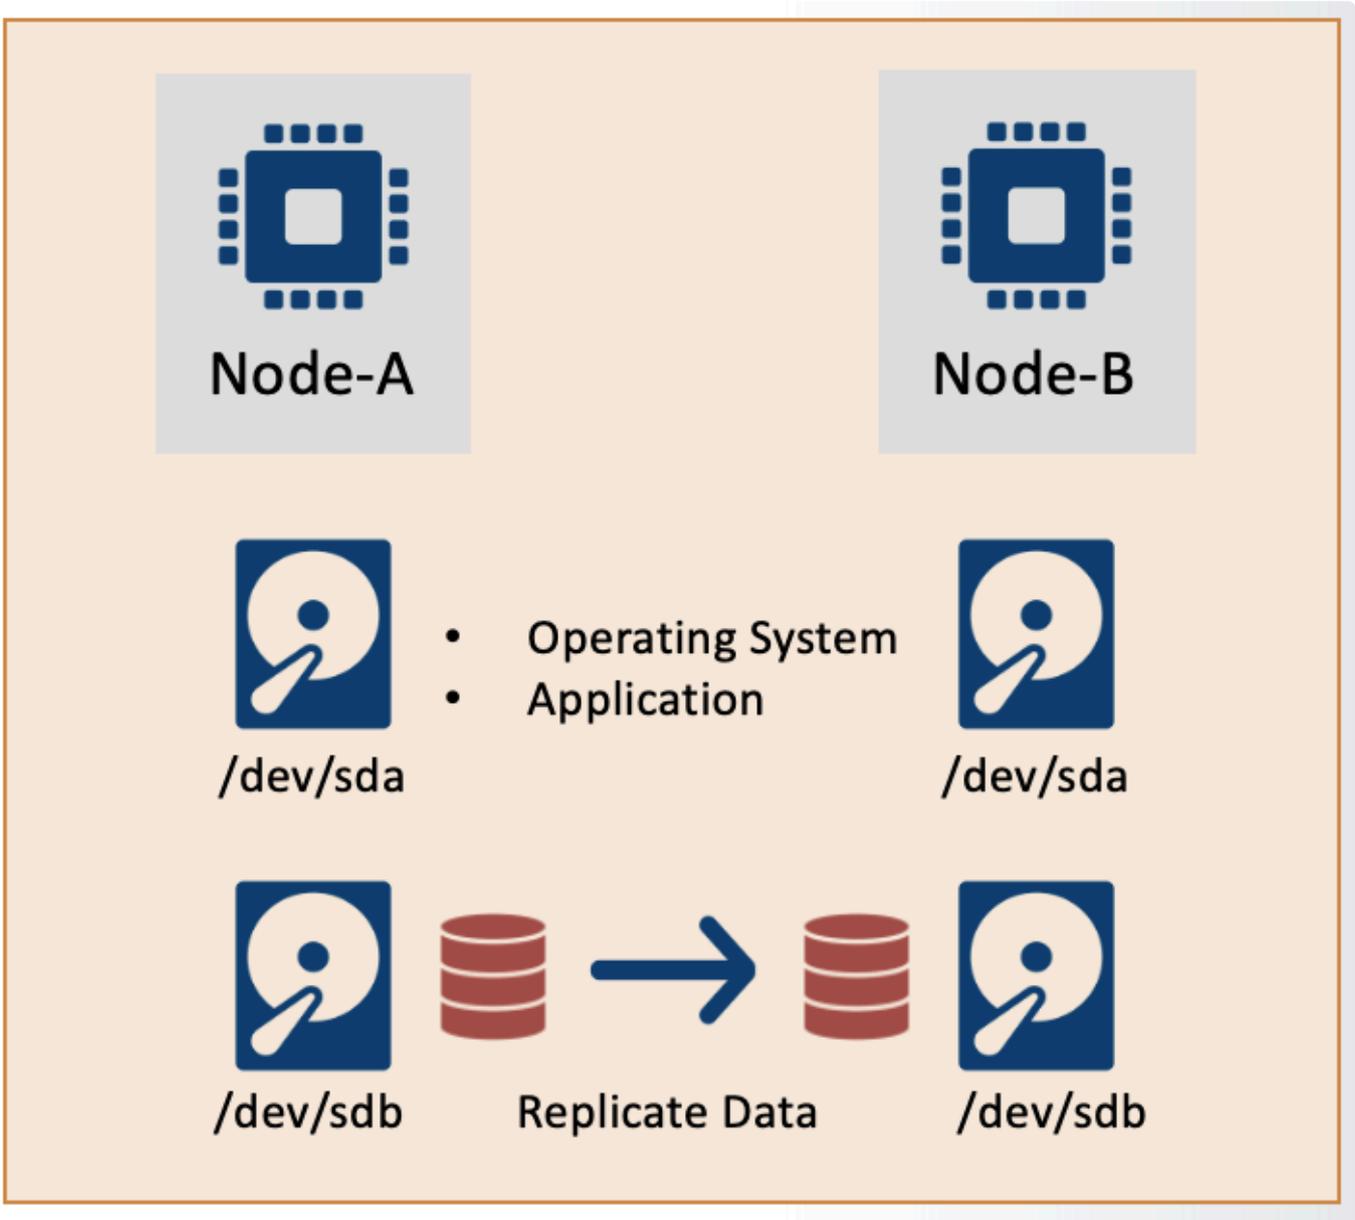
---

As discussed earlier, two nodes will be configured in a network to complete tasks in this tutorial.

Please refer to the LifeKeeper for Linux Release Notes as well as the resource plan for the resource(s) to be protected (e.g., Oracle). Check to ensure that the correct operating system is selected and that the appropriate resources (e.g., disk space, memory, CPU, etc.) are available.

 **Note:** Although this guide focuses on clusters that use local replicated storage, LifeKeeper clusters can also be built using certified shared storage solutions. If shared storage is to be leveraged in the cluster please refer to <http://docs.us.sios.com/spslinux/9.5.1/en/topic/storage-and-adapter-configuration>.

It is important to understand that two disks will be needed for each node. The first disk is used to boot the system (this has an operating system and the applications including LifeKeeper installed), while the other disk will be configured to replicate data from the “active” node to the “standby” node (note that the direction of data replication will change once a “switchover” occurs).



**\* Note:** The device name such as `/dev/sda` may vary based on the environment.

This guide assumes that the following resources are attached to each of the nodes:

Resource	Minimum Requirement	Evaluating with Oracle
Storage device for operating system & installing applications	10 Gb	30 Gb
Storage device for data replication between nodes	8 Gb	20 Gb
Memory	1 Gb	2 Gb

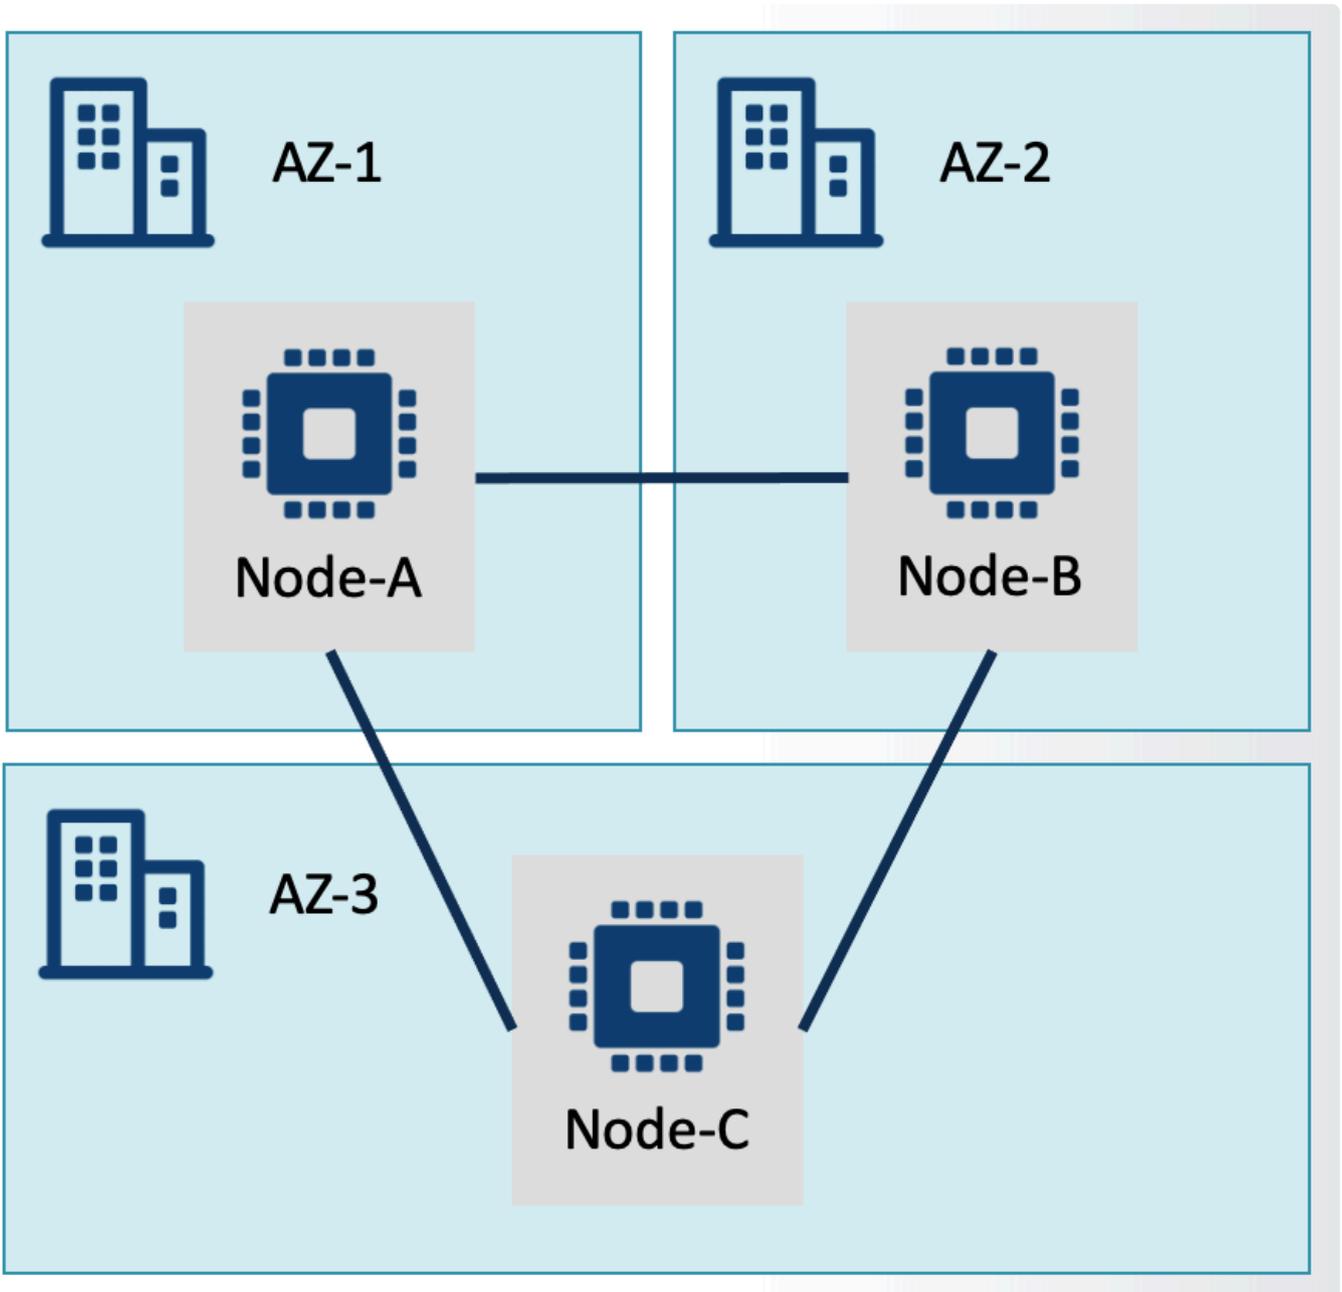
**!** As SIOS is not able to make resource recommendations, please refer to the application

vendor (e.g., SAP) in order to ensure that the correct resources are provisioned to run a particular application.

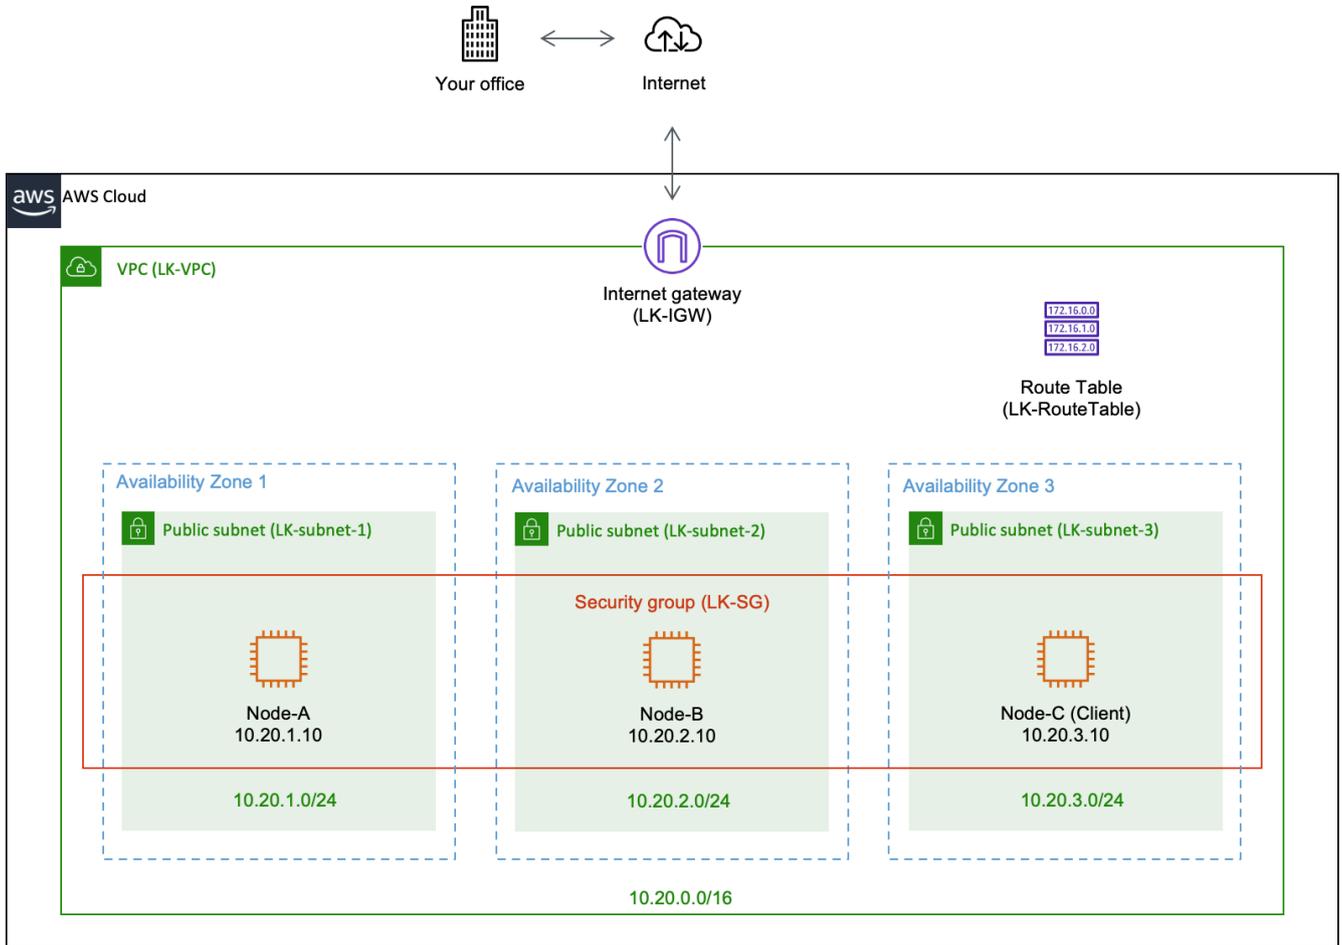
## 11.2.3.3. Creating an Instance in AWS from Scratch

**!** **Disclaimer:** The user interface may vary between regions or change over time. Please refer to the documentation provided by AWS if screenshots shown below are different from your experience.

This Evaluation Guide uses the following network structure and instances.



In AWS, these components can be defined as per the diagram below. One important aspect in AWS is that you need to create a subnet for each Availability Zone. Because the instances are distributed on different Availability Zones, these instances are assigned to different subnets.



Components listed in this diagram are listed in the following table:

Component	Name	Parameter	Value
VPC	LK-VPC	IPv4 CIDRs	10.20.0.0/16
Subnet	<b>Common values across subnets</b>	<b>VPC assignment</b>	<b>LK-VPC</b>
		<b>Auto-assign public IP</b>	<b>Yes</b>
	LK-subnet-1	IPv4 CIDR	10.20.1.0/24
	LK-subnet-2	IPv4 CIDR	10.20.2.0/24
	LK-subnet-3	IPv4 CIDR	10.20.3.0/24
Internet Gateway	LK-IGW	VPC association	LK-VPC
Route Table	LK-RouteTable	Subnet association	LK-public-subnet
		Destination=10.20.0.0/16	local
		Destination=0.0.0.0/0	LK-IGW
Security Group	LK-SG	Type=All traffic	Source=LK-SG
		Type=SSH	Source=Your

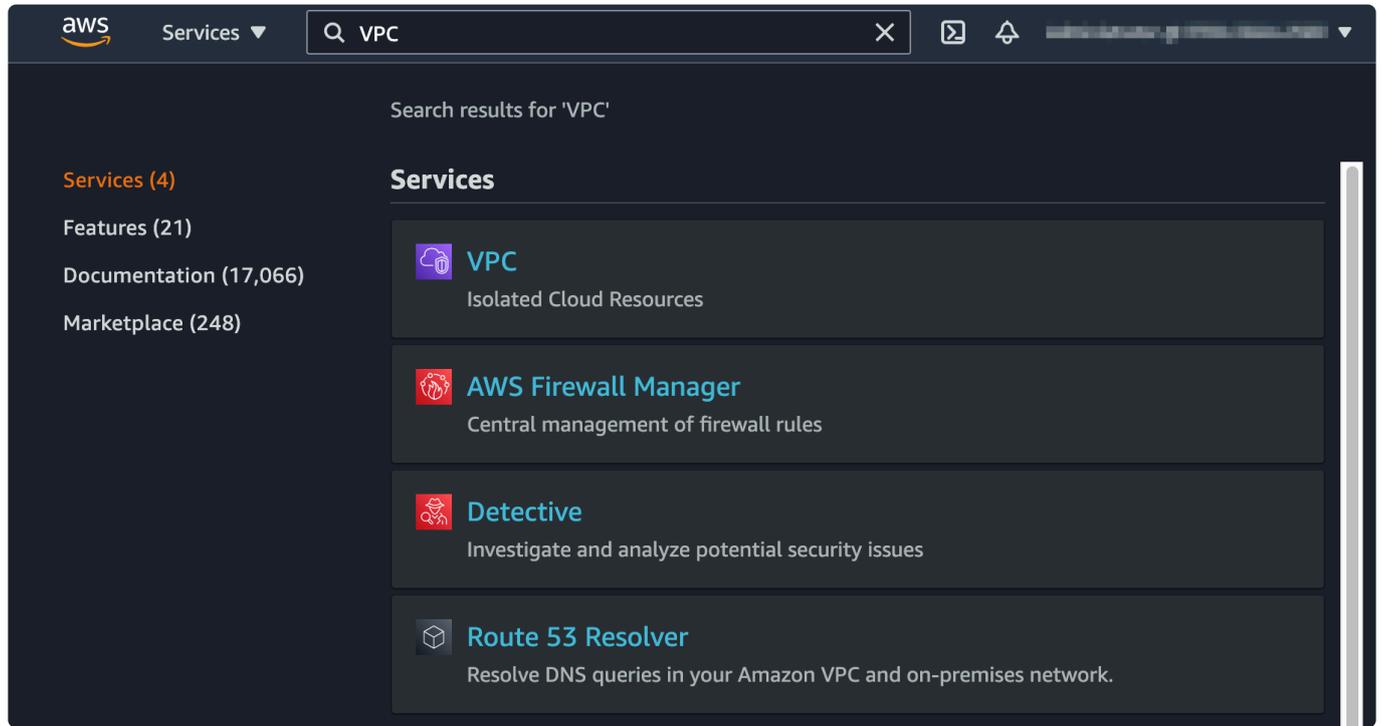
			Office's WAN IP
Instances	 <b>Common values across instances</b>	<b>VPC</b>	<b>LK-VPC</b>
		<b>Security Group</b>	<b>LK-SG</b>
		<b>Source/Dest Checking on a network interface (ENI)</b>	<b>Disabled</b>
	node-a	Subnet	LK-subnet-1
		Private IP Address	10.20.1.10
	node-b	Subnet	LK-subnet-2
		Private IP Address	10.20.2.10
	node-c	Subnet	LK-subnet-3
		Private IP Address	10.20.3.10

The following sections step through the processes necessary for creating these network components, beginning with the first instance.

- [Switching between AWS Services](#)
- [Deciding on an AWS Region](#)
- [Creating the VPC](#)
- [Creating a Subnet](#)
- [Creating an Internet Gateway and Assigning it to the VPC](#)
- [Creating the Route Table](#)
- [Creating a Security Group](#)
- [Creating the First EC2 Instance](#)
- [Creating the Second and Third Instances](#)

## 11.2.3.3.1. Switching between AWS Services

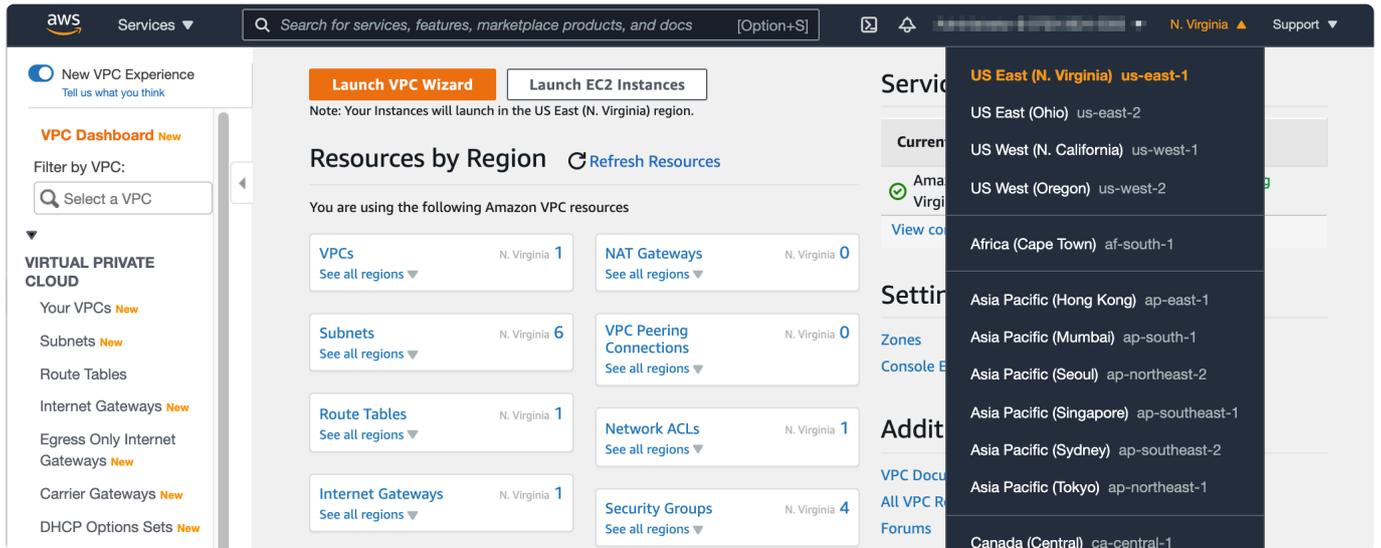
AWS has a significant and growing number of services available. To move to a different service you can type in the name of the service in the search box at the top of the AWS console. Select the service from the list. The screenshot below shows the VPC (Virtual Private Cloud) service as a selection choice.



## 11.2.3.3.2. Deciding on an AWS Region

AWS has regions in many geographic locations. It may be beneficial to select a region in close geographic proximity to the workplace location. Note that not all regions have 3+ Availability Zones, so exercise caution when making a selection in order to guarantee that the chosen region supports the intended configurations. When VPC is selected from the service menu, the current region can be seen at the top right corner of the AWS console.

Refer to the [AWS documentation](#) for more information.



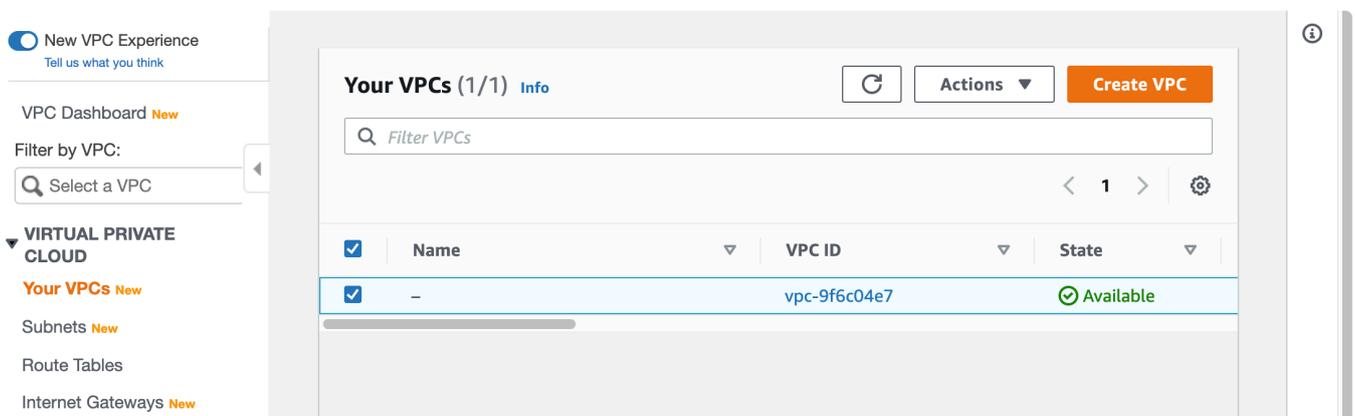
## 11.2.3.3.3. Creating the VPC

A VPC (Virtual Private Cloud) is an AWS resource that represents a local network. Different VPCs can be defined within the AWS cloud to logically separate different systems.

Refer to the [AWS documentation](#) for more information.

In this section we will create a VPC for testing LifeKeeper (named LK-VPC) as follows:

1. Select “Your VPCs” from the navigation pane located at the left side of the console.
2. There is a VPC created by AWS (default VPC). However, we are going to create a custom VPC based on the specification defined in this evaluation guide. Select “Create VPC” in the top right corner.



3. In the “Create VPC” wizard, enter the name of the VPC (LK-VPC) and the IPv4 CIDR block `10.20.0.0/16`.

VPC > Your VPCs > Create VPC

## Create VPC [Info](#)

A VPC is an isolated portion of the AWS cloud populated by AWS objects, such as Amazon EC2 instances.

### VPC settings

**Name tag - optional**  
Creates a tag with a key of 'Name' and a value that you specify.

  
**IPv4 CIDR block [Info](#)**  
  
**IPv6 CIDR block [Info](#)**  
 No IPv6 CIDR block  
 Amazon-provided IPv6 CIDR block  
 IPv6 CIDR owned by me  
**Tenancy [Info](#)**

### Tags

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

Key	Value - optional	
<input type="text" value="Name"/>	<input type="text" value="LK-VPC"/>	<input type="button" value="Remove"/>

You can add 49 more tags.

4. Now the VPC LK-VPC is created with the parameters specified.

New VPC Experience  
Tell us what you think

VPC Dashboard **New**

Filter by VPC:

**VIRTUAL PRIVATE CLOUD**

**Your VPCs** **New**

Subnets **New**

Route Tables

Internet Gateways **New**

Egress Only Internet Gateways **New**

Carrier Gateways **New**

DHCP Options Sets **New**

Elastic IPs **New**

Managed Prefix Lists **New**

Endpoints

Endpoint Services

NAT Gateways **New**

Peering Connections

**SECURITY**

Network ACLs **New**

Security Groups **New**

**REACHABILITY**

Reachability Analyzer

You successfully created vpc-0a7c93f9db7cf1dd2 / LK-VPC ✕ ?

VPC > Your VPCs > vpc-0a7c93f9db7cf1dd2

## vpc-0a7c93f9db7cf1dd2 / LK-VPC

Actions ▾

### Details [Info](#)

VPC ID vpc-0a7c93f9db7cf1dd2	State Available	DNS hostnames Disabled	DNS resolution Enabled
Tenancy Default	DHCP options set dopt-5716772f	Route table rtb-015a6633d04164f6d	Network ACL acl-0abe3a7fa3370b661
Default VPC No	IPv4 CIDR 10.20.0.0/16	IPv6 pool -	IPv6 CIDR (Network border group) -
Owner ID [REDACTED]			

**CIDRs** | Flow logs | Tags

### IPv4 CIDRs [Info](#)

CIDR	Status
10.20.0.0/16	Associated

## 11.2.3.3.4. Creating a Subnet

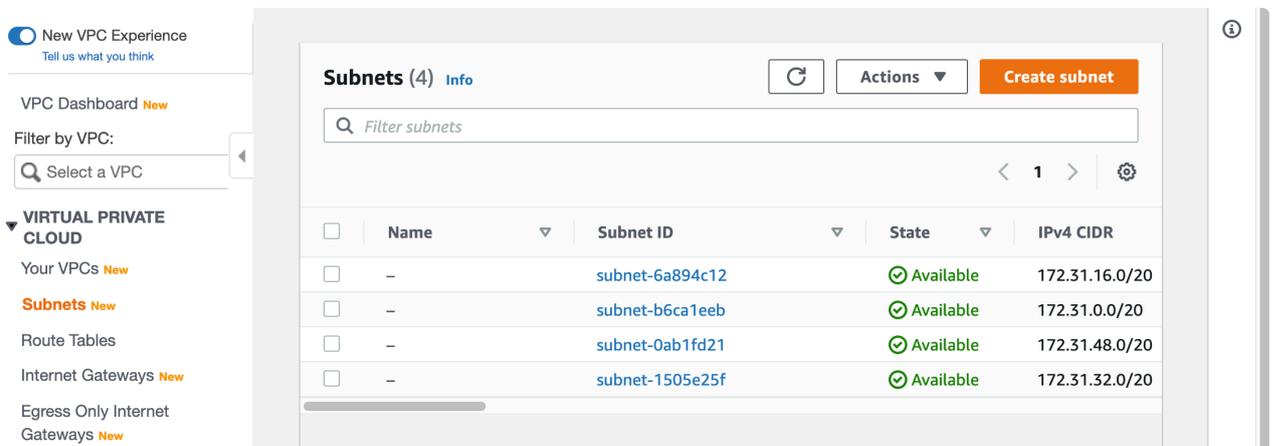
The basic concept of a subnet in AWS is the same as in on-premise environments. Namely it is a network within a network. Separate subnets can be defined to create a logical or physical boundary between components. Unique AWS-specific concepts are as follows:

- If a system is deployed across different Availability Zones (different datacenters), a separate subnet must be defined for each Availability Zone.
- If instances within a subnet need direct access from/to the Internet, it should be configured as a “public” subnet. Note that if instances need to be separated from/to the Internet, it should be configured as a “private” subnet. Consequently, there may need to be more than 2 subnets defined per Availability Zone.

Refer to the [AWS documentation](#) for more information.

In this tutorial, we will use three (3) Availability Zones and use only the “public” pattern to simplify the process. A “public” subnet is created as follows:

1. Go to the Subnet page by selecting “Subnets” from the navigation pane at the left side.
2. As per the VPC, there are pre-defined subnets for the default VPC. Select “Create Subnet” from the top right corner.



The screenshot shows the AWS Management Console interface for the Subnets page. On the left, the navigation pane is open, showing 'VIRTUAL PRIVATE CLOUD' with 'Subnets' selected. The main content area displays a table of subnets for a VPC. The table has columns for Name, Subnet ID, State, and IPv4 CIDR. Four subnets are listed, all in an 'Available' state.

	Name	Subnet ID	State	IPv4 CIDR
<input type="checkbox"/>	-	subnet-6a894c12	Available	172.31.16.0/20
<input type="checkbox"/>	-	subnet-b6ca1eeb	Available	172.31.0.0/20
<input type="checkbox"/>	-	subnet-0ab1fd21	Available	172.31.48.0/20
<input type="checkbox"/>	-	subnet-1505e25f	Available	172.31.32.0/20

3. Specify following parameters as shown in the screenshot:

- VPC ID: LK-VPC
- Subnet Name: LK-subnet-1
- Availability Zone: Select the first option.
- IPv4 CIDR block: 10.20.1.0/24
- Tag: Name = LK-subnet-1

While working in this wizard, a list of Availability Zones can be seen by selecting the dropdown box for the Availability Zones:

### Availability Zone [Info](#)

Choose the zone in which your subnet will reside, or let Amazon choose one for you.

No preference		▲
<input type="text" value="Q  "/>		
No preference		
US West (Oregon) / us-west-2a		us-west-2
ID: usw2-az1    Network border group: us-west-2		
US West (Oregon) / us-west-2b		us-west-2
ID: usw2-az2    Network border group: us-west-2		
US West (Oregon) / us-west-2c		us-west-2
ID: usw2-az3    Network border group: us-west-2		
US West (Oregon) / us-west-2d		us-west-2
ID: usw2-az4    Network border group: us-west-2		

Select the first Availability Zone (in this case `us-west-2a`)

VPC > Subnets > Create subnet ?

## Create subnet Info

**VPC**

VPC ID  
Create subnets in this VPC.

vpc-0a7c93f9db7cf1dd2 (LK-VPC) ▼

**Associated VPC CIDRs**

IPv4 CIDRs  
10.20.0.0/16

**Subnet settings**  
Specify the CIDR blocks and Availability Zone for the subnet.

**Subnet 1 of 1**

**Subnet name**  
Create a tag with a key of 'Name' and a value that you specify.

LK-subnet-1

The name can be up to 256 characters long.

**Availability Zone** Info  
Choose the zone in which your subnet will reside, or let Amazon choose one for you.

US West (Oregon) / us-west-2a ▼

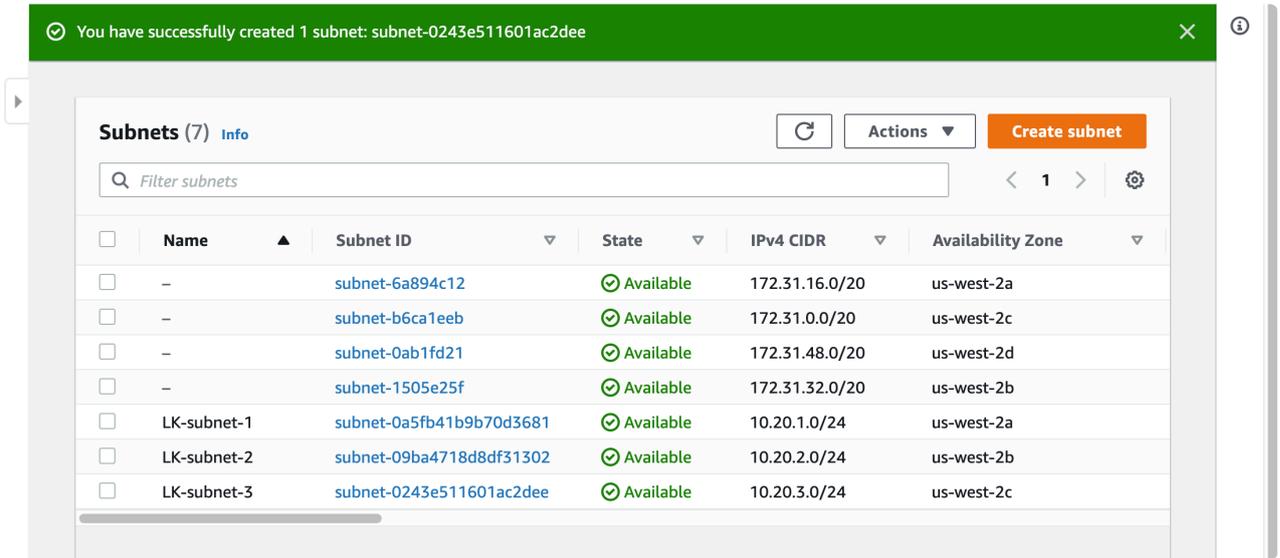
**IPv4 CIDR block** Info

10.20.1.0/24 ✕

▼ **Tags - optional**

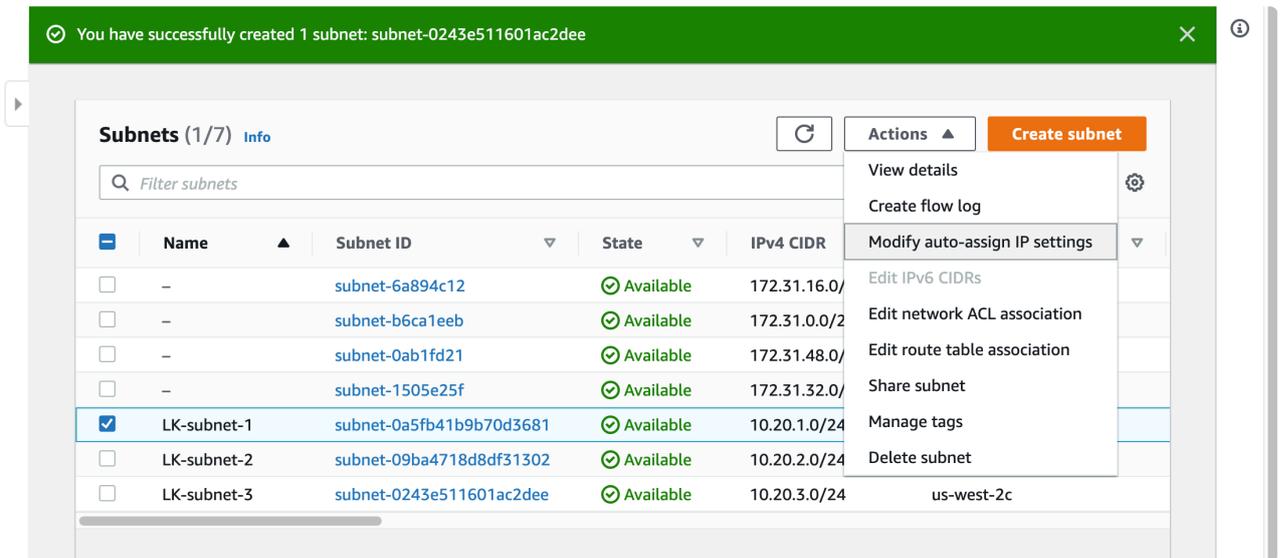
Key	Value - optional	
Q Name ✕	Q LK-subnet-1 ✕	Remove
<input type="button" value="Add new tag"/>		
You can add 49 more tags.		
<input type="button" value="Remove"/>		

4. Now the first subnet on Availability Zone `us-west-2a` is created. Return to the “Subnet Page”.
5. Create the second subnet `LK-subnet-2` with CIDR block `10.20.2.0/24` in another Availability Zone `us-west-2b`.
6. Create the third subnet `LK-subnet-3` with a CIDR block `10.20.3.0/24` on the Availability Zone `us-west-2c`.
7. Once the three subnets have been created, the list of subnets should be as shown in the screenshot below. The first subnets in the picture are the “default” subnets created for “default” VPC followed by the three subnets we have just created.



8. In this tutorial we will create a public subnet. This signifies that we want to assign a public IP address for instances we create under this subnet. To do so, select the subnet we have just created and then select “Modify auto-assign IP settings” from the Actions menu.

 **Note:** This change must be made for all 3 subnets, one by one.



9. Select “Enable auto-assign public IPv4 address”, and save the change.

VPC > Subnets > subnet-0a5fb41b9b70d3681 > Modify auto-assign IP settings i

## Modify auto-assign IP settings Info

Enable the auto-assign IP address setting to automatically request a public IPv4 or IPv6 address for a new network interface in this subnet.

### Settings

Subnet ID  
📄 subnet-0a5fb41b9b70d3681

Auto-assign IPv4 Info

Enable auto-assign public IPv4 address

Auto-assign customer-owned IPv4 address Info

Enable auto-assign customer-owned IPv4 address  
Option disabled because no customer owned pools found.

[Cancel](#) [Save](#)

10. Now the subnets are ready to use.

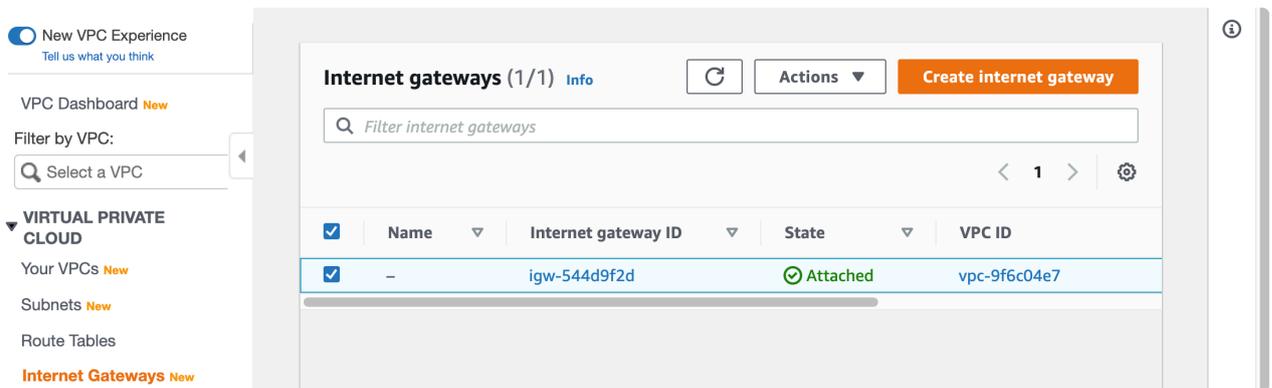
# 11.2.3.3.5. Creating an Internet Gateway and Assigning it to the VPC

The next step is to create an internet gateway. An internet gateway is a VPC component that allows communication between the VPC and the internet.

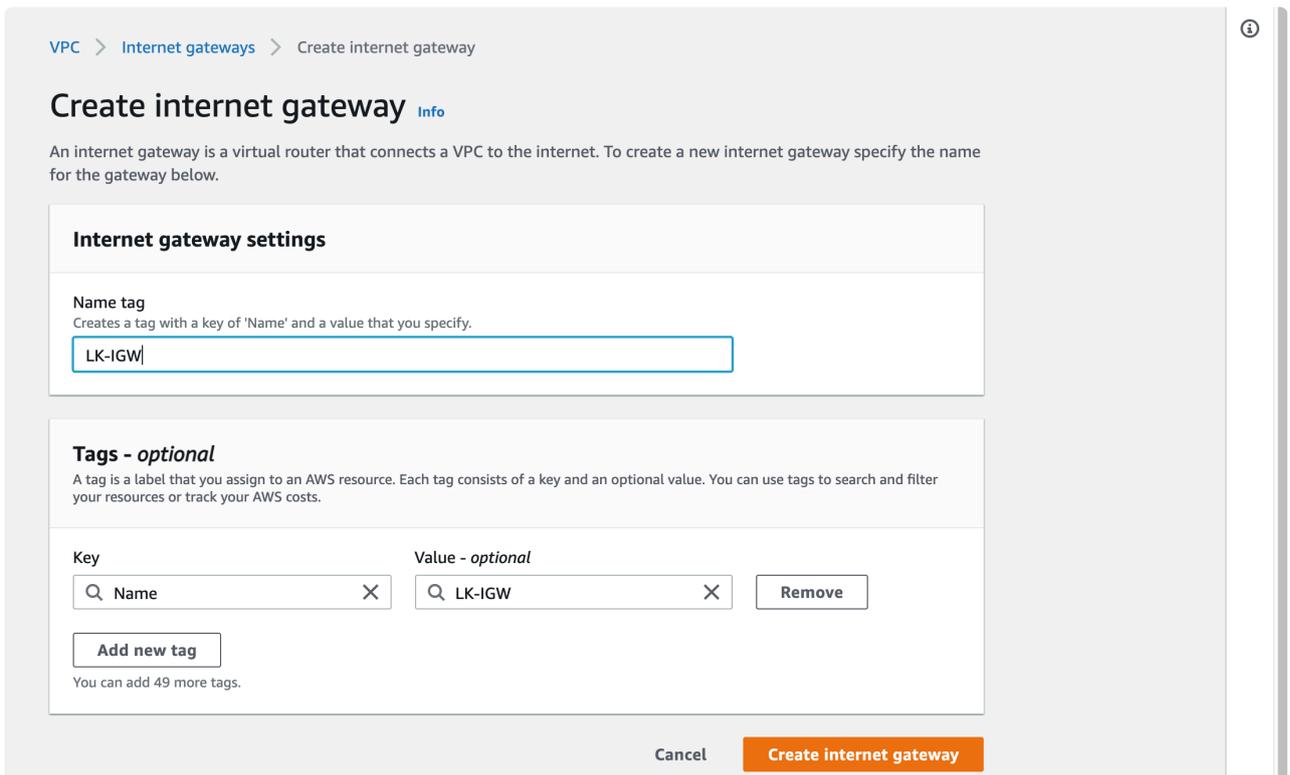
Refer to the [AWS documentation](#) for more information.

Let's create it and associate it with the VPC:

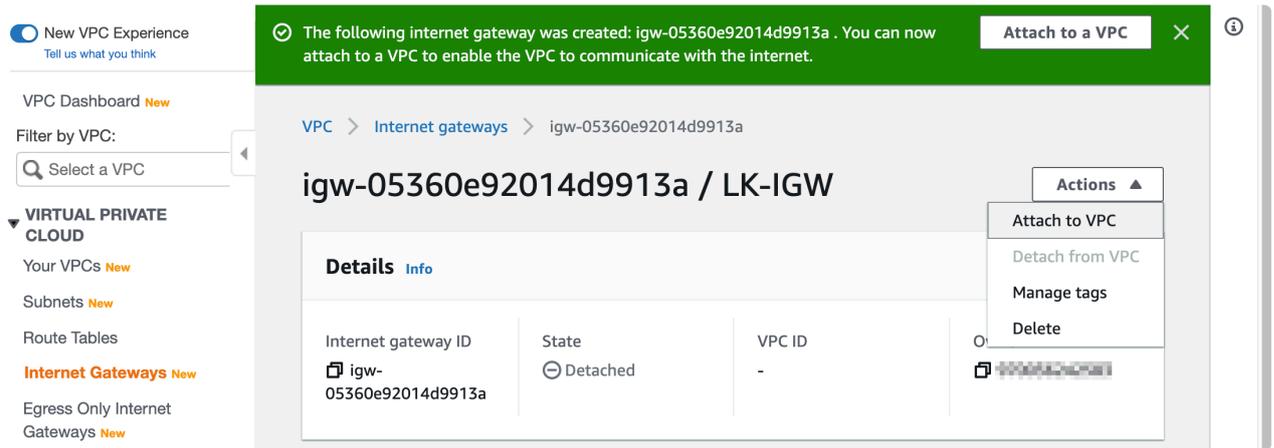
1. Select "Internet Gateways" from the navigation pane located at the left side.
2. You may see an existing Internet Gateway as a part of default VPC. Select "Create Internet Gateway" from the top right corner.



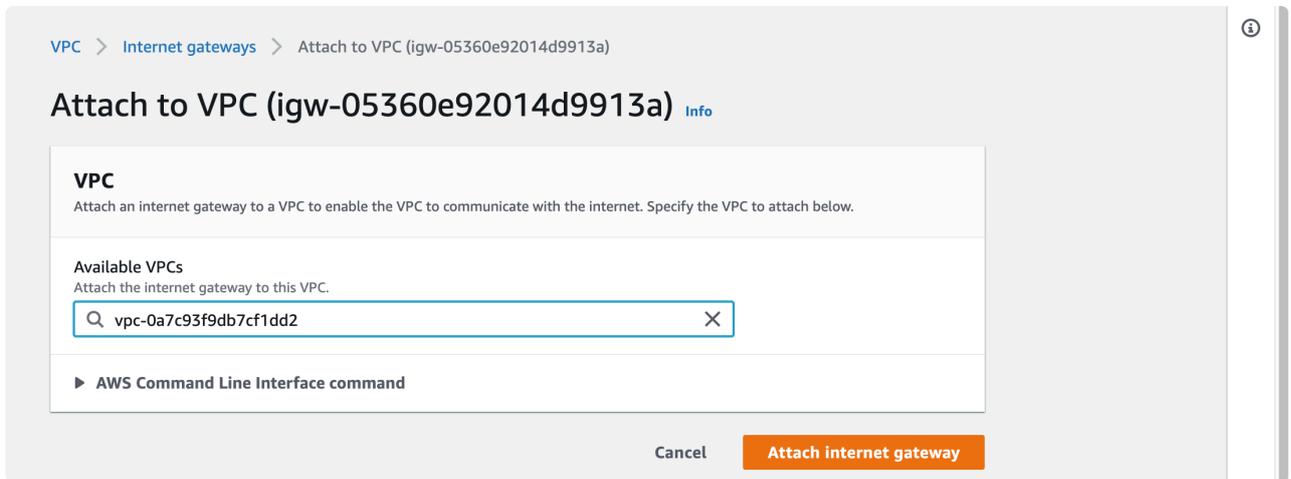
3. Specify the name of the Internet Gateway as follows.



4. Once the Internet Gateway is created, select “Attach to VPC”.



5. Select LK-VPC as the VPC to associate with.



6. The Internet Gateway is now configured.

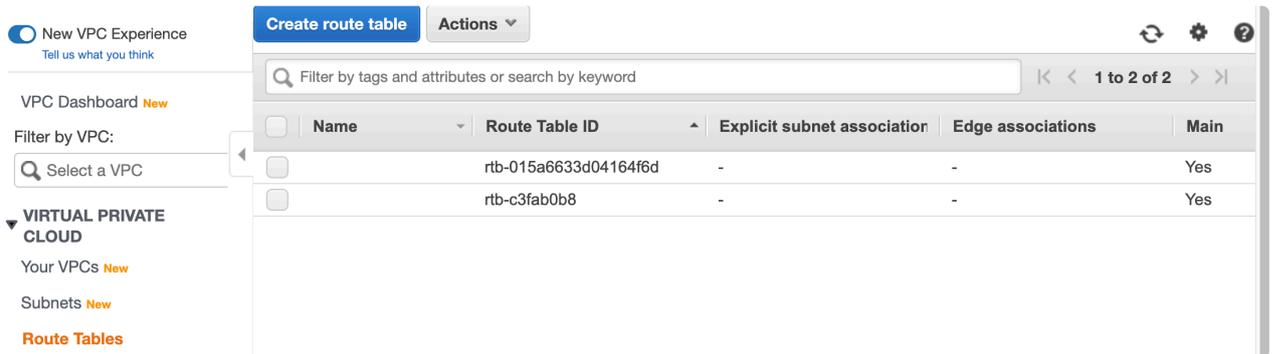
# 11.2.3.3.6. Creating the Route Table

The route table defines how the traffic from instances within a subnet should be transferred.

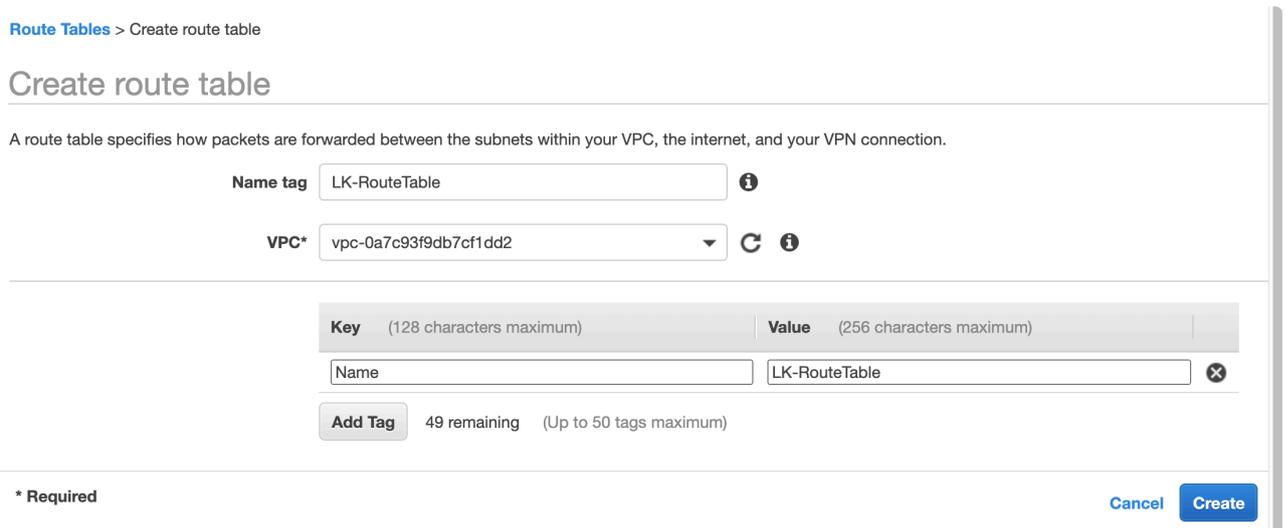
Refer to the [AWS documentation](#) for more information.

In this section we will create a route table and add a route to the internet via an Internet Gateway.

1. Select “Route Tables” from the left navigation pane.
2. Select “Create route table” at the top left corner.



3. Set the name of this route table as LK-RouteTable. Also, define the “Name” Tag. Please ensure that LK-VPC is selected as the associated VPC.



Once these fields are defined, select “Create”.

4. Go to the “Routes” page and click “Edit Routes”.

**Route Table: rtb-037dfb56d9950e373**

Destination	Target	Status
10.20.0.0/16	local	active

5. Local traffic is already defined. Now click “Add route”.

**Edit routes**

Destination	Target	Status	Propagated
10.20.0.0/16	local	active	No

**Add route**

\* Required Cancel Save routes

6. Enter 0.0.0.0/0 (meaning any IPv4 address, i.e., the internet) as the destination, and select “Internet Gateway” as Target.

**Edit routes**

Destination	Target	Status	Propagated
10.20.0.0/16	local	active	No
0.0.0.0/0	<b>Internet Gateway</b>		No

**Add route**

\* Required Cancel Save routes

7. Select “LK-IGW” as the Internet Gateway.

[Route Tables](#) > Edit routes

## Edit routes

Destination	Target	Status	Propagated
10.20.0.0/16	local	active	No
0.0.0.0/0	igw-05360e92014d9913a		No

**Add route**

**igw-05360e92014d9913a LK-IGW**

\* Required Cancel **Save routes**

Once a route is defined for 0.0.0.0/0, click “Save Routes” to close the wizard.

8. Now the routes are defined.

**Route Table: rtb-037dfb56d9950e373**

Summary | **Routes** | Subnet Associations | Edge Associations | Route Propagation | Tags

**Edit routes**

View: All routes

Destination	Target	Status
10.20.0.0/16	local	active
0.0.0.0/0	igw-05360e92014d9913a	active

9. The next step is to associate the route table with newly created subnets (LK-subnet-1, LK-subnet-2 and LK-subnet-3). Select the “Subnet Associations” tab.

The screenshot shows the AWS Management Console interface for configuring a route table. On the left, there is a navigation menu with categories like 'VIRTUAL PRIVATE CLOUD', 'SECURITY', and 'REACHABILITY'. The main content area shows the 'Route Table: rtb-037dfb56d9950e373' configuration page. The 'Subnet Associations' tab is active, displaying a table with no associations. Below this, a message states: 'The following subnets have not been explicitly associated with any route tables and are therefore associated with the main route table:'. A table lists these subnets:

Subnet ID	IPv4 CIDR	IPv6 CIDR
subnet-0243e511601ac2d...	10.20.3.0/24	-
subnet-0a5fb41b9b70d36...	10.20.1.0/24	-
subnet-09ba4718d8df313...	10.20.2.0/24	-

10. Click “Edit subnet associations” and select “LK-subnet-1”, “LK-subnet-2”, and “LK-subnet-3”.

The screenshot shows the 'Edit subnet associations' dialog box in the AWS console. At the top, it identifies the route table as 'Route table rtb-037dfb56d9950e373 (LK-RouteTable)'. Below this, there are three input fields for 'Associated subnets' containing the IDs: 'subnet-0243e511601ac2dee', 'subnet-0a5fb41b9b70d3681', and 'subnet-09ba4718d8df31302'. A modal window is open, displaying a table of subnets with their current associations:

Subnet ID	IPv4 CIDR	IPv6 CIDR	Current Route Table
subnet-0243e511601ac2dee   LK-subnet-3	10.20.3.0/24	-	Main
subnet-0a5fb41b9b70d3681   LK-subnet-1	10.20.1.0/24	-	Main
subnet-09ba4718d8df31302   LK-subnet-2	10.20.2.0/24	-	Main

At the bottom of the dialog, there are 'Cancel' and 'Save' buttons, and a note that asterisks (\*) denote required fields.

Now the new Route Table (“LK-RouteTable”) is defined and associated with the LK-subnet-1/2/3 subnets.

## 11.2.3.3.7. Creating a Security Group

A Security Group works as a firewall, and “allow” rules can be defined from each source. A source can be an internet address or security group, and a security group can be assigned to an Elastic Network Interface. As a network interface is attached to an Instance, this works as a firewall rule for each instance.

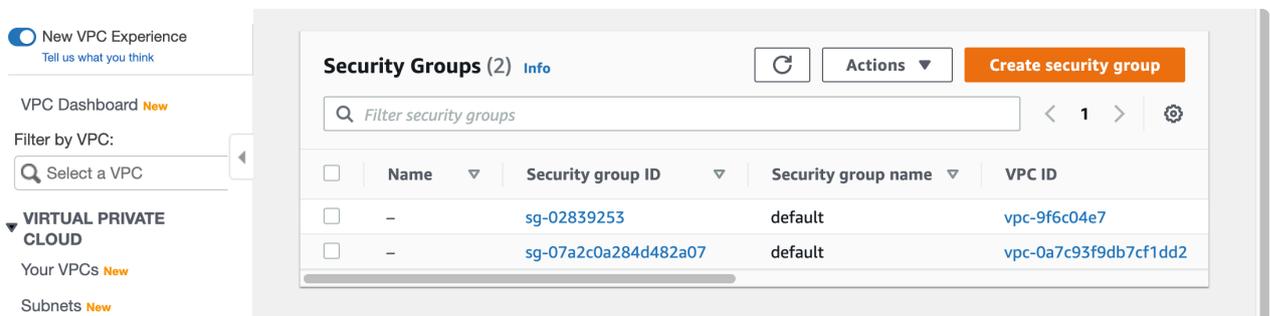
Refer to the [AWS documentation](#) for more information.

In this section, we will define two “allow” rules as follows:

- Allow access from the remote work location.
- Allow access between instances (node-a and node-b). This is done by allowing the same security group as the source and then assigning the same security group to both instances.

**Note:** Once a production environment is defined, the instances may belong to different security groups. In this case, a corresponding “allow” rule between security groups should be defined.

1. Select “Security Groups” from the left navigation pane. As previously mentioned, default security groups will already exist. Click “Create security group”.



2. Enter the values for the new security group as shown in the screenshot using the following parameters:

- Name: LK-SG
- VPC: LK-VPC
- Create the first Inbound Rule. Select “Add Rule” and select “All Traffic” as the type (for now). Then select “My IP” as Source type. This automatically selects the user’s WAN address as source. If required, update the source address range based on the in use WAN address (it may be appropriate to change this from a specific IP address to a range of IP addresses).

Once the values are confirmed, click “Create security group” at the bottom right.

## Create security group [Info](#)

A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. To create a new security group, complete the fields below.

### Basic details

Security group name [Info](#)

Name cannot be edited after creation.

Description [Info](#)

VPC [Info](#)

### Inbound rules [Info](#)

Inbound rule 1

Delete

Type [Info](#)

Protocol [Info](#)

Port range [Info](#)

Source type [Info](#)

Source [Info](#)

Description - optional [Info](#)

Add rule

### Outbound rules [Info](#)

Outbound rule 1

Delete

Type [Info](#)

Protocol [Info](#)

Port range [Info](#)

Destination type [Info](#)

Destination [Info](#)

Description - optional [Info](#)

Add rule

### Tags - optional

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

No tags associated with the resource.

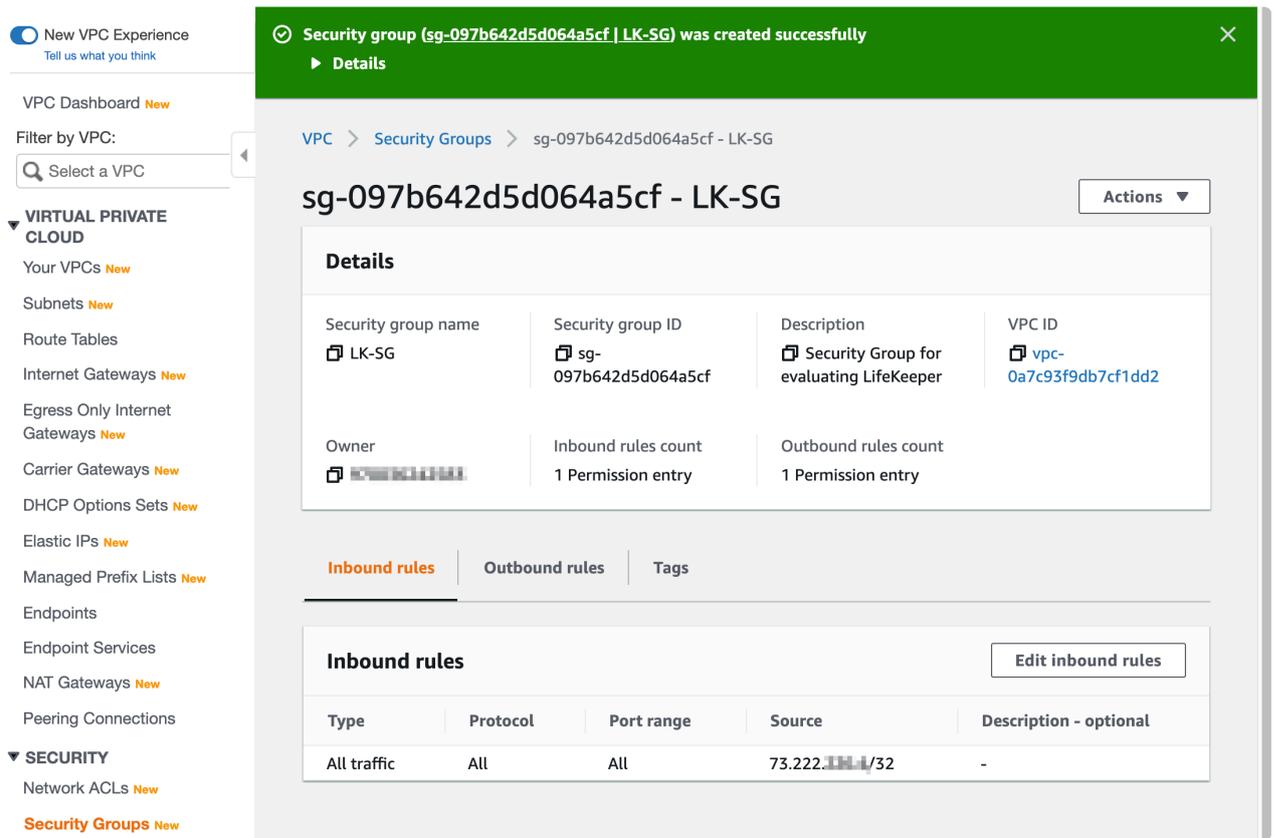
Add new tag

You can add up to 50 more tag

Cancel

Create security group

- Once new security group has been created, click “Edit inbound rules” to allow traffic within the same Security Group.



- Once the “Edit inbound rules” page appears, select “Add rule”, then select “LK-SG” (which we just created). With this change, any instances that are associated with “LK-SG” can communicate with each other.

VPC > Security Groups > sg-097b642d5d064a5cf - LK-SG > Edit inbound rules

## Edit inbound rules [Info](#)

Inbound rules control the incoming traffic that's allowed to reach the instance.

### Inbound rules [Info](#)

**Inbound rule 1** Delete

Type [Info](#): All traffic  
Protocol [Info](#): All  
Port range [Info](#): All  
Source type [Info](#): Custom  
Source [Info](#): 73.222.255.0/32  
Description - optional [Info](#)

**Inbound rule 2** Delete

Type [Info](#): All traffic  
Protocol [Info](#): All  
Port range [Info](#): All  
Source type [Info](#): Custom  
Source [Info](#):  
Description - optional [Info](#)

[Add rule](#)

**NOTE:** Any edits made on existing rules will cause traffic that depends on that rule to be interrupted. A new rule created with the new details. This will allow the new rule to be created.

[Cancel](#) [Preview changes](#) [Save rules](#)

Source
CIDR blocks
0.0.0.0/0
0.0.0.0/8
0.0.0.0/16
0.0.0.0/24
0.0.0.0/32
::/0
::/16
::/32
::/48
::/64
Security Groups
default   sg-07a2c0a284d482a07
LK-SG   sg-097b642d5d064a5cf
Prefix lists
LK-SG   sg-097b642d5d064a5cf
com.amazonaws.us-wes...   pl-00a54069
com.amazonaws.us-wes...   pl-68a54001

Once the change is confirmed, click "Save rules". The Security Group is now created.

5. Confirm the parameters for the newly created Security Group.

New VPC Experience  
Tell us what you think

VPC Dashboard **New**

Filter by VPC:

**VIRTUAL PRIVATE CLOUD**

Your VPCs **New**

Subnets **New**

Route Tables

Internet Gateways **New**

Egress Only Internet Gateways **New**

Carrier Gateways **New**

DHCP Options Sets **New**

Elastic IPs **New**

Managed Prefix Lists **New**

Endpoints

Endpoint Services

NAT Gateways **New**

Peering Connections

**SECURITY**

Network ACLs **New**

**Security Groups **New****

**REACHABILITY**

**Inbound security group rules successfully modified on security group (sg-097b642d5d064a5cf | LK-SG)**

**Details**

VPC > Security Groups > sg-097b642d5d064a5cf - LK-SG

# sg-097b642d5d064a5cf - LK-SG

Actions ▾

## Details

Security group name LK-SG	Security group ID sg-097b642d5d064a5cf	Description Security Group for evaluating LifeKeeper	VPC ID vpc-0a7c93f9db7cf1dd2
Owner [Avatar]	Inbound rules count 2 Permission entries	Outbound rules count 1 Permission entry	

**Inbound rules** | Outbound rules | Tags

## Inbound rules

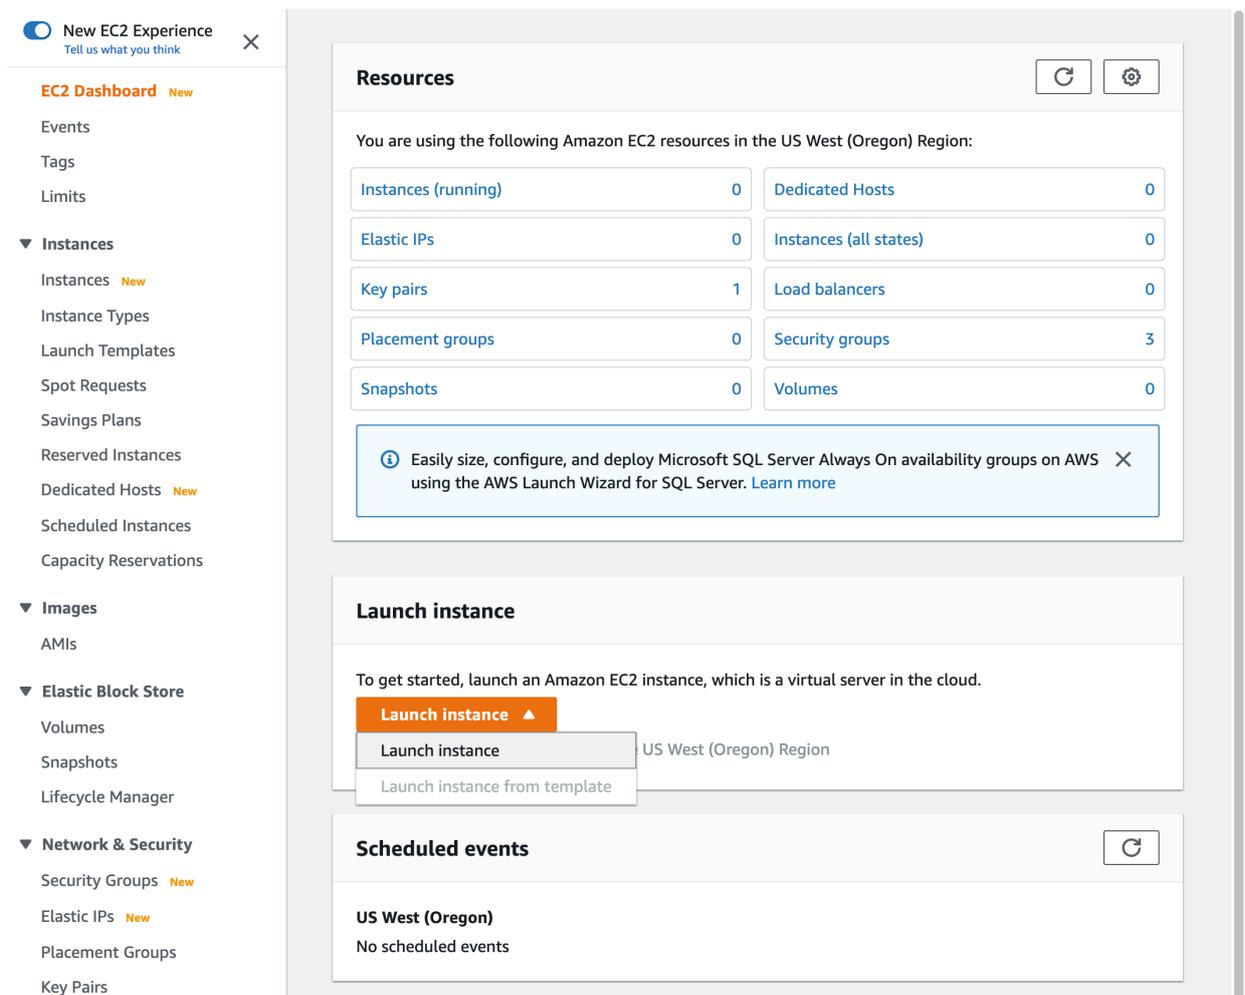
Edit inbound rules

Type	Protocol	Port range	Source	Description - optional
All traffic	All	All	73.222.111.1/32	-
All traffic	All	All	sg-097b642d5d064a5cf (LK-SG)	-

# 11.2.3.3.8. Creating the First EC2 Instance

In previous sections we covered configuration of the network. Now we are going to create the first instance. As discussed in [Computing Resources Used in this Tutorial](#), we need two disks. This section also describes how to create the second disk.

1. Go to “EC2” service. It may be necessary to enter “EC2” in the top search box of the AWS Console to select EC2.
2. Click “Launch Instance”.



3. Click “Operating System”. Enter the name of the operating system and select the specific Machine Image. Before selecting the operating system and its version, please review the system requirements for both LifeKeeper and the application you are going to protect.

- 1. Choose AMI
- 2. Choose Instance Type
- 3. Configure Instance
- 4. Add Storage
- 5. Add Tags
- 6. Configure Security Group
- 7. Review

### Step 1: Choose an Amazon Machine Image (AMI)

[Cancel and Exit](#)

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.

X

Search by Systems Manager parameter

Quick Start (1)

---

My AMIs (0)

---

AWS Marketplace (204)

---

Community AMIs (112)

**Red Hat Enterprise Linux (RHEL) 7 (HVM)**

★★★★★ (0) | 7.7\_HVM | By [Amazon Web Services](#)

Linux/Unix, Red Hat Enterprise Linux 7.7\_HVM | 64-bit (x86) Amazon Machine Image (AMI) | Updated: 10/15/19

Red Hat Enterprise Linux 7.7 (HVM)

[More info](#)

Select

Free tier eligible

1 to 10 of 204 Products

4. Click “Instance Size”. This tutorial uses `t2.micro` since it is defined as the minimum system requirement for evaluation of LifeKeeper and it may qualify for Free Tier usage. Again, it may be necessary to select a larger instance size depending on the application to be protected.

- 1. Choose AMI
- 2. Choose Instance Type
- 3. Configure Instance
- 4. Add Storage
- 5. Add Tags
- 6. Configure Security Group
- 7. Review

### Step 2: Choose an Instance Type

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instances are virtual servers that can run applications. They have varying combinations of CPU, memory, storage, and networking capacity, and give you the flexibility to choose the appropriate mix of resources for your applications. [Learn more](#) about instance types and how they can meet your computing needs.

Filter by: All instance families All generations [Show/Hide Columns](#)

Currently selected: t2.micro (- ECUs, 1 vCPUs, 2.5 GHz, -, 1 GiB memory, EBS only)

Note: The vendor recommends using a **m3.large** instance (or larger) for the best experience with this product.

	Family	Type	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance	IPv6 Support
⊘	t1	t1.micro <span style="font-size: 8px; background-color: #90EE90; padding: 1px;">Free tier eligible</span>	1	0.612	EBS only	-	Very Low	-
⊘	t2	t2.nano	1	0.5	EBS only	-	Low to Moderate	Yes
<input checked="" type="checkbox"/>	t2	t2.micro <span style="font-size: 8px; background-color: #90EE90; padding: 1px;">Free tier eligible</span>	1	1	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	t2	t2.small	1	2	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	t2	t2.medium	2	4	EBS only	-	Low to Moderate	Yes
<input type="checkbox"/>	t2	t2.large	2	8	EBS only	-	Low to Moderate	Yes

Cancel
Previous
Review and Launch
Next: Configure Instance Details

Once the instance size has been chosen, click “Next: Configure Instance Details”.

5. In the “Configure Instance Details” wizard, ensure that the following parameters are used:

- VPC: LK-VPC
- Subnet: LK-subnet-1
- Network Interface: Please enter 10.20.1.10 for the first instance (node-a).

**Note:** When creating node-b and node-c, please refer to [parameters related to node-b & node-c](#).

- 1. Choose AMI
- 2. Choose Instance Type
- 3. Configure Instance
- 4. Add Storage
- 5. Add Tags
- 6. Configure Security Group
- 7. Review

### Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

**Number of instances**  [Launch into Auto Scaling Group](#)

**Purchasing option**  Request Spot instances

**Network**  [Create new VPC](#)

**Subnet**  [Create new subnet](#)  
251 IP Addresses available

**Auto-assign Public IP**

**Placement group**  Add instance to placement group

**Capacity Reservation**

**Domain join directory**  [Create new directory](#)

**IAM role**  [Create new IAM role](#)

**CPU options**  Specify CPU options

**Shutdown behavior**

**Stop - Hibernate behavior**  Enable hibernation as an additional stop behavior

**Enable termination protection**  Protect against accidental termination

**Monitoring**  Enable CloudWatch detailed monitoring  
[Additional charges apply.](#)

**Tenancy**  [Additional charges will apply for dedicated tenancy.](#)

**Elastic Inference**  Add an Elastic Inference accelerator  
[Additional charges apply.](#)

**Credit specification**  Unlimited  
[Additional charges may apply](#)

**File systems**  [Create new file system](#)

▼ **Network interfaces**

Device	Network Interface	Subnet	Primary IP	Secondary IP addresses	IPv6 IPs
eth0	New network interface	subnet-0a5fb41b	10.20.1.10	Add IP	Add IP

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Storage](#)

Once these values are confirmed, click “Next: Add Storage”.

6. The wizard has already created the first storage device (volume). Click “Add New Volume” and create a new disk. Per the minimum requirements previously discussed in this guide, an 8GiB disk should be sufficient. However, this may differ depending on the volume of the data to be protected (replicated).

- 1. Choose AMI
- 2. Choose Instance Type
- 3. Configure Instance
- 4. Add Storage**
- 5. Add Tags
- 6. Configure Security Group
- 7. Review

### Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type <i>i</i>	Device <i>i</i>	Snapshot <i>i</i>	Size (GiB) <i>i</i>	Volume Type <i>i</i>	IOPS <i>i</i>	Throughput (MB/s) <i>i</i>	Delete on Termination <i>i</i>	Encryption <i>i</i>
Root	/dev/sda1	snap-0ea68ed57f7d8b84e	10	General Purpose S	100 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypte
<input type="text" value="EBS"/>	<input type="text" value="/dev/sdb"/>	<input type="text" value="Search (case-insensit)"/>	<input type="text" value="8"/>	General Purpose S	100 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypte

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Once the selection is confirmed, click “Next: Add Tags”.

- 7. Add a tag with key “Name” and value “Node-A” to make it easier to identify the instance from the AWS EC2 Console.

- 1. Choose AMI
- 2. Choose Instance Type
- 3. Configure Instance
- 4. Add Storage
- 5. Add Tags**
- 6. Configure Security Group
- 7. Review

### Step 5: Add Tags

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. A copy of a tag can be applied to volumes, instances or both. Tags will be applied to all instances and volumes. [Learn more](#) about tagging your Amazon EC2 resources.

Key (128 characters maximum)	Value (256 characters maximum)	Instances <i>i</i>	Volumes <i>i</i>	
<input type="text" value="Name"/>	<input type="text" value="Node-A"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="button" value="x"/>

(Up to 50 tags maximum)

Once the “Name” tag is defined, click “Next: Configure Security Group”.

- 8. Select the Security Group previously created in an earlier section (LK-SG):

- 1. Choose AMI
- 2. Choose Instance Type
- 3. Configure Instance
- 4. Add Storage
- 5. Add Tags
- 6. Configure Security Group
- 7. Review

### Step 6: Configure Security Group

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. [Learn more](#) about Amazon EC2 security groups.

Assign a security group:  Create a new security group  
 Select an existing security group

Security Group ID	Name	Description	Actions
<input type="checkbox"/> sg-07a2c0a284d482a07	default	default VPC security group	<a href="#">Copy to new</a>
<input checked="" type="checkbox"/> sg-097b642d5d064a5cf	LK-SG	Security Group for evaluating LifeKeeper	<a href="#">Copy to new</a>

Inbound rules for sg-097b642d5d064a5cf (Selected security groups: sg-097b642d5d064a5cf)

Type	Protocol	Port Range	Source	Description
All traffic	All	All	73.222.111.1/32	
All traffic	All	All	sg-097b642d5d064a5cf (LK-SG)	

[Cancel](#)
[Previous](#)
[Review and Launch](#)

9. Review and confirm all selections, then click “Launch”.

- 1. Choose AMI
- 2. Choose Instance Type
- 3. Configure Instance
- 4. Add Storage
- 5. Add Tags
- 6. Configure Security Group
- 7. Review

### Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

▼ AMI Details [Edit AMI](#)

**Red Hat Enterprise Linux (RHEL) 7 (HVM)**  
 Provided by Red Hat, Inc.  
 Free tier eligible  
 Root Device Type: ebs    Virtualization type: hvm

**Hourly Software Fees: \$0.00 per hour** on t2.micro instance. Additional taxes or fees may apply. Software charges will begin once you launch this AMI and continue until you terminate the instance.

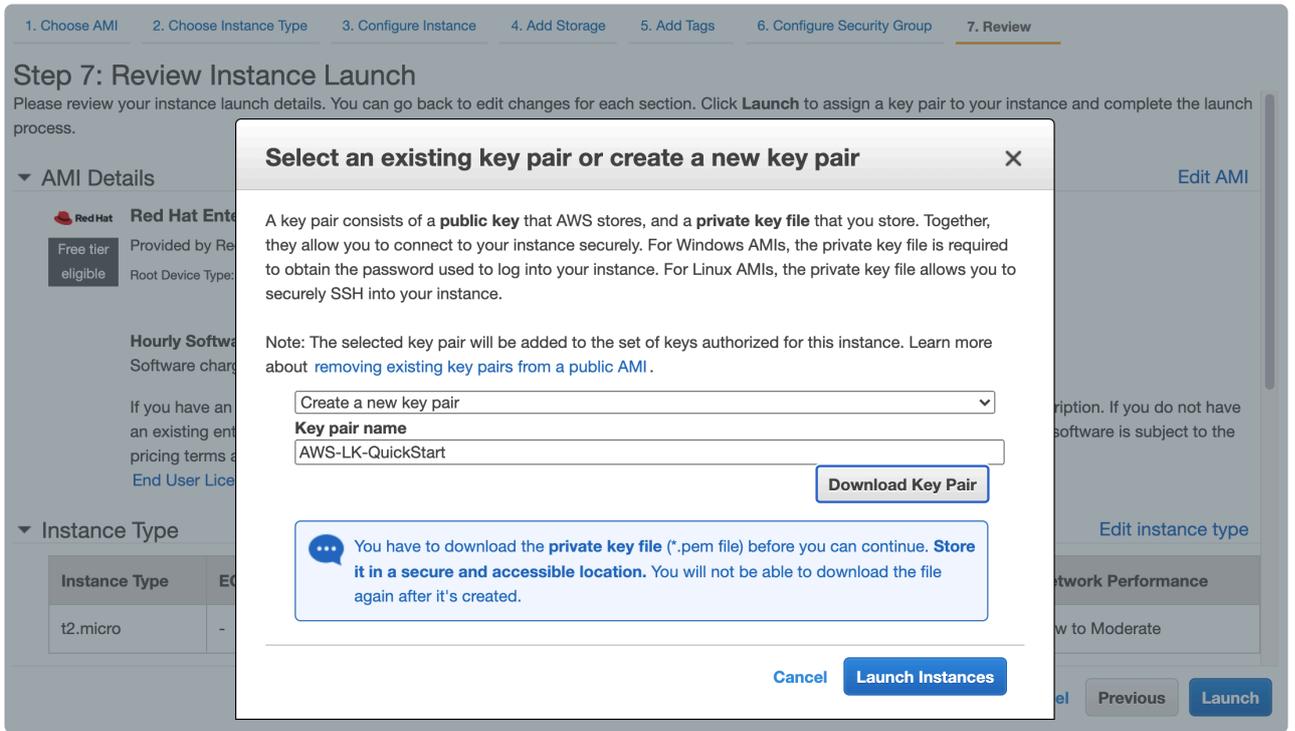
If you have an existing license entitlement to use this software, then you can launch this software without creating a new subscription. If you do not have an existing entitlement, then by launching this software, you will be subscribed to this software and agree that your use of this software is subject to the pricing terms and the seller's [End User License Agreement](#)

▼ Instance Type [Edit instance type](#)

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance
t2.micro	-	1	1	EBS only	-	Low to Moderate

[Cancel](#)
[Previous](#)
[Launch](#)

10. In order to access the newly created Linux EC2 instance, we will connect via ssh. AWS uses a key pair to authenticate user ssh sessions. Create a name for the key pair and download it to your local system. The private key will be needed to access the EC2 instance. Click “Launch Instances”.



11. The “Launch Status” page appears. Select the instance ID to view its details.

## Launch Status

**Your instances are now launching**  
The following instance launches have been initiated: [i-0f24161b65ff230f5](#) [View launch log](#)

**Get notified of estimated charges**  
Create [billing alerts](#) to get an email notification when estimated charges on your AWS bill exceed an amount you define (for example, if you exceed the free usage tier).

### How to connect to your instances

Your instances are launching, and it may take a few minutes until they are in the **running** state, when they will be ready for you to use. Usage hours on your new instances will start immediately and continue to accrue until you stop or terminate your instances.

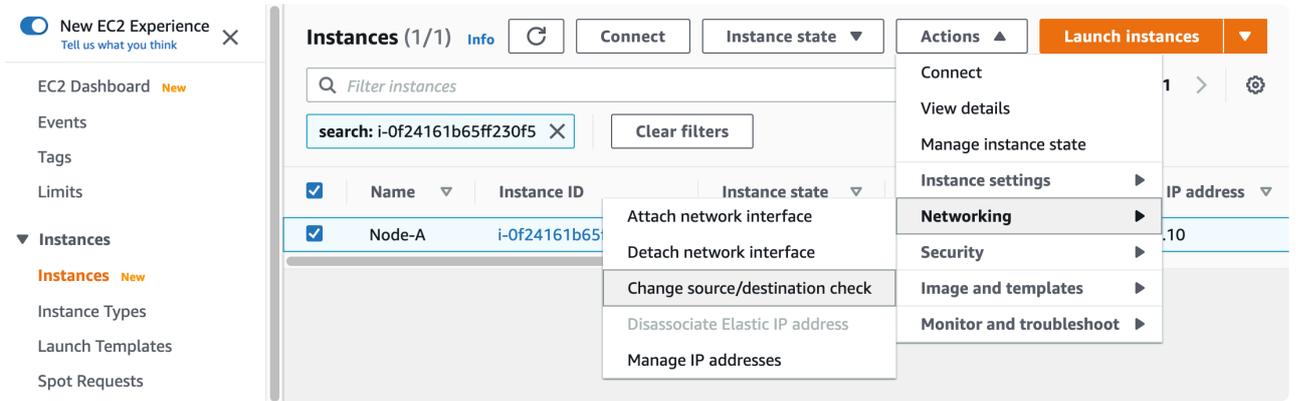
Click **View Instances** to monitor your instances' status. Once your instances are in the **running** state, you can **connect** to them from the Instances screen. [Find out](#) how to connect to your instances.

### Getting started with your software

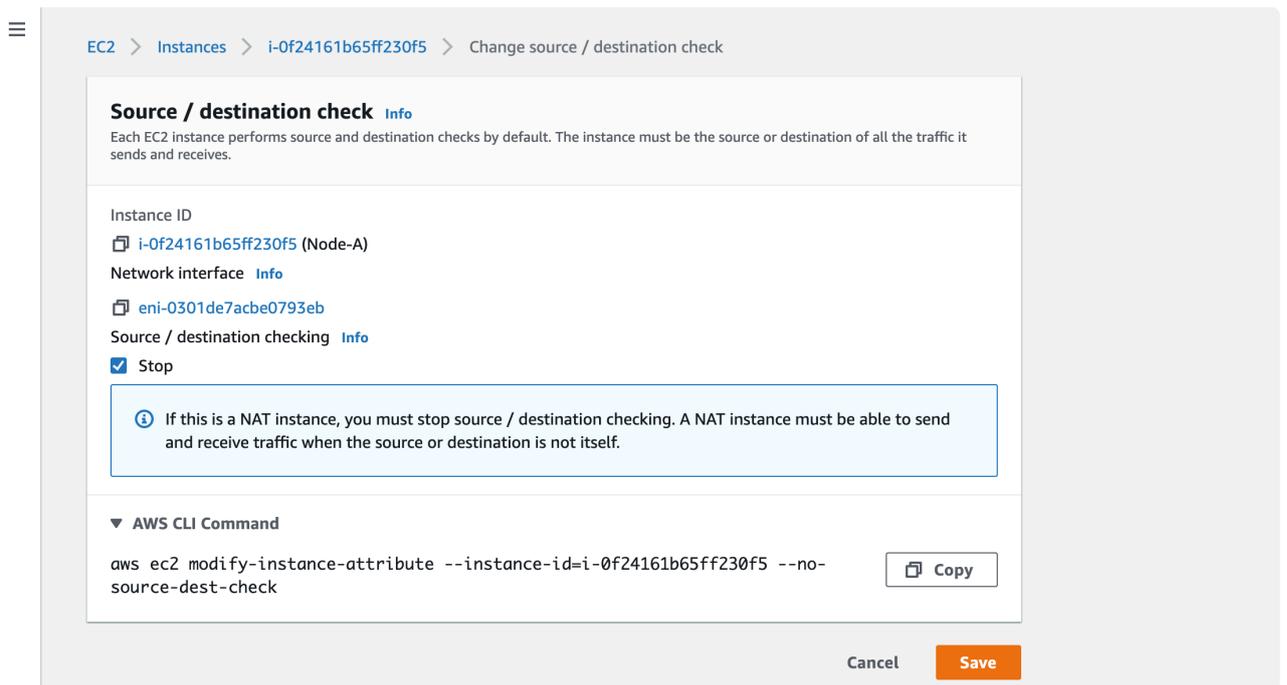
To get started with Red Hat Enterprise Linux (RHEL) 7 (HVM) [View Usage Instructions](#)

To manage your software subscription [Open Your Software on AWS Marketplace](#)

12. Once the Instance is created we need to change the network configuration. As the active node changes from time to time and we are using a Virtual IP address, we should disable the source/destination check on the network interface. To do this, select the Instance, then select Actions > Networking > Change source/destination check.

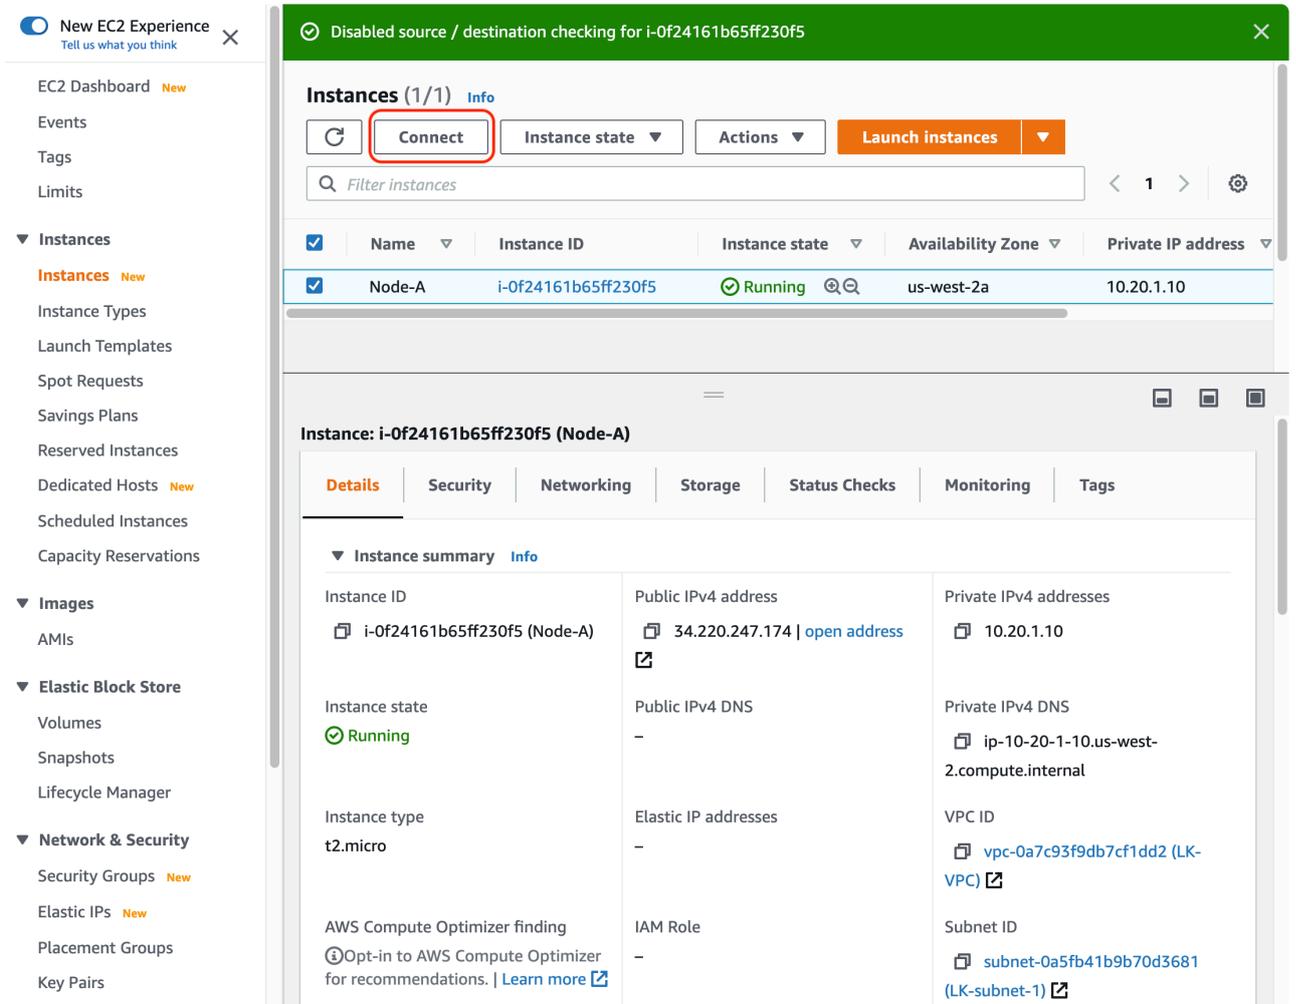


13. On the “Source / destination check” page, check “Stop” and save the change.

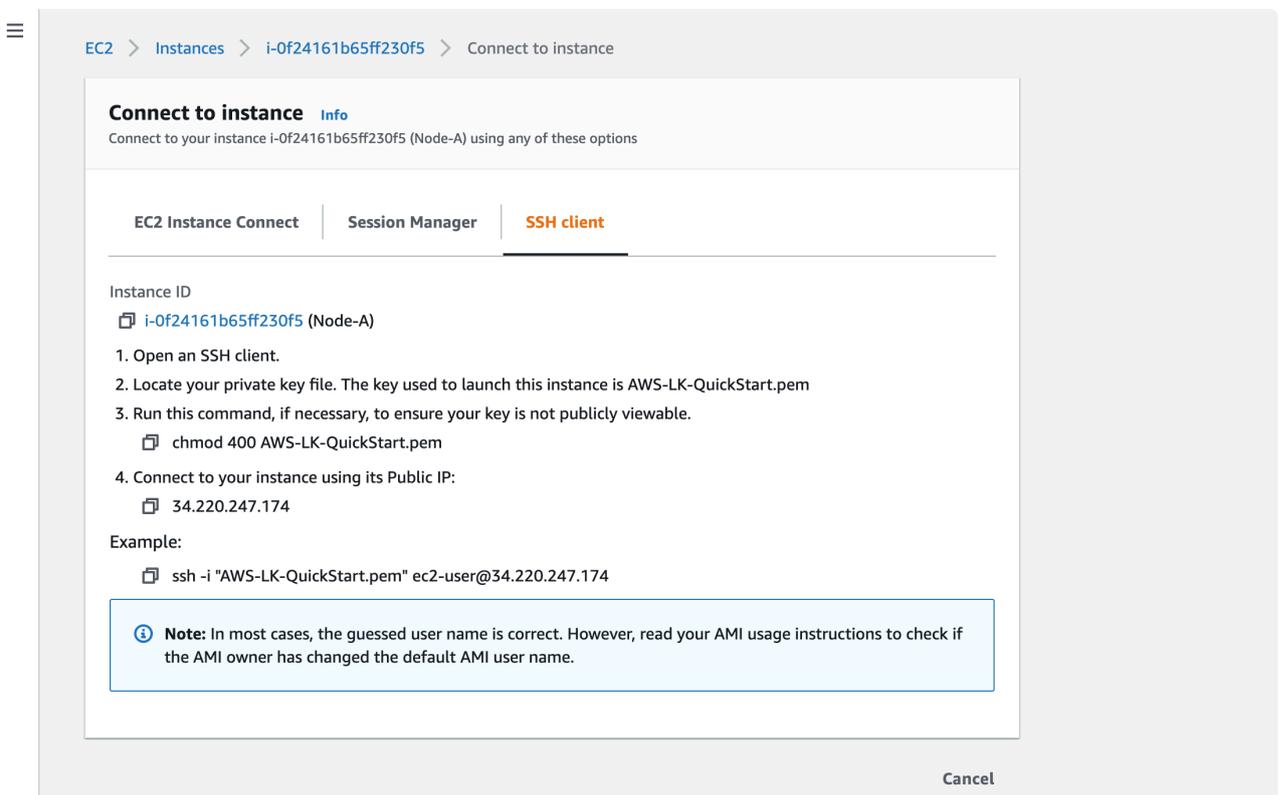


Now the instance is created and ready for us to connect.

14. The details of the instance may be reviewed on this page. Select “Connect” at the top of the page to see instructions on how to connect.



15. Here we can view the instructions that explain how to connect to the instance. On a Windows client, please refer to [Setup X Window client software on Microsoft Windows](#) for details.

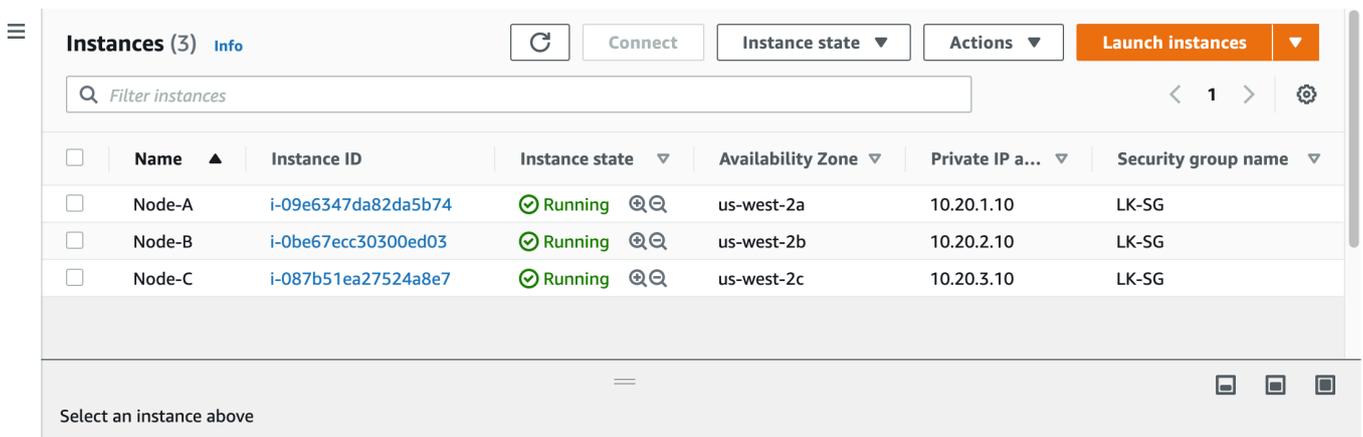


# 11.2.3.3.9. Creating the Second and Third Instances

The steps required to create the second node (node-b) and third node (node-c) are almost the same as previously described for the first node (node-a). The following table illustrates the differences between the three nodes and the details required to create these instances.

Name	Parameter	Value
 <b>Common values across instances</b>	VPC	LK-VPC
	Security Group	LK-SG
	Source/Dest Checking on a network interface (ENI)	Disabled
node-a	Subnet	LK-subnet-1
	Private IP Address	10.20.1.10
	Second Disk (Storage Device)	You need to have a second disk
node-b	Subnet	LK-subnet-2
	Private IP Address	10.20.2.10
	Second Disk (Storage Device)	You need to have a second disk
node-c	Subnet	LK-subnet-3
	Private IP Address	10.20.3.10
	Second Disk (Storage Device)	You do not need to attach a second disk

Once created, the list of instances should look like this:



The screenshot shows the AWS Management Console 'Instances' page. At the top, there are buttons for 'Refresh', 'Connect', 'Instance state', 'Actions', and 'Launch instances'. A search bar labeled 'Filter instances' is present. Below the search bar is a table with the following columns: Name, Instance ID, Instance state, Availability Zone, Private IP address, and Security group name. Three instances are listed:

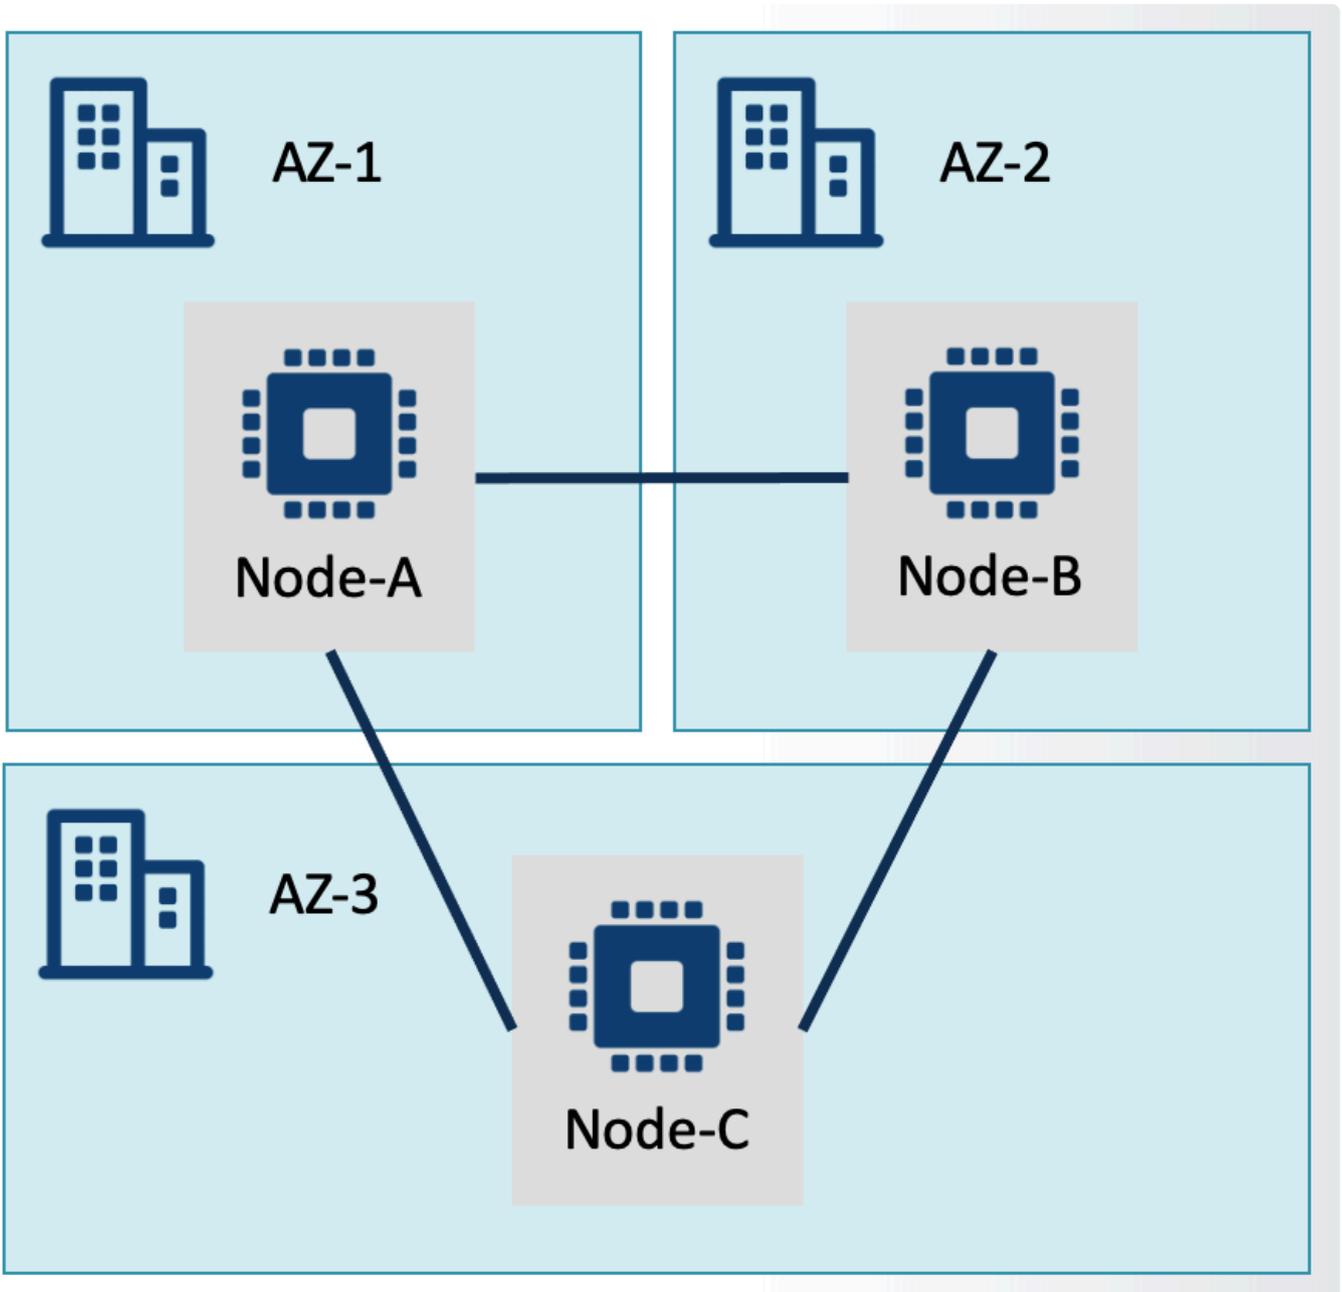
Name	Instance ID	Instance state	Availability Zone	Private IP address	Security group name
Node-A	i-09e6347da82da5b74	Running	us-west-2a	10.20.1.10	LK-SG
Node-B	i-0be67ecc30300ed03	Running	us-west-2b	10.20.2.10	LK-SG
Node-C	i-087b51ea27524a8e7	Running	us-west-2c	10.20.3.10	LK-SG

At the bottom of the console, there is a message: 'Select an instance above'.

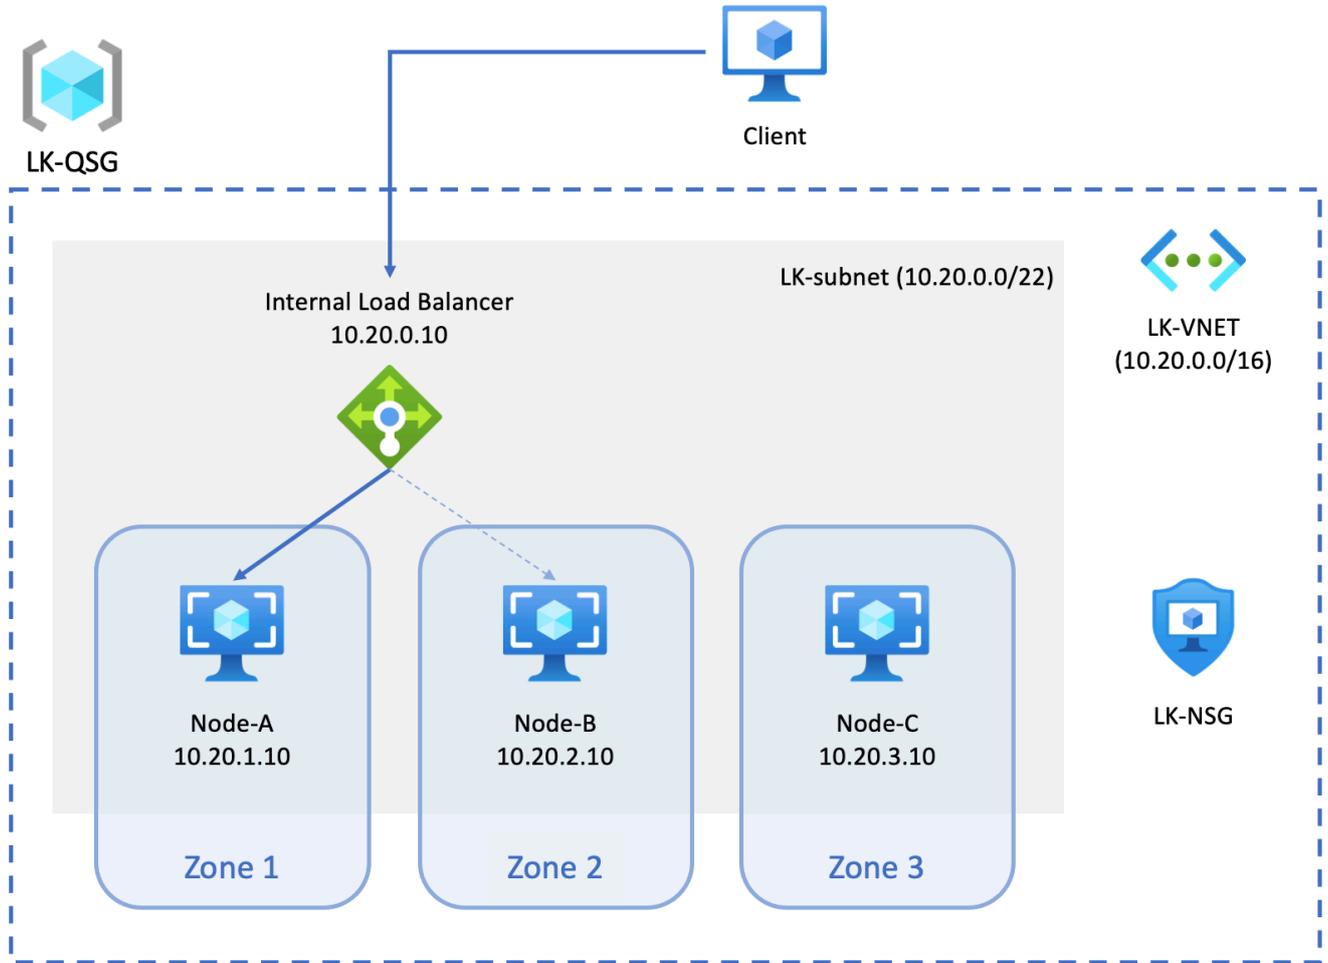
## 11.2.3.4. Creating an Instance in Azure from Scratch

**!** **Disclaimer:** The user interface may vary between regions or change over time. Please refer to the documentation provided by Azure if the screenshots shown below are different from your experience.

The Evaluation Guide uses the following network structure and instances.



On Azure, these components can be defined as seen in the diagram below.



In this section we will create these components with the exception of the Internal Load Balancer. We will create the Internal Load Balancer [later](#).

Components listed in this diagram are described in the following table.

To create this network structure, the following components need to be implemented:

Component	Name	Parameter	Value
Resource Group	LK-QSG		
Virtual network	LK-VNET	IPv4 CIDRs	10.20.0.0/16
		Subnet (LK-subnet)	10.20.0.0/22
Network Security	LK-NSG	Additional Inbound Rule Type=SSH	Source=Your Office's WAN IP Destination=Any

		Additional Outbound Rule Type=Any	Source=VirtualNetwork Destination=AzureLoadBalancer
Virtual Machine	 Common values across VM	Resource Group	LK-QSG
		Virtual network	LK-VNET
		Subnet	LK-subnet
		Network Security Group	LK-NSG
	node-a	Availability Zone (*1)	1
		Private IP Address (*2)	10.20.1.10
	node-b	Availability Zone (*1)	2
		Private IP Address (*2)	10.20.2.10
	node-c	Availability Zone (*1)	3
		Private IP Address (*2)	10.20.3.10

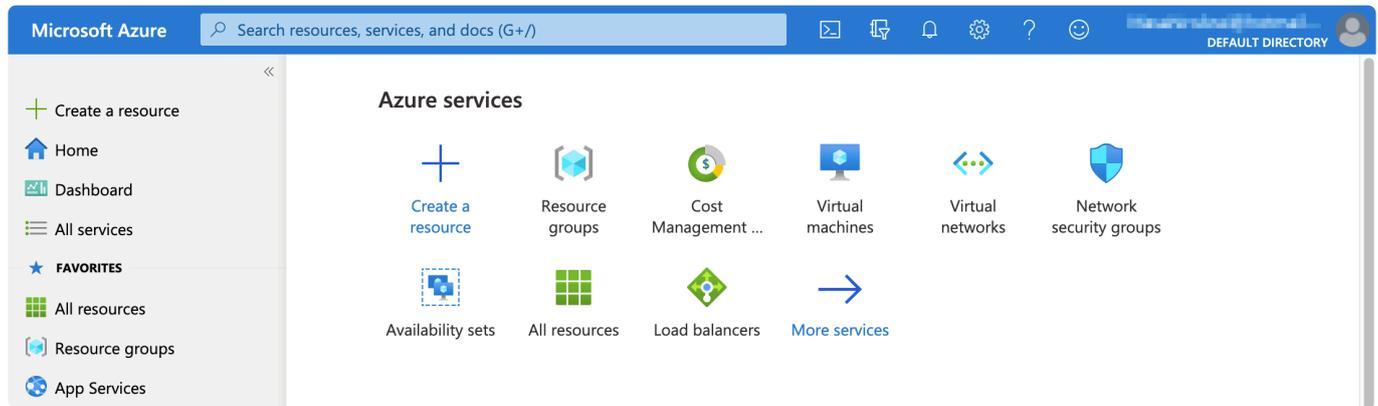
- (\*1) This MUST be chosen at the time of creating a virtual machine.
- (\*2) Wizard will automatically assign a DHCP based IP address. This will need to be modified after creating a virtual machine.

The following sections describe the steps necessary for creating these network components, beginning with the first instance.

- [Switching between Azure Services](#)
- [Deciding on an Azure Region](#)
- [Creating the Resource Group](#)
- [Creating a Virtual Network](#)
- [Creating a Network Security Group](#)
- [Creating the First Azure Virtual Machine](#)
- [Creating the Second and Third Virtual Machines](#)

## 11.2.3.4.1. Switching between Azure Services

Azure has a significant and growing number of services available. Typical services that you would use are accessible from the navigation bar located on the left hand side or listed in the home screen (select “Home” at the navigation bar).



If you are unable to find a particular service, you may type in the name of the service in the search box at the top.

## 11.2.3.4.2. Deciding on an Azure Region

---

Azure has regions in many geographic locations. It may be beneficial to select a region in close geographic proximity to your workplace location. Note that not all regions have 3 Availability Zones, so exercise caution when making a selection in order to guarantee that the chosen region supports the intended configurations. When you create a Resource Group or other resources in later steps, please make sure to select the correct region.

Refer to the [Azure documentation](#) for more information.

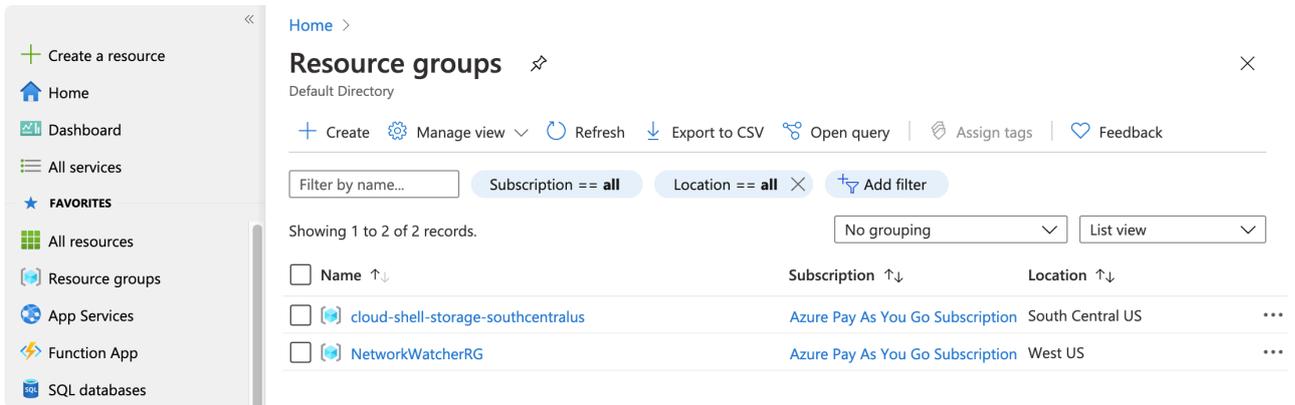
# 11.2.3.4.3. Creating the Resource Group

A Resource Group is the logical container in which you can create resources such as network components, computer resources, storage, etc.

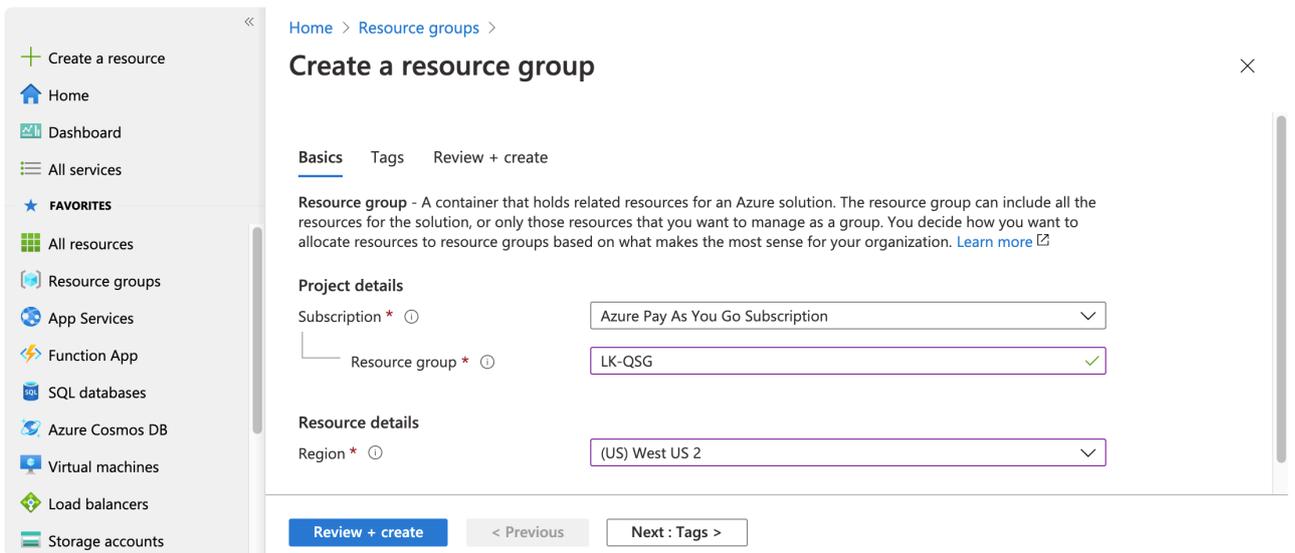
Refer to the [Azure documentation](#) for more information.

In this section, we will create a Resource Group for testing LifeKeeper (we will name it LK-QSG) as follows:

1. Select “Resource groups” from the home screen to browse the list of existing resource groups, then click “Create” located at the top.

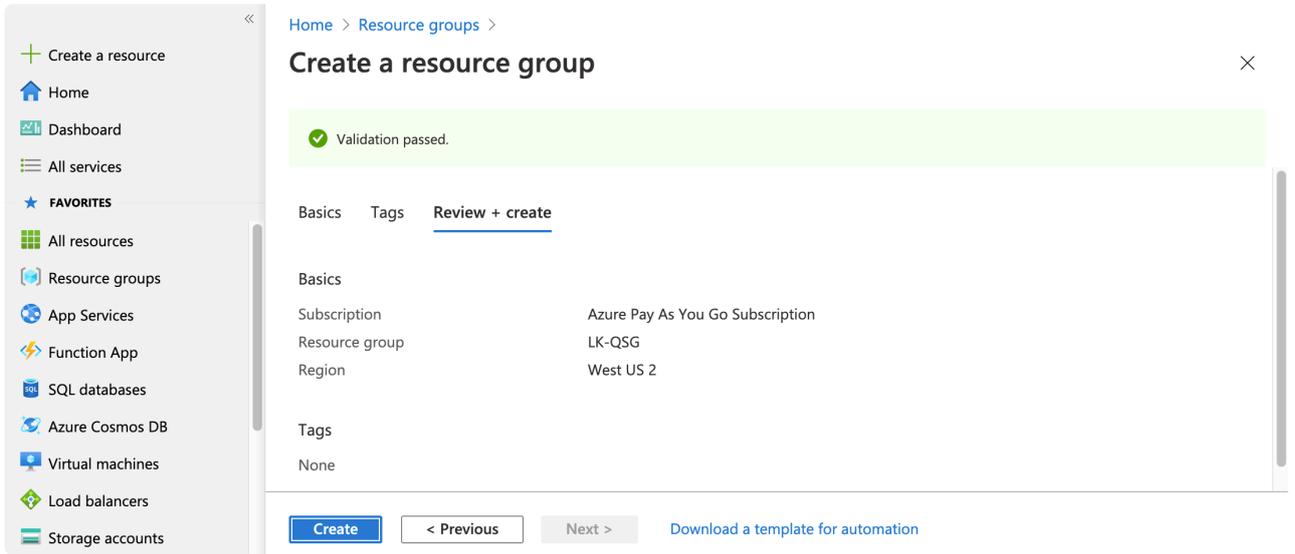


2. Specify the name of Resource group as LK-QSG, then select the Region to create the resource group. Note that the region used to create the resource group must have 3 availability zones.



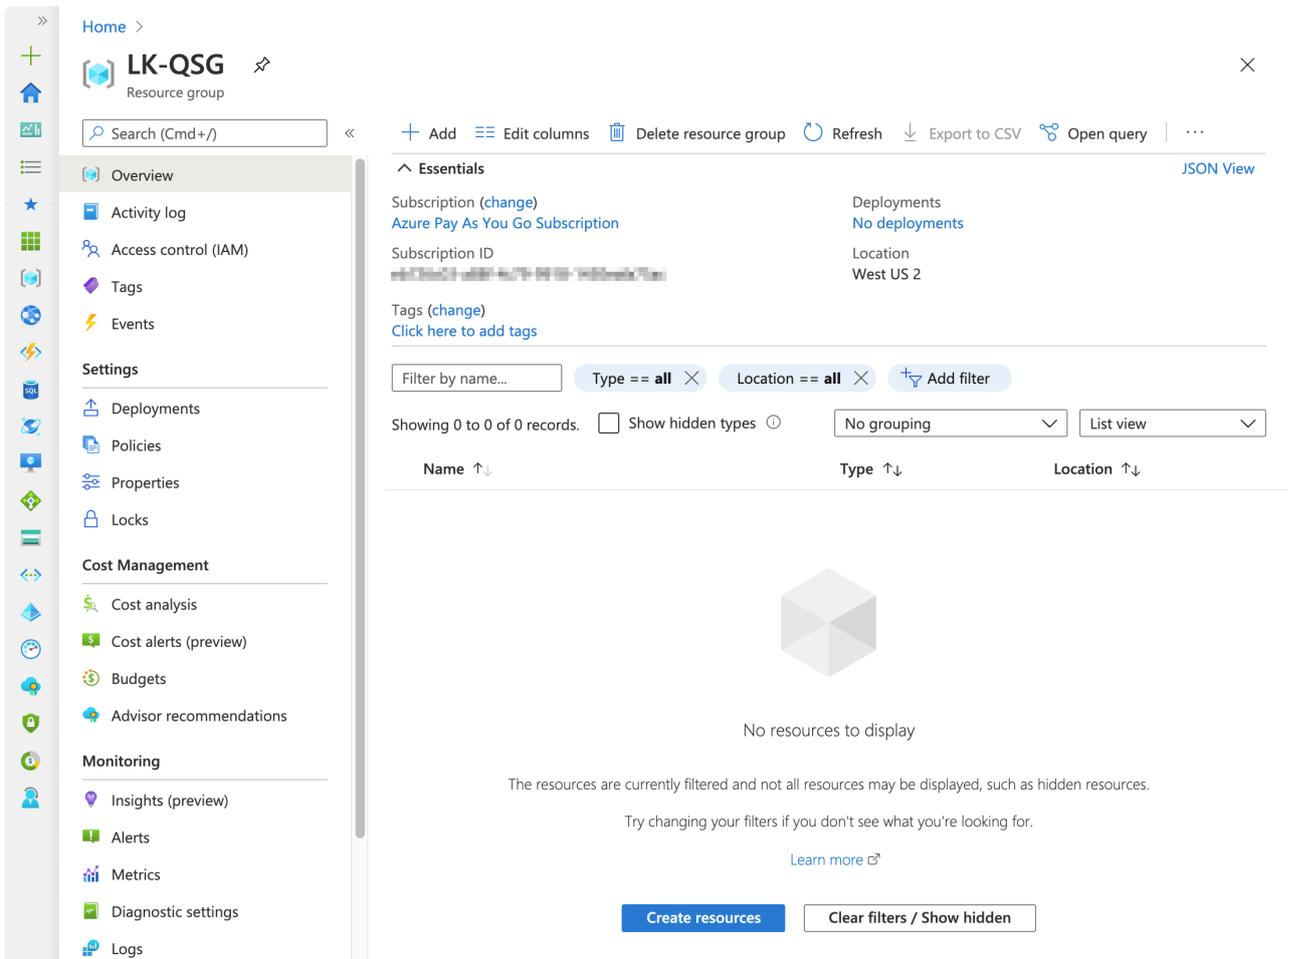
Once you make the selections, click “Review + Create”.

3. The wizard evaluates these values and you can now create the resource if the validation passes.



Click "Create" to create the resource group.

4. Now the resource group is created.



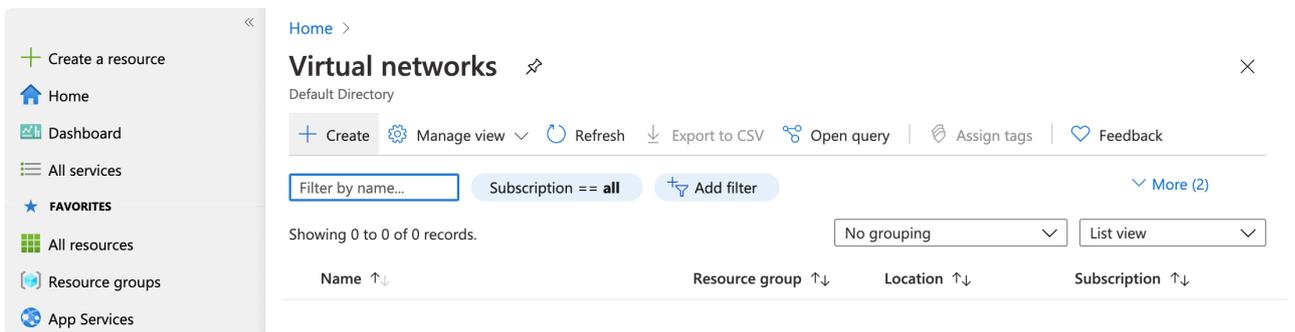
# 11.2.3.4.4. Creating a Virtual Network

The Virtual Network is an Azure resource that represent a local network. Different Virtual Networks can be defined within the Azure cloud to logically separate different systems.

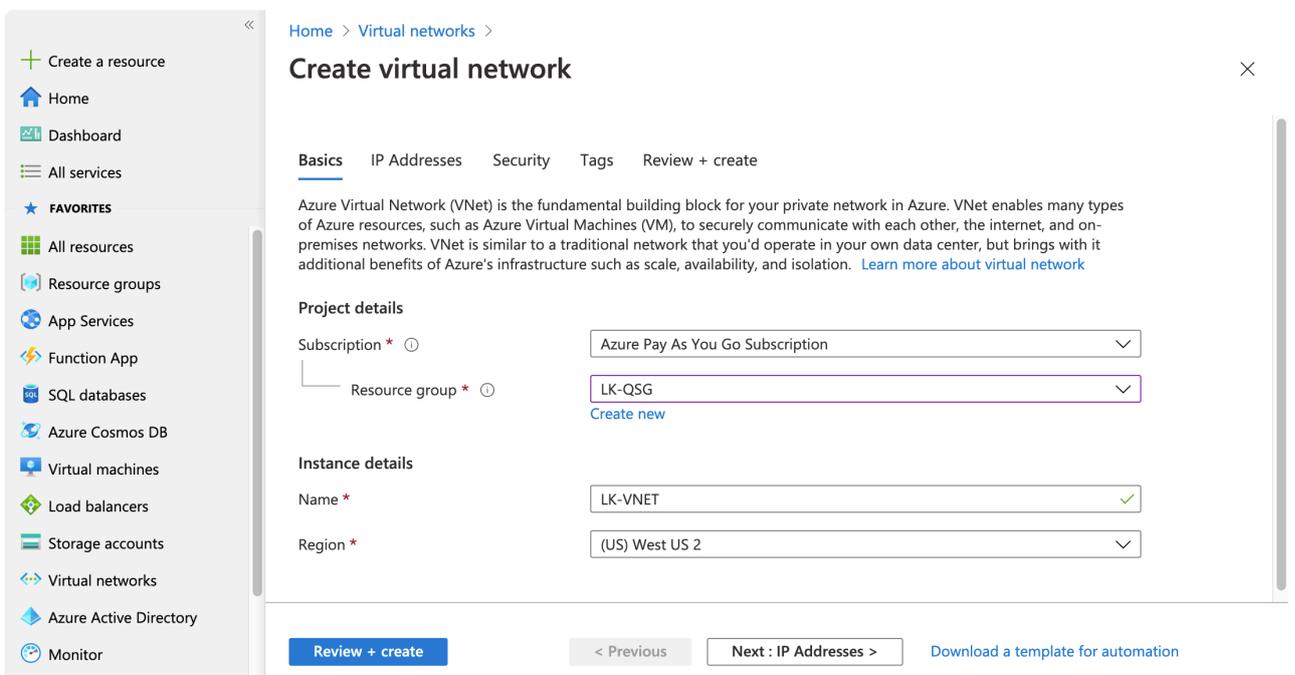
Refer to the [Azure documentation](#) for more information.

In this section we will create a Virtual Network for testing LifeKeeper (we will name it LK-VNET) as follows:

1. Select “Virtual networks” from the home screen to see current list of virtual networks, then click “Create” located at the top.

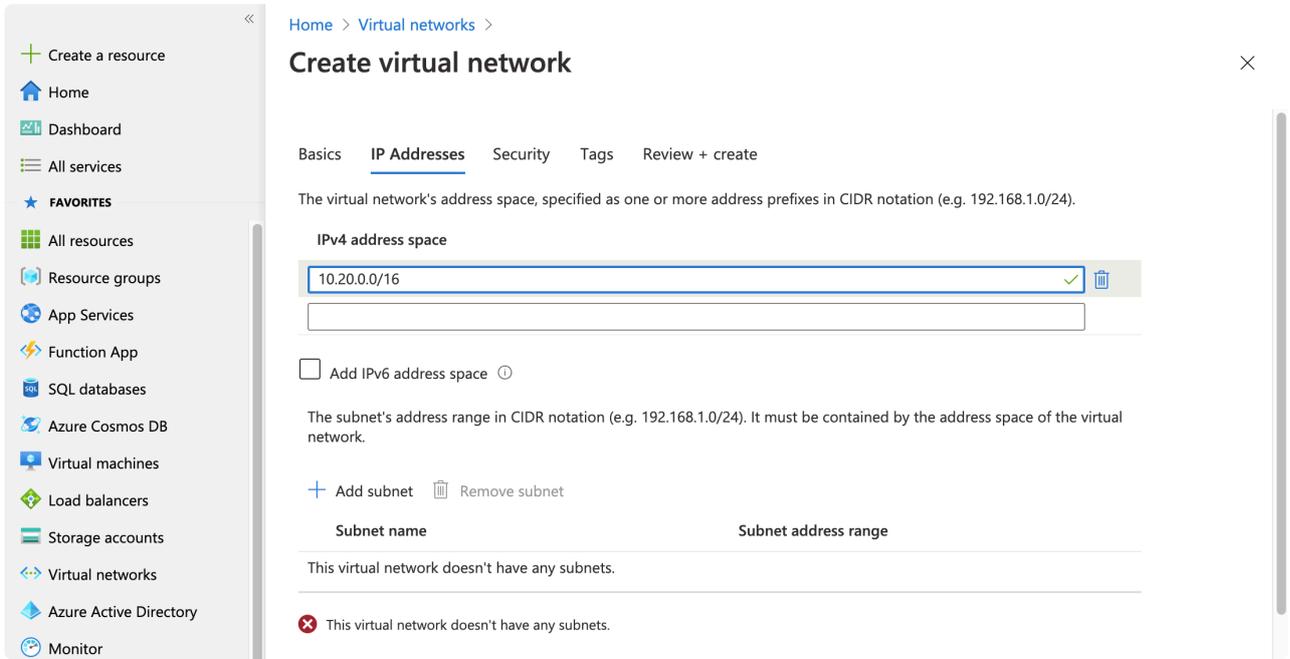


2. In the “Virtual networks” wizard, select the resource group LK-QSG, the virtual network LK-VNET and the region to create the virtual network.

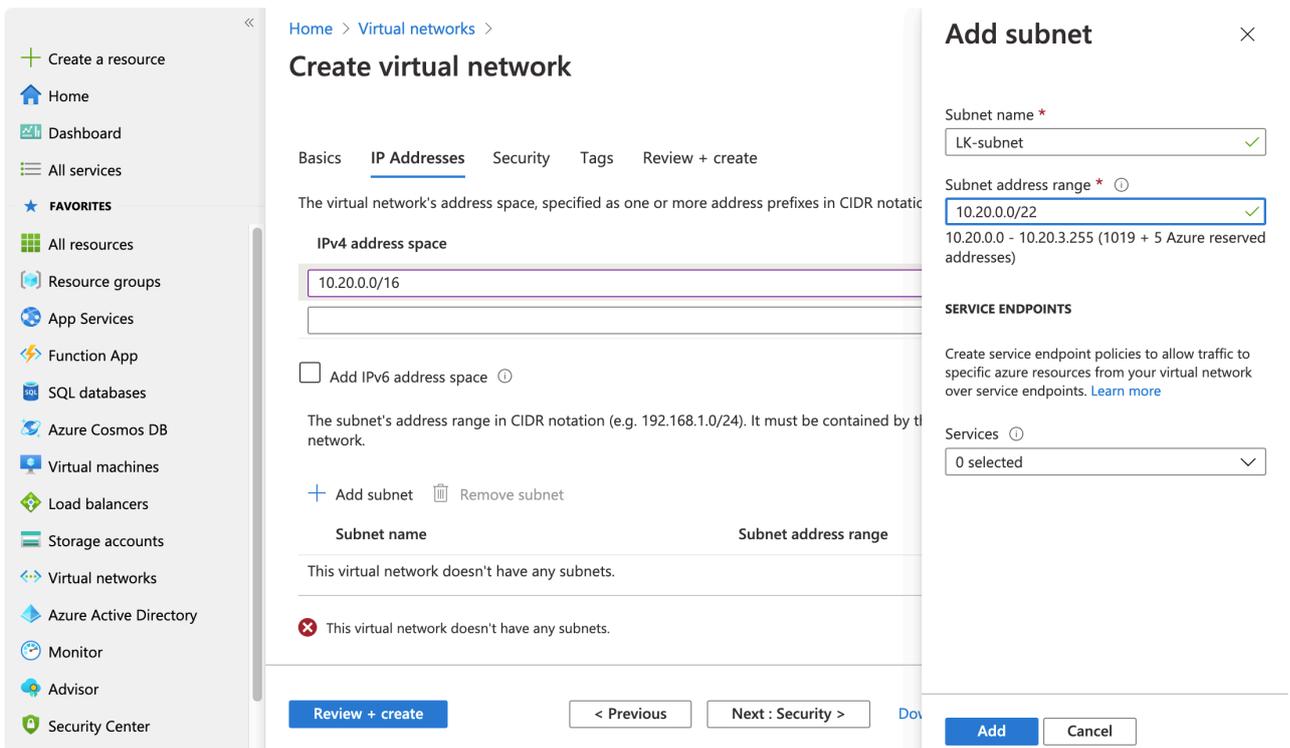


Once you enter the values, click “Next: IP Addresses”.

3. Enter the CIDR block for this virtual network as 10.20.0.0/16. Once you enter the CIDR block, click “Add subnet” to define a subnet.



4. Specify the name of the subnet LK-subnet and its CIDR block 10.20.0.0/22.



Once you specify them, click “Add” at the bottom.

5. Now parameters for both the virtual network and the subnet are defined. Click “Review + create”.

The screenshot shows the 'Create virtual network' wizard in the 'IP Addresses' step. The left sidebar contains navigation options like 'Home', 'Dashboard', and 'All services'. The main content area has tabs for 'Basics', 'IP Addresses', 'Security', 'Tags', and 'Review + create'. A text box contains the IPv4 address space '10.20.0.0/16'. Below it, there is an unchecked checkbox for 'Add IPv6 address space'. A table lists a subnet named 'LK-subnet' with an address range of '10.20.0.0/22'. At the bottom, there are buttons for 'Review + create', '< Previous', 'Next : Security >', and a link to 'Download a template for automation'.

6. Once the wizard validates the parameters, click “Create” to create a virtual network.

The screenshot shows the 'Create virtual network' wizard in the 'Review + create' step. A green banner at the top indicates 'Validation passed'. The 'Review + create' tab is active. The 'Basics' section shows: Subscription (Azure Pay As You Go Subscription), Resource group (LK-QSG), Name (LK-VNET), and Region (West US 2). The 'IP addresses' section shows: Address space (10.20.0.0/16) and Subnet (LK-subnet (10.20.0.0/22)). The 'Tags' section shows 'None'. At the bottom, there is a 'Create' button, '< Previous', 'Next >', and a link to 'Download a template for automation'.

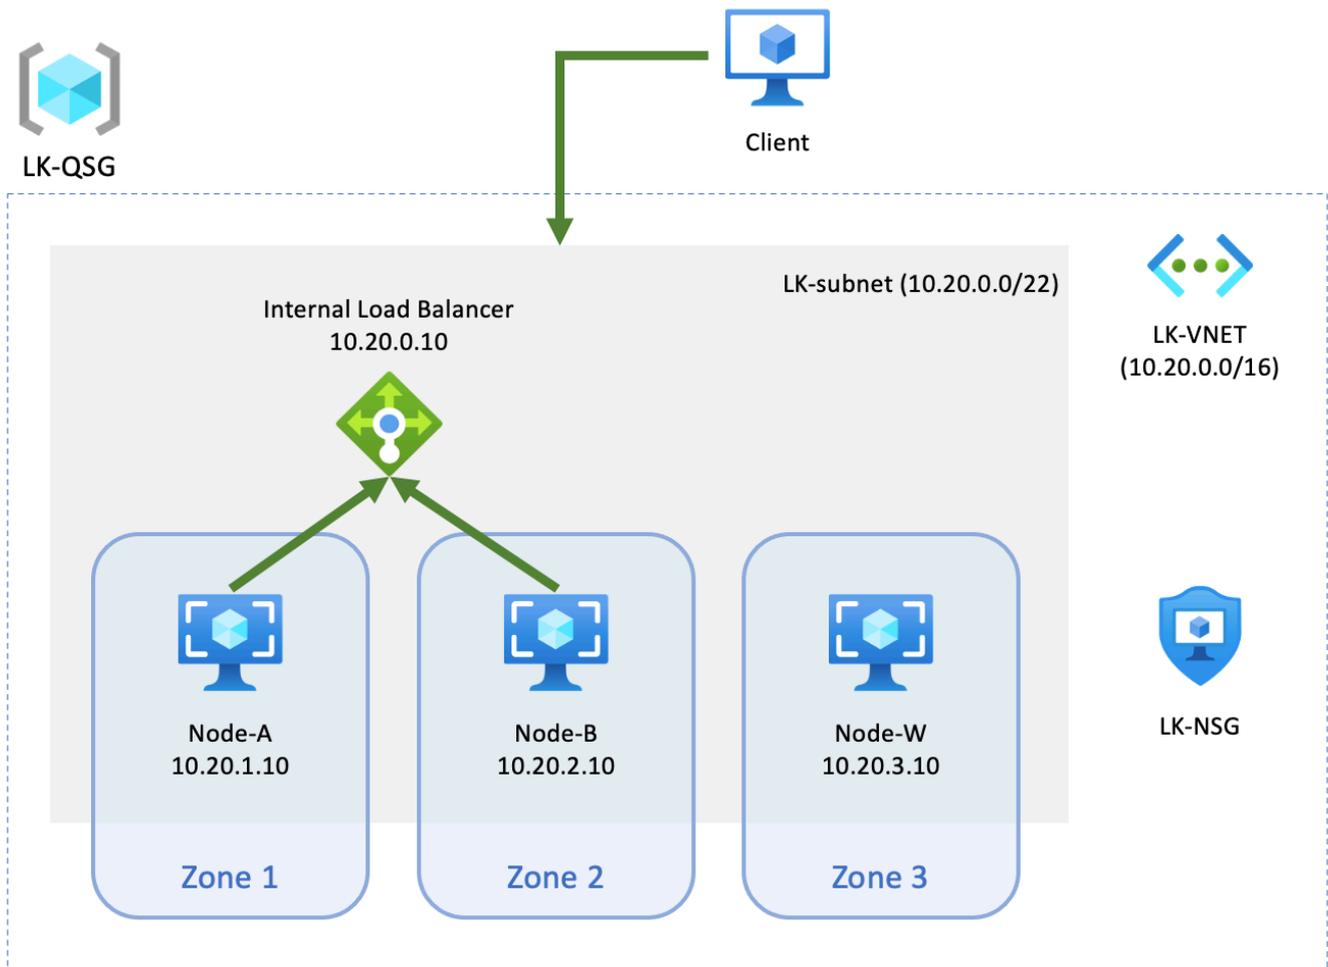
# 11.2.3.4.5. Creating a Network Security Group

A Security Group works as a firewall, allowing you to define both “allow” and “deny” rules. A source can be an internet address or a security group, and a security group can be assigned to a Virtual Machine or a subnet.

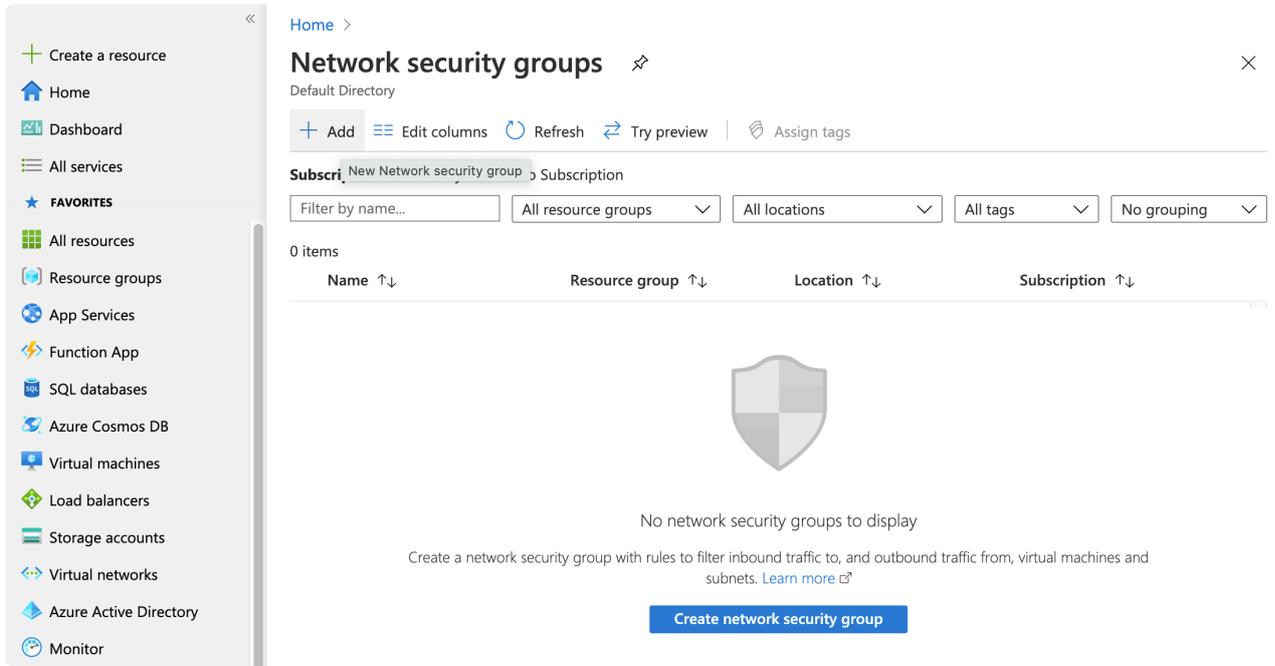
Refer to the [Azure documentation](#) for more information.

A new security group already has several “allow” and “deny” rules (those are defined with priority 65xxx). We will add following additional rules to “allow” access (see the green arrows).

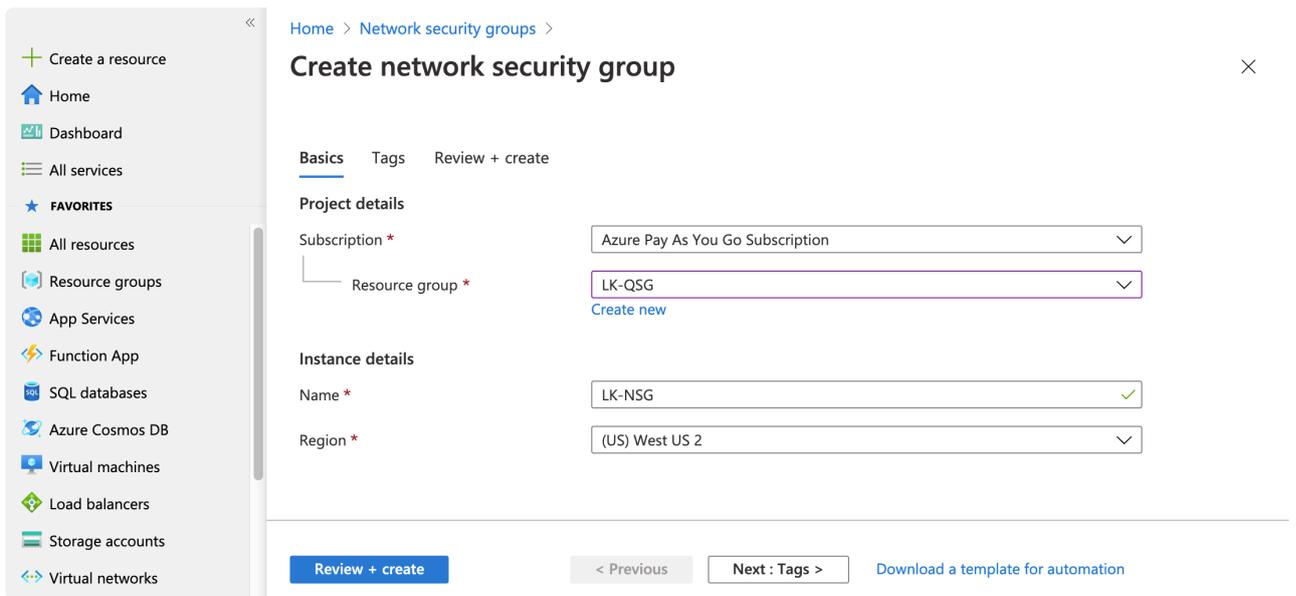
- Allow traffic from the work location to LK-subnet
- Allow traffic from Virtual Machines to Internal Load Balancer



1. Select “Network Security Group” from the home screen and click “Add” to create a new one.

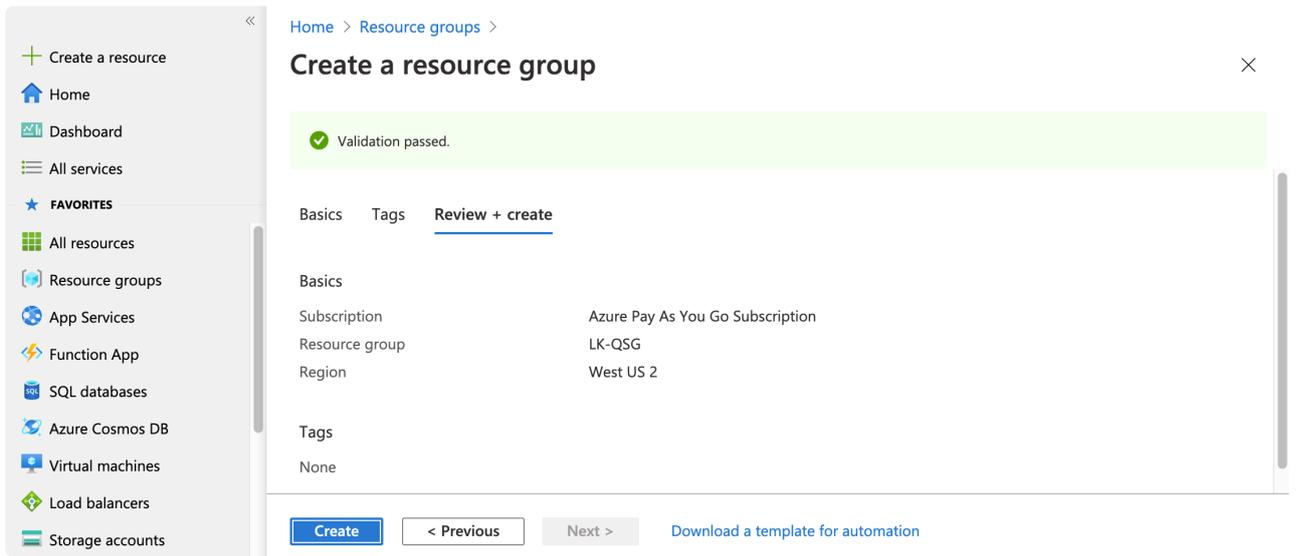


2. Select the Resource Group LK-QSG, specify the name of new security group as LK-NSG, then select a region.

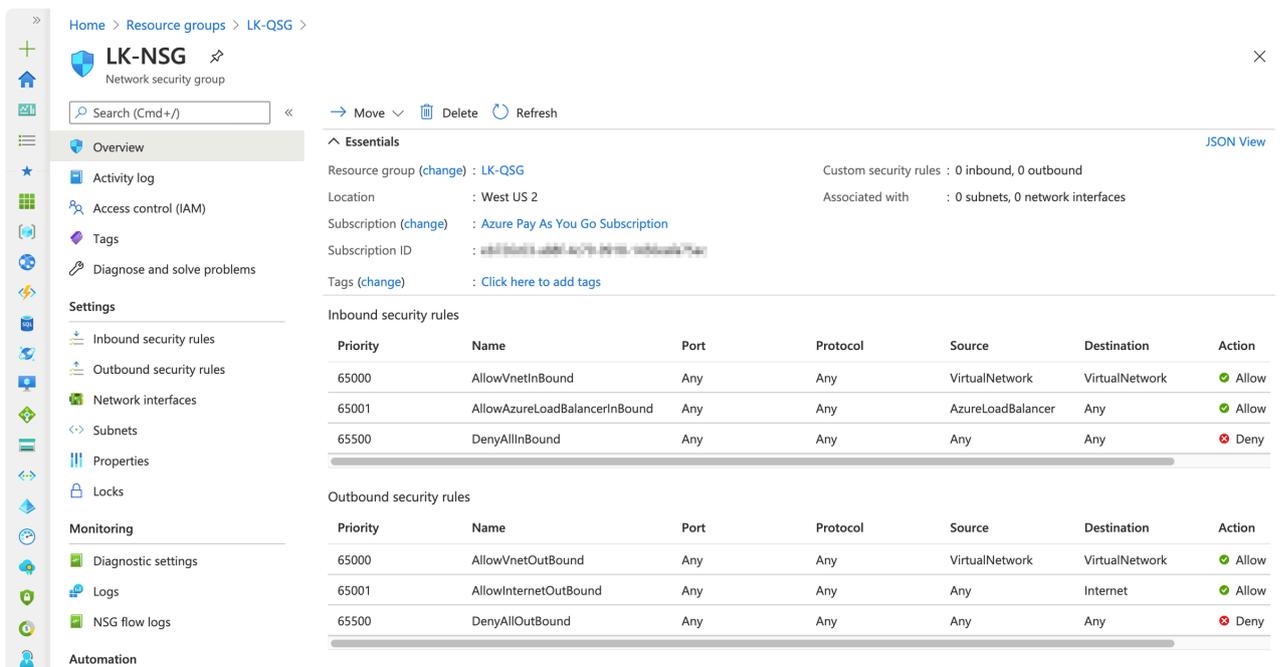


Once you specify the values, click “Review + create”.

3. Once the wizard validates your input, click “Create”.



4. The new security group LK-NSG is now created. You can see that there are several inbound and outbound rules already defined.



5. Select the Inbound security rules from the left side and click “Add” to enter the new rule. Enter the following values:

- Source IP: Your WAN IP address(es)
- Source Port Range: \*
- Destination: Virtual Network
- Destination Port Ranges: \*
- Name: AllowAccessFromWork (any name that represents this configuration is fine)

The screenshot displays the Azure portal interface for configuring a Network Security Group (NSG) rule. On the left, a table lists existing inbound security rules:

Priority	Name	Port
65000	AllowVnetInBound	Any
65001	AllowAzureLoadBalancerInBound	Any
65500	DenyAllInBound	Any

The right pane shows the 'Add inbound security rule' configuration form for LK-NSG. The configuration details are as follows:

- Source:** IP Addresses
- Source IP addresses/CIDR ranges:** 73.222.0.0/16
- Source port ranges:** \*
- Destination:** VirtualNetwork
- Destination port ranges:** \*
- Protocol:** Any (with options for TCP, UDP, ICMP)
- Action:** Allow (with option for Deny)
- Priority:** 100
- Name:** AllowAccessFromWork
- Description:** (empty)

An 'Add' button is located at the bottom of the configuration form.

After you enter the values, click “Add”.

6. Select Outbound security rules from the left side and click “Add” to enter the new rule. Enter the following values:

- Source: VirtualNetwork
- Source port range: \*
- Destination: Service Tag
- Destination Service Tag: AzureLoadBalancer
- Destination port ranges: \*
- Name: AllowTrafficFromInstancesToLoadBalancer (any name that represents this configuration is fine)

The screenshot displays the Azure portal interface for configuring a Network Security Group (NSG) rule. On the left, a table lists existing outbound security rules:

Priority	Name	Port
65000	AllowVnetOutBound	Any
65001	AllowInternetOutBound	Any
65500	DenyAllOutBound	Any

The right pane shows the 'Add outbound security rule' configuration form for LK-NSG. The configuration includes:

- Source:** VirtualNetwork
- Source port ranges:** \*
- Destination:** Service Tag
- Destination service tag:** AzureLoadBalancer
- Destination port ranges:** \*
- Protocol:** Any (selected), TCP, UDP, ICMP
- Action:** Allow (selected), Deny
- Priority:** 100
- Name:** AllowTrafficFromInstancesToLoadBalancer
- Description:** This allows to send traffic from VM to ILB

An 'Add' button is located at the bottom of the configuration pane.

After you enter the values, click “Add”.

7. Now the new Security group is defined.



Home >

# LK-NSG

Network security group

» [Move](#) [Delete](#) [Refresh](#)

## Essentials

Resource group (change) : LK-QSG

Custom security rules : 1 inbound, 1 outbound

Location : West US 2

Associated with : 0 subnets, 0 network interfaces

Subscription (change) : Azure Pay As You Go Subscription

Subscription ID : sub-2706c0d8-4a088-4c70-8088-63f0b6e4d70a

Tags (change) : [Click here to add tags](#)

### Inbound security rules

Priority	Name	Port	Protocol	Source	Destination	Action
100	AllowAccessFromWork	Any	Any	73.222.111.111	VirtualNetwork	✔ Allow
65000	AllowVnetInBound	Any	Any	VirtualNetwork	VirtualNetwork	✔ Allow
65001	AllowAzureLoadBalancerInBou...	Any	Any	AzureLoadBalancer	Any	✔ Allow
65500	DenyAllInBound	Any	Any	Any	Any	✘ Deny

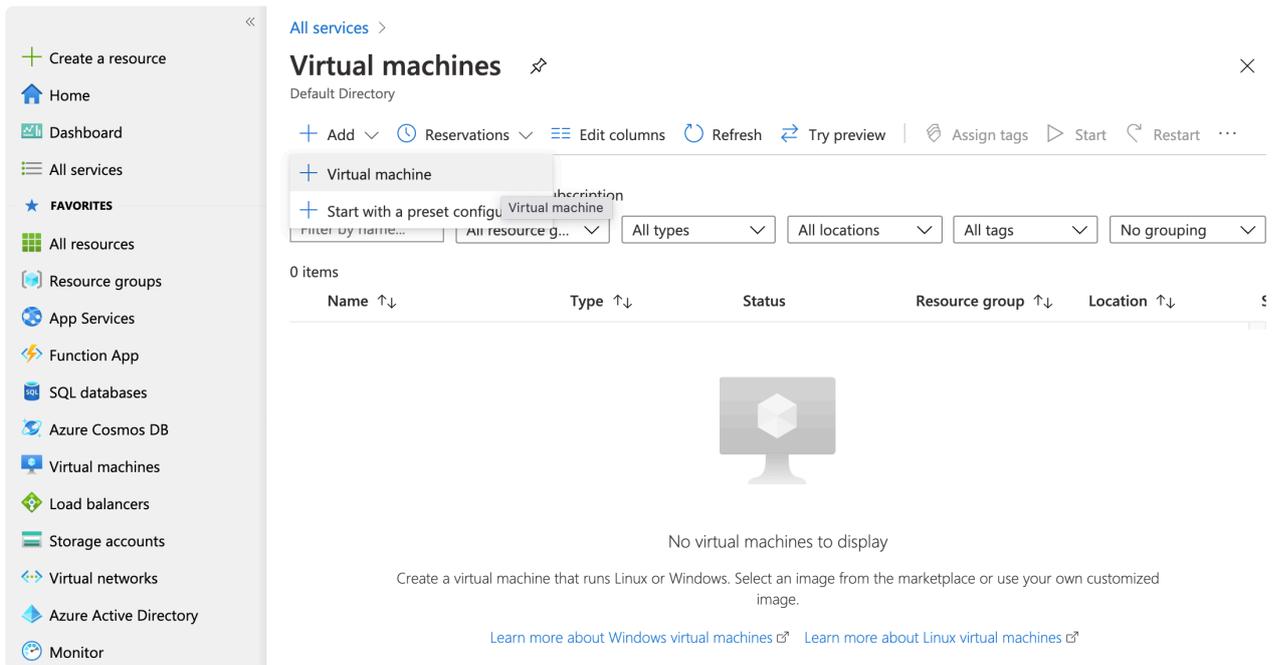
### Outbound security rules

Priority	Name	Port	Protocol	Source	Destination	Action
100	AllowTrafficFromInstancesToL...	Any	Any	VirtualNetwork	AzureLoadBalancer	✔ Allow
65000	AllowVnetOutBound	Any	Any	VirtualNetwork	VirtualNetwork	✔ Allow
65001	AllowInternetOutBound	Any	Any	Any	Internet	✔ Allow
65500	DenyAllOutBound	Any	Any	Any	Any	✘ Deny

## 11.2.3.4.6. Creating the First Azure Virtual Machine

In previous sections we covered configuration of the network. Now we are going to create the first instance. As discussed in [Computing Resources Used in this Tutorial](#), we need two disks. This section also discusses how to create the second disk.

1. Go to “Virtual Machine” and click “Add” to create a new VM.



2. Enter the following values:

- Resource group: LK-QSG
- Virtual machine name: `node-a` (when you create 2nd or 3rd nodes, this would be `node-b` and `node-c`)
- Region: Specify region based on your decision earlier
- Availability options: Availability Zone
- Availability Zone: 1 (when you create 2nd or 3rd nodes, this would be 2 and 3)
- Image: Select supported operating system
- Size: You can select image size based on the application you are going to protect. Minimum Requirement here is `Standard_B1s`
- Authentication type: You can select each option. In general, cloud providers recommend using SSH public key.
- username: `azureuser` (you can leave the default value)
- SSH public key source: `Generate new key pair` if you have not created a key before
- Key pair name: specify name (this screenshot uses `Azure-LK-QuickStart`)
- Public inbound ports: `None`

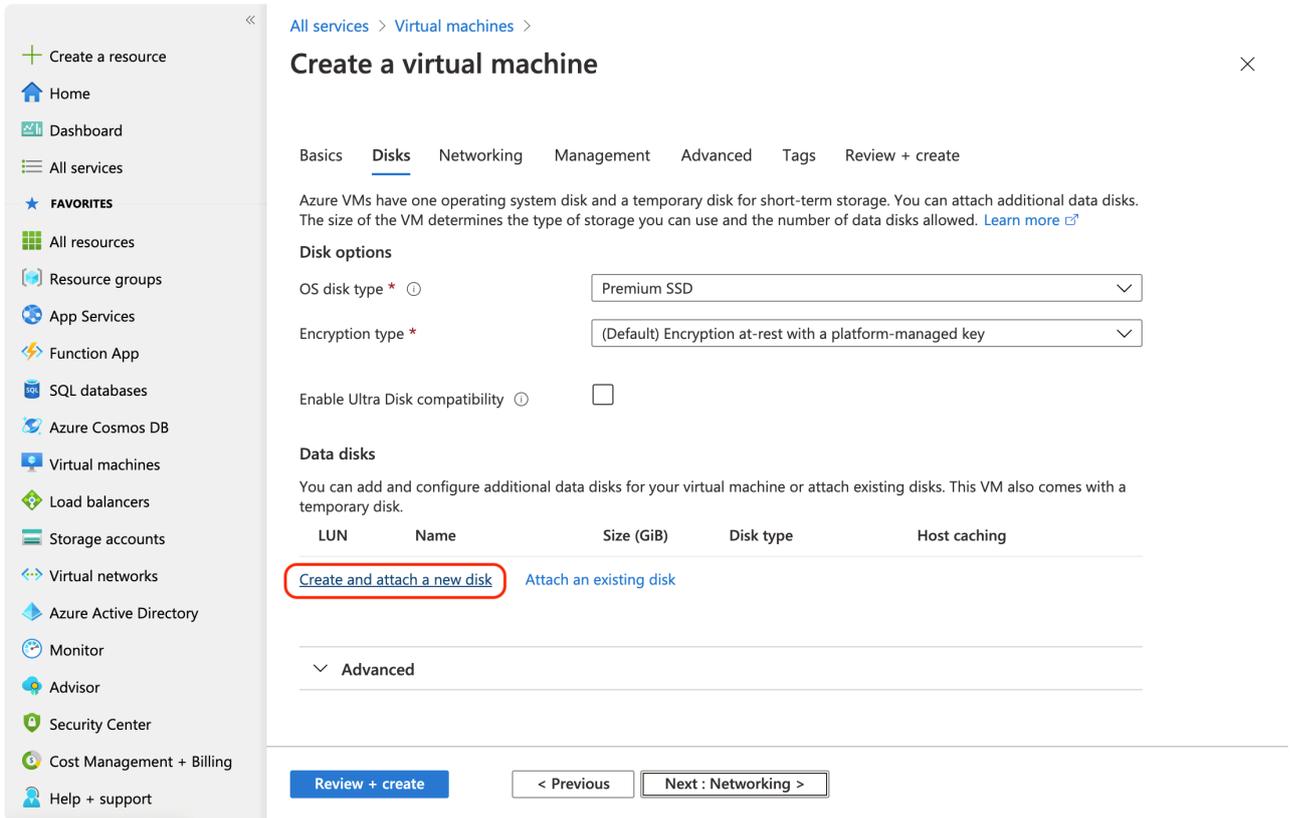
The screenshot shows the 'Create a virtual machine' wizard in the Azure portal. The 'Basics' tab is selected, and the following configuration is shown:

- Project details:** Subscription: Azure Pay As You Go Subscription; Resource group: LK-QSG.
- Instance details:** Virtual machine name: node-a; Region: (US) West US 2; Availability options: Availability zone; Availability zone: 1.
- Image:** Red Hat Enterprise Linux 8.2 (LVM) - Gen1.
- Size:** Standard\_B1s - 1 vcpu, 1 GiB memory (\$7.59/month).
- Administrator account:** Authentication type: SSH public key; Username: azureuser; SSH public key source: Generate new key pair; Key pair name: Azure-LK-QuickStart.
- Inbound port rules:** Public inbound ports: None.

At the bottom, the 'Next: Disks >' button is highlighted, indicating the next step in the wizard.

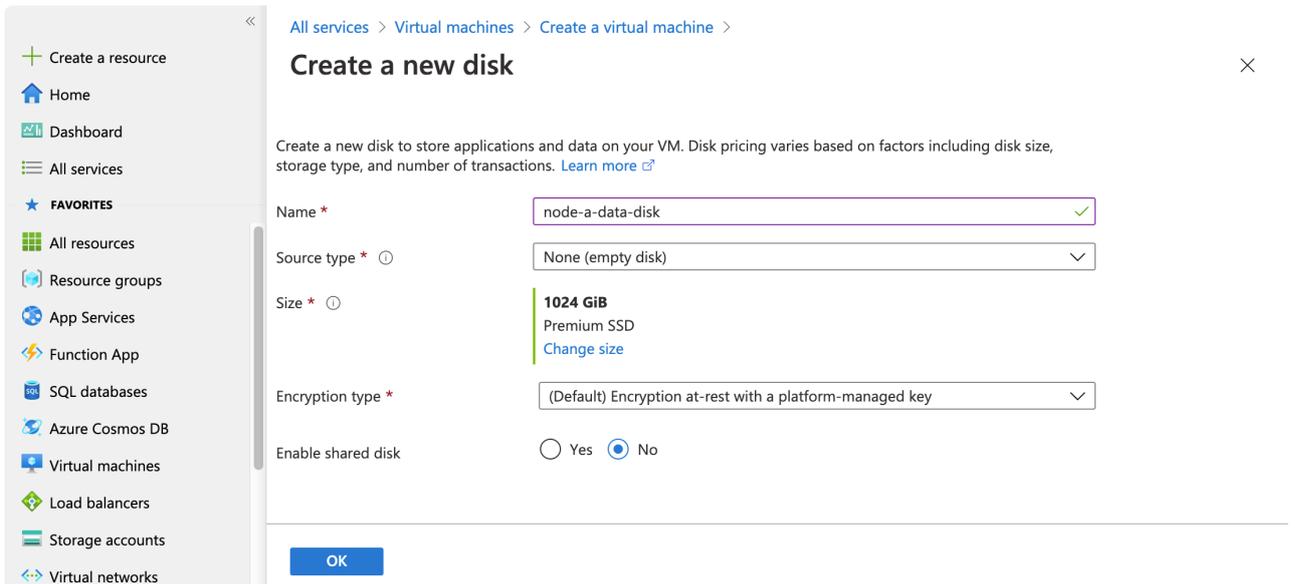
Once you enter or select these values, click “Next: Disks”.

- By default, the wizard creates one disk to boot the operating system. You need an additional disk for data replication.



Click "Create and attach a new disk".

4. Select the disk size. To evaluate LifeKeeper, please create at least an 8 GiB disk.



Click "Change size" to specify the disk size.

5. Select the disk size (select at least 8 GiB if you are evaluating LifeKeeper), then select "OK".

**Select a disk size**

Browse available disk sizes and their features.

Disk SKU

Size	Disk tier	Provisioned IOPS	Provisioned thro...	Max Shares	Max burst IOPS
4 GiB	P1	120	25	-	3500
8 GiB	P2	120	25	-	3500
16 GiB	P3	120	25	-	3500
32 GiB	P4	120	25	-	3500
64 GiB	P6	240	50	-	3500
128 GiB	P10	500	100	-	3500
256 GiB	P15	1100	125	2	3500

**OK**

6. Now the disk size is specified. Click “OK” to confirm.

**Create a new disk**

Create a new disk to store applications and data on your VM. Disk pricing varies based on factors including disk size, storage type, and number of transactions. [Learn more](#)

Name \*

Source type \*

Size \* **8 GiB**  
Premium SSD  
[Change size](#)

Encryption type \*

Enable shared disk  Yes  No  
Shared disk not available for the selected size.

**OK**

7. Once the disk configuration is finished, click “Next: Networking”.

[All services](#) > [Virtual machines](#) > **Create a virtual machine** ✕

[Basics](#) **[Disks](#)** [Networking](#) [Management](#) [Advanced](#) [Tags](#) [Review + create](#)

Azure VMs have one operating system disk and a temporary disk for short-term storage. You can attach additional data disks. The size of the VM determines the type of storage you can use and the number of data disks allowed. [Learn more](#)

**Disk options**

OS disk type \*

Encryption type \*

Enable Ultra Disk compatibility

**Data disks**

You can add and configure additional data disks for your virtual machine or attach existing disks. This VM also comes with a temporary disk.

LUN	Name	Size (GiB)	Disk type	Host caching
<input type="text" value="0"/>	node-a-data-disk	8	Premium SSD	<input type="text" value="None"/>

[Create and attach a new disk](#) [Attach an existing disk](#)

---

▼ **Advanced**

Review + create
< Previous
Next : Networking >

8. Select the following parameters for the network:

- Virtual network: LK-VNET
- Subnet: LK-subnet
- Public IP: (new) node-a-ip
- NIC network security group: Advanced
- Configure network security group: LK-NSG

The screenshot shows the 'Create a virtual machine' wizard in the 'Networking' step. The left sidebar contains navigation options like 'Home', 'Dashboard', 'All services', and 'FAVORITES'. The main content area has tabs for 'Basics', 'Disks', 'Networking' (selected), 'Management', 'Advanced', 'Tags', and 'Review + create'. Below the tabs, there is a description of network connectivity and a 'Network interface' section. This section includes dropdown menus for 'Virtual network' (LK-VNET), 'Subnet' (LK-subnet (10.20.0.0/22)), and 'Public IP' ((new) node-a-ip). There are also radio buttons for 'NIC network security group' (None, Basic, Advanced) and a dropdown for 'Configure network security group' (LK-NSG). A note states 'Accelerated networking' is not supported for the selected VM size. At the bottom, there are buttons for 'Review + create', '< Previous', and 'Next : Management >'.

Once you specify the values, click “Review + create” to complete.

9. Once the wizard validates your input, click “Create”.

The screenshot shows the 'Create a virtual machine' wizard in the 'Review + create' step. A green banner at the top indicates 'Validation passed'. The 'Review + create' tab is selected in the top navigation. The 'PRODUCT DETAILS' section shows 'Standard B1s by Microsoft' with a price of '0.0104 USD/hr'. Below this, the 'TERMS' section contains a legal disclaimer. At the bottom, the 'Basics' section shows 'Subscription' and 'Azure Pay As You Go Subscription'.

Resource group	LK-QSG
Virtual machine name	node-a
Region	West US 2
Availability options	Availability zone
Availability zone	1
Image	Red Hat Enterprise Linux 8.2 (LVM) - Gen1
Size	Standard B1s (1 vcpu, 1 GiB memory)
Authentication type	SSH public key
Username	azureuser
Key pair name	Azure-LK-QuickStart
Already have a Red Hat Enterprise Linux subscription?	No
Azure Spot	No
<b>Disks</b>	
OS disk type	Premium SSD
Use managed disks	Yes
Data disks	1
Use ephemeral OS disk	No
<b>Networking</b>	
Virtual network	LK-VNET
Subnet	LK-subnet (10.20.0.0/22)
Public IP	(new) node-a-ip
NIC network security group	LK-NSG
Accelerated networking	Off
Place this virtual machine behind an existing load balancing solution?	No

[Create](#)   [< Previous](#)   [Next >](#)   [Download a template for automation](#)

10. If this is the first time you are using the key pair, download the private key file. You will need the key file to connect to the instance later.

## Generate new key pair

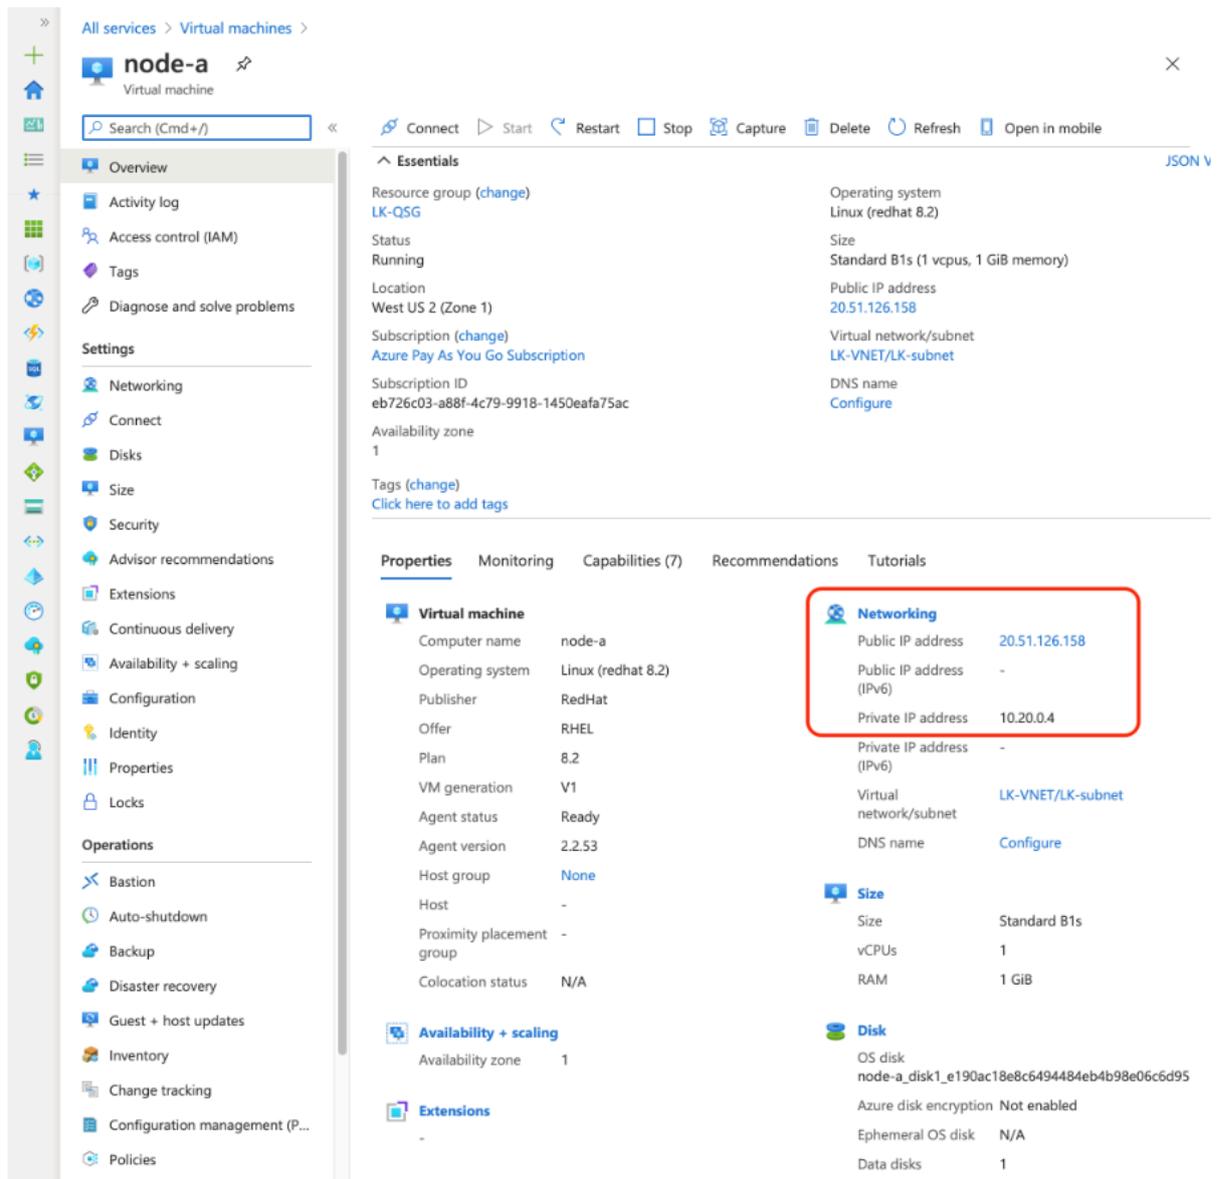
**i** An SSH key pair contains both a public key and a private key. **Azure doesn't store the private key.** After the SSH key resource is created, you won't be able to download the private key again. [Learn more](#)

[Download private key and create resource](#)

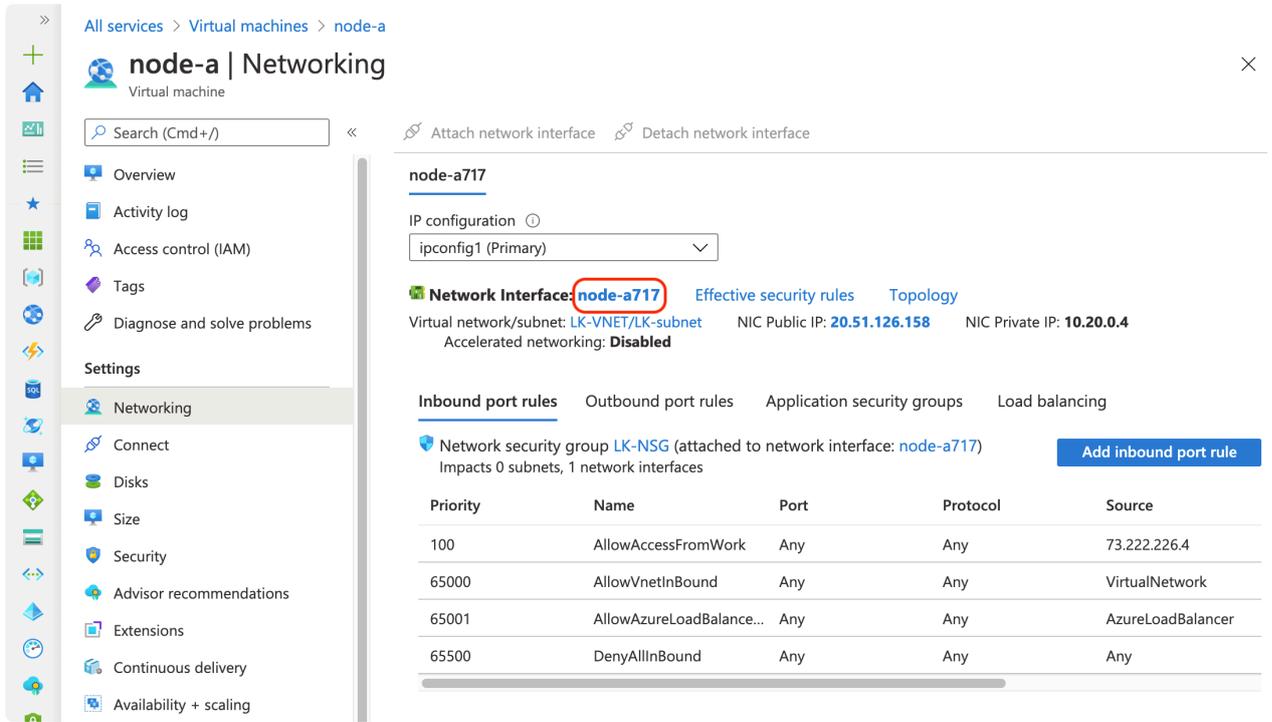
[Return to create a virtual machine](#)

11. Once the virtual machine is created, you can see the configuration of the VM. You may notice that

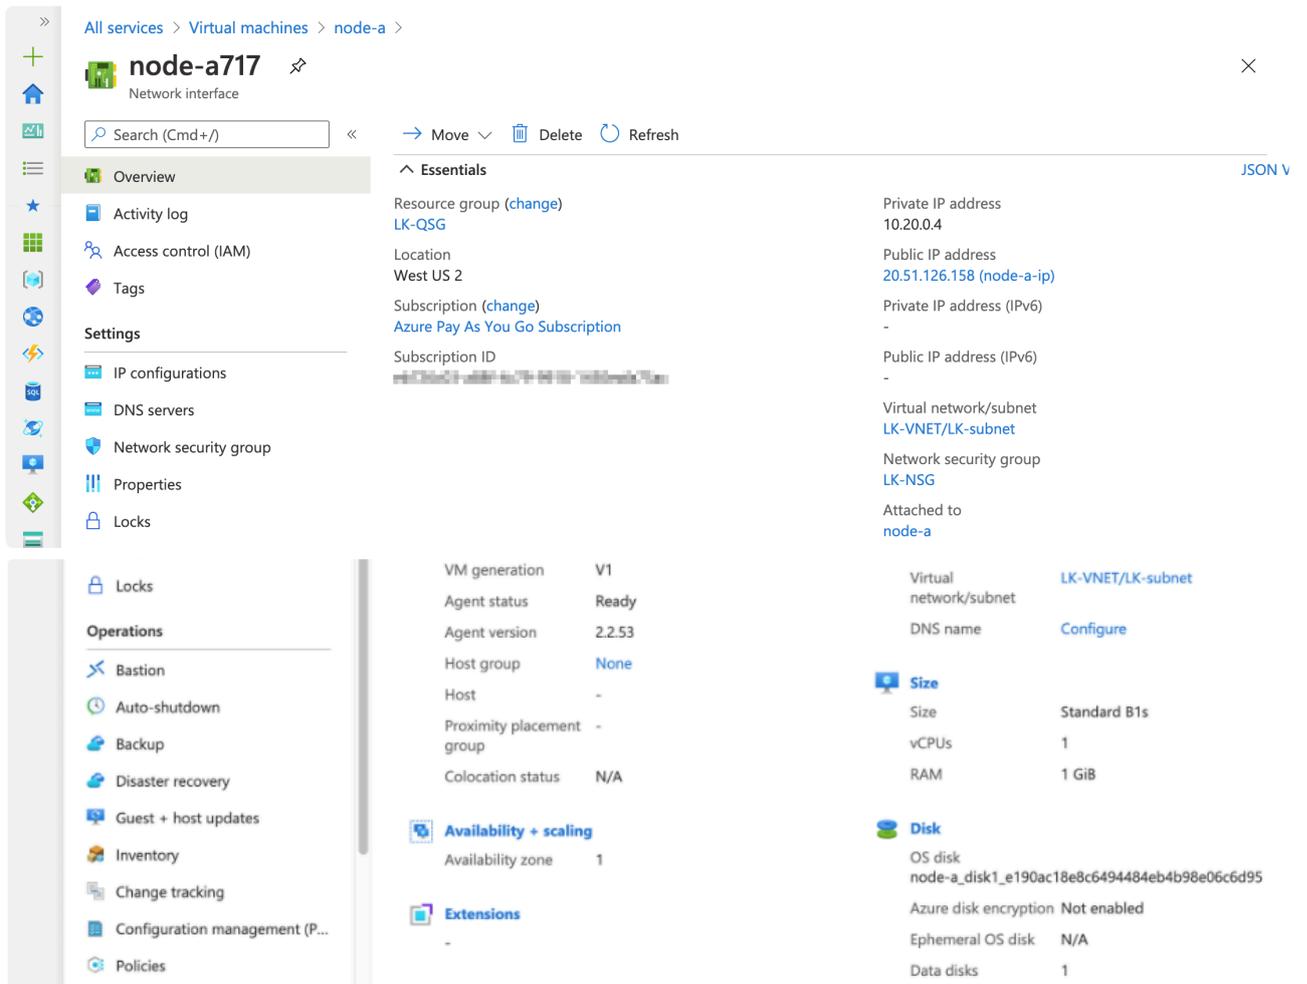
we haven't specified the IP address yet and the DHCP assigned address 10.20.0.4 was chosen automatically. Click "Networking" (the blue link) to change the network configuration.



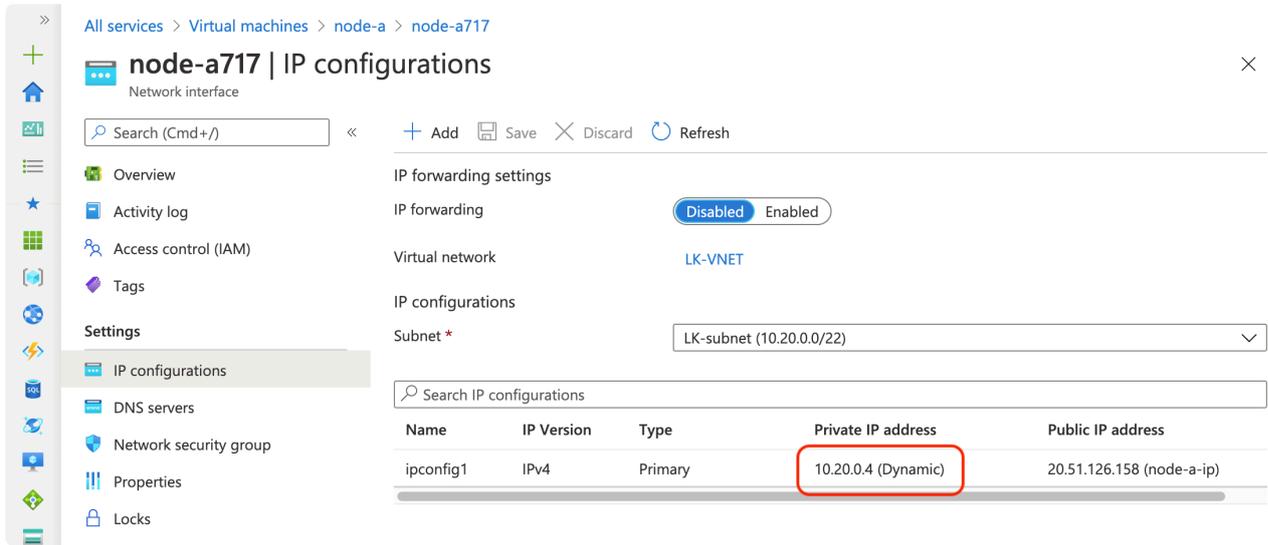
- You can see in the screenshot below that the virtual machine has a network interface (node-a717). Click `node-a717` to change configuration of this network card.



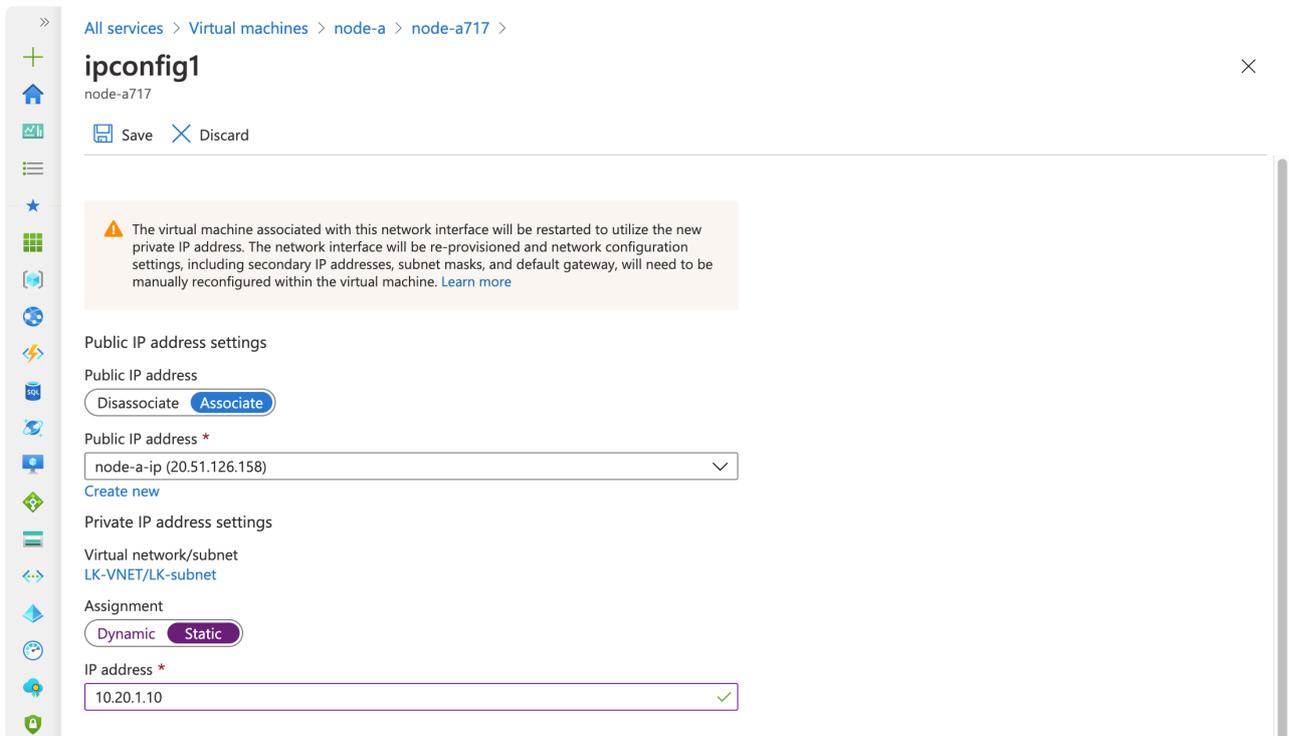
13. Now you can see the details of the network card. Click “IP configurations”.



14. Select the IP address to go to configuration page.



15. Change Assignment from Dynamic to Static and set the IP address to 10.20.1.10.



Once you specify the IP address, click “Save”.

16. The configuration of the virtual machine is done.

## 11.2.3.4.7. Creating the Second and Third Virtual Machines

The steps required to create the second node (node-b) and third node (node-c) is almost the same as previously described for the first node (node-a). The following table illustrates the differences between the three nodes and the details required to create these instances.

Name	Parameter	Value
 Common values across VM	Resource Group	LK-QSG
	Virtual network	LK-VNET
	Subnet	LK-subnet
	Network security Group	LK-NSG
node-a	Availability Zone (*1)	1
	Private IP Address (*2)	10.20.1.10
	Second Disk (Storage Device)	You need to have a second disk
node-b	Availability Zone (*1)	2
	Private IP Address (*2)	10.20.2.10
	Second Disk (Storage Device)	You need to have a second disk
node-c	Availability Zone (*1)	3
	Private IP Address (*2)	10.20.3.10 (node-c)
	Second Disk (Storage Device)	You do not need to create a second disk

Once created, the list of instances should look like the following.



All services >

## Virtual machines ↗



Default Directory

+ Add ▾ ⌚ Reservations ▾ ≡ Edit columns ↻ Refresh ↺ Try preview | 🏷️ Assign tags ▶ Start ↺ Restart ☐ Stop 🗑️ Delete ⋮

**Subscriptions:** Azure Pay As You Go Subscription

Filter by name... All resource groups ▾ All types ▾ All locations ▾ All tags ▾ No grouping ▾

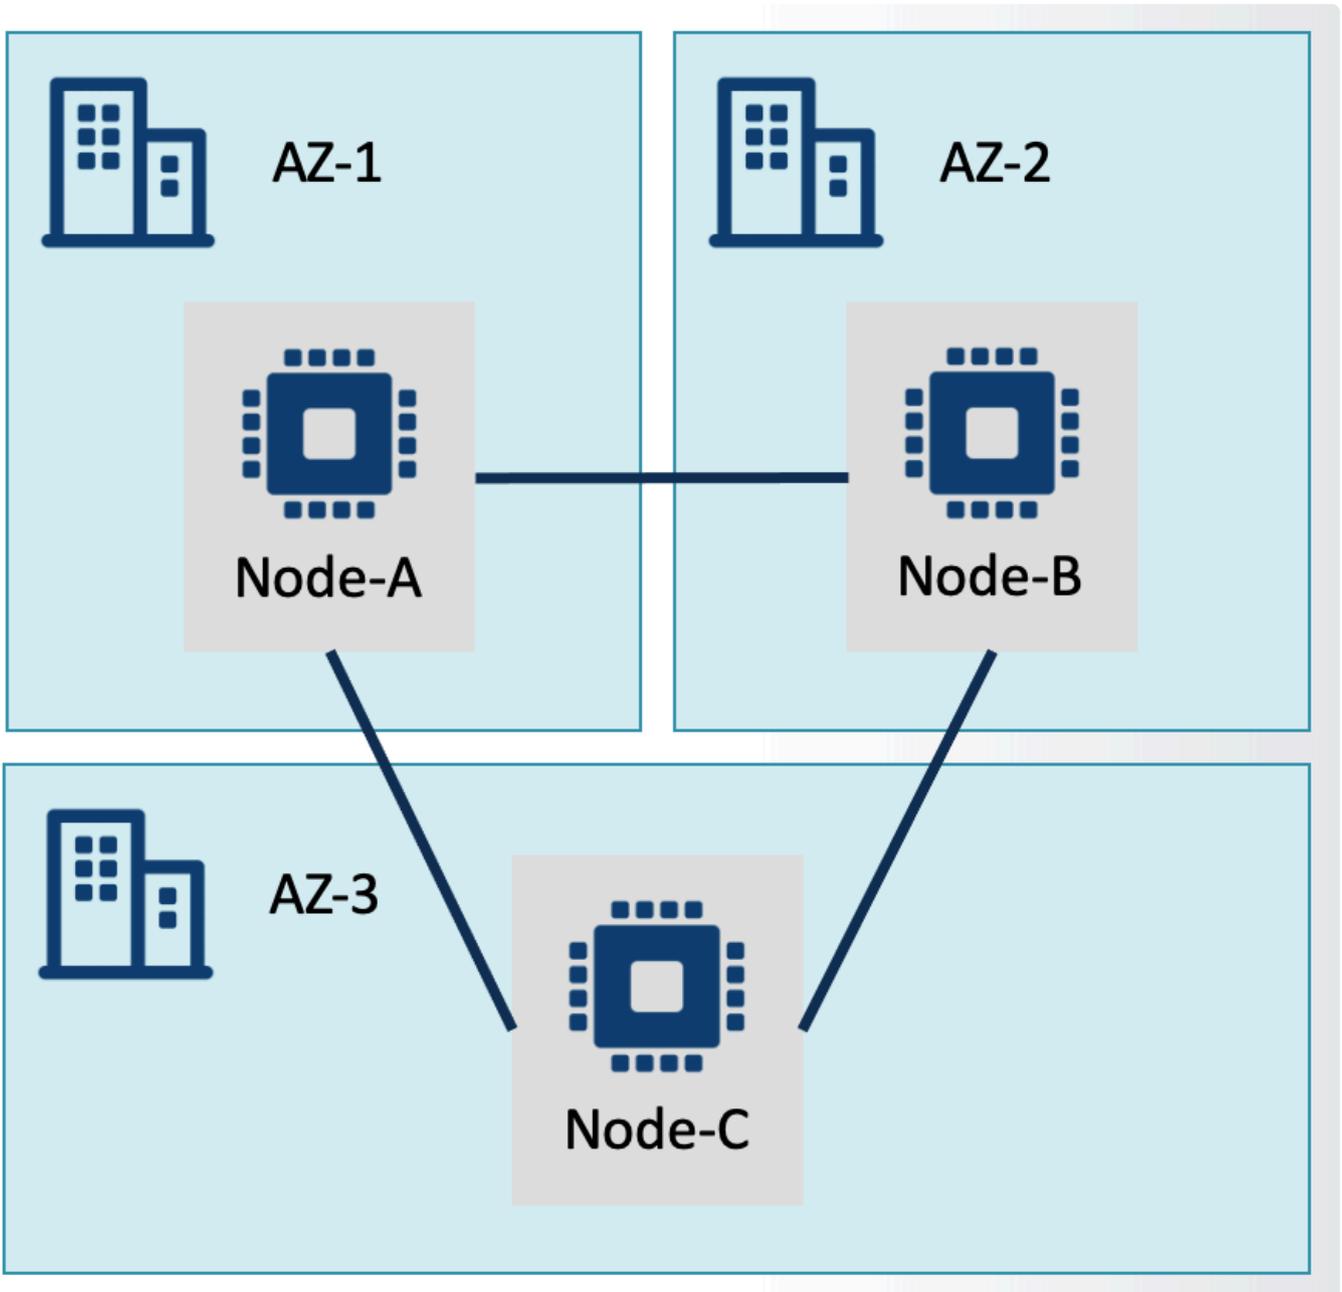
3 items

<input type="checkbox"/>	Name ↑↓	Status	Resource group ↑↓	Availability zone	Subnet	Private IP address	
<input type="checkbox"/>	node-a	Running	LK-QSG	1	LK-subnet	10.20.1.10	⋮
<input type="checkbox"/>	node-b	Running	LK-QSG	2	LK-subnet	10.20.2.10	⋮
<input type="checkbox"/>	node-c	Running	LK-QSG	3	LK-subnet	10.20.3.10	⋮

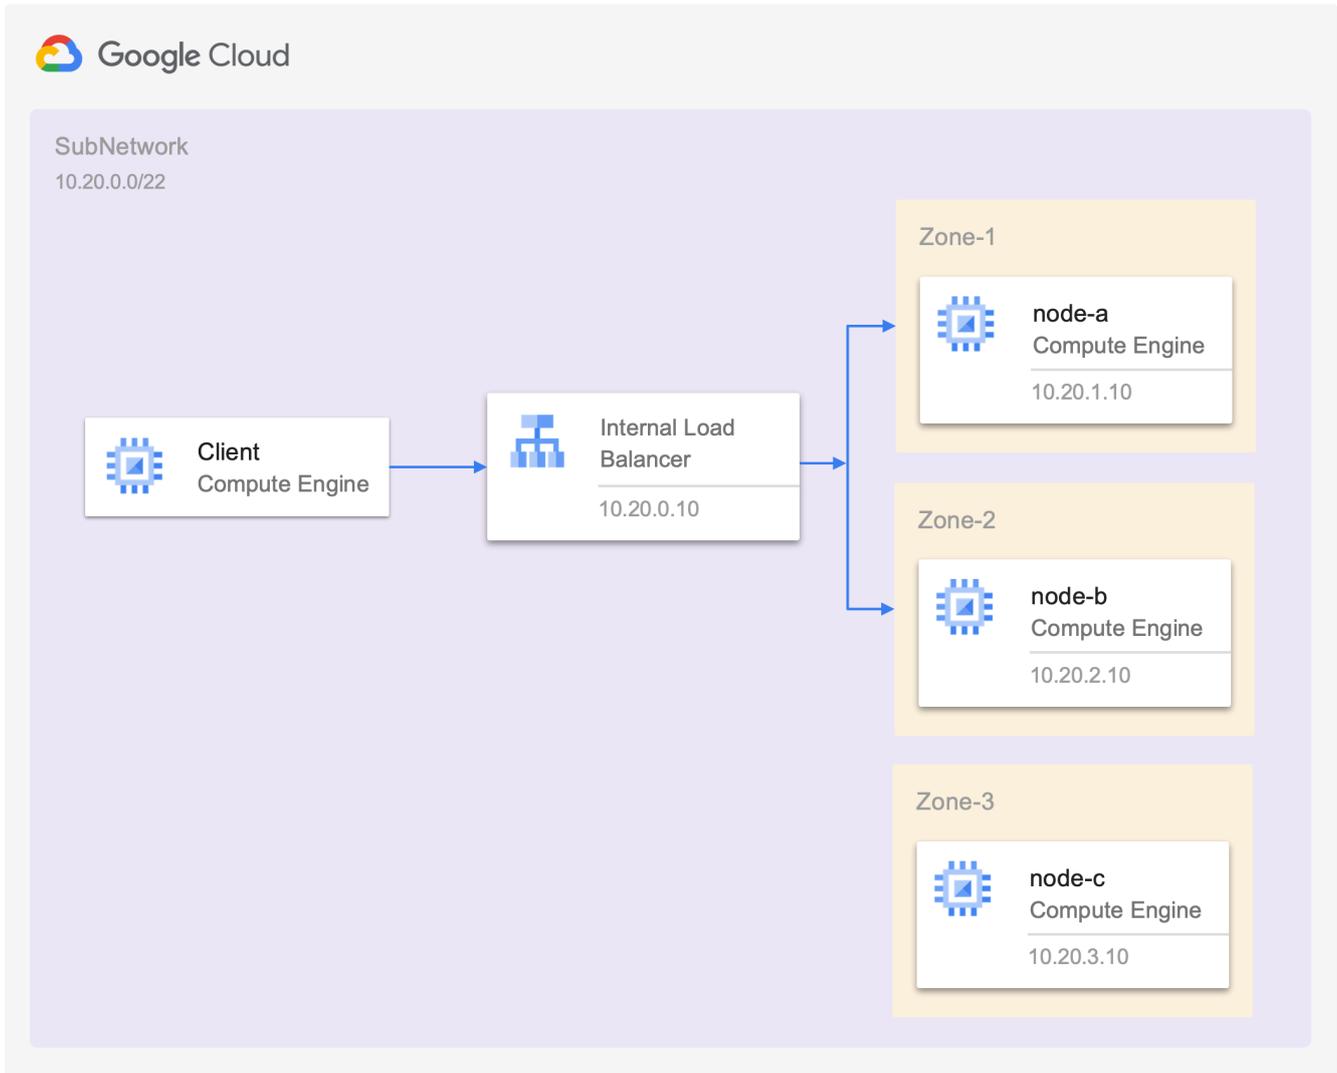
## 11.2.3.5. Creating an Instance in Google Cloud from Scratch

**!** **Disclaimer:** The user interface may vary between regions or change over time. Please refer to the documentation provided by Google Cloud if the screenshots shown below are different from your experience.

This Evaluation Guide uses the following network structure and instances.



On Google Cloud, these components can be defined as per the diagram below.



In this section we will create these components, with the exception of the Internal Load Balancer. We will create the Internal Load Balancer [later](#).

The components shown in this diagram are described in the following table.

To create this network structure, the following components must be implemented:

Component	Name	Parameter	Value
VPC network	lk-vpc	Region	us-west1 (*1)
		Subnet (lk-subnet)	10.20.0.0/22
Virtual Machine	 Common values across VM	VPC network	lk-vpc
		Subnet	lk-subnet
		Network Tags	lk-node
		External IP	Ephemeral
	node-a	Zone	us-west1-a (*1)

		Hostname	node-a.internal	
		Internal IP	10.20.1.10	
	node-b	Zone	us-west1-b (*1)	
		Hostname	node-b.internal	
	node-c	Internal IP	10.20.2.10	
		Zone	us-west1-c (*1)	
		Hostname	node-b.internal	
	Firewall Rules	 Common values across Rules	<b>Network</b>	<b>lk-vpc</b>
			<b>Type</b>	<b>Ingress</b>
<b>Target</b>			<b>lk-node</b>	
<b>Action</b>			<b>Allow</b>	
fw-allow-ssh		Source	WAN IP Address of work location	
		Protocols / ports	all	
fw-allow-lk-node-connection		Source	Tags: lk-node	
		Protocols / ports	all	
fw-allow-health-check (*2)		Source	130.211.0.0/22, 35.191.0.0/16	
		Protocols / ports	all	

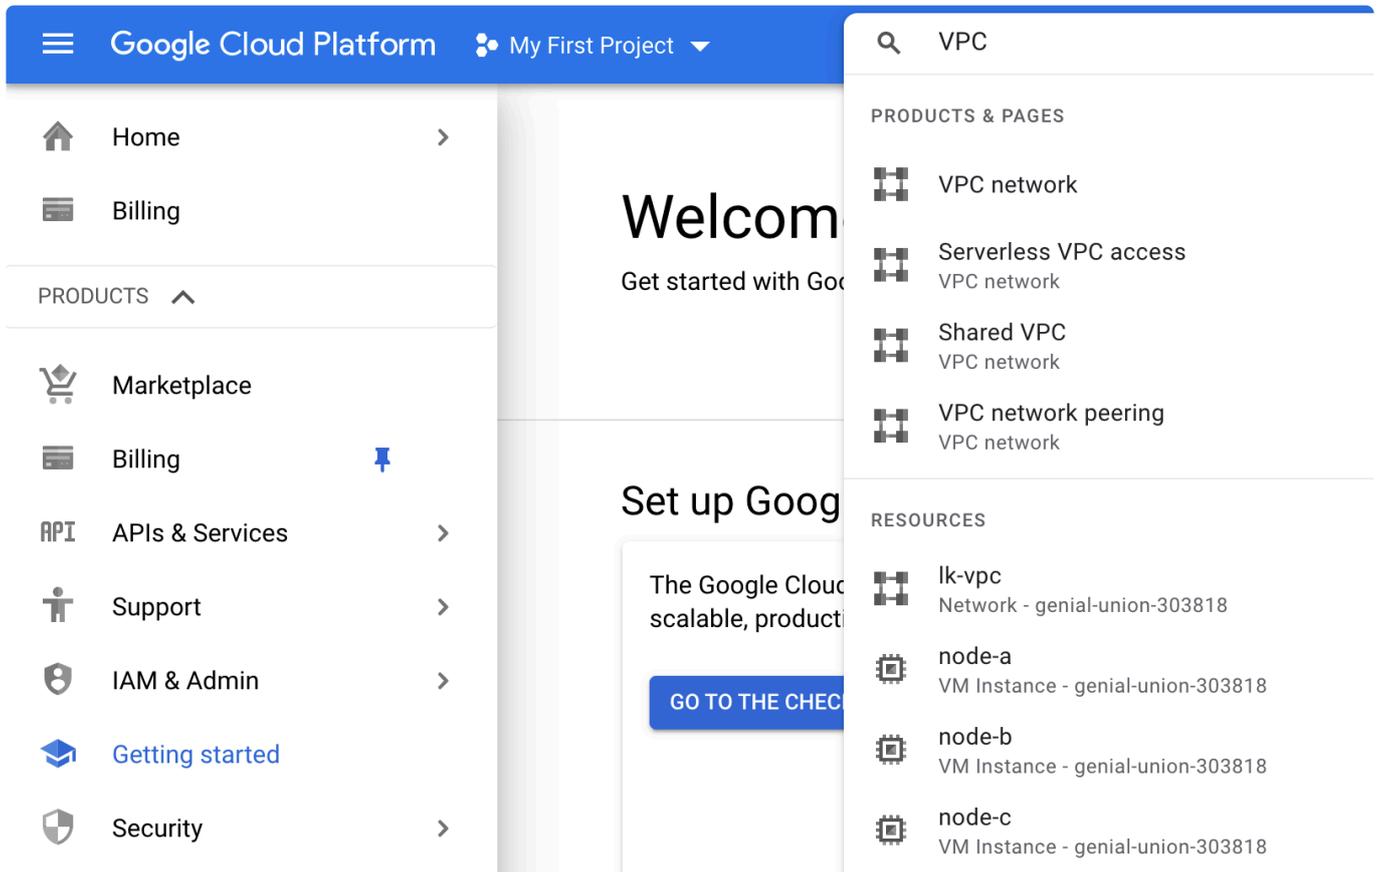
- (\*1) This example uses us-west1 as Region.
- (\*2) This is required to allow the Internal Load Balancer to check the status of each nodes.

The following sections step through the processes necessary for creating these network components, beginning with the first instance.

- [Switching between Google Cloud Services](#)
- [Deciding on a Google Cloud Region](#)
- [Creating the Project](#)
- [Creating a VPC Network](#)
- [Creating a New SSH Key](#)
- [Creating the First Google Cloud VM](#)
- [Configuring the Firewall Rules](#)
- [Creating the Second and Third VM](#)

# 11.2.3.5.1. Switching between Google Cloud Services

Google Cloud has a significant and growing number of services available. To move to a different service you may want to type the name of the service in the search box at the top of the Google Cloud console. Select the service from the list. The screenshot below shows how to select the VPC Network service.



## 11.2.3.5.2. Deciding on a Google Cloud Region

Google Cloud has regions in many geographic locations and it may be beneficial to select a region in close geographic proximity to the workplace location. Each region has three or more Zones, but available machine types are slightly different between regions.

Refer to the [Google Cloud documentation](#) for more information.

When you create a VM, select the Region and Zone. These cannot be changed once you create a VM.

☰ Google Cloud Platform My First Project ▾

← Create an instance

To create a VM instance, select one of the options:

- New VM instance**  
Create a single VM instance from scratch
- New VM instance from template**  
Create a single VM instance from an existing template

**Name** ⓘ  
Name is permanent

**Labels** ⓘ (Optional)

**Region** ⓘ  
Region is permanent

**Zone** ⓘ  
Zone is permanent

## 11.2.3.5.3. Creating the Project

A Project is the logical container in which you can create resources such as network components, computer resources, storage, etc.

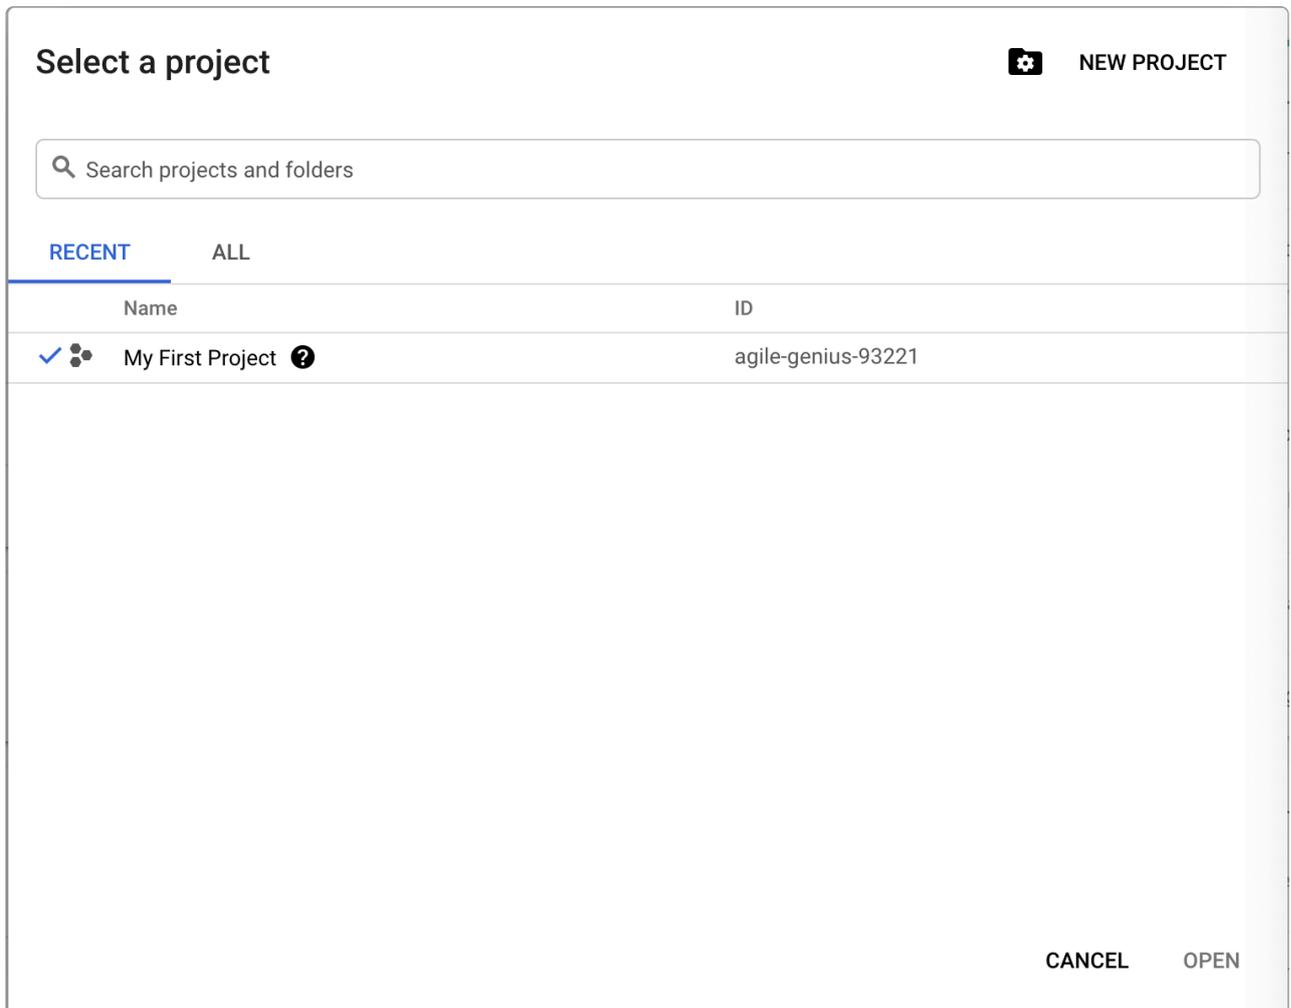
Refer to the [Google Cloud documentation](#) for more information.

If you have just signed up to Google Cloud, you have a default project called “My First Project”. You can use this project. However, you can also create a new project (we will name it LK Quick Start Project) as follows:

1. Select “My First Project” (or any other active project you are working on) located at the top.



2. On the “Select a project” screen, click “NEW PROJECT” located at the top right side.



3. Specify the name of the Project as `LK Quick Start Project`.

**Project name \***  
LK Quick Start Project ?

Project ID: lk-quick-start-project. It cannot be changed later. [EDIT](#)

**Organization \***  
[Redacted] ▼ ?

Select an organization to attach it to a project. This selection can't be changed later.

**Location \***  
[Redacted] [BROWSE](#)

Parent organization or folder

**CREATE** **CANCEL**

Once you fill in the fields, click “Create”.

4. Now the new Project is created.

Google Cloud Platform LK Quick Start Project Search products and resources

DASHBOARD ACTIVITY RECOMMENDATIONS

How Google Cloud is helping during COVID-19. [Learn more](#)

### Project info

Project name  
LK Quick Start Project

Project ID  
lk-quick-start-project

Project number  
921312571473

[ADD PEOPLE TO THIS PROJECT](#)

[Go to project settings](#)

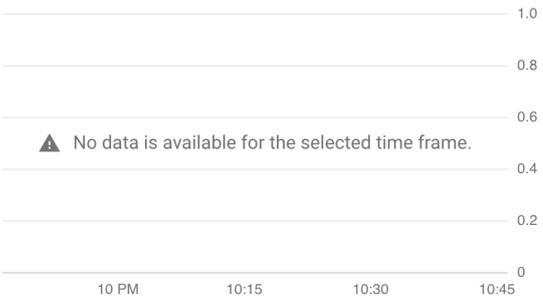
### Resources

This project has no resources

[CREATE RESOURCE](#)

### API APIs

Requests (requests/sec)



[Go to APIs overview](#)

## 11.2.3.5.4. Creating a VPC Network

A VPC network is a Google Cloud resource that represents a local network. Different VPC networks can be defined within the Google Cloud to logically separate different systems.

Refer to the [Google Cloud documentation](#) for more information.

In this section we will create a VPC network for testing LifeKeeper (we will name it `lk-vpc`) as follows:

1. Select “VPC network” to see the current list of VPC networks. There is a “default” VPC network and its subnets for each region defined. Click “CREATE VPC NETWORK” located at the top.

Name ↑	Region	Subnets	MTU ?	Mode
▼ default		24	1460	Auto ▼
	us-central1	default		
	europa-west1	default		
	us-west1	default		
	asia-east1	default		

2. In the “Create a VPC network” wizard, set the name of the VPC network as `lk-vpc`, the region to create the VPC network in, the name of the subnet, and the CIDR block `10.20.0.0/22`.

- Once you review the values you entered, click “CREATE” at the bottom. Now the new VPC network is created.

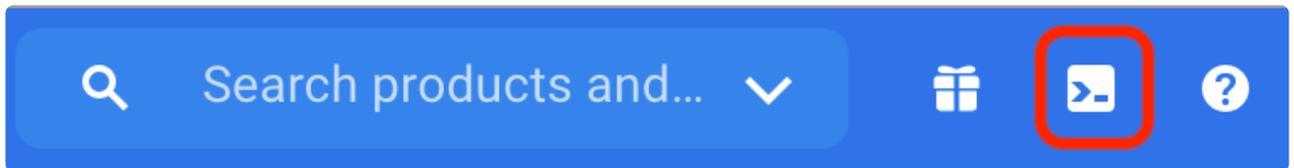
Name	Region	Subnets	MTU	Mode	IP address ranges	Gateways	Firewall Rules
default		24	1460	Auto			4
lk-vpc	us-west1	1	1460	Custom	10.20.0.0/22	10.20.0.1	0

## 11.2.3.5.5. Creating a New SSH Key

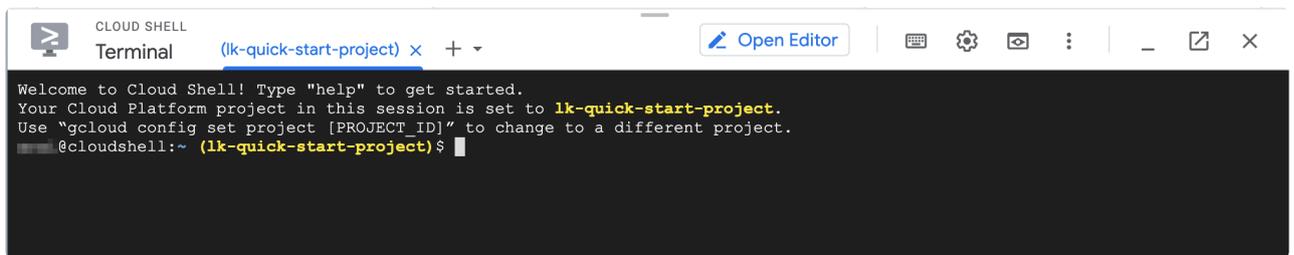
You will need to use an ssh key to connect to VMs. This section outlines how you can create a new ssh key on the Google Cloud console and then download the private key to your local system.

Refer to the [Google Cloud documentation](#) for more information.

1. Activate the Cloud Shell by clicking [**>\_**] at the top of the console.



2. Once you activate the Cloud Shell you will see the shell screen at the bottom of the browser.



3. Create a new ssh key by using the `ssh-keygen` tool.

```
1 ssh-keygen -t rsa -C "gcp-user" -f gcp-lk-quickstart
```

4. Now a new key `gcp-lk-quickstart` is generated with a user name of `gcp-user`.

```

CLOUD SHELL
Terminal (lk-quick-start-project) x +
Open Editor

Welcome to Cloud Shell! Type "help" to get started.
Your Cloud Platform project in this session is set to lk-quick-start-project.
Use "gcloud config set project [PROJECT_ID]" to change to a different project.
@cloudshell:~ (lk-quick-start-project)$ ssh-keygen -t rsa -C "gcp-user" -f gcp-lk-quickstart
Generating public/private rsa key pair.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in gcp-lk-quickstart.
Your public key has been saved in gcp-lk-quickstart.pub.
The key fingerprint is:
SHA256:/1X4HS8jaM+... gcp-user
The key's randomart image is:
+---[RSA 2048]-----+
|
|             ..
|      .S . .+.
|     + .. + ooE|
|    + + .+.o.+*+|
|   . o @o*B.+ =o+|
|  =.%*=B+.oo|
+-----[SHA256]-----+
@cloudshell:~ (lk-quick-start-project)$
    
```

- In a later step you will need to register a public key `gcp-lk-quickstart.pub` to a new VM. Type `cat gcp-lk-quickstart.pub` to copy the public key.

```

@cloudshell:~ (lk-quick-start-project)$ cat gcp-lk-quickstart.pub
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDA... gcp-user
@cloudshell:~ (lk-quick-start-project)$
    
```

- Download the private key with `cloudshell dl gcp-lk-quickstart`.

```

CLOUD SHELL
Terminal (lk-quick-start-project) x +
Open Editor

@cloudshell:~ (lk-quick-start-project)$ cloudshell dl gcp-lk-quickstart
@cloudshell:~ (lk-quick-start-project)$
    
```

**Download File**

Please confirm that you wish to download the following files:

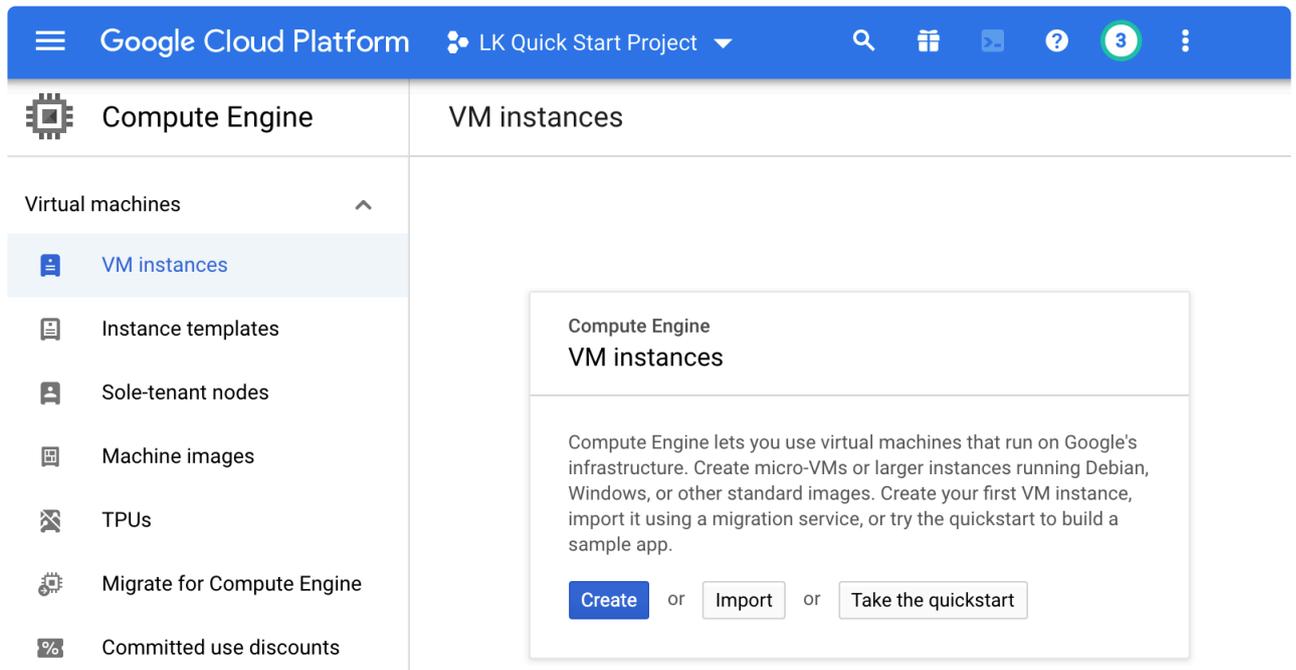
/home/.../gcp-lk-quickstart

You will need this private key to connect to the VM in a later step.

## 11.2.3.5.6. Creating the First Google Cloud VM

In previous sections we covered the configuration of the network. Now we are going to create the first VM. As discussed in [Computing Resources Used in this Tutorial](#), we need two disks. This section also discusses how to create the second disk.

1. Go to “Compute Engine” > “VM Instance” and click “Create” to create a new VM.



2. The “Create an instance” wizard appears.

- 3. Enter the name of the new VM: node-a (when you create 2nd or 3rd nodes, this would be node-b and node-c).

**Name** ?  
Name is permanent

- 4. Select the Region and Zone of the VM. Enter a nearby region, then select the Zone. Select the zone ending with -a for node-a, -b for node-b and -c for node-c. Note that the zones in some regions may be named differently.

**Region** ?  
Region is permanent

**Zone** ?  
Zone is permanent

us-west1 (Oregon) ▼      us-west1-a ▼

- 5. Select the machine configuration. Select the required CPU and Memory requirements based on the application you are protecting.

## Machine configuration

**Machine family**

General-purpose
Compute-optimized
Memory-optimized

Machine types for common workloads, optimized for cost and flexibility

**Series**

E2
▼

CPU platform selection based on availability

**Machine type**

e2-micro (2 vCPU, 1 GB memory)
▼



vCPU	Memory	GPUs
1 shared core	1 GB	-

6. Select an operating system. By default, Debian GNU/Linux 10 is selected. Select an operating system that is supported by both LifeKeeper as well as the database and/or the application you are planning to use. Click “Change” and select Red Hat Enterprise Linux 7.

**!** If planning to create an SAP or SAP HANA cluster, select the “Red Hat Enterprise Linux 8.2 for SAP Applications” boot image in this step.

## Boot disk

Select an image or snapshot to create a boot disk; or attach an existing disk. Can't find what you're looking for? Explore hundreds of VM solutions in [Marketplace](#).

[Public images](#) [Custom images](#) [Snapshots](#) [Existing disks](#)

### Operating system

Red Hat Enterprise Linux

### Version

Red Hat Enterprise Linux 7

x86\_64 built on 20210122, supports Shielded VM features [?](#)

### Boot disk type [?](#)

Standard persistent disk

### Size (GB) [?](#)

20

7. Before selecting "Create", expand the drop down list for "Management, security, disks, networking, sole tenancy".

### Firewall [?](#)

Add tags and firewall rules to allow specific network traffic from the Internet

- Allow HTTP traffic
- Allow HTTPS traffic

Management, security, disks, networking, sole tenancy

Your free trial credit will be used for this VM instance. [GCP Free Tier](#) [↗](#)

8. Go to the Security Tab and enter the public key that was downloaded in the [Creating a New SSH Key](#) section.

 Make sure there is no whitespace if you copy and paste the public key.



Management Security **Disks** Networking Sole Tenancy

**Boot disk**

**Deletion rule**

Delete boot disk when instance is deleted

**Encryption**

Data is encrypted automatically. Select an encryption key management solution.

- Google-managed key**  
No configuration required
- Customer-managed key**  
Manage via Google Cloud Key Management Service
- Customer-supplied key**  
Manage outside of Google Cloud

**Device name** ?

Used to reference the device for mounting or resizing.

Based on instance name (default) ▼

node-a

**Additional disks** ? (Optional)

+ Add new disk + Attach existing disk

⤴ Less

10. Specify the "Name" and the "Size" of the disk.

**Additional disks** ? (Optional)

New disk (node-a-datadisk, Blank, 10 GB) 

11. Go to the "Networking" tab. Specify "Network tags" and "Hostname". The "Network tags" are used to identify the VM in the Firewall configuration (discussed in the next step).

Management Security Disks **Networking** Sole Tenancy

**Network tags** ? (Optional)

lk-node 

**Hostname** ?

Set a custom hostname for this instance or leave it default. Choice is permanent

node-a.internal

12. Edit the network interface. Select the “Edit” icon.

**Network interfaces** ?

Network interface is permanent

default default (10.138.0.0/20)



[+ Add network interface](#)

13. Enter the parameters of the network interface as follows:

- Network: lk-vpc
- Subnetwork: lk-subnet (10.20.0.0/22)
- Primary internal IP: Ephemeral (Custom)
- Custom ephemeral IP address: 10.20.1.10 for node-a, 10.20.2.10 for node-b, 10.20.3.10 for node-c.
- External IP: Ephemeral
- IP forwarding: Off

### Network interface

**Network** ?

lk-vpc

**Subnetwork** ?

lk-subnet (10.20.0.0/22)

**Primary internal IP** ?

Ephemeral (Custom)

**Custom ephemeral IP address**

10.20.1.10

⌵ Show alias IP ranges

**External IP** ?

Ephemeral

**Network Service Tier** ?

Premium (Current project-level tier, [change](#)) ?

Standard (us-west1) ?

**IP forwarding** ?

Off

**Public DNS PTR Record** ?

Enable

PTR domain name

Done Cancel

Click "Done" once you enter the values.

Confirm the values you entered, then click "Create" at the bottom.

14. Congratulations! You have completed the configuration of the virtual machine.

The screenshot shows the Google Cloud Platform interface for VM instances. The top navigation bar includes the Google Cloud Platform logo, the project name 'LK Quick Start Project', a search bar, and notification icons. The left sidebar shows the 'Compute Engine' section with 'Virtual machines' expanded, listing 'VM instances', 'Instance templates', 'Sole-tenant nodes', and 'Machine images'. The main content area is titled 'VM instances' and features a toolbar with icons for adding, downloading, refreshing, playing, pausing, and deleting instances. Below the toolbar is a table of VM instances. A search filter 'Filter VM instances' and a 'Columns' dropdown are at the top of the table. The table has columns for Name, Zone, Recommendation, In use by, Internal IP, External IP, and Connect. One instance, 'node-a', is listed in the 'us-west1-a' zone with an internal IP of 10.20.1.10 (nic0) and an external IP of 35.233.141.70. The 'Connect' column shows 'SSH' and a menu icon.

<input type="checkbox"/>	Name ^	Zone	Recommendation	In use by	Internal IP	External IP	Connect
<input type="checkbox"/>	node-a	us-west1-a			10.20.1.10 (nic0)	35.233.141.70	SSH ▾ ⋮

## 11.2.3.5.7. Configuring the Firewall Rules

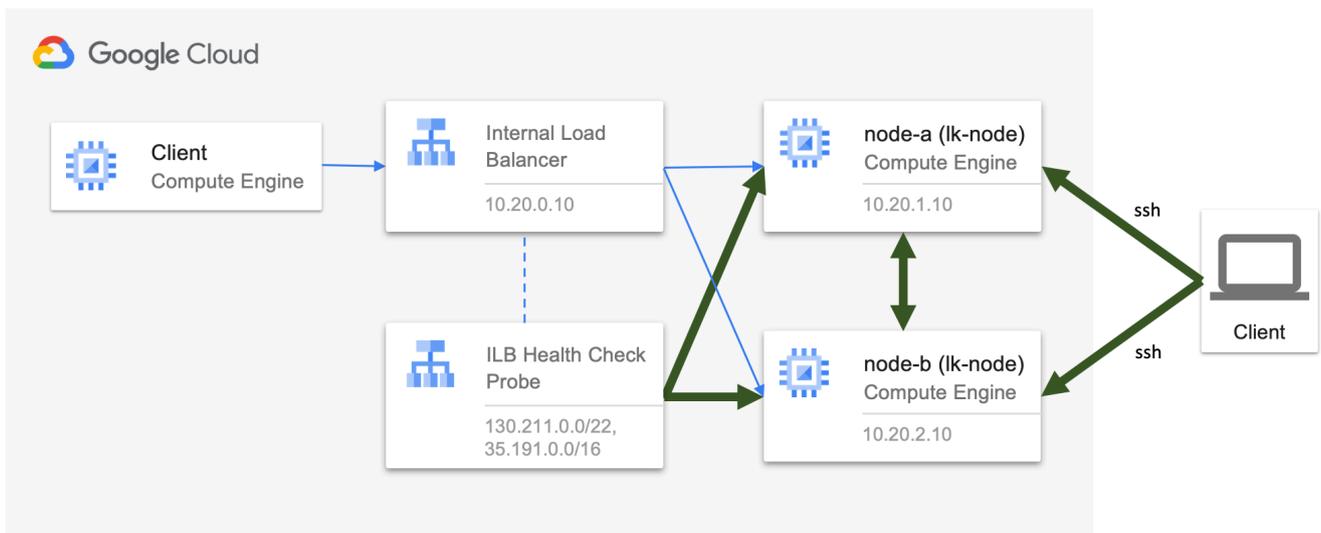
With a set of firewall rules, you can define both “allow” and “deny” rules for each type of traffic. A source can be an internet address, an Internal Load Balancer, or a group of VMs. A group of VMs is identified by a tag assigned to a VM.

Refer to the [Google Cloud documentation](#) for more information.

Once you go to “VPC network” > “Firewall”, a set of firewall rules are already defined for the “default” VPC. Now we are going to define the following rules for `lk-vpc`.

- Allow ssh traffic from your remote work location to VMs with `lk-node` tag. (`fw-allow-ssh`)
- Allow all traffic between VMs with `lk-node` tag. (`fw-allow-lk-node-connection`)
- Allow all health check probes from the Internal Load Balancer to VMs with `lk-node` tag. (`fw-allow-health-check`)

The traffic allowed by these firewall rules is highlighted by the thick arrows in the diagram below.



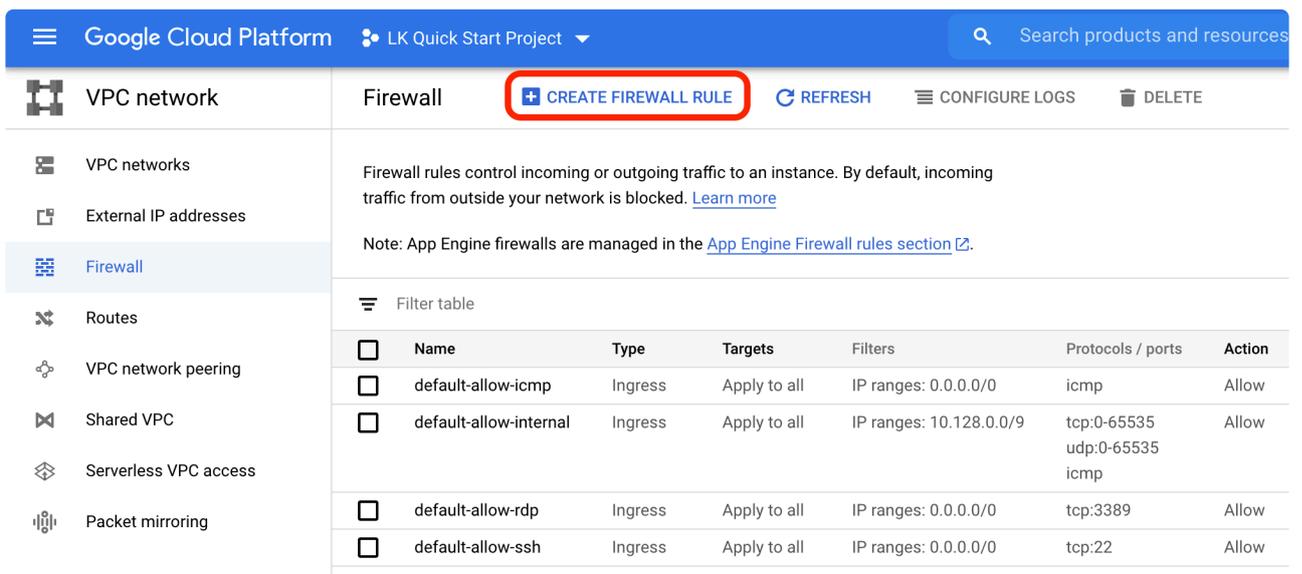
The following table outlines how we should configure these rules.

Name	Parameter	Value
 Common values across VM	Network	lk-vpc
	Type	Ingress
	Target	lk-node
	Action	Allow
fw-allow-ssh	Priority	1000

	Source	WAN IP Address of work location
	Protocols / ports	tcp:22
fw-allow-lk-node-connection	Priority	1100
	Source	Tags: lk-node
	Protocols / ports	all
fw-allow-health-check (*1)	Priority	1200
	Source	130.211.0.0/22, 35.191.0.0/16
	Protocols / ports	all

(\*1) [Google Cloud Documentation on Load Balancer Health Check Probe](#)

1. Select “VPC network” > “Firewall” from the home screen. Now you see a list of firewall rules for the default vpc. Click “CREATE FIREWALL RULE” located at the top and create the first rule “fw-allow-ssh”.



2. The following steps describe how to allow an ssh connection from your work location so that you will be able to configure the node from your work location. In the “Create a firewall rule” wizard, specify the name as fw-allow-ssh.

## ← Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. [Learn more](#)

Name \*

fw-allow-ssh



Lowercase letters, numbers, hyphens allowed

3. Select Network as lk-vpc.

Network \*

lk-vpc



4. Priority can be any value between 0 and 65535. Leave the default value of 1000. For other rules, you can use 1100, 1200, etc. This will allow for other priority values to be injected if additional rules are needed in the future.

Priority \*

1000

[CHECK PRIORITY OF OTHER FIREWALL RULES](#)



Priority can be 0 - 65535

5. Select “Ingress” for “Direction of traffic” and “Allow” for “Action on match”. With these selections incoming traffic matched with patterns specified in the following steps are “Allowed”.

Direction of traffic

Ingress

Egress

Action on match

Allow

Deny

6. Select Targets (destination of the traffic). Make sure to select the `lk-node` tag that is assigned to the VM you created.

**Targets**  
Specified target tags ▼ ?

**Target tags \***  
lk-node ×

7. Select source (WAN IP address(es) of your work location).

**Source filter**  
IP ranges ▼ ?

**Source IP ranges \***  
50.100.100.100/32 × 73.100.100.100/32 × for example, 0.0.0.0/0, 192.168 ?

8. Finally, select the protocol to allow. Confirm all values you have entered and create a rule.

**Protocols and ports** ?

- Allow all
- Specified protocols and ports

tcp :

udp :

Other protocols

[▼ DISABLE RULE](#)

**CREATE**

CANCEL

9. Now a rule is defined. Use the same steps to create other rules.

Firewall

[+ CREATE FIREWALL RULE](#)

[REFRESH](#)

[CONFIGURE LOGS](#)

[DELETE](#)

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. [Learn more](#)

Note: App Engine firewalls are managed in the [App Engine Firewall rules section](#).

lk-vpc Filter table

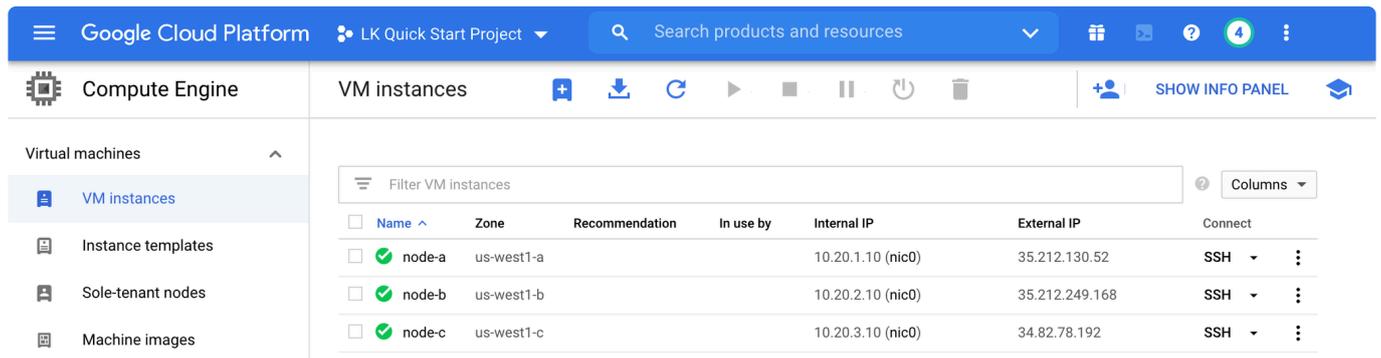
<input type="checkbox"/>	Name	Type	Targets	Filters	Protocols / ports
<input type="checkbox"/>	fw-allow-ssh	Ingress	lk-node	IP ranges: 50.0.0.0/32, 73.0.0.0/32	tcp:22
<input type="checkbox"/>	fw-allow-lk-node-connection	Ingress	lk-node	Tags: lk-node	all
<input type="checkbox"/>	fw-allow-health-check	Ingress	lk-node	IP ranges: 130.211.0.0/22, 35.191.0.0/16	all

# 11.2.3.5.8. Creating the Second and Third VM

The steps required to create the second node (node-b) and third node (node-c) are almost the same as previously described for the first node (node-a). The following table illustrates the differences between the three nodes and the details required to create these instances.

Name	Parameter	Value
 Common values across VM	VPC network	lk-vpc
	Subnet	lk-subnet
	Network Tags	lk-node
	External IP	Ephemeral
node-a	Zone	us-west1-a (*1)
	Hostname	node-a.internal
	Internal IP	10.20.1.10
	Second Disk (Storage Device)	You need to have a second disk
node-b	Zone	us-west1-b (*1)
	Hostname	node-b.internal
	Internal IP	10.20.2.10
	Second Disk (Storage Device)	You need to have a second disk
node-c	Zone	us-west1-c (*1)
	Hostname	node-c.internal
	Internal IP	10.20.3.10
	Second Disk (Storage Device)	You do not need to create a second disk

Once created, the list of VMs should look like this:



The screenshot shows the Google Cloud Platform interface for VM instances. The left sidebar shows 'Virtual machines' with 'VM instances' selected. The main area displays a table of VM instances:

Name	Zone	Recommendation	In use by	Internal IP	External IP	Connect
<input type="checkbox"/> node-a	us-west1-a			10.20.1.10 (nic0)	35.212.130.52	SSH
<input checked="" type="checkbox"/> node-b	us-west1-b			10.20.2.10 (nic0)	35.212.249.168	SSH
<input checked="" type="checkbox"/> node-c	us-west1-c			10.20.3.10 (nic0)	34.82.78.192	SSH

## 11.2.4. Configure Linux Nodes to Run LifeKeeper for Linux

---

This section outlines steps to configure the Linux nodes before installing LifeKeeper for Linux.

In order to connect a Linux node, it may be necessary to install the ssh client software on Windows.

- [Connecting to a Linux Node from Windows Client Using ssh](#)

### Configuring the Nodes

The following steps should be executed for each node:

- [Set a Hostname for Each Instance](#)
- [Disable SELinux](#)
- [Disable the Firewall](#)
- [Set a Password for the Root User](#)
- [Install x11](#)

## 11.2.4.1. Connecting to a Linux Node from Windows Client Using ssh

---

Linux nodes can accept a login using ssh and an ssh client is therefore required. This page explains how to use PuTTY to connect to the Linux node using a public/private key pair (instead of a password). It also covers how to create a basic configuration on the Linux node and install LifeKeeper for Linux. After installing LifeKeeper, we will add more tools to enable x11 to use the GUI based applications.

### Download and Install Tools

Download and install the following software:

- PuTTY (ssh client on Microsoft Windows)

Download the latest PuTTY software and install it.

### Converting the Private Key (.pem)

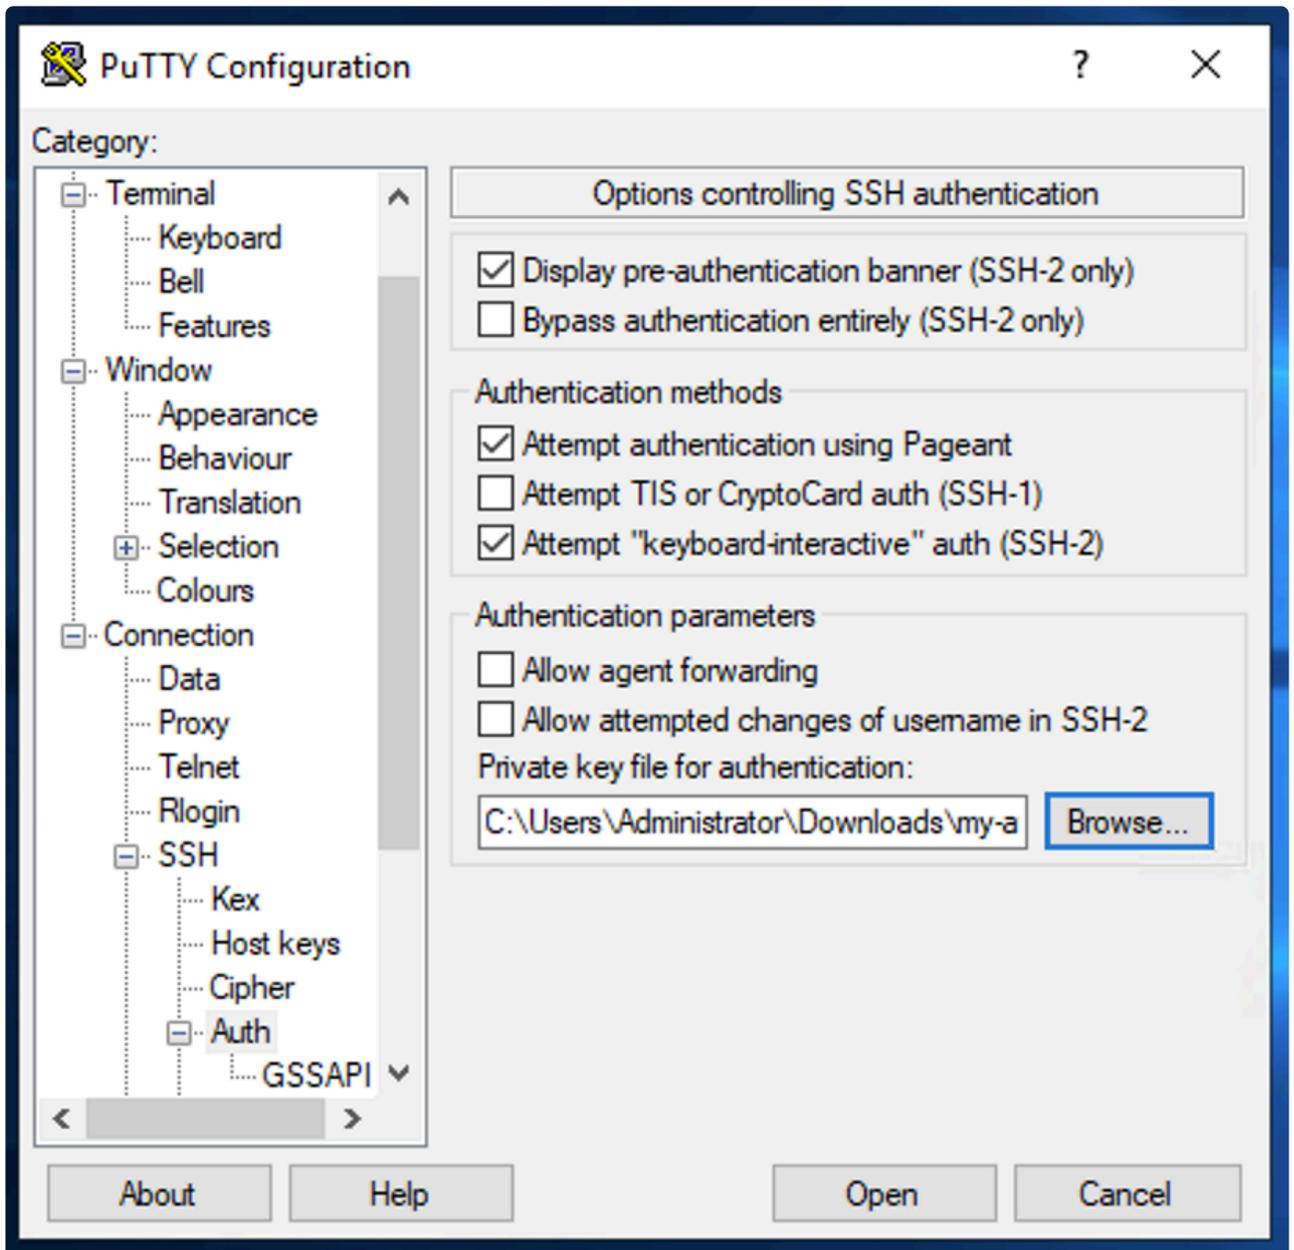
If you have a Private Key file in a `<filename>.pem` format, convert it to a PuTTY Private Key File (`*.ppk`) file.

To convert the file:

1. Start the 'PuTTY Key Generator'.
2. Select Load, then specify your `<filename>.pem` file.
3. Select 'Save private key' as `<filename>.ppk`.

### Configure PuTTY

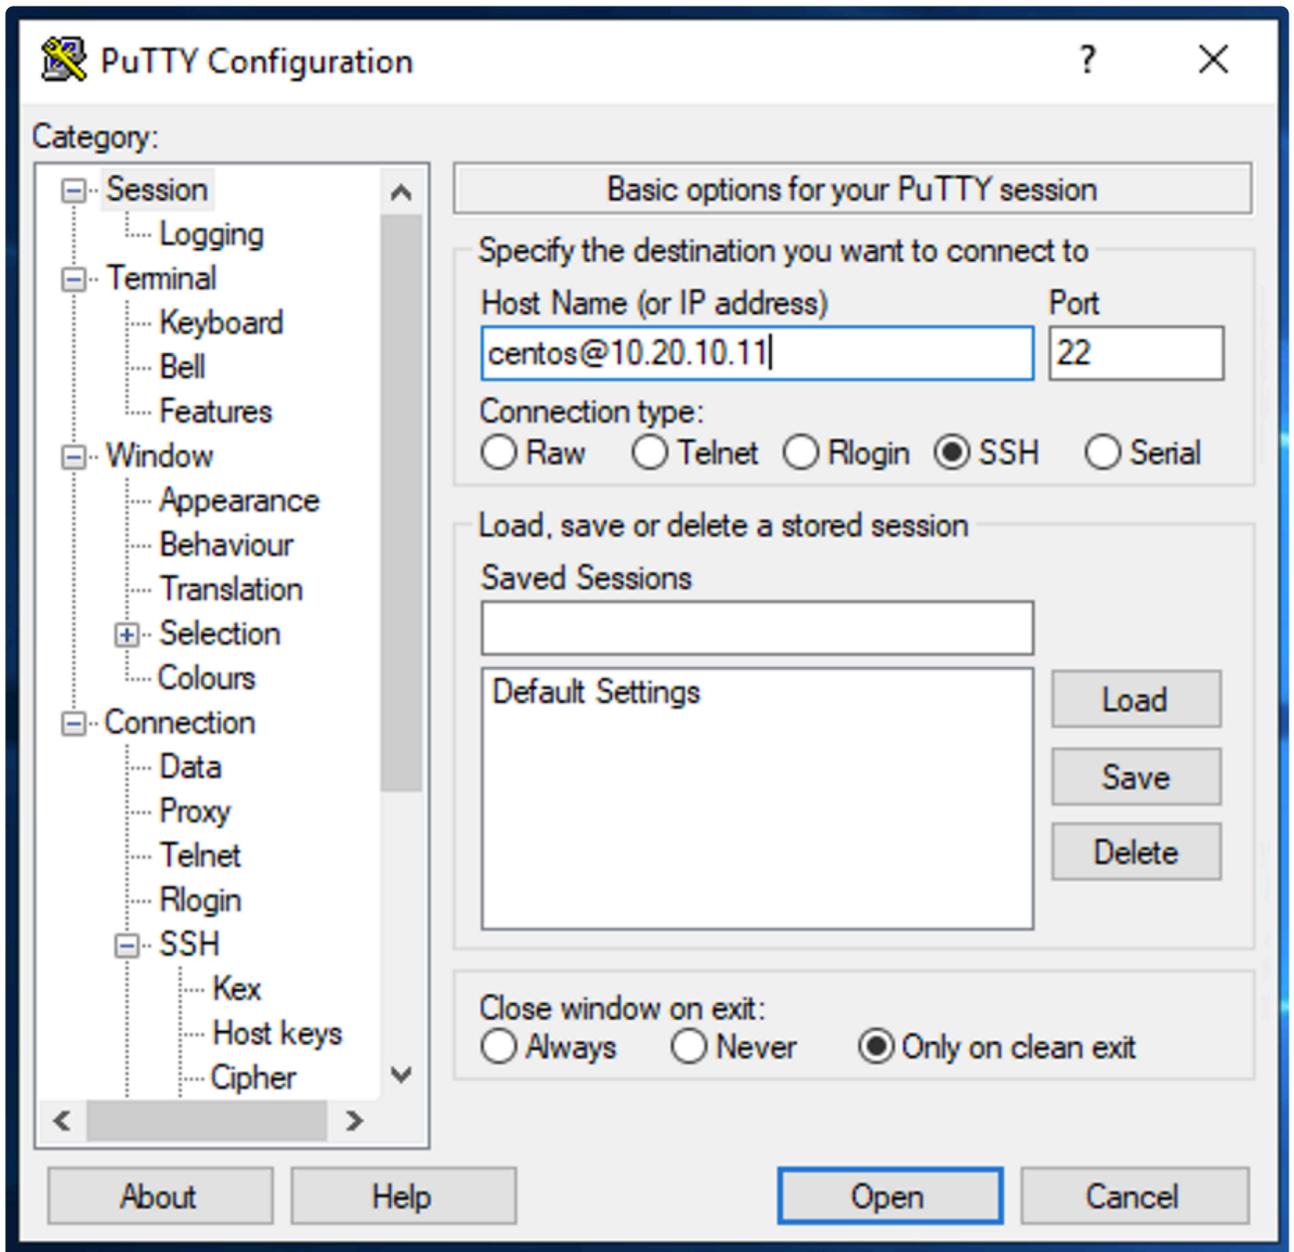
1. Open PuTTY and select Connection > SSH > Auth.
2. Select the private key file at 'Private key file for authentication'.



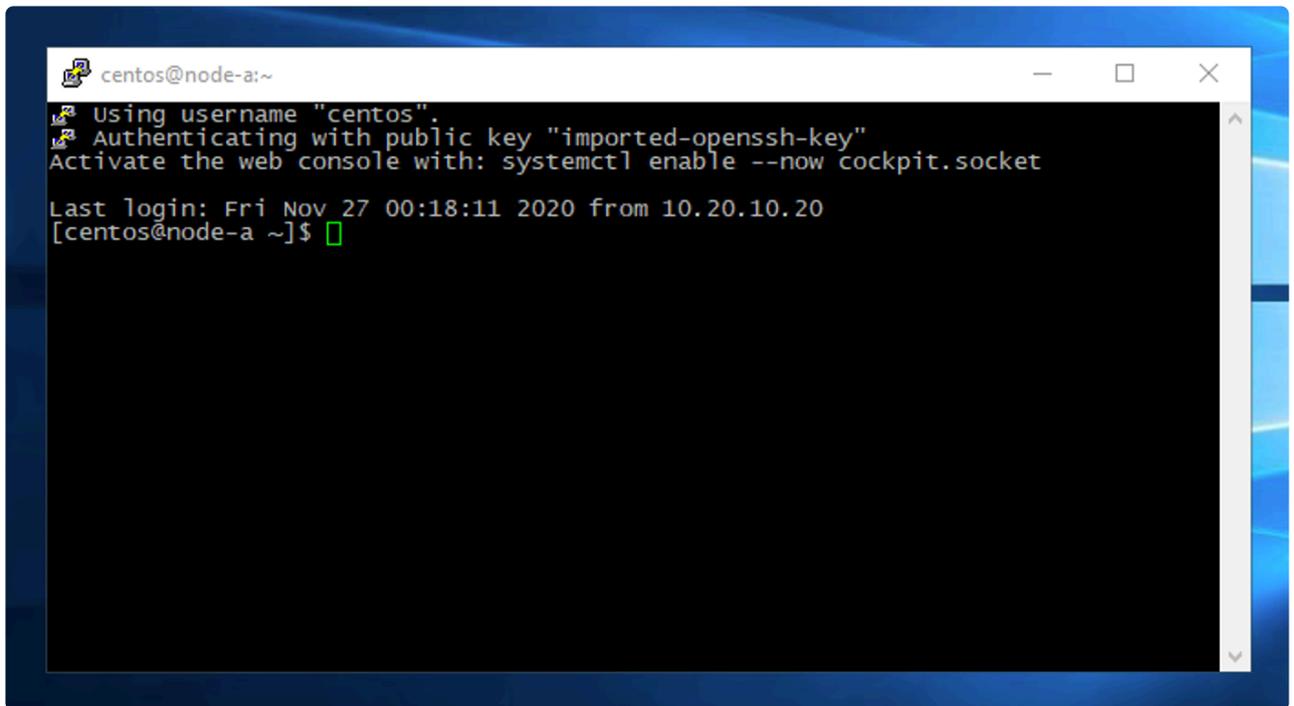
3. You may want to save your configuration as a 'Session' page to avoid reconfiguring each time.

## Connect to a Linux Instance

1. Start the session from PuTTY.



2. Confirm the security alert that indicates the server has not been connected to before. Select "Yes" to proceed.
3. You are now connected!

A terminal window titled 'centos@node-a:~' with standard window controls (minimize, maximize, close). The terminal output shows an SSH login process for the user 'centos'. It indicates authentication with a public key 'imported-openssh-key' and provides instructions to activate the web console using 'systemctl enable --now cockpit.socket'. The last login information is 'Fri Nov 27 00:18:11 2020 from 10.20.10.20'. The prompt is '[centos@node-a ~]\$' with a green cursor.

```
centos@node-a:~  
Using username "centos".  
Authenticating with public key "imported-openssh-key"  
Activate the web console with: systemctl enable --now cockpit.socket  
Last login: Fri Nov 27 00:18:11 2020 from 10.20.10.20  
[centos@node-a ~]$
```

## 11.2.4.2. Set a Hostname for Each Instance

✿ Execute the following commands for each Node.

Once the instances are created, set a hostname for each node and then add an entry to `/etc/hosts` file so that the nodes can locate each other. You can also use DNS if it's available in your environment.

1. use `hostnamectl` command to set a hostname:

```
1 $ sudo su -
2 # hostname
3 ip-10-20-1-10.us-west-2.compute.internal
4 # hostnamectl set-hostname node-a
```

2. If a DNS server is available, register your hostname(s) using the server. If there is not a DNS server available, using `/etc/hosts` is recommended to resolve the hostnames.

```
1 # cat /etc/hosts
2 127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
3 ::1        localhost localhost.localdomain localhost6 localhost6.localdomain6
```

3. Add the following lines to the `/etc/hosts` file so that `node-a` and `node-b` can be resolved to an IP address(es).

```
1 10.20.1.10 node-a
2 10.20.2.10 node-b
3 10.20.3.10 node-c
```

4. Create a backup of `/etc/hosts`, edit the file and confirm the changes.

```
1 # cp -p /etc/hosts /etc/hosts.org
2 # vi /etc/hosts
3 # cat /etc/hosts
4 127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
5 ::1        localhost localhost.localdomain localhost6 localhost6.localdomain6
6 10.20.1.10 node-a
7 10.20.2.10 node-b
8 10.20.3.10 node-c
```

5. Make sure you can ping these nodes by name.

```
1 # ping node-b
2 PING node-b (10.20.2.10) 56(84) bytes of data.
3 64 bytes from node-b (10.20.2.10): icmp_seq=1 ttl=64 time=0.868 ms
4 64 bytes from node-b (10.20.2.10): icmp_seq=2 ttl=64 time=0.892 ms
5 ^C
```

## 11.2.4.3. Disable SELinux

Security-Enhanced Linux (SELinux) is a Linux kernel security module that provides a mechanism for supporting access control security policies, including mandatory access controls (MAC).

[https://selinuxproject.org/page/Main\\_Page](https://selinuxproject.org/page/Main_Page)

✿ You will need to disable it in order to use LifeKeeper for Linux.

### Steps to Disable SELinux

1. Check `sestatus`.

```
1 $ sudo su
2 # sestatus
3 SELinux status:                enabled
4 SELinuxfs mount:              /sys/fs/selinux
5 SELinux root directory:       /etc/selinux
6 Loaded policy name:           targeted
7 Current mode:                 enforcing
8 Mode from config file:        enforcing
9 Policy MLS status:            enabled
10 Policy deny_unknown status:   allowed
11 Memory protection checking:   actual (secure)
12 Max kernel policy version:    32
```

2. Set `enforce` to 0.

```
1 # set enforce 0
```

3. Confirm current value specified at `/etc/selinux/config`.

```
1 # grep -e SELINUX= /etc/selinux/config
2 # SELINUX= can take one of these three values:
3 SELINUX=enforcing
```

4. In step 3, SELinux is set to `enforcing`. Replace the value with `disabled`.

```
1 # cp -p /etc/selinux/config /etc/selinux/config.org
2 # sed -i -e 's/SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config
```

## 5. Confirm the change.

```
1 # grep -e SELINUX= /etc/selinux/config
2 # SELINUX= can take one of these three values:
3 SELINUX=disabled
```

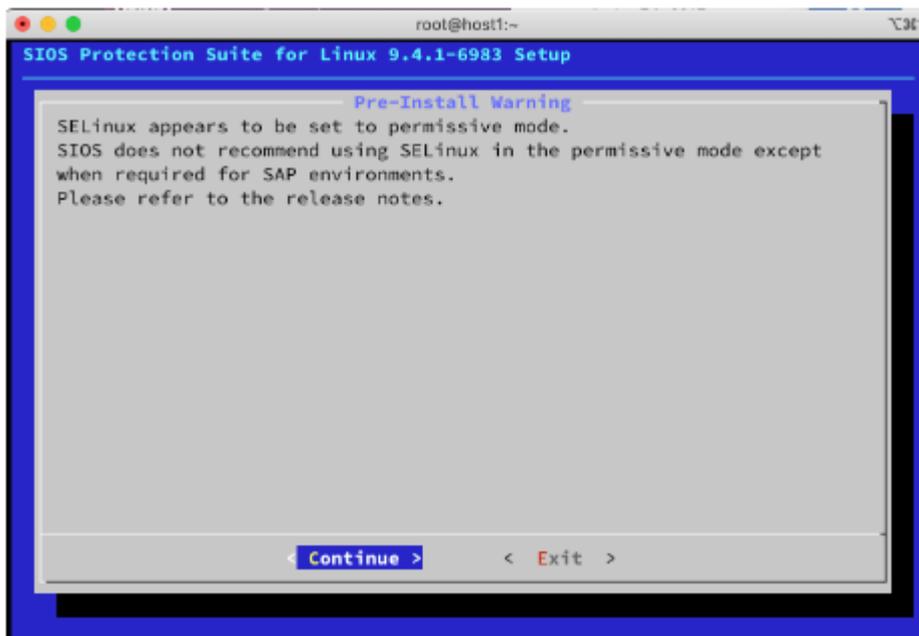
## 6. Restart the machine.

```
1 # shutdown -r now
```

## 7. After the restart, check the status again.

```
1 $ sestatus
2 SELinux status:                disabled
```

! In the latest LifeKeeper-Linux versions “permissive” mode is now supported. The installer will still warn if set to permissive but we can safely continue.



## 11.2.4.4. Disable the Firewall

\* The following commands must be executed for each node.

In this section we will disable the `firewalld` service to install LifeKeeper for Linux. The firewall can be left active but the ports required by LifeKeeper and the application you are protecting will have to be configured (not shown here).

### Steps to Disable `firewalld`

Disable `firewalld` using the following steps.

1. Check the current status.

```
1 # firewall-cmd --state
2 running
```

! If the following is displayed (indicating that `firewalld` is not installed), steps 2 and 3 can be skipped.

```
1 # firewall-cmd --state
2 bash: firewall-cmd: command not found
```

2. Stop and disable `firewalld`.

```
1 # systemctl stop firewalld
2 # systemctl disable firewalld
3 Removed symlink /etc/systemd/system/multi-user.target.wants/firewalld.service.
4 Removed symlink /etc/systemd/system/dbus-org.fedoraproject.FirewallD1.service.
5
6 # systemctl mask --now firewalld
7 Created symlink from /etc/systemd/system/firewalld.service to /dev/null.
```

3. Confirm that it is disabled.

```
1 # firewall-cmd --state
2 not running
```

## 11.2.4.5. Set a Password for the Root User

\* The following commands must be executed for each node.

This tutorial uses the `root` user to login into the LifeKeeper for Linux GUI and it requires a password to perform operations. Please define a password for root user.

### 1. Set a `root` password.

```
1 # passwd
2 Changing password for user root.
3 New password:
4 Retype new password:
5 passwd: all authentication tokens updated successfully.
```

## 11.2.4.6. Install x11

The LifeKeeper for the Linux GUI uses the X Window System. We need to install it on each node.

- Install x11 to the machine.
- Install x11 related packages and enable it with the following steps:

### 1. Install x11 related packages.

```
# yum install xorg-x11-server-Xorg xorg-x11-xauth -y
Last metadata expiration check: 0:00:51 ago on Wed 18 Nov 2020 12:36:11 AM UT
C.
Dependencies resolved.
=====
Package Architecture Version Re
pository Size
=====
Installing:
xorg-x11-server-Xorg x86_64 1.20.8-6.el8 rh
el-8-appstream-rhui-rpms 1.5 M
xorg-x11-xauth x86_64 1:1.0.9-12.el8 rh
el-8-appstream-rhui-rpms 39 k
(snip)
xorg-x11-drv-libinput-0.29.0-1.el8.x86_64 xorg-x11-drv-ves
a-2.4.0-3.el8.x86_64
xorg-x11-server-Xorg-1.20.8-6.el8.x86_64 xorg-x11-server-co
mmon-1.20.8-6.el8.x86_64
xorg-x11-xauth-1:1.0.9-12.el8.x86_64 xorg-x11-xkb-util
s-7.7-28.el8.x86_64
Complete!
```

- ### 2. Check the X11Forwarding parameter in /etc/ssh/sshd\_config to confirm that it is set to yes. If it is not set to yes, edit the file.

```
1 # grep -e X11Forwarding /etc/ssh/sshd_config
2 X11Forwarding yes
```

- ### 3. Assuming the system is booted to X11, start GUI sessions at start-up. If not, this step can be ignored.

```
# systemctl set-default graphical.target
```

```
Removed /etc/systemd/system/default.target.  
Created symlink /etc/systemd/system/default.target → /usr/lib/systemd/system/graphical.target.
```

#### 4. Restart the machine.

```
1 # shutdown -r now
```

## 11.2.5. Install LifeKeeper for Linux

\* The following commands should be executed for each node.

### Mount ISO Image for Install

1. Download the ISO image and license key files to the local environment.
2. Copy the ISO image to each of the nodes.
3. Copy the license file to each of the nodes.
4. Log into each node and mount the installer image.

```

1 $ ls
2 evalkeys.txt sps951.img
3 $ sudo su
4 # mount -o loop -t iso9660 sps951.img /media
5 mount: /media: WARNING: device write-protected, mounted read-only.
```

5. Check the content of the image.

```

1 # cd /media
2 # ls
3 CentOS      core                java      README      setuplibs
4 Chef        create_response_file  kits     RHAS        shfuncs
5 common      HADR-9.5.1-7154.src.rpm  OEL     sapconnector SuSE
6 COPYRIGHT  HADR-generic-9.5.1-7154.x86_64.rpm  quorum  setup      TRANS.TBL
```

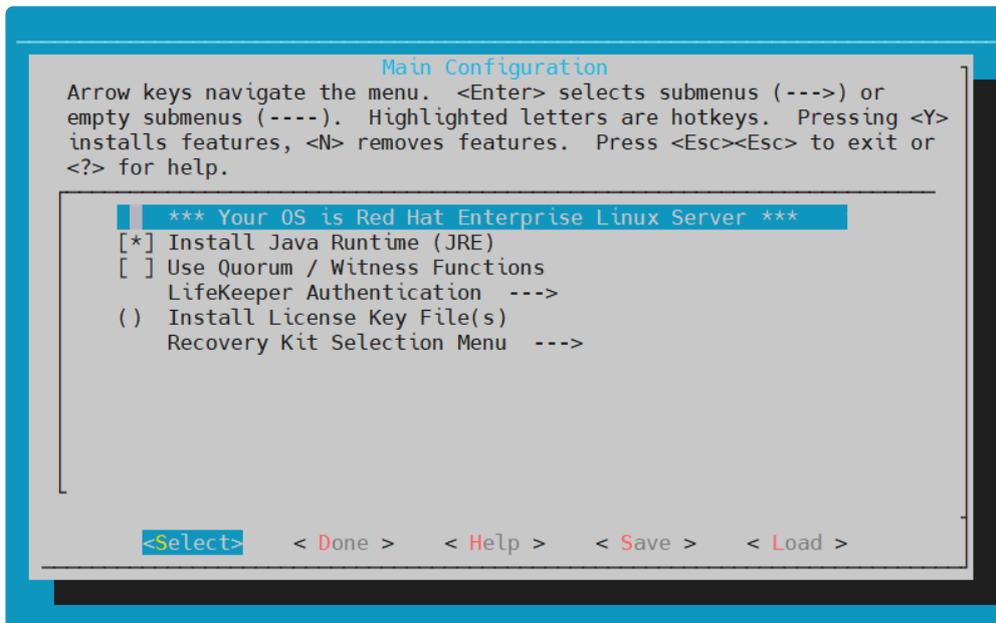
### Install LifeKeeper for Linux

The installer should be located at `/media/setup`. Execute the setup script to install LifeKeeper for Linux.

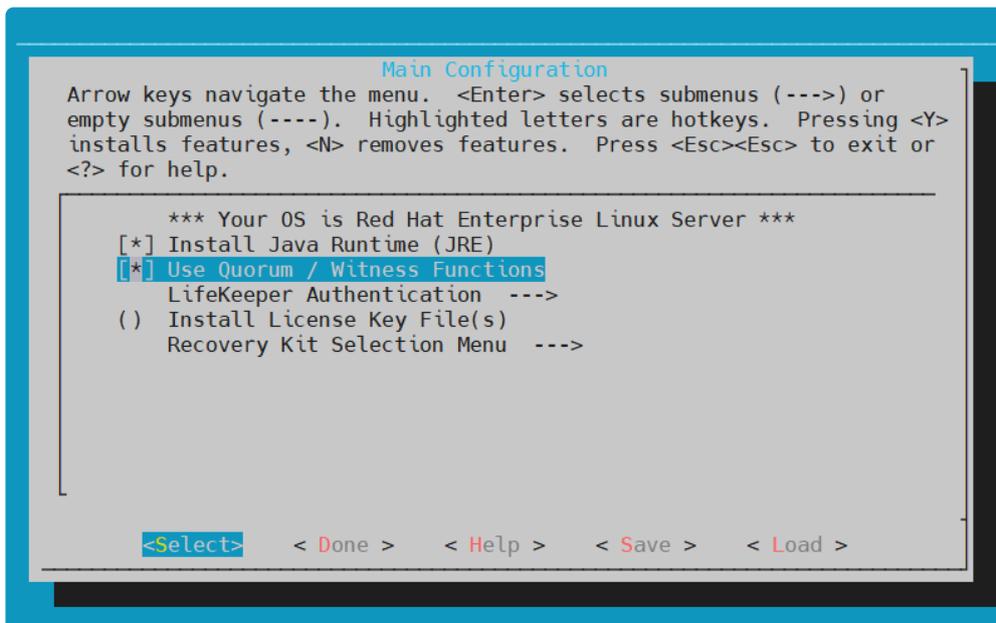
```

1 # ./setup
2 SIOS Protection Suite for Linux Setup
3 Validating files.....OK
4 Collecting system information.....done.
5 Preparing configuration information.....done.
```

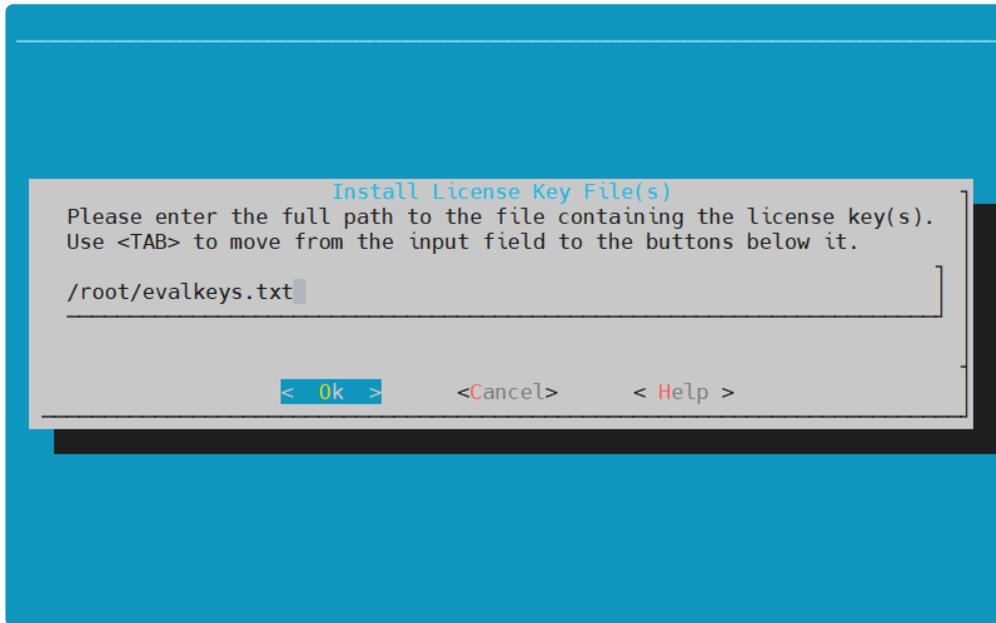
1. The GUI Installer starts as shown below.



2. Check "Use Quorum / Witness Functions".

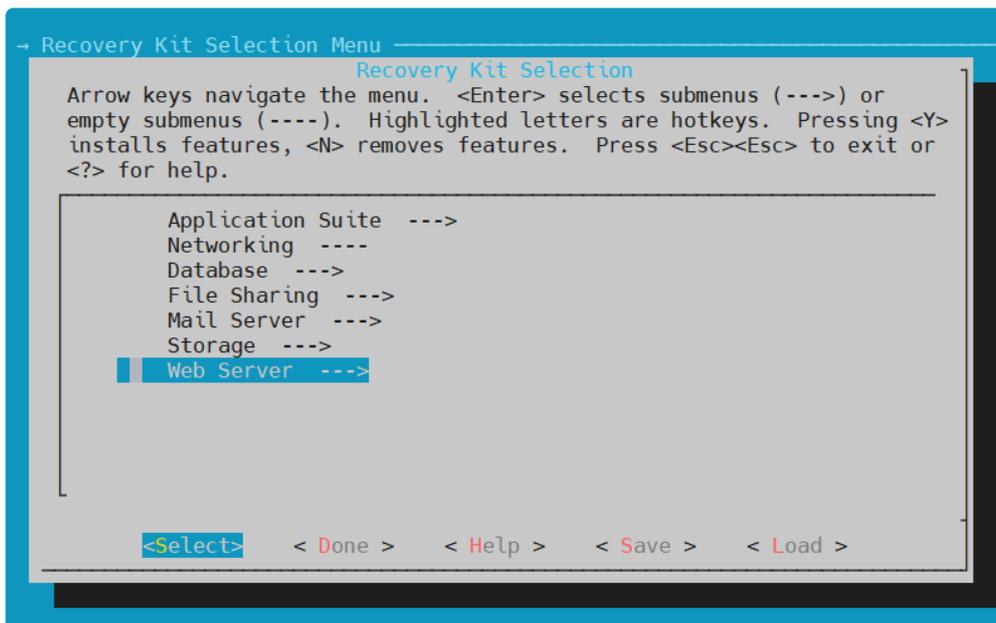


3. Specify the location of the license file.

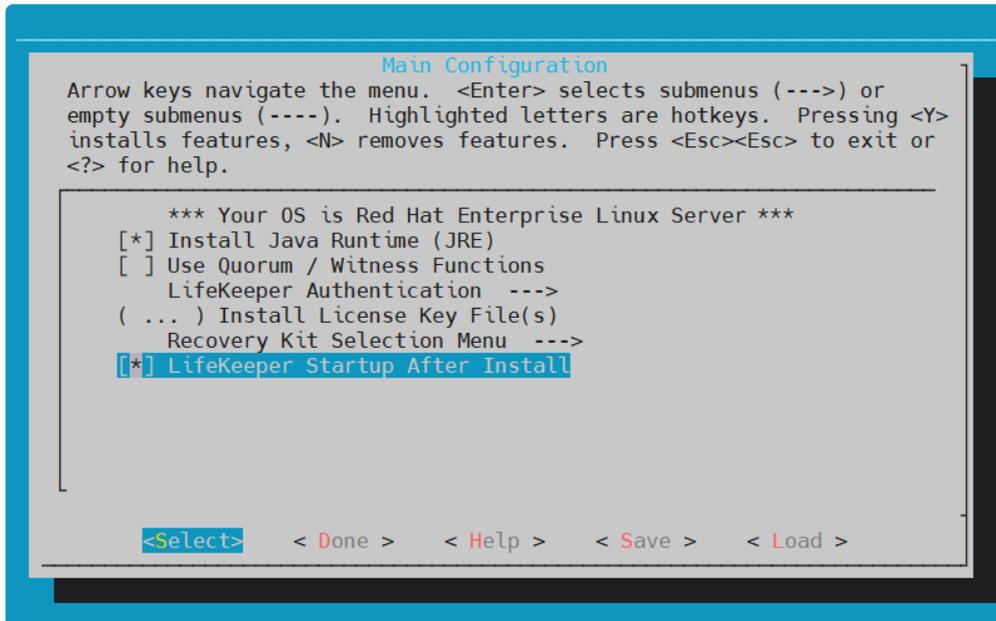


4. Select the Recovery Kits to be installed.

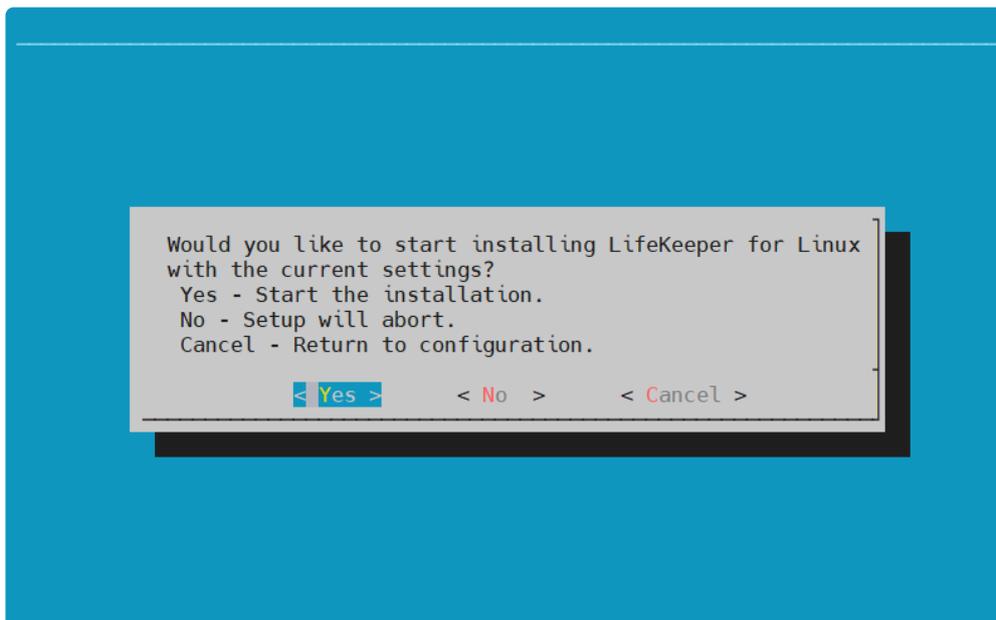
**!** On node-c (the Witness Node), **NO** recovery kit should be selected.



5. Check "LifeKeeper Startup After Install", then select "Done".



6. Start the Install process.



7. Install is now completed. Repeat the steps for node-b and node-c.

```
Preparing configuration information.....done.
Performing package installation and updating configuration information for SPS f
or Linux.
Install LifeKeeper and dependent packages done.
Setup high availability data replication features.. done.
Configure LifeKeeper management group
Install licenses.
Starting LifeKeeper...

Broadcast message from systemd-journald@node-a (Wed 2020-11-18 01:51:08 UTC):

lcdinit[12776]: EMERG:lcd.lcdchksem1:::011138:The LifeKeeper product on this sys
tem is using an evaluation license key which will expire at midnight on 02/08/21
. To continue functioning beyond that time, a permanent license key must be obta
ined.

Message from syslogd@node-a at Nov 18 01:51:08 ...
lcdinit[12776]:EMERG:lcd.lcdchksem1:::011138:The LifeKeeper product on this sys
tem is using an evaluation license key which will expire at midnight on 02/08/21
. To continue functioning beyond that time, a permanent license key must be obta
ined.
Setup complete.
[root@node-a media]# █
```

## Additional Environment Specific Tasks to Complete

If AWS is being used, complete the following steps:

- [Disable PING Broadcasting](#)
- [Install AWS CLI](#)
- [Assign Permission to Use EC2 Recovery Kit](#)

If Azure is being used, complete the following step:

- [Disable PING Broadcasting](#)

## 11.2.5.1. Install AWS CLI

\* The following commands must be executed on each node.

First, the AWS CLI must be installed. Detailed steps are documented at [Installing, updating, and uninstalling the AWS CLI version 2 on Linux](#).

Once the CLI is installed, add the location of the CLI to the `/etc/default/LifeKeeper` file.

1. Check the location of `aws` CLI.

```
1 # curl "https://awscli.amazonaws.com/awscli-exe-linux-x86_64.zip" -o "awscliv2.zip"
2 # unzip awscliv2.zip
3 # sudo ./aws/install
4 You can now run: /usr/local/bin/aws --version
```

2. Update `/etc/default/LifeKeeper` to include this location in the `PATH`.

```
1 # vi /etc/default/LifeKeeper
2 # grep -e PATH= /etc/default/LifeKeeper
3 PATH=/opt/LifeKeeper/bin:/bin:/usr/bin:/usr/sbin:/sbin:/usr/local/bin
```

## 11.2.5.2. Assign Permission to Use EC2 Recovery Kit



In order to complete the following steps, the AWS CLI must be installed on each node. See [Additional Environment Specific Tasks to Complete](#) for more information.

To use the EC2 Recovery Kit, the instance must have Roles that are allowed to update the RouteTable entries or reassign ENI (Elastic Network Interface).

- ec2:DisassociateAddress
- ec2:DescribeAddresses
- ec2:AssociateAddress
- ec2:DescribeRouteTables
- ec2:ReplaceRoute

To achieve this, create a policy as seen below (note that it might be desirable to limit the resources that may be accessed), then assign it to a Role.

```
1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Sid": "VisualEditor0",
6       "Effect": "Allow",
7       "Action": [
8         "ec2:DisassociateAddress",
9         "ec2:DescribeAddresses",
10        "ec2:AssociateAddress",
11        "ec2:DescribeRouteTables",
12        "ec2:ReplaceRoute"
13      ],
14      "Resource": "*"
15    }
16  ]
17 }
```

Once a Role is defined, assign it to the EC2 instances.

## 11.2.5.3. Disable PING Broadcasting

---

When LifeKeeper is used in a cloud environment, it may not be possible to use broadcast ping to find other machines on the subnet. Therefore, this should be disabled by changing the configuration of LifeKeeper as follows:

1. Check the current configuration.

```
1 # grep -e NOBCASTPING /etc/default/LifeKeeper
2 NOBCASTPING=0      # Can be used to disable the broadcast ping mechanism
```

2. NOBCASTPING must be set to 1.

```
1 # cp -p /etc/default/LifeKeeper /etc/default/LifeKeeper.org
2 # sed -i -e 's/NOBCASTPING=0/NOBCASTPING=1/g' /etc/default/LifeKeeper
```

3. Confirm the change.

```
1 # grep -e NOBCASTPING /etc/default/LifeKeeper
2 NOBCASTPING=1      # Can be used to disable the broadcast ping mechanism
```

# 11.2.5.4. AWS – Disable Source/Destination Checking

Source/destination checking must be disabled for each AWS EC2 instance on which an EC2 Recovery Kit resource will be created.

**Note:** These manual steps are not required in LifeKeeper for Linux v9.6.0 and later as source/destination checking is automatically disabled on each server by the EC2 Recovery Kit when creating or extending LifeKeeper EC2 resources. In order to allow the EC2 Recovery Kit to automatically disable source/destination checks, `ec2:DescribeNetworkInterfaceAttribute` and `ec2:ModifyNetworkInterfaceAttribute` permissions must be added to the IAM roles for the EC2 instances.

In the AWS Console, navigate to **Services** → **EC2** → **Instances**. Right click the EC2 instance for which you would like to disable source/destination checks, select **Networking**, then click **Change source/destination check**.

In the resulting dialog, click the check box next to **Stop** and click **Save** to update the settings for the instance.

### Source / destination check [Info](#)

Each EC2 instance performs source and destination checks by default. The instance must be the source or destination of all the traffic it sends and receives.

Instance ID  
 i-000c007d6a62841d5 (Node-A)

Network interface [Info](#)  
 eni-0c76a1a1925f6597e (Node-A)

Source / destination checking [Info](#)  
 Stop

 If this is a NAT instance, you must stop source / destination checking. A NAT instance must be able to send and receive traffic when the source or destination is not itself.

▼ **AWS CLI Command**

```
aws ec2 modify-instance-attribute --instance-id=i-000c007d6a62841d5 --no-source-dest-check
```

 Copy

Cancel Save

## 11.2.6. Login and Basic Configuration Tasks

By now you should have 3 nodes, each with LifeKeeper for Linux installed and ready for use.

### Connecting to Linux Node with the GUI

LifeKeeper has a GUI application to configure the cluster. Please refer to [Connecting to a Linux Node with x11 Forwarding](#) for more details about accessing the LifeKeeper GUI.

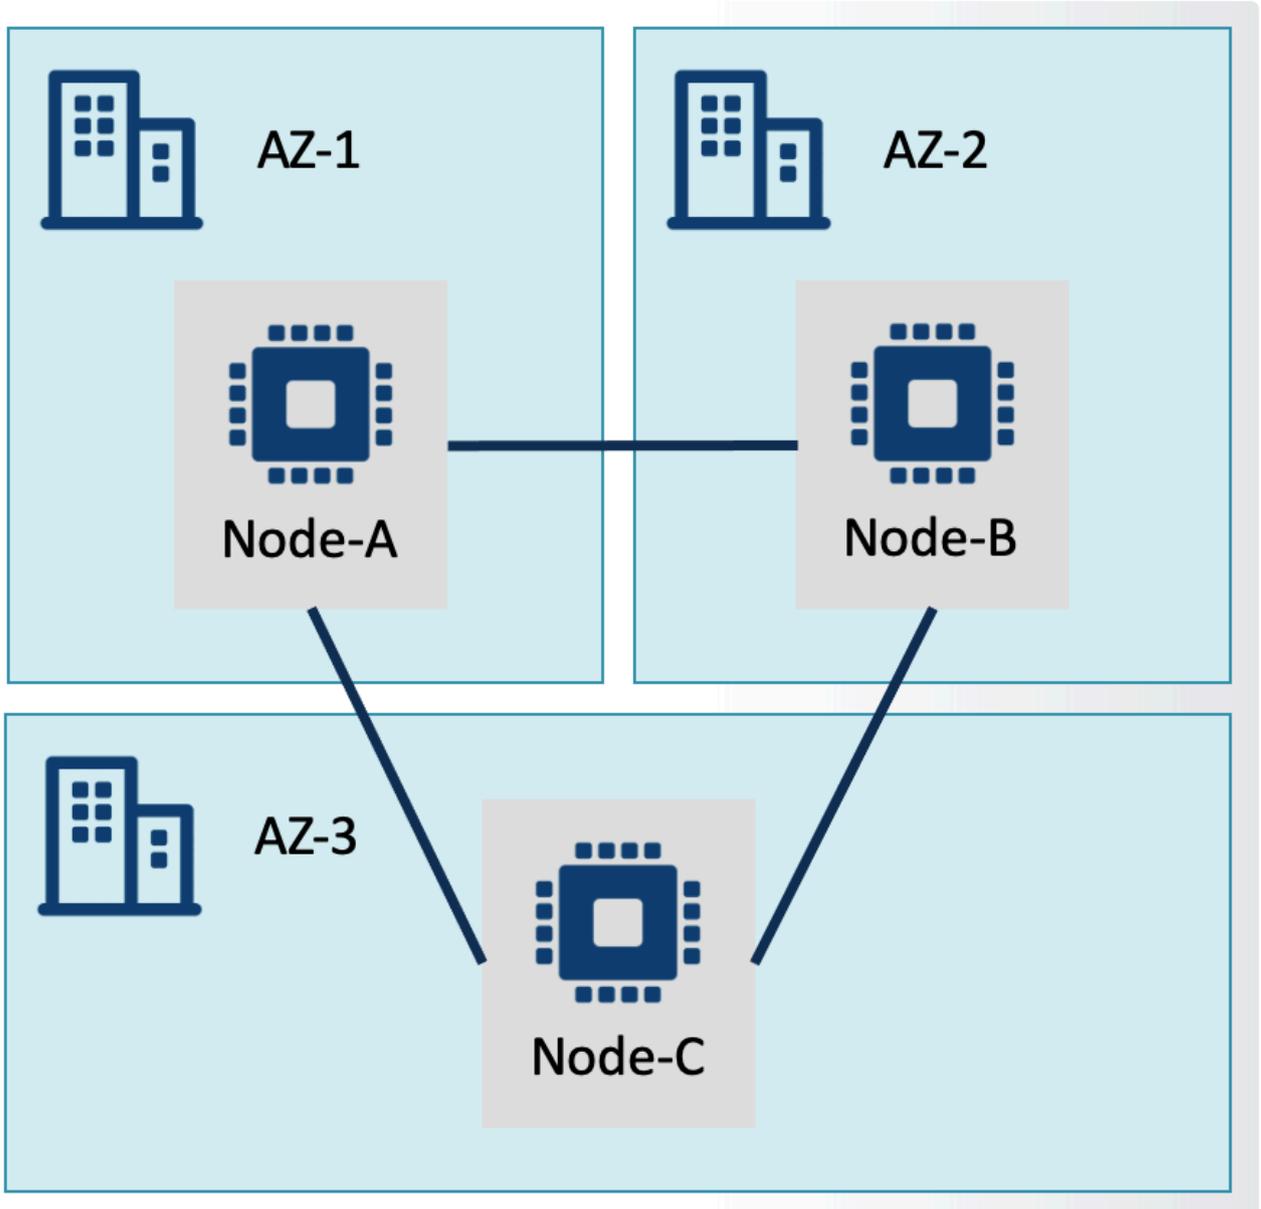
### Basic Configuration

In this tutorial there are three nodes with the following IP address range:

Node	IP Address
node-a (Active)	10.20.1.10
node-b (Standby)	10.20.2.10
node-c (Quorum / Witness)	10.20.3.10

The basic configuration is completed using the following steps:

- [Connecting to the First node](#)
- [Connecting to Other Nodes](#)
- [Define the Communication Path](#) **Note:** There are three nodes in the cluster including node-c (Quorum/Witness). Therefore, 3 communication paths must be defined in order to create bidirectional communication paths between all pairs of nodes.
  - node-a <—> node-b
  - node-a <—> node-c
  - node-b <—> node-c



## 11.2.6.1. Connecting to a Linux Node with “X11 Forwarding”

---

X Window System is used to enable bitmap based Graphical User Interface on Unix based operating systems. This is commonly referred to as “X11”. One of the interesting usages of X11 is the ability to transfer “a window” to the terminal (a machine which an operator uses). The original X11 process runs on the server but the X11 User Interface appears on the client machine terminal, This is called “X11 forwarding”.

In order to use X11 forwarding an X11 graphic server must be run (this is called an X11 server even though it actually runs on the client machine) on the local machine.

The LifeKeeper GUI runs on X Window Platform and therefore a client that runs X11 software is required (unless a GUI login is used on these nodes directly).

### Install Client Software to Microsoft Windows

To use Microsoft Windows as an “x11 server” it is necessary to install software on the Microsoft Windows client. The following page explains the basic steps required to install the software on a Windows client.

- [Setup X Window Client Software on Microsoft Windows](#)

### Setup Linux Nodes to Accept X11 Forwarding

1. Find the current value of the DISPLAY environment variable as the regular user that was set up in the previous steps.

```
$ echo $DISPLAY
localhost:10.0
```

2. View the current list of xauth keys for the user and identify the one associated with the current display that X11 is being forwarded on (10 in this example).

```
$ xauth list
node-a/unix:11 MIT-MAGIC-COOKIE-1 7289ba26871a37b94ff359df829e2686
node-a/unix:10 MIT-MAGIC-COOKIE-1 bae592c842916f23e3ba066ba594c5d0
```

3. Become the root user.

```
$ sudo -i
```

4. View the DISPLAY environment variable for the root user.

```
# echo $DISPLAY
```

```
localhost:10.0
```



**Note:** If the DISPLAY variable is not set for the root user, set it to the same value found in step 1:

```
# export DISPLAY=localhost:10.0
```

5. View the current list of xauth keys for the root user.

```
# xauth list
node-a/unix:1 MIT-MAGIC-COOKIE-1 5b249d176f8da2f7803343f6830002eb
node-a/unix:2 MIT-MAGIC-COOKIE-1 e62f7abaeda1d200576aac40aec2ba83
```

If there is no entry for the display found in step 1, add the relevant entry from the xauth list output in step 2 to the list of keys for the root user.

```
# xauth add node-a/unix:10 MIT-MAGIC-COOKIE-1
bae592c842916f23e3ba066ba594c5d0
```

6. Once the DISPLAY variable and xauth keys are set up correctly for the root user, it is now possible to run graphical applications as root from an SSH session.

```
# /opt/LifeKeeper/bin/lkGUIapp &
[1] 23957
openjdk version "12.0.2" 2019-07-16
OpenJDK Runtime Environment (build 12.0.2+10)
OpenJDK 64-Bit Server VM (build 12.0.2+10, mixed mode, sharing)
Setting up secure random number generator
Random number setup completed
Connecting to rmi://node-a:82/LKRemoteInterface at:
    Sat Mar 13 04:23:19 UTC 2021
Connection to node-a succeeded at:
    Sat Mar 13 04:23:20 UTC 2021
```

## 11.2.6.2. Setup X Window Client Software on Microsoft Windows

---

After following the steps in [Connecting to a Linux Node from Windows Client Using ssh](#), PuTTY has been installed and configured. This following explains how to enable X11 forwarding so that it is possible to connect with nodes and use GUI based applications.

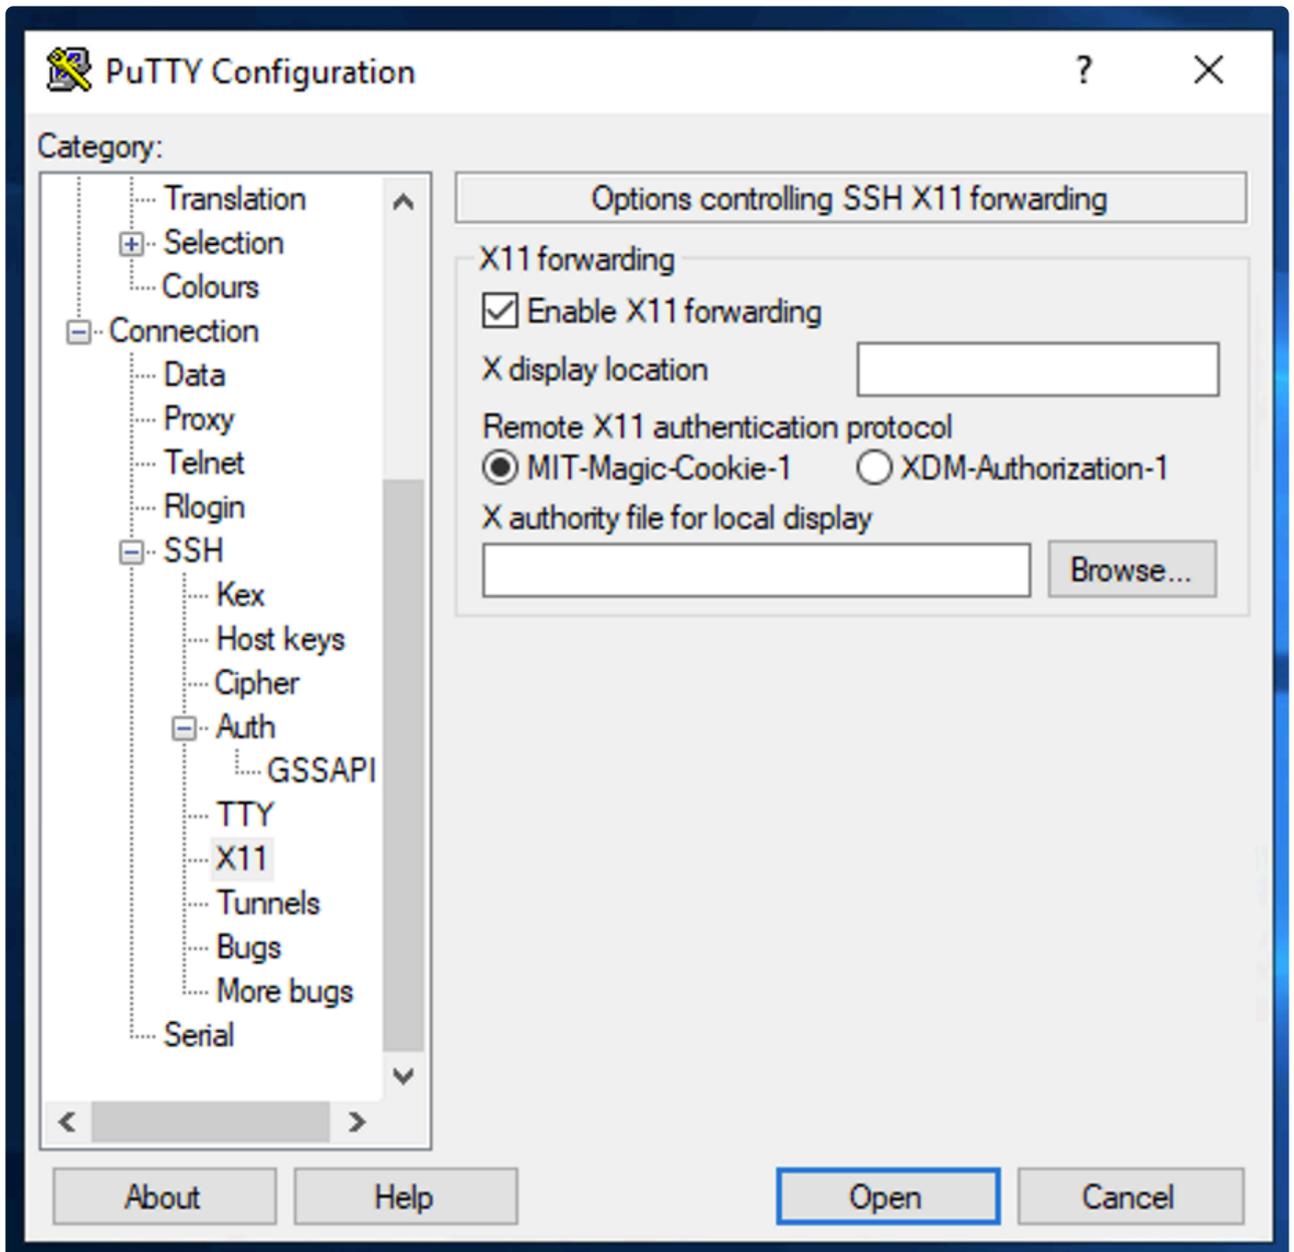
### Download and Install Tools

Download and install the following additional software:

- Xming (X11 server on Microsoft Windows, other X11 windows clients are available)

### Configure PuTTY

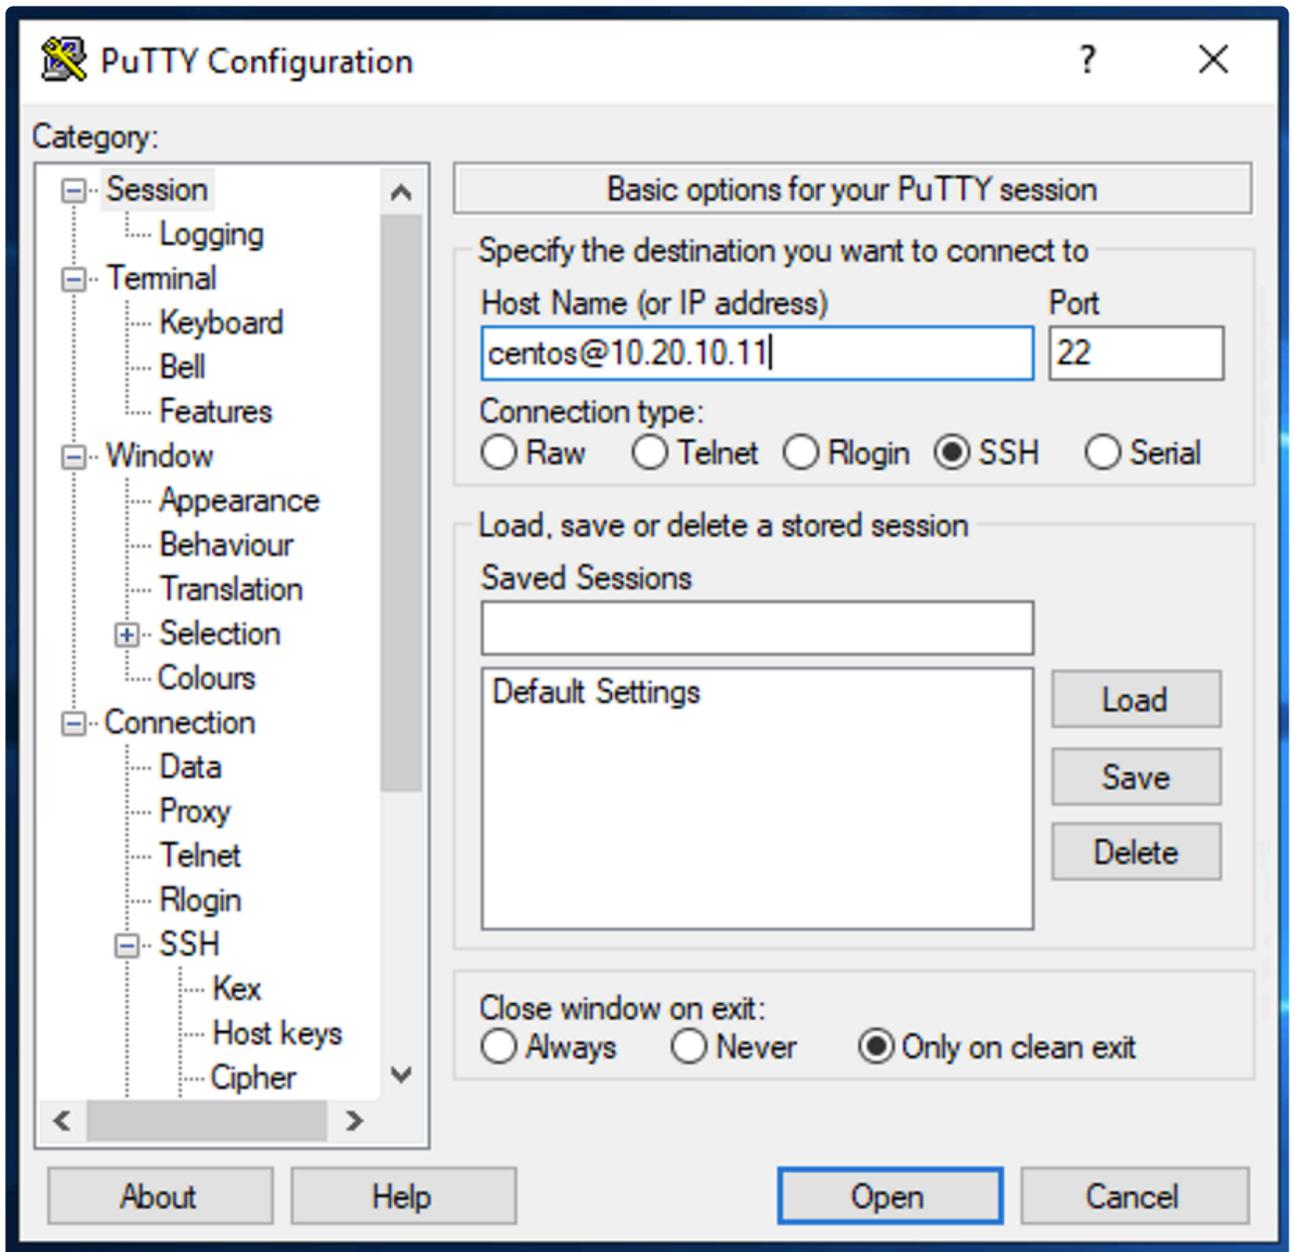
1. Open PuTTY and display Connection > SSH > Auth.
2. Display Connection > SSH > X11 and check 'Enable X11 forwarding'.



3. It is recommended that the configuration is saved on the 'Session' page (otherwise these steps will need to be repeated each time).

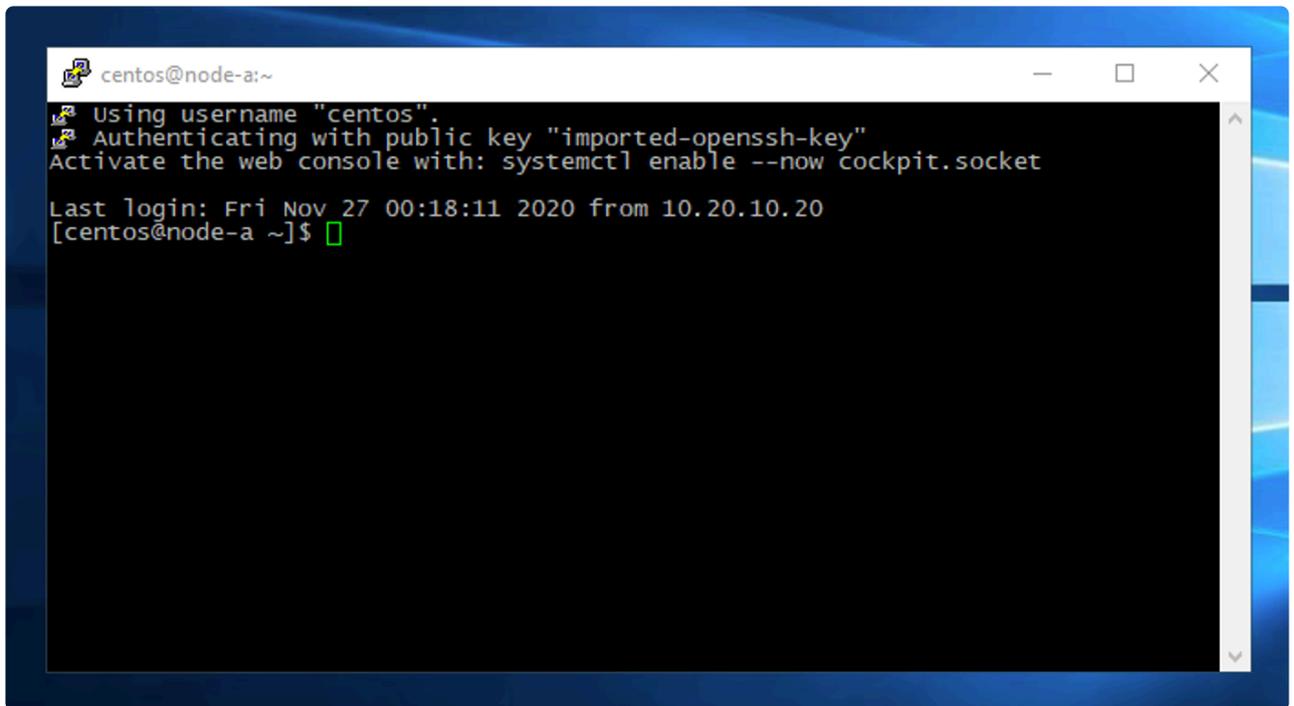
## Connect to a Linux Instance

1. Start Xming (ensure that the Ming icon is shown in the system tray).
2. Start the session from PuTTY.



3. Confirm the security alert indicating the server has not connected before. Select "Yes" to proceed.
4. If the warning `~/.Xauthority` does not exist is seen, use touch to create the file, then reconnect to the node.

You should now be connected.



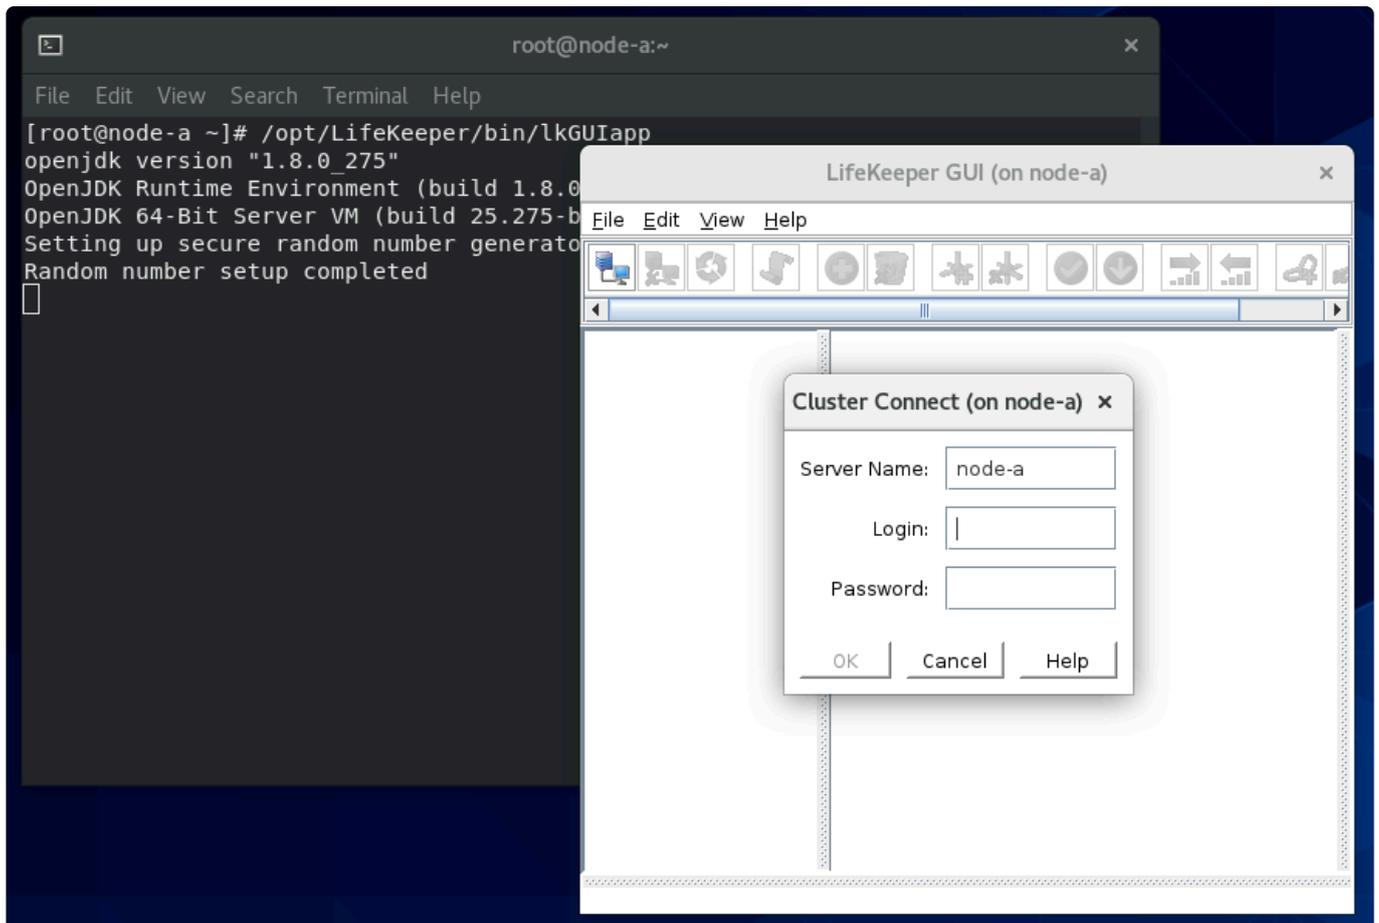
```
centos@node-a:~  
Using username "centos".  
Authenticating with public key "imported-openssh-key"  
Activate the web console with: systemctl enable --now cockpit.socket  
Last login: Fri Nov 27 00:18:11 2020 from 10.20.10.20  
[centos@node-a ~]$
```

## 11.2.6.3. Connecting to the First Node (node-a)

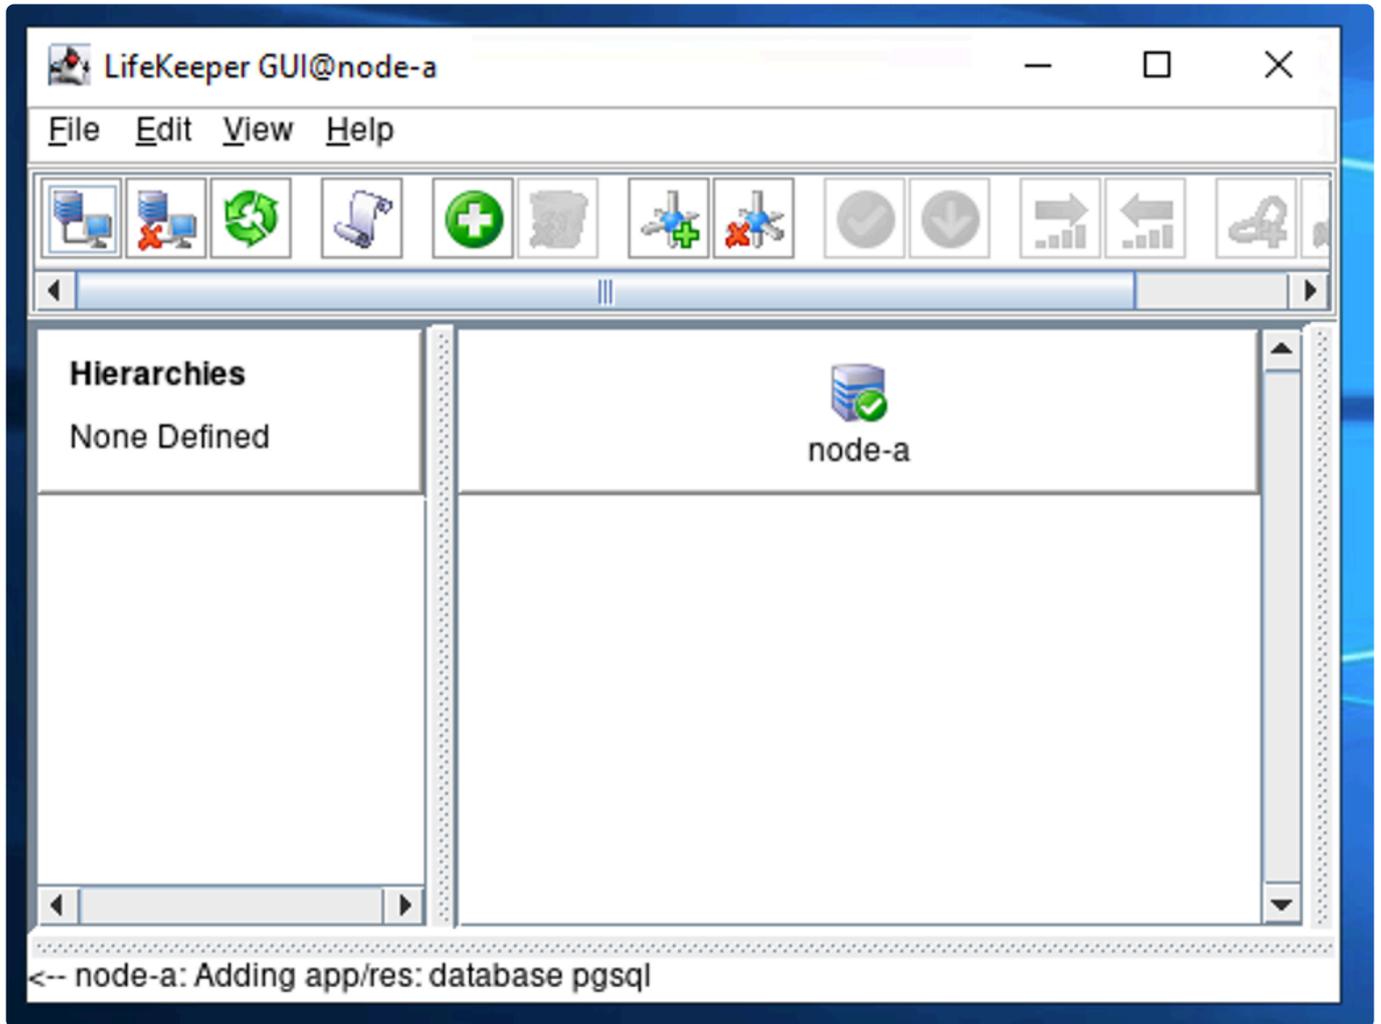
Once logged in to the node with “X11 forwarding” enabled, type the following command to start the LifeKeeper for Linux GUI:

```
1 # /opt/LifeKeeper/bin/lkGUIapp &
```

In the login dialog pop-up, enter the root username and password to login:

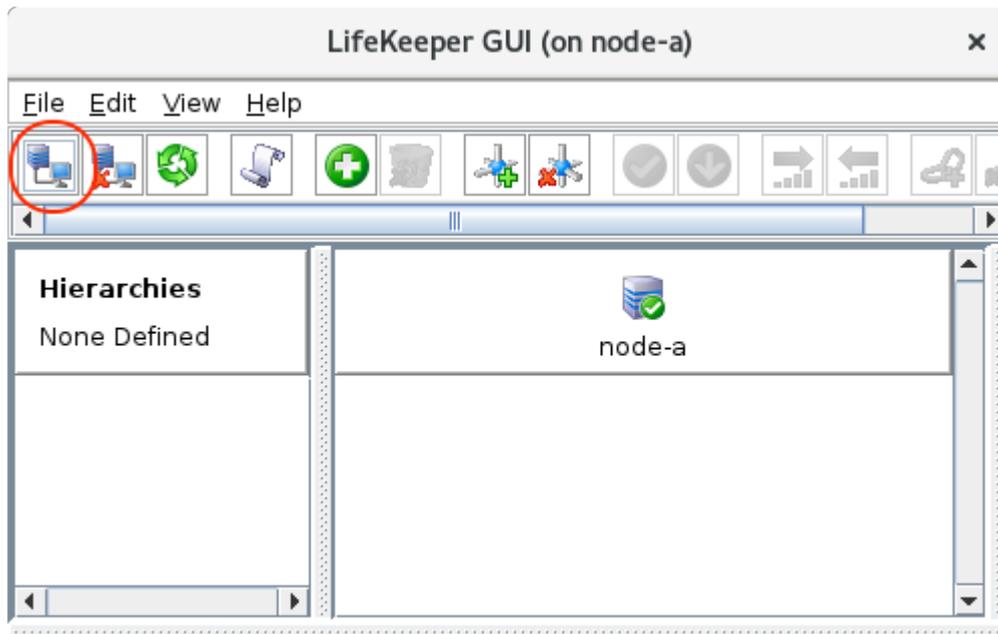


After successfully logging in, you should see node-a in the GUI.



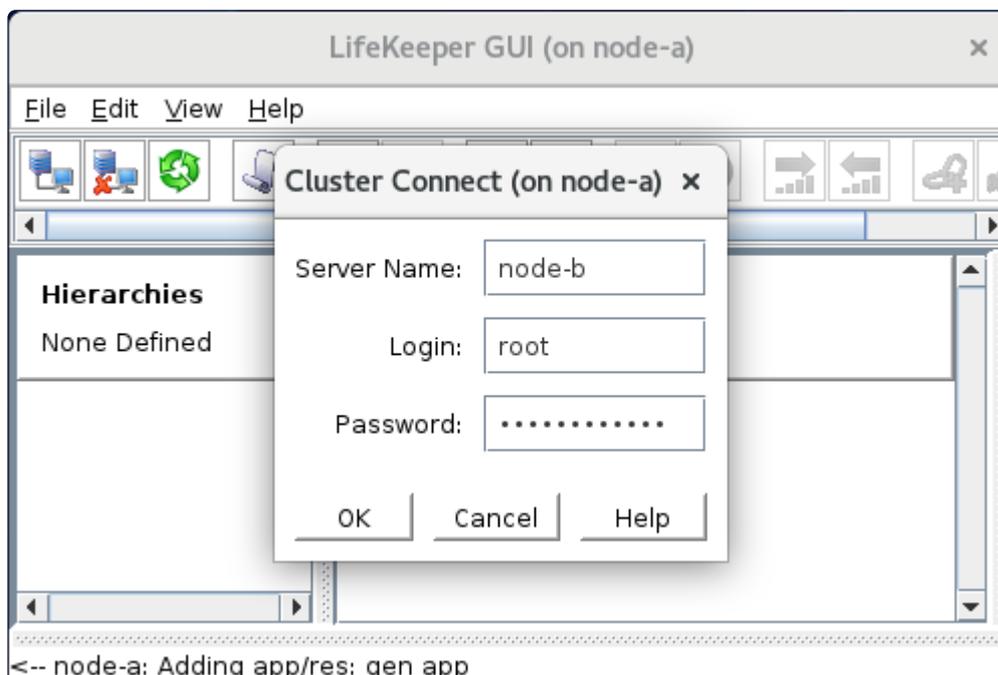
## 11.2.6.4. Connecting to Other Nodes (node-b and node-c)

1. Select "Connect" on LifeKeeper (the first icon on the toolbar).



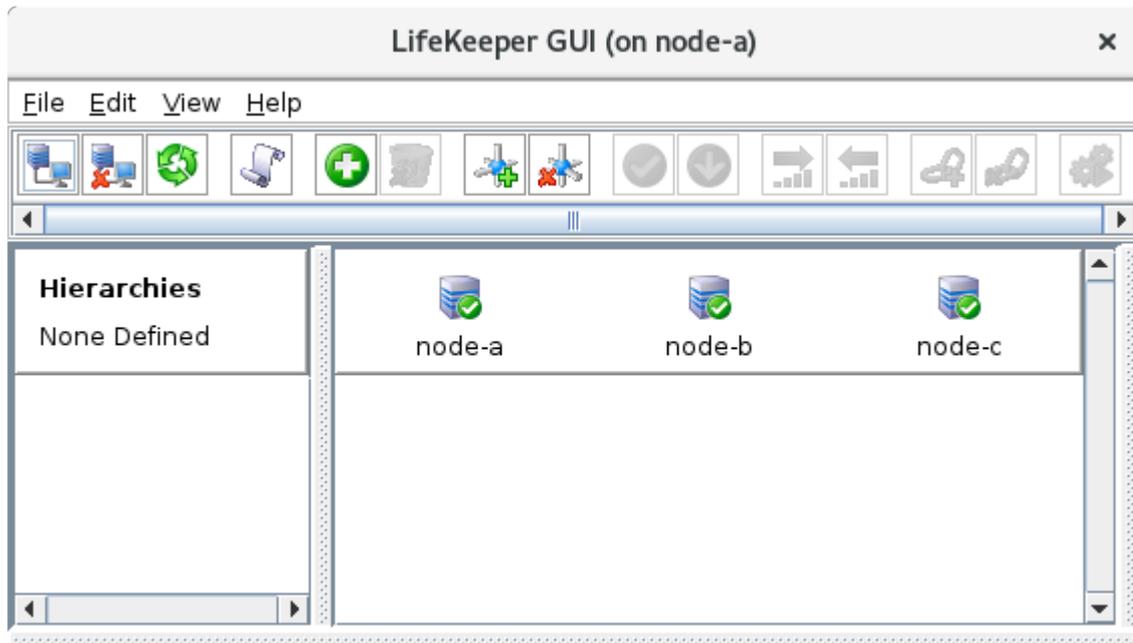
<-- node-a: Adding app/res: gen app

2. The login dialog is displayed again. Enter node-b and the root username and password to connect to the second node.



<-- node-a: Adding app/res: gen app

3. Connect to node-c the same way. Now you have three nodes displayed in the LifeKeeper GUI.



<-- node-c: Adding app/res: scsi device

## 11.2.6.5. Define Communication Paths

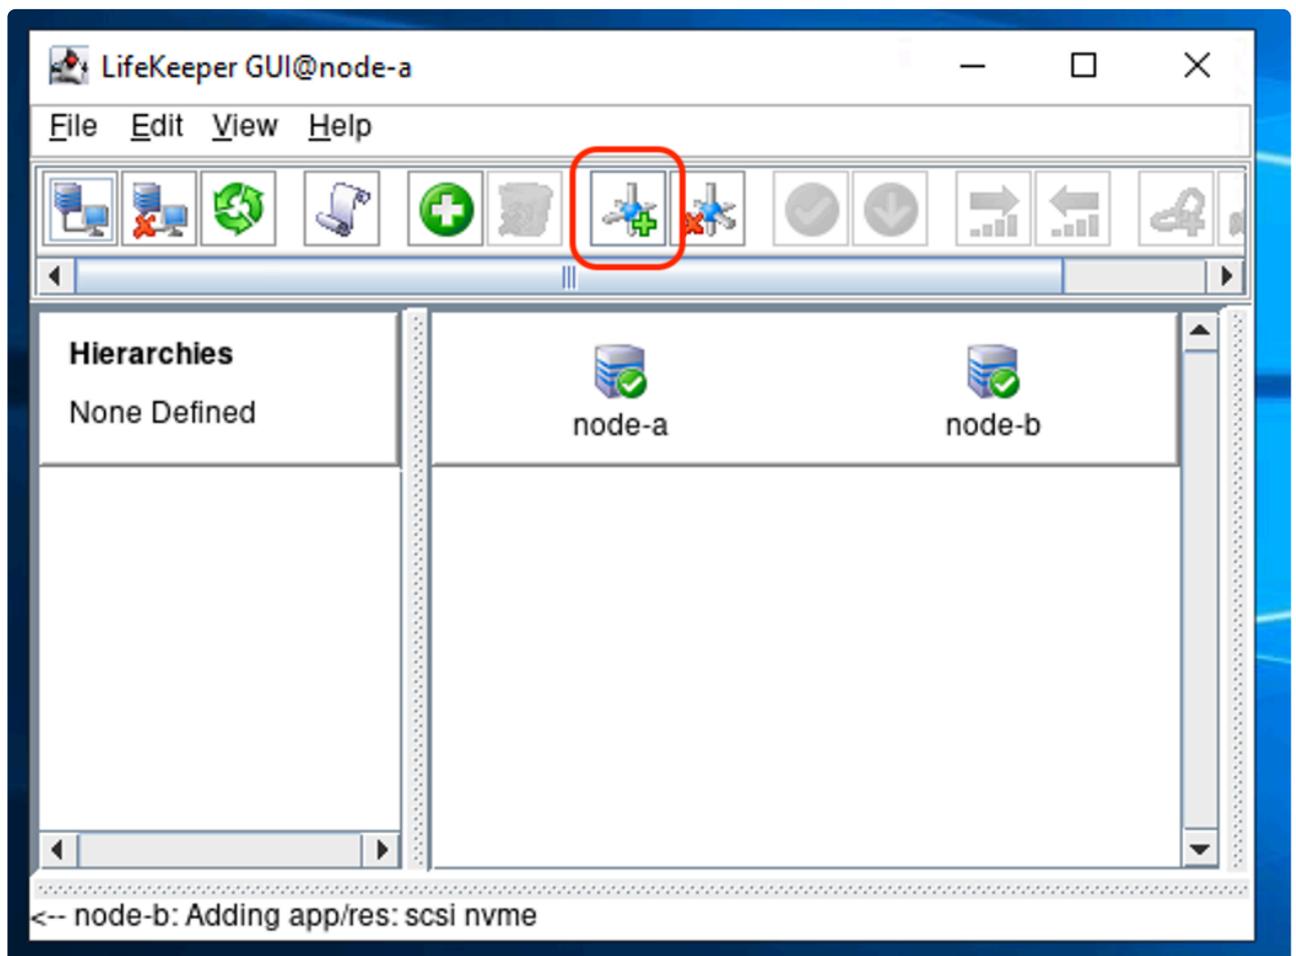
A Communication Path is a mechanism that allows the three nodes to communicate with each other and to transfer the data (for data replication purposes) between nodes. In a typical setup and especially in an on-premise environment, multiple communication paths are required to avoid a single point of failure. In this example, the network card in each node, or network routers switching between these nodes can be a point of failure.

In a cloud environment, however, having one communication path between nodes may be good enough as the underlying network layer is abstracted (we don't have visibility of these layers).

The next step is to define communication paths between all pairs of nodes.

**\* Important Note:** In a cluster consisting of three or more nodes, communication paths must be defined between **all pairs** of nodes. In a three-node cluster, this means that three communication paths must be defined (node-a ↔ node-b, node-a ↔ node-c, and node-b ↔ node-c).

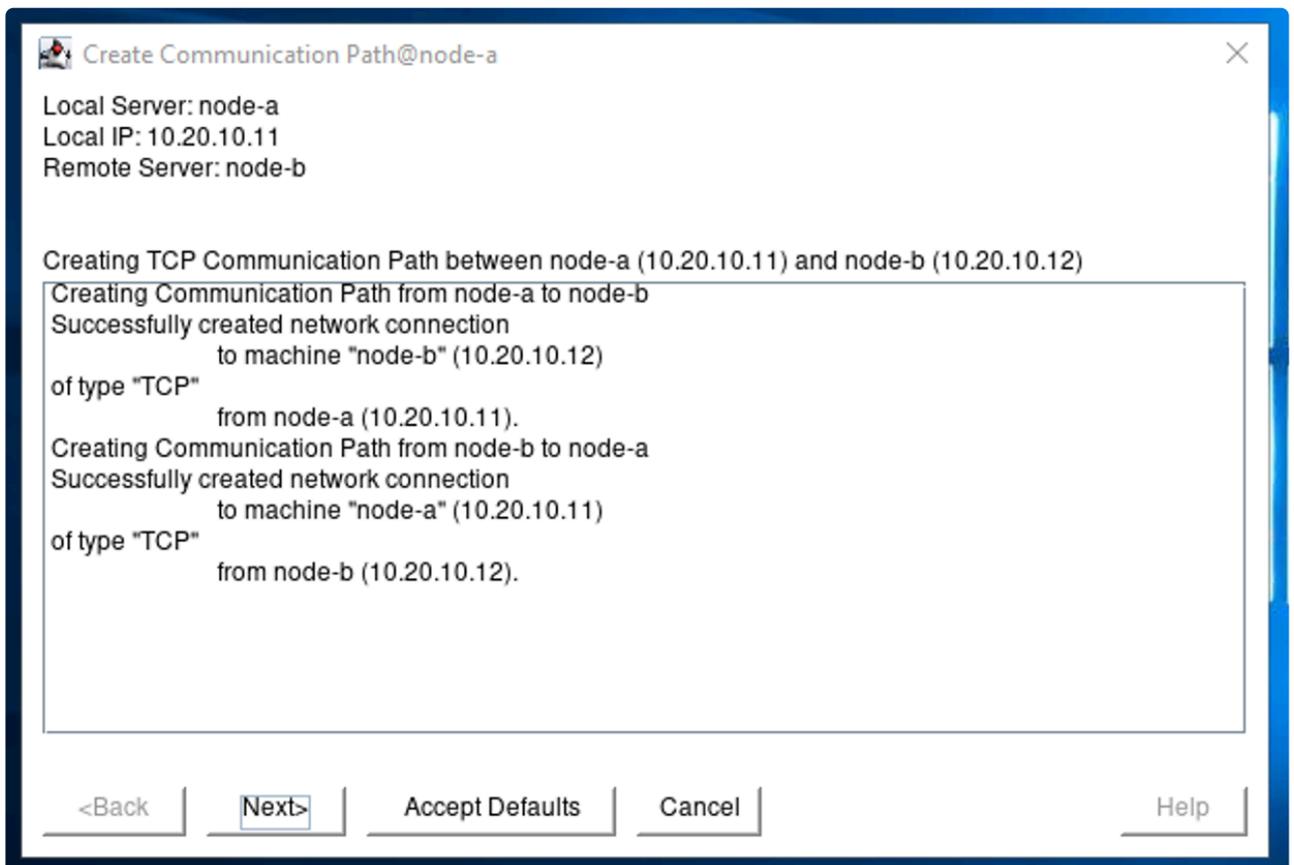
1. Click "Add communication path".



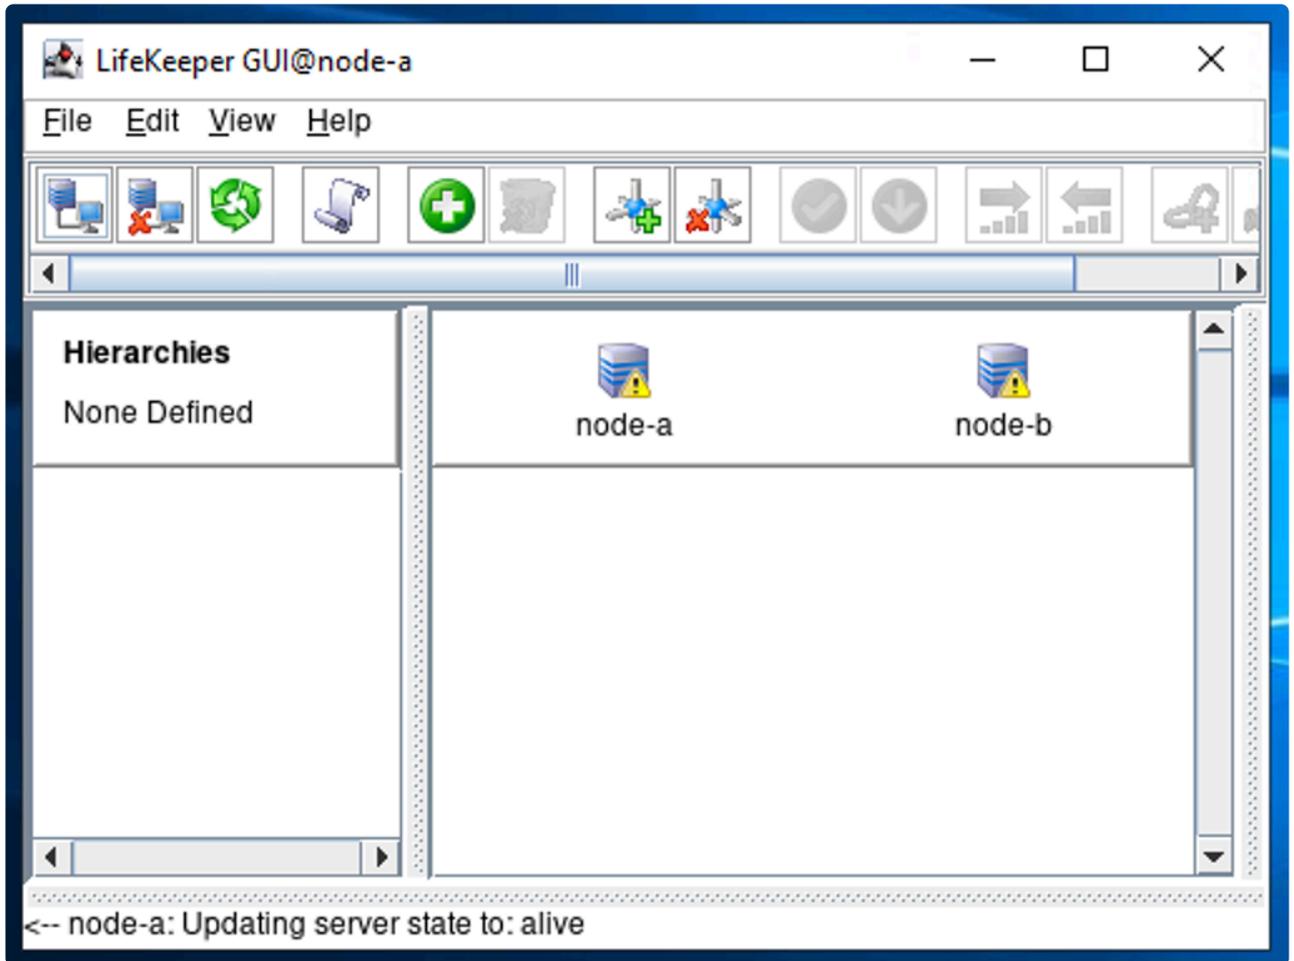
2. Select the following values. In most cases, these values should already be filled and selected.

Field	Value
Local Server	node-a
Remote Server(s)	node-b
Device Type	TCP
Local IP Address(es)	10.20.1.10
Remote IP Address	10.20.2.10
Priority	1

3. The summary of what has been configured will be displayed. Select "Next".



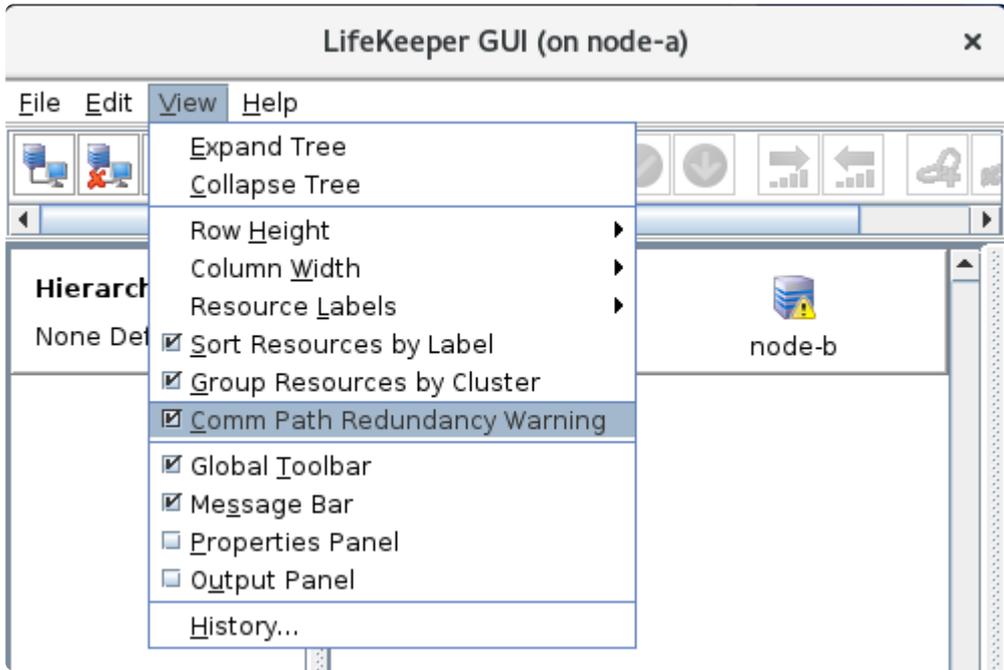
4. A communication path has now been defined.



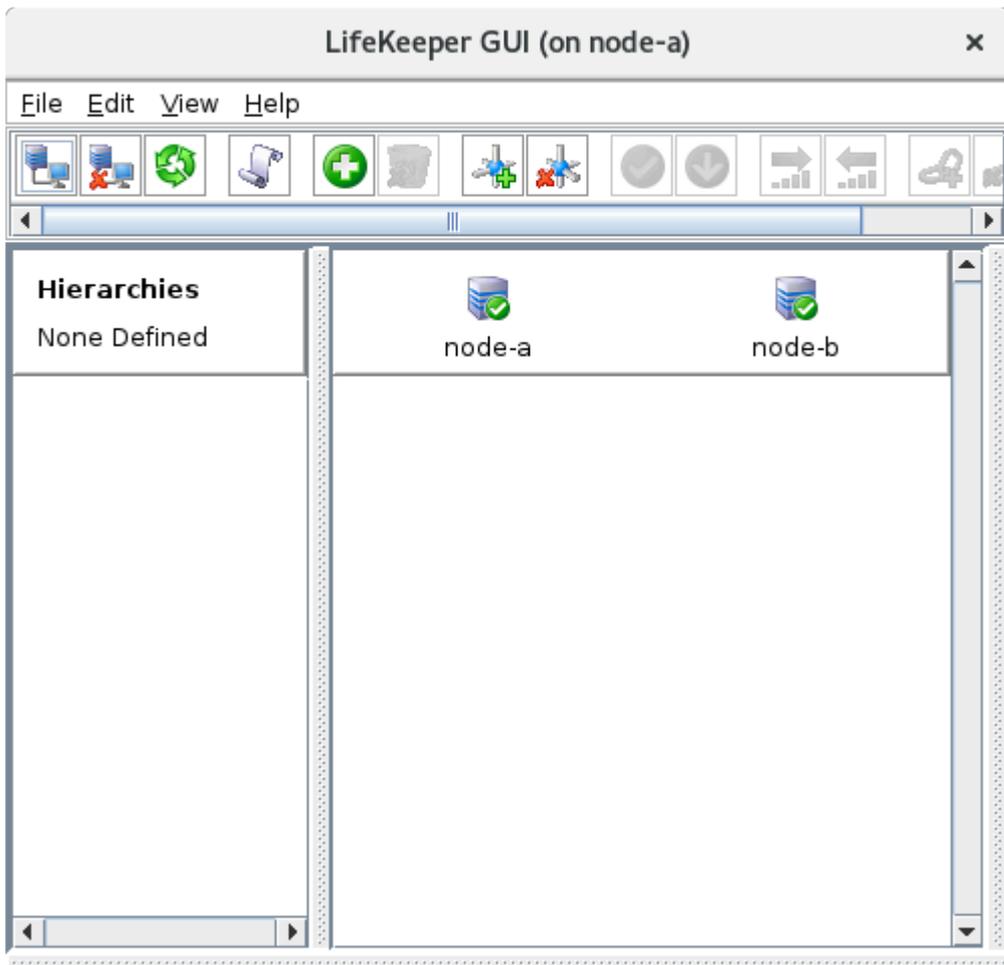
Once the communication path has been defined, you may notice a yellow triangle (  ) on each node. This is because we have only defined a single communication path. If a second communication path is defined, this changes to a green checkmark (  ).

## Turn off the Comm Path Redundancy Warning

This warning may be ignored if you are deploying in the cloud. The warning sign can be turned off by unchecking the "Comm Path Redundancy Warning" check box in the View menu of the LifeKeeper GUI.



Once unchecked, the icon will change from  to .



<-- node-b: Updating server state to: alive

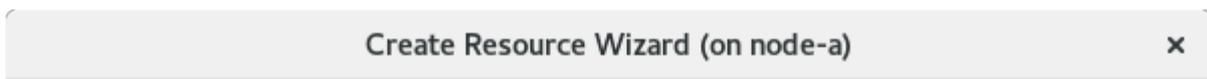
## 11.2.7. Protecting Our Resources

### A Typical Resource Protection Scenario with LifeKeeper

With LifeKeeper, most of the configuration can be completed using wizards. The wizard collects information about the resource(s) that are available for protection, and these resources can be protected by answering a series of simple questions. In most cases it is only necessary to select the default parameters. A standard example of LifeKeeper resource creation is given below.

1. Create Resource Wizard

- a. Select Resource Type (IP Resource, Data Replication, Oracle, etc.)



Please Select Recovery Kit



- b. Select Switchback Type

- c. Select the "active" server (where the application is currently running)

- d. Select the IP address to protect, configuration file location, and disk to protect

- e. Select the Resource Tag (name)

- f. Review the selections

2. Pre-Extend Wizard

- a. Select the target server (which will take over as the application host in the event of

switchover or failover)

b. Select Switchback Type

c. Select the Priority

d. Review the selections

### 3. Extend the Resource Hierarchy

a. Select the IP address of the target, configuration file location, and disk to replicate data to

b. Select the Resource Tag (name)

c. Review the selections

### 4. Hierarchy Integrity Verification

Once the configuration is completed, the wizard also automatically defines the dependencies between resources (for example, to protect a database resource it is necessary to mount the underlying disk(s) before starting the database).

## Switching Between Nodes

We are now ready to protect our resources.

As discussed earlier, cloud environments work differently than on-premise environments. If you are using a cloud environment, please read through [Switching between Nodes in a Cloud Environment](#) for relevant information about routing traffic to the active node in a cloud environment. We will configure the necessary resources after creating the IP resource, and the parameters selected for the IP resource in the next section will vary based on the cloud environment used.

If working in an on-premise environment, simply going through the following page is sufficient.

- [Creating an IP Resource](#)

If working in a cloud environment, select the relevant scenario from the following page:

- [Switching between Nodes in a Cloud Environment](#)

Once the IP resource is protected, it is recommended that you initiate a switchover (where the “standby” node becomes the “active” node) to test the functionality.

- [Switch to Standby Node to Confirm Switchover is Working](#)

## Setup Disk Replication

In most scenarios, data replication should be defined between nodes (especially when evaluating in the cloud). Follow the steps below to define data replication between nodes.

- [How does Data Replication between Nodes Work?](#)
- [How to Create Data Replication of a File System](#)

## Protect our Resources

Once the basic preparation is completed, you can now protect resources such as databases and applications.

- [Protecting an Oracle Resource](#)
- [Protecting MSSQL Using Quick Service Protection](#)
- [Protecting a PostgreSQL Resource](#)
- [Protecting an NFS Resource](#)
- [Protecting SAP Resources](#)
- [Protecting SAP HANA Resources](#)

# 11.2.7.1. Creating an IP Resource

The first resource to be created on LifeKeeper should be an IP resource. As discussed [earlier](#), a client can find the Active Node using a Virtual IP Address. Below is the list of IP addresses that are referenced in this section.

## IP Addresses for Nodes

Field	Value
node-a	10.20.1.10
node-b	10.20.2.10

## IP Address for the IP Resource Based on the Scenario

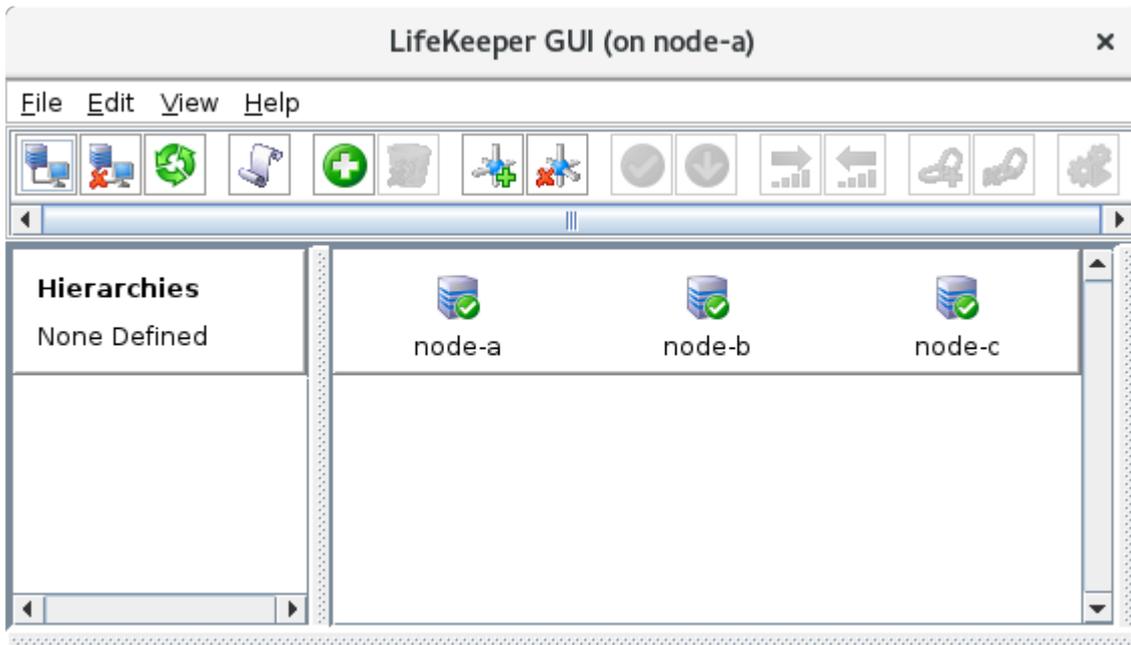
In a cloud environment, please review [Switching between Nodes in a Cloud Environment](#) to identify the most applicable scenario before selecting an IP address.

Platform	Scenario	How to Select an IP Address	Sample IP Addresses Used in this Guide	Default Name of the IP Address
AWS	Route Table Scenario	IP address that is outside of the VPC CIDR	10.10.10.10	ip-10.10.10.10
	Route 53 Scenario	IP address that is assigned to the “active” node	10.20.1.10	realip
	Elastic IP Scenario	 Not required to create an IP Resource. Skip this page.		
Google Cloud				
Azure	On-premise environment	An available IP address on the subnet	10.20.0.10	ip-10.20.0.10

## Configure an IP Resource

This example uses 10.10.10.10 as the IP address to be used for “Route Table Scenario” in AWS.

1. Click the  button in the LifeKeeper User Interface (or select Edit > Server > Create Resource Hierarchy).



<-- node-c: Adding app/res: scsi device

- 2. The Create Resource Wizard (on node-a) will be displayed and the type of resource to be protected can be selected. Select "IP" and then click "Next>".



Please Select Recovery Kit

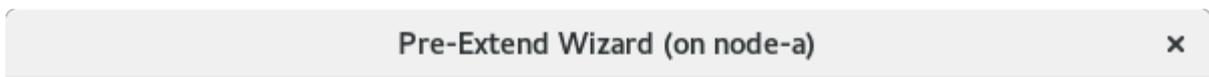
- 3. Select the default values. The only field you need to enter is the virtual IP address 10.10.10.10 in the IP Resource field.

 **Note:** For items with a checkmark (  ), review and use the default values.

Field	Value
Switchback Type	Intelligent ✓
Server	node-a ✓
IP Resource	10.10.10.10
Net Mask	255.255.255.0 ✓ Note: This value may vary based on the network.
Network Interface	eth0 ✓
IP Resource Tag	ip-10.10.10.10 ✓

Values such as Net Mask are displayed based on how the local system is configured. Therefore the default values seen may vary.

- Once you confirm the resource is successfully created, click “Next>”.
- The next step is the “Pre-Extend Wizard (on node-a)”.



Target Server

You have successfully created the resource hierarchy ip-10.10.10.10 on node-a. Select a target server to which the hierarchy will be extended.

If you cancel before extending ip-10.10.10.10 to at least one other server, LifeKeeper will provide no protection for the applications in the hierarchy.

Select all of the defaults as shown below.

Field	Value
-------	-------

Target Server	node-b
Switchback Type	Intelligent
Template Priority	1
Target Priority	10

The Pre-Extend check should be successfully completed. Click “Next”.

- The next step is “Extend comm/IP Resource Hierarchy (on node-a)”. Click “Next” and fill in the fields.

**Extend comm/ip Resource Hierarchy (on node-a)** ✕

Template Server: node-a  
 Tag to Extend: ip-10.10.10.10  
 Target Server: node-b

IP Resource

The IP address or symbolic name to be protected by the IP resource on the target server. The same value that was used on the template server is used for the IP resource on the target server. Therefore, this value cannot be changed. The IP resource is used by client applications to login into the parent application over a specific network interface. If a symbolic name is used, it must exist in the local /etc/hosts file or be accessible via a Domain Name Server (DNS). Any valid hosts file entry, including aliases, is acceptable. If the address cannot be determined or if it is found to be already in use, it will be rejected. If a symbolic name is given, it is used for translation to an IP address and is not retained by LifeKeeper. Both IPv4 and IPv6 style addresses are supported.

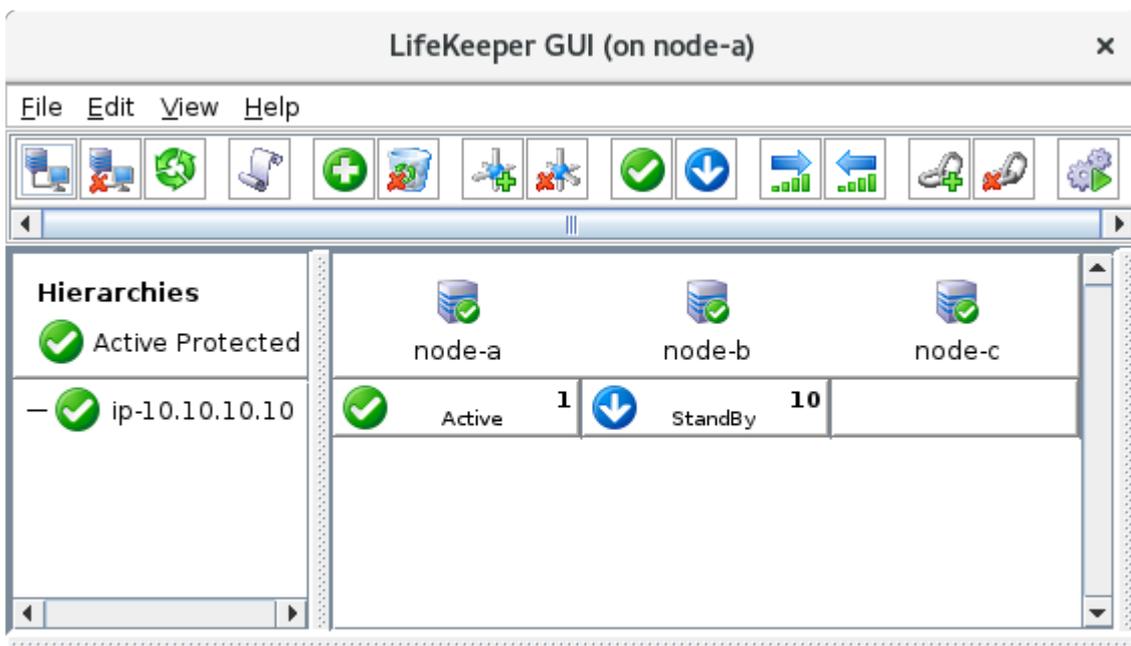
<Back
Next>
Accept Defaults
Cancel
Help

Field	Value
Net Mask	255.255.255.0 Note: This value may vary based on the network.
Network Interface	eth0
IP Resource Tag	ip-10.10.10.10

- Now the resource is extended. Select “Finish”.



- 8. The verification is now complete and the wizard can be closed.
- 9. The IP Resource (ip-10.10.10.10) will be displayed on the LifeKeeper User Interface.



<-- node-b: ip-10.10.10.10: Updating equivalency list

As we have seen, resources may be easily protected by answering the questions in the resource creation wizard. Most of the values are prepopulated based on an understanding of the environment and therefore minimal user interaction is required.

## 11.2.7.2. Switching between Nodes in a Cloud Environment

---

If you are operating in a cloud environment, there is an extra step to switch a virtual IP address between nodes. The following topics explain the background information and general approaches for each cloud environment:

- [AWS Route Table Scenario](#)
- [AWS Elastic IP Scenario](#)
- [AWS Route53 Scenario](#)
- [Azure Internal Load Balancer Scenario](#)
- [Google Cloud Internal Load Balancer Scenario](#)

The following sections explain the configuration steps required in each scenario:

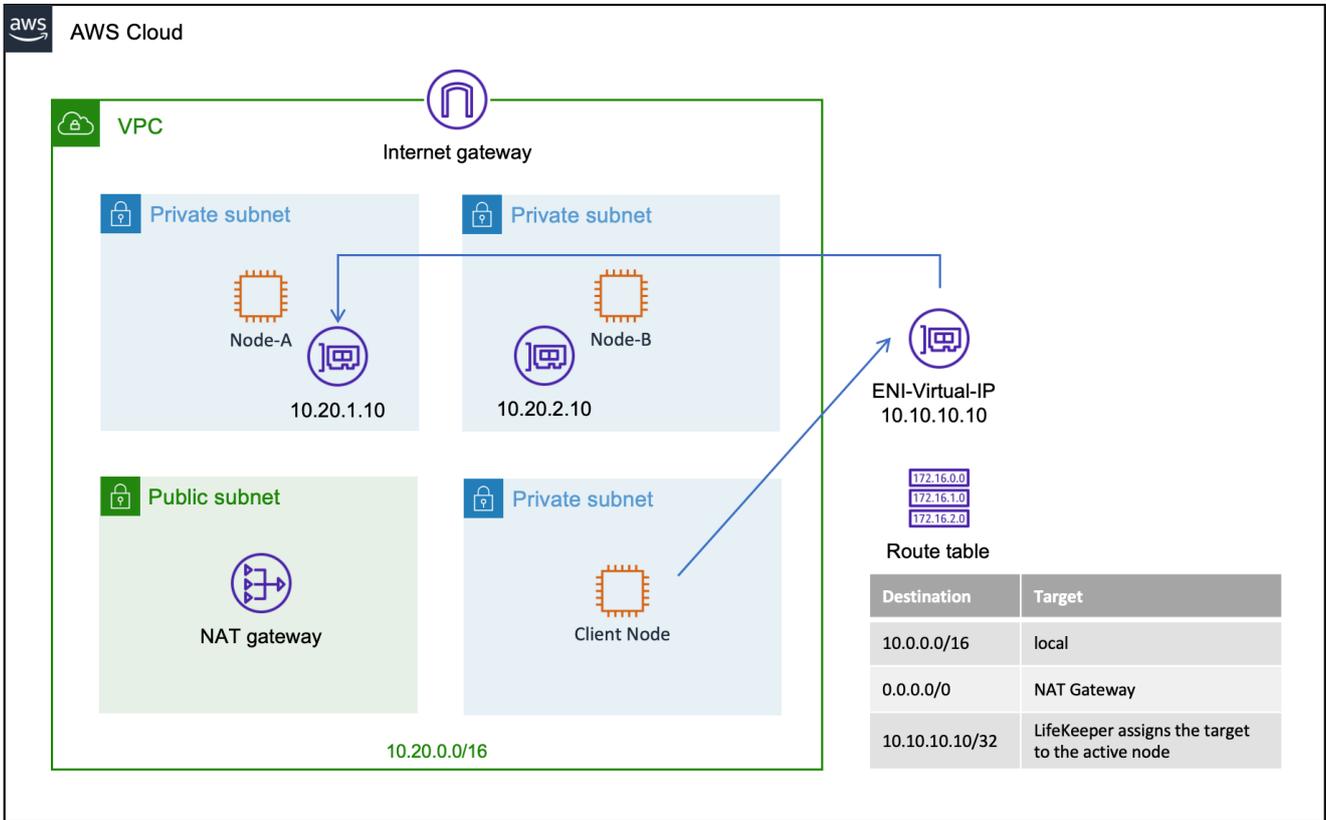
- [Creating an AWS EC2 Resource \(RouteTable Scenario\)](#)
- [Creating an AWS EC2 Resource \(Elastic IP Scenario\)](#)
- [Creating an AWS Route53 Resource](#)
- [Azure – Using an Internal Load Balancer](#)
- [Google Cloud – Using an Internal Load Balancer](#)

LifeKeeper also provides a resource type which will respond to TCP health check probes from an internal load balancer in order to route traffic to the current application host:

- [Responding to Load Balancer Health Checks](#)

# 11.2.7.2.1. Creating an AWS EC2 Resource (RouteTable Scenario)

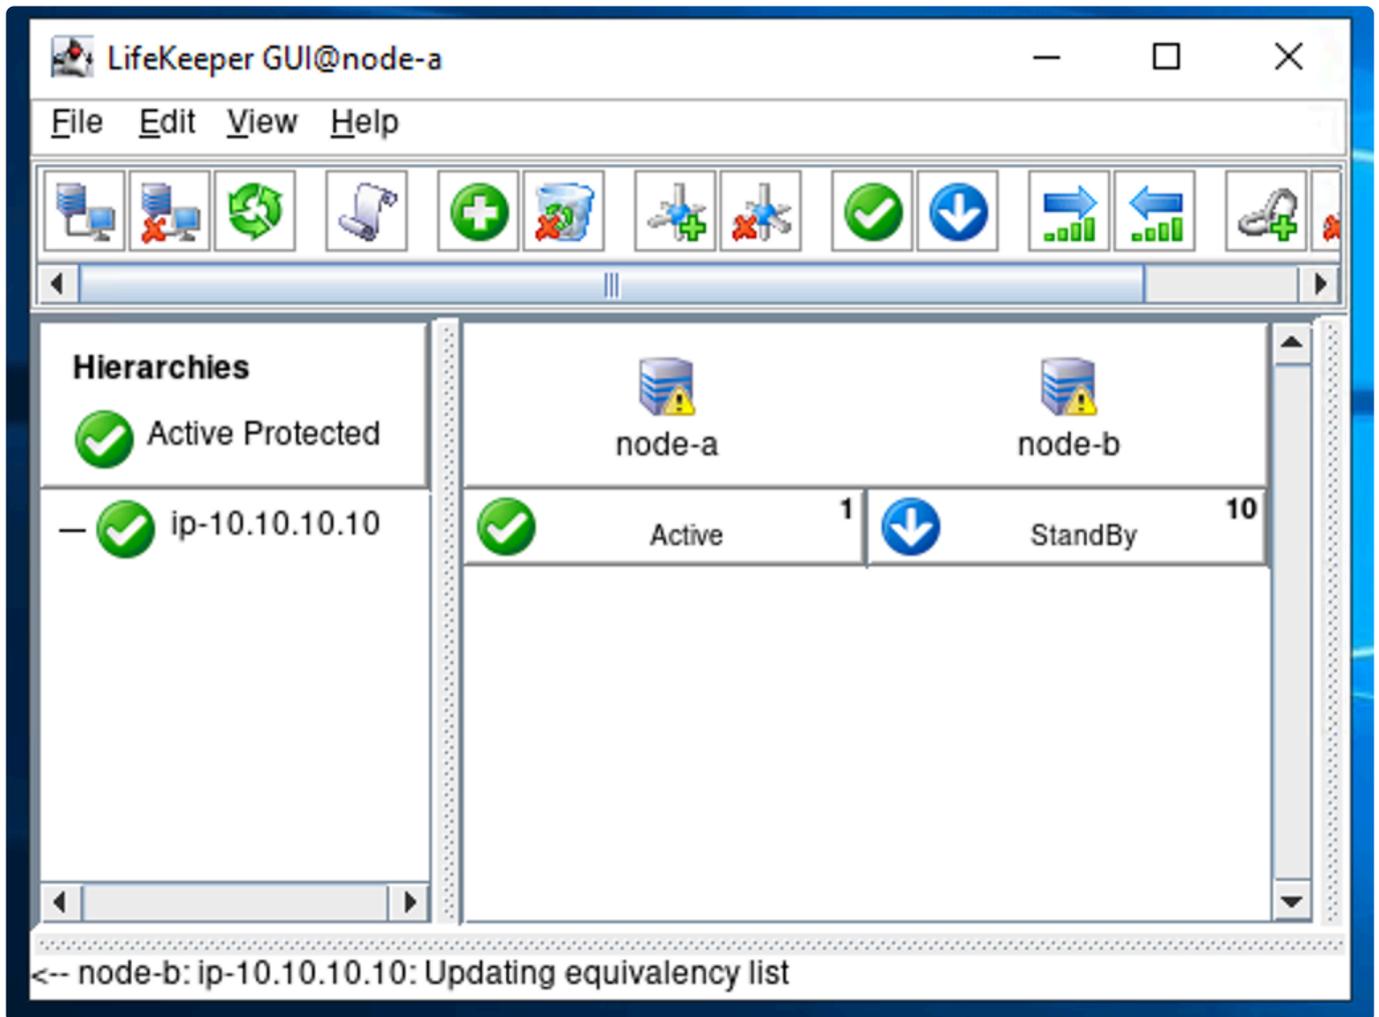
Here we will discuss how the Route Table can be used to switch between nodes by creating an entry in the Route Table that points to the active node.



## Create an IP Resource

Now you can create an IP resource. Note that the IP address of the resource must be located outside of the VPC. In this example, the IP resource is created using 10.10.10.10 (the VPC is defined as 10.20.0.0/16). Refer to [Creating an IP Resource](#) for more information.

Once the IP resource has been created, the LifeKeeper User Interface should look like this.



## Assign Permission to Update Route Table Entries to Instances

! To complete the steps below, the AWS CLI must be installed on each node.

- [Install AWS CLI](#)
- [Assign Permission to Use EC2 Recovery Kit](#)

## Add Initial Route Table Entry for the Primary Host

Before creating the EC2 resource, we must first add the initial route to the route table for LK-VPC (i.e., LK-RouteTable).

1. In the AWS Console, navigate to **Services** → **VPC** → **Route Tables**, click the route table associated with LK-VPC (LK-RouteTable), click the **Routes** tab, then click **Edit Routes**.
2. In the resulting dialog, add a new route with the following parameters:

Field	Value
Destination	10.10.10.10/32
Target	Select 'Instance', then select the instance ID for Node-A

### Edit routes

Destination	Target	Status	Propagated
10.20.0.0/16	Q local	Active	No
Q 0.0.0.0/0	Q igw-002cd16f5bf492263	Active	No <span>Remove</span>
Q 10.10.10.10/32	Q i-	-	No <span>Remove</span>

Add route

Q i-  
 i-000c007d6a62841d5 (Node-A)  
 i-08868c97b794b854b (Node-B)  
 i-0e320b83da7a002e7 (Node-C)

Cancel Preview Save changes

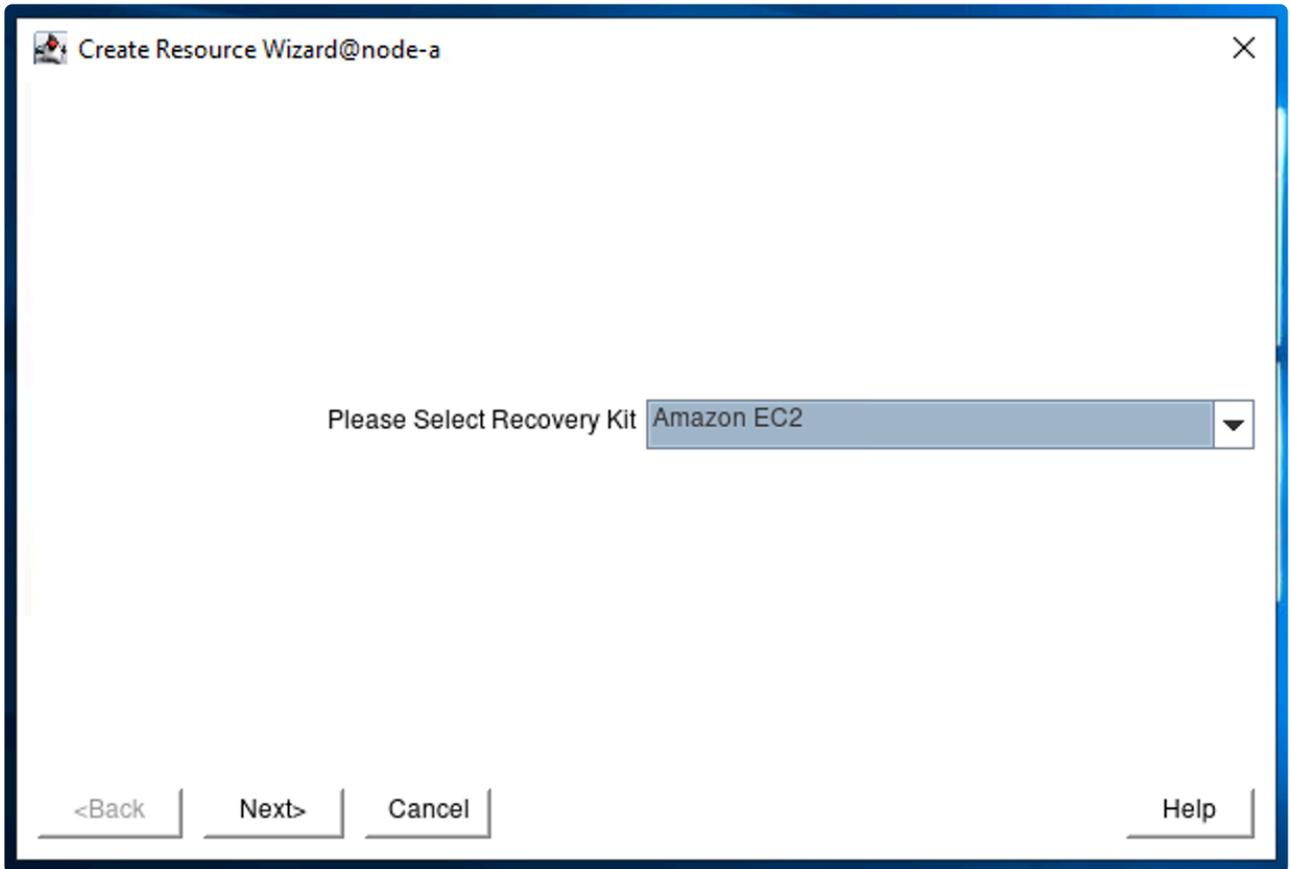
3. Click **Save changes** to add the route to the route table.

## Creating an EC2 Resource (Route Table Scenario)

! Before creating the EC2 resource, ensure that source/destination checking has been disabled for all cluster nodes by following the steps given in [AWS – Disable Source/Destination Checking](#).

Creating an EC2 Resource should be straightforward once an IP Resource has been created.

1. Select the + icon to start the “Create Resource Wizard”.
2. Select “Amazon EC2” as the Recovery Kit.



3. On the “Create Resource Wizard @ node-a”, specify the following values.

 **Note:** For items with a checkmark (  ) review the default value and use the value suggested.

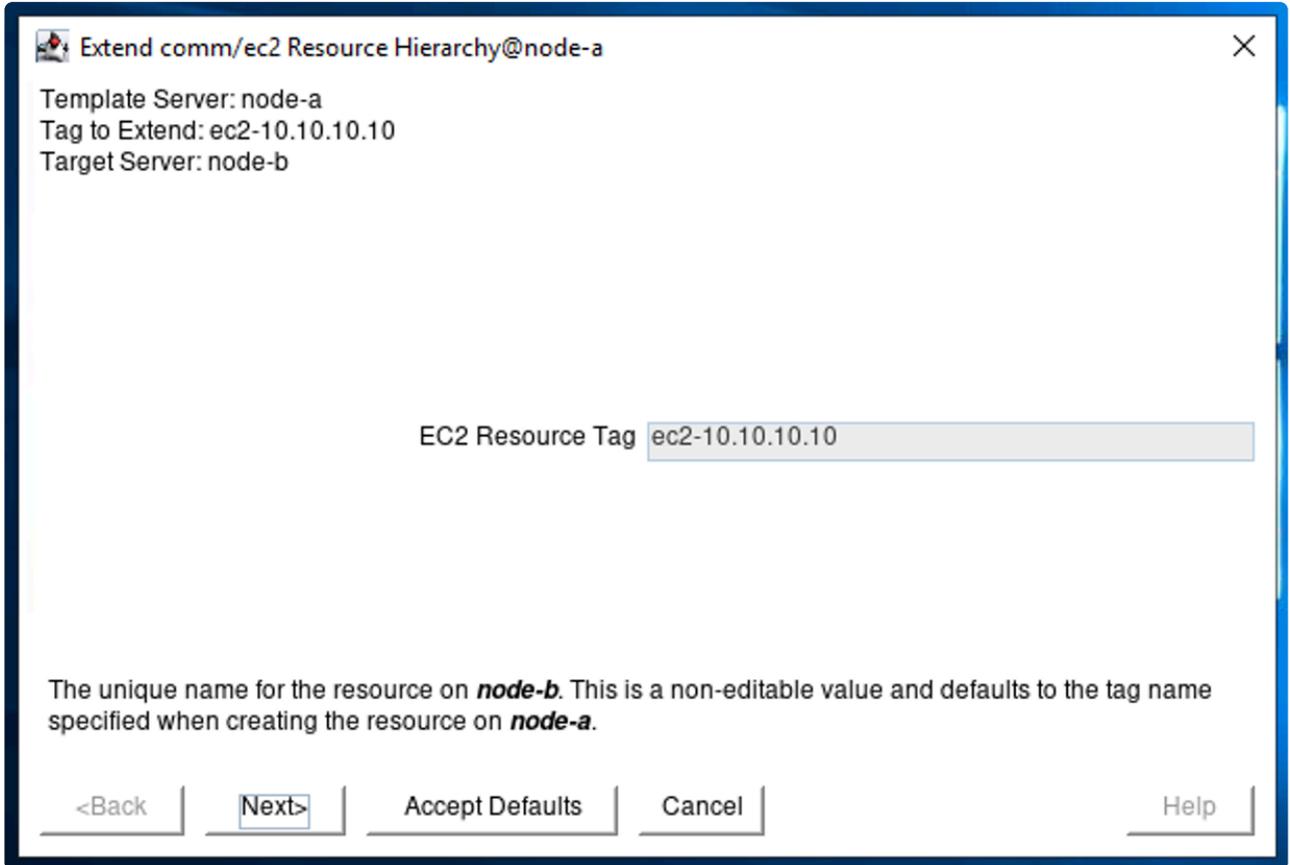
Field	Value
Switchback Type	Intelligent 
Server	node-a
EC2 Resource Type	Route Table (Backend Cluster)
IP Resource	ip-10.10.10.10 
EC2 Resource Tag	ec2-10.10.10.10 

4. On the “Pre-Extend Wizard” specify the following values:

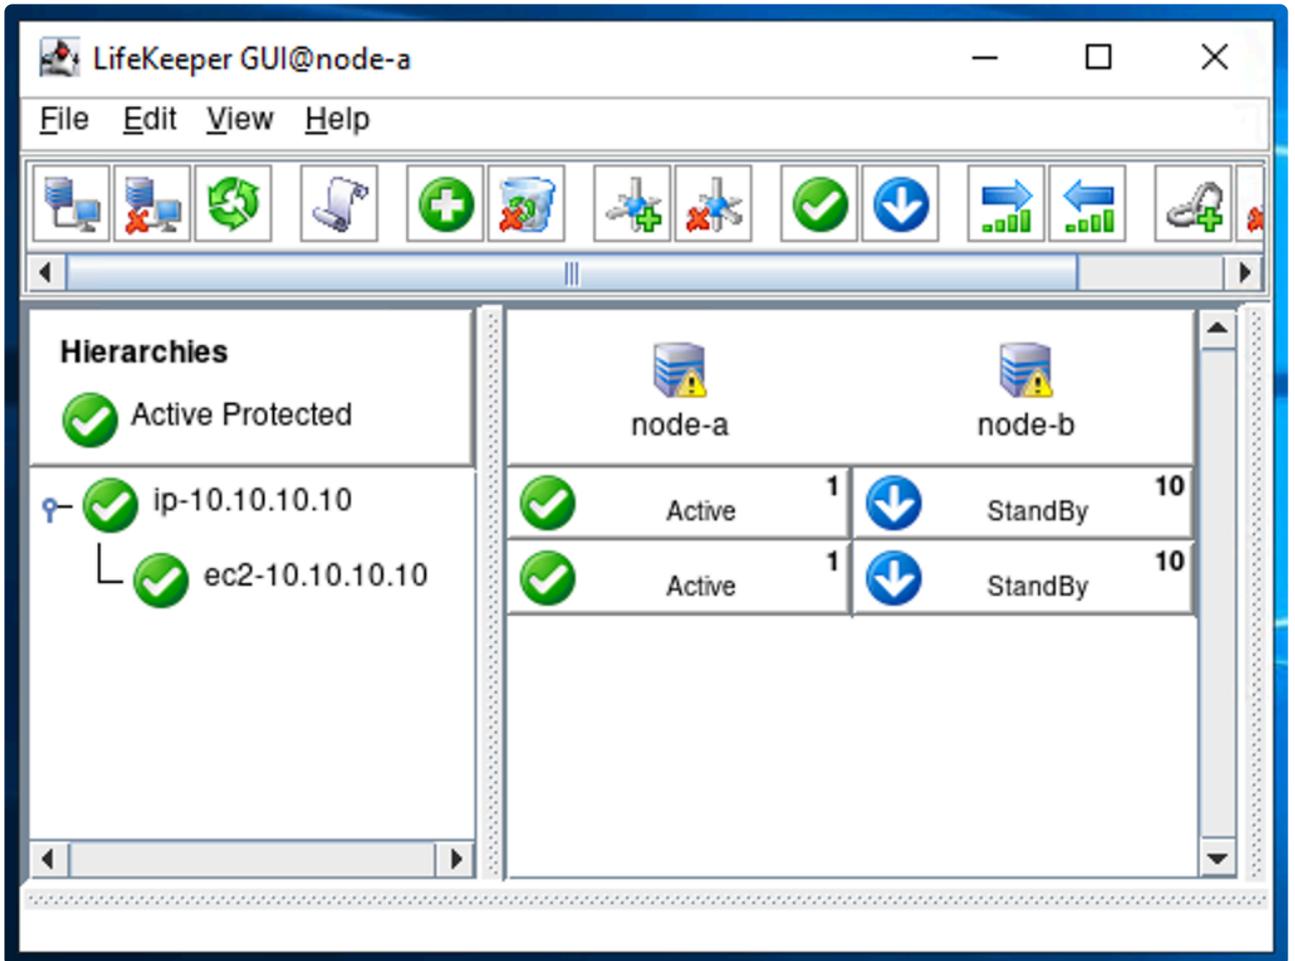
Field	Value
Target Server	node-b
Switchback Type	Intelligent 

Template Priority	1 
Target Priority	10 

- Once Pre-Extend is completed, move on to “Extend comm/ec2 Resource”. Select “ec2-10.10.10.10”.

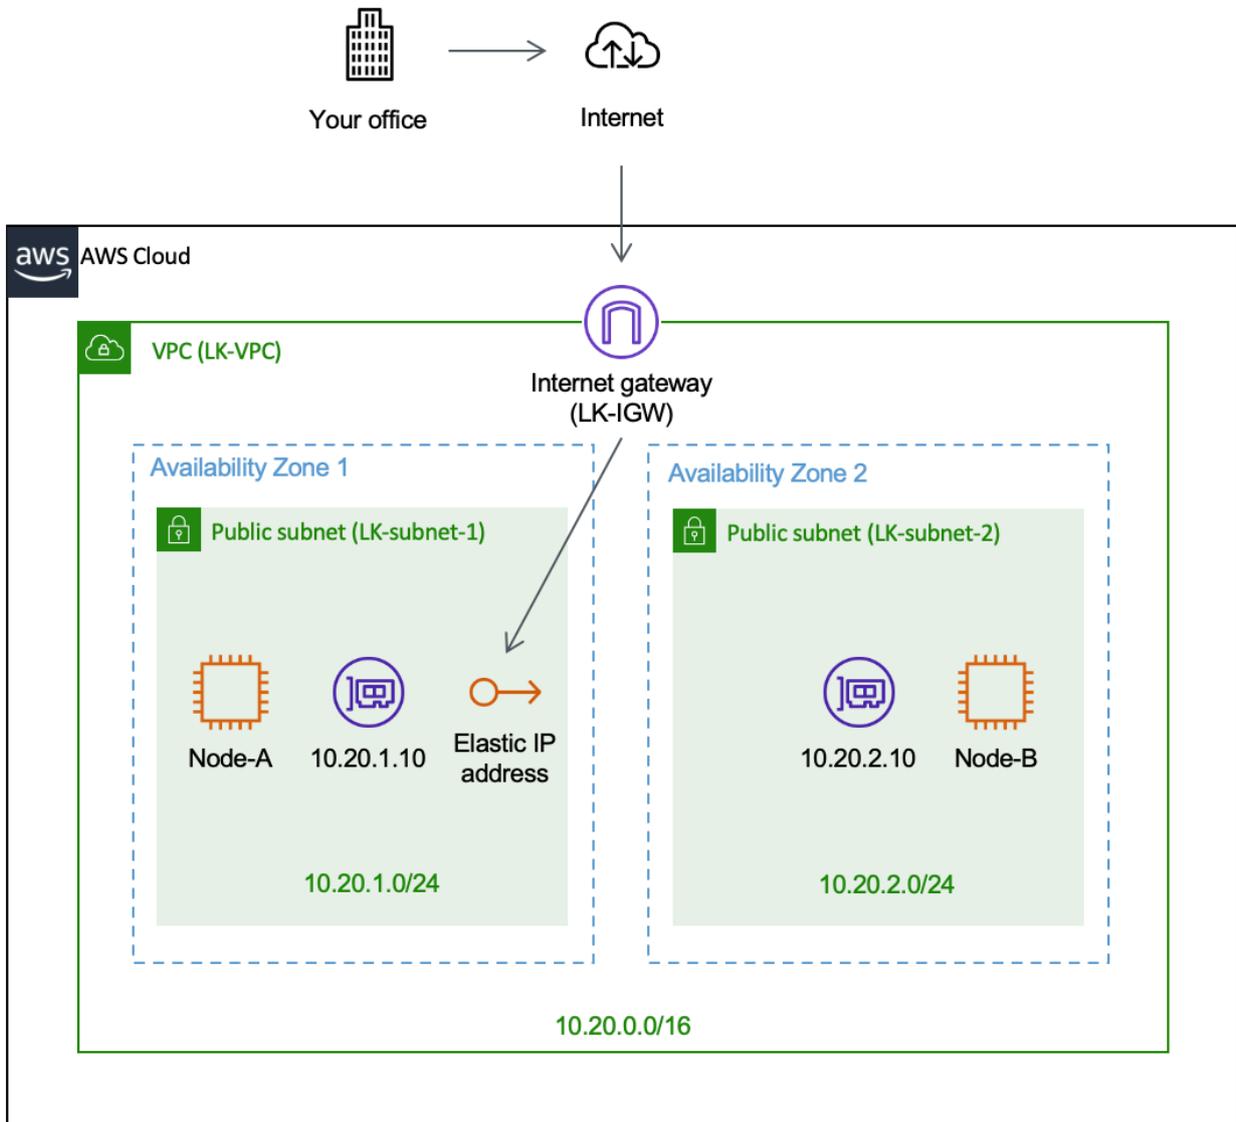


- Once “Hierarchy successfully extended” is displayed, the creation of the resource is now complete.
- The EC2 resource is created as shown below.



# 11.2.7.2.2. Creating an AWS EC2 Resource (Elastic IP Scenario)

Here we will discuss how to switch between nodes using an Elastic IP address by creating an entry that points to the active node.



## Assign Permission to Update Route Table Entries to Instances

! To complete the steps below, the AWS CLI must be installed on each node.

- [Install AWS CLI](#)

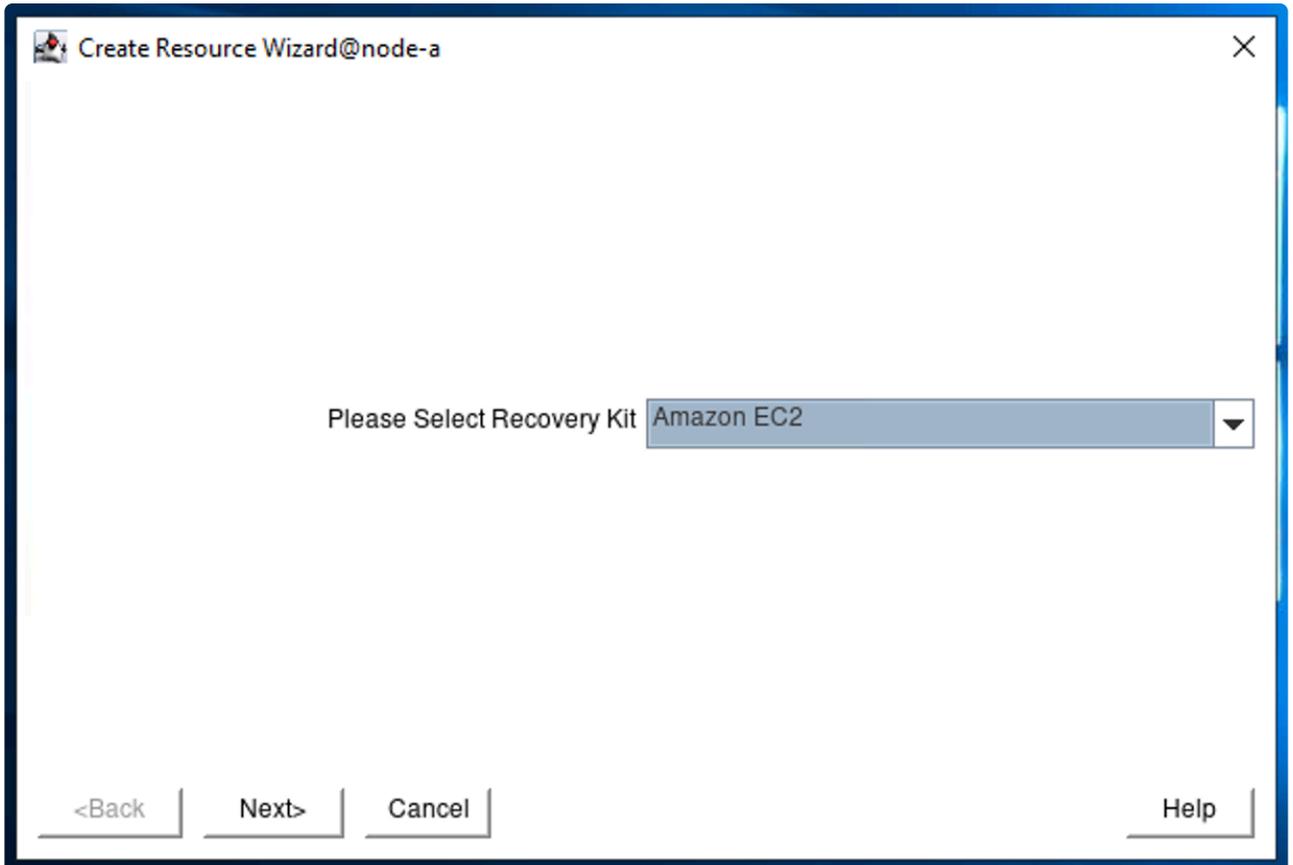
- [Assign Permission to Use EC2 Recovery Kit](#)

You will also need to acquire an Elastic IP. In the example below the Elastic IP address is 50.18.115.213. Note that this IP address is an example value.

## Creating an EC2 Resource (Elastic IP Scenario)

Complete the following steps to create an EC2 resource using the Elastic IP (Frontend Cluster) scenario.

1. Select “Amazon EC2” as the Recovery Kit.



2. On the “Create Resource Wizard @ node-a”, specify the following values.

 **Note:** For items with a checkmark (  ) review the default value and use the value suggested.

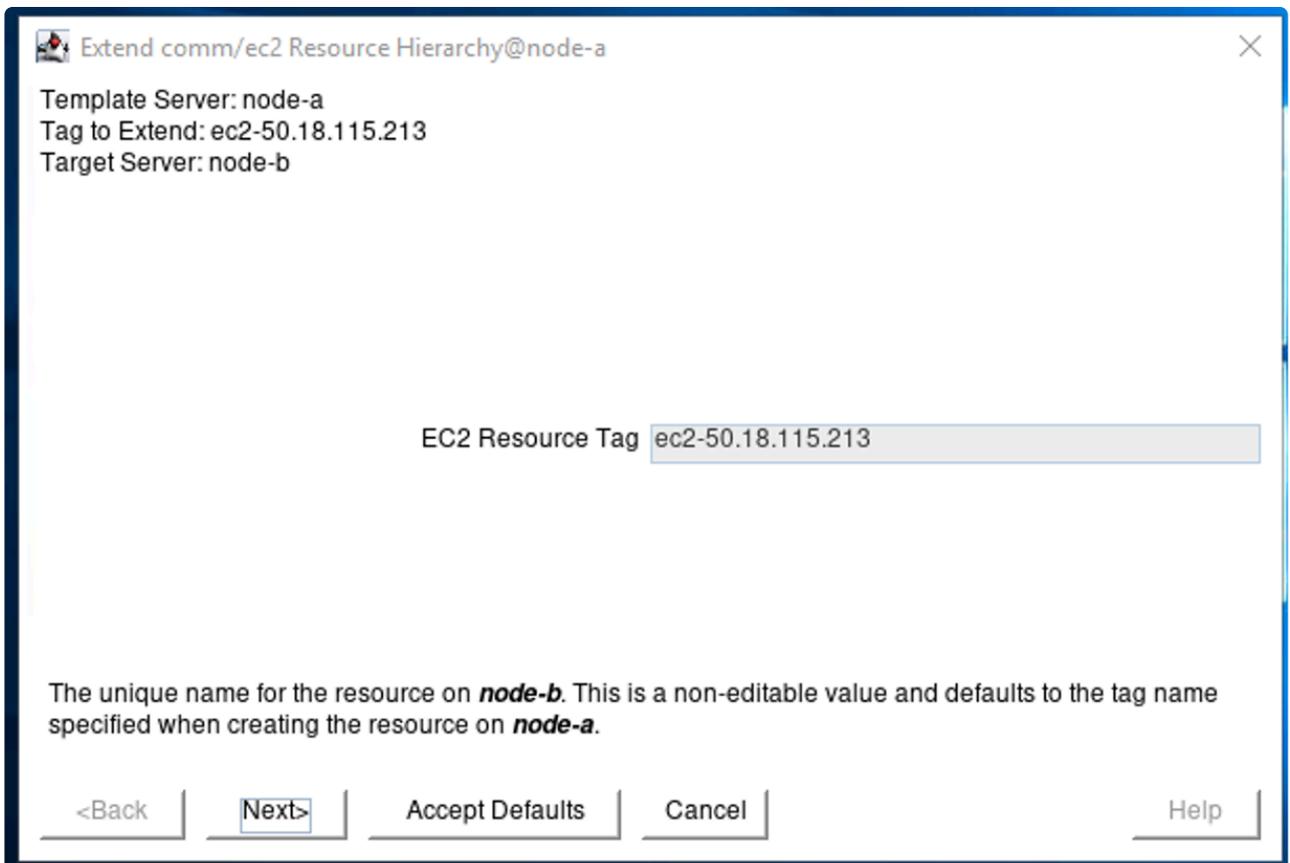
Field	Value
Switchback Type	Intelligent 
Server	node-a 
EC2 Resource Type	Elastic IP (Frontend Cluster)

Network Interface	eith0 
Elastic IP	Select an available Elastic IP from the list. For example, <input type="text" value="50.18.115.213"/>
EC2 Resource Tag	ec2-50.18.115.213 

3. On “Pre-Extend Wizard” specify the following values:

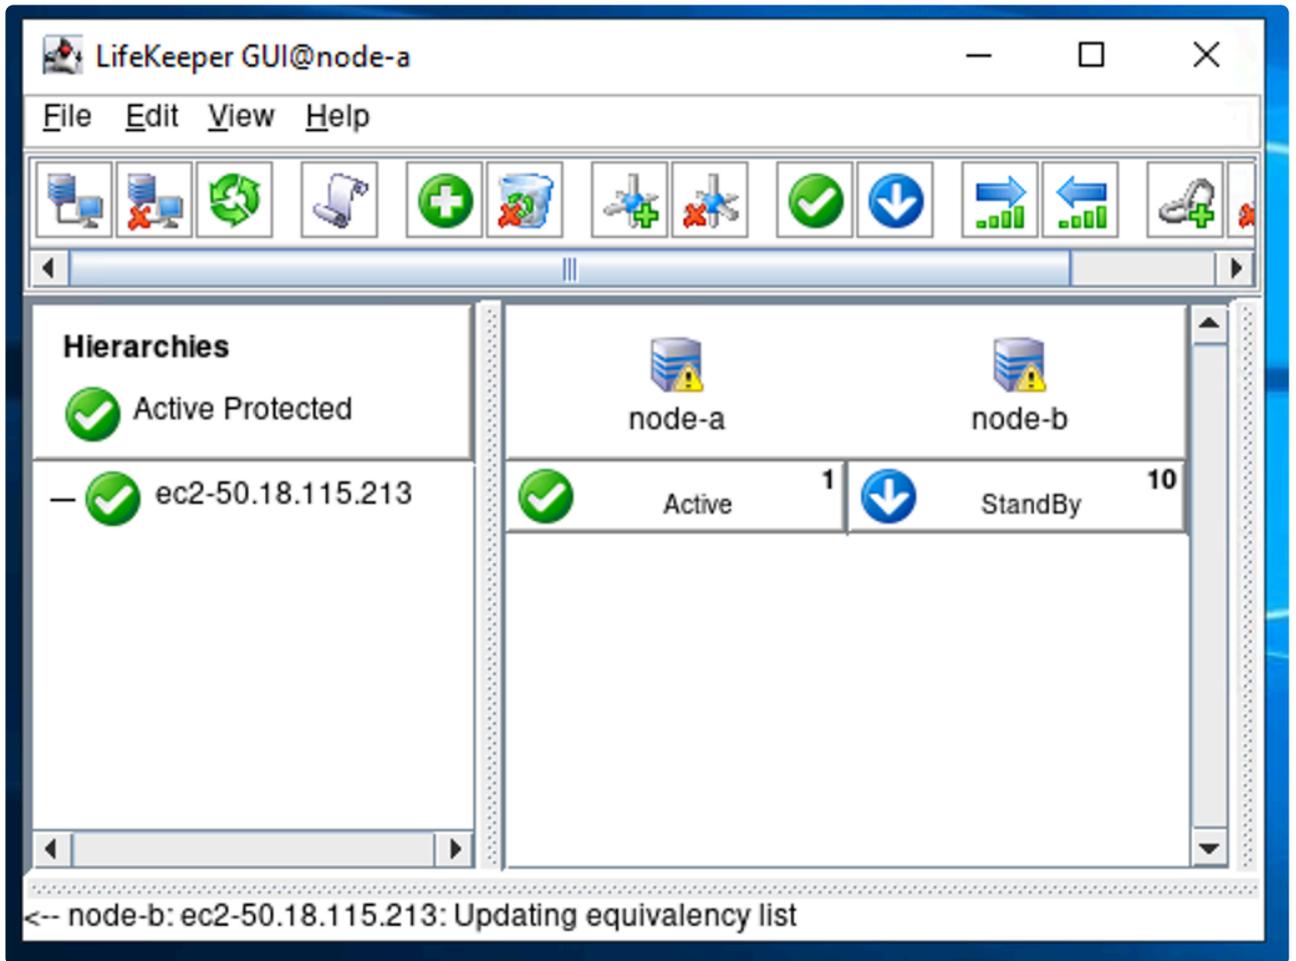
Field	Value
Target Server	node-b 
Switchback Type	Intelligent 
Template Priority	1 
Target Priority	10 

4. Once Pre-Extend is completed, move on to “Extend comm/ec2 Resource”. Select “ec2-50.18.115.213”.



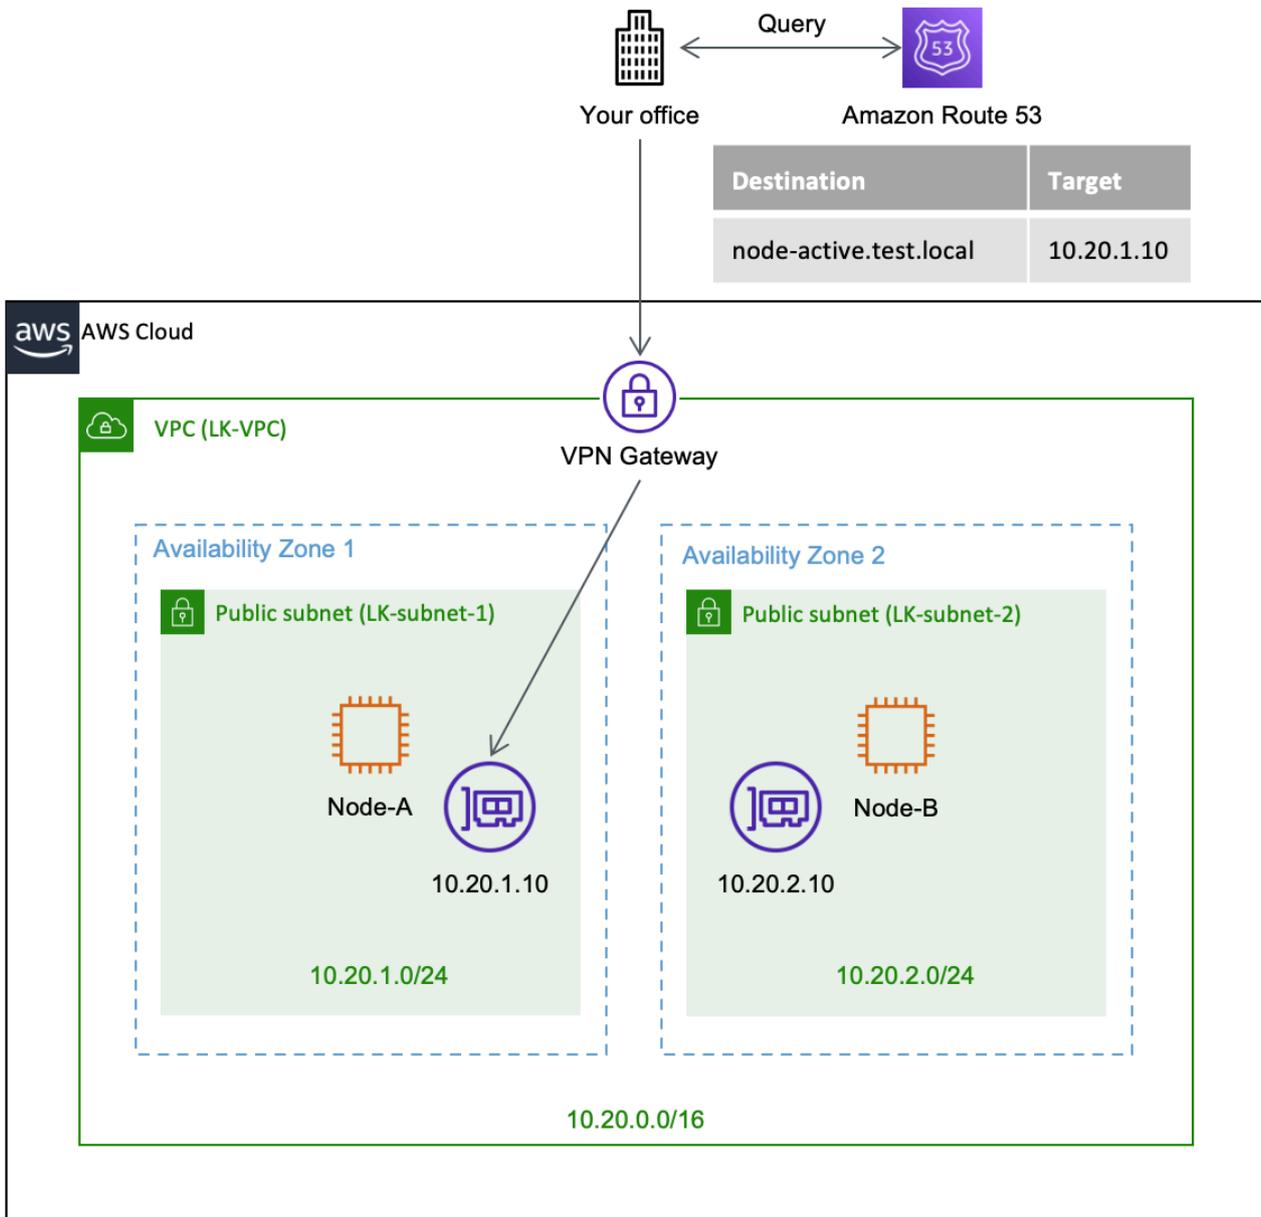
5. Once “Hierarchy successfully extended” is displayed, the creation of the resource is complete.

6. The EC2 resource is now created as shown below.



# 11.2.7.2.3. Creating an AWS Route53 Resource

Here we will discuss how to switch between nodes using Route53 by creating a DNS entry that points to the active node.



## Record Set on Route 53 to be Created

Name	Type	Value
node-a	A-IPv4 Address	10.20.1.10
node-b	A-IPv4 Address	10.20.2.10

node-active	A-IPv4 Address	10.20.1.10 (same as node-a)
-------------	----------------	-----------------------------

## Create a Hosted Zone

First, you need a Hosted Zone (this should look like a DNS domain such as example.com).

 If you already have a Hosted Zone, skip this section and use your Hosted Zone.

1. On the AWS management console, go to Route53.
2. Select [Hosted Zone](#).
3. Select "Create Hosted Zone".
4. Enter the following values:

Field	Value
Domain	test.local
Type	Private hosted zone
VPC Association	Select region and name of VPC (e.g., LK-VPC)

[Route 53](#) > [Hosted zones](#) > Create hosted zone

## Create hosted zone [Info](#)

### Hosted zone configuration

A hosted zone is a container that holds information about how you want to route traffic for a domain, such as example.com, and its subdomains.

**Domain name** [Info](#)  
This is the name of the domain that you want to route traffic for.

Valid characters: a-z, 0-9, ! " # \$ % & ' ( ) \* + , - / : ; < = > ? @ [ \ ] ^ \_ ` { | } . ~

**Description - optional** [Info](#)  
This value lets you distinguish hosted zones that have the same name.

The description can have up to 256 characters. 37/256

**Type** [Info](#)  
The type indicates whether you want to route traffic on the internet or in an Amazon VPC.

**Public hosted zone**  
A public hosted zone determines how traffic is routed on the internet.

**Private hosted zone**  
A private hosted zone determines how traffic is routed within an Amazon VPC.

### VPCs to associate with the hosted zone [Info](#)

To use this hosted zone to resolve DNS queries for one or more VPCs, choose the VPCs. To associate a VPC with a hosted zone when the VPC was created using a different AWS account, you must use a programmatic method, such as the AWS CLI.

[?](#) For each VPC that you associate with a private hosted zone, you must set the Amazon VPC settings [enableDnsHostnames](#) and [enableDnsSupport](#) to true. ✕

**Region** [Info](#) US West (Oregon) [us-west-2] ▼

**VPC ID** [Info](#) 🔍 Choose VPC Remove VPC

- vpc-9f6c04e7
- vpc-0a7c93f9db7cf1dd2 (LK-VPC)

vpc-0a7c93f9db7cf1dd2 (LK-VPC)

5. The Hosted Zone is now created. At this point, there are only two entries (NS & SOA) created.

Route 53 > Hosted zones > test.local

## test.local Info

Delete zone Test record Configure query logging

▶ **Hosted zone details** Edit

**Records (2)** | Hosted zone tags (0)

---

**Records (2) Info**  
 Automatic mode is the current search behavior optimized for best filter results. To change modes go to settings.

<input type="checkbox"/>	Record name	Type	Routing policy	Differentiator	Value/Route traffic to
<input type="checkbox"/>	test.local	NS	Simple	-	ns-1536.awsdns-00.co.uk. ns-0.awsdns-00.com. ns-1024.awsdns-00.org. ns-512.awsdns-00.net.
<input type="checkbox"/>	test.local	SOA	Simple	-	ns-1536.awsdns-00.co.uk. awsdns-hostmaster.amazon.com. 1 7200 900 1209600 86400

6. To connect from another VPC, **the other VPC** must be associated.

**!** In order to use DNS on VPC, enable **DNS Hostnames** and **DNS Resolution** in the VPC settings.

Go to the VPC settings and ensure that these two values are set to `enabled`.

VPC > Your VPCs > vpc-0a7c93f9db7cf1dd2

## vpc-0a7c93f9db7cf1dd2 / LK-VPC

**Actions ▲**

- Create flow log
- Edit CIDRs
- Edit DHCP options set
- Edit DNS hostnames
- Edit DNS resolution
- Manage tags
- Delete VPC

**Details** [Info](#)

<p>VPC ID   vpc-0a7c93f9db7cf1dd2</p> <p>Tenancy Default</p> <p>Default VPC No</p> <p>Owner ID   970036242583</p>	<p>State  <span style="color: green; font-weight: bold;">✔ Available</span></p> <p>DHCP options set  <a href="#">dopt-5716772f</a></p> <p>IPv4 CIDR 10.20.0.0/16</p>	<p>DNS hostnames Disabled</p> <p>Route table  <a href="#">rtb-015a6633d04164f6d</a></p> <p>IPv6 pool -</p>
-----------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------

## Create Entries on the Hosted Zone

1. Create records as shown below:

Route 53 > Hosted zones > test.local

**test.local** Info Delete zone Test record Configure query logging

▶ **Hosted zone details** Edit

**Records (5)** | Hosted zone tags (0)

**Records (5)** Info  
Automatic mode is the current search behavior optimized for best filter results. [To change modes go to settings.](#)

Refresh Edit Delete record Import zone file Create record

🔍 *Filter records by property or value* Type ▼ Routing policy ▼ Alias ▼ < 1 > ⚙️

<input type="checkbox"/>	Record name ▼	Type ▼	Routing policy ▼	Differentiator ▼	Value/Route traffic to ▼
<input type="checkbox"/>	test.local	NS	Simple	-	ns-1536.awsdns-00.co.uk. ns-0.awsdns-00.com. ns-1024.awsdns-00.org. ns-512.awsdns-00.net.
<input type="checkbox"/>	test.local	SOA	Simple	-	ns-1536.awsdns-00.co.uk. awsdns-hostmaster.amazon.com. 1 7200 900 1209600 86400
<input type="checkbox"/>	node-a.test.local	A	Simple	-	10.20.1.10
<input type="checkbox"/>	node-active.test.local	A	Simple	-	10.20.1.10
<input type="checkbox"/>	node-b.test.local	A	Simple	-	10.20.2.10

2. Confirm the hosts can be resolved by name.

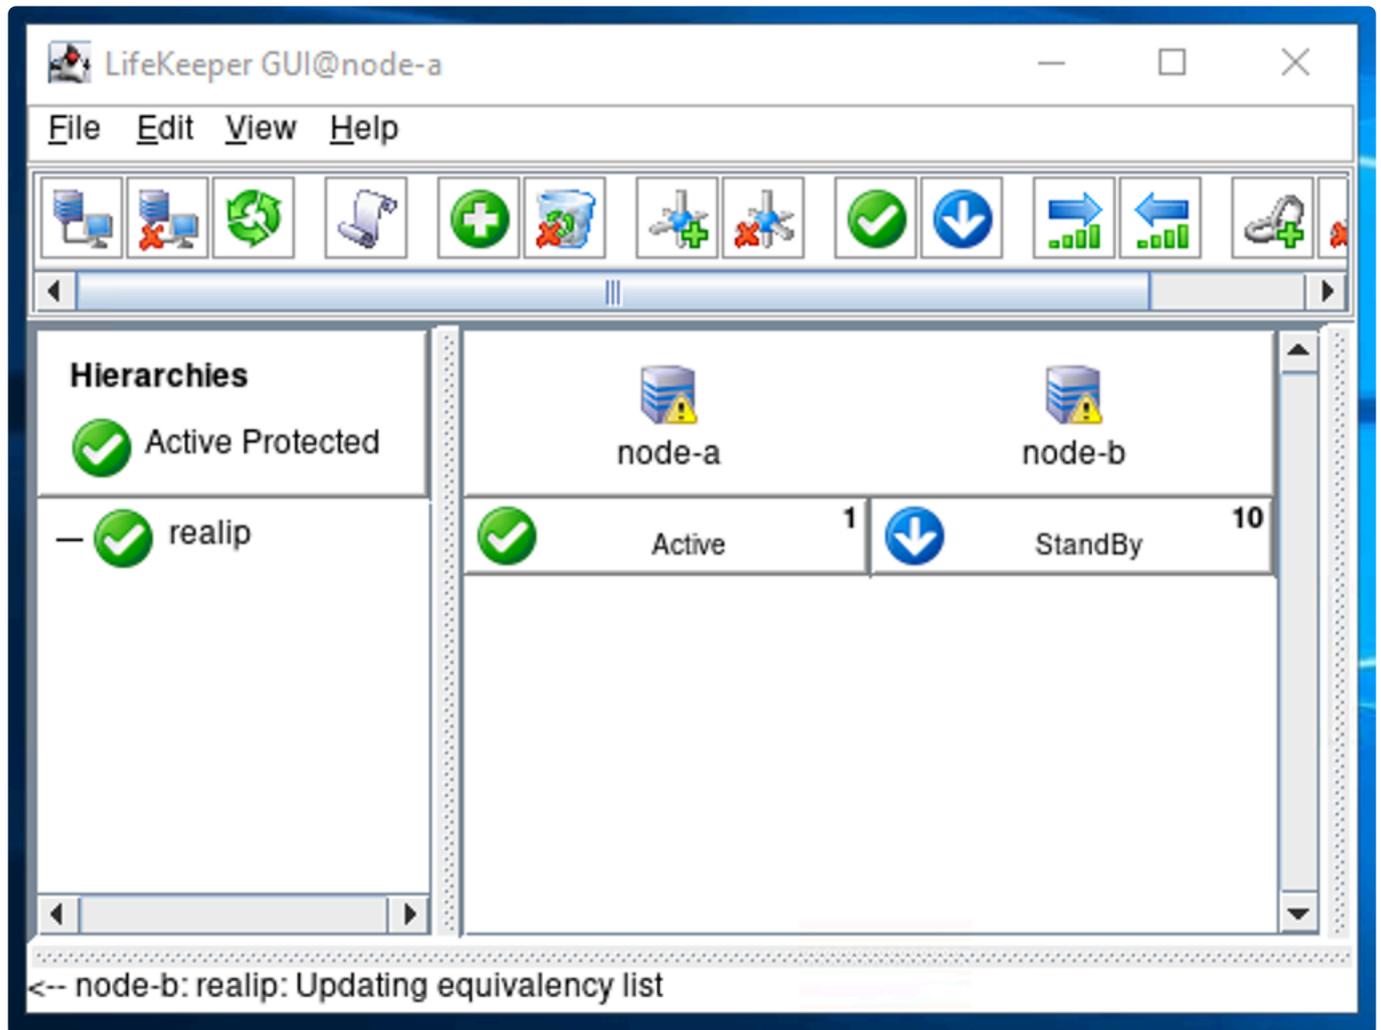
```

1 [root@node-a ~]# ping node-a.test.cl
2 PING node-a.test.cl (10.20.1.10) 56(84) bytes of data.
3 64 bytes from node-a (10.20.1.10): icmp_seq=1 ttl=64 time=0.019 ms
4 64 bytes from node-a (10.20.1.10): icmp_seq=2 ttl=64 time=0.022 ms
    
```

## Create an IP Resource

Now you can create an IP resource. Note that the IP address of the resource has to be an active node (the IP address of node-active.test.cl: 10.20.1.10). Refer to [Creating an IP Resource](#) for more information. Once the IP address of the active node (10.20.1.10) has been selected, the “IP Resource Tag” field is set to “realip” and the IP address of node-b (secondary) is set to 10.20.2.10 (the actual IP address of the second node). Those are automatically selected by selecting the “realip” id for the first node.

Once the IP resource is created, the LifeKeeper User Interface should look like this.



## Assign Permission to Update Route 53 Entries to Instances

! To complete the steps below, the AWS CLI must be installed on each node.

- [Install AWS CLI](#)

To allow LifeKeeper (running on the EC2 Instance) to update a Route53 entry, the instance must have roles that allow DNS updates.

- route53:ListHostedZones
- route53:GetChange
- route53:ChangeResourceRecordSets
- route53:ListResourceRecordSets

To achieve this, create a policy like the one seen below (note that it may be necessary to limit the resources that may be accessed), then assign it to a Role.

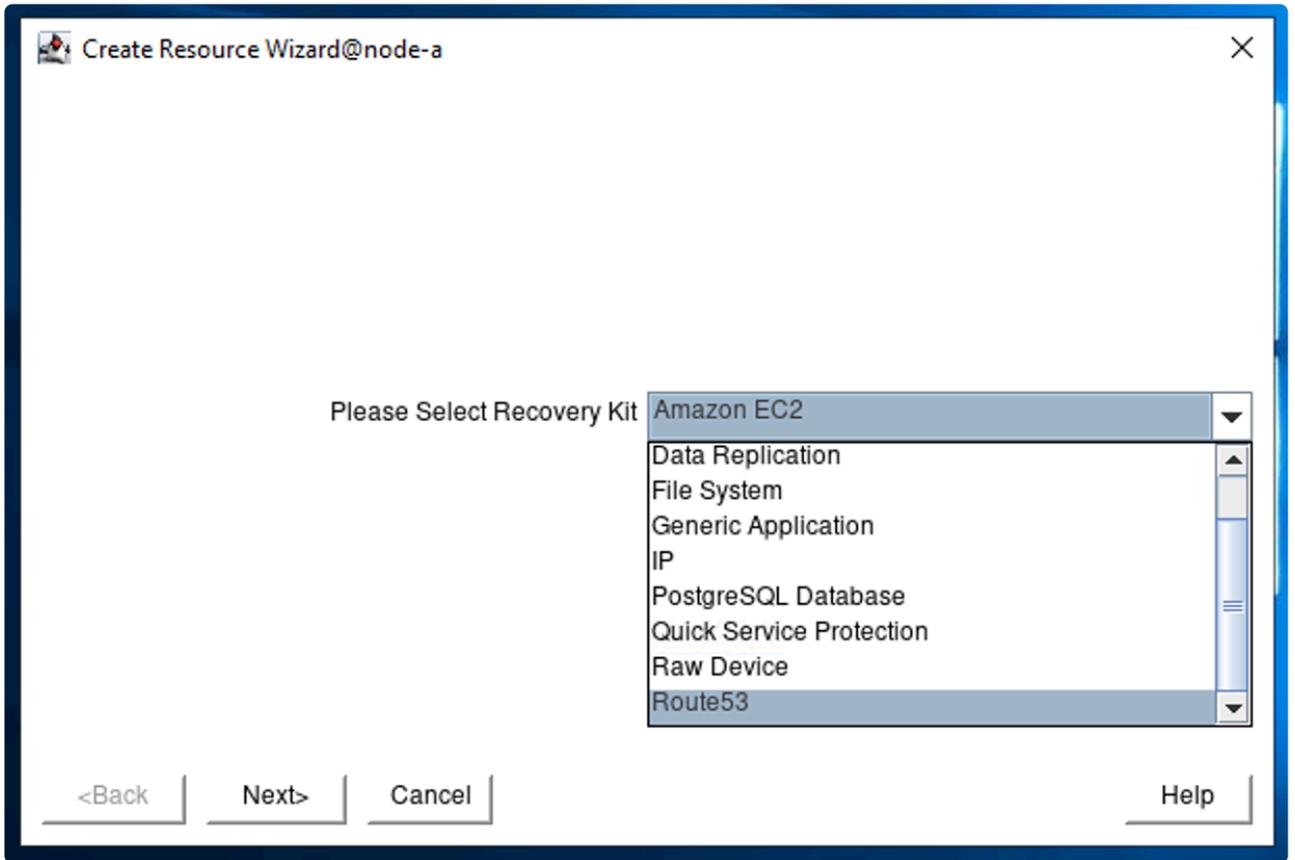
```
1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Sid": "VisualEditor0",
6       "Effect": "Allow",
7       "Action": [
8         "route53:GetChange",
9         "route53:ChangeResourceRecordSets",
10        "route53:ListResourceRecordSets"
11      ],
12      "Resource": [
13        "arn:aws:route53::hostedzone/*",
14        "arn:aws:route53::change/*"
15      ]
16    },
17    {
18      "Sid": "VisualEditor1",
19      "Effect": "Allow",
20      "Action": "route53:ListHostedZones",
21      "Resource": "*"
22    }
23  ]
24 }
```

Once a Role is defined, assign it to these EC2 Instances.

## Create the Route53 Resource

Creating the Route53 Resource should be straightforward once we have created the IP Resource.

1. Select Route53 as Recovery Kit.



2. On the “Create Resource Wizard @ node-a”, specify the following values.

 **Note:** For items with a checkmark (  ) review the default value and use the value suggested.

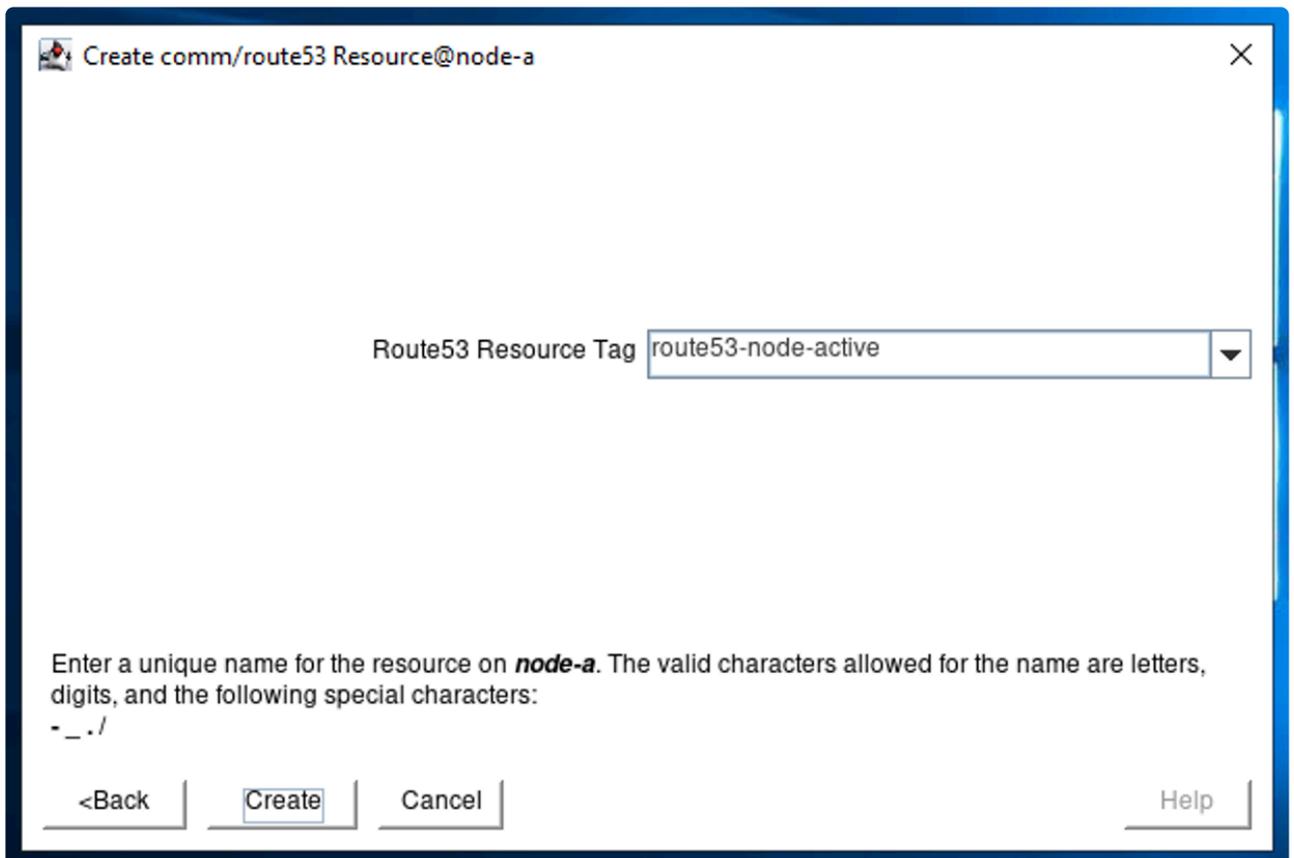
Field	Value
Switchback Type	Intelligent 
Server	node-a 
Domain Name (Route53 hosted zone)	test.local
Host Name	node-active
IP Resource	realip 
Route53 Resource Tag	reoute53-node-active 

3. On the “Pre-Extend Wizard” specify the following values:

Field	Value
Target Server	node-b 

Switchback Type	Intelligent 
Template Priority	1 
Target Priority	10 

- Once Pre-Extend is completed, move on to “Extend comm/route53 Resource”. Select “route53-node-active”.



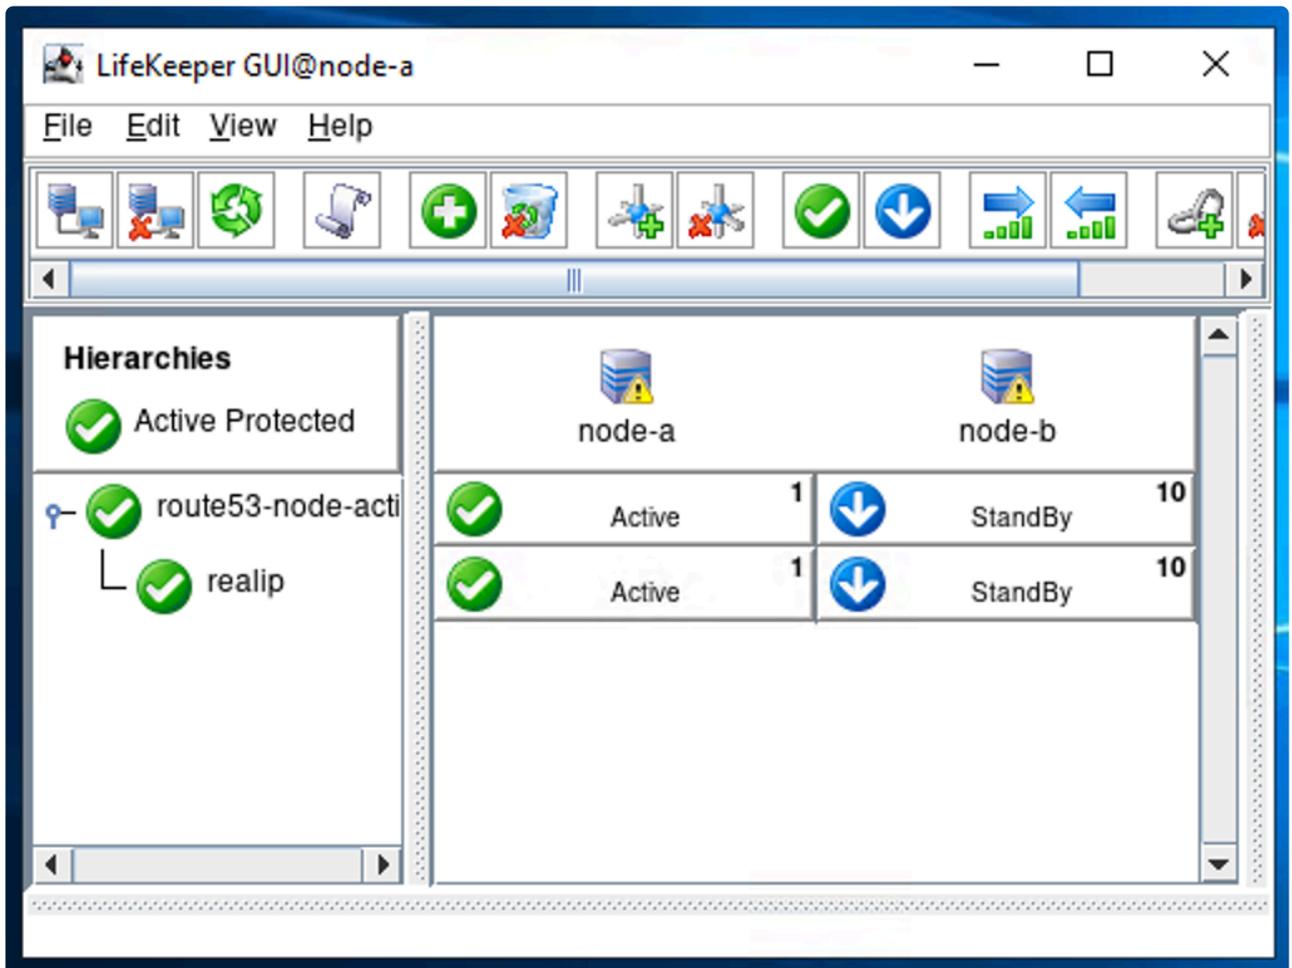
Create comm/route53 Resource@node-a

Route53 Resource Tag

Enter a unique name for the resource on **node-a**. The valid characters allowed for the name are letters, digits, and the following special characters:  
- \_ . /

<Back Create Cancel Help

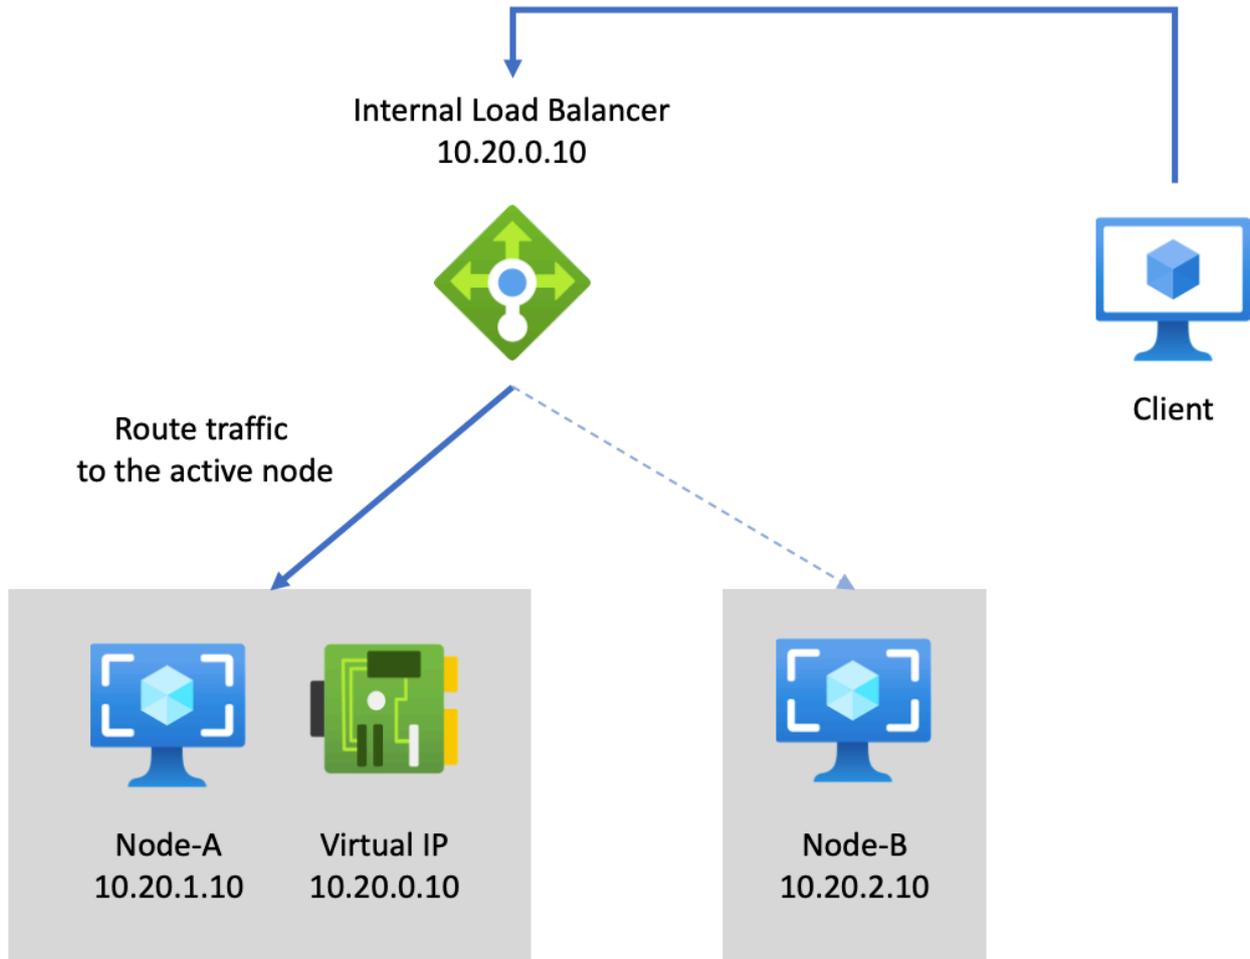
- Once “Hierarchy successfully extended” is displayed, the creation of the resource is complete.
- The Route53 resource is now created as shown below.



Notice that realip is located under route53-node-active. This indicates that the “route53-node-active” resource depends on the “realip” resource.

## 11.2.7.2.4. Azure – Using an Internal Load Balancer

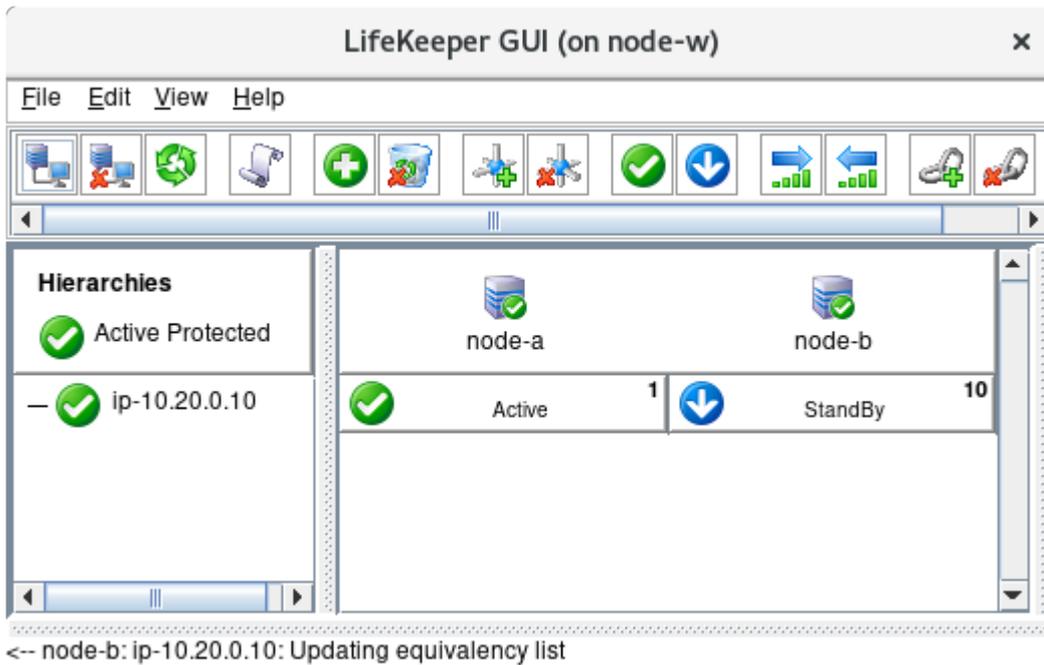
Here we will discuss how to switch between nodes using an Azure Internal Load Balancer by creating an Internal Load Balancer that points to the active node. Clients connect to the Frontend IP address provided by the Internal Load Balancer. The VMs are checked regularly by the “Health Check Probe” function and the Internal Load Balancer routes requests to the active node.



### Create an IP Resource

On Azure, a client connects to a load balancer so that it can distribute the traffic to the active node. In theory, an IP Resource (that represents a Virtual IP address) is not needed. However, as this section describes later, it is still necessary to create a Virtual IP address.

In Azure’s case, the Virtual IP address should be the same as the IP address for a load balancer in front of these nodes. In the example below we will use the IP address `10.20.0.10`.



**!** Before configuring a load balancer, it is important to check which port the application listens to (for example, port 80 for http). To ensure that the application is running, the Internal Load Balancer periodically checks the status of the application (this is referred to as a “health check probe”). It then sends incoming requests to the active node. The following sections use port 80 as an example, but this should be modified as appropriate for the application being protected.

## Create an Internal Load Balancer

An internal load balancer distributes traffic to the active node and can be created using the following steps. In order to configure the load balancer, start the application on node-a to ensure the load balancer is working before configuring it through LifeKeeper.

1. On the Azure Portal, go to Load Balancer, click “Add”, select the following parameters and then click “Review + Create”. Once the parameters are confirmed, the Load Balancer can be created.

Item	Value
Resource Group	LK-QSG
Name	LK-ILB
Region	(same as the Virtual Machines)
Type	Internal
SKU	Standard (select Standard as the workload is distributed across Availability Zones)
Virtual Network	LK-VNET
Subnet	LK-subnet

IP address assignment	Static
Private IP Address	10.20.0.10
Availability zone	Zone-redundant

>> All services > Load balancers >

## Create load balancer ×

Basics Tags Review + create

Azure load balancer is a layer 4 load balancer that distributes incoming traffic among healthy virtual machine instances. Load balancers uses a hash-based distribution algorithm. By default, it uses a 5-tuple (source IP, source port, destination IP, destination port, protocol type) hash to map traffic to available servers. Load balancers can either be internet-facing where it is accessible via public IP addresses, or internal where it is only accessible from a virtual network. Azure load balancers also support Network Address Translation (NAT) to route traffic between public and private IP addresses. [Learn more.](#)

**Project details**

Subscription \* Azure Pay As You Go Subscription

Resource group \* LK-QSG [Create new](#)

**Instance details**

Name \* LK-ILB ✓

Region \* (US) West US 2

Type \*  Internal  Public

SKU \*  Basic  Standard

**Tier**

Regional  Global

**Configure virtual network.**

Virtual network \* LK-VNET

Subnet \* LK-subnet (10.20.0.0/22) [Manage subnet configuration](#)

IP address assignment \*  Static  Dynamic

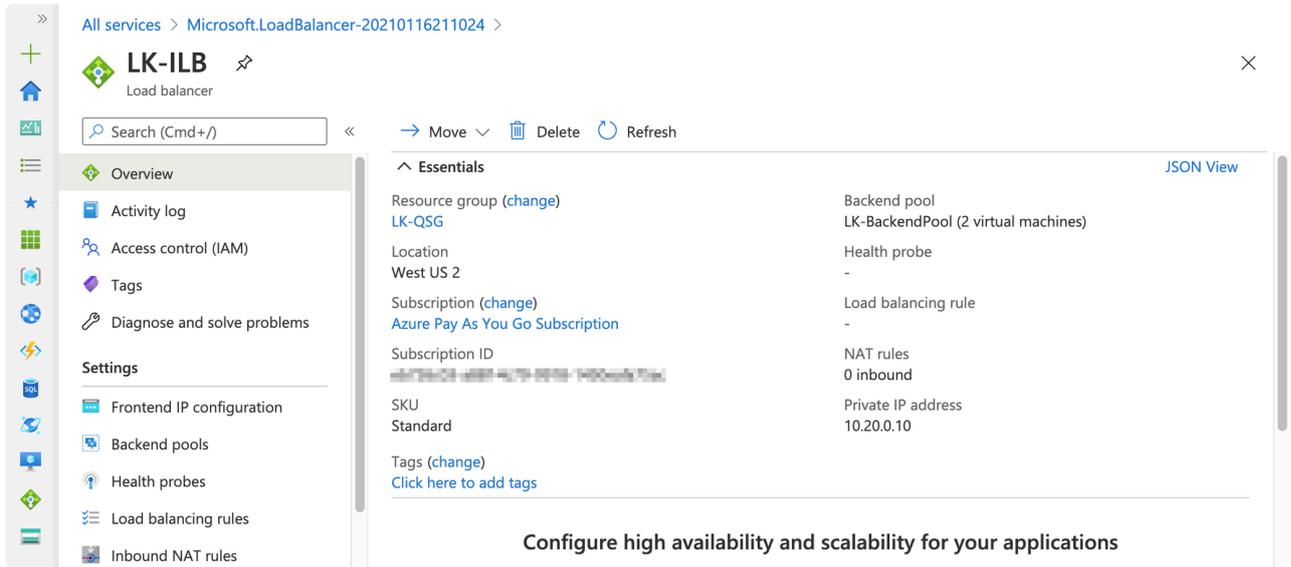
Private IP address \* 10.20.0.10 ✓

Availability zone \* Zone-redundant

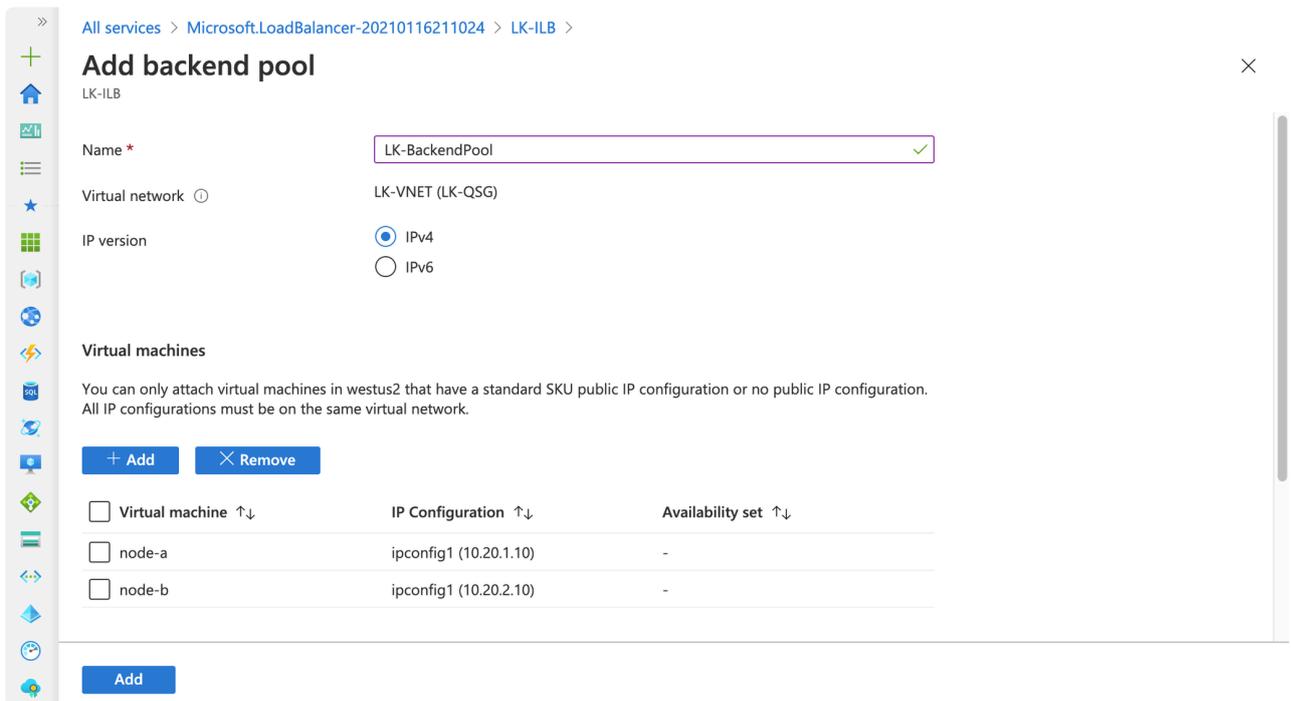
[Review + create](#) < Previous [Next: Tags >](#) [Download a template for automation](#)

2. Once the load balancer is created, go to the resource page (LK-ILB). Configure the following sections of the load balancer (located in the Setting section on the left side of the Load Balancer configuration page).

- Backend pools (step 3)
- Health probes (step 4)
- Load Balancing Rules (step 5)



3. Go to the Backend pools page and click “Add”. Once the Backend Pool Page is open, select node-a and node-b from the list of available virtual machines and provide a name for the pool (for example, LK-BackendPool).



4. Go to the Health Probes page and click “Add”. Select the TCP port number used by the application (e.g., 80 if we are going to protect `httpd`), then provide a name for the probe (for example, LK-Probe).

>> All services > Microsoft.LoadBalancer-20210116211024 > LK-ILB >

### Add health probe

LK-ILB

Name \*

Protocol

Port \*

Interval \*  seconds

Unhealthy threshold \*  consecutive failures

OK

5. Go to the Load Balancing Rule page and click “Add” to create a new rule. Enter the following values:

Item	Value
Name	LK-ILB-Rule
Frontend IP address	10.20.0.10
Protocol	TCP (assuming your app uses TCP)
Port #	Select the Port # in use by the application
Backend Port	Same as Port
Backend Pool	LK-BackendPool
Health Probe	LK-Probe
Session Persistence	None
Floating IP	Enabled (you can also choose Disabled). See following section for details.

>> All services > Microsoft.LoadBalancer-20210116211024 > LK-ILB >

## Add load balancing rule

LK-ILB

**i** A load balancing rule distributes incoming traffic that is sent to a selected IP address and port combination across a group of backend pool instances. Only backend instances that the health probe considers healthy receive new traffic.

Name \*  
LK-ILB-Rule ✓

IP Version \*  
 IPv4  IPv6

Frontend IP address \* ⓘ  
10.20.0.10 (LoadBalancerFrontEnd) ✓

HA Ports ⓘ

Protocol  
 TCP  UDP

Port \*  
80 ✓

Backend port \* ⓘ  
80 ✓

Backend pool ⓘ  
LK-BackendPool (2 virtual machines) ✓

Health probe ⓘ  
LK-Probe (TCP:80) ✓

Session persistence ⓘ  
None ✓

Idle timeout (minutes) ⓘ  
○ ————— □

TCP reset  
 Disabled  Enabled

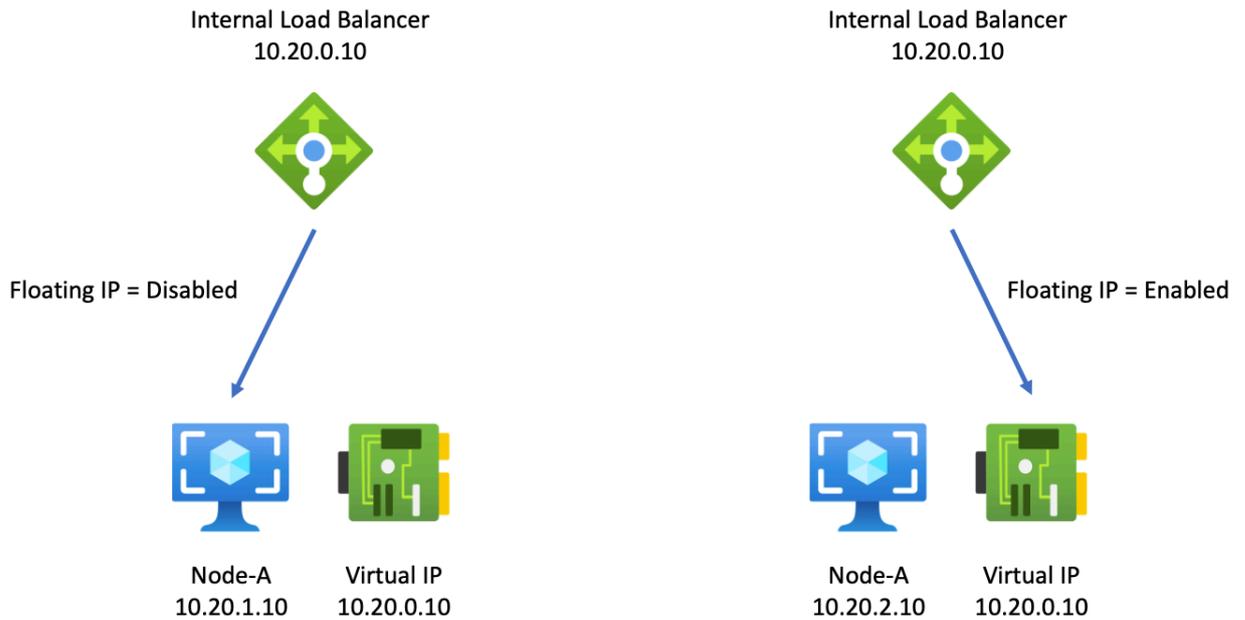
Floating IP ⓘ  
Disabled **Enabled**

OK

## The Difference between Floating IP Options

The difference between Floating IP options determines where the Internal Load Balancer sends packets to.

- If Disabled is selected, the Load Balancer sends packets to the node’s primary IP.
- If Enabled is selected, the Load Balancer sends packets to the Virtual IP.

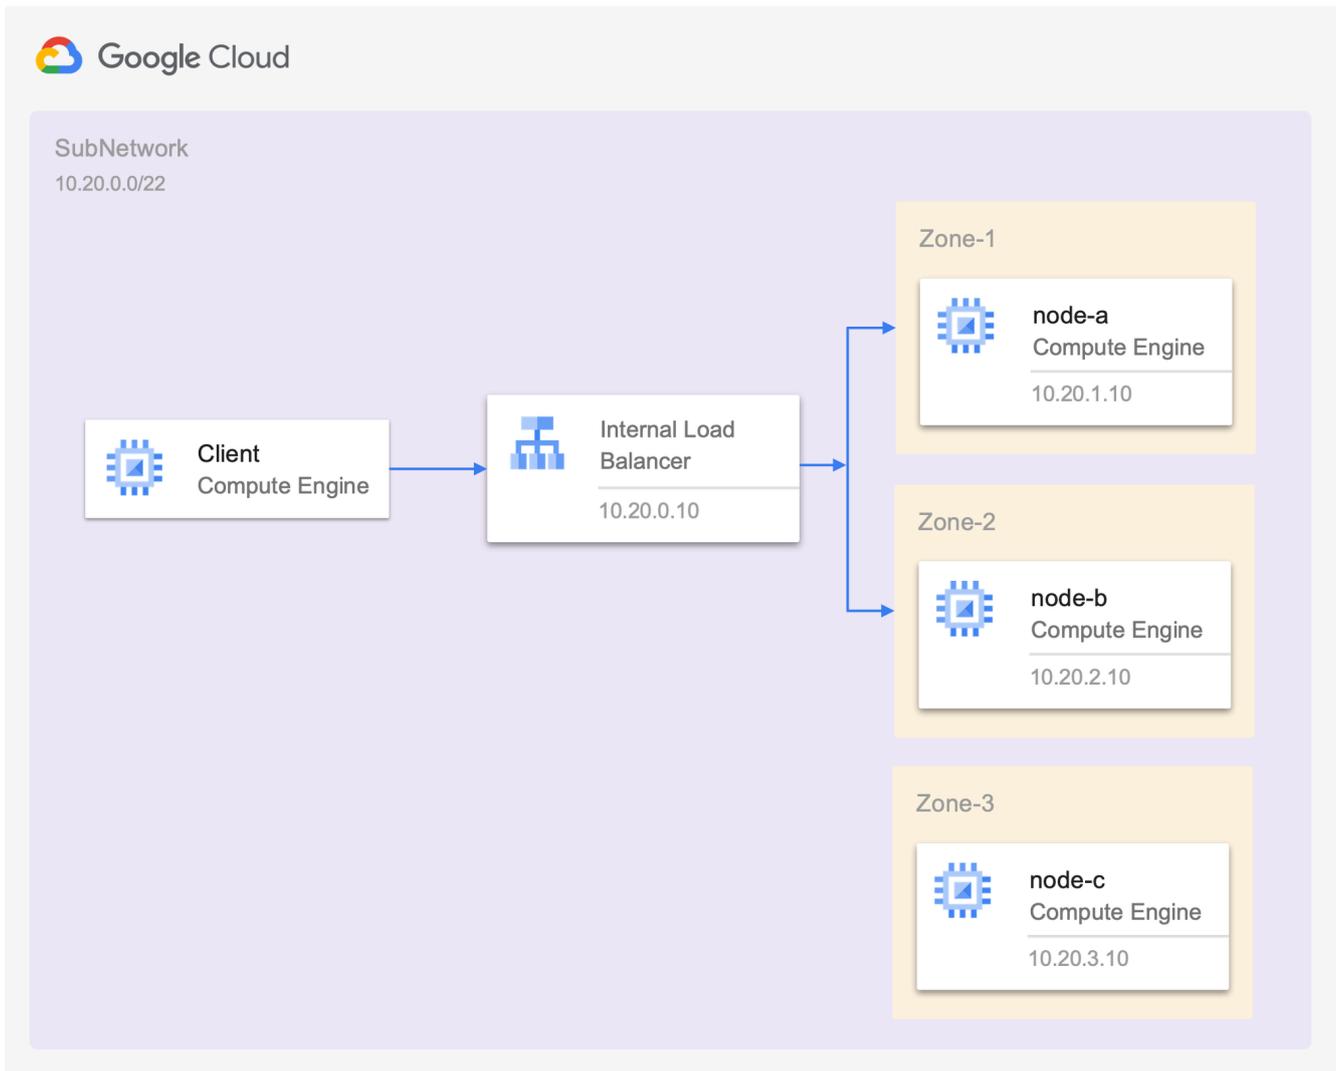


Because the Virtual IP is attached to the active node, it appears to be the same operation. However, if the response from the active node is required to be sent from the Virtual IP, “Enabled” should be selected. This is a more natural way for LifeKeeper to handle these nodes (in the same way as in an on-premise environment). If the application listens only to the primary IP address (for example, 10.20.1.10), the Floating IP parameter should be set to “Disabled”.

**Note:** The Health Probe checks the status of the node by sending requests to the primary IP address (for example, 10.20.1.10). This is true regardless of the “Floating IP” option value.

# 11.2.7.2.5. Google Cloud – Using an Internal Load Balancer

Here we will discuss how to switch between nodes using Google Cloud’s Internal Load Balancer by creating an Internal Load Balancer that routes traffic to the active node. Clients connect to the Frontend IP address provided by Internal Load Balancer. The Internal Load Balancer regularly checks the health of each VM in the backend pool using a user-defined “Health Check Probe” function, and then routes client requests to the active node.



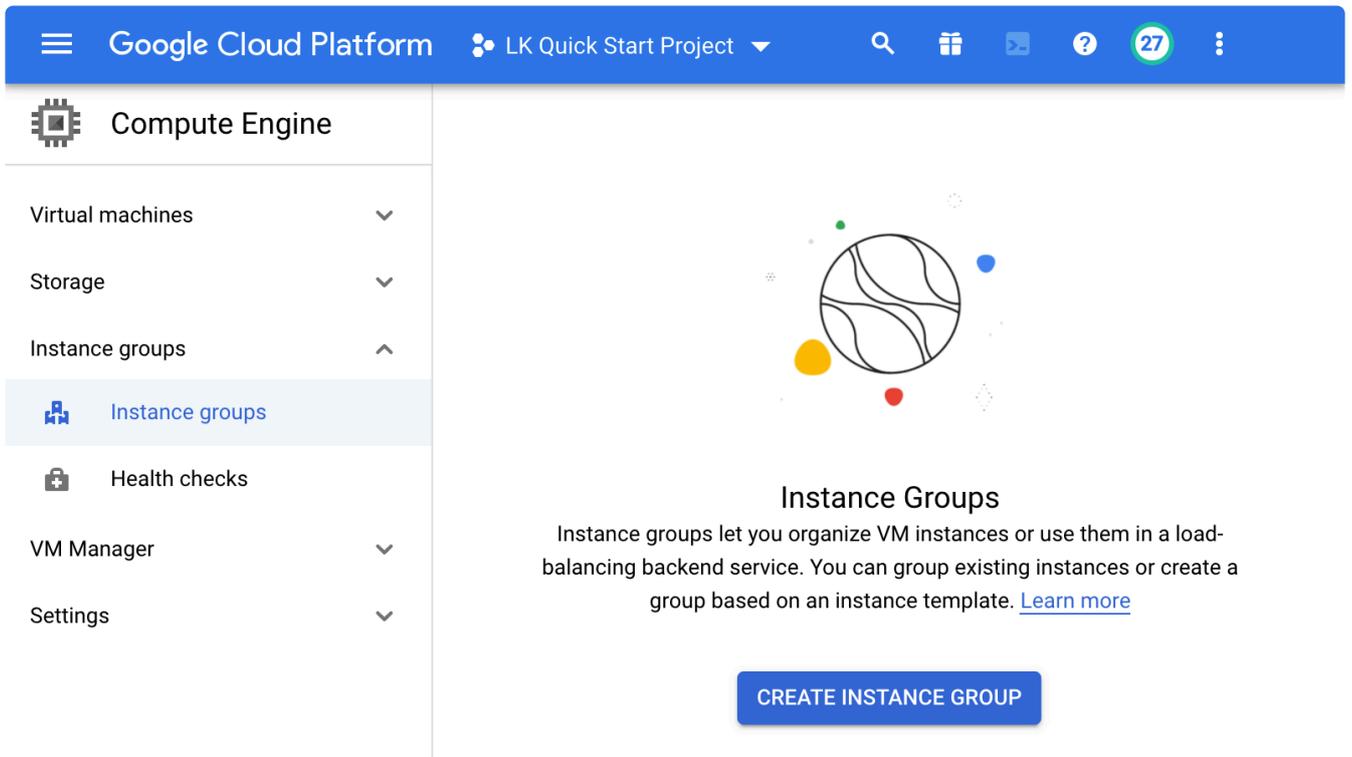
## Create Instance Groups

In order to use an Internal Load Balancer, the first step is to create Instance Groups containing each of our nodes. Create Instance Groups with the following parameters:

Name	Region	Zone	Network	Subnetwork	VM Instance
lk-ig01	<Deployment region>	a	lk-vpc	lk-subnet	node-a

lk-ig02		b			node-b
---------	--	---	--	--	--------

1. Select "Compute Engine" > "Instance Group" from the navigation menu.



2. Select "CREATE INSTANCE GROUP", "New Unmanaged instance group", then enter the parameters listed in the table.

Google Cloud Platform LK Quick Start Project

### Create an instance group

To create an instance group, select one of the options:

- New managed instance group (stateless)**  
 For stateless serving and batch workloads.  
 Supports:
  - autoscaling, autohealing, auto-updating
  - multi-zone deployment
  - load balancing
- New managed instance group (stateful)**  
 For stateful workloads such as databases.  
 Supports:
  - disk and metadata preservation
  - autohealing and updating
  - multi-zone-deployment
  - load balancing
- New unmanaged instance group**  
 A group of VMs that you manage yourself.
  - load balancing

Organize VM instances in a group to manage them together. [Instance groups](#)

**Name** ?  
Name is permanent  
lk-ig01

**Description** (Optional)

**Location**  
**Region** ? Region is permanent  
us-west1 (Oregon)  
**Zone** ? Zone is permanent  
us-west1-a

**Specify port name mapping** (Optional)

**Network** ?  
lk-vpc

**Subnetwork** ?  
lk-subnet (10.20.0.0/22)

**VM instances**

node-a ×

No available instances ▼

Your free trial credit will be used for VM instances in this group. [GCP Free Tier](#)

**Create** **Cancel**

Equivalent [REST](#) or [command line](#)

3. Create lk-ig02 following the same steps.

4. Two Instance groups are now created.

Google Cloud Platform LK Quick Start Project Search products and resources

Compute Engine Instance Groups **CREATE INSTANCE GROUP** **REFRESH** **DELETE**

Virtual machines ▼

Storage ▼

Instance groups ▲

**Instance groups**

Health checks

Instance groups are collections of VM instances that use load balancing and automated services, like autoscaling and autohealing. [Learn more](#)

Filter table

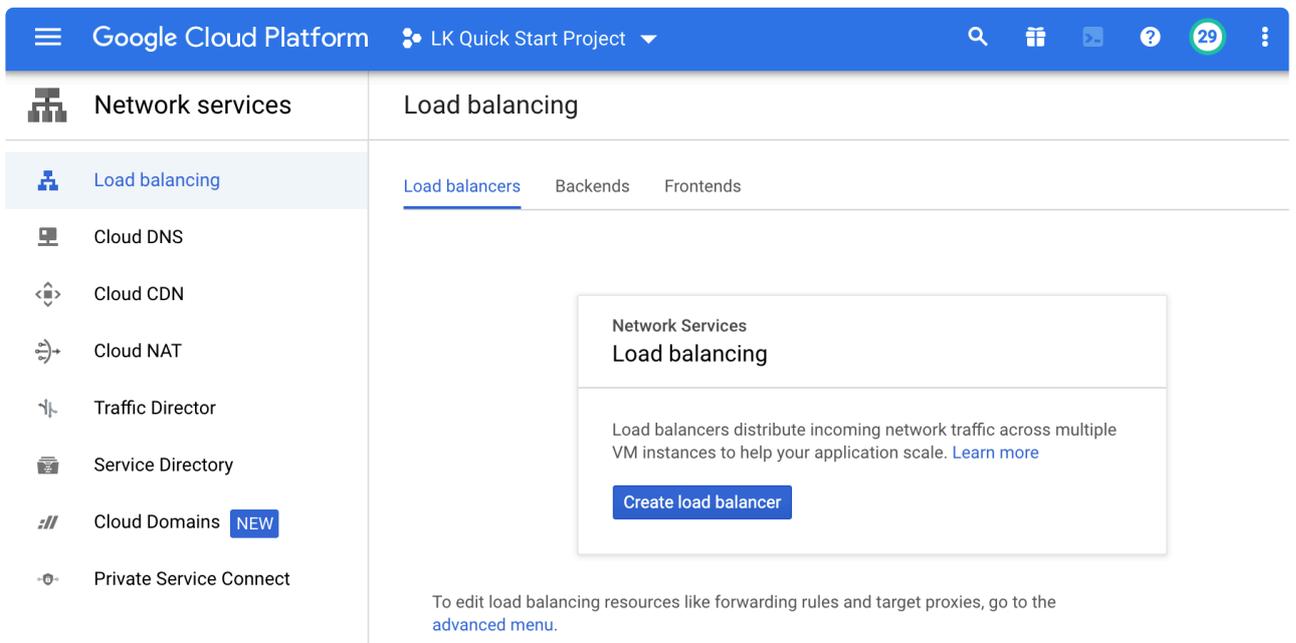
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Name <span>↑</span>	Instances	Group Type	Zone	In Use By
<input type="checkbox"/>	<input checked="" type="checkbox"/>	lk-ig01	1	Unmanaged	us-west1-a	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	lk-ig02	1	Unmanaged	us-west1-b	

# Create an Internal Load Balancer

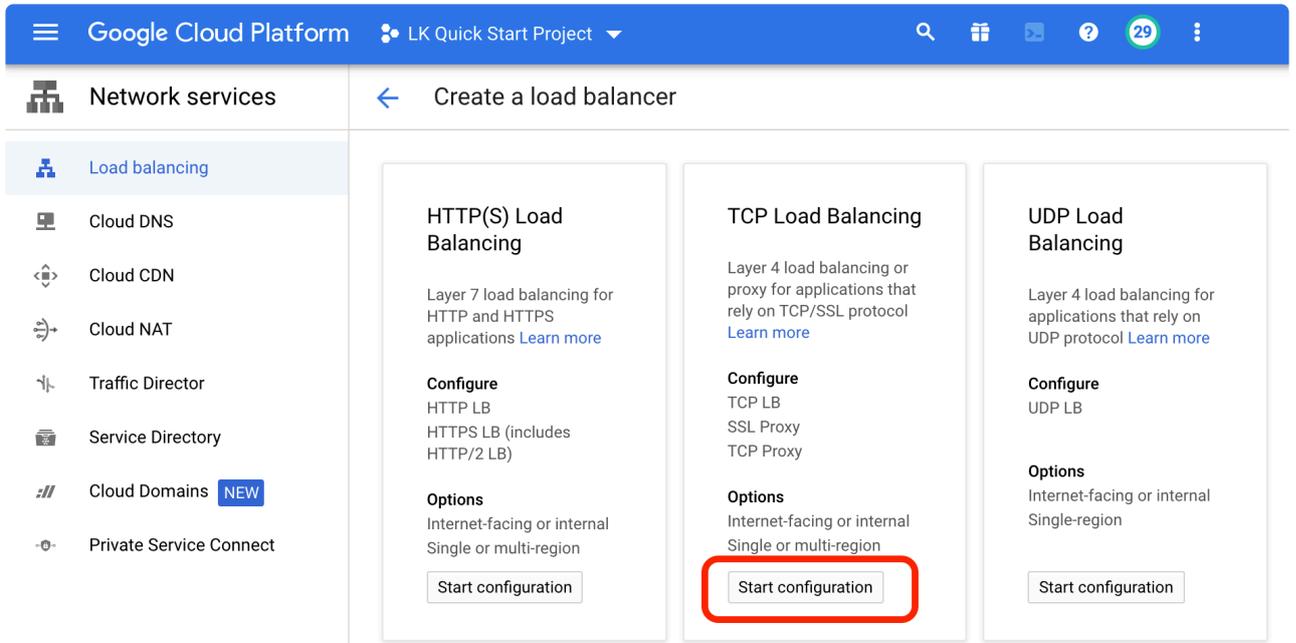
**!** Before configuring a load balancer, it is important to check which port the application listens to (for example, port 80 for http). To ensure that the application is running, the Internal Load Balancer periodically checks the status of the application (this is referred to as a “health check probe”). It then sends incoming requests to the active node. The following sections use port 80 as the example (which would be appropriate when setting up a health probe on a set of backend targets running web servers), but you should modify the port based on the application that you are protecting.

An internal load balancer distributes traffic to the active node and can be created using the following steps. In order to configure the load balancer, start the application on node-a (to ensure the load balancer is working before configuring it through LifeKeeper).

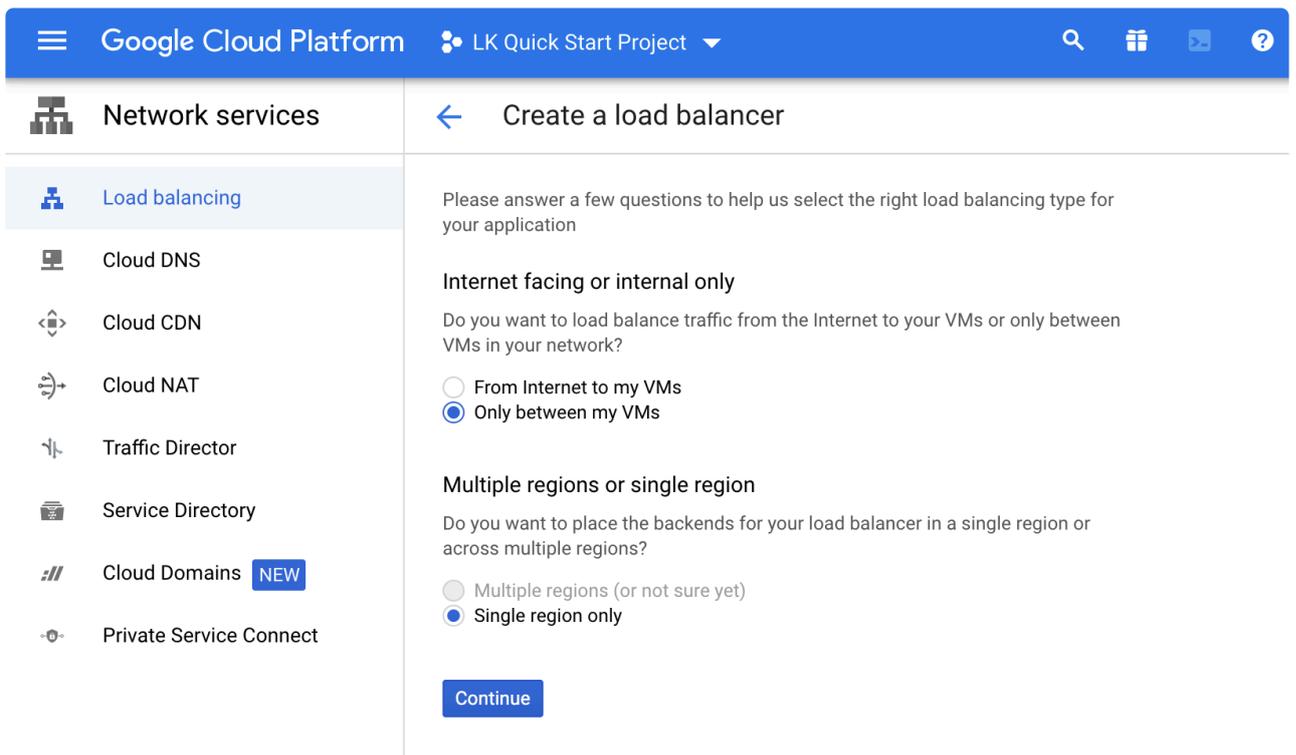
1. On the navigation menu, select “Network Service” > “Load Balancing”.



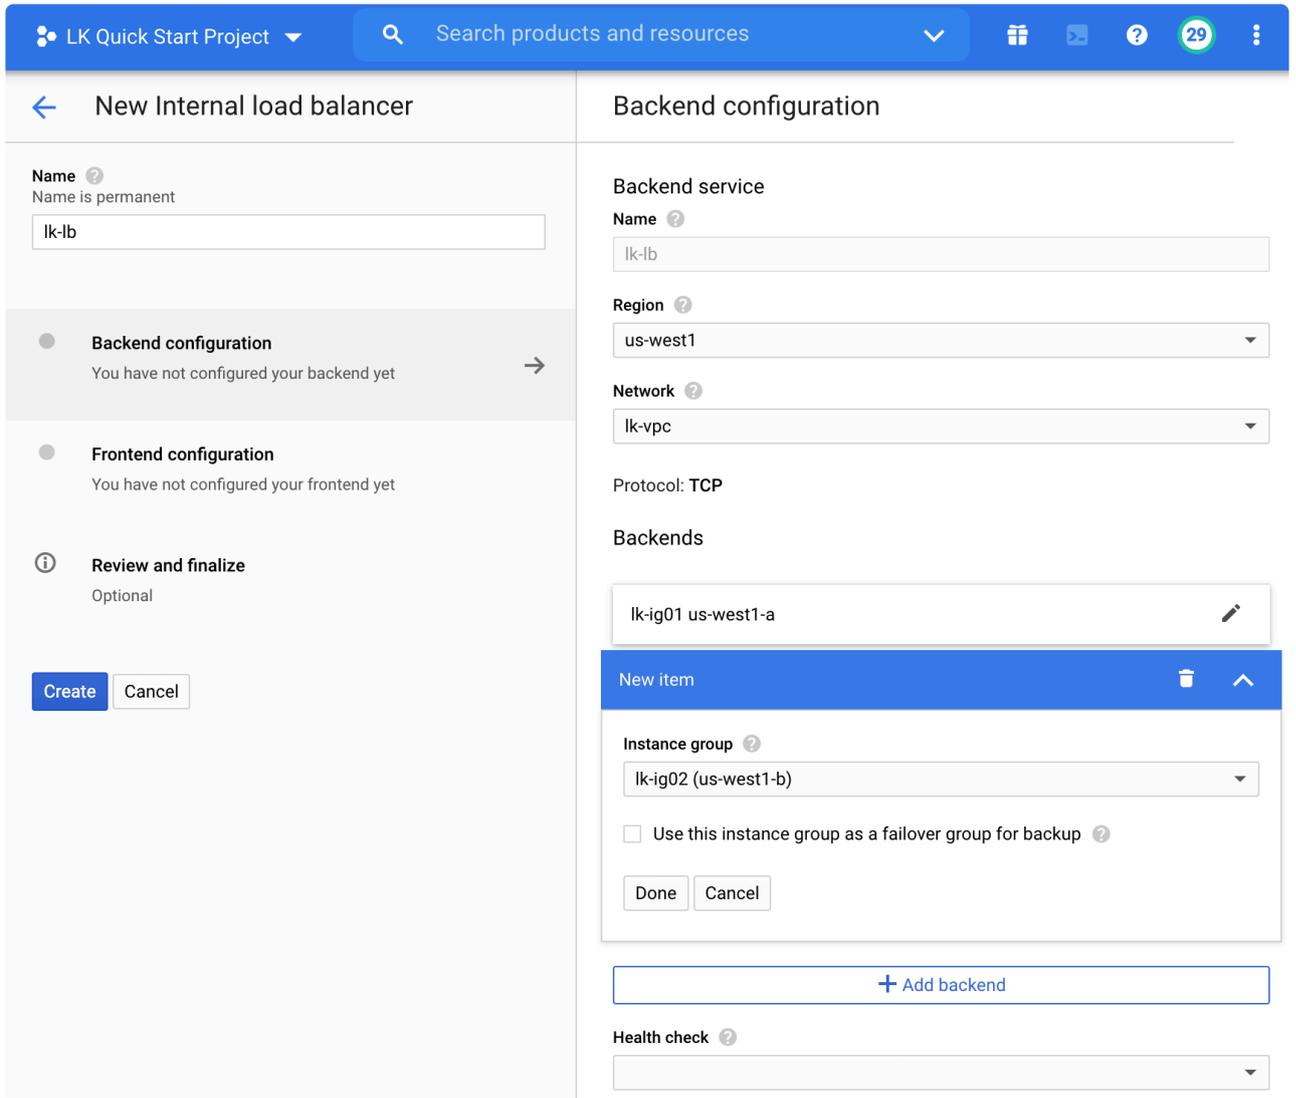
2. Click “Start Configuration” under “TCP Load Balancing”.



- 3. Select "Only between my VMs" (meaning Internal Load Balancer), and Single region only. Click "Continue".



- 4. Specify the name as `lk-lb`, then configure the backend parameters. Select Region, Network as `lk-vpc`, then select two instance groups (`lk-ig01` and `lk-ig02`) from the dropdown menu to be added to the backend pool.



- 5. Move on to “Health check” located at the bottom of the Backend configuration page. Select “Create a health check”.



- 6. The health check configuration screen appears. Select the name `lk-health-check`, select Port (make sure it matches with the service you are going to use), then review the other parameters and click “Save and continue”.

**Name \***  
lk-health-check ?  
Lowercase, no spaces.

Description

**Scope**

Global  
 Regional

**Protocol**  
TCP ?

**Port**  
80 ?

**Proxy protocol**  
NONE ?

**Request** ?      **Response** ?

**Logs**  
Turning on Health check logs can increase costs in Cloud Logging.  
 On  
 Off

### Health criteria

Define how health is determined: how often to check, how long to wait for a response, and how many successful or failed attempts are decisive

**Check interval**  
5 seconds ?

**Timeout**  
5 seconds ?

**Healthy threshold**  
2 consecutive successes ?

**Unhealthy threshold**  
2 consecutive failures ?

7. Move to the “Frontend configuration” tab. Provide the name and subnetwork for the frontend. In this example we will use the name `lk-frontend` and select the `lk-subnet` subnetwork. Under “IP address”, choose “Ephemeral (Custom)” and specify the custom ephemeral IP address to be used by the load balancer frontend. In this example we will use IP address `10.20.0.10`. Under “Ports”, select either “Single”, “Multiple”, or “All” as appropriate, based on the application that traffic is being routed to, and enter any required ports. In this example we will forward traffic on TCP port 80.

### New Frontend IP and port

**Name** (Optional) ?  
Name is permanent

[Add a description](#)

**Protocol**  
TCP

**Subnetwork**

**Internal IP**

**Purpose** ?

Non-shared  
 Shared

**IP address**

**Custom ephemeral IP address**

**Ports** ?

Single  
 Multiple  
 All

**Port number**

**Global access** ?

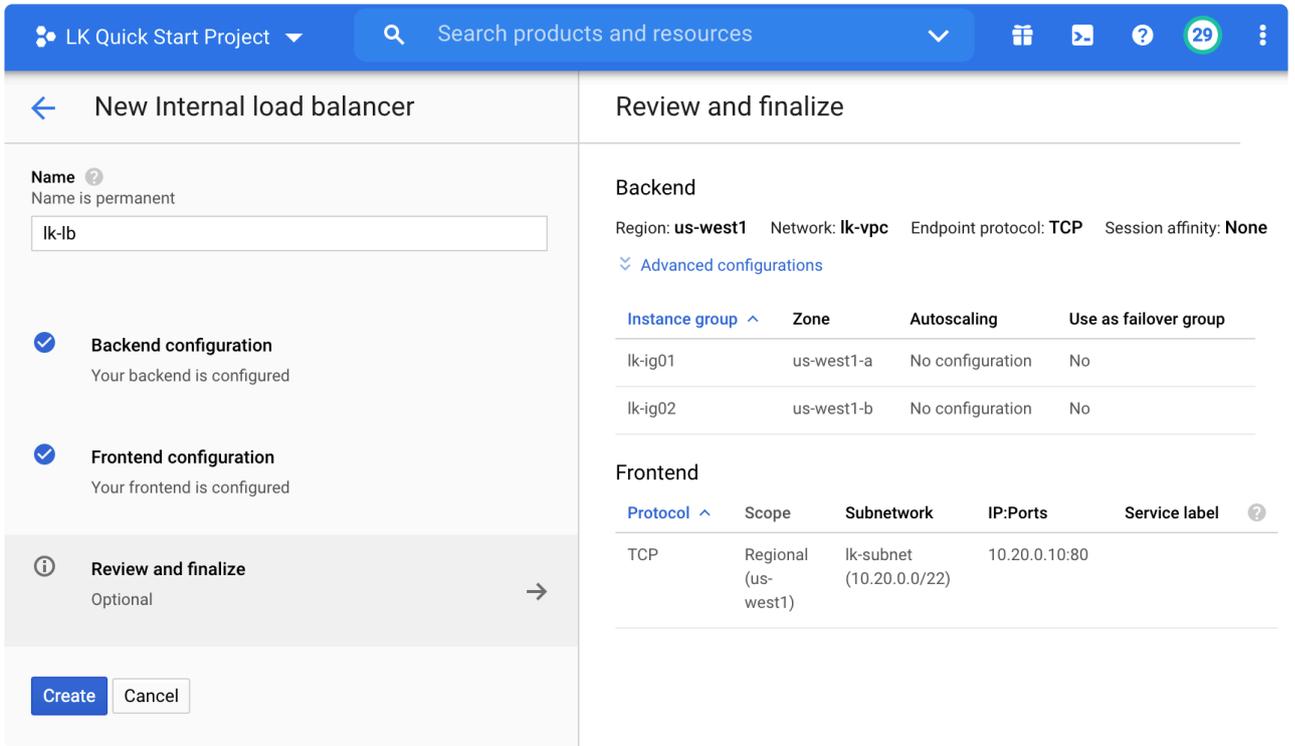
Disable  
 Enable

**Service label** ? (Optional)

[Advanced options](#)

[+ Add frontend IP and port](#)

8. Move on to the “Review and Finalize” tab and click “Create”.



**Name** <sup>?</sup>  
Name is permanent  
lk-lb

**Backend configuration**  
Your backend is configured

**Frontend configuration**  
Your frontend is configured

**Review and finalize**  
Optional

**Backend**  
Region: **us-west1** Network: **lk-vpc** Endpoint protocol: **TCP** Session affinity: **None**  
Advanced configurations

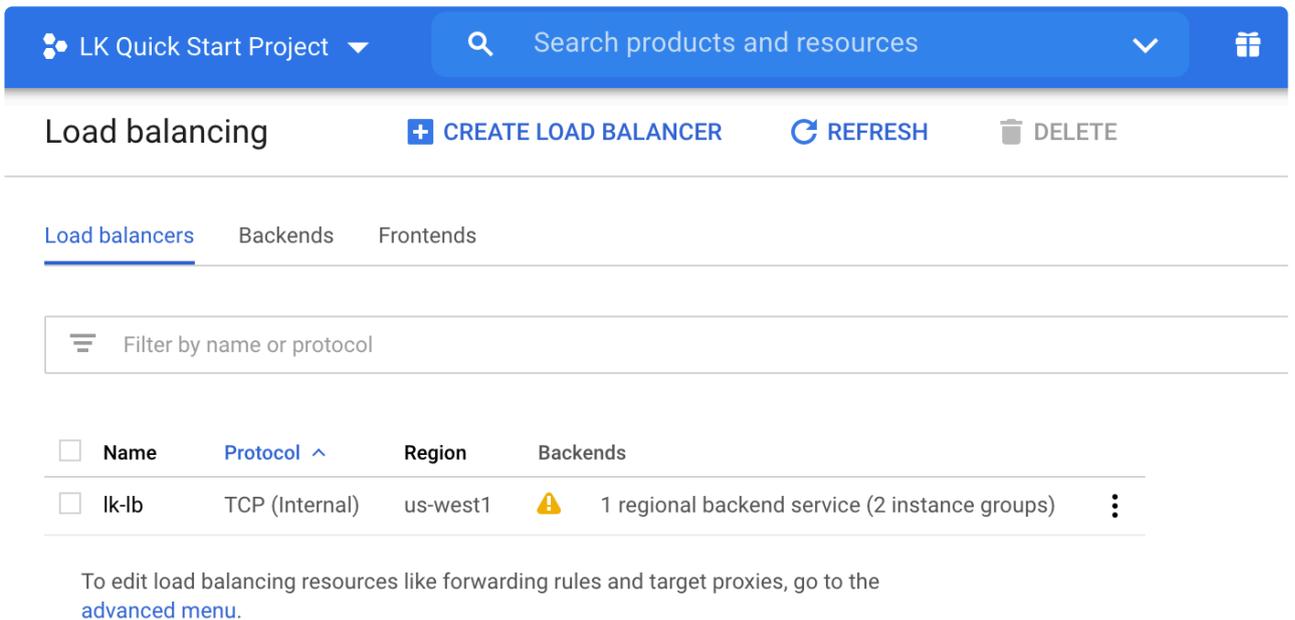
Instance group	Zone	Autoscaling	Use as failover group
lk-ig01	us-west1-a	No configuration	No
lk-ig02	us-west1-b	No configuration	No

**Frontend**

Protocol	Scope	Subnetwork	IP:Ports	Service label
TCP	Regional (us-west1)	lk-subnet (10.20.0.0/22)	10.20.0.10:80	

**Create** **Cancel**

9. Once the Load Balancer is created, the status will be displayed. As the application (such as httpd) runs on only one of these nodes, you may see the  icon, but this behavior is expected.



**Load balancing** **+ CREATE LOAD BALANCER** **REFRESH** **DELETE**

**Load balancers** Backends Frontends

Filter by name or protocol

<input type="checkbox"/>	Name	Protocol	Region	Backends
<input type="checkbox"/>	lk-lb	TCP (Internal)	us-west1	 1 regional backend service (2 instance groups)

To edit load balancing resources like forwarding rules and target proxies, go to the [advanced menu](#).

10. Now the Internal Load Balancer is configured. Once you install the application to protect on node-a, you can connect to it through the frontend IP (10.20.0.10) of the ILB you have just configured. The following example shows how to check the current target for http traffic via the ILB.

```
1 [gcp-user@node-c ~]$ curl http://10.20.0.10/
2 <html>
3 <body>
4 Test Page from node-a
5 </body>
6 </html>
```

## Disable IP Forwarding

In certain configurations, two applications which are hosted on different backend servers of an internal load balancer may need to communicate through the frontend IP of the load balancer itself. This happens, for example, when using internal load balancers to manage floating IP failover for ASCS and ERS instances in an SAP AS ABAP deployment.

As a consequence of the route-based load balancer implementation in Google Cloud, traffic that is sent from a backend server of a load balancer to the frontend IP of the load balancer will by default be routed back to the same backend server that it originated from, regardless of whether or not it is considered healthy by the load balancer's health checks. See the **Traffic is sent to unexpected backend VMs** section of [Troubleshooting Internal TCP/UDP Load Balancing](#) for more details.

The following steps provide a method for disabling IP forwarding on Google Cloud VMs, which also resolves the intra-backend traffic issue described above. These steps must be performed on **all** VMs which will be added as backend targets for the load balancer.

1. Set `ip_forwarding = false` in the `[NetworkInterfaces]` section of the `/etc/default/instance_configs.cfg` file.

```
# vi /etc/default/instance_configs.cfg
# cat /etc/default/instance_configs.cfg
[...]
[NetworkInterfaces]
[...]
ip_forwarding = false
[...]
```

2. Once the changes have been saved successfully, reboot the VM to allow the changes to take effect.

When IP forwarding is disabled, the LifeKeeper GenLB resource that responds to the load balancer's health checks must have a child IP resource protecting the frontend IP address of the load balancer and using net mask /32 (equivalently, 255.255.255.255). Without this child IP resource, the GenLB resource will fail to come in-service. See the **Create Frontend IP Resource(s)** and **Add Frontend IP Resources as Dependencies of GenLB Resource** sections of [Responding to Load Balancer Health Checks](#) for details on creating these resources and adding them as dependencies of the GenLB Resource.

See the **Test GenLB Resource Switchover and Failover** section of [Responding to Load Balancer Health Checks](#) for details on how to test for successful operation of the load balancer and corresponding GenLB resource hierarchy.

## 11.2.7.2.6. Responding to Load Balancer Health Checks

The **LifeKeeper Generic Application Recovery Kit for Load Balancer Health Checks** (“**GenLB Recovery Kit**”) may be used as part of a LifeKeeper resource hierarchy to help route load balancer traffic to the cluster node where a particular resource is currently in-service. This is achieved by maintaining a listener on a user-specified TCP port on the cluster node where the resource is in-service.

### Install the GenLB Recovery Kit

The GenLB Recovery Kit is only supported on LifeKeeper for Linux version 9.5.1 and later, and an rpm installation file can be obtained from the same FTP directory where the LifeKeeper for Linux installation media was downloaded. The filename format is `steeleye-lkHOTFIX-Gen-LB-PL-7172-x.x.x-xxxx.x86_64.rpm`. No additional SIOS licenses are required in order to create GenLB resources.

1. Install the rpm:

```
[root@node-a ~]# rpm -ivh steeleye-lkHOTFIX-Gen-LB-PL-7172-x.x.x-xxxx.x86_64.r
pm
Verifying...                               ##### [100%]
Preparing...                               ##### [100%]
Updating / installing... 1:steeleye-lkHOTFIX-Gen-LB-PL-7172-# [100%]
```

2. Verify that the GenLB resource action scripts have been successfully installed to the `/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172` directory.

```
[root@node-a ~]# ls -l /opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/
total 12
-r-x----- 1 root root 2579 Jan 01 00:00 quickCheck
-r-x----- 1 root root 1734 Jan 01 00:00 remove.pl
-r-x----- 1 root root 3909 Jan 01 00:00 restore.pl
```

### Create Frontend IP Resource(s)

 **Note:** This section only applies when using a Google Cloud internal load balancer where IP forwarding is disabled. For other configurations, you may proceed directly to the ‘Create a GenLB Resource’ section below.

When using a Google Cloud load balancer with IP forwarding disabled (see the ‘Disable IP Forwarding’ section of [Google Cloud – Using an Internal Load Balancer](#)), the IP address(es) associated to the load balancer frontend(s) must be added to a local network interface on each backend server using network mask /32 (equivalently, 255.255.255.255). In these configurations, the load balancer frontend IP address

is **not** automatically added to a network interface by any cloud agent process. Instead, it must be added manually within the guest operating system on each backend server. The simplest way to achieve this is by creating a LifeKeeper IP resource for each load balancer frontend IP address, which will then be added as a dependency of the GenLB resource that will be created in the next section.

Following the steps given in [Creating an IP Resource](#), create and extend an IP resource for each load balancer frontend IP address using the following parameters:

**!** The use of 255.255.255.255 as the value of the Netmask parameter (both during create and extend) is important. Be careful not to select “Accept Defaults” during the extend process as this will typically cause 255.255.255.0 to be incorrectly used as the network mask on the standby node, which may lead to unexpected networking issues.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent
Server	node-a
IP Resource	<Frontend IP Address>
Netmask	255.255.255.255
Network Interface	<Network Interface>
IP Resource Tag	<Resource Tag>
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent
Template Priority	1
Target Priority	10
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	<Frontend IP Address>
Netmask	255.255.255.255
Network Interface	<Network Interface>
IP Resource Tag	<Resource Tag>

Now that IP resources have been created for each load balancer frontend IP address, we may now create the GenLB resources that will respond to the load balancer health check probes.

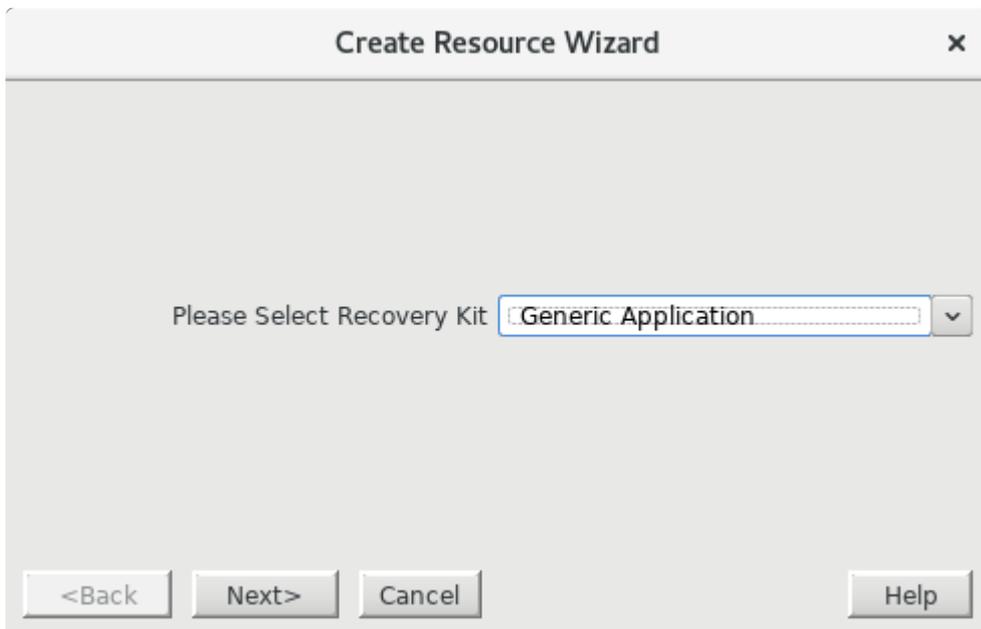
# Create a GenLB Resource

In this example we will create a sample GenLB resource on server node-a listening on TCP port 54321.

**!** Make sure to modify the server, port, and/or resource tag name as appropriate for the load balancer health check that this resource is responding to. For example, if you have set up a health check which is probing TCP port 41098 on each backend server, then the GenLB resource will contain “41098” in its application info field rather than the example value “54321” given here. When choosing a port to use for a GenLB resource, any open port in the range 1024-65535 may be used.

**\* Note:** In order for the resource to come in-service successfully, the load balancer must be actively sending health check probes to the server during the resource creation process. This means that the load balancer and health checks must be fully configured before the GenLB resource is created. See the subsections under [Switching Between Nodes in a Cloud Environment](#) for more details on how to create internal load balancers in Microsoft Azure and Google Cloud environments.

1. In the LifeKeeper GUI, click  to open the **Create Resource Wizard**. Select the “Generic Application” Recovery Kit.



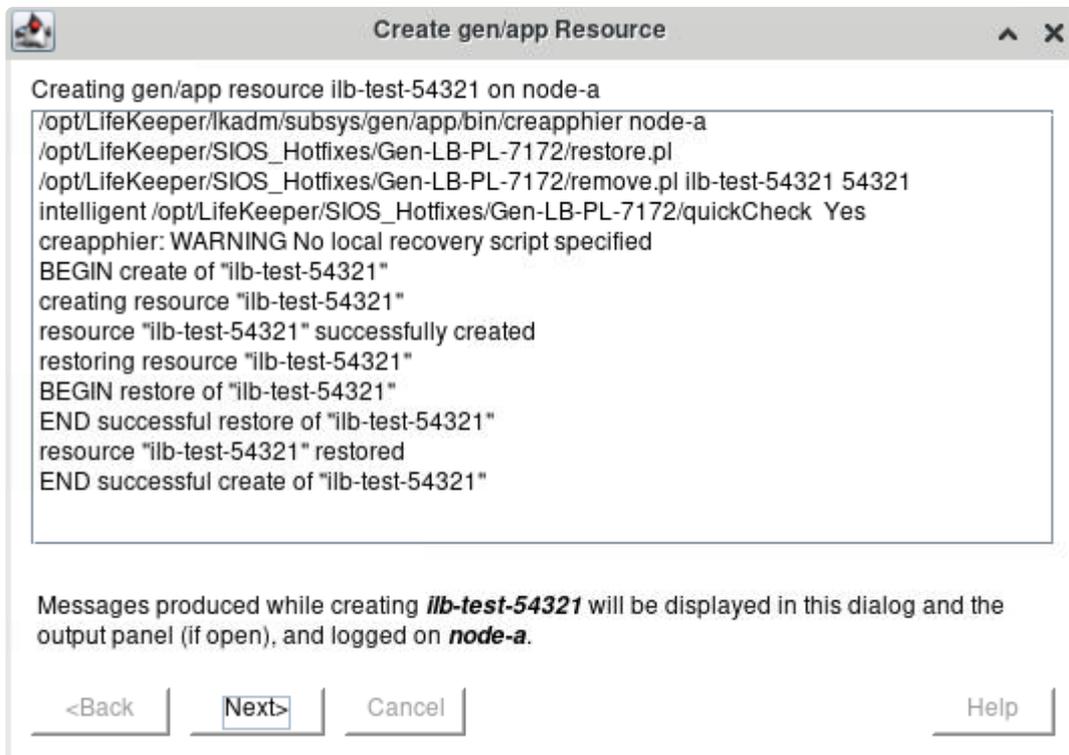
2. Enter the following values into the **Create Resource Wizard** and click **Create** when prompted.

The  icon indicates that the default option is chosen. The “Application Info” field has the format “<TCP Port> <Response>”. The response is optional, and no whitespace is allowed in the response text.

Field	Value
-------	-------

Switchback Type	Intelligent 
Server	node-a
Restore Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/restore.pl
Remove Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/remove.pl
QuickCheck Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/quickCheck ( <b>Note:</b> Although the quickCheck script may be optional for some 'Generic Application' resource types, it is required for GenLB resources.)
Local Recovery Script	None (Empty)
Application Info	54321
Bring Resource In Service	Yes
Resource Tag	ilb-test-54321

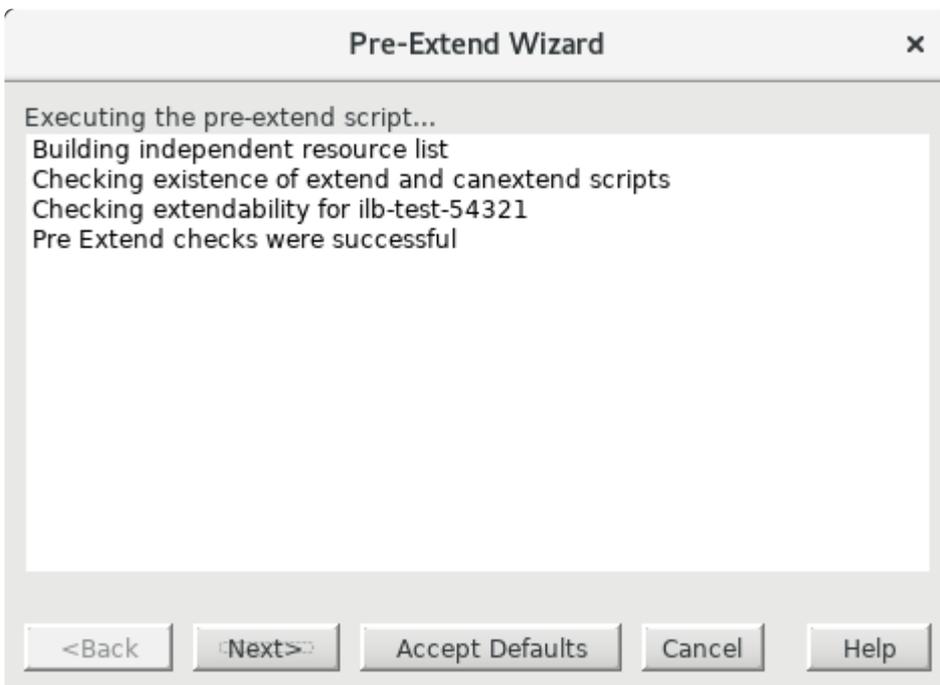
Once the resource has been created and brought in-service successfully, click **Next>** to proceed to the **Pre-Extend Wizard**.



3. Enter the following values into the **Pre-Extend Wizard**. The  icon indicates that the default option is chosen.

Field	Value
Target Server	node-b
Switchback Type	Intelligent 
Template Priority	1 
Target Priority	10 

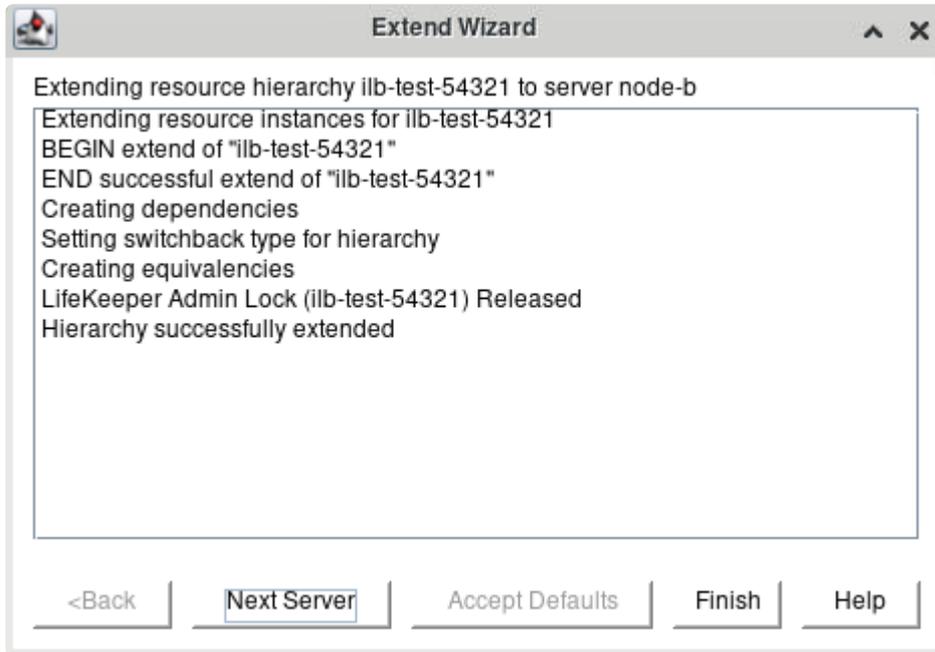
Once the pre-extend checks have passed, click **Next>** to proceed to the **Extend gen/app Resource Hierarchy Wizard**.



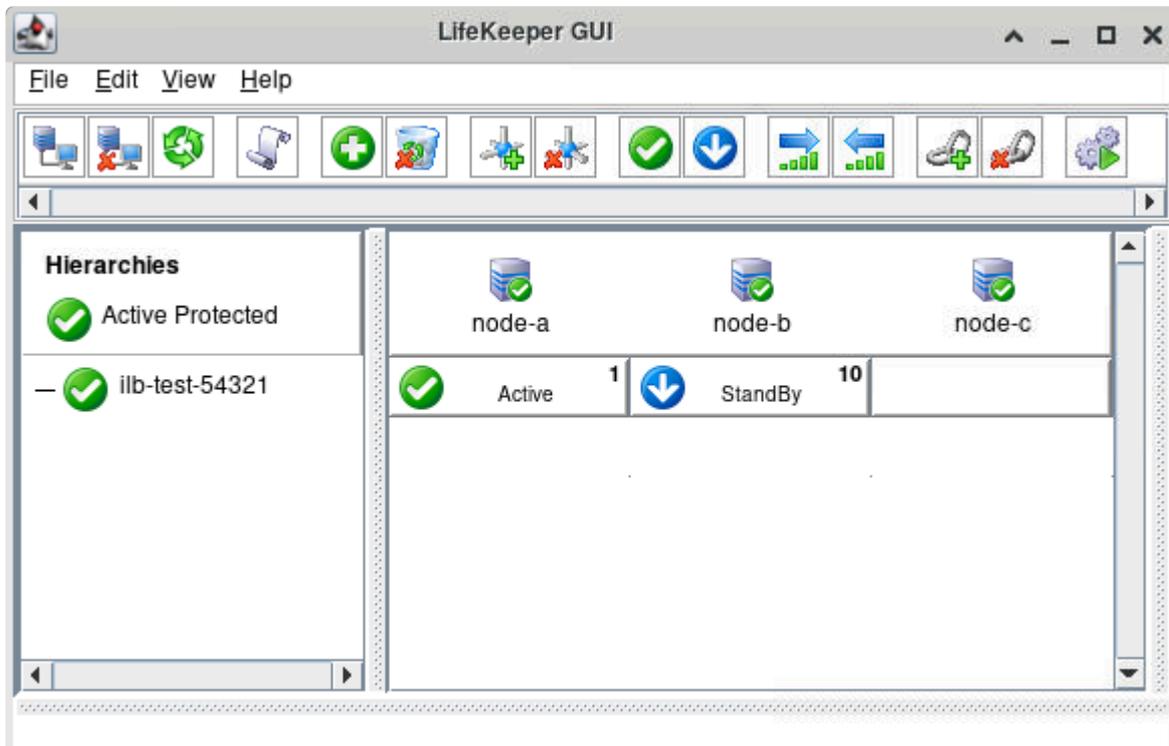
4. Enter the following values into the **Extend gen/app Resource Hierarchy Wizard** and click **Extend** when prompted. The  icon indicates that the default option is chosen.

Field	Value
Resource Tag	ilb-test-54321 
Application Info	54321 

Once the resource has been extended successfully, click **Finish**.



- 5. Back in the LifeKeeper GUI, we see that the newly created **ilb-test-54321** resource is Active on node-a and Standby on node-b. In this state, a TCP load balancer with a TCP health check on port 54321 will treat node-a as healthy and node-b as unhealthy, causing all load balancer traffic to be routed to node-a. When placed in a resource hierarchy with a protected application, this resource will ensure that load balancer traffic is always routed to the server on which the application is currently running.



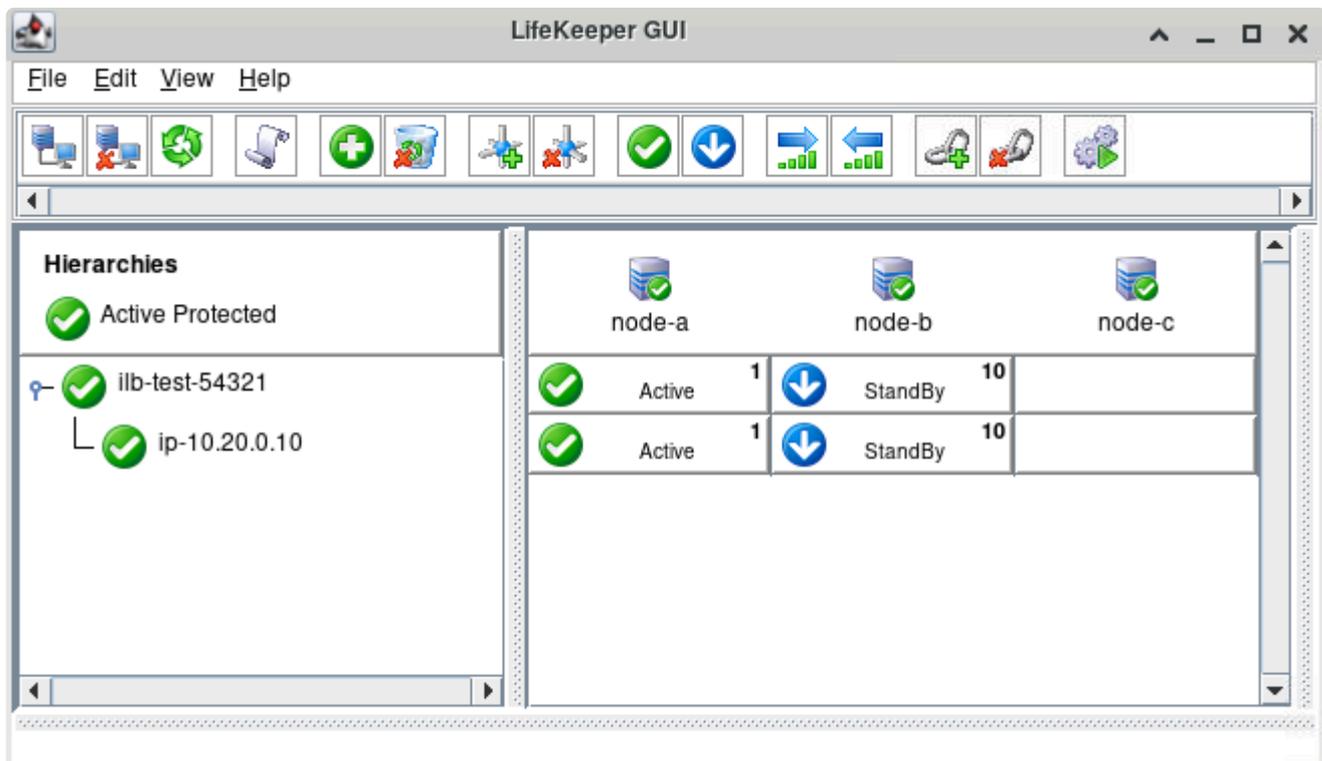
# Add Frontend IP Resources as Dependencies of GenLB Resource

**Note:** This section only applies when using a Google Cloud internal load balancer where IP forwarding is disabled. For other configurations, you may proceed directly to the 'Test GenLB Resource Switchover and Failover' section below.

When using a Google Cloud load balancer with IP forwarding disabled (see the 'Disable IP Forwarding' section of [Google Cloud – Using an Internal Load Balancer](#)), IP resource(s) protecting the IP address(es) associated to the load balancer frontend(s) (using network mask 255.255.255.255) must be added as dependencies of the GenLB resource. Complete the following steps for each of the IP resources created in the 'Create Frontend IP Resource(s)' section above.

1. Right-click on the **ilb-test-54321** resource and select **Create Dependency...** from the drop-down menu.
2. For **Child Resource Tag**, specify the resource protecting the frontend IP address of the load balancer.
3. Click **Next>** to continue, then click **Create Dependency** to create the dependency.

Once the IP resource has been added as a dependency, the resulting hierarchy will look similar to the following:



## Test GenLB Resource Switchover and Failover

In this section we will assume that we have created an internal load balancer with node-a and node-b as backend targets which has the following properties:

- Front-end internal IP: 10.20.0.10
- TCP health check on port 54321

and that the **ilb-test-54321** GenLB resource that was created in the previous section is currently Active on node-a.

For convenience we will set up a temporary Apache web server that will simply return the hostname of each server. Execute the following commands **on both node-a and node-b**. Adjust the commands accordingly (e.g., to use `zypper install`) if installing on a SLES server.

```
# yum install -y httpd
# systemctl start httpd
# echo $(hostname) > /var/www/html/index.html
```

Before continuing, verify that traffic is allowed on TCP port 80 for node-a and node-b.

We will now test the switchover and failover capabilities of the **ilb-test-54321** GenLB resource.

1. With the **ilb-test-54321** resource Active on node-a and Standby on node-b, verify the output of the following command on each server.

```
[root@node-a ~]# curl http://10.20.0.10
node-a
[root@node-b ~]# curl http://10.20.0.10
node-a
```



**Note:** If the `curl` command returns 'node-b' when run on node-b and the servers are located in Google Cloud, see the **Disable IP Forwarding** section of [Google Cloud – Using an Internal Load Balancer](#). Once the steps from that section have been completed, return here to verify the expected behavior.

2. Execute the following command on node-a:

```
[root@node-a ~]# while true; do curl http://10.20.0.10; sleep 1; done
```

and initiate a switchover of the **ilb-test-54321** resource to node-b. Once the switchover has completed successfully, use Ctrl-C (SIGINT) to terminate the running command on node-a.

The output of the command should be similar to:

```

...
node-a
node-a
node-a
[switchover occurs]
node-b
node-b
node-b
...

```

In particular, the load balancer should cleanly stop routing traffic to node-a before beginning to route it to node-b. If the output near the switchover point looks like the following:

```

...
node-a
[switchover occurs]
node-b
node-a
node-b
node-a
node-b
node-a
node-b
node-b
node-b
node-b
...

```

then you may need to edit the health check properties to decrease the time between health check probes and/or decrease the minimum number of unsuccessful health check probes before a backend instance is marked unhealthy and removed from the load balancer pool. See the **Tuning Load Balancer Health Check Parameters** section below for more details.

3. With the **ilb-test-54321** resource Active on node-b, execute the following command on node-a:

```
[root@node-a ~]# while true; do curl http://10.20.0.10; sleep 1; done
```

and forcefully reboot node-b to initiate a failover of the **ilb-test-54321** resource back to node-a:

```
[root@node-b ~]# echo b > /proc/sysrq-trigger
```

After the failover has completed successfully, use Ctrl-C (SIGINT) to terminate the running command on node-a.

The output of the command on node-a should be similar to:

```

...
node-b

```

```
node-b
node-b
[failover occurs]
node-a
node-a
node-a
...
```

At this point basic verification of the GenLB resource behavior is complete. Execute additional tests as necessary to verify the interaction between the GenLB resource and your protected application on switchover and failover. Once finished testing the GenLB resource functionality, the temporary Apache web servers may be removed by executing the following commands on both node-a and node-b:

```
# systemctl stop httpd
# rm -f /var/www/html/index.html
# yum remove -y httpd
```

## Tuning Load Balancer Health Check Parameters

While the default load balancer health check parameters should work in most common situations, it may be necessary to tune them in order to achieve the desired switchover behavior. There are two typical issues that might require a user to tune these parameters:

1. The values are set too low, causing the load balancer to be too sensitive to temporary resource constraints or network interruptions.
2. The values are set too high, causing the previous resource host to still be marked as healthy as the load balancer begins routing traffic to the new resource host during a switchover.

There are typically four primary health check parameters that may be tuned in a cloud load balancing environment:

- **Health Check Interval** – How often the health check servers send health check probes to the backend target VMs.
- **Timeout** – How long a health check server will wait to receive a response before considering the health probe failed.
- **Healthy Threshold** – The number of consecutive health check probes that must receive successful responses in order for a backend target VM to be marked as healthy.
- **Unhealthy Threshold** – The number of consecutive health check probes that must fail in order for a backend target VM to be marked as unhealthy.

See [Microsoft Azure – Load Balancer Health Probes](#) and [Google Cloud – Health Checks Overview](#) for more details about these parameters.

From these parameters, we can also derive the following values:

- **Total Time to Mark Healthy = Health Check Interval × (Healthy Threshold – 1)** – The total amount of time after the initial health check probe of a healthy server before it is marked healthy by the load balancer, assuming low network latency between the server and the health probe servers.
- **Total Time to Mark Unhealthy = Timeout × Unhealthy Threshold** – The total amount of time after the initial health check probe of a failed server before it is marked unhealthy by the load balancer.

While the exact values for these parameters will vary depending on each user's particular environment, some general guidelines are given below.

- Tuning either the Timeout or Unhealthy Threshold too low may result in a load balancer configuration which is not resilient against temporary VM resource constraints or transient network issues. For example, setting extreme values such as **Timeout = 1 second** and **Unhealthy Threshold = 1 failure** would cause a backend VM to be marked unhealthy even if the network became unresponsive for only a few seconds, which is not uncommon in cloud environments. It is recommended to leave the Timeout at a reasonable value (e.g., 5 seconds) to allow the load balancer configuration to be more resilient against minor transient issues.
- If the combination of Health Check Interval and Healthy Threshold is set too high, the load balancer may take an unnecessarily long time to begin routing traffic to the new resource host after a switchover or failover, prolonging the recovery time for the application. Setting the value **Healthy Threshold = 2 consecutive successes** should be appropriate for most situations.
- If the combination of Timeout and Unhealthy Threshold is set too high, the load balancer will not react quickly enough to mark the previous resource host node as unhealthy after a switchover, and there may be a period where load balancer traffic is being routed to both the active and standby servers in a round-robin format. In order to avoid this situation, it is recommended to gather data from several switchovers of the resource hierarchy to determine the minimum time between the GenLB resource being taken out-of-service on the previous host and being brought in-service on the new host. Assuming that the time is synchronized between the cluster servers, this can be found by inspecting `/var/log/lifekeeper.log` on each server and determining the amount of time between the end of the GenLB **remove** script on the previous host and the end of the GenLB **restore** script on the new host. We will denote this amount of time (in seconds) as **Minimum GenLB Switchover Time**. The recommendation then is to set the load balancer health check parameters such that:

**Total Time to Mark Unhealthy < Minimum GenLB Switchover Time + Total Time to Mark Healthy**

In the common situation where the Timeout value is chosen to be the same as the Health Check Interval (e.g., both are set to 5 seconds), this recommendation is equivalent to:

**Unhealthy Threshold < (Minimum GenLB Switchover Time / Timeout) + Healthy Threshold – 1**

As an example, suppose that we have configured our load balancer health check parameters with **Health Check Interval = Timeout = 5 seconds** and **Healthy Threshold = 2 consecutive successes**.

By gathering data from repeated switchover tests with our resource hierarchy, we find empirically that **Minimum GenLB Switchover Time = 20 seconds**. Using the recommendation given above, we should choose Unhealthy Threshold so that:

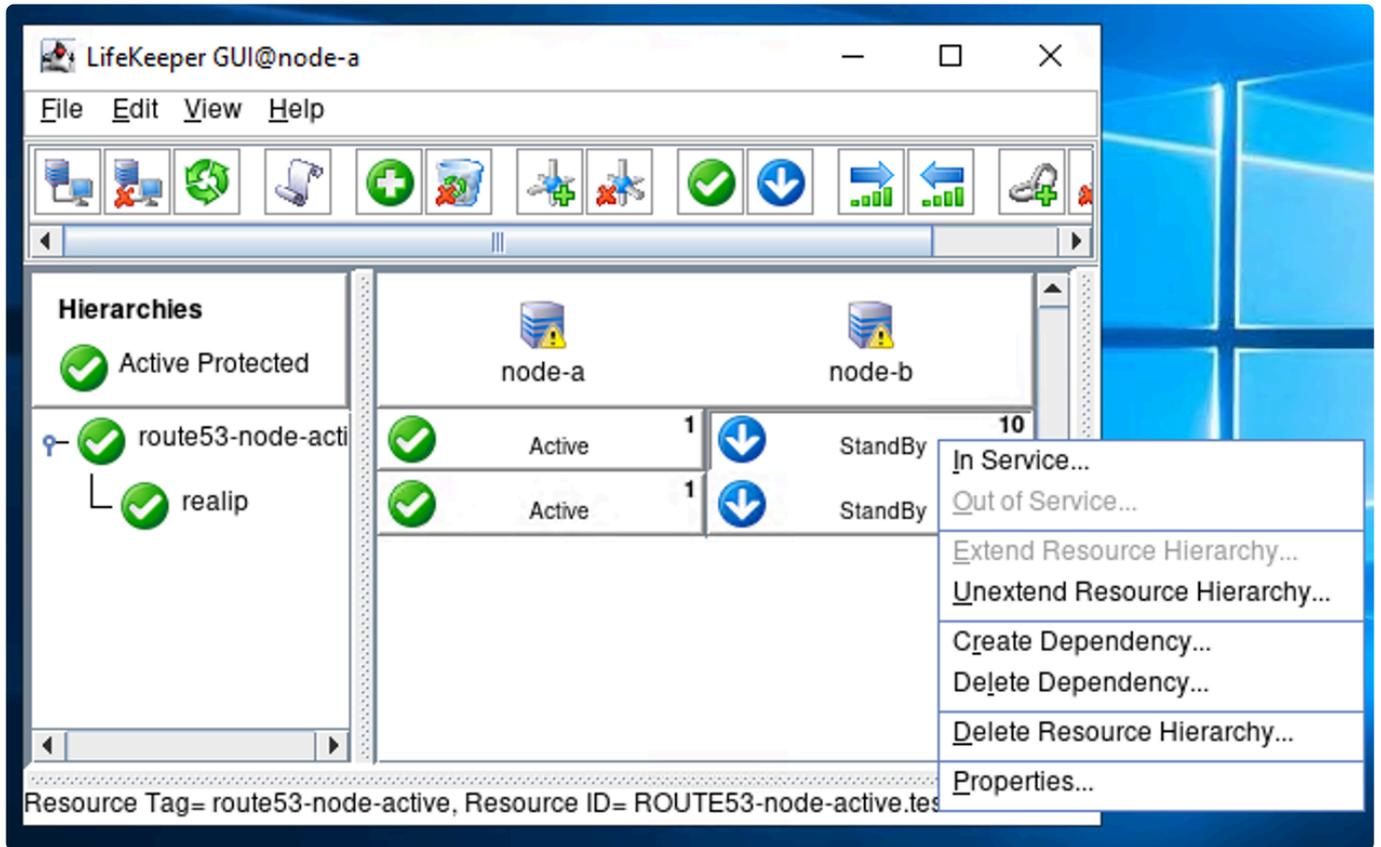
$$\text{Unhealthy Threshold} < (20 / 5) + 2 - 1 = 5$$

Based on this, any value from 2 to 4 consecutive failures could be a reasonable choice for Unhealthy Threshold. If we find during switchover testing that the load balancer is not marking the previous host as unhealthy quickly enough after switchover then we would select a lower value (e.g., 2 or 3). If we find during testing that the load balancer regularly marks the current resource host as unhealthy due to transient server or network issues, then we would select a higher value (e.g., 3 or 4). Choosing **Unhealthy Threshold = 3 consecutive failures** could be a reasonable compromise in this example.

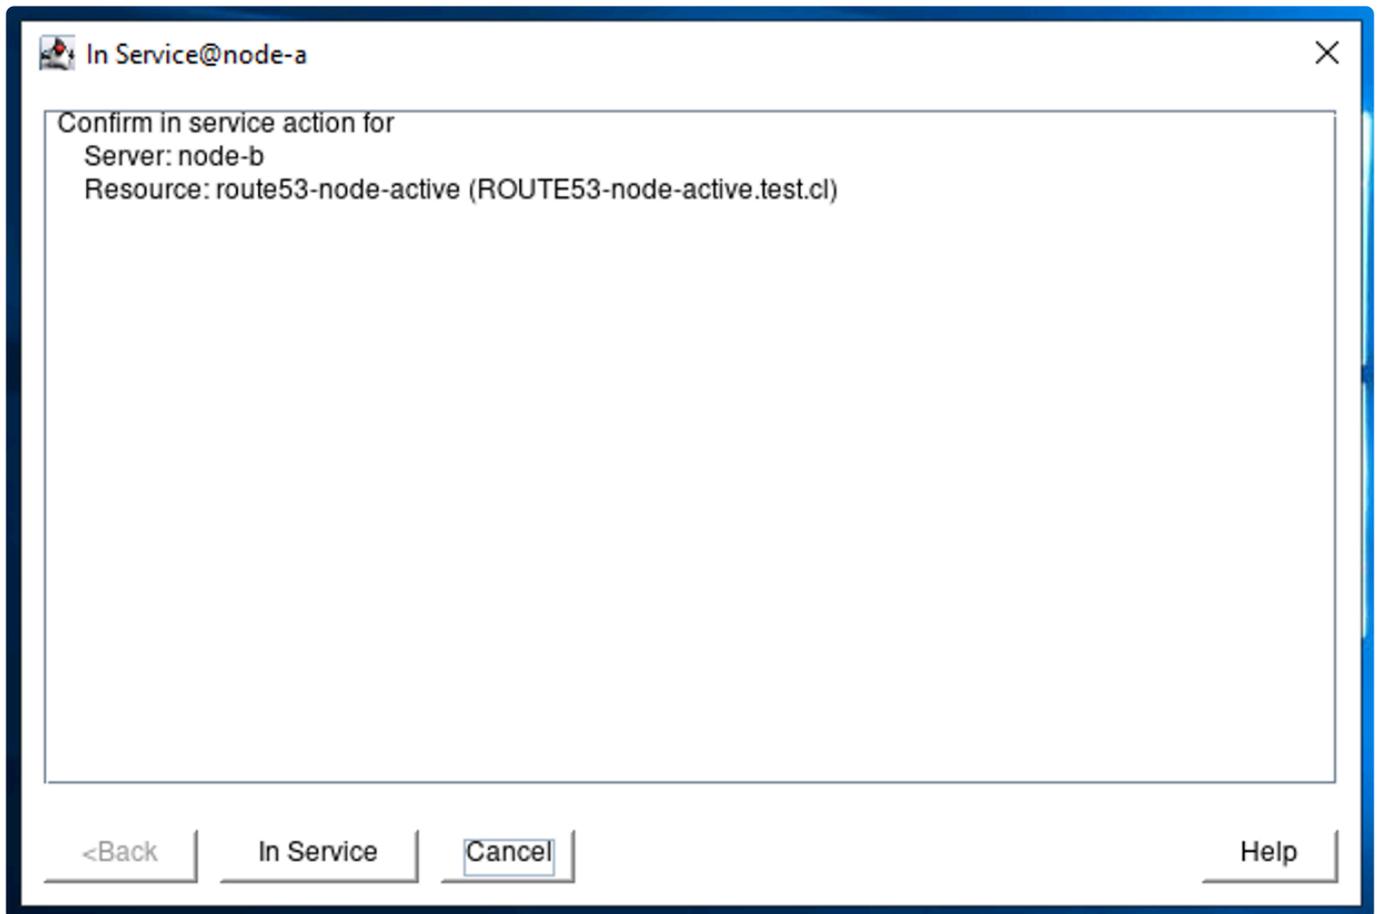
# 11.2.7.3. Switch to Standby Node to Confirm Switchover is Working

It is recommended to test each cluster resource as it is created. This helps to identify misconfigurations or other issues early in the cluster creation process.

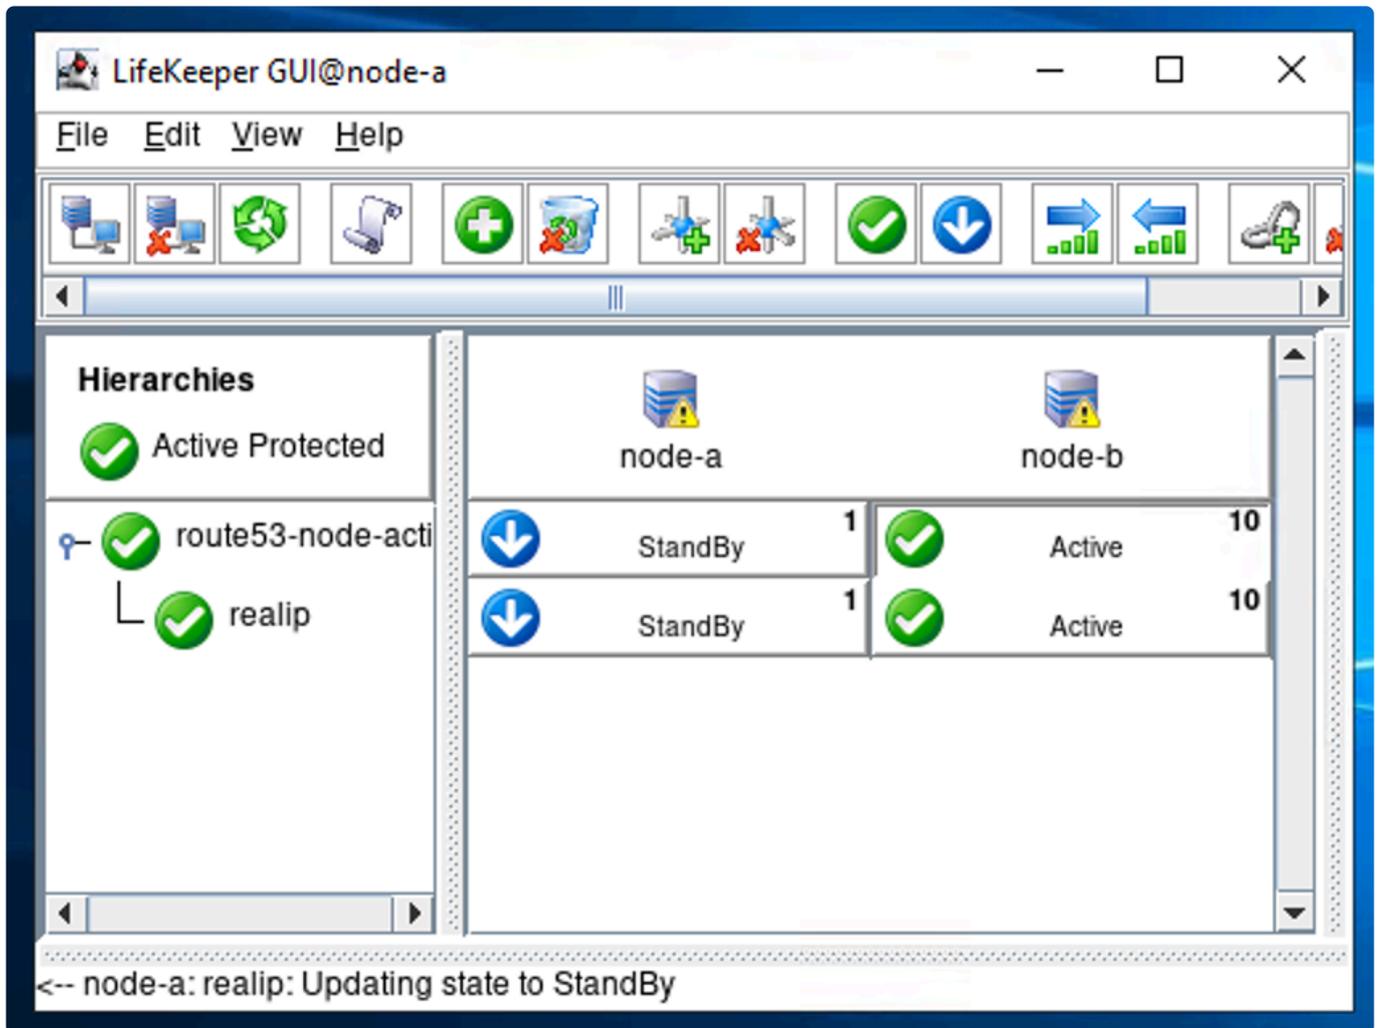
In the LifeKeeper User Interface, right-click a resource on the standby node and select “In Service...” from the context menu. This forces a switchover to the standby node.



Once the switchover is confirmed, LifeKeeper will switch the resource (including any dependent resources) to the standby node.



Once the switchover has completed successfully, node-b will be listed as the active node.



\* Switch back to node-a once you confirm that the switchover works.

## 11.2.7.4. How to Create Data Replication of a File System

As discussed in [How Does Data Replication between Nodes Work](#), DataKeeper creates a NetRAID Device which works as RAID1 device.

This guide uses the following parameters as examples. Replace these parameters based on the local environment that DataKeeper is being installed on.

Item	Value
Disk Device Name for Replication	/dev/xvdb (the second storage device attached)  <b>Note:</b> The Storage Device name may vary based on your environment.
Mount point	/datakeeper

1. See one the following topics to prepare a disk for replication.

- [How to Prepare Disks for Replication on AWS](#)
- [How to Prepare Disks for Replication on Azure](#)
- [How to Prepare Disks for Replication on Google Cloud](#)

2. Mount the disk on the primary node (node-a).



Be sure to check the name of the storage device to mount (such as /dev/xvdb1 or /dev/sdc1).

```
1 # mkdir /datakeeper
2 # mount /dev/xvdb1 /datakeeper
```

3. On the LifeKeeper User Interface, define a new resource. Select  to start the Create Resource Wizard @ node-a. Select "Data Replication" as the Recovery Kit.

## Create Resource Wizard (on node-a)



Please Select Recovery Kit

Data Replication



&lt;Back

Next&gt;

Cancel

Help

- Enter the following values. Select the default values.



**Note:** For items with a checkmark (  ) review the default values and use the value suggested.

Field	Value
Switchback Type	intelligent 
Server	node-a 
Hierarchy Type	Replicate Existing Filesystem

- Select Mount Point `/datakeeper`. The wizard will have already scanned the system and will show `/datakeeper` as the candidate for this field.

**Create Data Replication Resource Hierarchy (on node-a)** ✕

Existing Mount Point  ▼

Select the desired mount point to be replicated. The mount point must already be mounted.

6. The wizard asks for confirmation of the selection. Select 'Continue' to move on.

**Create Data Replication Resource Hierarchy (on node-a)** ✕

**ATTENTION!**

/datakeeper is not shareable with any other server.

Using this choice will result in a data replication hierarchy that cannot be extended to other servers to form a shared-storage configuration.

To confirm the selection of this entry press **Continue**.

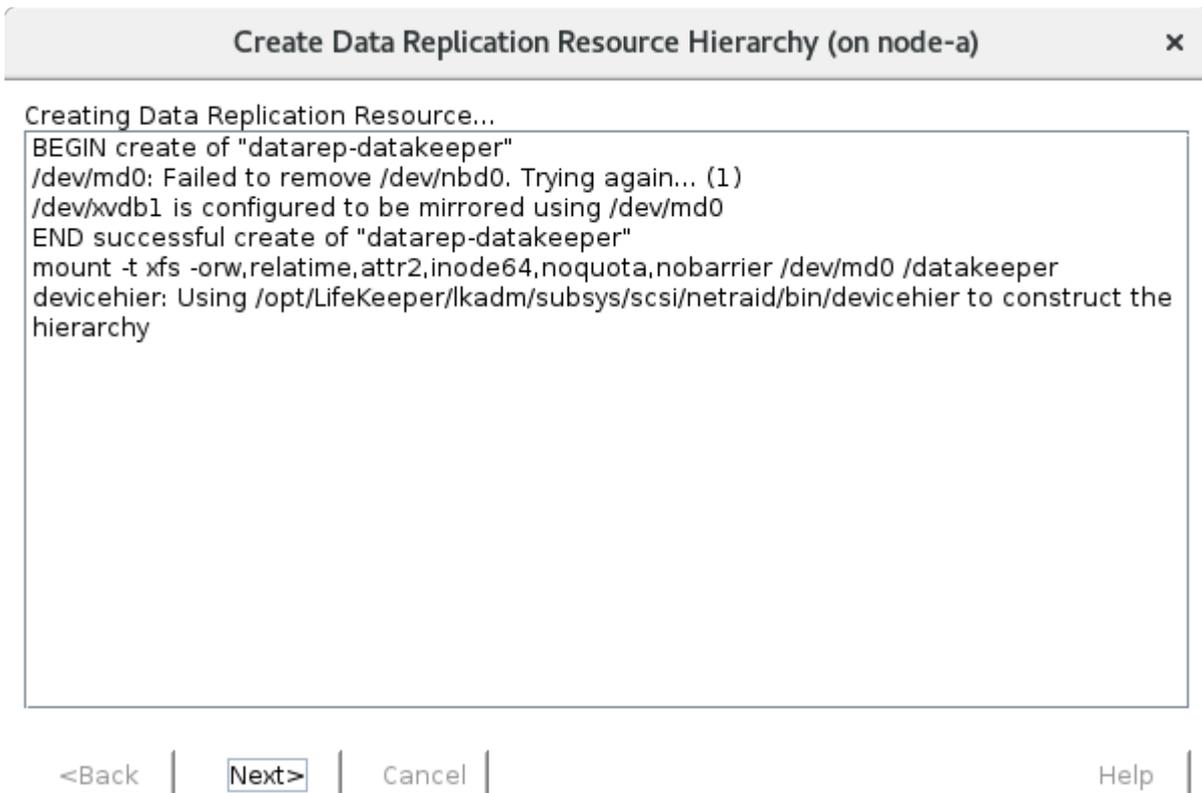
Press **Back** to select a different entry from the list.

7. Enter the following values in the fields. Select the default values.

Field	Value

Data Replication Resource Tag	datarep-datakeeper ✓
File System Resource Tag	/datakeeper ✓
Bitmap File	/opt/LifeKeeper/bitmap__datakeeper ✓
Enable Asynchronous Replication?	no ✓

8. The 'Create Data Replication Resource Hierarchy' wizard displays these values. Select "Next" to continue.



9. The next step is the "Pre-Extend Wizard". Select the default values.

Pre-Extend Wizard (on node-a) x

Target Server

You have successfully created the resource hierarchy datarep-datakeeper on node-a. Select a target server to which the hierarchy will be extended.

If you cancel before extending datarep-datakeeper to at least one other server, LifeKeeper will provide no protection for the applications in the hierarchy.

Field	Value
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 

Once the checks in the 'Pre-Extend Wizard' have passed, select 'Next' to continue.

- On the "Extend gen/filesys Resource Hierarchy" wizard, the first choice is selecting the disk as the standby node. Select the disk from the dropdown list.

Extend Data Replication Resource (on node-a) ✕

Template Server: node-a  
 Tag to Extend: datarep-datakeeper  
 Target Server: node-b

Target Disk /dev/xvdb1 (20.0 GB) ▼

Select a disk on **node-b**. The selection must not be mounted and must be at least as large as the source disk on **node-a**.

<Back
Next>
Accept Defaults
Cancel
Help

11. On the following pages, select the default values.

Field	Value
Mount point	/datakeeper ✓
Root Tag	/datakeeper ✓
Target Disk	/dev/xvdb1 (20G)
Data Replication Resource Tag	datarep-datakeeper ✓
Bitmap File	/opt/LifeKeeper/bitmap__datakeeper ✓
Replication Path	10.20.1.10/10.20.2.10

Once the 'Extend Wizard' successfully completes, select 'Finish'.

Now the disk is configured to replicate the data. As seen below, a full initial sync of the data is required. Once the data is synced, the label will change to "Target".

**LifeKeeper GUI (on node-a)**

File Edit View Help

Hierarchies		node-a		node-b	
✓ Active Protected		✓			
✓ /datakeeper		✓ Active	1	↓ StandBy	10
└─ ✓ datarep-datakeeper		Source	1	Resyncing	10
✓ ip-10.10.10.10		✓ Active	1	↓ StandBy	10
└─ ✓ ec2-10.10.10.10		✓ Active	1	↓ StandBy	10

## 11.2.7.4.1. How to Prepare Disks for Replication on AWS

\* The following steps must be performed on each node.

Before setting up disk replication using DataKeeper, the disks must be prepared for replication. Please note that this tutorial assumes the availability of an empty disk for each node, both of equal size.

This guide uses the following parameters as examples. You will need to replace these parameters based on the local environment DataKeeper is being installed on.

Items	Value
Disk Device Name for Replication	/dev/xvdb (as the second disk)
Mount point	/datakeeper

### Create /datakeeper Folder

```
1 # mkdir /datakeeper
```

\* **Note:** After performing `mkdir /datakeeper`, **change ownership to oracle.oinstall** so that Oracle setup can actually create database files. **Run command:** `sudo chown oracle.oinstall /datakeeper`

### Check the Available Disks on node-a

Confirm the disk to mount using `lsblk`. As seen in the screenshot below, the storage device `xvdb` is not mounted to the instance yet.

```
1 [root@node-a media]# lsblk
2 NAME      MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
3 xvda      202:0    0    10G  0 disk
4 └─xvda1   202:1    0     1M  0 part
5 └─xvda2   202:2    0    10G  0 part /
6 xvdb      202:16   0     8G   0 disk
7 loop0     7:0      0 376.3M 0 loop /media
```

## Use `gdisk` to Create a Partition

Use `gdisk /dev/xvdb` to create a partition on this disk.

Type these commands in the following order:

- `n` to create a new partition
- `<Enter>` to select default (Partition number)
- `<Enter>` to select default (First Sector)
- `<Enter>` to select default (Last Sector)
- `<Enter>` to select default (HEX code)
- `w` to write down the changes
- `y` to confirm the operation

```
1 # gdisk /dev/xvdb
2 GPT fdisk (gdisk) version 0.8.10
3
4 Partition table scan:
5   MBR: not present
6   BSD: not present
7   APM: not present
8   GPT: not present
9
10 Creating new GPT entries.
11
12 Command (? for help): n
13 Partition number (1-128, default 1):
14 First sector (34-41943006, default = 2048) or {+}size{KMGTP}:
15 Last sector (2048-41943006, default = 41943006) or {+}size{KMGTP}:
16 Current type is 'Linux filesystem'
17 Hex code or GUID (L to show codes, Enter = 8300):
18 Changed type of partition to 'Linux filesystem'
19
20 Command (? for help): w
21
22 Final checks complete. About to write GPT data. THIS WILL OVERWRITE EXISTING
23 PARTITIONS!!
24
25 Do you want to proceed? (Y/N): Y
26 OK; writing new GUID partition table (GPT) to /dev/xvdb.
27 The operation has completed successfully.
```

## Create a File System to the Partition using `mkfs`

```
1 # mkfs -t xfs /dev/xvdb1
2 meta-data=/dev/xvdb1          isize=512    agcount=4, agsize=1310655 blks
3      =                               sectsz=512   attr=2, projid32bit=1
4      =                               crc=1      finobt=0, sparse=0
5 data      =                       bsize=4096  blocks=5242619, imaxpct=25
6      =                               sunit=0     swidth=0 blks
7 naming    =version 2            bsize=4096  ascii-ci=0 ftype=1
8 log       =internal log        bsize=4096  blocks=2560, version=2
9      =                               sectsz=512   sunit=0 blks, lazy-count=1
10 realtime =none                 extsz=4096  blocks=0, rtextents=0
```

## 11.2.7.4.2. How to Prepare Disks for Replication on Azure

If working in Azure, review the following document describing how to attach a second disk to an instance.

<https://docs.microsoft.com/en-us/azure/virtual-machines/linux/attach-disk-portal>

### Create /datakeeper Folder

```
1 # mkdir /datakeeper
```

\* **Note:** After performing `mkdir /datakeeper`, **change ownership to oracle.oinstall** so that Oracle setup can actually create database files. **Run command:** `sudo chown oracle.oinstall /datakeeper`

### Check the Available Disks on node-a

Confirm the disk to mount using `lsblk`. As seen in the screenshot below, the storage device `sdc` is not mounted to the instance yet.

```
1 [root@node-a ~]# lsblk
2 NAME                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
3 sda                  8:0    0   64G  0 disk
4 └─sda1               8:1    0   500M  0 part /boot/efi
5 └─sda2               8:2    0   500M  0 part /boot
6 └─sda3               8:3    0     2M  0 part
7 └─sda4               8:4    0    63G  0 part
8   ├─rootvg-tmplv    253:0    0     2G  0 lvm  /tmp
9   ├─rootvg-usrlv    253:1    0    10G  0 lvm  /usr
10  ├─rootvg-optlv     253:2    0     2G  0 lvm  /opt
11  ├─rootvg-homelv    253:3    0     1G  0 lvm  /home
12  ├─rootvg-varlv     253:4    0     8G  0 lvm  /var
13  └─rootvg-rootlv    253:5    0     2G  0 lvm  /
14 sdb                  8:16    0     4G  0 disk
15 └─sdb1               8:17    0     4G  0 part /mnt/resource
16 sdc                  8:32    0     8G  0 disk
```

 Follow the Azure document above to mount it to the instance.

Use `parted` to prepare the partition, then use `mkfs` to prepare the file system.

```

1 [root@node-a ~]# parted /dev/sdc --script mklabel gpt mkpart xfspart xfs 0% 100%
2 [root@node-a ~]# mkfs.xfs /dev/sdc1
3 meta-data=/dev/sdc1          isize=512    agcount=4, agsize=524160 blks
4      =                       sectsz=4096  attr=2, projid32bit=1
5      =                       crc=1        finobt=0, sparse=0
6 data      =                   bsize=4096  blocks=2096640, imaxpct=25
7      =                       sunit=0       swidth=0 blks
8 naming   =version 2          bsize=4096  ascii-ci=0 ftype=1
9 log      =internal log      bsize=4096  blocks=2560, version=2
10     =                       sectsz=4096  sunit=1 blks, lazy-count=1
11 realtime =none              extsz=4096  blocks=0, rtextents=0
12 [root@node-a ~]# partprobe /dev/sdc1

```

Run `lsblk` again, now `/dev/sdc1` should be available to mount.

```

1 [root@node-a ~]# lsblk
2 NAME                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
3 sda                  8:0    0   64G  0 disk
4 └─sda1               8:1    0   500M  0 part /boot/efi
5 └─sda2               8:2    0   500M  0 part /boot
6 └─sda3               8:3    0     2M  0 part
7 └─sda4               8:4    0    63G  0 part
8   └─rootvg-tmplv    253:0    0     2G  0 lvm  /tmp
9   └─rootvg-usrlv    253:1    0    10G  0 lvm  /usr
10  └─rootvg-optlv     253:2    0     2G  0 lvm  /opt
11  └─rootvg-homelv    253:3    0     1G  0 lvm  /home
12  └─rootvg-varlv     253:4    0     8G  0 lvm  /var
13  └─rootvg-rootlv    253:5    0     2G  0 lvm  /
14 sdb                  8:16   0     4G  0 disk
15 └─sdb1              8:17   0     4G  0 part /mnt/resource
16 sdc                  8:32   0     8G  0 disk
17 └─sdc1              8:33   0     8G  0 part

```

Now you are ready to create the Data Replication Resource using LifeKeeper / DataKeeper.

## 11.2.7.4.3. How to Prepare Disks for Replication on Google Cloud

### Create /datakeeper Folder

```
1 # mkdir /datakeeper
```

\* **Note:** After performing `mkdir /datakeeper`, **change ownership to oracle.oinstall** so that Oracle setup can actually create database files. **Run command:** `sudo chown oracle.oinstall /datakeeper`

### Check the Available Disks on node-a

Confirm the disk to mount using `lsblk`. As seen in the screenshot below, the storage device `sdc` is not mounted to the instance yet.

```
1 [root@node-a ~]# lsblk
2 NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
3 sda 8:0 0 20G 0 disk
4 └─sda1 8:1 0 200M 0 part /boot/efi
5 └─sda2 8:2 0 19.8G 0 part /
6 sdb 8:16 0 10G 0 disk
```

\* Follow the Azure document above to mount it to the instance.

Use `parted` to prepare the partition, then use `mkfs` to prepare the file system.

```

1 [root@node-a ~]# parted /dev/sdb --script mklabel gpt mkpart xfs xfs 0% 100%
2 [root@node-a ~]# mkfs.xfs /dev/sdb1
3 meta-data=/dev/sdb1          isize=512    agcount=4, agsize=655232 blks
4      =                       sectsz=4096  attr=2, projid32bit=1
5      =                       crc=1        finobt=0, sparse=0
6 data      =                   bsize=4096  blocks=2620928, imaxpct=25
7      =                       sunit=0       swidth=0 blks
8 naming   =version 2          bsize=4096  ascii-ci=0 ftype=1
9 log      =internal log      bsize=4096  blocks=2560, version=2
10     =                       sectsz=4096  sunit=1 blks, lazy-count=1
11 realtime =none              extsz=4096  blocks=0, rtextents=0
12 [root@node-a ~]# partprobe /dev/sdb1

```

Run `lsblk` again, now `/dev/sdb1` should now be available to mount.

```

1 [root@node-a ~]# lsblk
2 NAME MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
3 sda   8:0    0   20G  0 disk
4 └─sda1 8:1    0   200M  0 part /boot/efi
5 └─sda2 8:2    0  19.8G  0 part /
6 sdb   8:16   0   10G  0 disk
7 └─sdb1 8:17   0   10G  0 part

```

Now you are ready to create the Data Replication Resource using LifeKeeper / DataKeeper.

## 11.2.7.5. How to Protect Other Resources (Databases or Applications)

---

This section includes the basic steps to create various resources such as databases and applications. Each resource can be created independently (unless noted otherwise).

- [Protecting an Oracle Resource](#)
- [Protecting MSSQL Using Quick Service Protection](#)
- [Protecting a PostgreSQL Resource](#)
- [Protecting an NFS Resource](#)
- [Protecting an SAP Resource](#)
- [Protecting an SAP HANA Resource](#)

\* This guide uses different types of IP Resources to demonstrate different patterns. Different IP resource tags may be seen including:

- ip-10.10.10.10
- ip-10.20.0.10
- ip-10.20.10.100
- realip

Although different tags (based on how IP resources are configured) may be seen, these differences don't affect how each application or service (such as Oracle) are protected.

# 11.2.7.5.1. Protecting an Oracle Resource (non-PDB)

This section outlines the steps to protect Oracle Resources (non-PDB).

To protect Oracle resources, the data needs to be replicated across nodes using DataKeeper (unless shared storage or SAN devices are used).

The following table outlines the location of each piece.

Items	Location
Oracle 19c (ORACLE_HOME)	/u01/app/oracle/product/19.3.0/dbhome_1
Database (SID: ORCL)	/datakeeper/oradata/ORCL

As discussed earlier in [How to Create Data Replication of a File System](#), this guide uses `/datakeeper` to replicate data between nodes. Therefore, the data on `/datakeeper/oradata/ORCL` is also replicated between nodes.

Please note that this guide uses computing resources with the following devices attached (Oracle resources require more space than other resources even for evaluation purposes).

Resource	Required Size
Boot Disk (such as <code>/dev/sda/</code> )	30 Gb
Data Disk (such as <code>/dev/sdb/</code> , configured and mounted at <code>/datakeeper</code> )	20 Gb
Memory	2 Gb

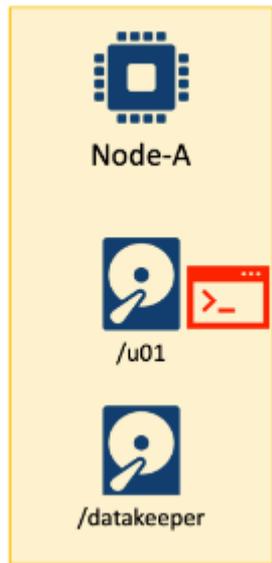
## Configure the Oracle Resource

The following table outlines the general steps to configure an Oracle Resource. The red ‘stacked disk’ shaped icons indicate a node that has Oracle Database instances at the time of each step. The grey ‘stacked disk’ icons indicate Oracle Databases that are not running.

Also, the grey “storage” icon indicates that the storage `/datakeeper` is not available for the node.

**Install Oracle 19c on both nodes**

- [Install Oracle on Your System](#)

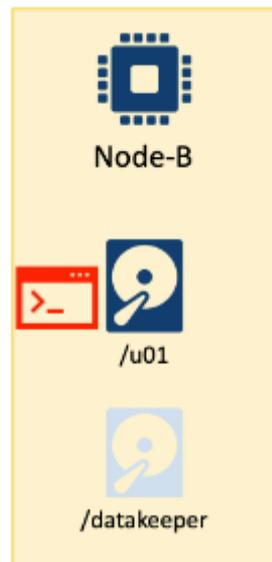


Node-A

/u01

/datakeeper

This diagram shows the configuration for Node-A. It features a central processor icon labeled 'Node-A'. Below it, a disk icon is labeled '/u01' and is accompanied by a red terminal window icon containing '>\_'. At the bottom, another disk icon is labeled '/datakeeper'.



Node-B

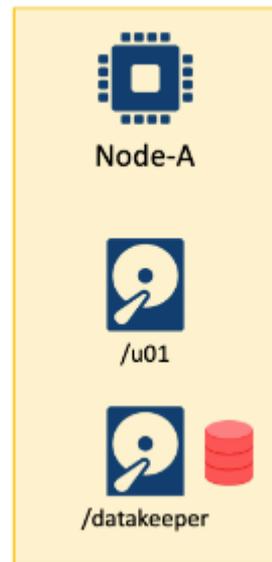
/u01

/datakeeper

This diagram shows the configuration for Node-B. It features a central processor icon labeled 'Node-B'. Below it, a disk icon is labeled '/u01' and is accompanied by a red terminal window icon containing '>\_'. At the bottom, another disk icon is labeled '/datakeeper'.

**Create Oracle Database on Node-A**

- Review [How to Confirm if the Data Storage is Available on a Node](#). The data storage should be available on node-a.
- [Create an Oracle Database](#) on node-a
- [Stop Oracle Instance](#) on node-a



Node-A

/u01

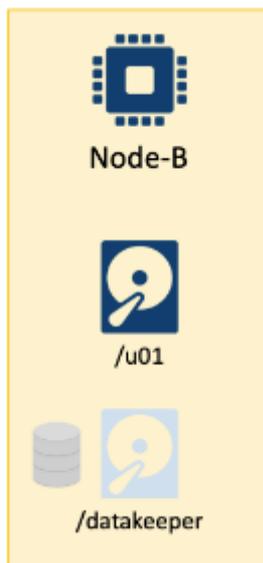
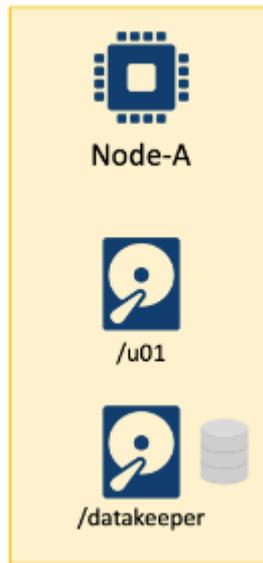
/datakeeper

This diagram shows the configuration for Node-A during database creation. It features a central processor icon labeled 'Node-A'. Below it, a disk icon is labeled '/u01'. At the bottom, another disk icon is labeled '/datakeeper' and is accompanied by a red database cylinder icon.

	 <p>Node-B</p>  <p>/u01</p>  <p>/datakeeper</p>
<p><b>Create Oracle Database on Node-B</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Switchover (the data storage)</a> to node-b</li> <li>• <a href="#">Rename /datakeeper/oradata/ORCL</a></li> <li>• <a href="#">Create an Oracle Database</a> on node-b</li> <li>• <a href="#">Stop Oracle Instance</a> on node-b</li> </ul>	 <p>Node-A</p>  <p>/u01</p>  <p>/datakeeper</p>  <p>Node-B</p>  <p>/u01</p>  <p>/datakeeper</p>

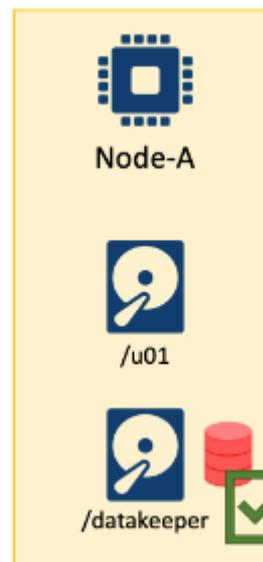
**Configure Oracle LISTENER on both nodes**

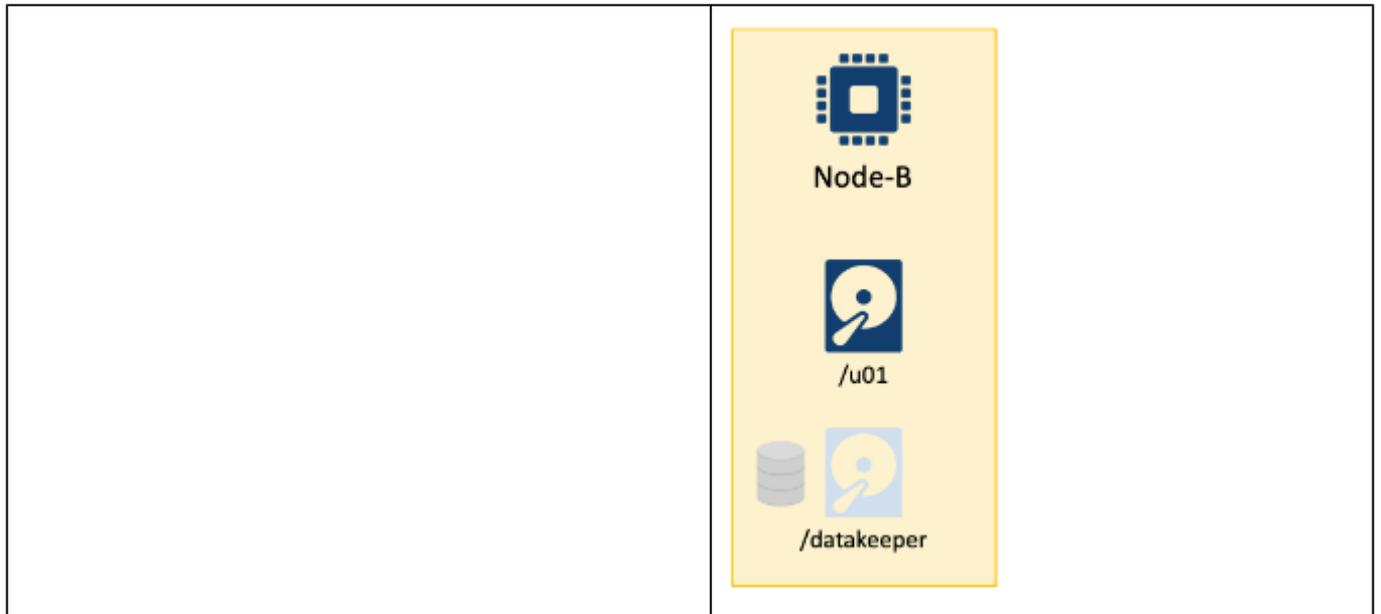
- [Switchover \(the data storage\)](#) back to node-a
- [Update Config File for Oracle Listener on Both Nodes](#)



**Complete Configuration using LifeKeeper**

- [Start Database and Listener on node-a](#)
- Configure resources on LifeKeeper
  - [Configure Oracle LISTENER Resource](#)
  - [Configure Oracle Resource](#)
- [Test Switchover of the Oracle Resource](#)





## 11.2.7.5.1.1. Install Oracle

 Install Oracle 19c on both nodes.

Install Oracle on a Linux system.

- Refer to the [Database Installation Guide for Linux](#)

Install Oracle 19c at the following location:

```
/u01/app/oracle/product/19.3.0/dbhome_1
```

The location should be specified in the environment variable `ORACLE_HOME`. Once you start the Oracle installer, specify the following values on the installer wizard pages.

Step #	Item	Value (  indicates accepting default value)
1	Install Option	Set Up Software Only
2	Type of database installation	Single Instance database installation
3	Database Edition	Standard Edition 2
4	Oracle Base	/u01/app/oracle 
5	Inventory Location	/u01/app/oraInventory 
6	Privileges Operating System Groups	
7	Root Script Execution Configuration	Enter "root" user credential
8	Perform Prerequisite Checks	Fix if the wizard indicates important issue 
9	Summary	Select "Install" 
10	Install Product	Wait for install 
11	Finish	Install should complete with no errors 

## 11.2.7.5.1.2. Create an Oracle Database (non-PDB)

\* Be sure to perform these steps on the nodes indicated in the documentation. Also, [confirm if the data storage is available on the node.](#)

1. On node-a, start installer as follows:

```

1 $ su - oracle
2 Password:
3 Last login: Sun Dec 20 04:33:22 UTC 2020 on pts/1
4 $ cd $ORACLE_HOME
5 $ ./bin/dbca

```

2. Once the installer starts, complete the following steps:

Step #	Item	Value (✓ indicates accepting default value)
1	Select Operation	Create a database ✓
2	Creation Mode	Advanced Configuration
3	Deployment Type	Database Type: Oracle Single Instance Database Template: General Purpose or Transaction Processing
4	Database Identification	orcl ✓
5	Storage Option	Storage Type: File System Location: /datakeeper/oradata
6	Fast Recovery Option	Select nothing ✓
7	Network Configuration	Create a new Listener (Name: LISTENER1)
8	Configuration Option	Leave as default values ✓

9	Management Option	Leave as default values 
10	User Credentials	Select "Use the same administrative password for all accounts" and then Specify Password
11	Creation Option	Create database 
12	Summary	Confirm the values 
13	Progress Page	Watch for progress 
14	Finish	Confirm the result 

## 11.2.7.5.1.3. Stop the Oracle Instance

\* Ensure these steps are completed on the correct node.

### Shutdown the Database

To shutdown the database, execute `sqlplus` and run the `shutdown immediate` command as follows:

```
1 [oracle]$ $ORACLE_HOME/bin/sqlplus / AS SYSDBA
2
3 SQL*Plus: Release 19.0.0.0.0 - Production on Sun Dec 20 05:50:40 2020
4 Version 19.3.0.0.0
5
6 Copyright (c) 1982, 2019, Oracle. All rights reserved.
7
8
9 Connected to:
10 Oracle Database 19c Standard Edition 2 Release 19.0.0.0.0 - Production
11 Version 19.3.0.0.0
12
13 SQL> shutdown immediate
14 Database closed.
15 Database dismounted.
16 ORACLE instance shut down.
17 SQL>
```

### Confirm the Database can be Restarted Later

The following is not required. However, it is recommended that the database be restarted using the `startup` command and then shut down with the `shutdown immediate` command.

```
1 SQL> startup
2 ORACLE instance started.
3
4 Total System Global Area 771748536 bytes
5 Fixed Size 8901304 bytes
6 Variable Size 490733568 bytes
7 Database Buffers 268435456 bytes
8 Redo Buffers 3678208 bytes
9 Database mounted.
10 Database opened.
11 SQL> shutdown immediate
12 Database closed.
13 Database dismounted.
14 ORACLE instance shut down.
15 SQL> exit
16 Disconnected from Oracle Database 19c Standard Edition 2 Release 19.0.0.0.0 - Production
17 Version 19.3.0.0.0
```

## 11.2.7.5.1.4. Rename /datakeeper/oradata/ORCL

So far the following tasks have been completed:

1. Created the database on Node-A
2. Switched over the `datarep-datakeeper` resource to Node-B

In step 1, the information about the newly created ORCL database was created in `/u01/app/oracle/product/19.3.0/dbhome_1/`.

Now we will create the same ORCL database on Node-B so the information about the ORCL database will be stored on Node-B in `/u01/app/oracle/product/19.3.0/dbhome_1/`. The database files themselves are located in `/datakeeper/oradata/ORCL`.

On Node-B, these files can be seen as shown below:

```
1 [ec2-user@node-b ~]$ su - oracle
2 Password:
3 Last login: Sun Dec 20 03:58:10 UTC 2020 on pts/0
4 [oracle@node-b ~]$ ls -al /datakeeper/oradata/ORCL/
5 total 2477920
6 drwxr-x--- 2 oracle oinstall      200 Dec 20 04:15 .
7 drwxrwxr-x 3 oracle oinstall      18 Dec 20 04:11 ..
8 -rw-r----- 1 oracle oinstall 10600448 Dec 20 05:15 control01.ctl
9 -rw-r----- 1 oracle oinstall 10600448 Dec 20 05:15 control02.ctl
10 -rw-r----- 1 oracle oinstall 209715712 Dec 20 05:15 redo01.log
11 -rw-r----- 1 oracle oinstall 209715712 Dec 20 05:15 redo02.log
12 -rw-r----- 1 oracle oinstall 209715712 Dec 20 05:15 redo03.log
13 -rw-r----- 1 oracle oinstall 555753472 Dec 20 05:15 sysaux01.dbf
14 -rw-r----- 1 oracle oinstall 943726592 Dec 20 05:15 system01.dbf
15 -rw-r----- 1 oracle oinstall 33562624 Dec 20 04:31 temp01.dbf
16 -rw-r----- 1 oracle oinstall 356524032 Dec 20 05:15 undotbs01.dbf
17 -rw-r----- 1 oracle oinstall 5251072 Dec 20 05:15 users01.dbf
```

Rename the ORCL folder so that creation of the ORCL database on Node-B can be completed successfully.

```
1 [oracle@node-b ~]$ cd /datakeeper/oradata/
2 [oracle@node-b oradata]$ mv ORCL/ ORCL_Backup/
3 [oracle@node-b oradata]$ ls -al
4 total 0
5 drwxrwxr-x 3 oracle oinstall 25 Dec 20 05:26 .
6 drwxr-xr-x 3 root   root      21 Dec 20 02:24 ..
7 drwxr-x--- 2 oracle oinstall 200 Dec 20 04:15 ORCL_Backup
```

## 11.2.7.5.1.5. Update Config File for Oracle Listener on Both Nodes

\* The following steps should be executed on both nodes.

### Check the Listener Status

Check the current status of the Listener with following command:

```
1 $ cd $ORACLE_HOME/bin/
2 $ ./lsnrctl status LISTENER1
```

### Stop Listener if it is Running

```
1 $ cd $ORACLE_HOME/bin/
2 $ ./lsnrctl stop LISTENER1
```

### Edit the Configuration File

The configuration file is located at `/u01/app/oracle/product/19.3.0/dbhome_1/network/admin/listener.ora`.

The content of this file should be similar to the example shown below:

```
# listener.ora Network Configuration File: /u01/app/oracle/product/19.3.0/dbhome_1/network/admin/listener.ora
# Generated by Oracle configuration tools.
LISTENER1 =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = node-a) (PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))
    )
  )
```

See [Creating an IP Resource](#) to confirm which IP address was previously set.

Also, an additional section for the `SID_LIST_LISTENER1` resource is needed. The updated file should look similar to the example below:

```
# listener.ora Network Configuration File: /u01/app/oracle/product/19.3.0/dbhome_1/network/admin/listener.ora
# Generated by Oracle configuration tools.
SID_LIST_LISTENER1 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = orcl)
    )
  )
LISTENER1 =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = 10.10.10.10) (PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))
    )
  )
)
```

## 11.2.7.5.1.6. Start Database and Listener on node-a

✿ The following steps should be completed on node-a.

As discussed in previous steps, start the database with `sqlplus` and the listener with `lsnrctl`.

Start both processes on node-a so that LifeKeeper can see these processes as the corresponding resources are created.

### Start the Database

```
1 [oracle@node-a ~]$ $ORACLE_HOME/bin/sqlplus / as sysdba
2
3 SQL*Plus: Release 19.0.0.0.0 - Production on Sun Dec 20 05:58:05 2020
4 Version 19.3.0.0.0
5
6 Copyright (c) 1982, 2019, Oracle. All rights reserved.
7
8 Connected to an idle instance.
9
10 SQL> startup
11 ORACLE instance started.
12
13 Total System Global Area 771748536 bytes
14 Fixed Size 8901304 bytes
15 Variable Size 490733568 bytes
16 Database Buffers 268435456 bytes
17 Redo Buffers 3678208 bytes
18 Database mounted.
19 Database opened.
20 SQL>EXIT
```

### Start the Listener

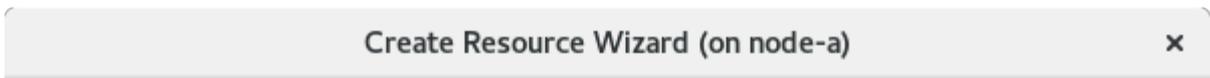
```
1 $ cd $ORACLE_HOME/bin/
2 $ ./lsnrctl start LISTENER1
```

# 11.2.7.5.1.7. Configure Oracle LISTENER Resource

As discussed in [Update Config File for Oracle LISTENER on Both Nodes](#), the name of the Listener to protect is LISTENER1.

 **Note:** If using a version **previous to 9.5.2**, do not specify the ip address for the listener.

1. In the LifeKeeper User Interface, define a new resource. Select  to start the Create Resource Wizard (on node-a). Select “Oracle Database Listener” as the Recovery Kit.



Please Select Recovery Kit Oracle Database Listener ▼

<Back | Next> | Cancel Help

2. Enter the following values. Select the default values.

Field	Value
Switchback Type	intelligent 
Server	node-a 
Listener Configuration File Path	/u01/app/oracle/product/19.3.0/dbhome_1/network/admin/listener.ora 
Listener Name(s)	LISTENER1 

Listener Executable(s)	/u01/app/oracle/product/19.3.0/dbhome_1/bin/lsnrctl 
------------------------	-----------------------------------------------------------------------------------------------------------------------------------------

3. Select Listener Protection Level. Select “Full Control” as the wizard recommends.

**Create Oracle Listener Resource (on node-a)** ✕

Listener Protection Level  ▼

Select:

- **Full Control (Start, Stop, Monitor, & Recover)** Allow this resource to start, stop, monitor, and recover the listener(s).
- **Intermediate Control (Start, Monitor, & Recover)** Allow this resource to start, monitor, and recover the specified listener(s). If Intermediate Control (Start, Monitor, & Recover) is selected, the specified listener(s) will not be stopped during remove operations.
- **Minimal Control (Start & Monitor Only)** Allow this resource to perform **only** the startup of the specified listener(s), and status logging. **If Minimal Control (Start & Monitor Only) is selected, the specified listener(s) will not be not be stopped, and will not be restarted on failures during LifeKeeper operation.**

4. Enter the following values. Select the default values.

Field	Value
Listener Recovery Level	Standard (On) 
IP address Name(s)	ip-10.10.10.10 
Listener Tag	LSNR.LISTENER1 

5. The wizard reviews all values specified. Select “Next >” to continue.

**Create Oracle Listener Resource (on node-a)** ✕

---

Creating database/listener resource...

```

BEGIN create of "LSNR.LISTENER1" on server "node-a"
BEGIN restore of "LSNR.LISTENER1" on server "node-a"
END successful restore of "LSNR.LISTENER1" on server "node-a"
END successful create of "LSNR.LISTENER1" on server "node-a"

```

<Back
Next>
Cancel
Help

6. The next step is the "Pre-Extend Wizard". Select the default values.

**Pre-Extend Wizard (on node-a)** ✕

---

Target Server node-b ▼

You have successfully created the resource hierarchy LSNR.LISTENER1 on node-a. Select a target server to which the hierarchy will be extended.

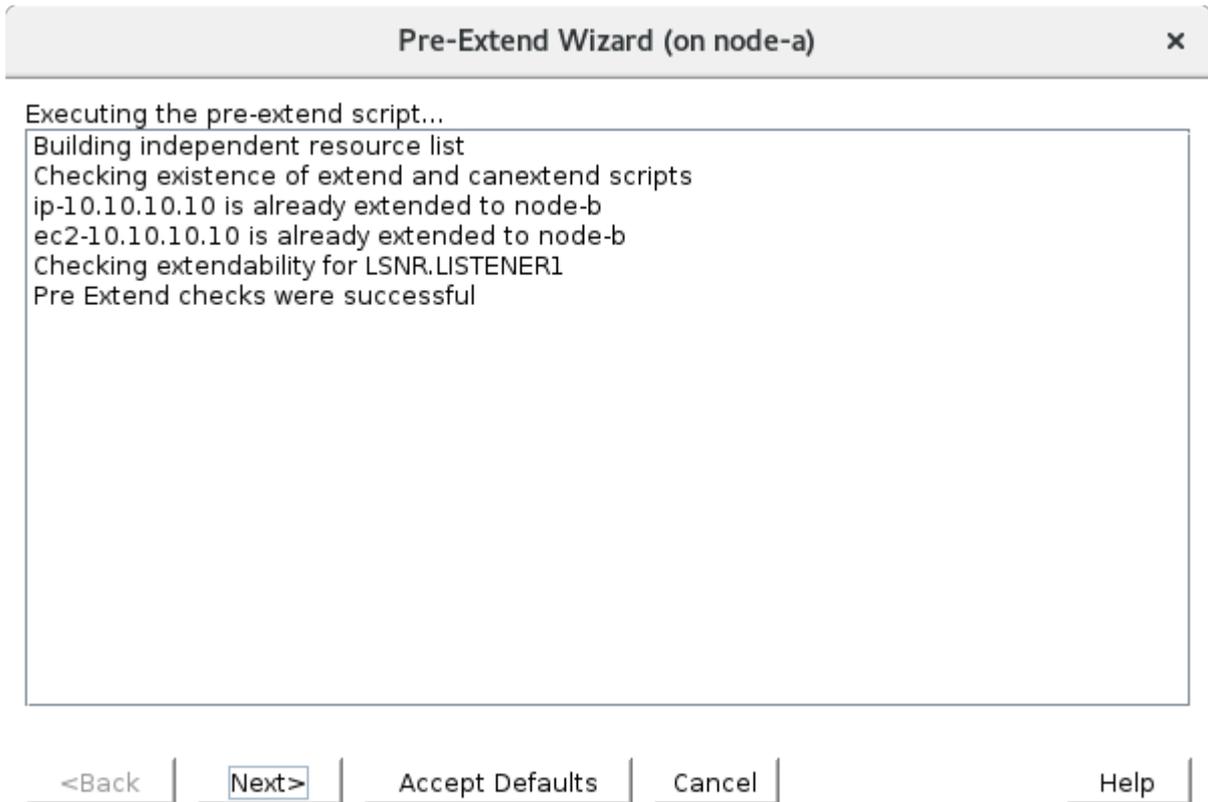
If you cancel before extending LSNR.LISTENER1 to at least one other server, LifeKeeper will provide no protection for the applications in the hierarchy.

<Back
Next>
Accept Defaults
Cancel
Help

Field	Value
Switchback Type	intelligent <span style="color: green; font-weight: bold;">✔</span>

Template Priority	1 
Target Priority	10 

Once the 'Pre-Extend Wizard' checks are passed, select 'Next' to continue.



7. On the "Extend Oracle Listener Resource" wizard, the first choice is selecting the Listener Configuration File Path. This is the same as step #2, use the path detected by the wizard.

**Extend Oracle Listener Resource (on node-a)**
✕

Template Server: node-a  
 Tag to Extend: LSNR.LISTENER1  
 Target Server: node-b

Enter the Listener Configuration File Path

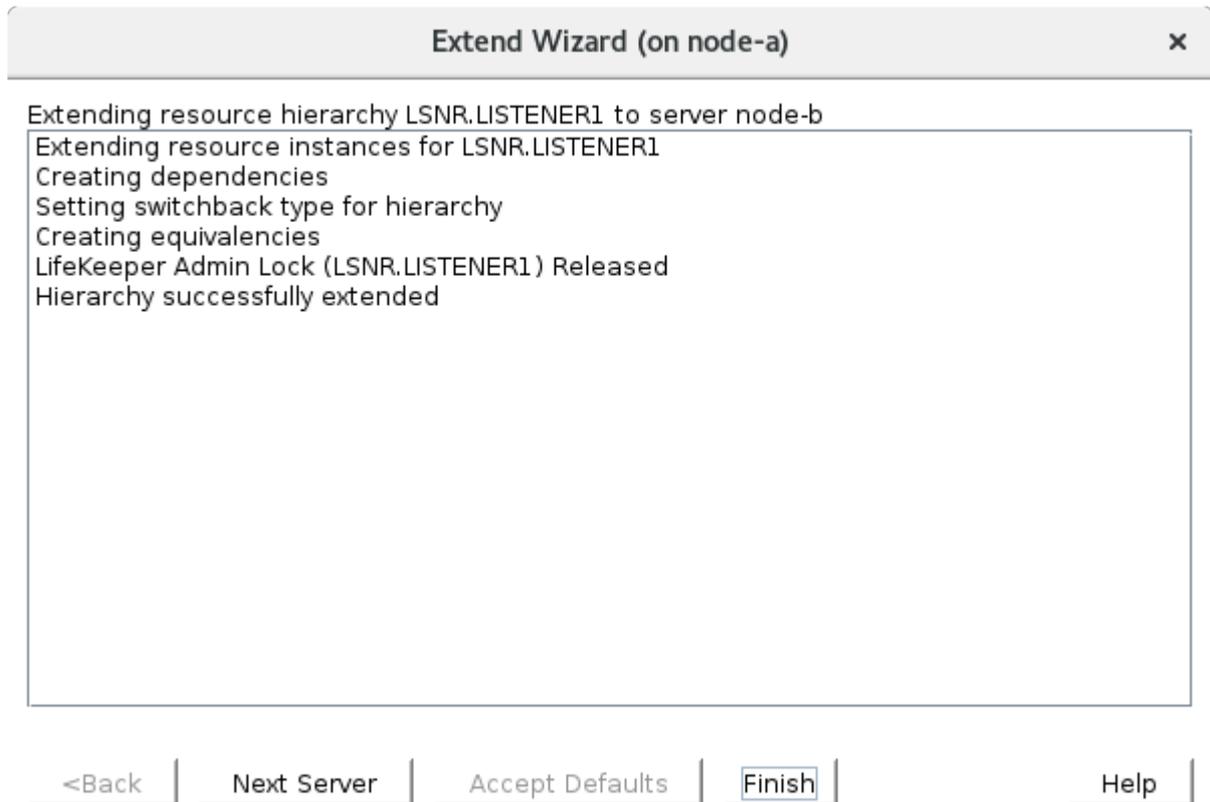
Enter or select the full path to the Oracle Listener configuration file.

<Back
Next>
Accept Defaults
Cancel
Help

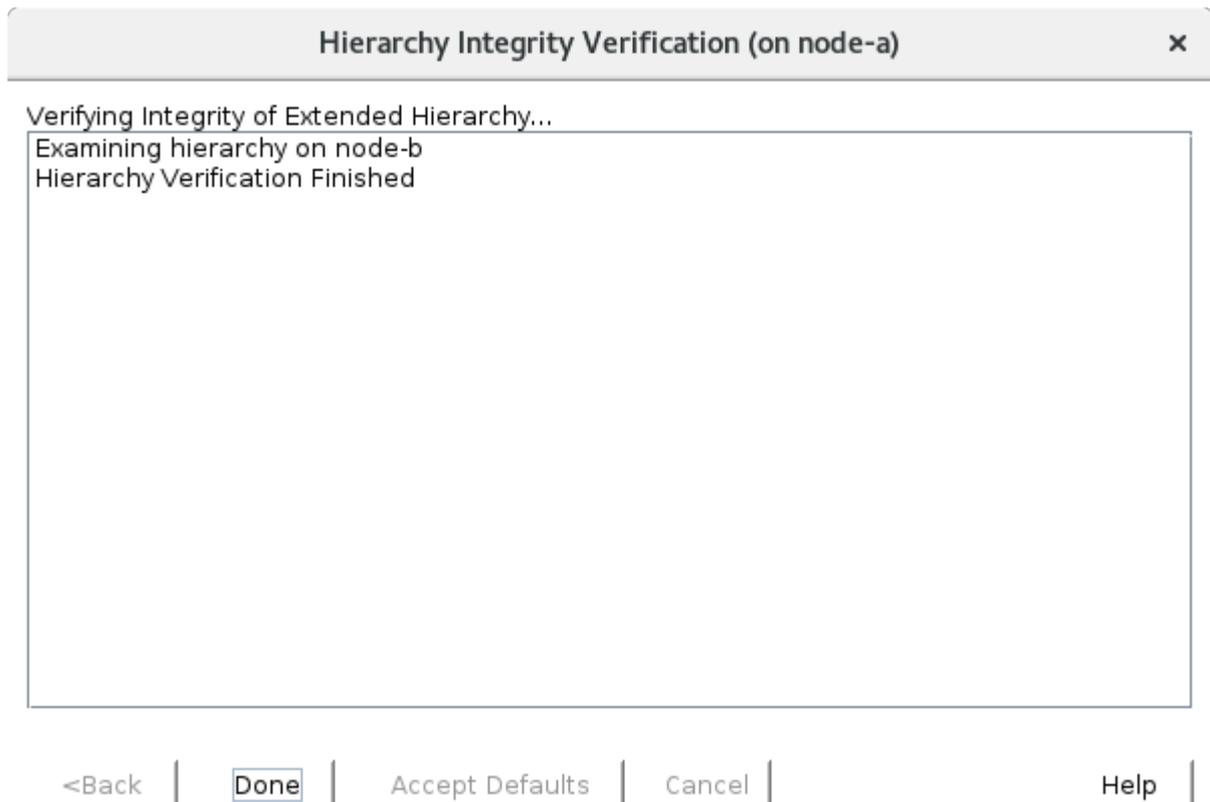
8. Enter the following values. Select the default values.

Field	Value
Listener Executable(s)	/u01/app/oracle/product/19.3.0/dbhome_1/bin/lsnrctl
Listener Tag	LSNR.LISTENER1

9. The Extend Wizard reviews the values and extends the Listener Resource to Node-B. Once the configuration is completed, select "Finish".



10. The Hierarchy is now verified. Select “Done” once the verification completes.



11. Now the Oracle Listener Resource is defined. The Oracle Listener Resource depends on the ip-10.10.10.10 resource as shown below.

LifeKeeper GUI (on node-a)

File Edit View Help

**Hierarchies**

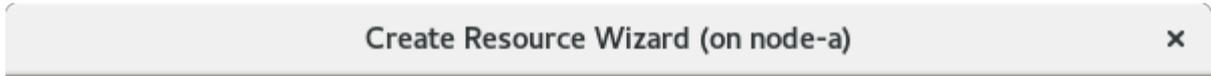
- ✓ Active Protected
- ✓ /datakeeper
  - ✓ datarep-datakeeper
  - ✓ LSNR.LISTENER1
    - ✓ ip-10.10.10.10
      - ✓ ec2-10.10.10.10

node-a		node-b	
✓	Active	1	StandBy 10
✓	Source	1	Target 10
✓	Active	1	StandBy 10
✓	Active	1	StandBy 10
✓	Active	1	StandBy 10

# 11.2.7.5.1.8. Configure the Oracle Resource

As discussed in [Create an Oracle Database \(non-PDB\)](#), the name of Oracle Instance to protect is ORCL.

1. In the LifeKeeper User Interface, define a new resource. Select  to start the Create Resource Wizard (on node-a). Select "Oracle Database" as the Recovery Kit.



Please Select Recovery Kit

2. Enter the following values. Select the default values.

Field	Value
Switchback Type	intelligent 
Server	node-a 
Oracle SID	ORCL 
Username	<Leave blank> 
Oracle Listener	LSNR.LISTENER1 
Database Tag	ORCL 

3. The wizard reviews the values. Once this is complete, click "Next >" to continue.

**Create database/oracle Resource (on node-a)** ✕

Creating database/oracle resource...

```

BEGIN create of "orcl" on server "node-a"
Creating resource instance "orcl" on server "node-a"
Setting resource state for "orcl" on server "node-a" to "ISP".
#####
ORACLE_HOME "/u01/app/oracle/product/19.3.0/dbhome_1" does not reside on a shared
file system. Please be sure that the ORACLE_HOME directory and associated files are
identical on all servers. Refer to the LifeKeeper Oracle Recovery Kit documentation for
more information.
#####
Creating dependency between Oracle database "orcl (orcl)" and the dependent resource
"/datakeeper" on "node-a".
Creating dependency between Oracle database "orcl (orcl)" and the listener resource
"LSNR.LISTENER1" on "node-a".
Performing in-service of new Oracle resource tag=< orcl > on "node-a".
END successful create of "orcl" on server "node-a"

```

<Back
Next>
Cancel
Help

4. The next step is the "Pre-Extend Wizard". Select the default values.

**Pre-Extend Wizard (on node-a)** ✕

Target Server

You have successfully created the resource hierarchy orcl on node-a. Select a target server to which the hierarchy will be extended.

If you cancel before extending orcl to at least one other server, LifeKeeper will provide no protection for the applications in the hierarchy.

<Back
Next>
Accept Defaults
Cancel
Help

Field	Value
Switchback Type	intelligent

Template Priority	1 
Target Priority	10 

Once the 'Pre-Extend Wizard' checks pass, click 'Next' to continue.

**Pre-Extend Wizard (on node-a)** ✕

Executing the pre-extend script...

```

Building independent resource list
Checking existence of extend and canextend scripts
ip-10.10.10.10 is already extended to node-b
ec2-10.10.10.10 is already extended to node-b
LSNR.LISTENER1 is already extended to node-b
datarep-datakeeper is already extended to node-b
Checking extendability for orcl
#####
ORACLE_HOME "/u01/app/oracle/product/19.3.0/dbhome_1" on "node-b" does not reside on
a shared file system. Please be sure that the ORACLE_HOME directory and associated files
are identical on all servers. Use the standard Linux utilities to create and copy directories
and files to "node-b". Refer to the LifeKeeper Oracle Recovery Kit documentation for more
information.
#####
Verifying the Oracle user and group IDs match between "node-a" and "node-b" for
resource "orcl".
Checking Oracle resource "orcl" dependent children extendability to "node-b".
Pre Extend checks were successful

```

<Back
**Next>**
Accept Defaults
Cancel
Help

- On the "Extend database/oracle Resource" wizard, the first choice is selecting the Database Tag. Confirm the default value and click "Extend".

**Extend database/oracle Resource (on node-a)**
✕

Template Server: node-a  
 Tag to Extend: orcl  
 Target Server: node-b

Database Tag

The unique name for the resource on **node-b**. This is a non-editable value and defaults to the tag name specified when creating the resource on **node-a**.

<Back
**Extend**
Accept Defaults
Cancel
Help

- Once the Extend Wizard configures the Oracle resource, click "Finish".

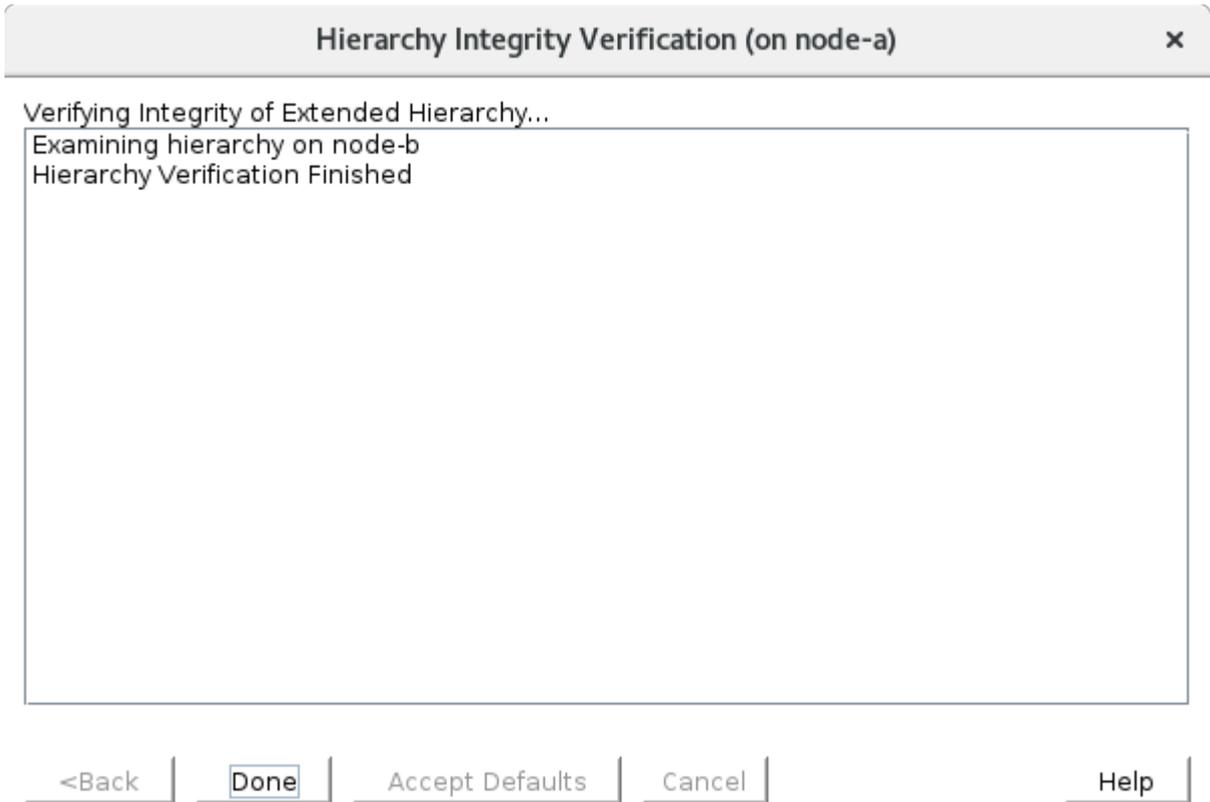
**Extend Wizard (on node-a)**
✕

Extending resource hierarchy orcl to server node-b

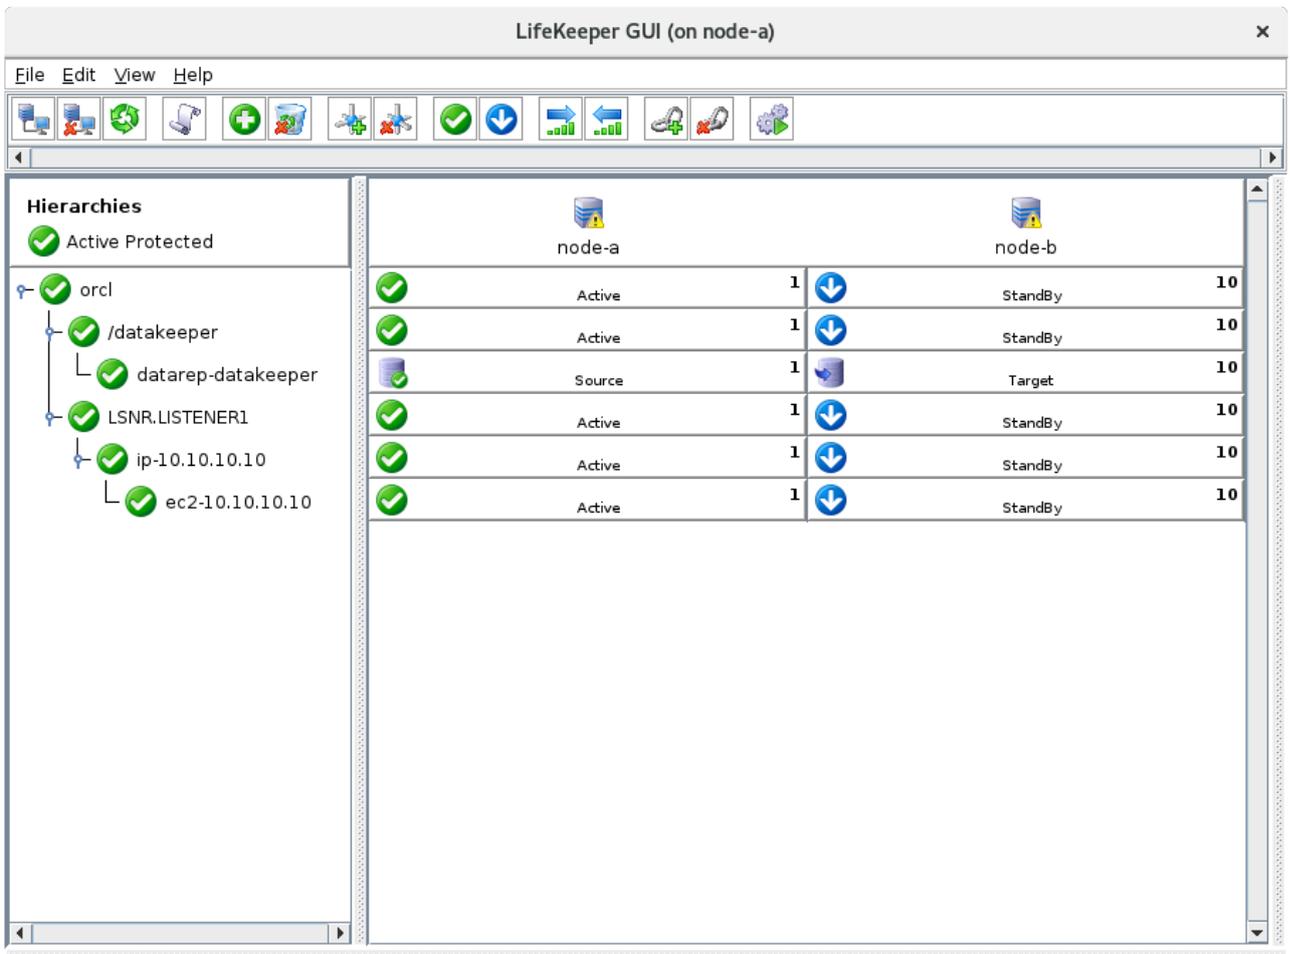
Extending resource instances for orcl  
 Creating resource instance "orcl" on server "node-b"  
 Setting resource state for "orcl" on server "node-b" to "OSU".  
 Creating dependencies  
 Setting switchback type for hierarchy  
 Creating equivalencies  
 LifeKeeper Admin Lock (orcl) Released  
 Hierarchy successfully extended

<Back
**Next Server**
Accept Defaults
Finish
Help

- The Hierarchy is now checked and complete, click "Done".



- 8. Return to the LifeKeeper GUI. The “orcl” resource depends on the /datakeeper resource and LSNR.LISTENER as shown below.



Resource Tag= orcl, Resource ID= orcl

## 11.2.7.5.1.9. Test Switchover of the Oracle Resource

### Understanding the Status from Command Line Tools

Before testing the switchover, review the status of each node.

#### Understand the Status of the File System

The easiest way to check the status is using `df`.

```

1 [oracle@node-b dbhome_1]$ df
2 Filesystem      1K-blocks      Used Available Use% Mounted on
3 devtmpfs         916888          0   916888   0% /dev
4 tmpfs            7340032    450560   6889472   7% /dev/shm
5 tmpfs            940184     17480   922704   2% /run
6 tmpfs            940184          0   940184   0% /sys/fs/cgroup
7 /dev/xvda2      31444972 17247248 14197724  55% /
8 tmpfs            188040          0   188040   0% /run/user/1000
9 /dev/md0        20960236 4997824 15962412  24% /datakeeper
10 tmpfs            188040          0   188040   0% /run/user/54321

```

The `/datakeeper` resource is attached to the node and is mounted as shown here.

#### Understand the Status of the Listener and Oracle Processes

We can use the `ps` command to review the status of each node. If the processes are running they will be listed as shown below.

```

$ ps afx | grep -v grep | grep tnslnsr
13806 ?          Ssl    0:00 /u01/app/oracle/product/19.3.0/dbhome_1/bin/tnslnsr LISTENER1 -inherit
$ ps afx | grep -v grep | grep ora_pmon
23091 ?          Ss     0:00 ora_pmon_orcl

```

#### Understand the Status of the IP Resource

As discussed in [How a Client Connects to the Active Node](#), if the IP resource is active, the virtual IP address can be found on the `ip` command as follows (see 10.10.10.10).

```

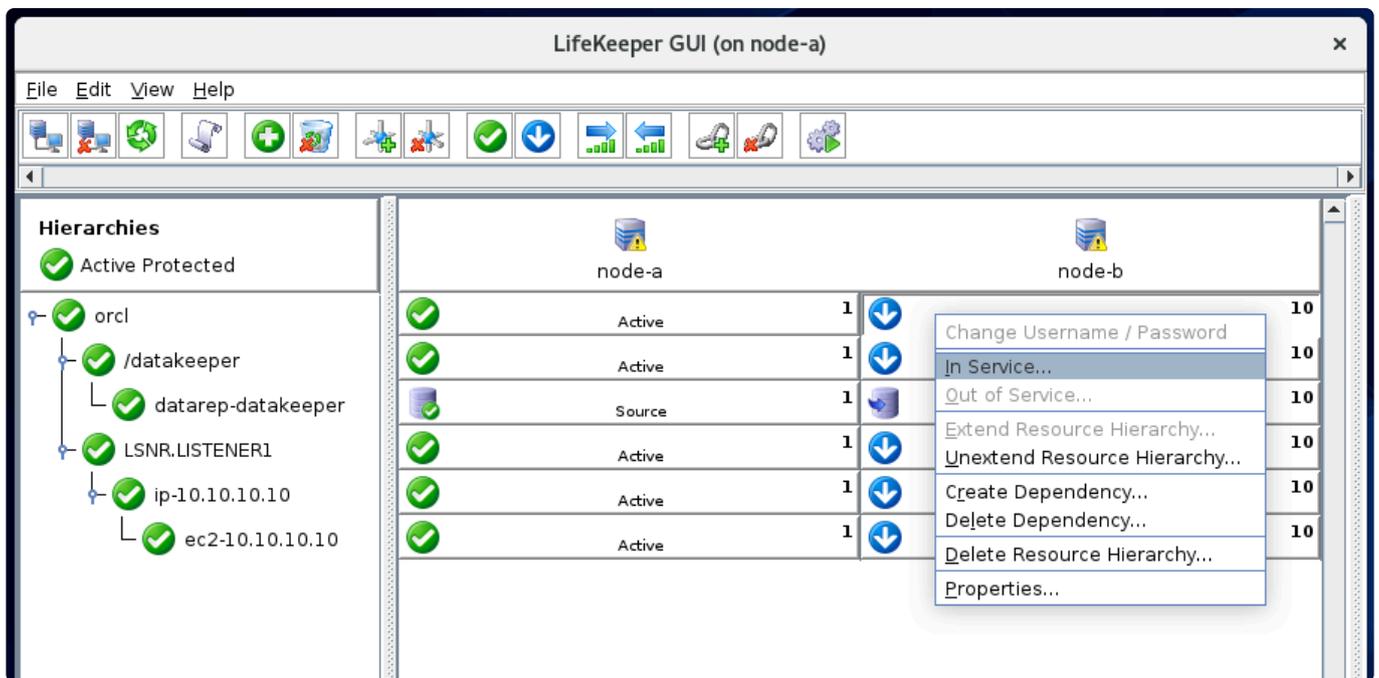
[oracle@node-b dbhome_1]$ ip addr show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9001 qdisc pfifo_fast state UP
group default qlen 1000

```

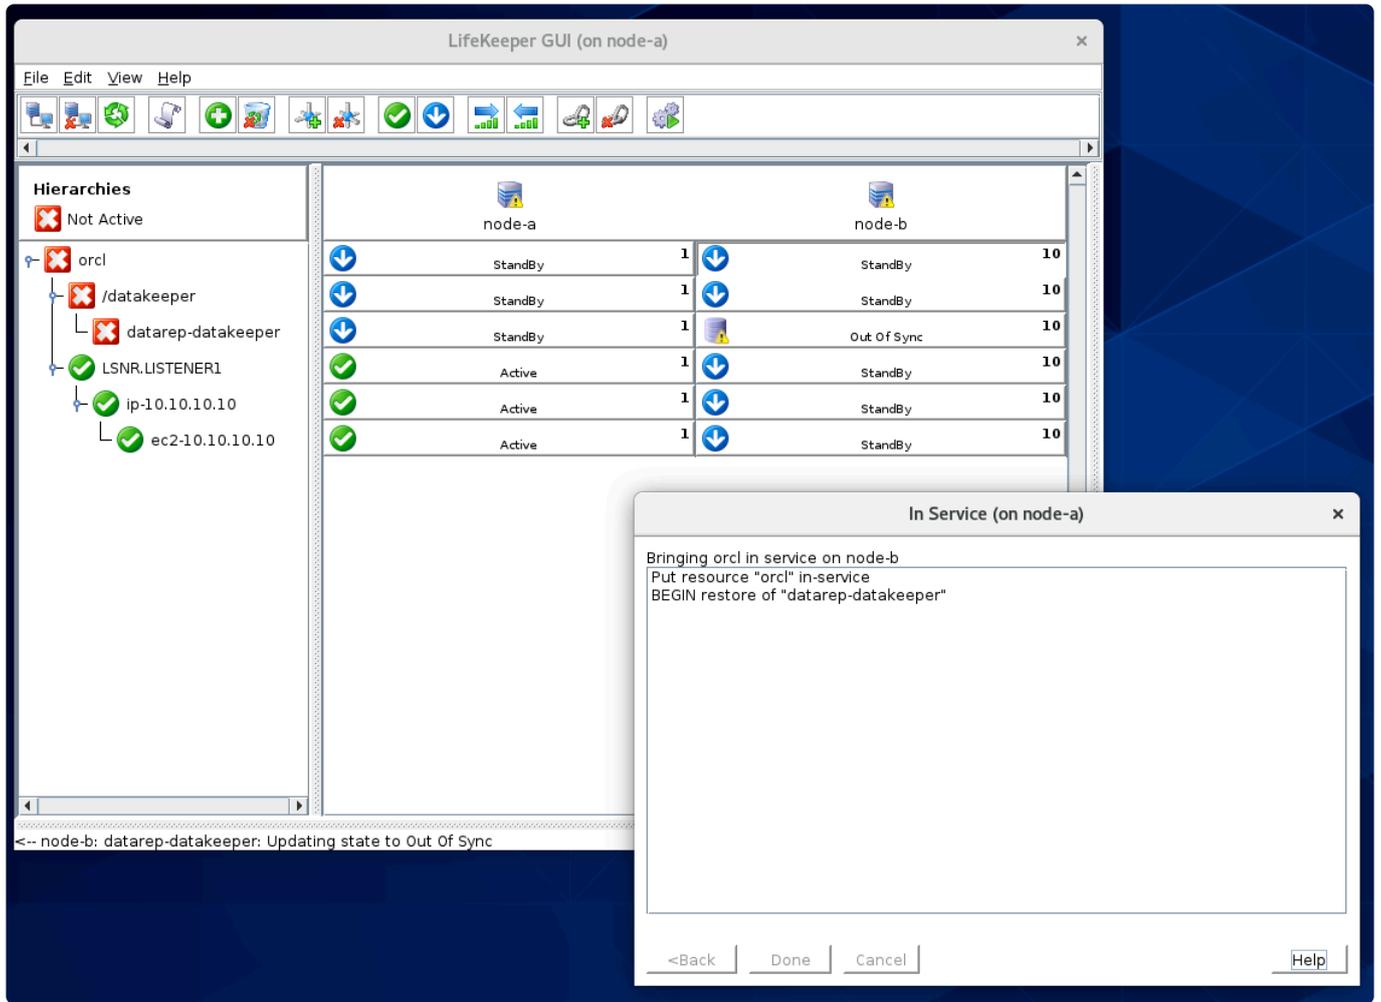
```
link/ether 02:a8:0c:57:53:0b brd ff:ff:ff:ff:ff:ff
inet 10.20.10.12/24 brd 10.20.10.255 scope global noprefixroute dynamic et
h0
    valid_lft 3585sec preferred_lft 3585sec
inet 10.10.10.10/32 scope global eth0
    valid_lft forever preferred_lft forever
inet6 fe80::a8:cff:fe57:530b/64 scope link
    valid_lft forever preferred_lft forever
```

## Switching the Resource Between Nodes

Show the context menu on the “orcl” resource on node-b (standby node) and select “In Service”.



First, the resource and all child dependencies will be transitioned to “StandBy” on node-a.



Once all resources are stopped on node-a, the dependencies and the “orcl” resource will be started on node-b.

The screenshot displays the LifeKeeper GUI interface. On the left, a 'Hierarchies' tree shows the following structure:

- Active Protected
  - orcl
    - /datakeeper
      - datarep-datakeeper
    - LSNR.LISTENER1
      - ip-10.10.10.10
        - ec2-10.10.10.10

The main table shows the status of these components across two nodes:

Component	node-a	node-b
orcl	StandBy	Active
/datakeeper	StandBy	Active
datarep-datakeeper	Target	Source
LSNR.LISTENER1	StandBy	Active
ip-10.10.10.10	StandBy	Active
ec2-10.10.10.10	StandBy	Active

An 'In Service (on node-a)' dialog box is open, displaying the following log output:

```
Bringing orcl in service on node-b
Partial resynchronization of component "/dev/nbd0" has begun for mirror "/dev/md0"
END successful restore of "datarep-datakeeper"
BEGIN restore of /datakeeper
mounting file system /datakeeper
mount -txfs -orw,relatime,attr2,nobarrier,inode64,noquota /dev/md0 /datakeeper
File system /datakeeper has been successfully mounted.
END successful restore of /datakeeper
BEGIN restore of "ec2-10.10.10.10"
END successful restore of "ec2-10.10.10.10"
BEGIN restore of "ip-10.10.10.10"
END successful restore of "ip-10.10.10.10"
BEGIN restore of "LSNR.LISTENER1" on server "node-b"
END successful restore of "LSNR.LISTENER1" on server "node-b"
BEGIN restore of "orcl" on server "node-b"
Begin the "start [ start.normal ]" of the database "orcl" on "node-b".
End successful "start [ start.normal ]" of the database "orcl" on "node-b".
END successful restore of "orcl" on server "node-b"
Put "orcl" in-service successful
```

At the bottom of the GUI, a status bar shows: <-- node-a: orcl: Updating state to StandBy

## 11.2.7.5.2. Protecting MSSQL Using Quick Service Protection

This section outlines the steps to protect Microsoft SQL Server 2017 in a Linux environment.

\* This document uses Microsoft SQL Server 2017 as an example. You can apply the same steps on Microsoft SQL Server 2012 R2 through Microsoft SQL Server 2019.

Although LifeKeeper doesn't have a specific Application Recovery Kit for MSSQL, LifeKeeper can still protect MSSQL as a general service. This feature is called Quick Service Protection (QSP) and you can protect MSSQL as well as other services running in a Linux environment.

To protect database resources, the data needs to be replicated across nodes using DataKeeper (unless you use shared storage or a SAN device).

The following table outlines the location of each component.

Items	Location
Master DataBase Files	/datakeeper/mssql/data
Master Log File	/datakeeper/mssql/xlog

As discussed earlier in [How to Create Data Replication of a File System](#), this guide uses `/datakeeper` to replicate data between nodes. Therefore, the data on `/datakeeper/mssql/data` is also replicated between nodes (this is the same for `/datakeeper/mssql/xlog` as well).

Please note that this guide uses the following computing resources (Microsoft SQL Server resource requires more memory than other resources even for evaluation purposes).

Resource	Required Size
Memory	4 GiB

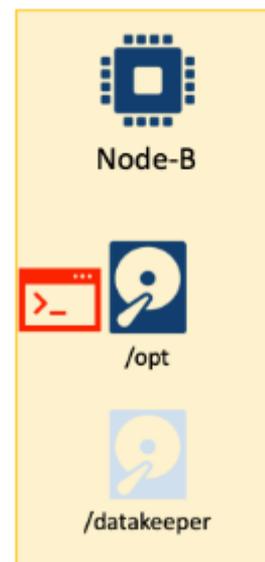
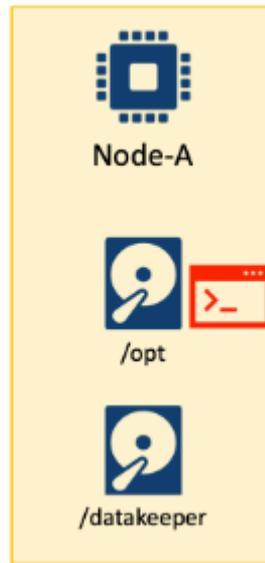
### Configure the MSSQL Resource

The following table outlines the general steps to configure a MSSQL Resource. The red 'stacked disk' shaped icons indicate a node that has the database instances at the time of each step. The grey 'stacked disk' icons indicate a database that is not running.

Also, the grey "storage" icon indicates that the replicated storage `/datakeeper` is not available for the node.

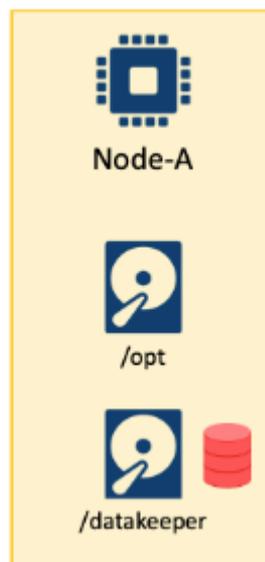
**Install MSSQL 2017 on both nodes**

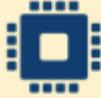
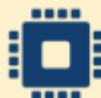
- [Install MSSQL 2017 on a System](#)



**Create Database on Node-A**

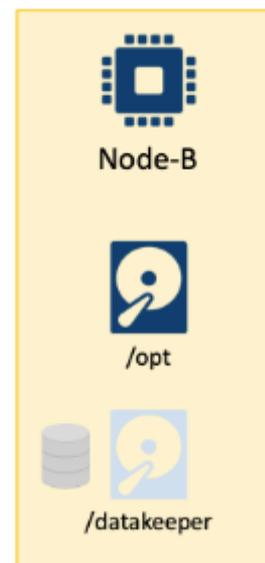
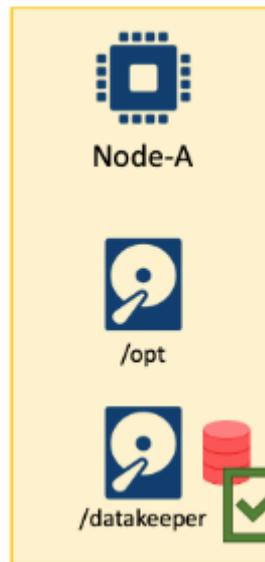
- Review [How to Confirm if the Data Storage is Available on a Node](#). The data storage should be available on node-a.
- [Relocate Master Database and Log Files to Replicated Storage](#) on node-a



	 <p>Node-B</p>  <p>/opt</p>  <p>/datakeeper</p>
<p><b>Create Database on Node-B</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Switchover (the data storage)</a> to node-b</li> <li>• <a href="#">Rename Folders Under /datakeeper/MSSQL</a></li> <li>• <a href="#">Relocate Master Database and Log Files to Replicated Storage</a> on node-b</li> </ul>	 <p>Node-A</p>  <p>/opt</p>  <p>/datakeeper</p>  <p>Node-B</p>  <p>/opt</p>  <p>/datakeeper</p>

**Complete Configuration using LifeKeeper**

- [Switchover \(the data storage\)](#) back to node-a
- [Configure Resources with LifeKeeper](#)
- [Customize LocalRecovery Parameter on Both Nodes](#)
- [Update Dependency between Resources](#)



## 11.2.7.5.2.1. Install MSSQL 2017

✿ Install MSSQL on both nodes.

1. Refer to Microsoft's [Installation guidance for SQL Server on Linux](#) guide. It contains instructions for different versions of operating systems as well as MSSQL servers.
2. Download the Microsoft SQL Server 2017 Red Hat repository configuration file.

```
# curl -o /etc/yum.repos.d/mssql-server.repo  
https://packages.microsoft.com/config/rhel/8/mssql-server-2017.repo
```

3. Run the following commands to install SQL Server:

```
1 # yum install -y mssql-server
```

4. Configure the MSSQL server.

```
1 # /opt/mssql/bin/mssql-conf setup
```

5. The server is now installed. Check the status of MSSQL server with the following command.

```
1 # systemctl status mssql-server
```



6. There are a few more steps before starting MSSQL Server. The MSSQL server is now installed successfully. Complete these steps on both nodes.

## 11.2.7.5.2.2. Relocate Master Database and Log Files to Replicated Storage

As discussed earlier, the data should be stored on the replicated file system. Therefore, we use the following locations for each component.

Items	Location
Master DataBase Files	/datakeeper/mssql/data
Master Log File	/datakeeper/mssql/xlog

\* The following steps should be carried out on node-a **AND** node-b.

There are some additional steps that need to be done in between. Please review the summary of the overall steps in [Protecting MSSQL Using Quick Service Protection](#).

First, create directories on node-a.

```
1 # mkdir -p /datakeeper/mssql/data
2 # mkdir -p /datakeeper/mssql/xlog
3 # chown mssql -R /datakeeper/mssql/
4 # chgrp mssql -R /datakeeper/mssql/
```

Relocate the Master database to the newly created directories with the `/opt/mssql/bin/mssql-conf` tool as follows:

```
1 # /opt/mssql/bin/mssql-conf set filelocation.masterlogfile /datakeeper/mssql/xlog/mastlog.ldf
2 SQL Server needs to be restarted in order to apply this setting. Please run
3 'systemctl restart mssql-server.service'.
4
5 # /opt/mssql/bin/mssql-conf set filelocation.masterdatafile /datakeeper/mssql/data/master.mdf
6 SQL Server needs to be restarted in order to apply this setting. Please run
7 'systemctl restart mssql-server.service'.
```

Start the MSSQL server.

```
1 # systemctl restart mssql-server.service
2 # systemctl status mssql-server.service
3 • mssql-server.service - Microsoft SQL Server Database Engine
4   Loaded: loaded (/usr/lib/systemd/system/mssql-server.service; enabled; vendor preset: disabled)
5   Active: active (running) since Wed 2020-12-30 18:55:54 UTC; 6s ago
6     Docs: https://docs.microsoft.com/en-us/sql/linux
7   Main PID: 15506 (sqlservr)
8   CGroup: /system.slice/mssql-server.service
9           └─15506 /opt/mssql/bin/sqlservr
10          └─15508 /opt/mssql/bin/sqlservr
```

Once the database is started, you will see some files created on the folders specified earlier.

```
1 # ls -al /datakeeper/mssql/data
2 total 51136
3 drwxr-xr-x 2 mssql mssql    137 Dec 30 18:56 .
4 drwxr-xr-x 6 mssql mssql    62 Dec 30 18:48 ..
5 -rw-rw---- 1 mssql mssql 4194304 Dec 30 18:56 master.mdf
6 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:56 modellog.ldf
7 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:56 model.mdf
8 -rw-rw---- 1 mssql mssql 14090240 Dec 30 18:56 msdbdata.mdf
9 -rw-rw---- 1 mssql mssql 524288 Dec 30 18:56 msdblog.ldf
10 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:56 tempdb.mdf
11 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:56 templog.ldf
12 # ls -al /datakeeper/mssql/xlog
13 total 2048
14 drwxr-xr-x 2 mssql mssql    25 Dec 30 18:55 .
15 drwxr-xr-x 6 mssql mssql    62 Dec 30 18:48 ..
16 -rw-rw---- 1 mssql mssql 2097152 Dec 30 18:56 mastlog.ldf
```

Now the instance can be stopped.

```
1 # systemctl stop mssql-server.service
```

## 11.2.7.5.2.3. Rename Folders Under /dataKeeper/MSSQL

---

So far the following tasks have been completed:

1. Created the database on node-a
2. Switched over the `datarep-datakeeper` resource to node-b

In step 1, the information about the newly created mssql database was created in `/var/opt/mssql/mssql.conf`.

Now we will create the same mssql database on node-b so the information about the mssql database will be stored on node-b in `/var/opt/mssql/mssql.conf`. The database files themselves are located in `/datakeeper/mssql/`.

On node-b, these files can be seen as shown below:

```
1 [root@node-b ~]# cd /datakeeper/mssql/
2 [root@node-b mssql]# ls -al data
3 total 51136
4 drwxr-xr-x 2 mssql mssql      137 Dec 30 18:38 .
5 drwxr-xr-x 6 mssql mssql       62 Dec 30 18:36 ..
6 -rw-rw---- 1 mssql mssql 4194304 Dec 30 18:38 master.mdf
7 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:38 modellog.ldf
8 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:38 model.mdf
9 -rw-rw---- 1 mssql mssql 14090240 Dec 30 18:38 msdbdata.mdf
10 -rw-rw---- 1 mssql mssql 524288 Dec 30 18:38 msdblog.ldf
11 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:38 tempdb.mdf
12 -rw-rw---- 1 mssql mssql 8388608 Dec 30 18:38 templog.ldf
13 [root@node-b mssql]# ls -al xlog_bak/
14 total 2048
15 drwxr-xr-x 2 mssql mssql      25 Dec 30 18:38 .
16 drwxr-xr-x 6 mssql mssql       62 Dec 30 18:36 ..
17 -rw-rw---- 1 mssql mssql 2097152 Dec 30 18:38 mastlog.ldf
```

Rename these folders so that the database can be created successfully on node-b.

```
1 [root@node-b mssql]# mv data data_bak
2 [root@node-b mssql]# mv xlog xlog_bak
3 [root@node-b mssql]# ls -al
4 total 0
5 drwxr-xr-x 4 mssql mssql 38 Dec 30 18:48 .
6 drwxr-xr-x 5 root root 57 Dec 30 18:36 ..
7 drwxr-xr-x 2 mssql mssql 137 Dec 30 18:38 data_bak
8 drwxr-xr-x 2 mssql mssql 25 Dec 30 18:38 xlog_bak
```

# 11.2.7.5.2.4. Configure MSSQL Resource

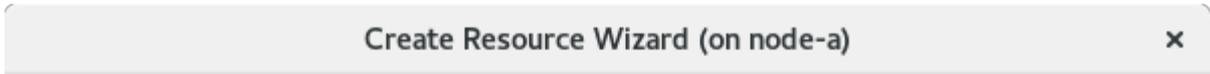
Now that the master database has been relocated to replicated storage, the MSSQL server resources are ready to be protected using LifeKeeper. Before configuring it with LifeKeeper, start MSSQL on node-a.

```

1 [root@node-a ~]# systemctl start mssql-server.service
2 [root@node-a ~]# systemctl status mssql-server.service
3 • mssql-server.service - Microsoft SQL Server Database Engine
4   Loaded: loaded (/usr/lib/systemd/system/mssql-server.service; enabled; vendor preset: disabled)
5   Active: active (running) since Wed 2020-12-30 18:59:02 UTC; 1s ago
6     Docs: https://docs.microsoft.com/en-us/sql/linux
7   Main PID: 17146 (sqlservr)
8   CGroup: /system.slice/mssql-server.service
9           └─17146 /opt/mssql/bin/sqlservr
10          └─17148 /opt/mssql/bin/sqlservr

```

1. In the LifeKeeper User Interface, define a new resource. Click the  icon to start the Create Resource Wizard (on node-a). Select "Quick Service Protection" as the Recovery Kit.



Please Select Recovery Kit Quick Service Protection ▼

<Back | Next> | Cancel Help

2. Enter the following values. Select the default values.

Field	Value
Switchback Type	intelligent 
Server	node-a 

3. Select “mssql-server” for the Service Name from the dropdown list.

**Create gen/qsp Resource (on node-a)** ✕

Service Name

- cloud-config
- cloud-final
- cloud-init
- cloud-init-local
- gssproxy
- mssql-server
- sshd
- temp-disk-swapfile

Select the service that you want to protect.

<Back
Next>
Cancel
Help

4. Enter the following values. Select the default values.

Field	Value
Enable or Disable Monitoring	enable <span style="color: green; font-weight: bold;">✔</span>
Resource Tag	QSP-mssql-server <span style="color: green; font-weight: bold;">✔</span>

5. The wizard reviews the values. Once it completes, click “Next >” to continue.

**Create gen/qsp Resource (on node-a)** ✕

Creating gen/qsp resource QSP-mssql-server on node-a

```

BEGIN create of "QSP-mssql-server"
info: cfgQuickCheck=enable

Performing in-service of QSP resource tag QSP-mssql-server.
BEGIN restore of "QSP-mssql-server"
END successful restore of "QSP-mssql-server"
END successful create of "QSP-mssql-server"
        
```

Messages produced while creating **QSP-mssql-server** will be displayed in this dialog and the output panel (if open), and logged on **node-a**.

<Back
Next>
Cancel
Help

6. The next step is the "Pre-Extend Wizard". Select the default values.

**Pre-Extend Wizard (on node-a)** ✕

Target Server node-b ▼

You have successfully created the resource hierarchy QSP-mssql-server on node-a. Select a target server to which the hierarchy will be extended.

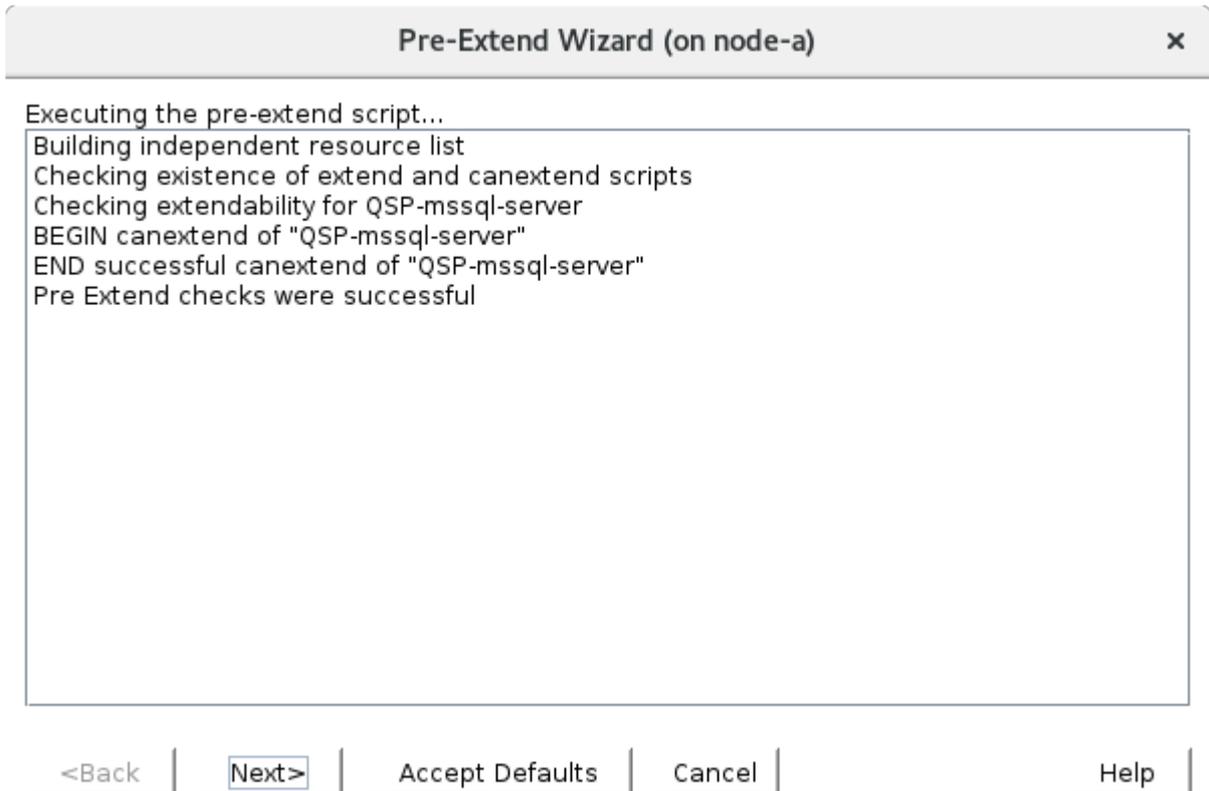
If you cancel before extending QSP-mssql-server to at least one other server, LifeKeeper will provide no protection for the applications in the hierarchy.

<Back
Next>
Accept Defaults
Cancel
Help

Field	Value
Switchback Type	intelligent <span style="color: green; font-weight: bold;">✔</span>

Template Priority	1 
Target Priority	10 

Once the checks on the 'Pre-Extend Wizard' have completed and passed, click 'Next' to continue.



7. On the "Extend gen/qsp Resource" wizard, the first choice is selecting the Database Tag. Confirm the default value and click "Extend".

**Extend gen/qsp Resource (on node-a)** ✕

Template Server: node-a  
 Tag to Extend: QSP-mssql-server  
 Target Server: node-b

Resource Tag

Enter a unique name for the resource instance on **node-b**.

The valid characters allowed for the tag are letters, digits, and the following special characters:  
 - \_ . /

<Back
Extend
Accept Defaults
Cancel
Help

8. Once the Extend Wizard configures the gen/qsp resource, click “Finish” to complete.

**Extend Wizard (on node-a)** ✕

Extending resource hierarchy QSP-mssql-server to server node-b

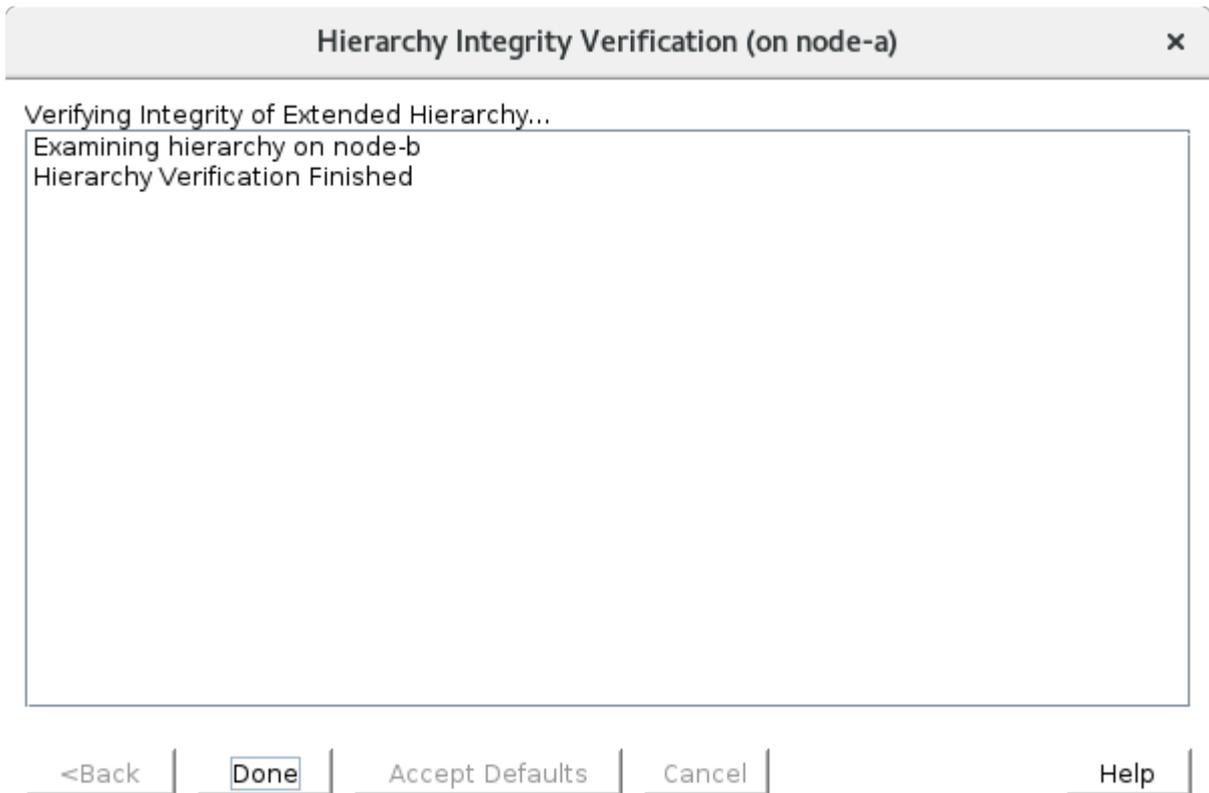
```

Extending resource instances for QSP-mssql-server
BEGIN extend of "QSP-mssql-server"

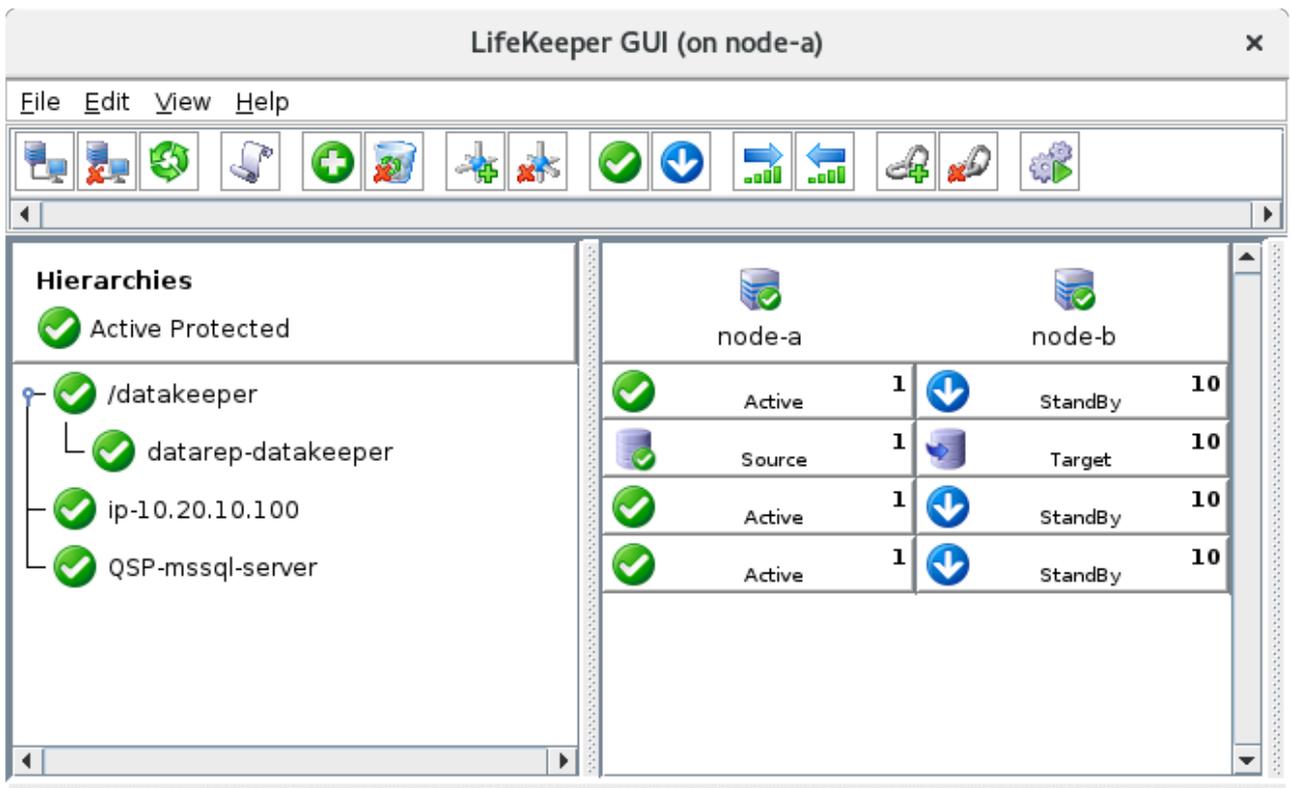
END successful extend of "QSP-mssql-server"
Creating dependencies
Setting switchback type for hierarchy
Creating equivalencies
LifeKeeper Admin Lock (QSP-mssql-server) Released
Hierarchy successfully extended
          
```

<Back
Next Server
Accept Defaults
Finish
Help

9. The hierarchy has now been successfully created and verified. Click “Done” to complete.



10. Return to the LifeKeeper GUI. The “QSP-mssql-server” resource is defined as shown below.



<-- node-b: QSP-mssql-server: Updating equivalency list

## 11.2.7.5.2.5. Customize LocalRecovery Parameter on Both Nodes

The Quick Service Protection Wizard completes most of the configuration. However, we still need to customize some parameters.

 Execute this command on both node-a **AND** node-b.

First, we'll review some policy parameters using the `lkipolicy` tool as follows:

```
1 # /opt/LifeKeeper/bin/lkipolicy -g -v
2 Server-level Policies:
3     Failover: On
4     LocalRecovery: On
5 Resource-level Policies:
6     QSP-mssql-server:
7         Failover: On
8         LocalRecovery: On    <--- This needs to be updated
9     /datakeeper:
10        Failover: On
11        LocalRecovery: On
12    datarep-datakeeper:
13        Failover: On
14        LocalRecovery: On
15    ip-10.20.10.100:
16        Failover: On
17        LocalRecovery: On
```

The `LocalRecovery` parameter for `QSP-mssql-server` is currently set to `On` and needs to be changed to `Off`. Update it with the following command.

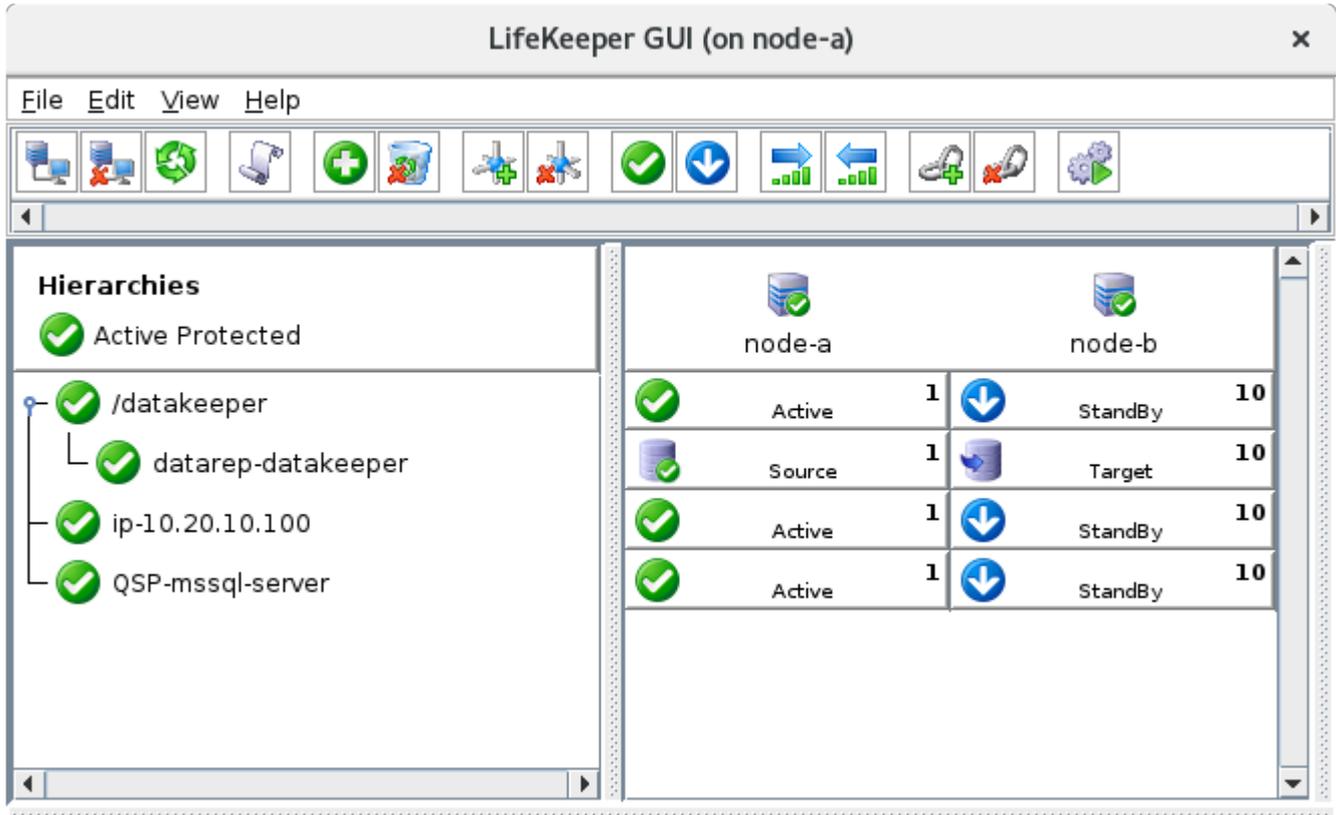
```
1 # /opt/LifeKeeper/bin/lkipolicy -s LocalRecovery -E tag="QSP-mssql-server"
2 Policy Disabled
```

Once the value is updated, confirm it as follows:

```
1 # /opt/LifeKeeper/bin/lkpolicy -g -v
2 Server-level Policies:
3     Failover: On
4     LocalRecovery: On
5 Resource-level Policies:
6     QSP-mssql-server:
7         Failover: On
8         LocalRecovery: Off
9     /datakeeper:
10         Failover: On
11         LocalRecovery: On
12     datarep-datakeeper:
13         Failover: On
14         LocalRecovery: On
15     ip-10.20.10.100:
16         Failover: On
17         LocalRecovery: On
```

# 11.2.7.5.2.6. Update Dependency between Resources

At this point, the LifeKeeper GUI shows the MSSQL related resources as shown below:



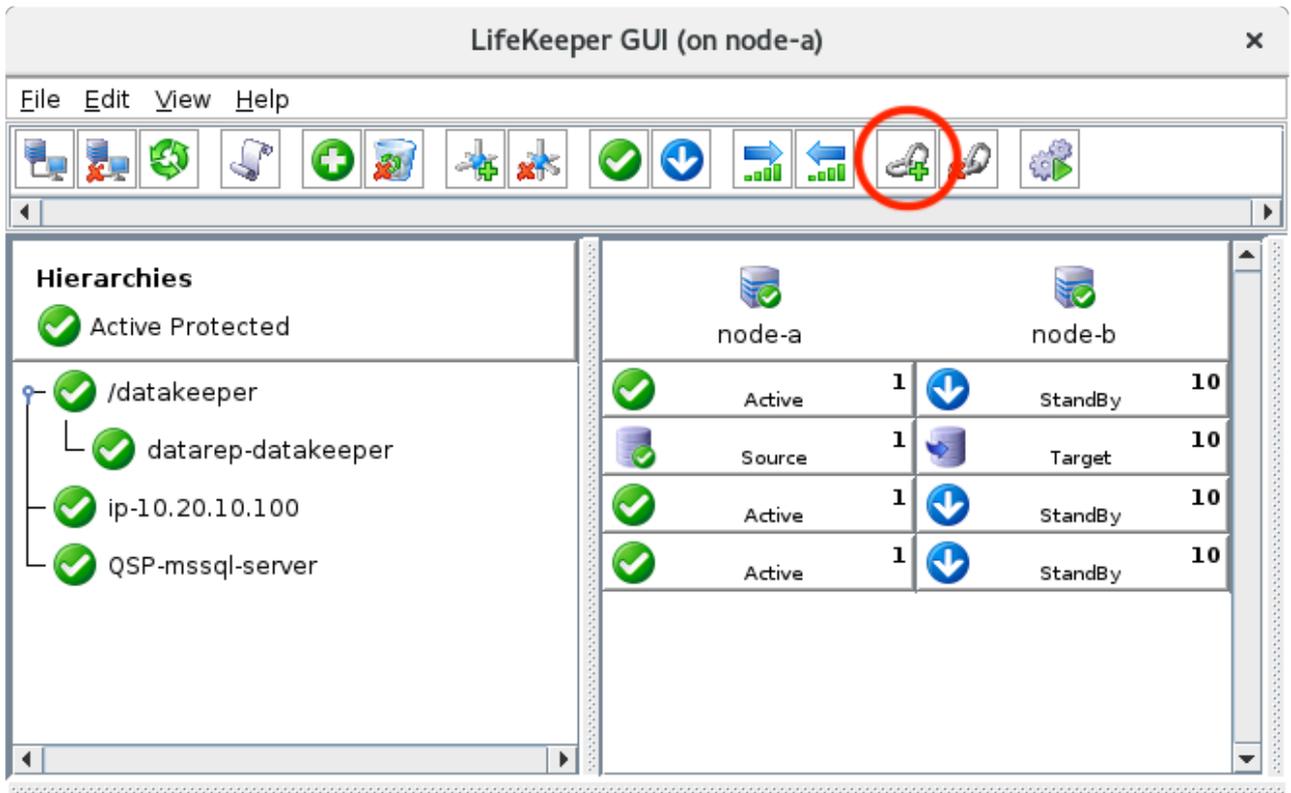
<-- node-b: QSP-mssql-server: Updating equivalency list

The QSP-mssql-server resource depends on the following resources:

- ip-10.20.10.100 (to connect to the database, we need the virtual IP to connect to)
- /datakeeper (to make the database work, we need to mount the underlying filesystem)

This section explains how to define the dependency between resources.

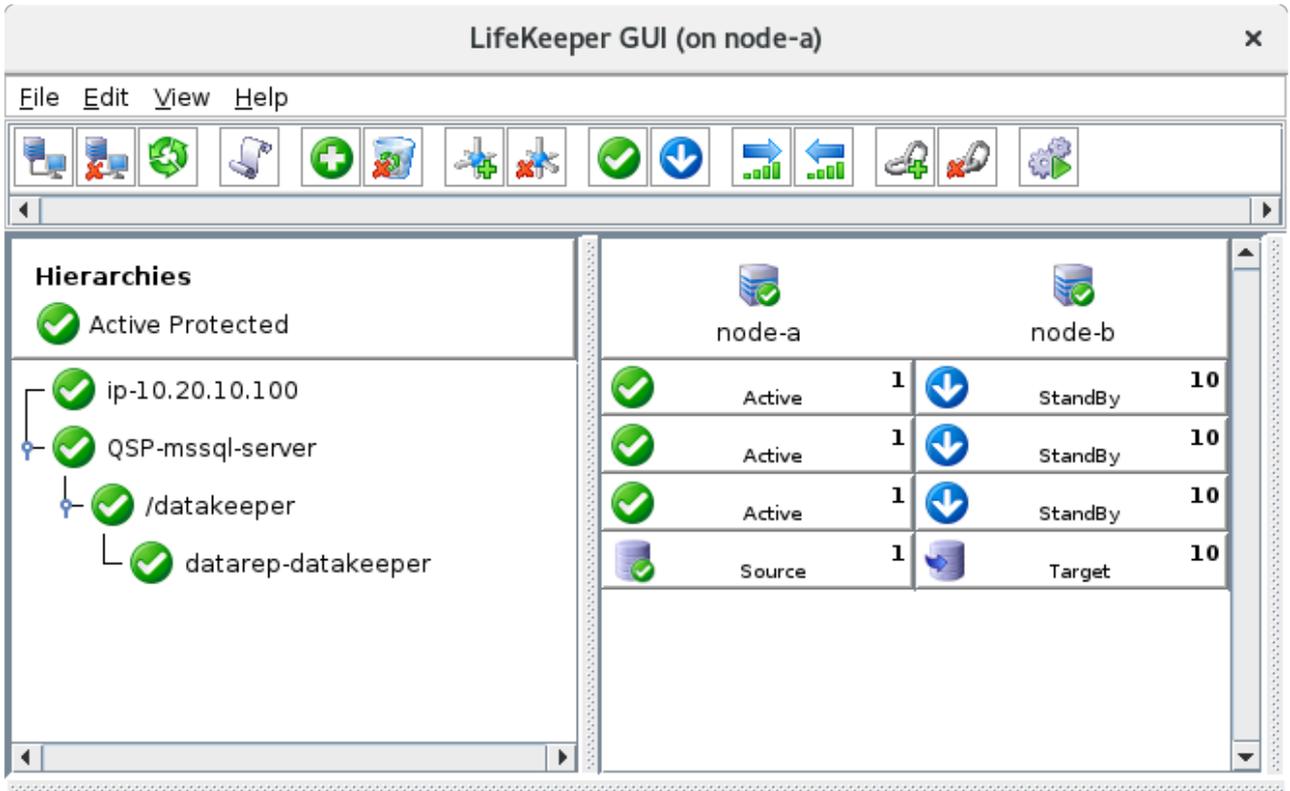
1. Select "Create Dependency".



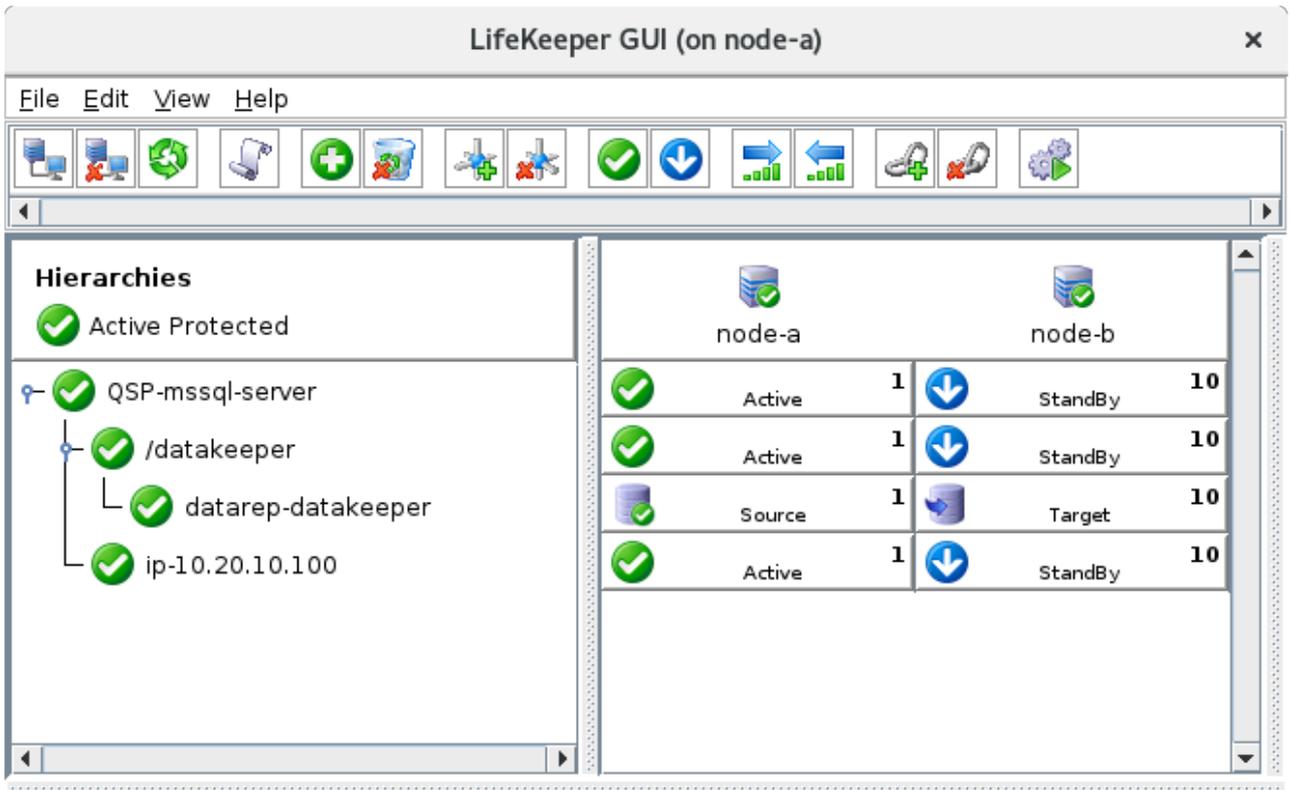
<-- node-b: QSP-mssql-server: Updating equivalency list

2. Select "node-a" as the server.
3. Select "QSP-mssql-server" as the Parent Resource Tag.
4. Select "datakeeper" as the Child Resource Tag.

Once the dependency is defined, the LifeKeeper GUI shows the relationship as follows:



5. Repeat the steps with "ip-10.20.10.100". Now the dependencies are defined.



<-- node-b: ip-10.20.10.100: Adding parent: QSP-mssql-server

## 11.2.7.5.3. Protecting a PostgreSQL Resource

! This section uses `datakeeper/pgsql/data/` as the location for Postgres data. In the steps below, this directory will be created on the replicated storage. If a different directory structure is created, be sure to account for the specific environment variances throughout the remainder of the steps described below.

Before configuring Postgres 12, please follow the instructions at [Install Postgres 12 to Linux Nodes](#).

\* Once PostgreSQL is installed on both the primary and standby nodes, execute the following steps on the primary node **only**.

### Create a Directory to Store Data on the Primary Node

```
1 [root@node-a]# mkdir -m 700 -p /datakeeper/pgsql/data
2 [root@node-a]# chown postgres:postgres -R /datakeeper/pgsql
```

### Initialize Postgres Database Server on the Primary Node

In this instance, `initdb` is located at `/usr/pgsql-12/bin/initdb`. Use the location of `initdb` to construct the parameters here.

```
1 [root@node-a ~]# su postgres -c "/usr/pgsql-12/bin/initdb -D /datakeeper/pgsql/data"
2 The files belonging to this database system will be owned by user "postgres".
3 This user must also own the server process.
4
5 (snip)
6
7 Success. You can now start the database server using:
8
9 /usr/pgsql-12/bin/pg_ctl -D /datakeeper/pgsql/data -l logfile start
```

\* Check the following locations to be used during the following steps.

Item	Location
Location of Postgres binaries	<code>/usr/pgsql-12/bin/</code>

Data	/datakeeper/pgsql/data
------	------------------------

## Start the Postgres Process on the Primary Node

```
[root@node-a ~]# su postgres -c "/usr/pgsql-12/bin/pg_ctl start -D /datakeeper/pgsql/data"
waiting for server to start....2020-12-30 03:40:20.925 UTC [6346] LOG:  starting PostgreSQL 12.5 on x86_64-pc-linux-gnu, compiled by gcc (GCC) 4.8.5 20150623 (Red Hat 4.8.5-39), 64-bit
2020-12-30 03:40:20.926 UTC [6346] LOG:  listening on IPv6 address ":::1", port 5432
2020-12-30 03:40:20.926 UTC [6346] LOG:  listening on IPv4 address "127.0.0.1", port 5432
2020-12-30 03:40:20.932 UTC [6346] LOG:  listening on Unix socket "/var/run/postgresql/.s.PGSQL.5432"
2020-12-30 03:40:20.946 UTC [6346] LOG:  listening on Unix socket "/tmp/.s.PGSQL.5432"
2020-12-30 03:40:20.960 UTC [6346] LOG:  redirecting log output to logging collector process
2020-12-30 03:40:20.960 UTC [6346] HINT:  Future log output will appear in directory "log".
done
server started
```

 Check the following locations to be used during the following steps.

Item	Location
Unix socket	/tmp/.s.PGSQL.5432

## Protecting PostgreSQL Resource with LifeKeeper

1. Click  in the LifeKeeper user interface.
2. The Create Resource Wizard at node-a will appear. Select PostgreSQL Database as the Recovery Kit.
3. Select the following parameters.

Item	Location
Switchback Type	intelligent 

Server	node-a ✓
PostgreSQL Executable Location	/usr/pgsql-12/bin/
PostgreSQL Client Executable Location	/usr/pgsql-12/bin/psql ✓
PostgreSQL Administration Executable Location	/usr/pgsql-12/bin/pg_ctl ✓
PostgreSQL Data Directory	/datakeeper/pgsql/data
PostgreSQL Port	5432 ✓
PostgreSQL Socket Path	/tmp/.s.PGSQL.5432 ✓
Enter Database Administrator User	postgres
PostgreSQL Logfile	/tmp/pgsql-5432.lk.log ✓
PostgreSQL Database Tag	pgsql-5432 ✓

The wizard checks these values. Once “pgsql-5432” is successfully created on node-a, continue to the next steps.

4. Select the following values in the Pre-Extend Wizard @ node-a.

Item	Location
Target Server	node-b ✓
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓

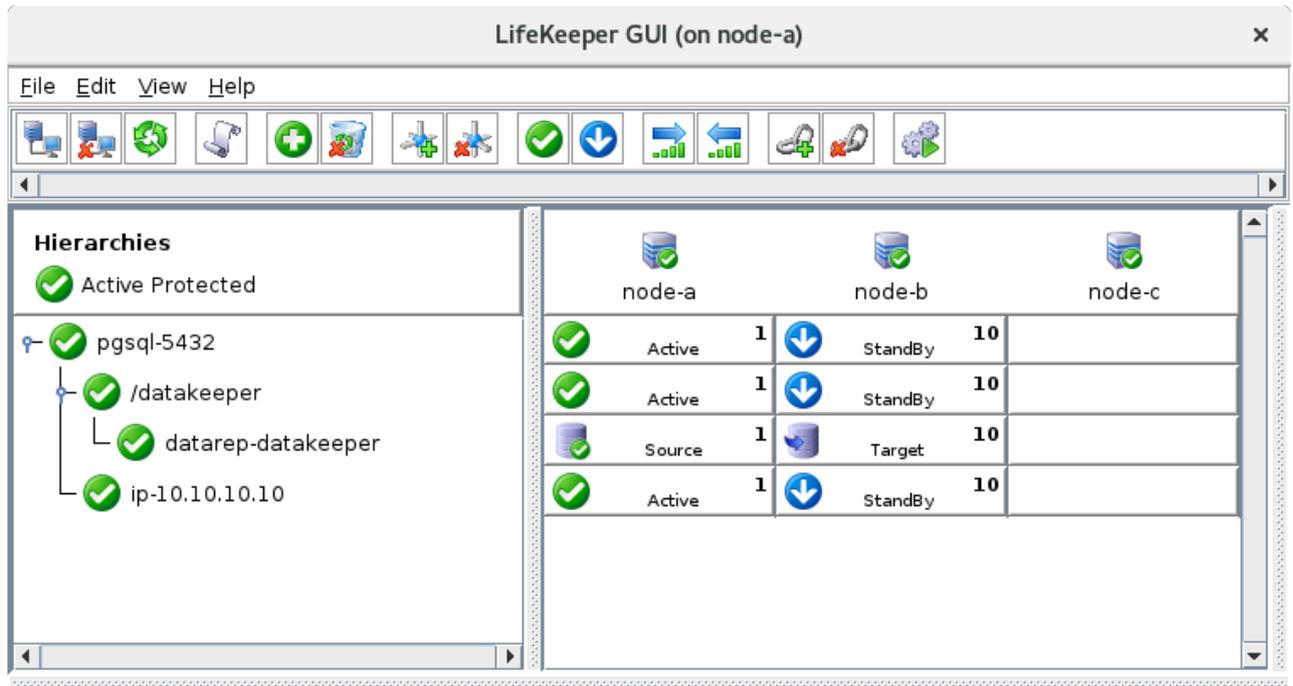
The wizard checks these values. Once the Pre-Extend check is complete, continue to the next steps.

5. Select the following values in the Extend database/pgsql Resource @ node-a wizard.

Field	Value
PostgreSQL Executable Location	/usr/pgsql-12/bin/ ✓
PostgreSQL Database Tag	pgsql-5432 ✓

Now the PostgreSQL resource is defined on LifeKeeper. The wizard automatically defines a

dependency between the PostgreSQL resource and the replicated disk (/datakeeper).



<-- node-b: ip-10.10.10.10: Adding parent: pgsq-5432

## 11.2.7.5.3.1. Install Postgres 12 on Linux Nodes

This section outlines the basic steps required to install Postgres 12 on Linux nodes.

✿ Postgres must be installed on both the active and standby nodes.

- Instructions for installing Postgres are available at <https://www.postgresql.org/download/>.

By selecting the Linux distribution, version, and architecture, customized installation instructions will be provided. Follow the instructions up to the point immediately before installing the product. Because `initdb` should be called differently between node-a and node-b, those steps are discussed on the parent page.

In the case of RedHat 8 or CentOS 8, execute the following sections:

```
# Install the repository RPM:
# dnf install -y https://download.postgresql.org/pub/repos/yum/reporpm/EL-8-x
86_64/pgdg-redhat-repo-latest.noarch.rpm
# Disable the built-in PostgreSQL module:
# dnf -qy module disable postgresql
# Install PostgreSQL:
# dnf install -y postgresql12-server
```

! Please stop at the step before executing `initdb`.

### Set Password to the postgres User

```
1 # passwd postgres
```

## Change Directory to Store Data to Under /datakeeper/pgsql/data

```
1 # cp -p /var/lib/pgsql/.bash_profile /var/lib/pgsql/.bash_profile.org
2 # vi /var/lib/pgsql/.bash_profile
3 # cat /var/lib/pgsql/.bash_profile
4 [ -f /etc/profile ] && source /etc/profile
5 PGDATA=/datakeeper/pgsql/data
6 export PGDATA
7 # If you want to customize your settings,
8 # Use the file below. This is not overridden
9 # by the RPMS.
10 [ -f /var/lib/pgsql/.pgsql_profile ] && source /var/lib/pgsql/.pgsql_profile
```

## 11.2.7.5.4. Protecting an NFS Resource

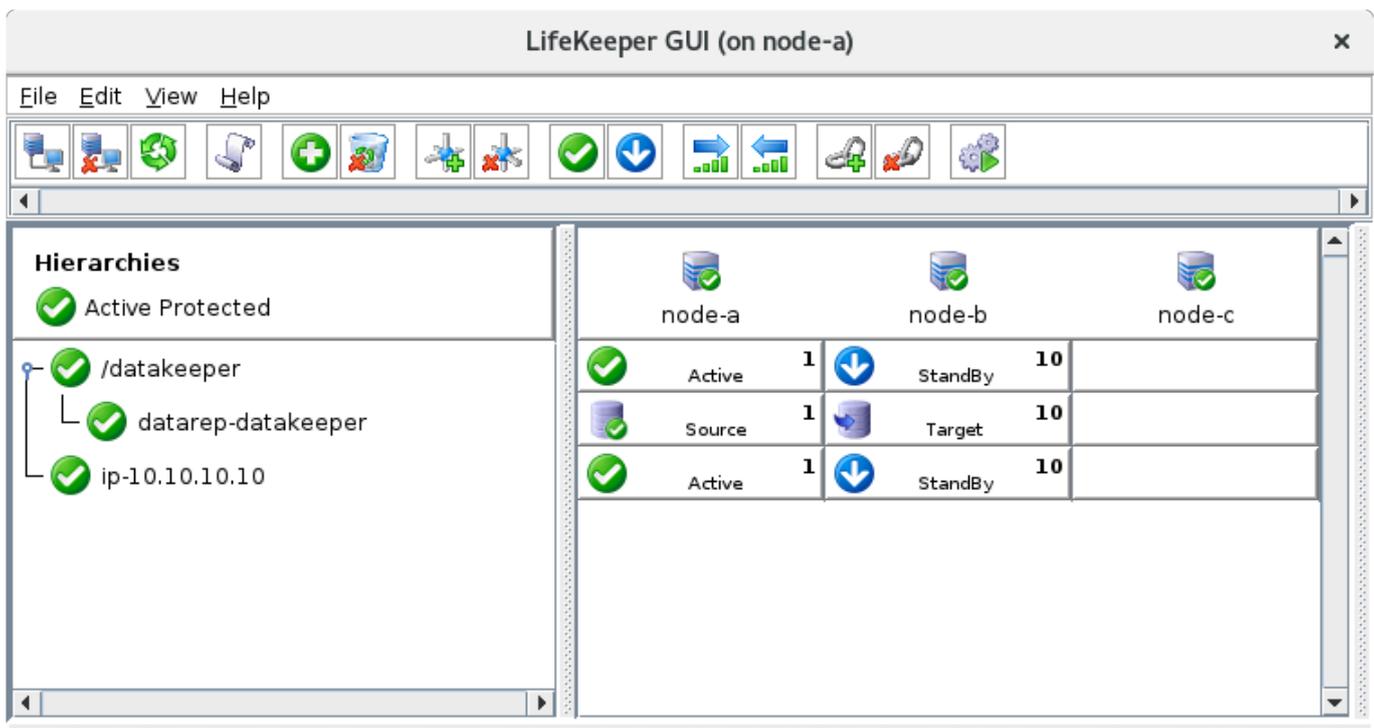
This section describes the steps required to protect a cluster of NFS servers.

### Before Beginning

This section assumes that the following 2 resources are configured on the cluster.

- [IP Resource](#)
- [Data Replication](#) at `/datakeeper`

The LifeKeeper GUI should look like the following prior to installing the NFS software.



<-- node-b: Updating app/res type: scsi i2o

## Install NFS Server Software

- ✿ Install NFS server and client utilities on both nodes. On RHEL 8.x these are provided by the `nfs-utils` package, and on SLES 15.x they are provided by the `nfs-kernel-server` package and its dependencies.

For example, on RHEL 8.x, install `nfs-utils` on both node-a and node-b:

```
1 # yum -y install nfs-utils
2 (snip)
3 Complete!
```

## Configure the NFS Server

Start the required NFS-related services on both node-a and node-b and enable them so that they start automatically on boot.

```
# systemctl enable --now rpcbind nfs-server
```

 The following steps should be performed on node-a **only**.

Edit `/etc/exports` to define the NFS export. We are using the shared storage (`/datakeeper`) that is protected by DataKeeper. Create `/datakeeper/nfs/data` and use it as the location of the export.

Once the `/etc/exports` file has been edited it should look like the following:

```
1 [root@node-a ~]# mkdir -p /datakeeper/nfs/data
2 [root@node-a ~]# vi /etc/exports
3 [root@node-a ~]# cat /etc/exports
4 /datakeeper/nfs/data 10.20.0.0/22(rw,no_root_squash)
```

Export the shared file system that was added `/etc/exports` on node-a:

```
[root@node-a ~]# exportfs -rav
```

## Confirm Access to the NFS Server from a Client

 The following steps should be done on a client machine.

On the client machine (e.g., node-c), complete the following steps to install the NFS software. Mount the NFS export from node-a to a local folder (e.g., mount to `/local/nfsclient`).

```
1 [root@node-c ~]# yum -y install nfs-utils
2 (snip)
3 Complete!
4 [root@node-c /]# mkdir -p /local/nfsclient
5 [root@node-c ~]# mount -t nfs -o nfsvers=4 node-a:/datakeeper/nfs/data /local/nfsclient
```

Now the content on the NFS server should be visible. Disconnect (unmount) from the NFS server to configure the NFS cluster.

```
1 [root@node-c ~]# umount /local/nfsclient
```

## Protecting an NFS Resource using LifeKeeper

\* Complete the follow these steps on node-a.

1. Select  on the LifeKeeper user interface.
2. The Create Resource Wizard at node-a screen will appear. Select the NFS Recovery Kit as the Recovery Kit.



Please Select Recovery Kit

3. Select the following parameters.

Item	Location
Switchback Type	intelligent 
Server	node-a 
Export Point	/datakeeper/nfs/data
IP Tag	ip-10.10.10.10 

NFS Tag	nfs-/datakeeper/nfs/data 
---------	------------------------------------------------------------------------------------------------------------

The wizard checks these values. Once “nfs-/datakeeper/nfs/data” is successfully created on node-a, continue to the next steps.

4. Select the following values on Pre-Extend Wizard @ node-a.

Field	Value
Target Server	node-b 
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 

The wizard checks these values. Once the Pre-Extend Check is completed, move on to the next steps.

5. Select the following values on the Extend gen/nfs Resource Hierarchy @ node-a wizard.

Field	Value
NFS Tag	nfs-/datakeeper/nfs/data 

The NFS resource hierarchy is now defined in LifeKeeper. The wizard automatically defines a dependency between the NFS resource and required resources (IP resource and /datakeeper resource).

**LifeKeeper GUI (on node-a)**

File Edit View Help

**Hierarchies**

- Active Protected
- nfs-/datakeeper/nfs/data
  - ip-10.10.10.10
    - hanfs-/datakeeper/nfs/data
      - /datakeeper
        - datarep-datakeeper

	node-a	node-b	node-c
Active	1	StandBy	10
Active	1	StandBy	10
Active	1	StandBy	10
Active	1	StandBy	10
Source	1	Target	10

<-- node-w: Updating server state to: alive

## 11.2.7.5.5. Protecting SAP Resources

---

In this section we will deploy an SAP S/4HANA 1909 AS ABAP environment and use the LifeKeeper SAP Recovery Kit to protect the ABAP SAP Central Services (ASCS) and corresponding Enqueue Replication Server (ERS) instances. While it is possible to protect a redundant Primary Application Server (PAS) instance using LifeKeeper, we will instead take the approach of creating an external application server pool consisting of one PAS instance and one Additional Application Server (AAS) instance spread across two availability zones.

Since the PAS and AAS instances will not be protected by LifeKeeper, this guide only details the installation and configuration of the highly-available Central Services (ASCS/ERS) cluster. This section also does not describe the configuration of a highly-available database cluster (e.g., using SAP HANA). Please see the section of the guide relevant to the database that is being installed and adapt it to fit your planned SAP landscape.

The details of the Central Services cluster nodes that will be created are given below. The instance provisioning described in this guide gives the **minimum** required specifications to run the applications in an evaluation context, and is not generally sufficient to handle productive workloads. Consult the documentation provided by SAP and your cloud provider for best practices when provisioning VM's in a production environment. For example, see [Hardware Requirements for SAP System Hosts](#).

The operating system and configuration used in the deployment must be supported by SAP, SIOS, and your cloud provider. Consult SAP's Product Availability Matrix (PAM) in order to determine which operating systems are supported for various versions of SAP NetWeaver and S/4HANA.

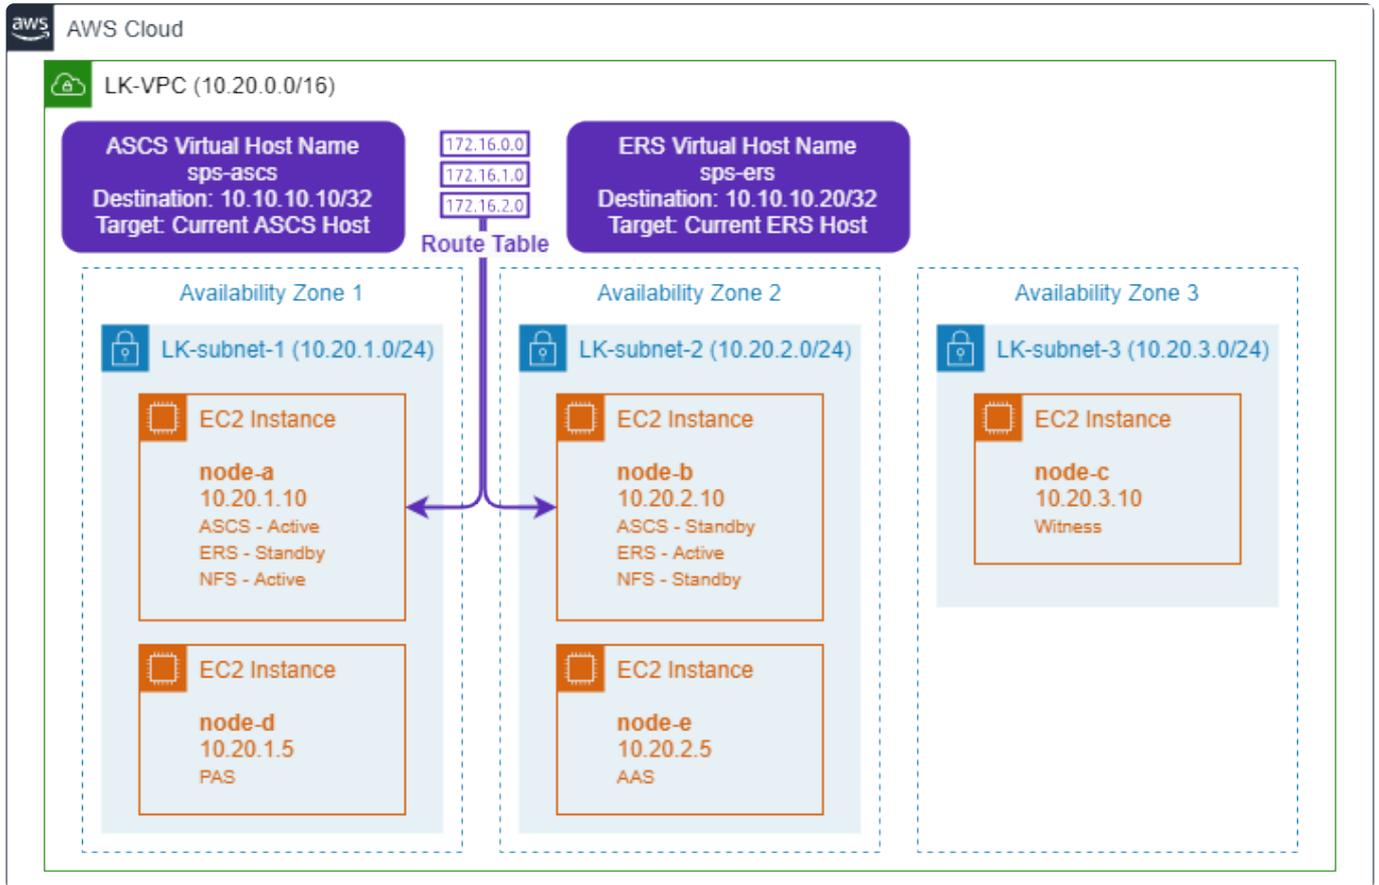
### Cluster Architecture for Evaluation Deployment

- node-a (Primary ASCS instance host, Standby ERS instance host, Primary NFS server)
  - Private IP: 10.20.1.10
  - vCPU's: 2
  - Memory: 4GB
  - SAP NFS shares mounted
  - SAP instances ASCS10 and ERS20 installed under SID 'SPS'
  - LifeKeeper for Linux installed with quorum/witness functionality enabled and the following recovery kits installed:
    - SAP Recovery Kit
    - SIOS DataKeeper
    - NFS Recovery Kit
    - [AWS] EC2 Recovery Kit
    - [Azure, Google Cloud] Generic Application Recovery Kit for Load Balancer Health Checks (GenLB Recovery Kit)
  
- node-b (Primary ERS instance host, Standby ASCS instance host, Standby NFS server)
  - Private IP: 10.20.2.10
  - vCPU's: 2
  - Memory: 4GB
  - SAP NFS shares mounted

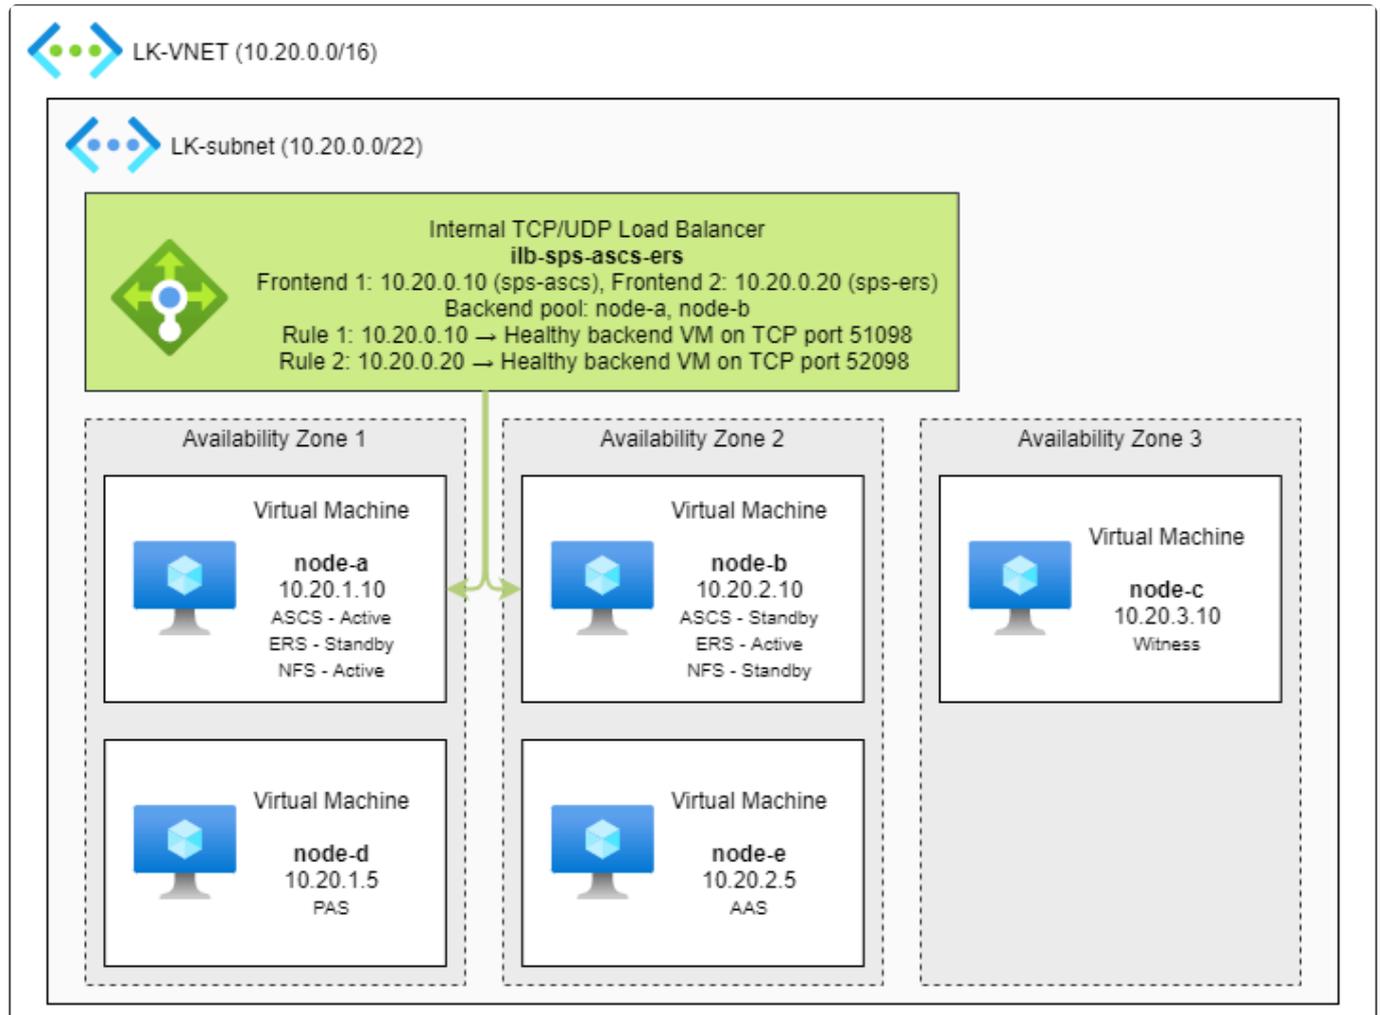
- SAP instances ASCS10 and ERS20 installed under SID 'SPS'
- LifeKeeper for Linux installed with quorum/witness functionality enabled and the following recovery kits installed:
  - SAP Recovery Kit
  - SIOS DataKeeper
  - NFS Recovery Kit
  - [AWS] EC2 Recovery Kit
  - [Azure, Google Cloud] Generic Application Recovery Kit for Load Balancer Health Checks (GenLB Recovery Kit)
- node-c (Quorum/Witness node)
  - Private IP: 10.20.3.10
  - vCPU's: 1
  - Memory: 2GB
  - LifeKeeper for Linux installed with quorum/witness functionality enabled
- node-d (Primary Application Server host)
  - Private IP: 10.20.1.5
  - vCPU's: 2
  - Memory: 4GB
  - SAP NFS shares mounted
  - SAP PAS instance D01 installed under SID 'SPS'
- node-e (Additional Application Server host)
  - Private IP: 10.20.2.5
  - vCPU's: 2
  - Memory: 4GB
  - SAP NFS shares mounted
  - SAP AAS instance D02 installed under SID 'SPS'
- ASCS10 virtual hostname: sps-asc
- ERS20 virtual hostname: sps-ers
- NFS shared file systems:
  - /sapmnt/SPS
  - /usr/sap/trans
- SIOS DataKeeper replicated file systems:
  - /usr/sap/SPS/ASCS10
  - /usr/sap/SPS/ERS20

# Cloud-Specific Architecture Diagrams

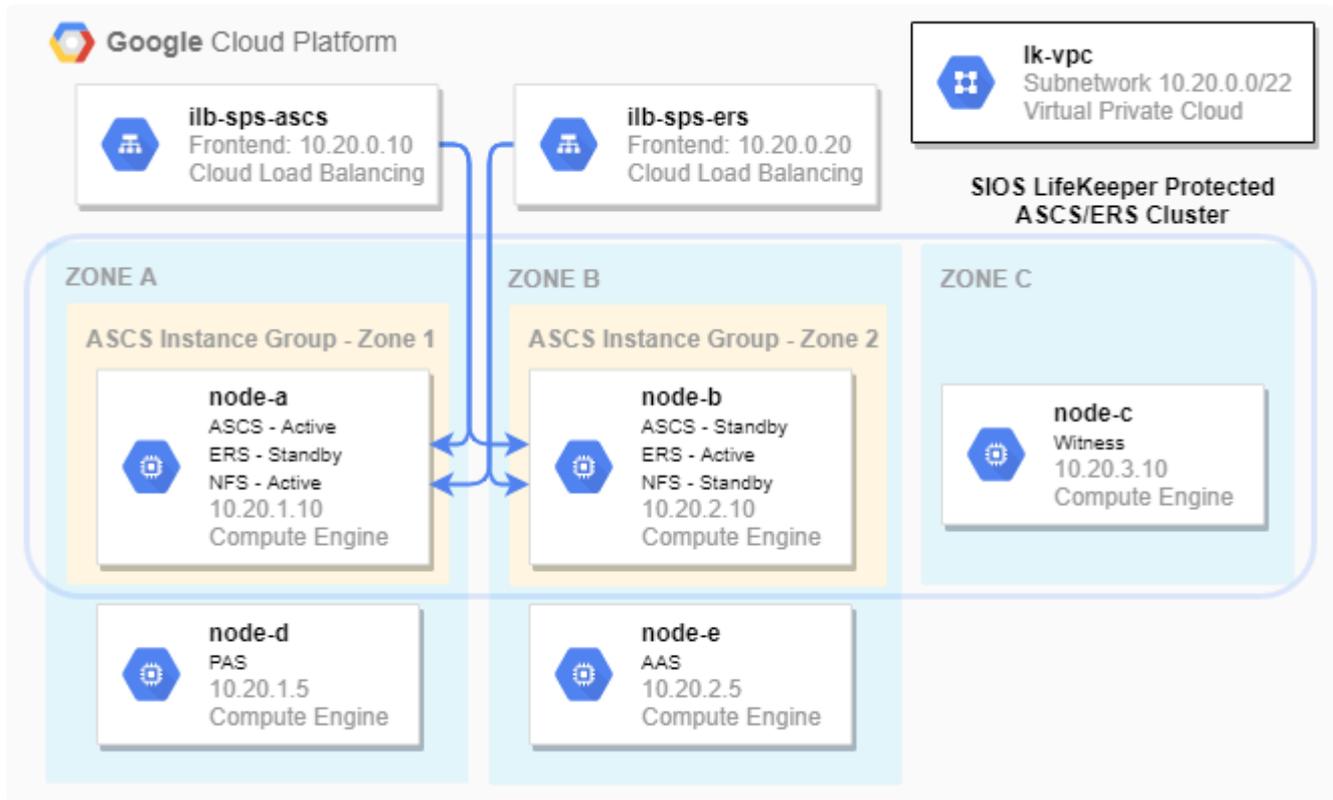
## SIOS Protected ASCS/ERS Cluster on Amazon Web Services (AWS)



## SIOS Protected ASCS/ERS Cluster on Microsoft Azure



## SIOS Protected ASCS/ERS Cluster on Google Cloud



### Prerequisites

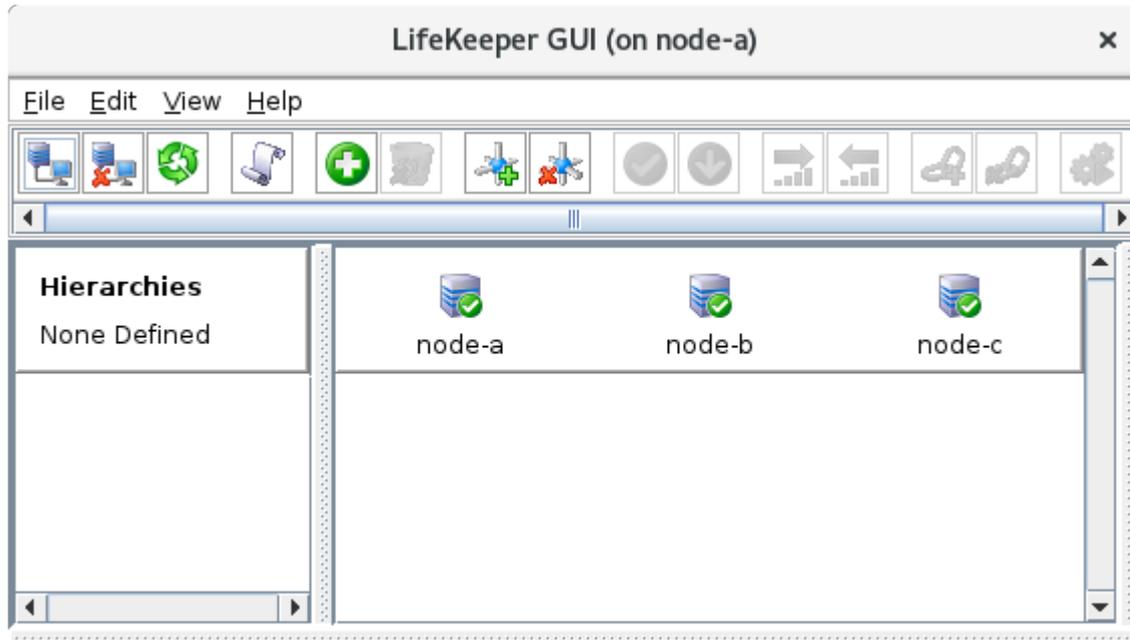
In order to follow this guide, the following prerequisite steps must already be complete:

1. Three VM instances (node-a, node-b, and node-c) have been created using the networking conventions described earlier in the guide. Firewall rules are in place to allow inter-node communication as well as SSH connections. See [Configuring Network Components and Creating Instances](#) for details.
2. All three VM instances have been configured to run LifeKeeper for Linux. In particular, SELinux is disabled, the firewall is disabled, the /etc/hosts file on each node contains entries to resolve each node's hostname to its private IP, and the root user has the same password on each node. See [Configure Linux Nodes to Run LifeKeeper for Linux](#) for details.
3. LifeKeeper for Linux has been installed on all three nodes with quorum/witness functionality enabled, and all additional required recovery kits have been installed on node-a and node-b. On node-c (the witness node), no additional recovery kits beyond the core LifeKeeper installation are required. All necessary SIOS licenses have been installed on each node. See [Install LifeKeeper for Linux](#) for details.

**Note:** Since the required recovery kits are installed only on node-a and node-b, all steps in this guide that are performed through the LifeKeeper GUI must be performed on either node-a or node-b.

- LifeKeeper communication paths have been defined between all pairs of cluster nodes. Note that this requires creation of **three** bi-directional communication paths (node-a ↔ node-b, node-a ↔ node-c, node-b ↔ node-c). See [Login and Basic Configuration Tasks](#) for details.

After completing all of these tasks, the LifeKeeper GUI should resemble the following image.



<-- node-c: Adding app/res: scsi device

## 11.2.7.5.5.1. Create ASCS and ERS Virtual IPs

---

As of SAP AS ABAP release 7.53 (ABAP Platform 1809) and the introduction of the [Enqueue Standalone Framework v2](#) (ENSAv2), the ASCS and ERS instances communicate with each other via associated virtual host names. The virtual host names sps-ascs and sps-ers will be configured to resolve to the current hosts of the ASCS10 and ERS20 instances, respectively.

The recommended implementation of these virtual IP addresses varies by cloud platform. Please follow the steps provided in the section corresponding to your cloud platform:

- [AWS – Create ASCS and ERS Virtual IPs](#)
- [Azure – Create ASCS and ERS Internal Load Balancer](#)
- [Google Cloud – Create ASCS and ERS Internal Load Balancers](#)

# 11.2.7.5.5.1.1. AWS – Create ASCS and ERS Virtual IPs

Before creating resources to protect the virtual IP addresses associated to the ASCS and ERS instances, the following steps must be completed on node-a and node-b:

- [Install AWS CLI](#)
- [Assign Permission to Use EC2 Recovery Kit](#)
- [Disable PING Broadcasting](#)
- [AWS – Disable Source/Destination Checking](#)

We will now follow the process described in [Creating an AWS EC2 Resource \(RouteTable Scenario\)](#) to create resource hierarchies for the virtual IP addresses associated to the ASCS and ERS instances.

Since the IP addresses used with the EC2 RouteTable resource type must be located outside of the VPC CIDR (which for LK-VPC is 10.20.0.0/16), we will use the following virtual host names and IP addresses:

Instance	Virtual Host Name	IP Address
ASCS10	sps-ascs	10.10.10.10
ERS20	sps-ers	10.10.10.20

## Create the ASCS Virtual IP Resource Hierarchy

1. Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend a LifeKeeper IP resource (**ip-sps-ascs**) to protect the ASCS virtual IP address on node-a and node-b. Notice that the IP resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The

 icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
IP Resource	10.10.10.10
Netmask	255.255.255.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-ascs
<b>Pre-Extend Wizard</b>	

Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.10.10.10 
Netmask	255.255.255.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-ascs 

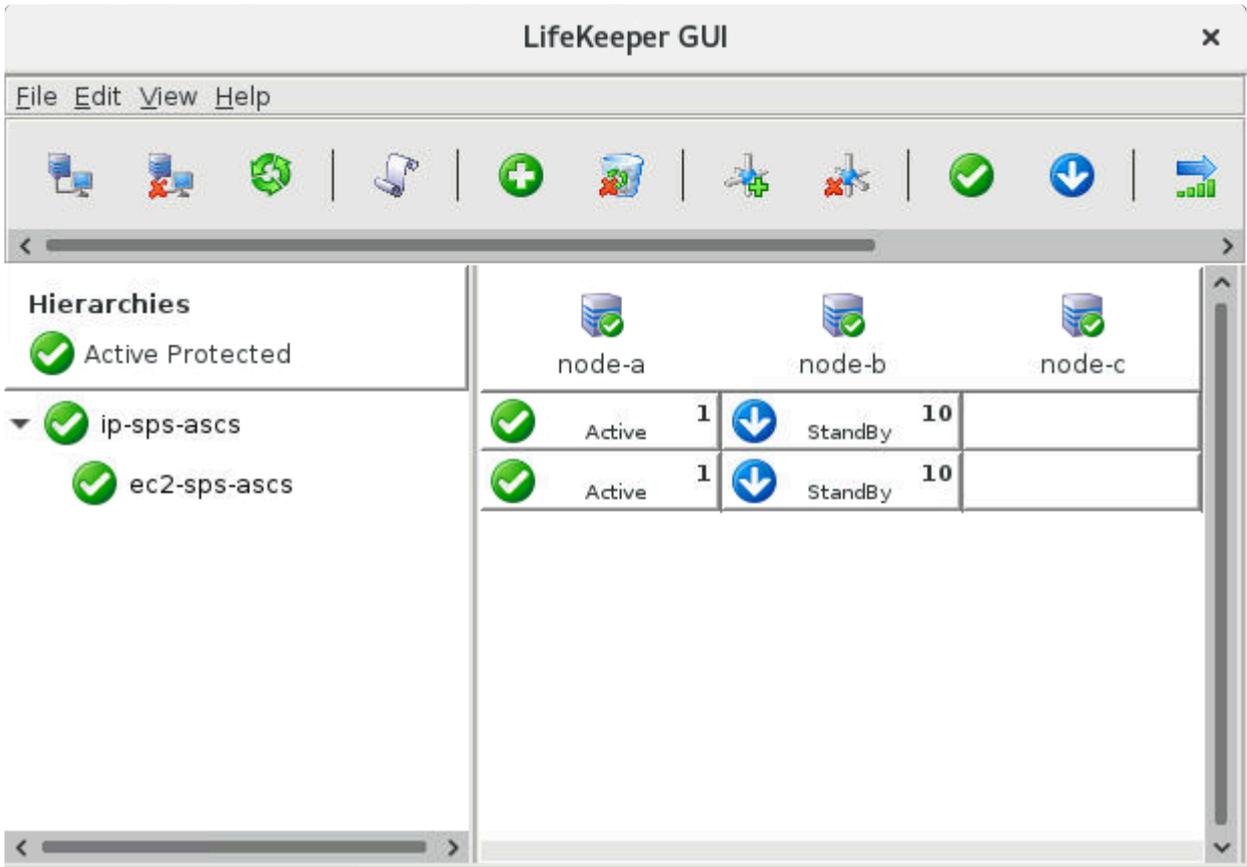
- Following the steps described in [Creating an AWS EC2 Resource \(RouteTable Scenario\)](#), use the following parameters to create and extend a LifeKeeper EC2 resource (**ec2-sps-ascs**) to manage the backend manipulation of the route table in AWS when the IP resource is switched over. Notice that the EC2 resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
EC2 Resource type	RouteTable (Backend cluster) 
IP Resource	ip-sps-ascs 
EC2 Resource Tag	ec2-sps-ascs
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ec2 Resource Hierarchy Wizard</b>	

EC2 Resource Tag	ec2-sps-ascs 
------------------	------------------------------------------------------------------------------------------------

 **Note:** If creation or extension of the EC2 resource fails, verify that (i) a route directing traffic from 10.10.10.10/32 to Node-A exists in the route table for LK-VPC in AWS, (ii) source/destination checks have been disabled on Node-A and Node-B, and (iii) that the priority values for the dependent IP resource **ip-sps-ascs** are 1 on Node-A and 10 on Node-B, respectively.

Once the IP resource hierarchy has been successfully created, the LifeKeeper GUI should resemble the following image:



## Create the ERS Virtual IP Resource Hierarchy

- Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend a LifeKeeper IP resource (**ip-sps-ers**) to protect the ERS virtual IP address on node-a and node-b. Notice that the IP resource is being created on node-b and extended to node-a. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
Create Resource Wizard	

Switchback Type	intelligent 
Server	node-b
IP Resource	10.10.10.20
Netmask	255.255.255.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-ers
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.10.10.20 
Netmask	255.255.255.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-ers 

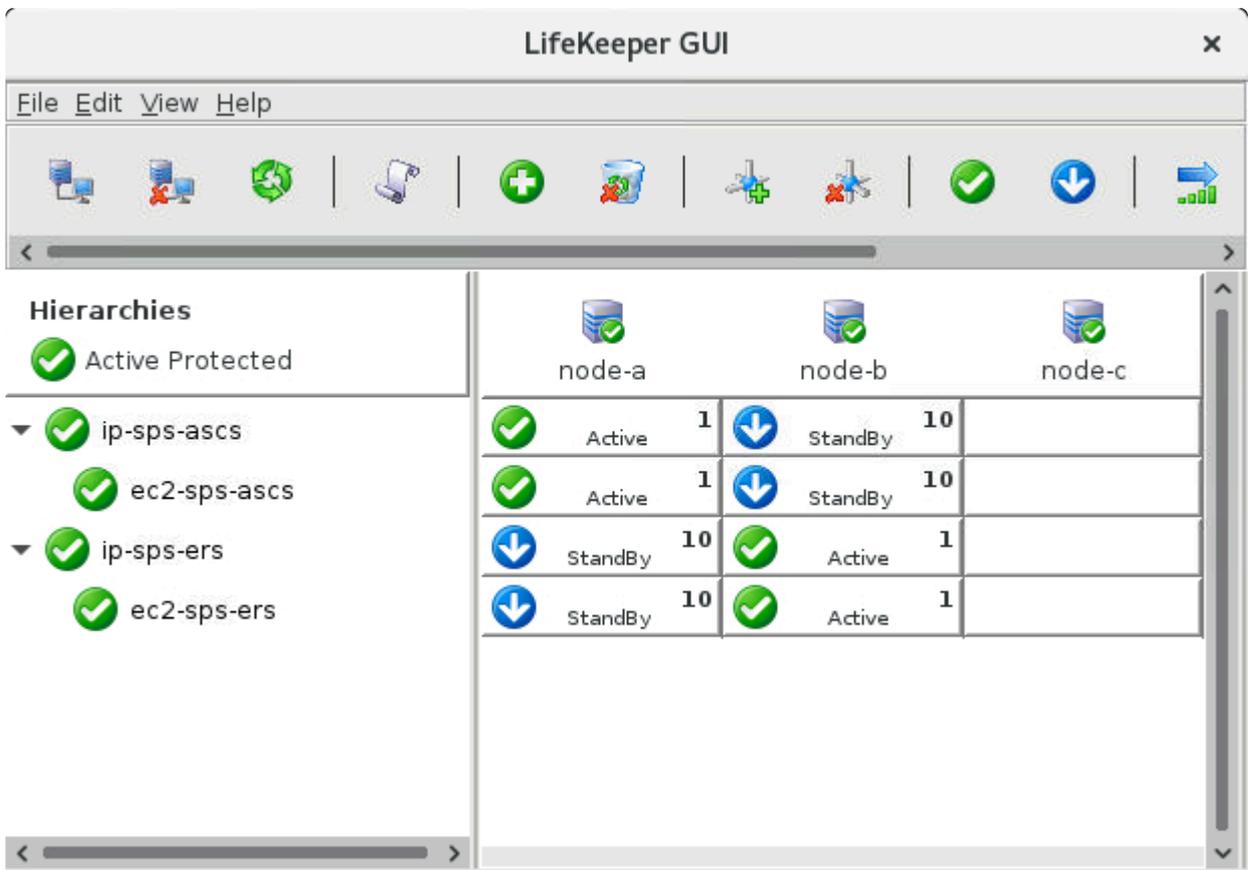
- Following the steps described in [Creating an AWS EC2 Resource \(RouteTable Scenario\)](#), use the following parameters to create and extend a LifeKeeper EC2 resource (**ec2-sps-ers**) to manage the backend manipulation of the route table in AWS when the IP resource is switched over. Notice that the EC2 resource is being created on node-b and extended to node-a. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
EC2 Resource type	RouteTable (Backend cluster) 
IP Resource	ip-sps-ers 

EC2 Resource Tag	ec2-sps-ers
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
<b>Extend comm/ec2 Resource Hierarchy Wizard</b>	
EC2 Resource Tag	ec2-sps-ers ✓

✿ **Note:** If creation or extension of the EC2 resource fails, verify that (i) a route directing traffic from 10.10.10.20/32 to Node-B exists in the route table for LK-VPC in AWS, (ii) source/destination checks have been disabled on Node-A and Node-B, and (iii) that the priority values for the dependent IP resource **ip-sps-ers** are 10 on Node-A and 1 on Node-B, respectively.

Once the IP resource hierarchy has been successfully created, the LifeKeeper GUI should resemble the following image:



## Add Entries to /etc/hosts

Add the following entries to /etc/hosts on node-a and node-b to allow resolution of the virtual host names to their corresponding IP addresses. The entry for sps-ascs must also be added on node-d and node-e (the PAS and AAS hosts), as this virtual host name will be used to access the /export/sapmnt/SPS and /export/usr/sap/trans shared file systems.

```
10.10.10.10    sps-ascs
10.10.10.20    sps-ers
```

## 11.2.7.5.5.1.2. Azure – Create ASCS and ERS Internal Load Balancer

In Microsoft Azure, a TCP Internal Load Balancer is used to facilitate failover of the ASCS and ERS virtual IP's. The load balancer consists of two frontends, each assigned an IP address from the subnet that it operates in. In this example, we will be using the following IP addresses for the load balancer frontends corresponding to each virtual host:

Instance	Virtual Host Name	Internal TCP Load Balancer	Frontend IP Address
ASCS10	sps-ascs	ilb-sps-ascs-ers	10.20.0.10
ERS20	sps-ers	ilb-sps-ascs-ers	10.20.0.20

Following the steps provided in [Azure – Using an Internal Load Balancer](#) and [Responding to Load Balancer Health Checks](#), create and configure the following Azure and LifeKeeper resources.

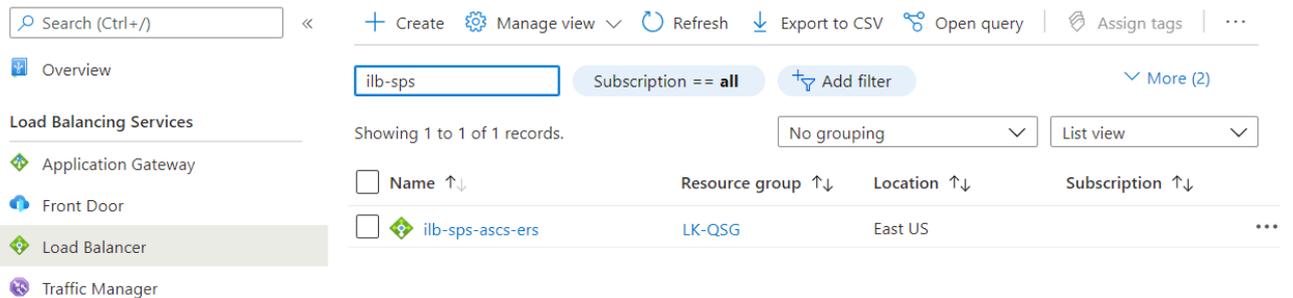
1. Create a single Load Balancer in Azure, **ilb-sps-ascs-ers**, with the following properties.

Name	ilb-sps-ascs-ers
<b>Create Load Balancer</b>	
Resource Group	LK-QSG
Name	ilb-sps-ascs-ers
Region	(same as the Virtual Machines)
Type	Internal
SKU	Standard (select Standard as the workload is distributed across Availability Zones)
Tier	Regional
<b>Frontend IP 1 Configuration</b>	
Name	SPSASCSFrontEnd
Virtual Network	LK-VNET
Subnet	LK-subnet (10.20.0.0/22)
Assignment	Static
IP address	10.20.0.10
Availability zone	Zone-redundant
<b>Frontend IP 2 Configuration</b>	
Name	SPSERSFrontEnd

Virtual Network	LK-VNET
Subnet	LK-subnet (10.20.0.0/22)
Assignment	Static
IP address	10.20.0.20
Availability zone	Zone-redundant
<b>Backend Pool</b>	
Name	backend-sps-ascs-ers
Backend Pool Configuration	NIC
IP Version	IPv4
Virtual machines	node-a, node-b
<b>Health Probe #1</b>	
Name	probe-sps-ascs
Protocol	TCP
Port	51098 (This must agree with the corresponding LifeKeeper GenLB resource created later)
Interval	5 seconds
Unhealthy threshold	2 consecutive failures
<b>Health Probe #2</b>	
Name	probe-sps-ers
Protocol	TCP
Port	52098 (This must agree with the corresponding LifeKeeper GenLB resource created later)
Interval	5 seconds
Unhealthy threshold	2 consecutive failures
<b>Load Balancing Rule #1</b>	
Name	ilb-rule-sps-ascs
IP Version	IPv4
Frontend IP address	SPSASCSFrontEnd (10.20.0.10)
HA Ports	Click (allow forwarding to all ports for simplicity of evaluation deployment)
Backend pool	backend-sps-ascs-ers
Health probe	probe-sps-ascs (TCP:51098)

Session persistence	None
Idle timeout	4 minutes
TCP reset	Disabled
Floating IP	Enabled
<b>Load Balancing Rule #2</b>	
Name	ilb-rule-sps-ers
IP Version	IPv4
Frontend IP address	SPSERSFrontEnd (10.20.0.20)
HA Ports	Click (allow forwarding to all ports for simplicity of evaluation deployment)
Backend pool	backend-sps-ascs-ers
Health probe	probe-sps-ers (TCP:52098)
Session persistence	None
Idle timeout	4 minutes
TCP reset	Disabled
Floating IP	Enabled

Once created, the **Load Balancing Services** → **Load Balancer** page in the Microsoft Azure Console will show the newly created load balancer.



- Following the steps in [Responding to Load Balancer Health Checks](#), install the LifeKeeper Generic Application Recovery Kit for Load Balancer Health Checks (“GenLB Recovery Kit”) and create two LifeKeeper GenLB resources, **ilb-sps-ascs** and **ilb-sps-ers**, with the following properties. Notice that **ilb-sps-ascs** is created on node-a and extended to node-b, while **ilb-sps-ers** is created on node-b and extended to node-a. The  icon indicates that the default option is chosen.

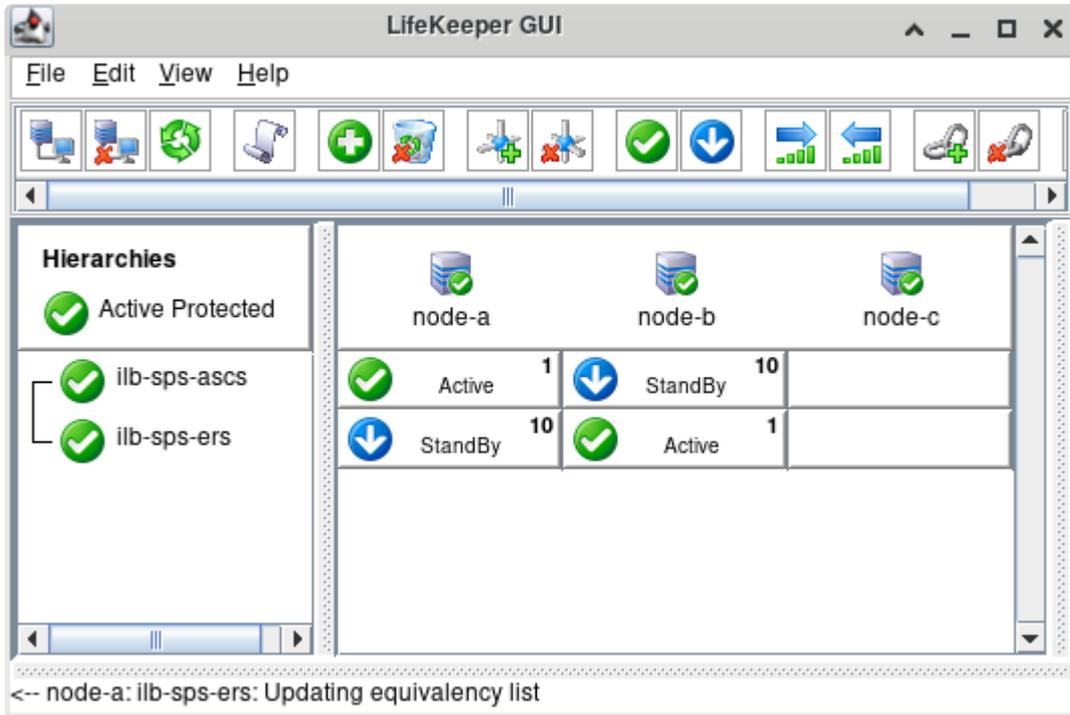
Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 

Server	node-a
Restore Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/restore.pl
Remove Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/remove.pl
QuickCheck Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/quickCheck <b>(Note:</b> Although the quickCheck script may be optional for some 'Generic Application' resource types, it is required for GenLB resources.)
Local Recovery Script	None (Empty)
Application Info	51098
Bring Resource In Service	Yes
Resource Tag	ilb-sps-ascs
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Resource Tag	ilb-sps-ascs 
Application Info	51098 

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
Restore Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/restore.pl
Remove Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/remove.pl
QuickCheck Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/quickCheck <b>(Note:</b> Although the quickCheck script may be optional for some 'Generic

	Application' resource types, it is required for GenLB resources.)
Local Recovery Script	None (Empty)
Application Info	52098
Bring Resource In Service	Yes
Resource Tag	ilb-sps-ers
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Resource Tag	ilb-sps-ers 
Application Info	52098 

The resources will appear in the LifeKeeper GUI resource pane once they have been created and extended successfully.



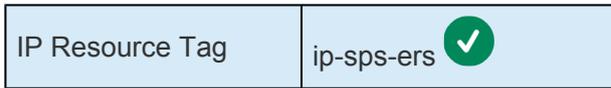
- Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend a LifeKeeper IP resource (**ip-sps-ascs**) to protect the ASCS virtual IP address on node-a and node-b. Notice that the IP resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
IP Resource	10.20.0.10
Netmask	255.255.252.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-ascs
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 

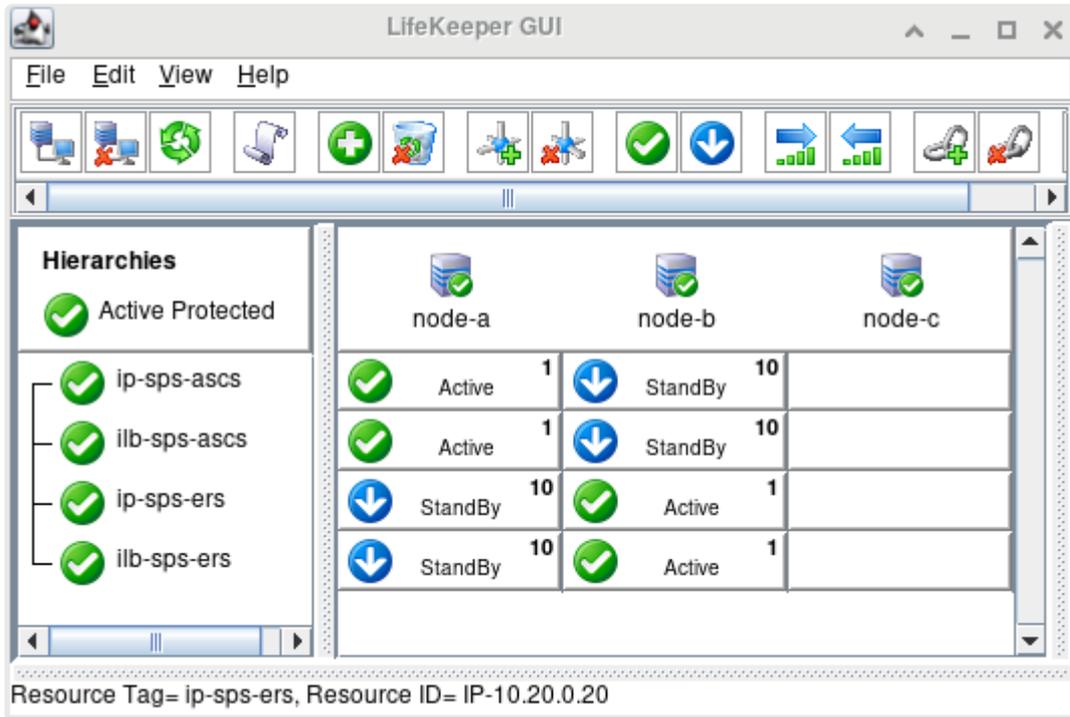
Extend comm/ip Resource Hierarchy Wizard	
IP Resource	10.20.0.10 
Netmask	255.255.252.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-ascs 

4. Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend a LifeKeeper IP resource (**ip-sps-ers**) to protect the ERS virtual IP address on node-a and node-b. Notice that the IP resource is being created on node-b and extended to node-a. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
IP Resource	10.20.0.20
Netmask	255.255.252.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-ers
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.20.0.20 
Netmask	255.255.252.0 
Network Interface	eth0 

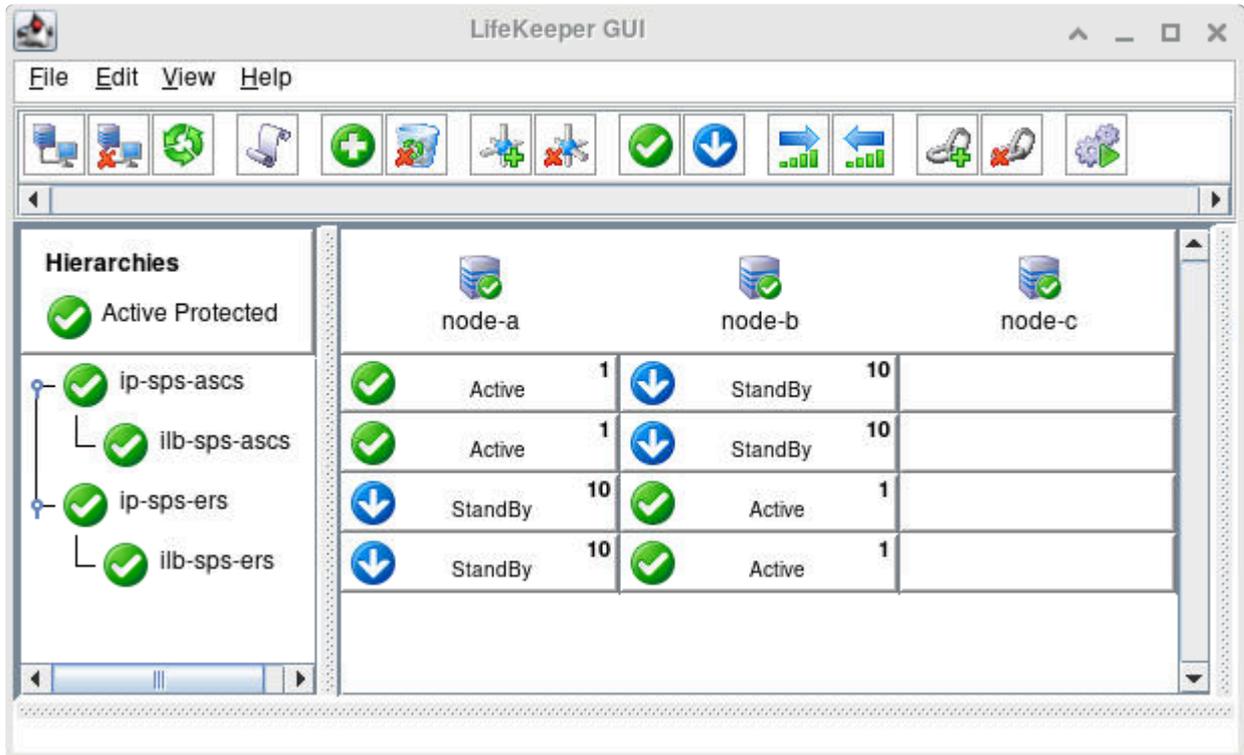


Once the IP resources have been created successfully, the LifeKeeper GUI should resemble the following image.



- Right-click the **ip-sps-ascs** resource on node-a and click “Create Dependency...” Specify **ilb-sps-ascs** as the Child Resource Tag and click **Create Dependency**.
- Right-click the **ip-sps-ers** resource on node-b and click “Create Dependency...” Specify **ilb-sps-ers** as the Child Resource Tag and click **Create Dependency**.

Once the dependencies have been created successfully, the LifeKeeper GUI should resemble the following image.



7. Add the following entries to `/etc/hosts` on node-a and node-b to allow resolution of each virtual host name to the frontend IP address of the corresponding load balancer. The entry for `sps-ascs` must also be added on node-d and node-e (the PAS and AAS hosts), as this virtual host name will be used to access the `/export/sapmnt/SPS` and `/export/usr/sap/trans` shared file systems.

```
10.20.0.10 sps-ascs
10.20.0.20 sps-ers
```

8. Test switchover and failover of the GenLB resource hierarchies as described in the **Test GenLB Resource Switchover and Failover** section of [Responding to Load Balancer Health Checks](#). Correct any issues found or tune the parameters of the load balancer health checks as required to achieve successful operation.

# 11.2.7.5.5.1.3. Google Cloud – Create ASCS and ERS Internal Load Balancers

In Google Cloud, TCP Internal Load Balancers are used to facilitate failover of the ASCS and ERS virtual IP's. The frontend of each load balancer is assigned an ephemeral IP from the subnet that it operates in. In this example, we will be using the following IP addresses for load balancers corresponding to each virtual host:

Instance	Virtual Host Name	Internal TCP Load Balancer	Frontend IP Address
ASCS10	sps-ascs	ilb-sps-ascs	10.20.0.10
ERS20	sps-ers	ilb-sps-ers	10.20.0.20

Following the steps provided in [Google Cloud – Using an Internal Load Balancer](#) and [Responding to Load Balancer Health Checks](#), create and configure the following Google Cloud and LifeKeeper resources.

1. Create two unmanaged instance groups: **ig-sps-ascs-zone1**, containing node-a, and **ig-sps-ascs-zone2**, containing node-b.

Instance groups CREATE INSTANCE GROUP REFRESH DELETE

Instance groups are collections of VM instances that use load balancing and automated services, like autoscaling and autohealing. [Learn more](#)

☰ Enter property name or value

<input type="checkbox"/>	<input checked="" type="checkbox"/>	Name ↑	Instances	Template	Group type	Creation time	Recommendation	Autoscaling	Zone
<input type="checkbox"/>	<input checked="" type="checkbox"/>	ig-sps-ascs-zone1	1	-	Unmanaged	Feb 28, 2021, 9:34:55 PM UTC-05:00			us-east1-b
<input type="checkbox"/>	<input checked="" type="checkbox"/>	ig-sps-ascs-zone2	1	-	Unmanaged	Feb 28, 2021, 9:35:25 PM UTC-05:00			us-east1-c

2. Create two TCP Internal Load Balancers, **ilb-sps-ascs** and **ilb-sps-ers**, with the following properties. The  icon indicates that the default option is chosen.

Name	ilb-sps-ascs
<b>Backend configuration</b>	
Region	<Deployment region> (e.g., us-east1)
Network	lk-vpc
Backends	ig-sps-ascs-zone1 ig-sps-ascs-zone2
<b>Health Check</b>	
Name	hc-sps-ascs

Protocol	TCP
Port	51098 (This must agree with the corresponding LifeKeeper GenLB resource created later)
Proxy protocol	NONE
Request	Leave empty
Response	Leave empty
Check interval	5 seconds
Timeout	5 seconds
Healthy threshold	2 consecutive successes
Unhealthy threshold	2 consecutive failures
Connection draining timeout	15 seconds
<b>Frontend configuration</b>	
Name	fe-sps-ascs
Subnetwork	lk-subnet
Purpose	Non-shared
IP address	Ephemeral (Custom)
Custom ephemeral IP address	10.20.0.10
Ports	All
Global access	Disable
Service label	Leave empty

Name	ilb-sps-ers
<b>Backend configuration</b>	
Region	<Deployment region> (e.g., us-west1)
Network	lk-vpc
Backends	ig-sps-ascs-zone1 ig-sps-ascs-zone2

Health Check	
Name	hc-sps-ers
Protocol	TCP 
Port	52098 (This must agree with the corresponding LifeKeeper GenLB resource created later)
Proxy protocol	NONE 
Request	Leave empty 
Response	Leave empty 
Check interval	5 seconds 
Timeout	5 seconds 
Healthy threshold	2 consecutive successes 
Unhealthy threshold	2 consecutive failures 
Connection draining timeout	15 seconds
Frontend configuration	
Name	fe-sps-ers
Subnetwork	lk-subnet
Purpose	Non-shared 
IP address	Ephemeral (Custom)
Custom ephemeral IP address	10.20.0.20
Ports	All
Global access	Disable 
Service label	Leave empty 

Once created, the **Network services** → **Load balancing** page in the Google Cloud Console will show the two load balancers.

☰ Filter by name or protocol

<input type="checkbox"/>	Name	Protocol ^	Region	Backends
<input type="checkbox"/>	ilb-sps-ascs	TCP (Internal)	us-east1	⚠️ 1 regional backend service (2 instance groups) ⋮
<input type="checkbox"/>	ilb-sps-ers	TCP (Internal)	us-east1	⚠️ 1 regional backend service (2 instance groups) ⋮

- Follow the steps described in the **Disable IP Forwarding** section of [Google Cloud – Using an Internal Load Balancer](#) to allow node-a and node-b to communicate through the frontend IP addresses of the load balancers. Reboot node-a and node-b for the changes to take effect.
- Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend two LifeKeeper IP resources (**ip-sps-ascs** and **ip-sps-ers**), each with net mask 255.255.255.255, to protect the ASCS and ERS virtual IP addresses on node-a and node-b. Notice that **ip-sps-ascs** is created on node-a and extended to node-b, while **ip-sps-ers** is created on node-b and extended to node-a. Also note that the resulting resources should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

 The use of 255.255.255.255 as the value of the Netmask parameter (both during create and extend) is important. Be careful not to select “Accept Defaults” during the extend process as this will typically cause 255.255.255.0 to be incorrectly used as the network mask on the standby node, which may lead to unexpected networking issues.

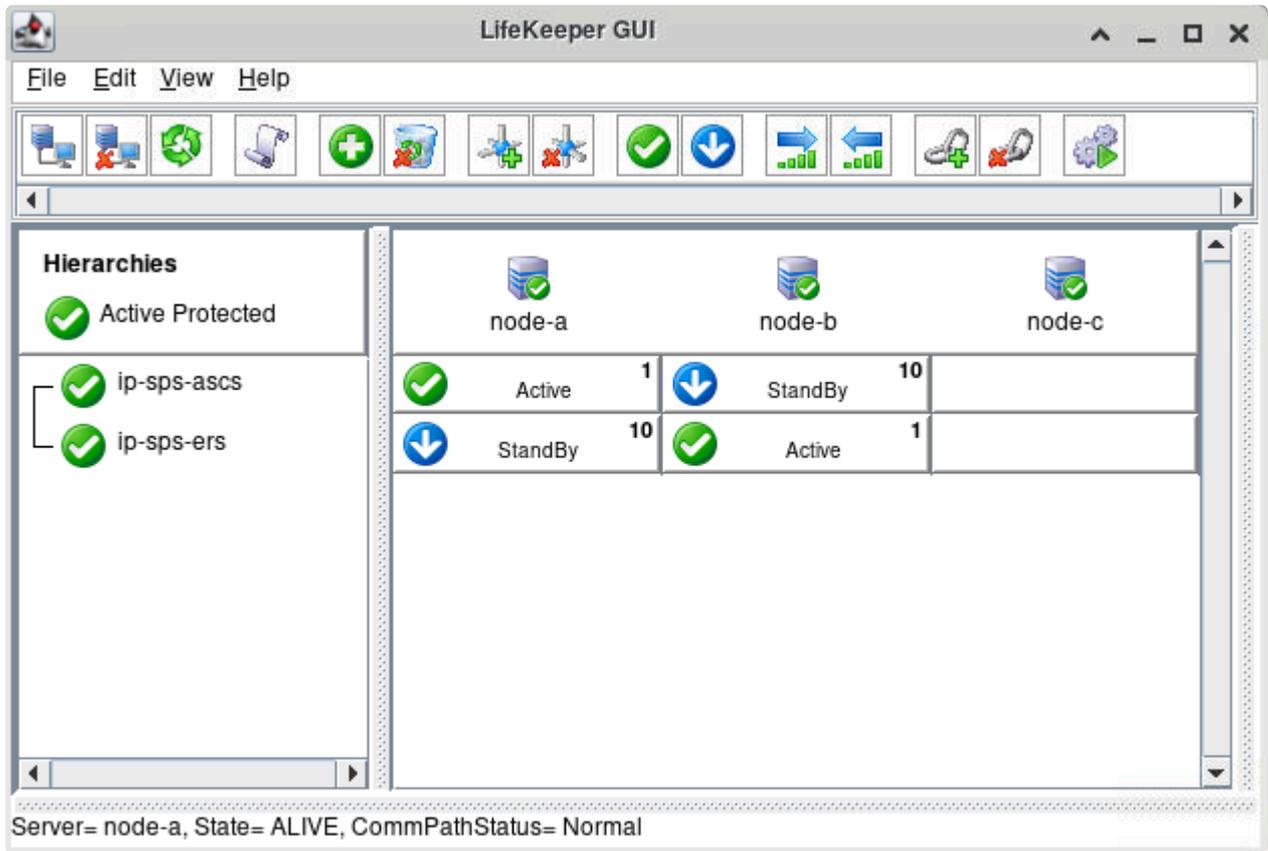
Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
IP Resource	10.20.0.10
Netmask	255.255.255.255
Network Interface	eth0 
IP Resource Tag	ip-sps-ascs
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 

Template Priority	1 
Target Priority	10 
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.20.0.10 
Netmask	255.255.255.255
Network Interface	eth0 
IP Resource Tag	ip-sps-ascs 

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
IP Resource	10.20.0.20
Netmask	255.255.255.255
Network Interface	eth0 
IP Resource Tag	ip-sps-ers
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.20.0.20 
Netmask	255.255.255.255
Network Interface	eth0 
IP Resource Tag	ip-sps-ers 

Once the IP resources have been created successfully, the LifeKeeper GUI should resemble the

following image.



- Following the steps in [Responding to Load Balancer Health Checks](#), install the LifeKeeper Generic Application Recovery Kit for Load Balancer Health Checks (“GenLB Recovery Kit”) and create two LifeKeeper GenLB resources, **ilb-sps-ascs** and **ilb-sps-ers**, with the following properties. Notice that **ilb-sps-ascs** is created on node-a and extended to node-b, while **ilb-sps-ers** is created on node-b and extended to node-a. The  icon indicates that the default option is chosen.

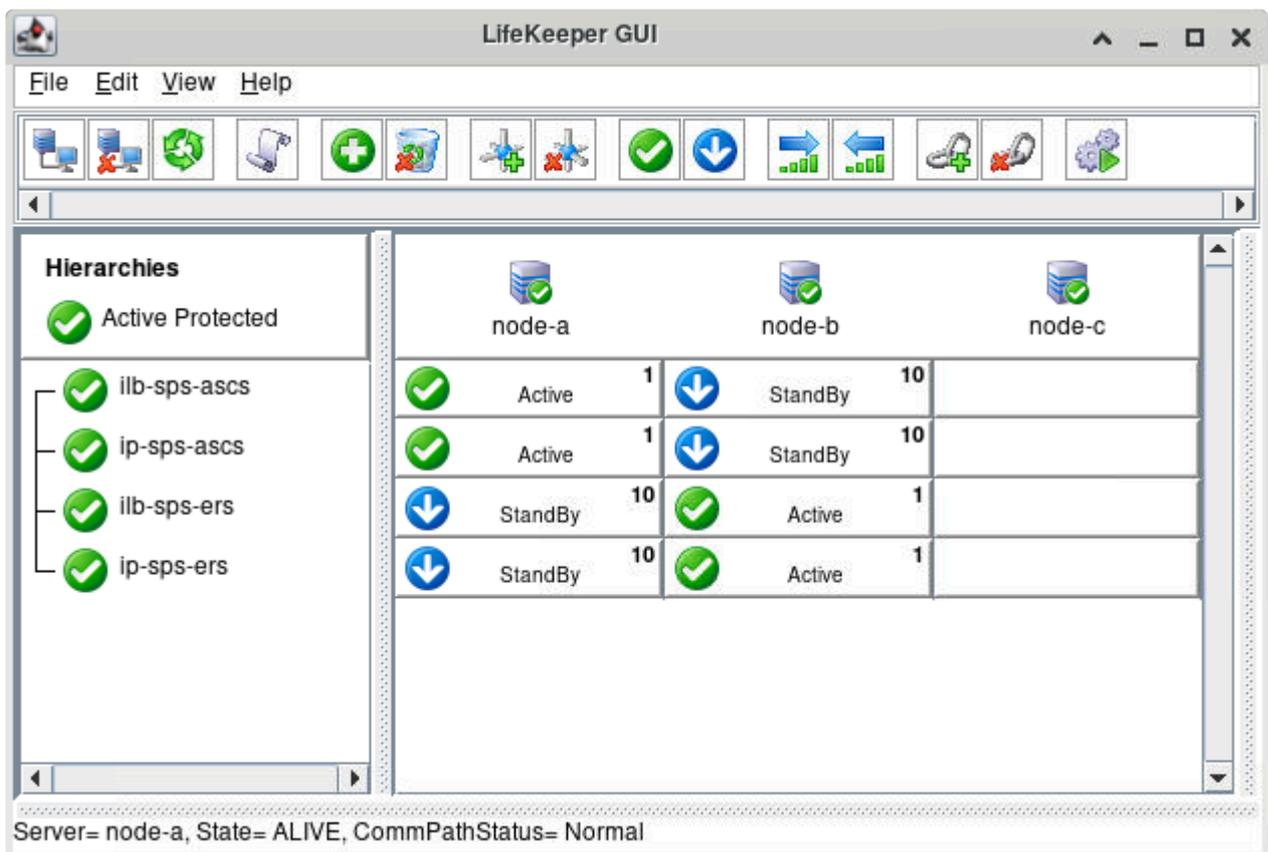
Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Restore Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/restore.pl
Remove Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/remove.pl
QuickCheck Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/quickCheck <b>(Note:</b> Although the quickCheck script may be optional for some ‘Generic Application’ resource types, it is required for GenLB resources.)
Local Recovery Script	None (Empty)

Application Info	51098
Bring Resource In Service	Yes
Resource Tag	ilb-sps-ascs
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Resource Tag	ilb-sps-ascs 
Application Info	51098 

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
Restore Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/restore.pl
Remove Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/remove.pl
QuickCheck Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/quickCheck <b>(Note:</b> Although the quickCheck script may be optional for some 'Generic Application' resource types, it is required for GenLB resources.)
Local Recovery Script	None (Empty)
Application Info	52098
Bring Resource In Service	Yes
Resource Tag	ilb-sps-ers

Pre-Extend Wizard	
Target Server	node-a
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
Extend gen/app Resource Hierarchy Wizard	
Resource Tag	ilb-sps-ers ✓
Application Info	52098 ✓

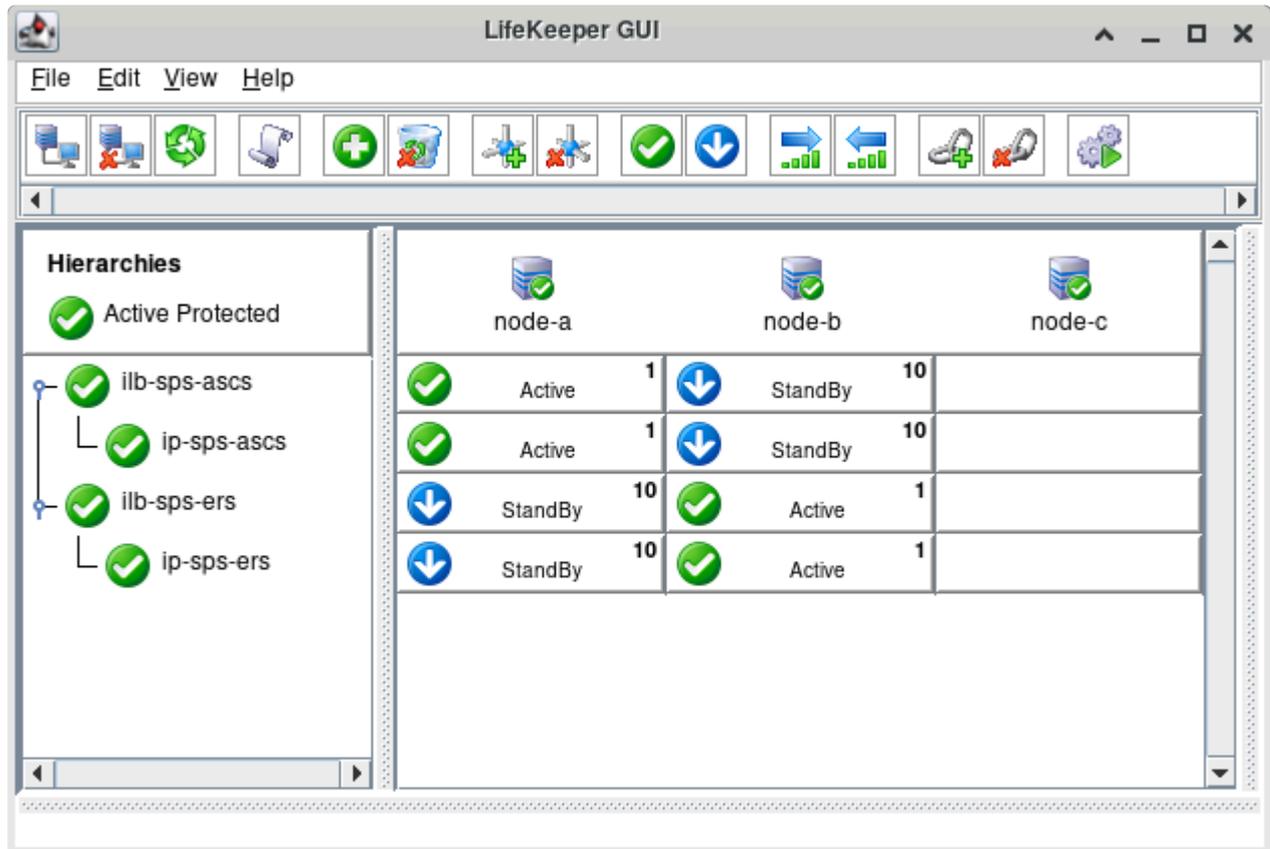
The resources will appear in the LifeKeeper GUI resource pane once they have been created and extended successfully.



- Right-click the **ilb-sps-ascs** resource on node-a and click “Create Dependency...” Specify **ip-sps-ascs** as the Child Resource Tag and click **Create Dependency**.
- Right-click the **ilb-sps-ers** resource on node-b and click “Create Dependency...” Specify **ip-sps-ers** as the Child Resource Tag and click **Create Dependency**.

Once both dependencies have been created, the LifeKeeper GUI should resemble the following

image.



8. Add the following entries to `/etc/hosts` on node-a and node-b to allow resolution of each virtual hostname to the frontend IP address of the corresponding load balancer:

```
10.20.0.10 sps-ascs
10.20.0.20 sps-ers
```

9. Test switchover and failover of the GenLB resources as described in the **Test GenLB Resource Switchover and Failover** section of [Responding to Load Balancer Health Checks](#). Correct any issues found or tune the parameters of the load balancer health checks as required to achieve successful operation.

## 11.2.7.5.5.2. Create SAP File Systems

In this section we will create highly-available NFS shared file systems which will be mounted on each node hosting an SAP instance. Typical file systems to be shared between systems in a highly-available SAP AS ABAP environment include:

- /sapmnt/<SID>
- /usr/sap/trans
- /usr/sap/<SID>/ASCS<InstNum>
- /usr/sap/<SID>/ERS<InstNum>

See [Setting up File Systems for a High-Availability System](#) for more details.

 **Note:** Since the instance-specific file systems (i.e., /usr/sap/<SID>/ASCS<InstNum> and /usr/sap/<SID>/ERS<InstNum>) only need to be mounted on the system where the corresponding SAP instance is running, it is possible to replicate them using SIOS DataKeeper rather than sharing them via NFS.

 **Note:** These shared file systems support the highly-available AS ABAP deployment. Additional database-related file systems will vary depending on the backend database platform chosen.

The recommended implementation for exporting the NFS shares varies by cloud platform. Please follow the steps provided in the section corresponding to your cloud platform:

- [AWS/Azure – Create SAP Shared and Replicated File Systems](#)
- [Google Cloud – Creating SAP Shared and Replicated File Systems](#)

## 11.2.7.5.5.2.1. AWS/Azure – Create SAP Shared and Replicated File Systems

\* **Note:** This section applies to deployments on AWS and Microsoft Azure. For deployments on Google Cloud, follow the steps given in [Google Cloud – Create SAP Shared File Systems](#).

In this section we will create highly-available NFS shared file systems which will be mounted on each node hosting an SAP instance. We will also create SIOS DataKeeper mirrors to replicate data for the ASCS10 and ERS20 instances between node-a and node-b. Typical file systems to be shared or replicated between systems in a highly-available SAP AS ABAP environment include:

- /sapmnt/<SID>
- /usr/sap/trans
- /usr/sap/<SID>/ASCS<InstNum>
- /usr/sap/<SID>/ERS<InstNum>

See [Setting up File Systems for a High-Availability System](#) for more details.

\* **Note:** For the purposes of this guide we have chosen to implement the file sharing and replication mechanisms in such a way that all resources are self-contained within the cluster itself. However, other configurations are also possible. For example, the NFS shared file systems could be hosted on a highly-available external NFS server cluster or could be hosted via a cloud-native NFS solution such as Amazon Elastic File System (EFS).

1. Create and attach four additional disks to node-a and node-b to support the highly-available SAP installation. The device names used in this example may vary depending on your environment (e.g., /dev/xvdb instead of /dev/sdb), so adjust the commands given in the section accordingly to use the appropriate device names. Note that these disk sizes are used for evaluation purposes only. Consult the relevant documentation from SAP, your cloud provider, and any third-party storage provider when provisioning resources in a production environment.

Device Name	Minimum Size	Mount Point
/dev/sdb	8GB	/export/sapmnt/SPS
/dev/sdc	8GB	/export/usr/sap/trans
/dev/sdd	8GB	/usr/sap/SPS/ASCS10
/dev/sde	4GB	/usr/sap/SPS/ERS20

\* **Note:** Other configurations may be used when attaching storage to support the SAP

installation. For example, it may be convenient to use a single physical volume partitioned with multiple logical partitions via LVM.

- Execute the following commands on both node-a and node-b to create a /dev/sdb1 partition with an xfs file system.

```
[root@node-a ~]# parted /dev/sdb --script mklabel gpt mkpart xfs part xfs 0% 10
0%

[root@node-a ~]# mkfs.xfs /dev/sdb1
meta-data=/dev/sdb1      isize=512    agcount=4, agsize=1310592 blks
           =             sectsz=4096   attr=2,   projid32bit=1
           =             crc=1        finobt=1, sparse=1, rmapbt=0
           =             reflink=1
data      =             bsize=4096   blocks=5242368, imaxpct=25
           =             sunit=0      swidth=0 blks
naming   =version 2     bsize=4096   ascii-ci=0,  ftype=1
log      =internal log bsize=4096   blocks=2560,  version=2
           =             sectsz=4096   sunit=1 blks, lazy-count=1
realtime =none         extsz=4096   blocks=0,    rtextents=0

[root@node-a ~]# partprobe /dev/sdb1
```

Repeat these commands for /dev/sdc, /dev/sdd, and /dev/sde on both node-a and node-b.

- Execute the following command on both node-a and node-b to create the mount points for the shared and replicated SAP file systems:

```
[root@node-a ~]# mkdir -p /export/{usr/sap/trans,sapmnt/SPS} /sapmnt/SPS /usr/
sap/{trans,SPS/{ASCS10,ERS20,SYS}}
```

## Create DataKeeper Data Replication Resource Hierarchies

- Following the steps described in [How to Create Data Replication of a File System](#), use the following parameters to create and extend a DataKeeper data replication resource (**datarep-/export/sapmnt/SPS**) to mirror the contents of the /export/sapmnt/SPS directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
Create Resource Wizard	

Switchback Type	intelligent 
Server	node-a
Hierarchy Type	Replicate New Filesystem 
Source Disk	/dev/sdb1 (8.0 GB)
New Mount Point	/export/sapmnt/SPS
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-/export/sapmnt/SPS
File System Resource Tag	/export/sapmnt/SPS 
Bitmap File	/opt/LifeKeeper/bitmap__export_sapmnt_SPS 
Enable Asynchronous Replication?	no 
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/export/sapmnt/SPS 
Root Tag	/export/sapmnt/SPS 
Target Disk	/dev/sdb1 (8.0 GB)
Data Replication Resource Tag	datarep-/export/sapmnt/SPS 
Bitmap File	/opt/LifeKeeper/bitmap__export_sapmnt_SPS 
Replication Path	10.20.1.10/10.20.2.10

2. Following the steps described in [How to Create Data Replication of a File System](#), use the following parameters to create and extend a DataKeeper data replication resource (**datarep-/export/usr/sap/trans**) to mirror the contents of the /export/usr/sap/trans directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent ✓
Server	node-a
Hierarchy Type	Replicate New Filesystem ✓
Source Disk	/dev/sdc1 (8.0 GB)
New Mount Point	/export/usr/sap/trans
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-/export/usr/sap/trans
File System Resource Tag	/export/usr/sap/trans ✓
Bitmap File	/opt/LifeKeeper/bitmap__export_usr_sap_trans ✓
Enable Asynchronous Replication?	no ✓
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/export/usr/sap/trans ✓
Root Tag	/export/usr/sap/trans ✓
Target Disk	/dev/sdc1 (8.0 GB)
Data Replication Resource Tag	datarep-/export/usr/sap/trans ✓
Bitmap File	/opt/LifeKeeper/bitmap__export_usr_sap_trans ✓
Replication Path	10.20.1.10/10.20.2.10

- Following the steps described in [How to Create Data Replication of a File System](#), use the following parameters to create and extend a DataKeeper data replication resource (**datarep-/usr/sap/SPS/ASCS10**) to mirror the contents of the /usr/sap/SPS/ASCS10 directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-a and extended to

node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Hierarchy Type	Replicate New Filesystem 
Source Disk	/dev/sdd1 (8.0 GB)
New Mount Point	/usr/sap/SPS/ASCS10
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-/usr/sap/SPS/ASCS10
File System Resource Tag	/usr/sap/SPS/ASCS10 
Bitmap File	/opt/LifeKeeper/bitmap__usr_sap_SPS_ASCS10 
Enable Asynchronous Replication?	no 
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/usr/sap/SPS/ASCS10 
Root Tag	/usr/sap/SPS/ASCS10 
Target Disk	/dev/sdd1 (8.0 GB)
Data Replication Resource Tag	datarep-/usr/sap/SPS/ASCS10 
Bitmap File	/opt/LifeKeeper/bitmap__usr_sap_SPS_ASCS10 
Replication Path	10.20.1.10/10.20.2.10

4. Following the steps described in [How to Create Data Replication of a File System](#), use the

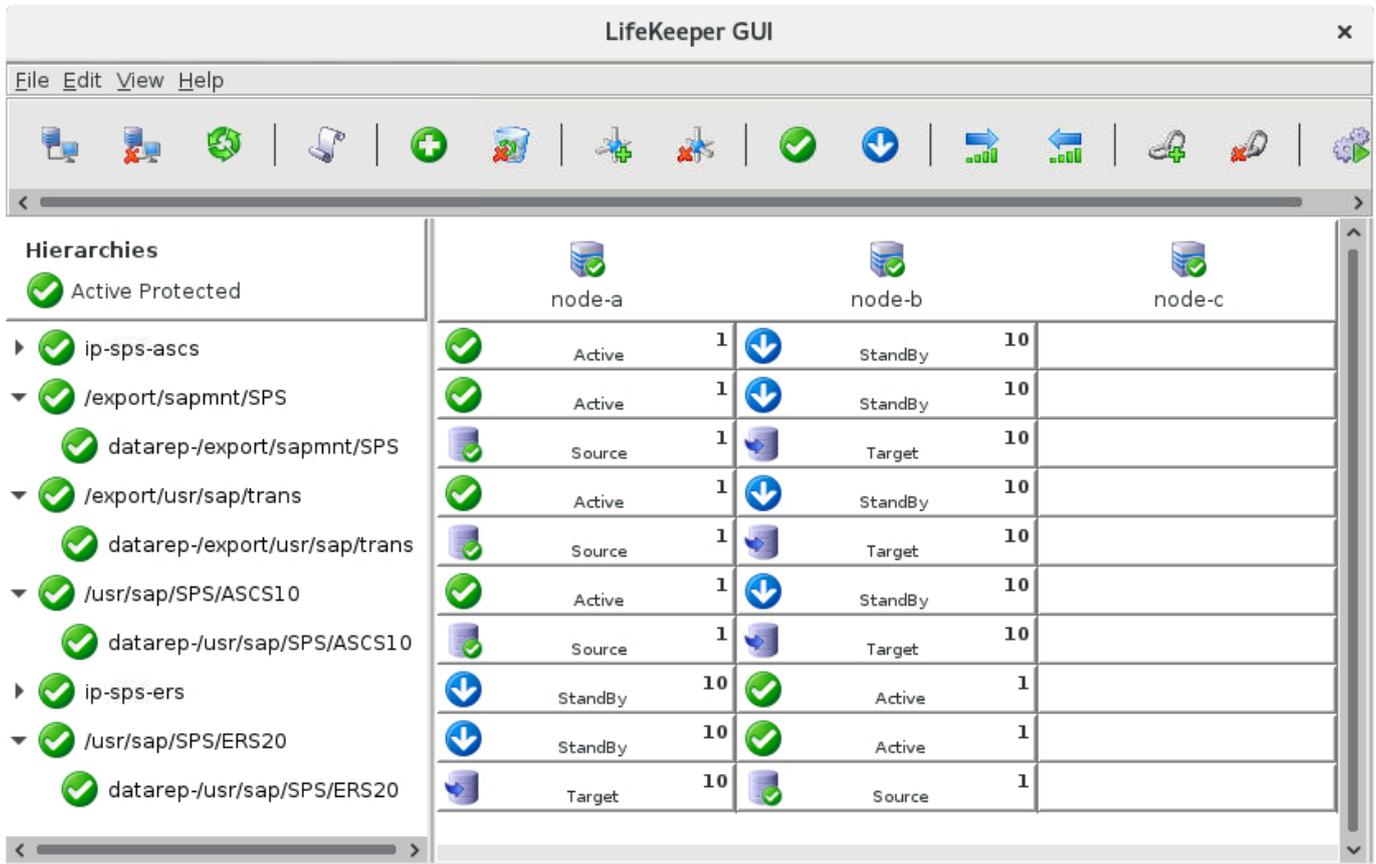
following parameters to create and extend a DataKeeper data replication resource (**datarep-/usr/sap/SPS/ERS20**) to mirror the contents of the /usr/sap/SPS/ERS20 directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-b and extended to node-a.

Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
Hierarchy Type	Replicate New Filesystem 
Source Disk	/dev/sde1 (4.0 GB)
New Mount Point	/usr/sap/SPS/ERS20
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-/usr/sap/SPS/ERS20
File System Resource Tag	/usr/sap/SPS/ERS20 
Bitmap File	/opt/LifeKeeper/bitmap__usr_sap_SPS_ERS20 
Enable Asynchronous Replication?	no 
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/usr/sap/SPS/ERS20 
Root Tag	/usr/sap/SPS/ERS20 
Target Disk	/dev/sde1 (4.0 GB)
Data Replication Resource Tag	datarep-/usr/sap/SPS/ERS20 
Bitmap File	/opt/LifeKeeper/bitmap__usr_sap_SPS_ERS20 

Replication Path	10.20.2.10/10.20.1.10
------------------	-----------------------

Once all of the DataKeeper data replication resource hierarchies have been created, the LifeKeeper GUI should look similar to the following image:



## Create NFS Resource Hierarchies

As described in [Protecting an NFS Resource](#), ensure that the NFS server and client utilities are installed on both node-a and node-b and that the `rpcbind` and `nfs-server` services have been started and enabled to run at boot on both nodes.

1. If deploying on RedHat Enterprise Linux 8.x, enable NFSv3 via UDP by adding the line 'udp=y' to the [nfsd] section of /etc/nfs.conf and restarting the nfs-server service. This step must be performed on both node-a and node-b.

```
# vi /etc/nfs.conf
# cat /etc/nfs.conf
<snip>
[nfsd]
<snip>
udp=y
<snip>
```

```
# systemctl restart nfs-server
```

2. Add the following entries to /etc/exports on node-a:

```
/export/sapmnt/SPS *(rw, sync, no_root_squash)
/export/usr/sap/trans *(rw, sync, no_root_squash)
```

 **Note:** We are allowing any client system to mount these shared file systems for simplicity, but these entries may be modified to restrict access to only the servers within the SAP environment that need it.

3. Execute the following command on node-a to export the shared file systems:

```
[root@node-a ~]# exportfs -rav
exporting */export/usr/sap/trans
exporting */export/sapmnt/SPS
```

4. Execute the following command to verify that the shared file systems are visible from both node-a and node-b, as well as from the servers that will host the PAS and AAS instances:

```
# showmount -e sps-ascs
Export list for sps-ascs:
/export/usr/sap/trans *
/export/sapmnt/SPS *
```

5. Following the steps described in [Protecting an NFS Resource](#), use the following parameters to create and extend an NFS resource (**nfs-/export/sapmnt/SPS**) to protect the shared /export/sapmnt/SPS file system. Notice that the NFS resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Export Point	/export/sapmnt/SPS
IP Tag	ip-sps-ascs
NFS Tag	nfs-/export/sapmnt/SPS 

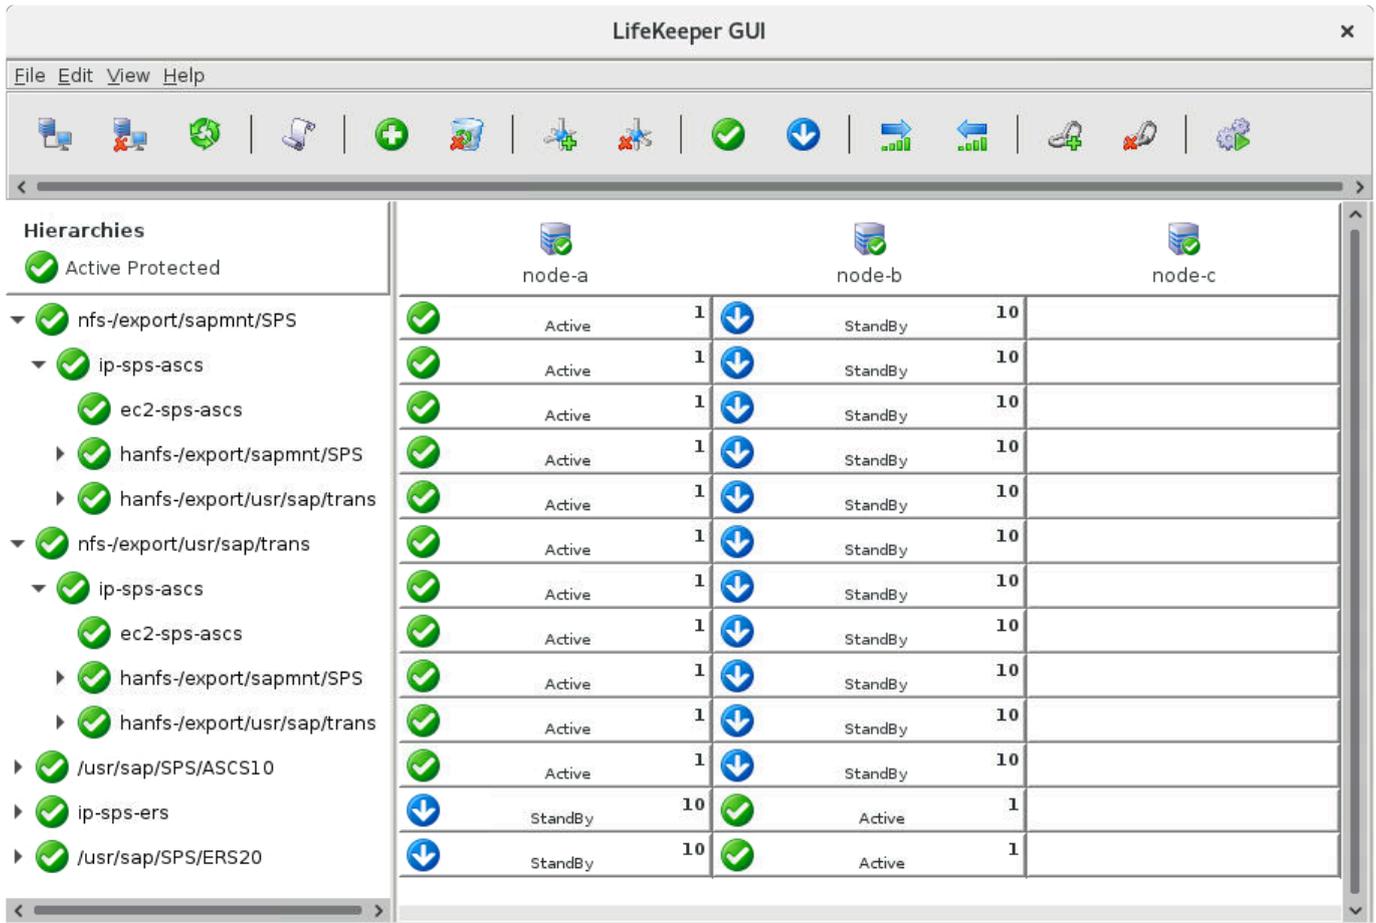
Pre-Extend Wizard	
Target Server	node-b
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
Extend gen/nfs Resource Hierarchy Wizard	
NFS Tag	nfs-/export/sapmnt/SPS ✓

6. Following the steps described in [Protecting an NFS Resource](#), use the following parameters to create and extend an NFS resource (**nfs-/export/usr/sap/trans**) to protect the shared /export/usr/sap/trans file system. Notice that the NFS resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The ✓ icon indicates that the default option is chosen.

Field	Value
Create Resource Wizard	
Switchback Type	intelligent ✓
Server	node-a
Export Point	/export/usr/sap/trans
IP Tag	ip-sps-ascs
NFS Tag	nfs-/usr/sap/trans ✓
Pre-Extend Wizard	
Target Server	node-b
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
Extend gen/nfs Resource Hierarchy Wizard	
NFS Tag	nfs-/export/usr/sap/trans ✓

Once the NFS resource hierarchies have been created successfully, the LifeKeeper GUI should resemble the following image. Note that the names of the resources used to protect the ASCS and ERS instance virtual host names may differ depending on which subsection of [Create ASCS and ERS Virtual](#)

IPs was followed.



## Add Mount Entries to /etc/fstab

- Execute the following commands on node-a, node-b, node-d, and node-e (the ASCS, ERS, PAS, and AAS instance hosts) to mount sps-ascs:/export/sapmnt/SPS at /sapmnt/SPS during boot:

```
# echo "sps-ascs:/export/sapmnt/SPS /sapmnt/SPS nfs nfsvers=3,proto=udp,rw,sync,bg 0 0" >> /etc/fstab
# mount /sapmnt/SPS
```

- Execute the following commands on node-d and node-e (the PAS and AAS instance hosts) to mount sps-ascs:/export/usr/sap/trans at /usr/sap/trans during boot:

```
# echo "sps-ascs:/export/usr/sap/trans /usr/sap/trans nfs nfsvers=3,proto=udp,rw,sync,bg 0 0" >> /etc/fstab
# mount /usr/sap/trans
```

## Modify /etc/default/LifeKeeper

Add the line 'SAP\_NFS\_CHECK\_DIRS=/sapmnt/SPS' to /etc/default/LifeKeeper on node-a and node-b to allow the SAP Recovery Kit to monitor the availability of the sps-ascs:/export/sapmnt/SPS NFS share

before performing administrative tasks that require access to files found on the shared file system.

**!** **Note:** The `SAP_NFS_CHECK_DIRS` tunable parameter should not be used when sharing the file systems via the Amazon Elastic File System (EFS) service, since the EFS service will not respond to a health ping in the same way as a local NFS server.

**\* Note:** Issues and unexpected hangs may arise during switchover and failover when using NFSv4 or TCP to share the file systems. However, if one or both of these must be used in the file-sharing configuration then changes must be made to parameters in `/etc/default/LifeKeeper`. If using NFSv4, set the tunable value `NFS_VERSION=4` in `/etc/default/LifeKeeper` on both node-a and node-b. If using TCP (either with NFSv3 or NFSv4), set the tunable value `NFS_RPC_PROTOCOL=tcp` in `/etc/default/LifeKeeper` on both node-a and node-b. For example, if the file systems are shared using NFSv4 over TCP, both tunable values `NFS_VERSION=4` and `NFS_RPC_PROTOCOL=tcp` must be set in `/etc/default/LifeKeeper` on both node-a and node-b.

The shared file systems are now prepared for installation of the SAP instances.

# 11.2.7.5.5.2.2. Google Cloud – Create SAP Shared and Replicated File Systems

**Note:** This section applies to deployments on Google Cloud. For deployments on AWS or Microsoft Azure, follow the steps given in [AWS/Azure – Create SAP Shared and Replicated File Systems](#).

In this section we will create highly-available NFS shared file systems which will be mounted on each node hosting an SAP instance. We will also create SIOS DataKeeper mirrors to replicate data for the ASCS10 and ERS20 instances between node-a and node-b. Typical file systems to be shared or replicated between systems in a highly-available SAP AS ABAP environment include:

- /sapmnt/<SID>
- /usr/sap/trans
- /usr/sap/<SID>/ASCS<InstNum>
- /usr/sap/<SID>/ERS<InstNum>

See [Setting up File Systems for a High-Availability System](#) for more details.

**Note:** For the purposes of this guide we have chosen to implement the file sharing and replication mechanisms in such a way that all resources are self-contained within the cluster itself. However, other configurations are also possible. For example, the NFS shared file systems could be hosted on a highly-available external NFS server cluster or could be hosted via a cloud-native NFS solution such as NetApp Cloud Volumes Service.

1. Create and attach four additional disks to node-a and node-b to support the highly-available SAP installation. The device names used in this example may vary depending on your environment (e.g., /dev/xvdb instead of /dev/sdb), so adjust the commands given in the section accordingly to use the appropriate device names. Note that these disk sizes are used for evaluation purposes only. Consult the relevant documentation from SAP, your cloud provider, and any third-party storage provider when provisioning resources in a production environment.

Device Name	Minimum Size	Mount Point
/dev/sdb	10GB	/export/sapmnt/SPS
/dev/sdc	10GB	/export/usr/sap/trans
/dev/sdd	10GB	/usr/sap/SPS/ASCS10
/dev/sde	10GB	/usr/sap/SPS/ERS20

**Note:** Other configurations may be used when attaching storage to support the SAP

installation. For example, it may be convenient to use a single physical volume partitioned with multiple logical partitions via LVM.

- Execute the following commands on both node-a and node-b to create a /dev/sdb1 partition with an xfs file system.

```
[root@node-a ~]# parted /dev/sdb --script mklabel gpt mkpart xfs part xfs 0% 10
0%

[root@node-a ~]# mkfs.xfs /dev/sdb1
meta-data=/dev/sdb1      isize=512    agcount=4, agsize=1310592 blks
           =             sectsz=4096   attr=2,   projid32bit=1
           =             crc=1        finobt=1, sparse=1, rmapbt=0
           =             reflink=1
data      =             bsize=4096   blocks=5242368, imaxpct=25
           =             sunit=0      swidth=0 blks
naming    =version 2     bsize=4096   ascii-ci=0,  ftype=1
log       =internal log bsize=4096   blocks=2560,  version=2
           =             sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none         extsz=4096   blocks=0,    rtextents=0

[root@node-a ~]# partprobe /dev/sdb1
```

Repeat these commands for /dev/sdc, /dev/sdd, and /dev/sde on both node-a and node-b.

- Execute the following command on both node-a and node-b to create the mount points for the shared and replicated SAP file systems:

```
[root@node-a ~]# mkdir -p /export/{usr/sap/trans,sapmnt/SPS} /sapmnt/SPS /usr/
sap/{trans,SPS/{ASCS10,ERS20,SYS}}
```

## Create DataKeeper Data Replication Resource Hierarchies

- Following the steps described in [How to Create Data Replication of a File System](#), use the following parameters to create and extend a DataKeeper data replication resource (**datarep-/export/sapmnt/SPS**) to mirror the contents of the /export/sapmnt/SPS directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	

Switchback Type	intelligent 
Server	node-a
Hierarchy Type	Replicate New Filesystem 
Source Disk	/dev/sdb1 (10.0 GB)
New Mount Point	/export/sapmnt/SPS
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-/export/sapmnt/SPS
File System Resource Tag	/export/sapmnt/SPS 
Bitmap File	/opt/LifeKeeper/bitmap__export_sapmnt_SPS 
Enable Asynchronous Replication?	no 
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/export/sapmnt/SPS 
Root Tag	/export/sapmnt/SPS 
Target Disk	/dev/sdb1 (10.0 GB)
Data Replication Resource Tag	datarep-/export/sapmnt/SPS 
Bitmap File	/opt/LifeKeeper/bitmap__export_sapmnt_SPS 
Replication Path	10.20.1.10/10.20.2.10

2. Following the steps described in [How to Create Data Replication of a File System](#), use the following parameters to create and extend a DataKeeper data replication resource (**datarep-/export/usr/sap/trans**) to mirror the contents of the /export/usr/sap/trans directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Hierarchy Type	Replicate New Filesystem 
Source Disk	/dev/sdc1 (10.0 GB)
New Mount Point	/export/usr/sap/trans
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-/export/usr/sap/trans
File System Resource Tag	/export/usr/sap/trans 
Bitmap File	/opt/LifeKeeper/bitmap__export_usr_sap_trans 
Enable Asynchronous Replication?	no 
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/export/usr/sap/trans 
Root Tag	/export/usr/sap/trans 
Target Disk	/dev/sdc1 (10.0 GB)
Data Replication Resource Tag	datarep-/export/usr/sap/trans 
Bitmap File	/opt/LifeKeeper/bitmap__export_usr_sap_trans 
Replication Path	10.20.1.10/10.20.2.10

- Following the steps described in [How to Create Data Replication of a File System](#), use the following parameters to create and extend a DataKeeper data replication resource (**datarep-/usr/sap/SPS/ASCS10**) to mirror the contents of the /usr/sap/SPS/ASCS10 directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-a and extended to

node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Hierarchy Type	Replicate New Filesystem 
Source Disk	/dev/sdd1 (10.0 GB)
New Mount Point	/usr/sap/SPS/ASCS10
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-/usr/sap/SPS/ASCS10
File System Resource Tag	/usr/sap/SPS/ASCS10 
Bitmap File	/opt/LifeKeeper/bitmap__usr_sap_SPS_ASCS10 
Enable Asynchronous Replication?	no 
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/usr/sap/SPS/ASCS10 
Root Tag	/usr/sap/SPS/ASCS10 
Target Disk	/dev/sdd1 (10.0 GB)
Data Replication Resource Tag	datarep-/usr/sap/SPS/ASCS10 
Bitmap File	/opt/LifeKeeper/bitmap__usr_sap_SPS_ASCS10 
Replication Path	10.20.1.10/10.20.2.10

4. Following the steps described in [How to Create Data Replication of a File System](#), use the

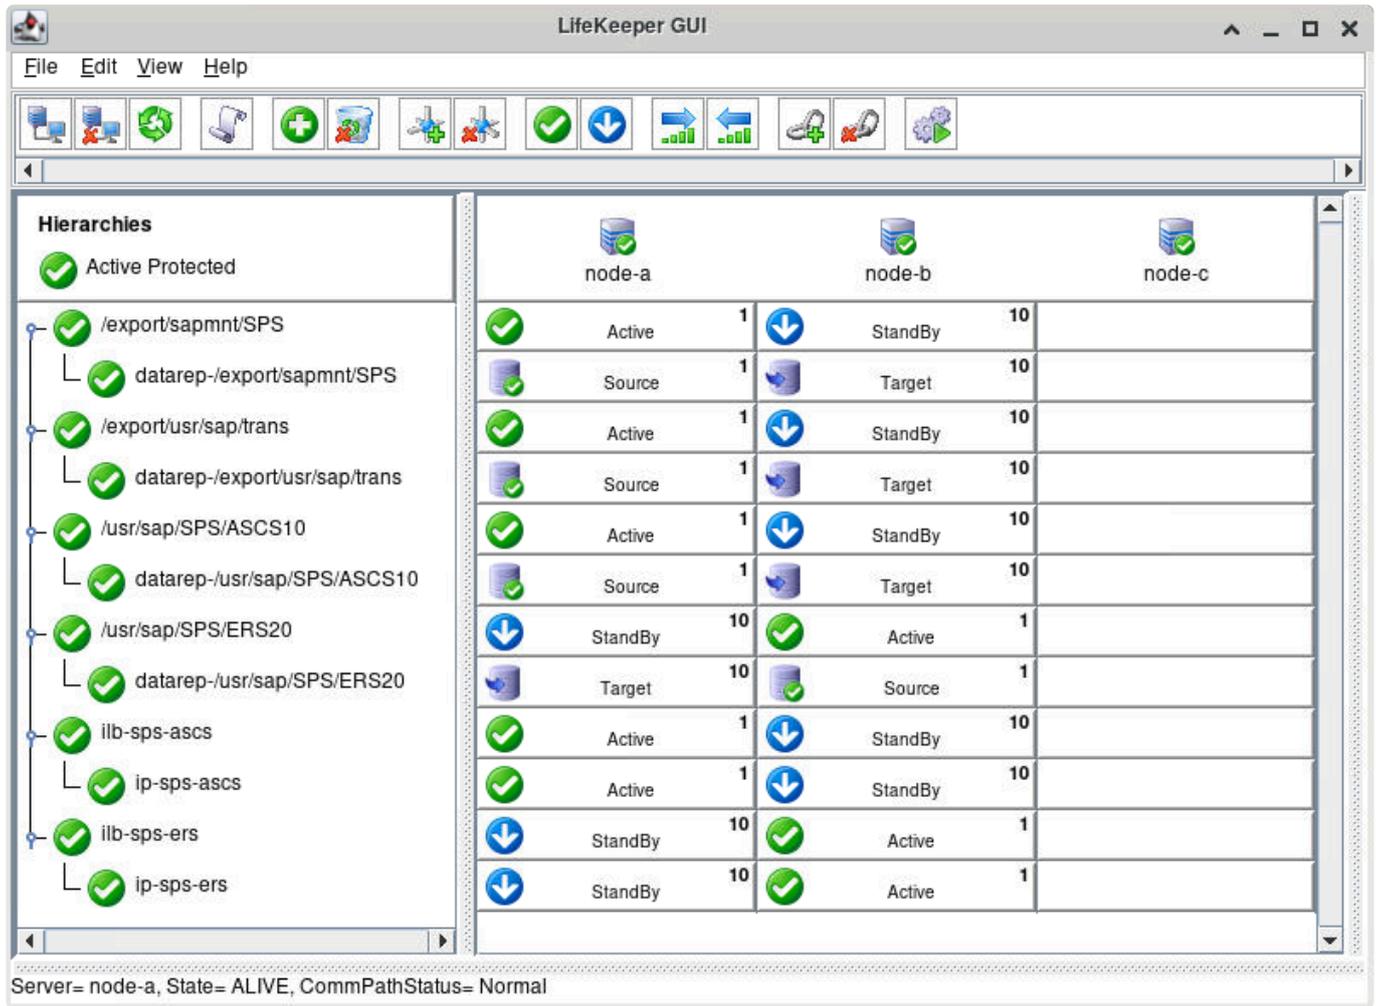
following parameters to create and extend a DataKeeper data replication resource (**datarep-usr/sap/SPS/ERS20**) to mirror the contents of the /usr/sap/SPS/ERS20 directory between node-a and node-b. Notice that the DataKeeper resource is being created on node-b and extended to node-a.

Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
Hierarchy Type	Replicate New Filesystem 
Source Disk	/dev/sde1 (10.0 GB)
New Mount Point	/usr/sap/SPS/ERS20
New Filesystem Type	xfs
Data Replication Resource Tag	datarep-usr/sap/SPS/ERS20
File System Resource Tag	/usr/sap/SPS/ERS20 
Bitmap File	/opt/LifeKeeper/bitmap_usr_sap_SPS_ERS20 
Enable Asynchronous Replication?	no 
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend scsi/netraid Resource Hierarchy Wizard</b>	
Mount Point	/usr/sap/SPS/ERS20 
Root Tag	/usr/sap/SPS/ERS20 
Target Disk	/dev/sde1 (10.0 GB)
Data Replication Resource Tag	datarep-usr/sap/SPS/ERS20 
Bitmap File	/opt/LifeKeeper/bitmap_usr_sap_SPS_ERS20 

Replication Path	10.20.2.10/10.20.1.10
------------------	-----------------------

Once all of the DataKeeper data replication resource hierarchies have been created, the LifeKeeper GUI should look similar to the following image:



## Create NFS Resource Hierarchies

✿ As described in [Protecting an NFS Resource](#), ensure that the NFS server and client utilities are installed on both node-a and node-b and that the `rpcbind` and `nfs-server` services have been started and enabled to run at boot on both nodes.

1. Add the following entries to `/etc/exports` on node-a:

```
/export/sapmnt/SPS *(rw,sync,no_root_squash)
/export/usr/sap/trans *(rw,sync,no_root_squash)
```

✿ **Note:** We are allowing any client system to mount these shared file systems for simplicity, but these entries may be modified to restrict access to only the servers within

the SAP environment that need it.

2. Execute the following command on node-a to export the shared file systems:

```
[root@node-a ~]# exportfs -rav
exporting */export/usr/sap/trans
exporting */export/sapmnt/SPS
```

3. Execute the following command to verify that the shared file systems are visible from both node-a and node-b, as well as from the servers that will host the PAS and AAS instances:

```
# showmount -e sps-ascscs
Export list for sps-ascscs:
/export/usr/sap/trans *
/export/sapmnt/SPS *
```

4. Following the steps described in [Protecting an NFS Resource](#), use the following parameters to create and extend an NFS resource (**nfs-/export/sapmnt/SPS**) to protect the shared /export/sapmnt/SPS file system. Notice that the NFS resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Export Point	/export/sapmnt/SPS
IP Tag	ip-sps-ascscs
NFS Tag	nfs-/export/sapmnt/SPS 
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/nfs Resource Hierarchy Wizard</b>	

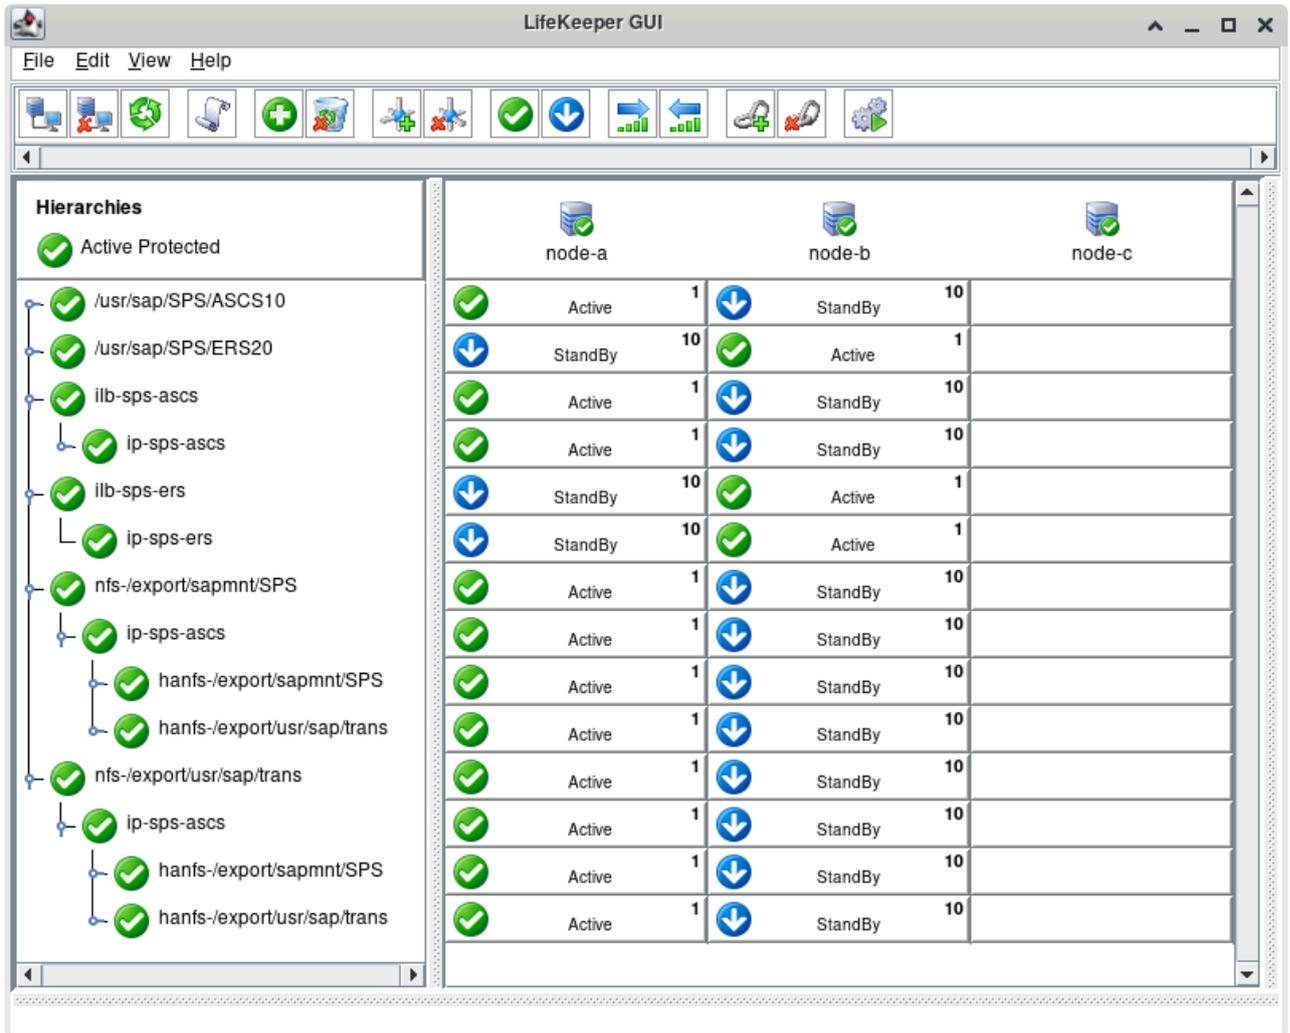
NFS Tag	nfs-/export/sapmnt/SPS 
---------	----------------------------------------------------------------------------------------------------------

- Following the steps described in [Protecting an NFS Resource](#), use the following parameters to create and extend an NFS resource (**nfs-/export/usr/sap/trans**) to protect the shared /export/usr/sap/trans file system. Notice that the NFS resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node.

The  icon indicates that the default option is chosen.

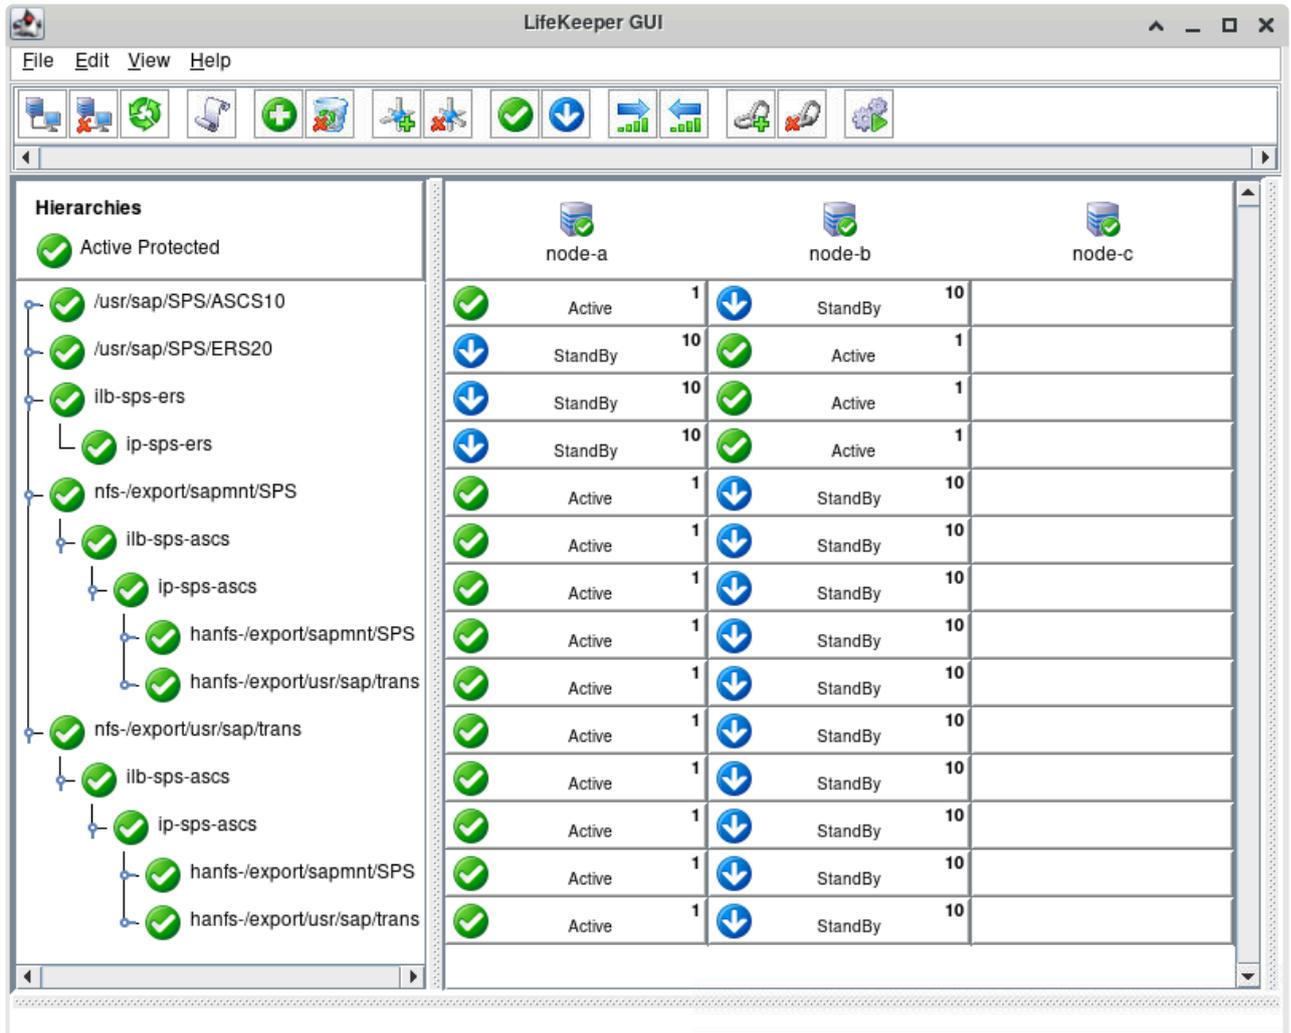
Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Export Point	/export/usr/sap/trans
IP Tag	ip-sps-ascs
NFS Tag	nfs-/usr/sap/trans 
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/nfs Resource Hierarchy Wizard</b>	
NFS Tag	nfs-/export/usr/sap/trans 

Once the NFS resource hierarchies have been created successfully, the LifeKeeper GUI should resemble the following image.



6. Right-click on the **nfs-/export/sapmnt/SPS** resource and select **Delete Dependency...** from the drop-down menu. For **Child Resource Tag**, select **ip-sps-ascsc**. Click **Next>** to continue, then click **Delete Dependency** to delete the dependency.
7. Right-click on the **nfs-/export/sapmnt/SPS** resource and select **Create Dependency...** from the drop-down menu. For **Child Resource Tag**, select **ilb-sps-ascsc**. Click **Next>** to continue, then click **Create Dependency** to create the dependency.
8. Right-click on the **nfs-/export/usr/sap/trans** resource and select **Delete Dependency...** from the drop-down menu. For **Child Resource Tag**, select **ip-sps-ascsc**. Click **Next>** to continue, then click **Delete Dependency** to delete the dependency.
9. Right-click on the **nfs-/export/usr/sap/trans** resource and select **Create Dependency...** from the drop-down menu. For **Child Resource Tag**, select **ilb-sps-ascsc**. Click **Next>** to continue, then click **Create Dependency** to create the dependency.

Once these steps are complete, the LifeKeeper GUI should resemble the following image.



## Add Mount Entries to /etc/fstab

1. Execute the following commands on node-a, node-b, node-d, and node-e (the ASCS, ERS, PAS, and AAS instance hosts) to mount sps-ascsc:/export/sapmnt/SPS at /sapmnt/SPS during boot:

```
# echo "sps-ascsc:/export/sapmnt/SPS /sapmnt/SPS nfs nfsvers=3,proto=tcp,rw,sync,bg 0 0" >> /etc/fstab
# mount /sapmnt/SPS
```

2. Execute the following commands on node-d and node-e (the PAS and AAS instance hosts) to mount sps-ascsc:/export/usr/sap/trans at /usr/sap/trans during boot:

```
# echo "sps-ascsc:/export/usr/sap/trans /usr/sap/trans nfs nfsvers=3,proto=tcp,rw,sync,bg 0 0" >> /etc/fstab
# mount /usr/sap/trans
```

## Modify /etc/default/LifeKeeper

Add the line 'SAP\_NFS\_CHECK\_DIRS=/sapmnt/SPS' to /etc/default/LifeKeeper on node-a and node-b to allow the SAP Recovery Kit to monitor the availability of the sps-ascsc:/export/sapmnt/SPS NFS share

before performing administrative tasks that require access to files found on the shared file system.

Also add the line 'NFS\_RPC\_PROTOCOL=tcp' to /etc/default/LifeKeeper on node-a and node-b to ensure that the pingnfs utility uses the TCP protocol when checking the availability of the NFS shares.

```
# vi /etc/default/LifeKeeper
# grep 'SAP_NFS_CHECK_DIRS=\\|NFS_RPC_PROTOCOL=' /etc/default/LifeKeeper
SAP_NFS_CHECK_DIRS=/sapmnt/SPS
NFS_RPC_PROTOCOL=tcp
```

The shared file systems are now prepared for installation of the SAP instances.

## 11.2.7.5.5.3. Install SAP Instances

Complete the following steps to install the ASCS10 and ERS20 instances on node-a and node-b.

### Prepare for Installation

1. On node-a and node-b, increase the available swap space to greater than 1MiB to satisfy the SAP system requirements. In this example we will create a 1GiB swap file to increase the available swap space.

 **Note:** Production environments will likely require more available swap space. Consult SAP Note 1597355 (Swap space recommendation for Linux) for more details.

```
# mkdir /swap
# dd if=/dev/zero of=/swap/swapfile.img bs=1024 count=1M
# chmod 600 /swap/swapfile.img
# mkswap /swap/swapfile.img
# echo "/swap/swapfile.img    swap    swap    sw    0 0" >> /etc/fstab
# swapon /swap/swapfile.img
```

2. Ensure that the `tcsh` package is installed on both node-a and node-b. Modify this command (e.g., to use `zypper install`) if installing on a SLES server.

```
# yum install -y tcsh
```

3. On node-a and node-b, create a `/sap-install` directory which will contain the SAP installation files.

```
# mkdir /sap-install
```

 **Note:** If the root file system for your chosen instance type does not have sufficient disk space to store the required SAP installation files once they are extracted, you may attach a temporary disk to each instance and mount it at `/sap-install` in order to complete installation of the SAP software.

4. Download SAPCAR, SAP Software Provisioning Manager 2.0 (SWPM20SPxx\_x-xxxx.SAR), and all necessary packages (e.g., SAPEXE\_xx-xxxx.SAR, SAPHOSTAGENTxx\_xx-xxxx.SAR, etc.) to the `/sap-install` directory on both nodes. Make sure that SAPCAR is executable.

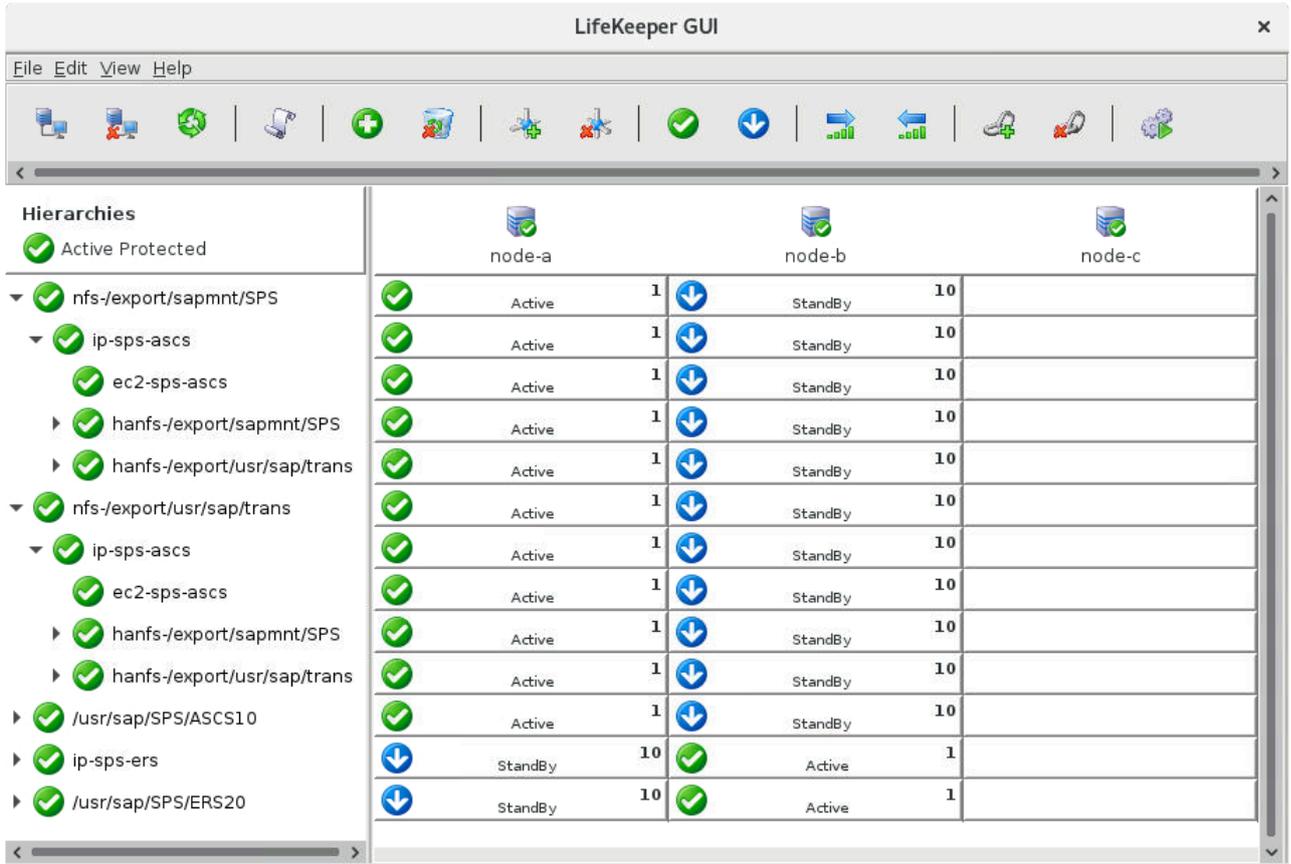
```
# chmod +x /sap-install/SAPCAR
```

5. Use SAPCAR to extract the files for SWPM 2.0.

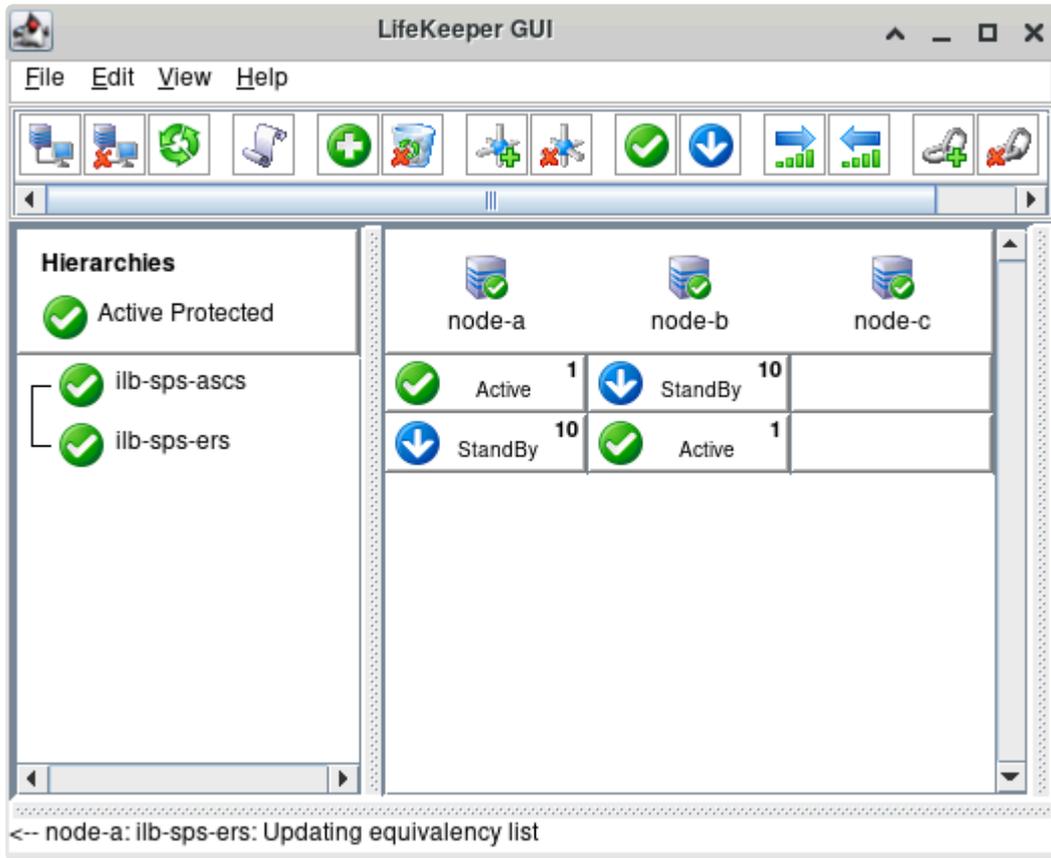
```
# cd /sap-install
# ./SAPCAR -xvf SWPM20SPxx_x-xxxx.SAR
```

Before installing the SAP instances, verify that all shared and replicated file systems described in the appropriate cloud-specific subsection of [Create SAP File Systems](#) are mounted correctly on node-a and/or node-b and that the LifeKeeper resources providing virtual hostname failover for the ASCS and ERS instances are in-service on node-a and node-b, respectively.

If deploying on AWS or Microsoft Azure, the LifeKeeper GUI should resemble the following image:



If deploying on Google Cloud, the LifeKeeper GUI should resemble the following image:



## Install the ASCS Instance on Node-A

Complete the following steps to install the ASCS10 instance on node-a.

1. If installing from a remote system, add a firewall rule to allow inbound traffic on TCP port 4237 (the port which exposes the sapinst graphical interface) on node-a.
2. Execute the following command on node-a to initiate sapinst with the virtual hostname associated to the ASCS instance:

```
[root@node-a ~]# /sap-install/sapinst SAPINST_USE_HOSTNAME=sps-ascs
```

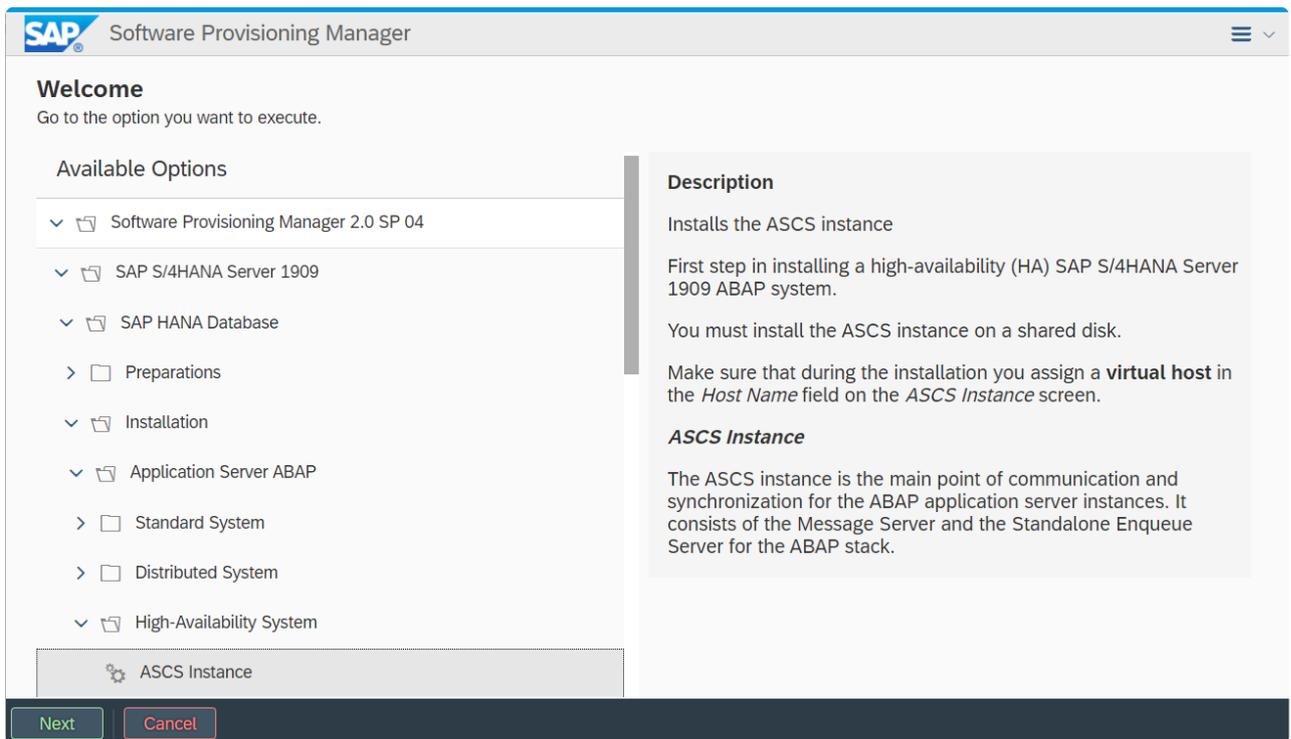
**Note:** The virtual host name used as the value of the `SAPINST_USE_HOSTNAME` parameter must resolve to the ASCS instance virtual IP that was created and protected in the section of [Create ASCS and ERS Virtual IPs](#) that you followed.

**Note:** The `SAPINST_REMOTE_ACCESS_USER=<user>` parameter may be used to allow login to the sapinst graphical interface by a non-root user.

**Note:** If sapinst displays an error due to a missing `dt_socket` transport library, install the `libnsl` package with `yum -y install libnsl`.

3. To access the sapinst graphical interface, use a web browser to navigate to `https://<node-a Public IP>:4237/sapinst/docs/index.html`.
4. Select **Software Provisioning Manager 2.0 SP XX** → **SAP S/4HANA Server 1909** → **SAP HANA Database** → **Installation** → **Application Server ABAP** → **High-Availability System** → **ASCS Instance** and click **Next**.

 **Note:** The screenshots shown in this guide are specific to Software Provisioning Manager 2.0 SP 04, and the interface may be different when using other versions of Software Provisioning Manager.



5. Provide the following parameters when installing the ASCS instance. The  icon indicates that the default option is chosen.

Field	Value
SAP System ID (SAPSID)	SPS
SAP Mount Directory	/sapmnt 
Password for All Users	<SAP User Password>
Password of SAP System Administrator	<SAP Admin Password>
User ID	Leave empty 
Group ID of sapsys	Leave empty 
Specify path to SAPEXE.SAR	/sap-install

Specify path to SAPHOSTAGENT.SAR	/sap-install
ASCS Instance Number	10
ASCS Instance Host Name	sps-ascs ✓
ABAP Message Server Port	3610 ✓
Internal ABAP Message Server Port	3910 ✓
Yes, clean up operating system users	Click

6. Review the parameters and proceed with the installation.

 **Note:** If the installation fails due to insufficient permissions on the SAP file systems, modify the permissions as described in the **Modify SAP File System Permissions** section below and retry the installation.

7. Execute the following command on node-a to clean up SAP installation files:

```
[root@node-a ~]# rm -rf /root/.sapinst
```

## Prepare Node-B For ERS Instance Installation

Complete the following steps to prepare node-b for installation of the ERS instance. In particular, this step will create the necessary users and groups, install SAP Host Agent, and install required binaries into the /usr/sap/SPS/SYS directory on node-b.

1. If installing from a remote system, add a firewall rule to allow inbound traffic on TCP port 4237 (the port which exposes the sapinst graphical interface) on node-b.
2. Execute the following command on node-b to initiate sapinst:

```
[root@node-b ~]# /sap-install/sapinst
```

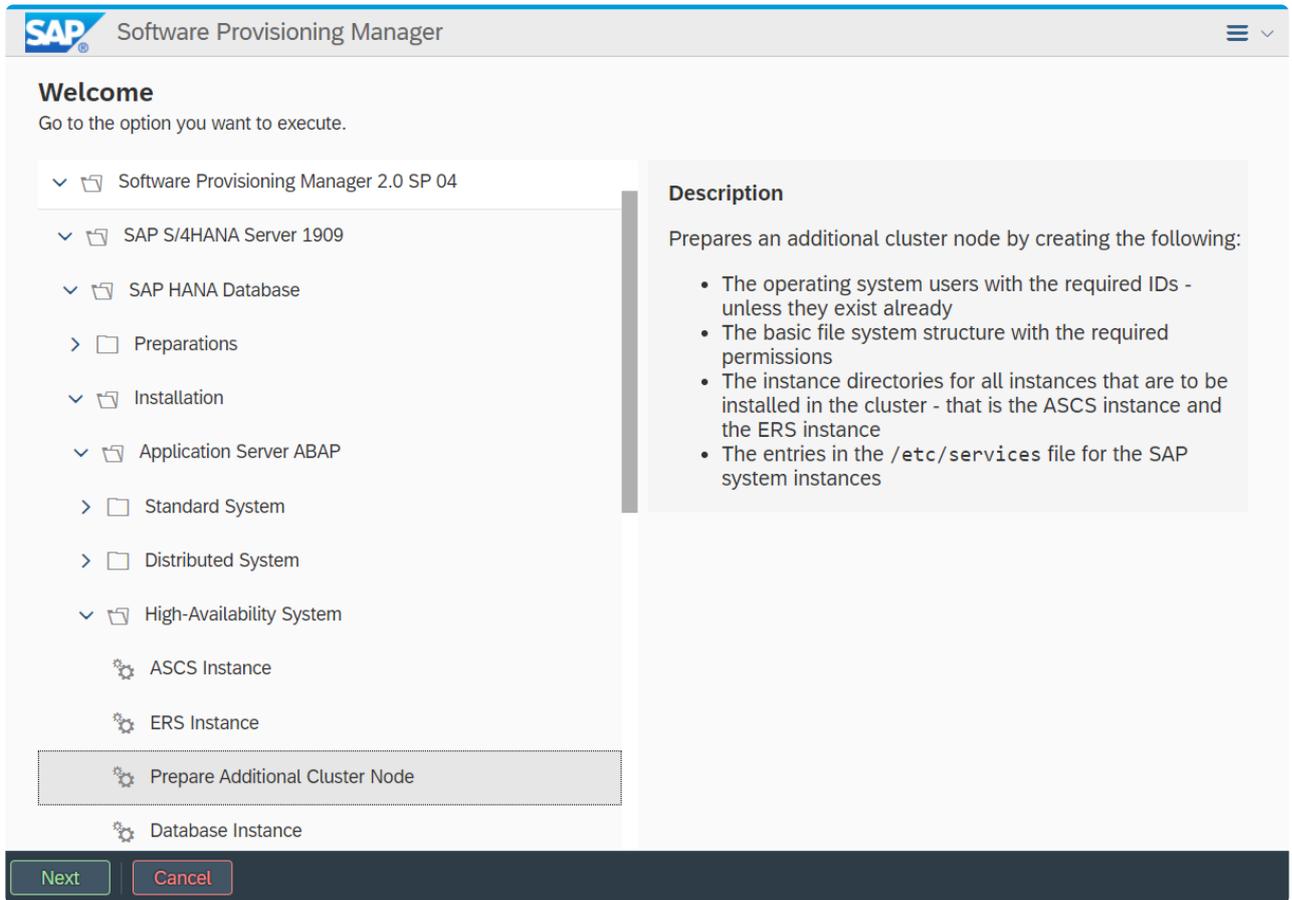
 **Note:** The `SAPINST_REMOTE_ACCESS_USER=<user>` parameter may be used to allow login to the sapinst graphical interface by a non-root user.

 **Note:** If sapinst displays an error due to a missing dt\_socket transport library, install the libnsl package with `yum install -y libnsl`.

3. To access the sapinst graphical interface, use a web browser to navigate to `https://<node-b Public`

IP>:4237/sapinst/docs/index.html.

4. **Select Software Provisioning Manager 2.0 SP XX → SAP S/4HANA Server 1909 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Prepare Additional Cluster Node** and click **Next**.



5. Provide the following parameters when preparing node-b for installation of the ERS instance.

Field	Value
Profile Directory	/sapmnt/SPS/profile
Password of SAP System Administrator	<SAP Admin Password>
Specify path to SAPHOSTAGENT.SAR	/sap-install
Yes, clean up operating system users	Click

6. Review the parameters and proceed with the installation.

**Note:** If the installation fails due to insufficient permissions on the SAP file systems, modify the permissions as described in the **Modify SAP File System Permissions** section below and retry the installation.

7. Execute the following command on node-b to clean up SAP installation files:

```
[root@node-b ~]# rm -rf /root/.sapinst
```

## Modify SAP File System Permissions

Execute the following commands on both node-a and node-b to set the correct permissions on the SAP file systems:

```
# chown -R spsadm /export/sapmnt/SPS /sapmnt/SPS /usr/sap/SPS
# chgrp -R sapsys /export/{usr/sap,sapmnt} /sapmnt /usr/sap
```

## Install the ERS Instance on Node-B

Complete the following steps to install the ERS20 instance on node-b.

1. If installing from a remote system, add a firewall rule to allow inbound traffic on TCP port 4237 (the port which exposes the sapinst graphical interface) on node-b.
2. Execute the following command on node-b to initiate sapinst with the virtual hostname associated to the ERS instance:

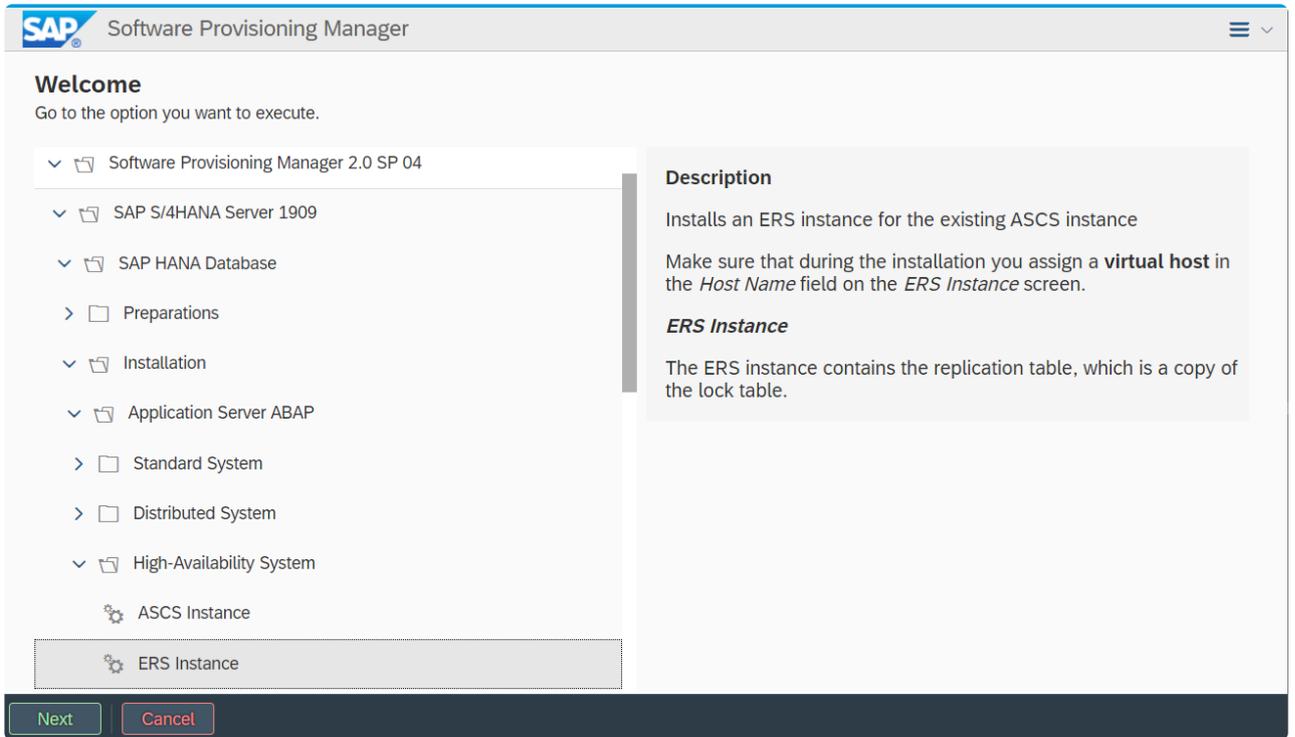
```
[root@node-b ~]# /sap-install/sapinst SAPINST_USE_HOSTNAME=sps-ers
```

 **Note:** The virtual host name used as the value of the `SAPINST_USE_HOSTNAME` parameter must resolve to the ERS instance virtual IP that was created and protected in the section of [Create ASCS and ERS Virtual IPs](#) that you followed.

 **Note:** The `SAPINST_REMOTE_ACCESS_USER=<user>` parameter may be used to allow login to the sapinst graphical interface by a non-root user.

 **Note:** If sapinst displays an error due to a missing `dt_socket` transport library, install the `libnsl` package with `yum install -y libnsl`.

3. To access the sapinst graphical interface, use a web browser to navigate to `https://<node-b Public IP>:4237/sapinst/docs/index.html`.
4. **Select Software Provisioning Manager 2.0 SP XX → SAP S/4HANA Server 1909 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → ERS Instance** and click **Next**.



5. Provide the following parameters when installing the ERS instance. The  icon indicates that the default option is chosen.

Field	Value
Profile Directory	/sapmnt/SPS/profile 
Password of SAP System Administrator (spsadm)	<spsadm User Password>
Specify path to SAPHOSTAGENT.SAR	/sap-install
Password of SAP System Administrator (sapadm)	<spsadm User Password>
User ID	Leave empty 
Group ID of sapsys	Use default value 
Name of the ASCS Instance to be Replicated	ASCS10 
Number of the ASCS Instance to be Replicated	10 
Number of the ERS Instance	20 
ERS Instance Host	sps-ers 
Yes, clean up operating system users	Click

6. Review the parameters and proceed with the installation.

 **Note:** If the installation fails due to insufficient permissions on the SAP file systems, modify the permissions as described in the **Modify SAP File System Permissions** section above and retry the installation.

7. Execute the following command on node-b to clean up SAP installation files:

```
[root@node-b ~]# rm -rf /root/.sapinst
```

## Modify /usr/sap/sapservices on Both Nodes

Since we only installed the ASCS10 instance on node-a and only installed the ERS20 instance on node-b, the /usr/sap/sapservices file is incomplete on each node. Modify this file on both node-a and node-b so that it contains entries for both the ASCS10 and ERS20 instances.

```
# vi /usr/sap/sapservices
# cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/SPS/ASCS10/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/sap/SPS/ASCS10/exe/sapstartsrv pf=/usr/sap/SPS/SYS/profile/SPS_ASCS10_sps-ascs -D -u spsadm
LD_LIBRARY_PATH=/usr/sap/SPS/ERS20/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/sap/SPS/ERS20/exe/sapstartsrv pf=/usr/sap/SPS/SYS/profile/SPS_ERS20_sps-ers -D -u spsadm
```

## Restart the ASCS Instance to Enable Enqueue Replication

Now that the ERS instance has been installed on node-b, the ASCS instance and its corresponding SAP Start Service process must be restarted on node-a in order to enable enqueue replication.

1. Verify that the following lines appear in the default profile for the LifeKeeper SAP installation (/usr/sap/SPS/SYS/profile/DEFAULT.PFL). If they do not exist, add them.

```
enq/replicatorhost = sps-ers
enq/replicatorinst = 20
```

2. Execute the following commands on node-a to restart the ASCS instance and verify that enqueue replication has been enabled.

```
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 10 -function StopWait 60 2"
15.03.2021 00:29:23
Stop
OK
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 10 -function RestartService SP
```

```

s"
15.03.2021 00:29:44
RestartService
OK
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 10 -function StartWait 60 2"
15.03.2021 00:30:07
Start
OK
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 10 -function EnqGetStatistic
| grep replication_state"
replication_state: GREEN

```

## Verify Successful ASCS and ERS Installation

Execute the following commands to verify that the ASCS10 and ERS20 instances are running successfully on node-a and node-b, respectively.

1. [root@node-a ~]# su - spsadm -c "sapcontrol -nr 10 -function GetProcessList"

```

01.01.2021 00:00:00
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
msg_server, MessageServer, GREEN, Running, 2021 01 01 00:00:00, 0:00:00, 8420
enq_server, Enqueue Server 2, GREEN, Running, 2021 01 01 00:00:00, 0:00:00, 8428

```

2. [root@node-b ~]# su - spsadm -c "sapcontrol -nr 20 -function GetProcessList"

```

01.01.2021 00:00:00
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
enq_replicator, Enqueue Replicator 2, GREEN, Running, 2021 01 01 00:00:00, 0:00:00, 25263

```

## Install PAS and AAS Instances

Once the database (e.g., SAP HANA) has been installed and configured for high-availability, use SWPM 2.0 to install one Primary Application Server (PAS) and one Additional Application Server (AAS) instance on two external application server nodes spread across two availability zones. As application server redundancy is provided by the existence of the AAS instance, the application server instances do not require LifeKeeper protection. Note that the previously created NFS shares for the /sapmnt/SPS and

`/usr/sap/trans` file systems must be mounted on both the PAS and AAS instance hosts before installation. Please consult SAP documentation for details on the installation of the PAS and AAS instances.

## 11.2.7.5.5.4. Create LifeKeeper SAP Resources

---

The steps required to create the LifeKeeper resources to protect the ASCS10 and ERS20 instances vary by cloud platform. Please follow the steps provided in the section corresponding to your cloud platform:

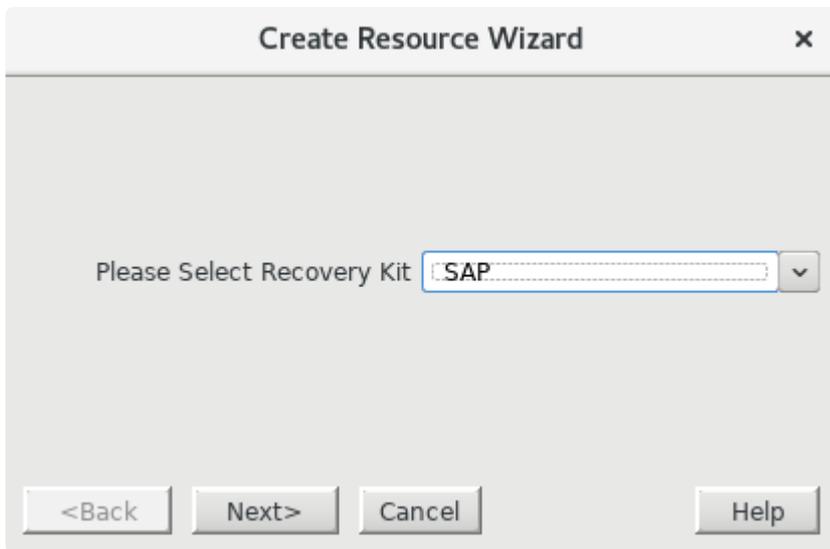
- [AWS/Azure – Create LifeKeeper SAP Resources](#)
- [Google Cloud – Create LifeKeeper SAP Resources](#)

# 11.2.7.5.5.4.1. AWS/Azure – Create LifeKeeper SAP Resources

\* **Note:** This section applies to deployments on AWS and Microsoft Azure. For deployments on Google Cloud, follow the steps given in [Google Cloud – Create LifeKeeper SAP Resources](#).

## Create the ASCS Resource Hierarchy

- In the LifeKeeper GUI, click the  icon to open the **Create Resource Wizard**. Select the “SAP” Recovery Kit.

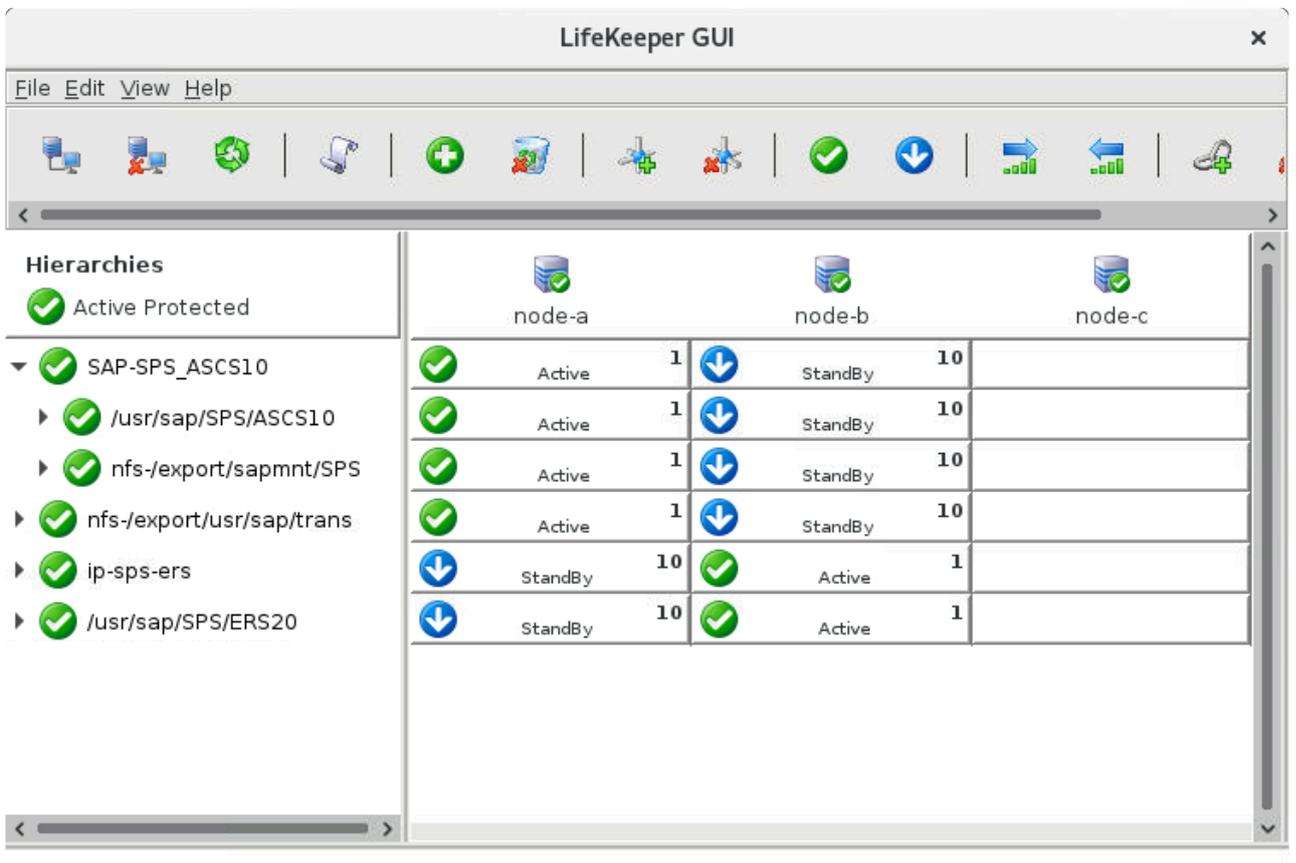


- Enter the following values to create and extend a LifeKeeper resource (**SAP-SPS\_ASCS10**) to protect the ASCS10 instance on node-a and node-b. Notice that the resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

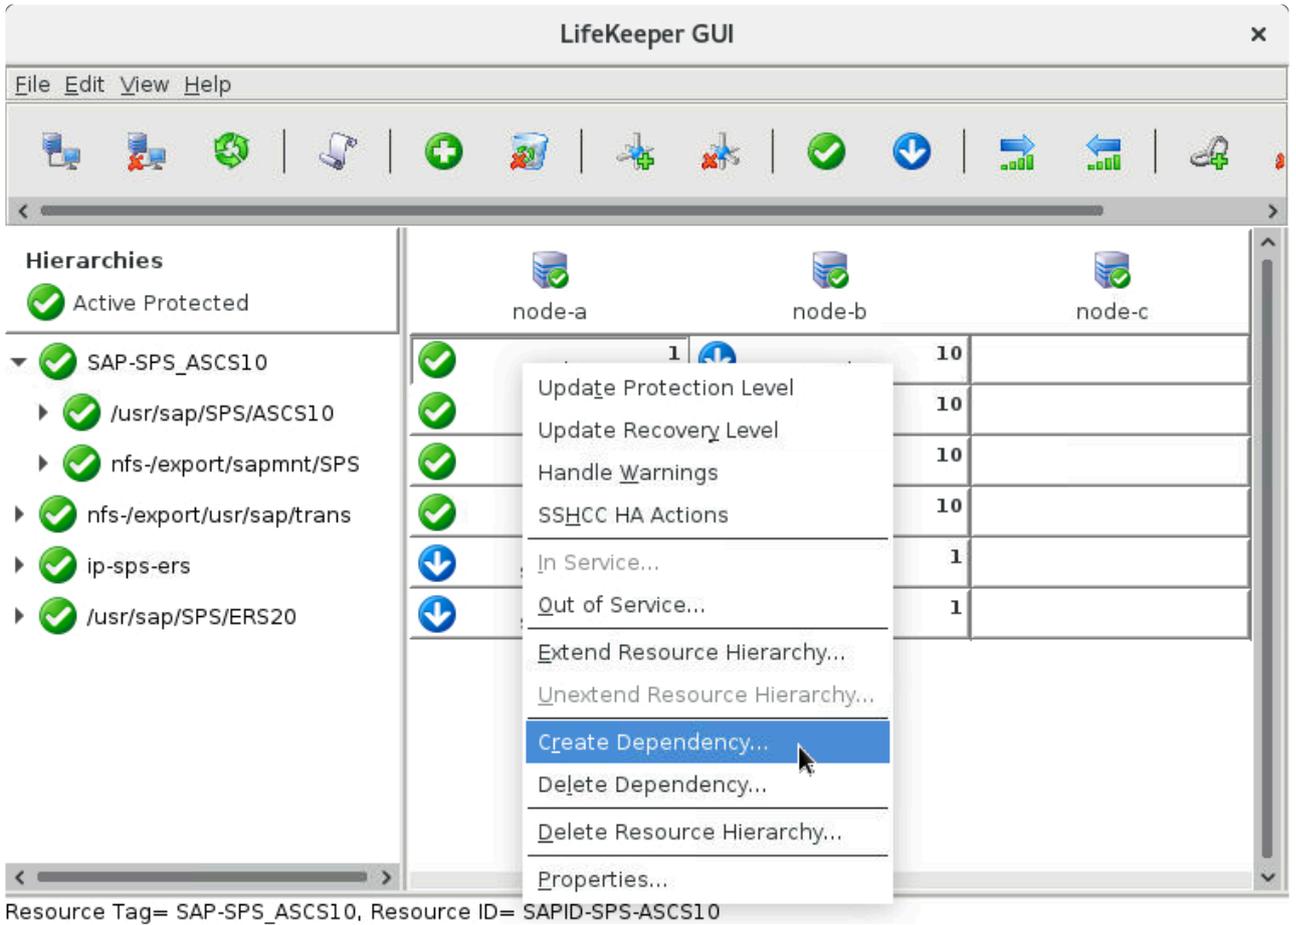
Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
SAP SID	SPS 
SAP Instance for SPS	ASCS10 
IP child resource	ip-sps-ascs

Automate dependent filesystem creation	no
Dependent filesystem resource	/usr/sap/SPS/ASCS10,nfs-/export/sapmnt/SPS
SAP Tag	SAP-SPS_ASCS10
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent
Template Priority	1
Target Priority	10
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Root Tag	SAP-SPS_ASCS10

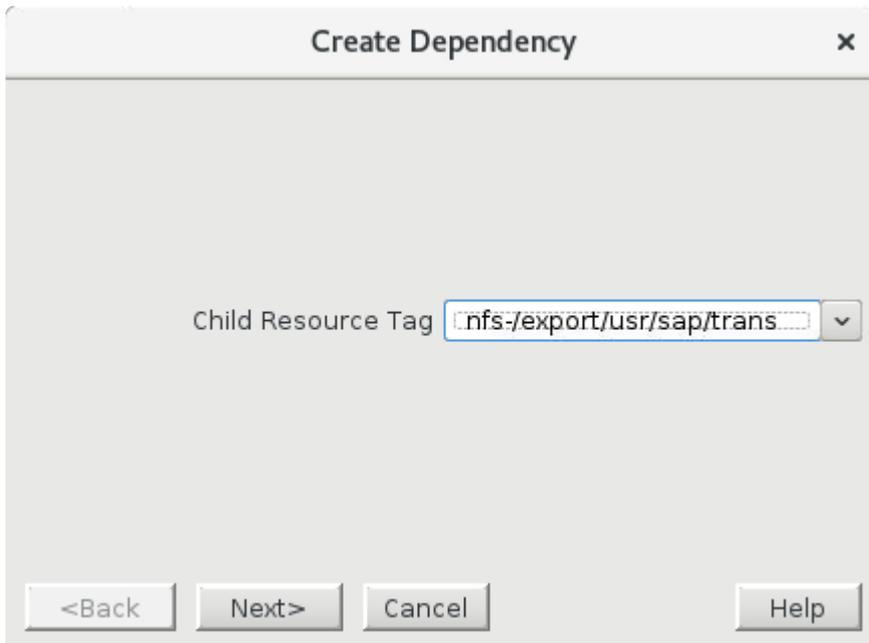
Once the **SAP-SPS\_ASCS10** resource has been successfully created and extended to node-b, the LifeKeeper resource panel should look similar to the following.



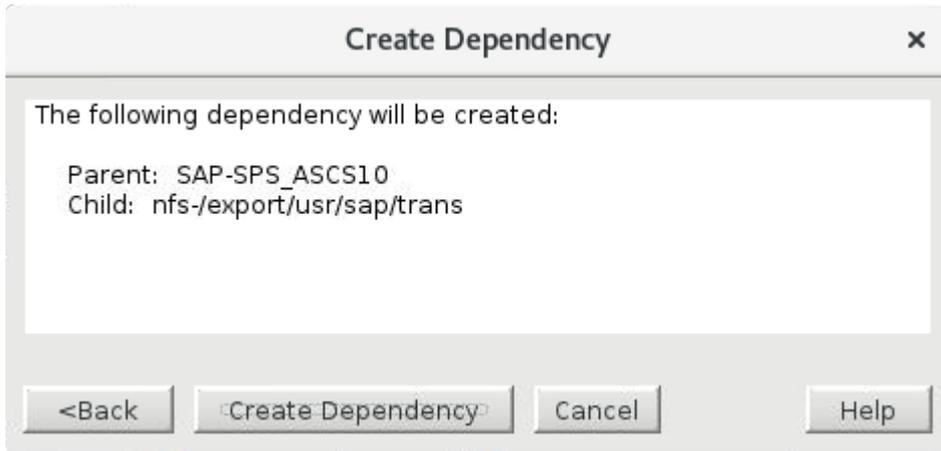
- Right-click on the **SAP-SPS\_ASCS10** resource and select **Create Dependency...** from the drop-down menu.



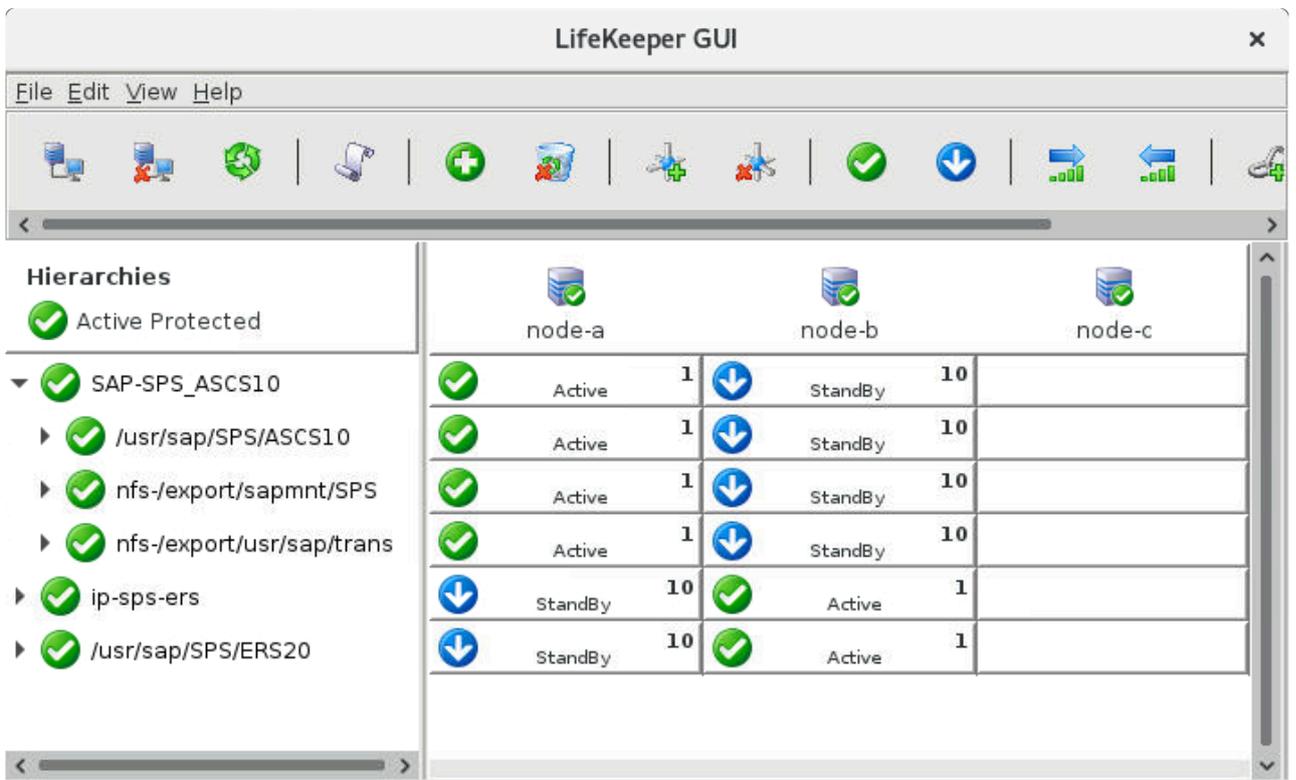
For **Child Resource Tag**, specify the **nfs-/export/usr/sap/trans** resource.



Click **Next>** to continue, then click **Create Dependency** to create the dependency.

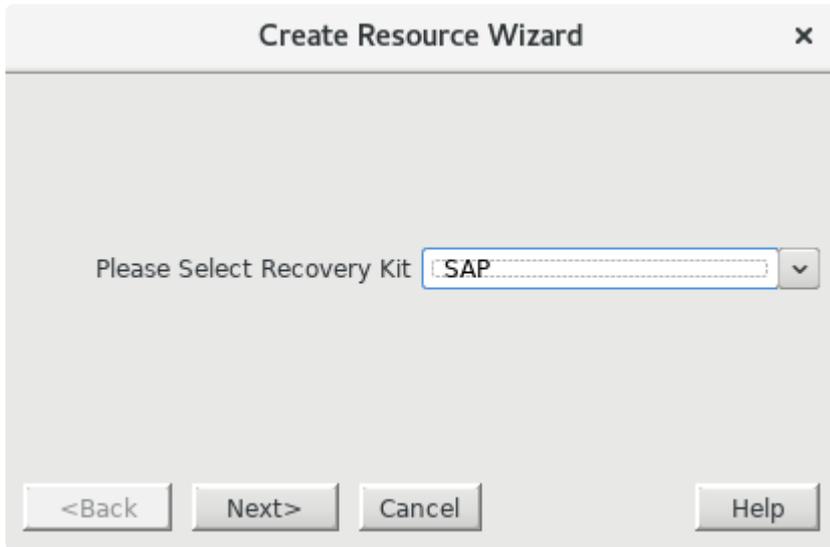


The resulting hierarchy will look similar to the following:



## Create the ERS Resource Hierarchy

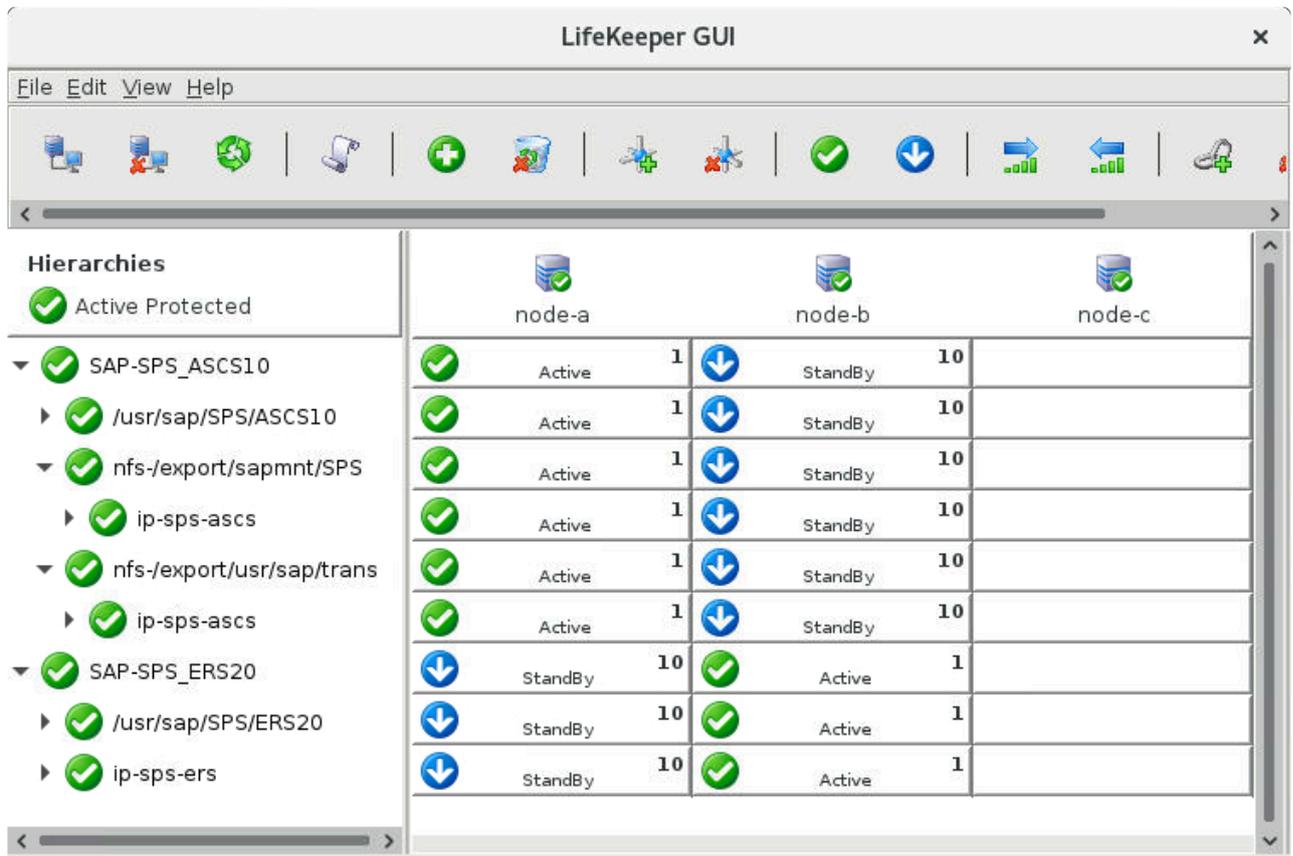
1. In the LifeKeeper GUI, click the  icon to open the **Create Resource Wizard**. Select the “SAP” Recovery Kit.



2. Enter the following values to create and extend a LifeKeeper resource (**SAP-SPS\_ERS20**) to protect the ERS20 instance on node-a and node-b. Notice that this resource is being created on node-b and extended to node-a. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
SAP SID	SPS 
SAP Instance for SPS	ERS20 
IP child resource	ip-sps-ers
Dependent filesystem resource	/usr/sap/SPS/ERS20
SAP Tag	SAP-SPS_ERS20 
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Root Tag	SAP-SPS_ERS20 

Once the **SAP-SPS\_ERS20** resource has been successfully created and extended to node-a, the LifeKeeper resource panel should look similar to the following.



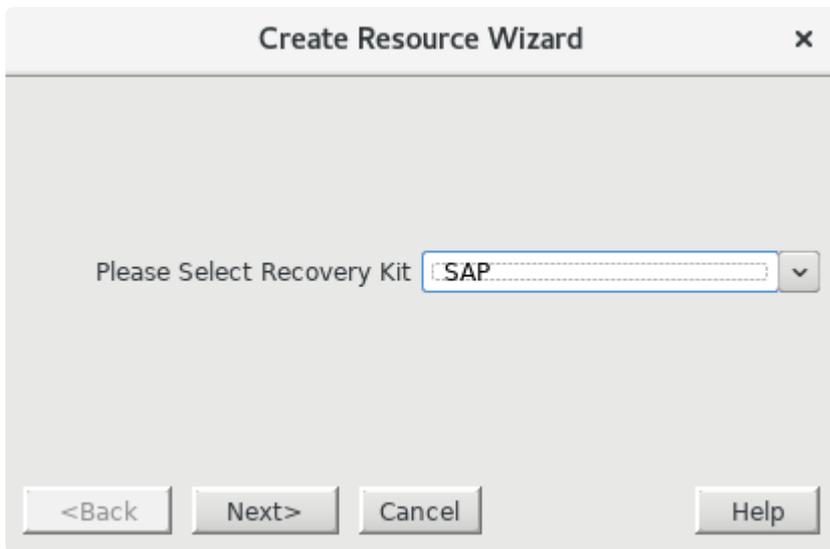
The ASCS10 and ERS20 instances have now been successfully protected by LifeKeeper.

# 11.2.7.5.5.4.2. Google Cloud – Create LifeKeeper SAP Resources

\* **Note:** This section applies to deployments on Google Cloud. For deployments on AWS or Microsoft Azure, follow the steps given in [AWS/Azure – Create LifeKeeper SAP Resources](#).

## Create the ASCS Resource Hierarchy

1. In the LifeKeeper GUI, click the  icon to open the **Create Resource Wizard**. Select the “SAP” Recovery Kit.

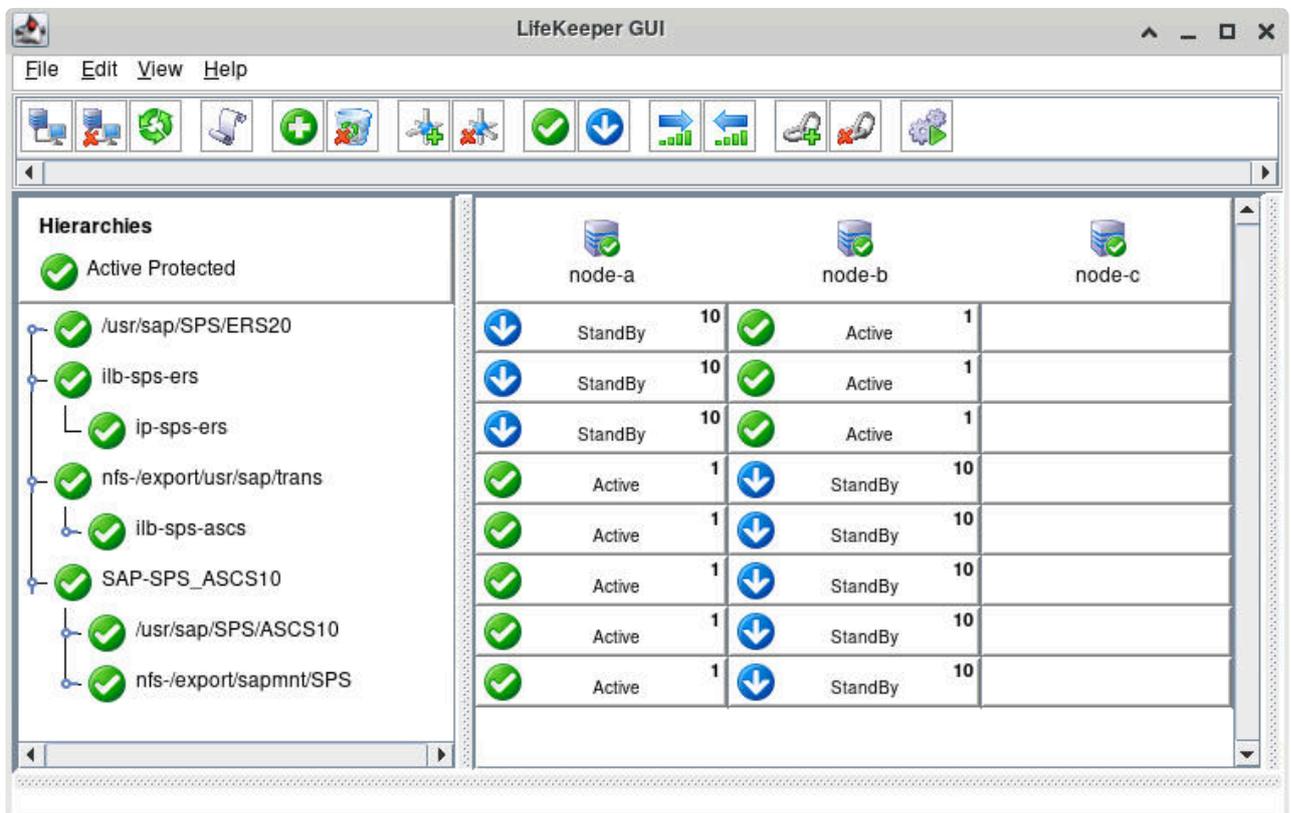


2. Enter the following values to create and extend a LifeKeeper resource (**SAP-SPS\_ASCS10**) to protect the ASCS10 instance on node-a and node-b. Notice that the resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
SAP SID	SPS 
SAP Instance for SPS	ASC10 

IP child resource	none ✓
Automate dependent filesystem creation	no ✓
Dependent filesystem resource	/usr/sap/SPS/ASCS10,nfs-/export/sapmnt/SPS
SAP Tag	SAP-SPS_ASCS10 ✓
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Root Tag	SAP-SPS_ASCS10 ✓

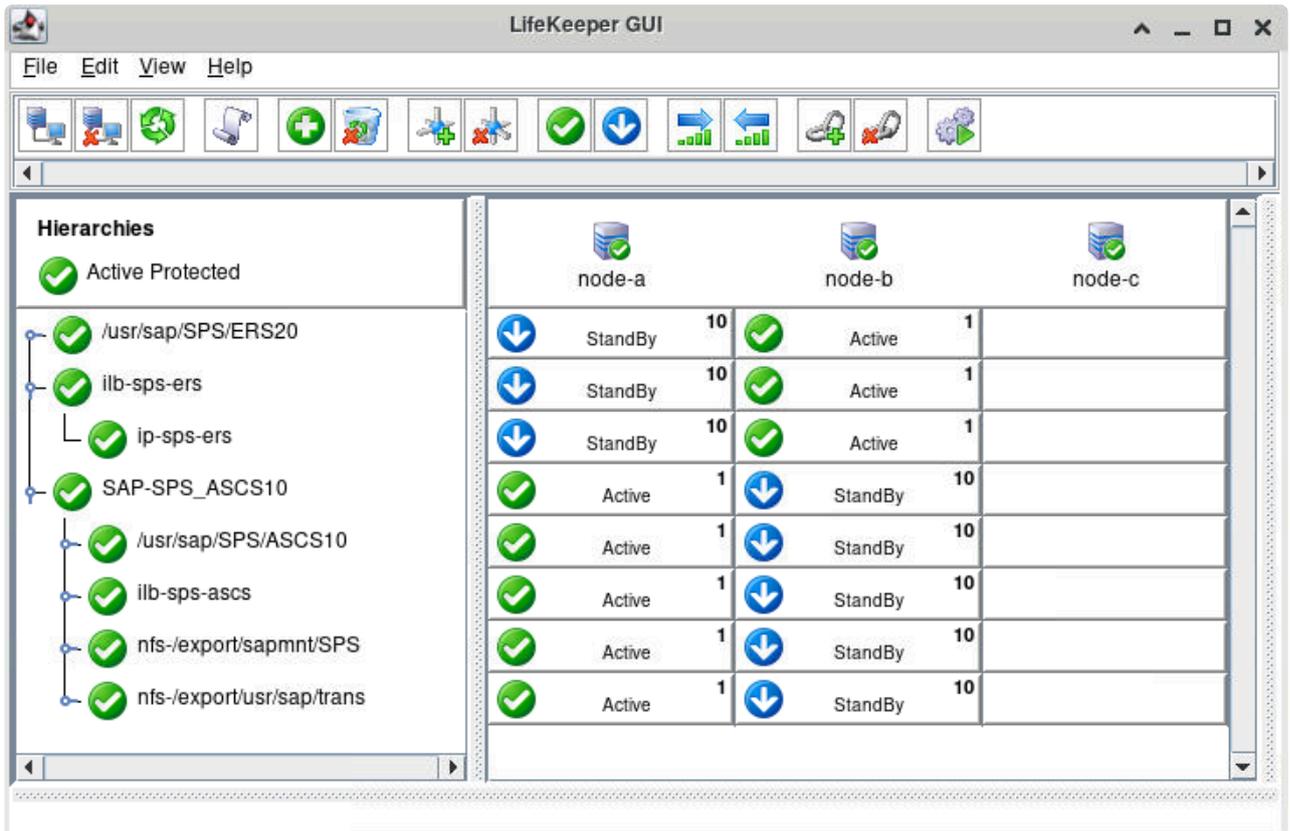
Once the **SAP-SPS\_ASCS10** resource has been successfully created and extended to node-b, the LifeKeeper resource panel should look similar to the following.



- Right-click on the **SAP-SPS\_ASCS10** resource and select **Create Dependency...** from the drop-down menu. For **Child Resource Tag**, select **nfs-/export/usr/sap/trans**. Click **Next>** to continue, then click **Create Dependency** to create the dependency.

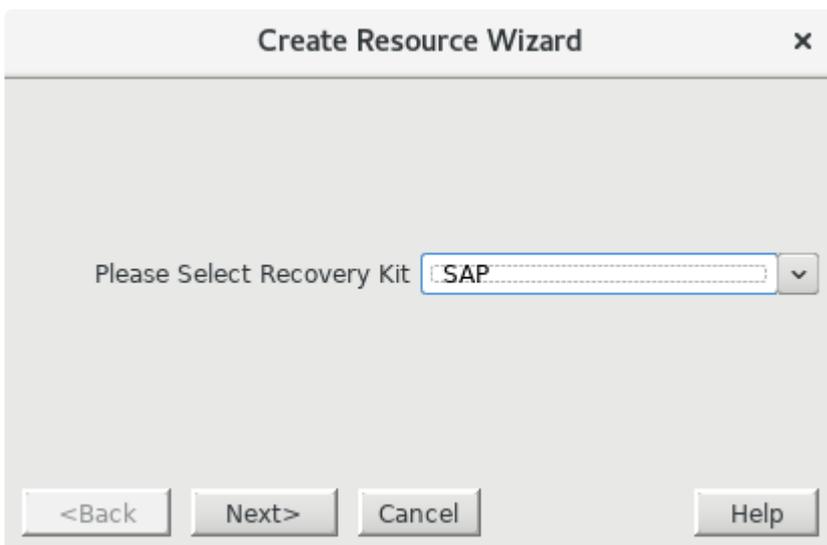
- Right-click on the **SAP-SPS\_ASCS10** resource and select **Create Dependency...** from the drop-down menu. For **Child Resource Tag**, select **ilb-sps-ascs**. Click **Next>** to continue, then click **Create Dependency** to create the dependency.

The resulting hierarchy will look similar to the following:



## Create the ERS Resource Hierarchy

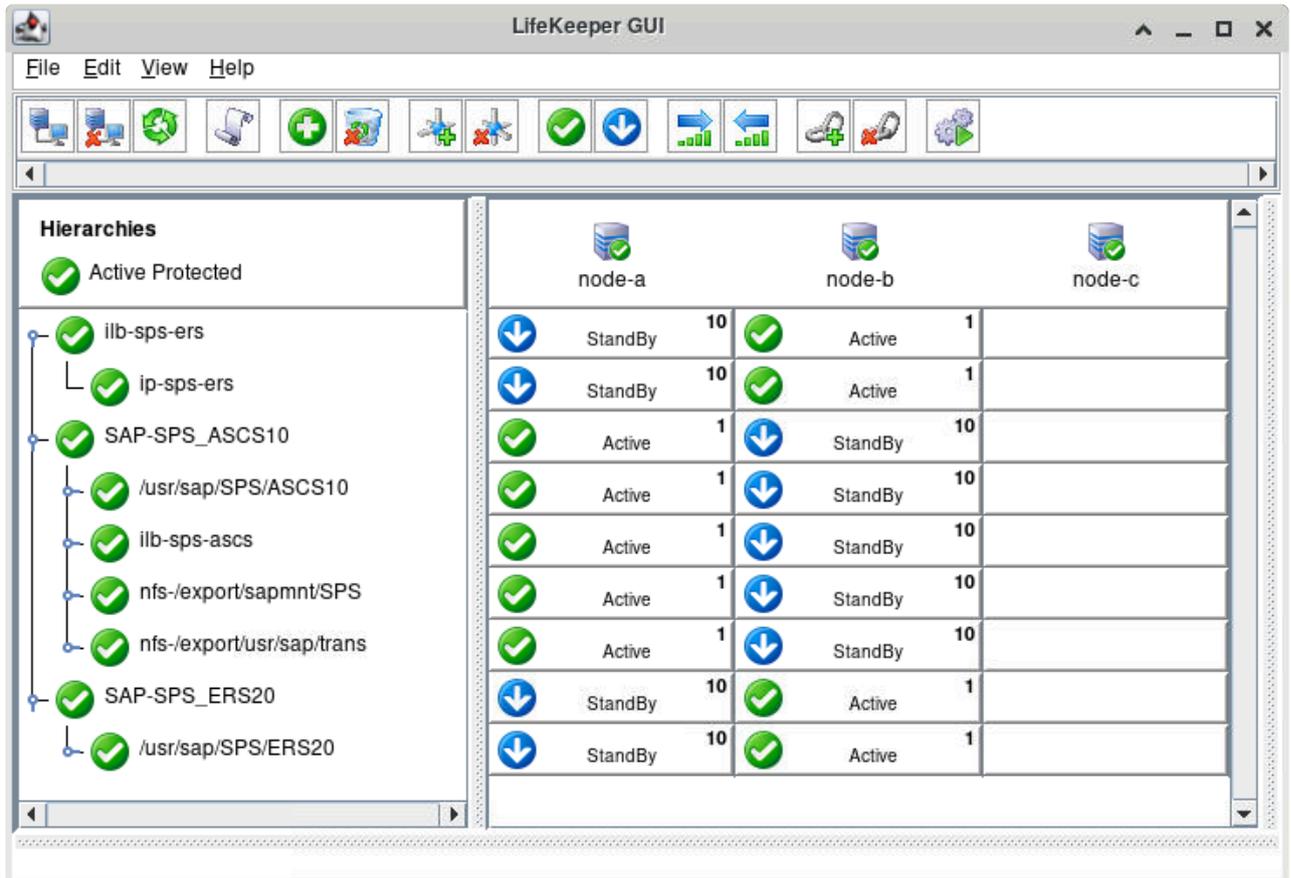
- In the LifeKeeper GUI, click the  icon to open the **Create Resource Wizard**. Select the "SAP" Recovery Kit.



2. Enter the following values to create and extend a LifeKeeper resource (**SAP-SPS\_ERS20**) to protect the ERS20 instance on node-a and node-b. Notice that this resource is being created on node-b and extended to node-a. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

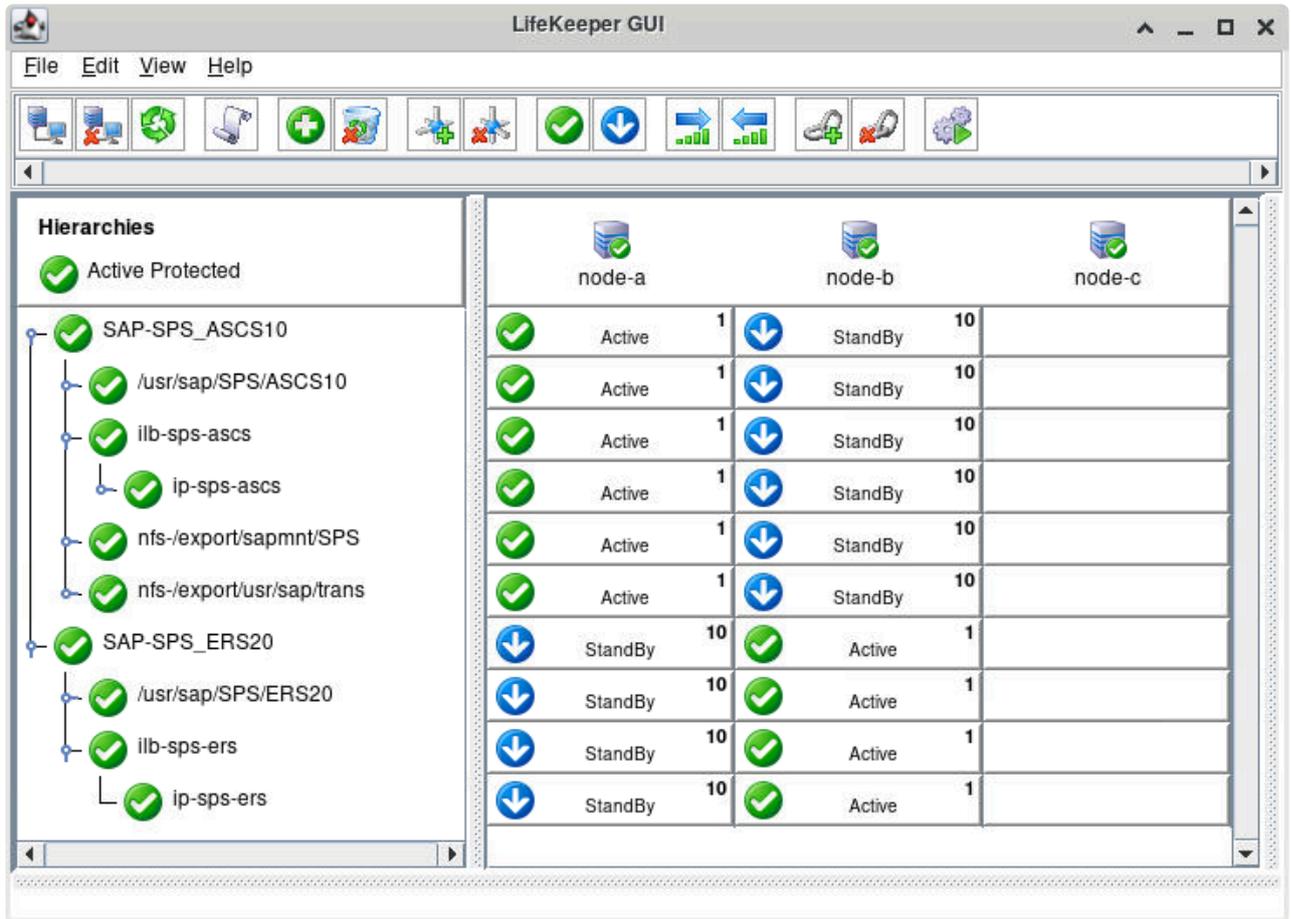
Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-b
SAP SID	SPS 
SAP Instance for SPS	ERS20 
IP child resource	none 
Automate dependent filesystem creation	no 
Dependent filesystem resource	/usr/sap/SPS/ERS20
SAP Tag	SAP-SPS_ERS20 
<b>Pre-Extend Wizard</b>	
Target Server	node-a
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Root Tag	SAP-SPS_ERS20 

Once the **SAP-SPS\_ERS20** resource has been successfully created and extended to node-a, the LifeKeeper resource panel should look similar to the following.



- Right-click on the **SAP-SPS\_ERS20** resource and select **Create Dependency...** from the drop-down menu. For **Child Resource Tag**, select **ilb-sps-ers**. Click **Next>** to continue, then click **Create Dependency** to create the dependency.

The resulting hierarchy will look similar to the following:



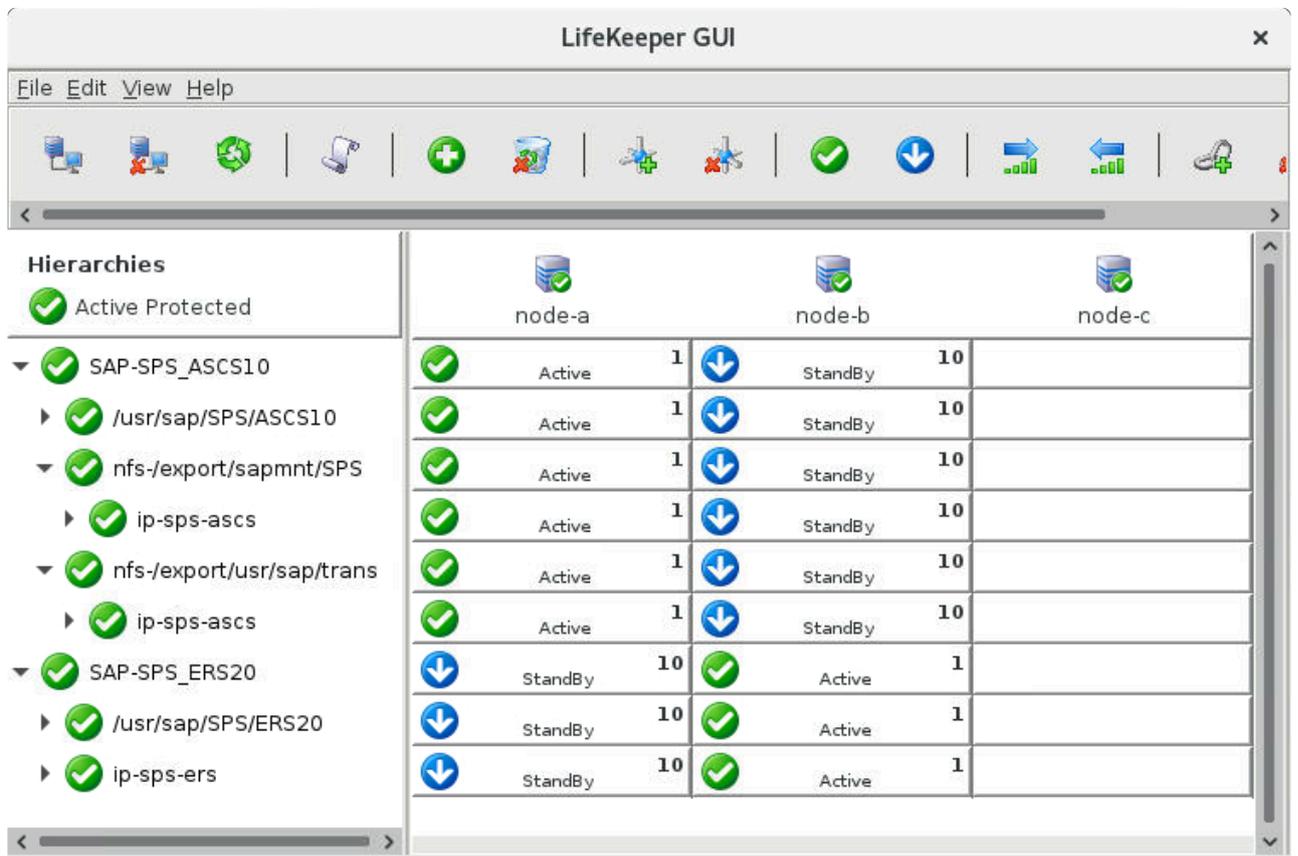
The ASCS10 and ERS20 instances have now been successfully protected by LifeKeeper.

# 11.2.7.5.5. Test Switchover and Failover

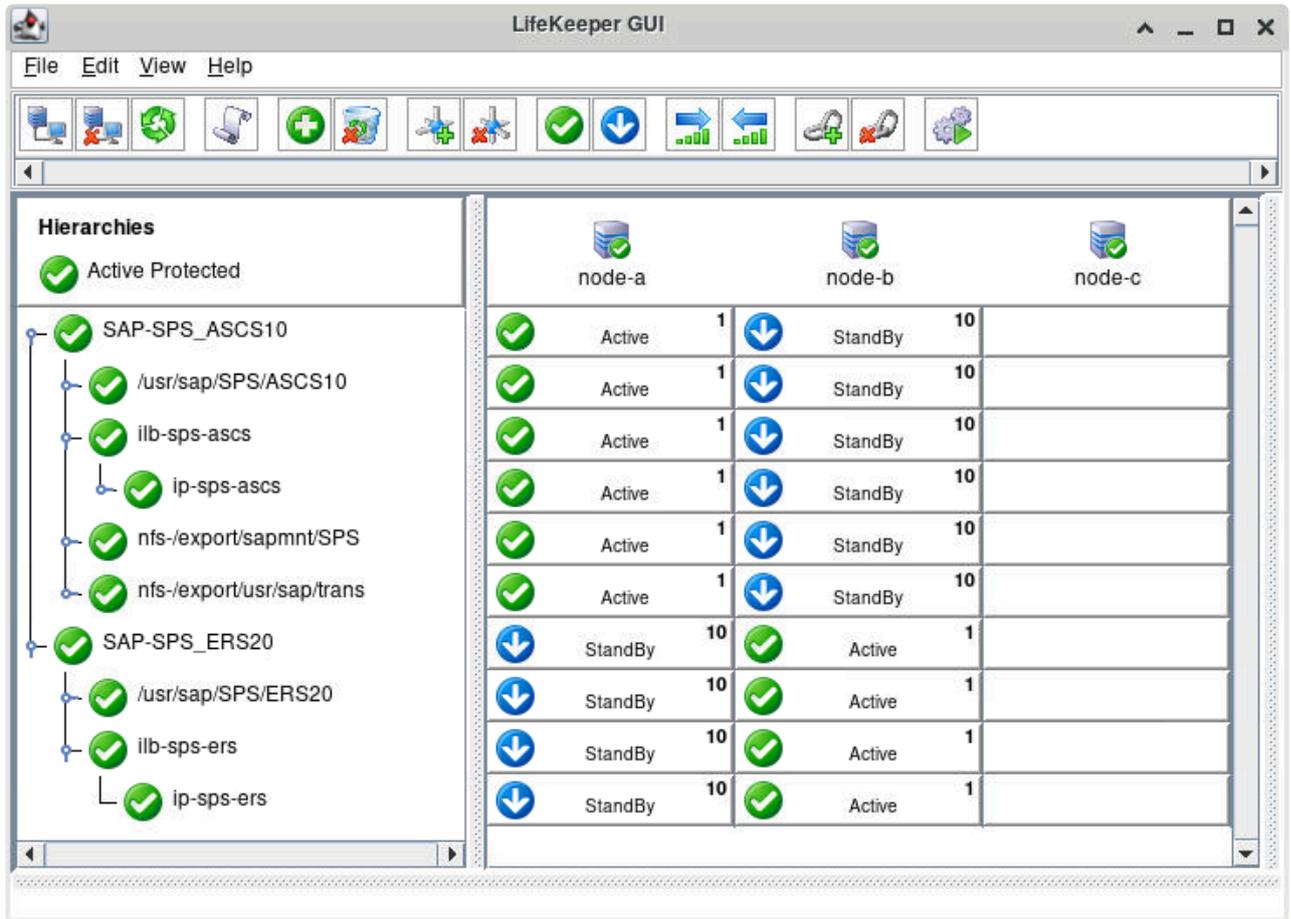
In this section we will perform basic tests to verify the expected behavior of the **SAP-SPS\_ASCS10** and **SAP-SPS\_ERS20** resource hierarchies on switchover and failover. It is important to test that the enqueue server process in the ASCS10 instance is able to successfully recover the enqueue lock table from the ERS20 instance after switchover or failover.

1. Verify that the **SAP-SPS\_ASCS10** resource state is currently Active on node-a and Standby on node-b, and that the **SAP-SPS\_ERS20** resource state is currently Active on node-b and Standby on node-a.

On AWS or Azure, the LifeKeeper GUI should resemble the following image:



On Google Cloud, the LifeKeeper GUI should resemble the following image:



- Execute the following commands to verify that the ASCS10 and ERS20 instances are running successfully on node-a and node-b, respectively:

```
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 10 -function GetProcessList"
04.03.2021 20:24:12
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
msg_server, MessageServer, GREEN, Running, 2020 12 21 16:53:00, 1755:31:12, 11497
enq_server, Enqueue Server 2, GREEN, Running, 2020 12 21 16:53:00, 1755:31:12, 11498

[root@node-b ~]# su - spsadm -c "sapcontrol -nr 20 -function GetProcessList"
04.03.2021 20:24:22
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
enq_replicator, Enqueue Replicator 2, GREEN, Running, 2021 02 22 16:55:17, 243:29:05, 30028
```

- Execute the following command on node-a to write 100 exclusive non-cumulative locks labeled 0-99 to the lock table maintained by the enqueue server.

```
[root@node-a ~]# su - spsadm -c "enq_admin --set_locks=100:X:DIAG::TAB:%u p
f=/usr/sap/SPS/SYS/profile/SPS_ASCS10_sps-ascs"
Enqueue Server 2

2021-03-04 20:32:16; OK; 'Set Locks'; Response=41496 usec
=====
```

Execute the following commands to verify that the locks have been successfully stored in the lock table on node-a and replicated to the enqueue replication server on node-b:

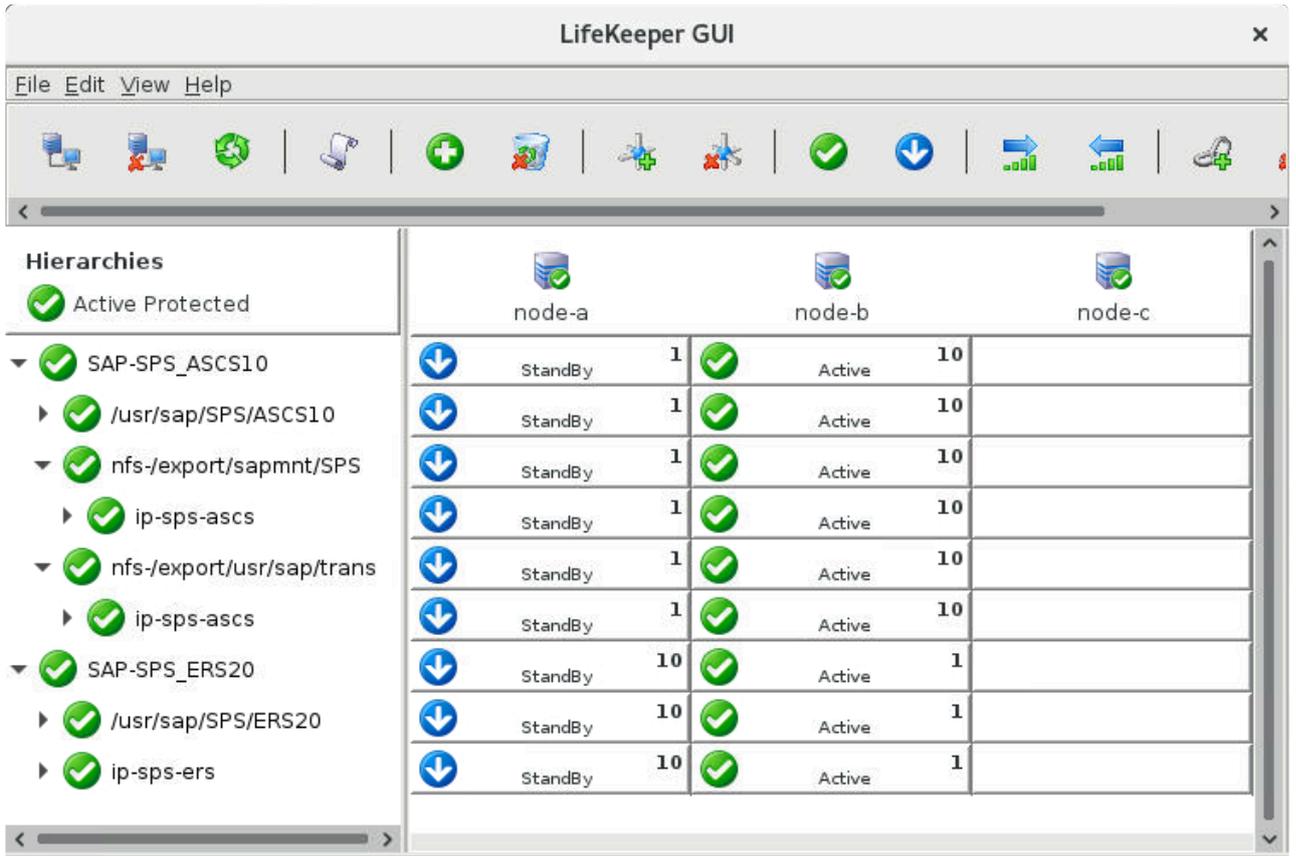
```
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 10 -function EnqGetStatistic"
| grep locks_now
locks_now: 100

[root@node-b ~]# su - spsadm -c "sapcontrol -nr 20 -function EnqGetStatistic"
| grep locks_now
locks_now: 100
```

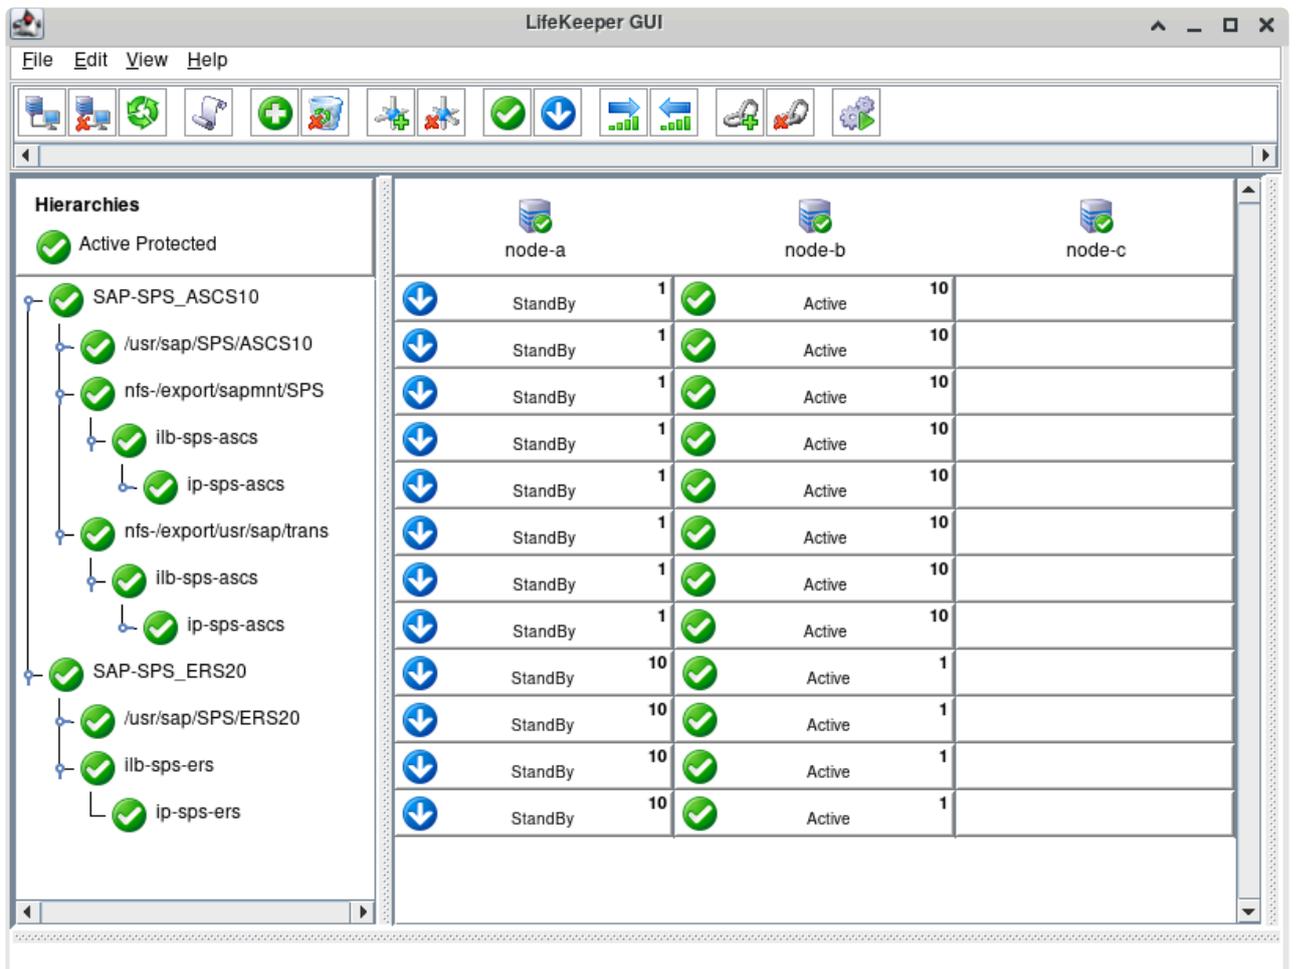
**!** If the locks are not being successfully replicated to the Enqueue Replication Server on node-b in a Google Cloud deployment, please verify that IP forwarding has been disabled on both node-a and node-b as described in the **Disable IP Forwarding** section of [Google Cloud – Using an Internal Load Balancer](#). In this configuration, also verify that each GenLB resource has a dependent IP resource protecting the frontend IP address of the corresponding load balancer using network mask 255.255.255.255. Without completing these configuration steps on Google Cloud, the Enqueue Server and Enqueue Replication Server will be unable to communicate with each other through the frontend IP addresses of their corresponding internal load balancers when running on different cluster nodes.

4. Perform a switchover of the ASCS resource hierarchy by right-clicking the **SAP-SPS\_ASCS10** resource on node-b and choosing the **In-Service...** operation. Click **In Service** to begin the switchover. Once the switchover is complete, the **SAP-SPS\_ASCS10** and **SAP-SPS\_ERS20** resources will both be in-service on node-b.

On AWS or Azure, the LifeKeeper GUI should resemble the following image:

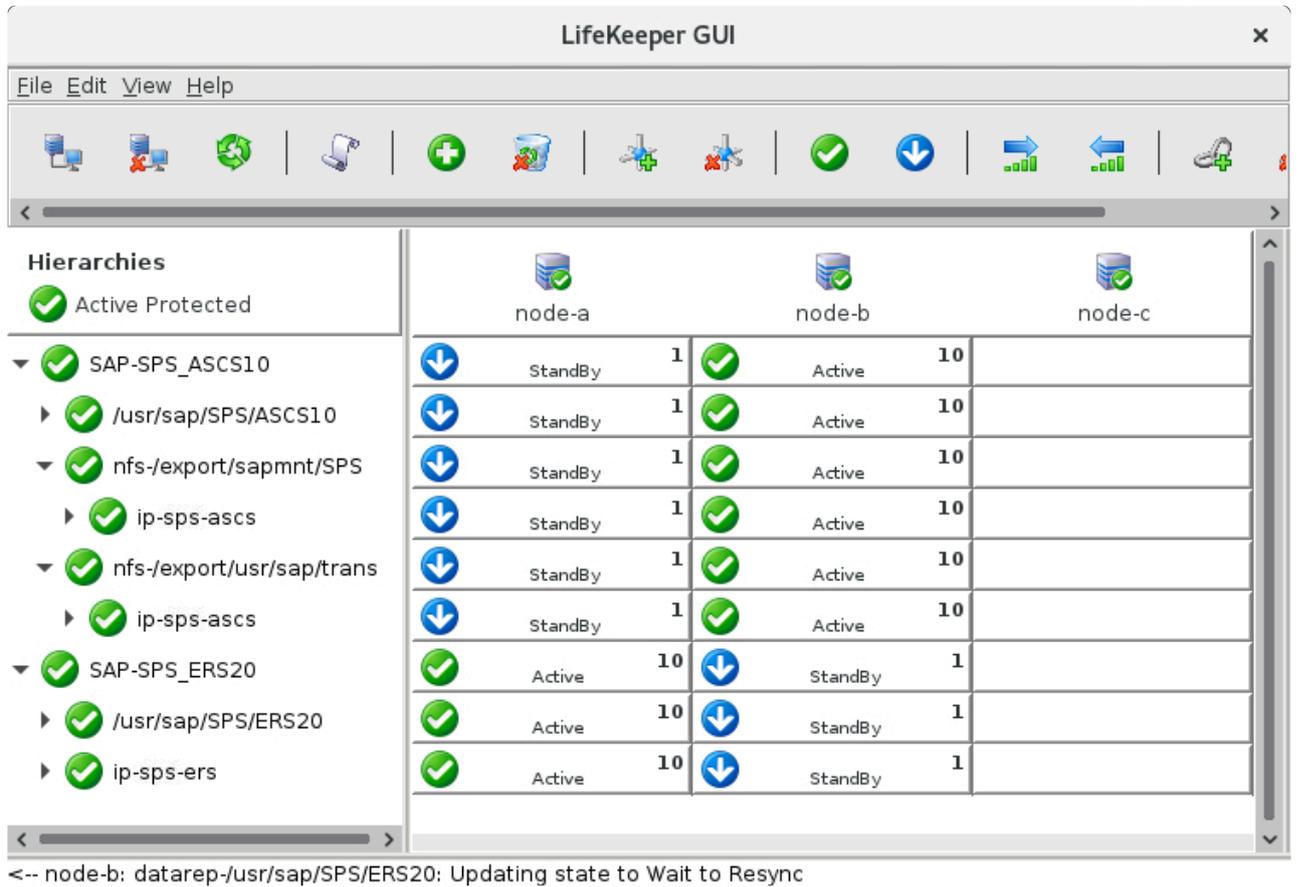


On Google Cloud, the LifeKeeper GUI should resemble the following image:

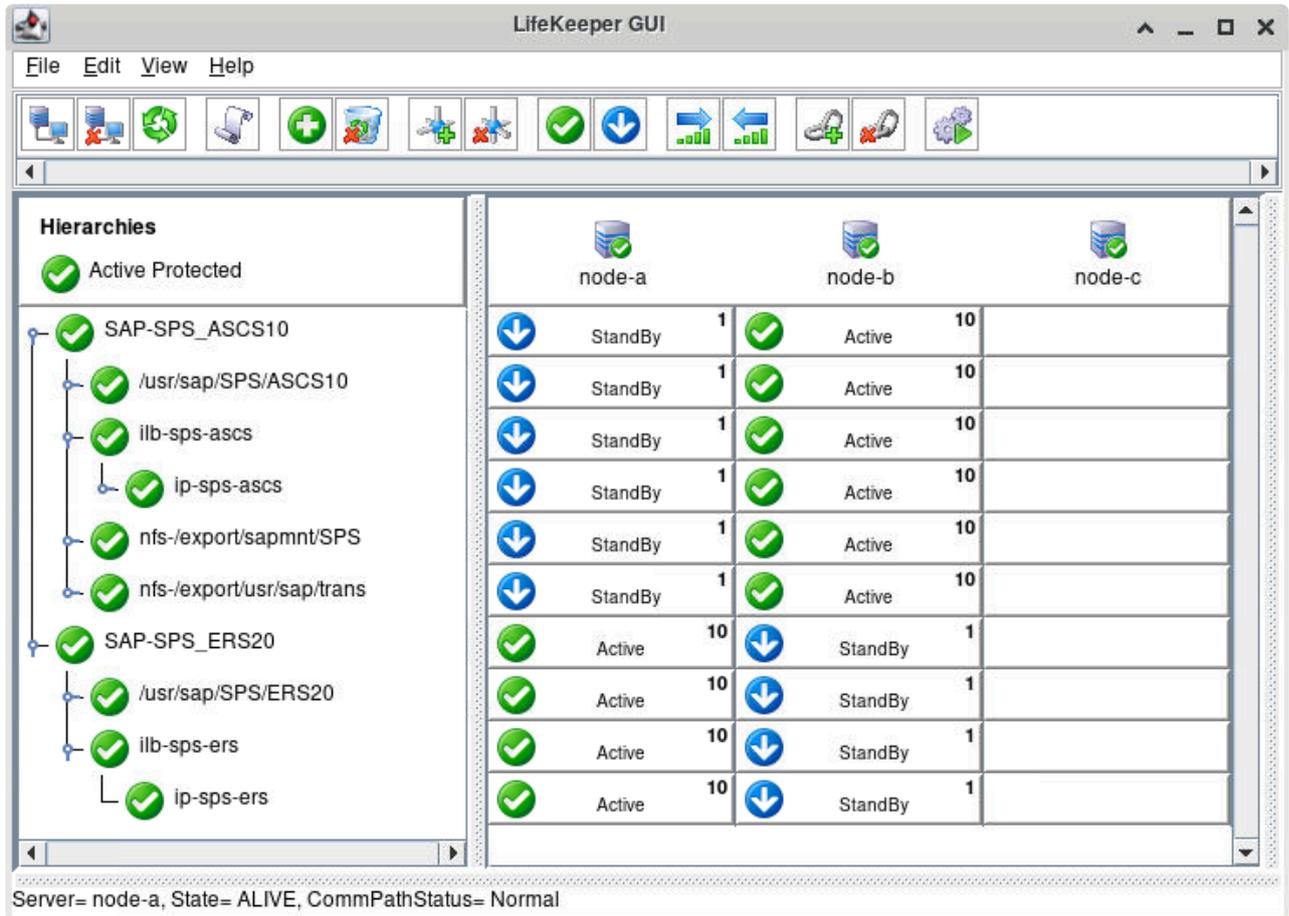


Once the ASCS resource hierarchy has successfully come in-service on node-b and the enqueue server process has obtained the copy of the backup enqueue lock table from the enqueue replication server process, LifeKeeper will automatically relocate the **SAP-SPS\_ERS20** resource to node-a to provide lock table redundancy across cluster nodes. This process may take several minutes to complete.

On AWS or Azure, the LifeKeeper GUI should resemble the following image:



On Google Cloud, the LifeKeeper GUI should resemble the following image:



- Once LifeKeeper has relocated the SAP-SPS\_ERS20 resource back to node-a, execute the following commands to verify that the ASCS10 and ERS20 instances are running successfully on node-b and node-a, respectively, and that they both still hold the 100 locks written in step 3.

```
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 20 -function GetProcessList"
04.03.2021 20:58:57
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
enq_replicator, Enqueue Replicator 2, GREEN, Running, 2021 03 04 20:57:34, 0:0
1:23, 21967

[root@node-a ~]# su - spsadm -c "sapcontrol -nr 20 -function EnqGetStatistic"
| grep locks_now
locks_now: 100

[root@node-b ~]# su - spsadm -c "sapcontrol -nr 10 -function GetProcessList"
04.03.2021 20:56:56
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
msg_server, MessageServer, GREEN, Running, 2021 03 04 20:54:47, 0:02:09, 17074
enq_server, Enqueue Server 2, GREEN, Running, 2021 03 04 20:54:47, 0:02:09, 17
075
```

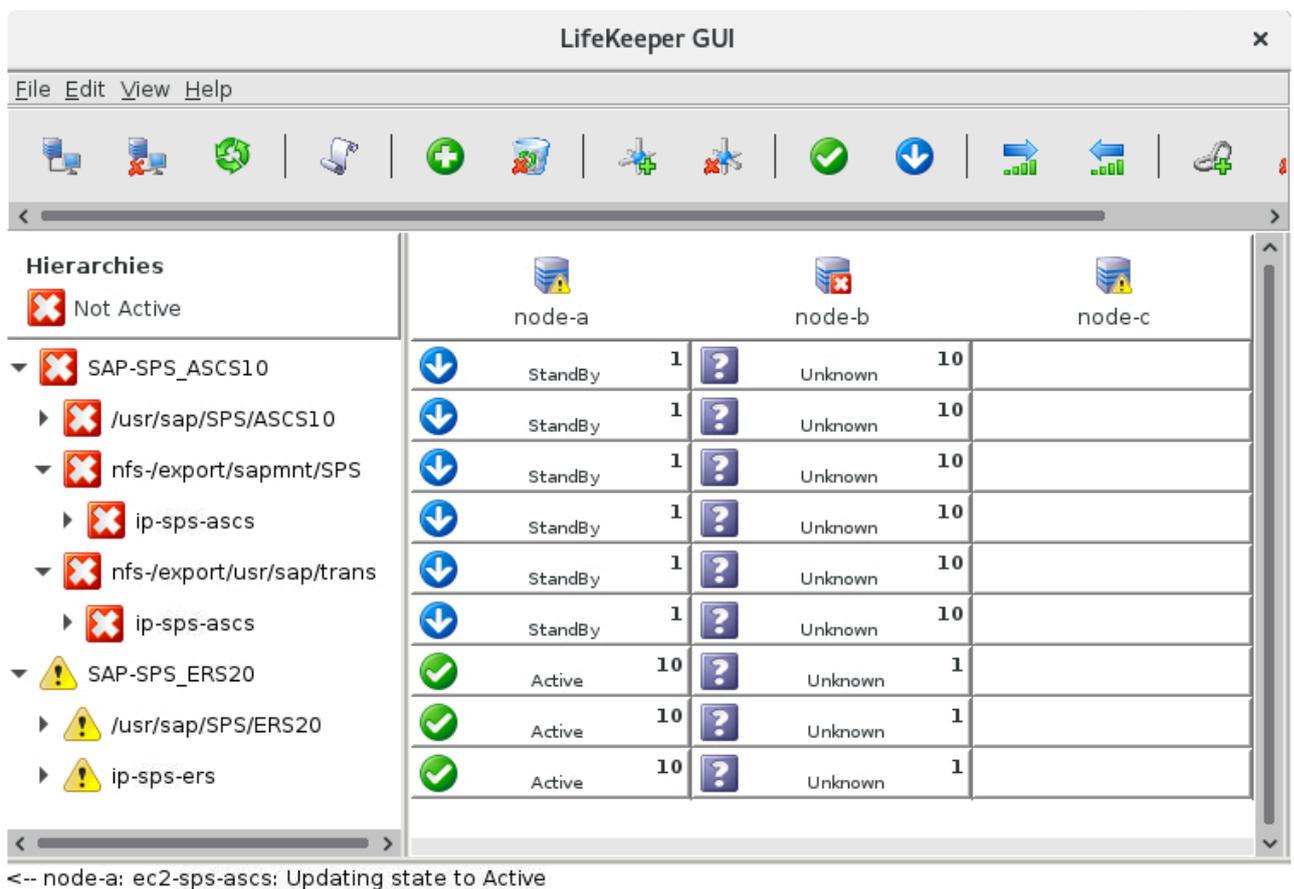
```
[root@node-b ~]# su - spsadm -c "sapcontrol -nr 10 -function EnqGetStatistic"
| grep locks_now
locks_now: 100
```

6. Execute the following command to forcefully reboot node-b:

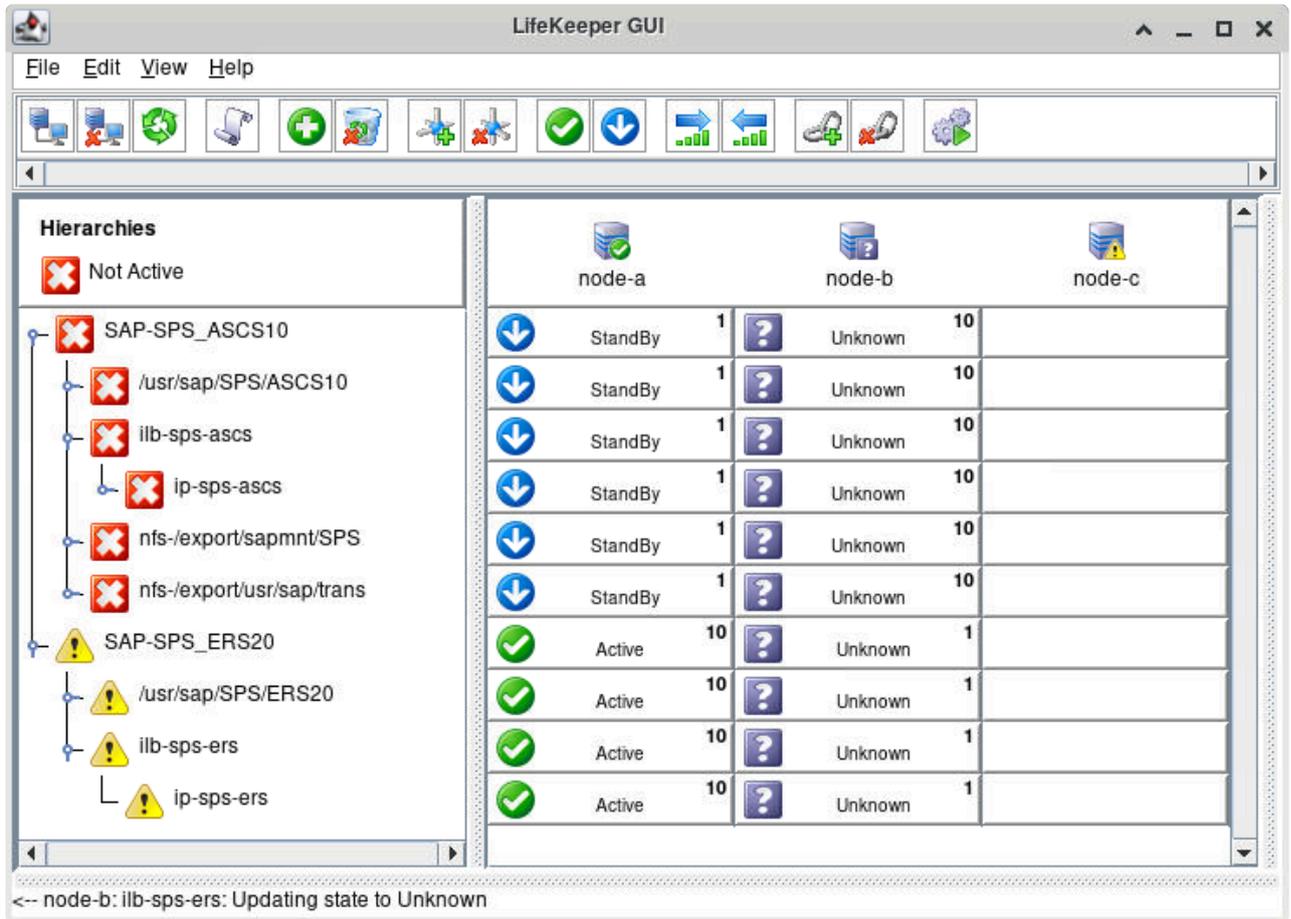
```
[root@node-b ~]# echo b > /proc/sysrq-trigger
```

Once LifeKeeper has detected that node-b has been powered off, the status of node-b updates to "Unknown" in the LifeKeeper GUI.

On AWS or Azure, the LifeKeeper GUI should resemble the following image:

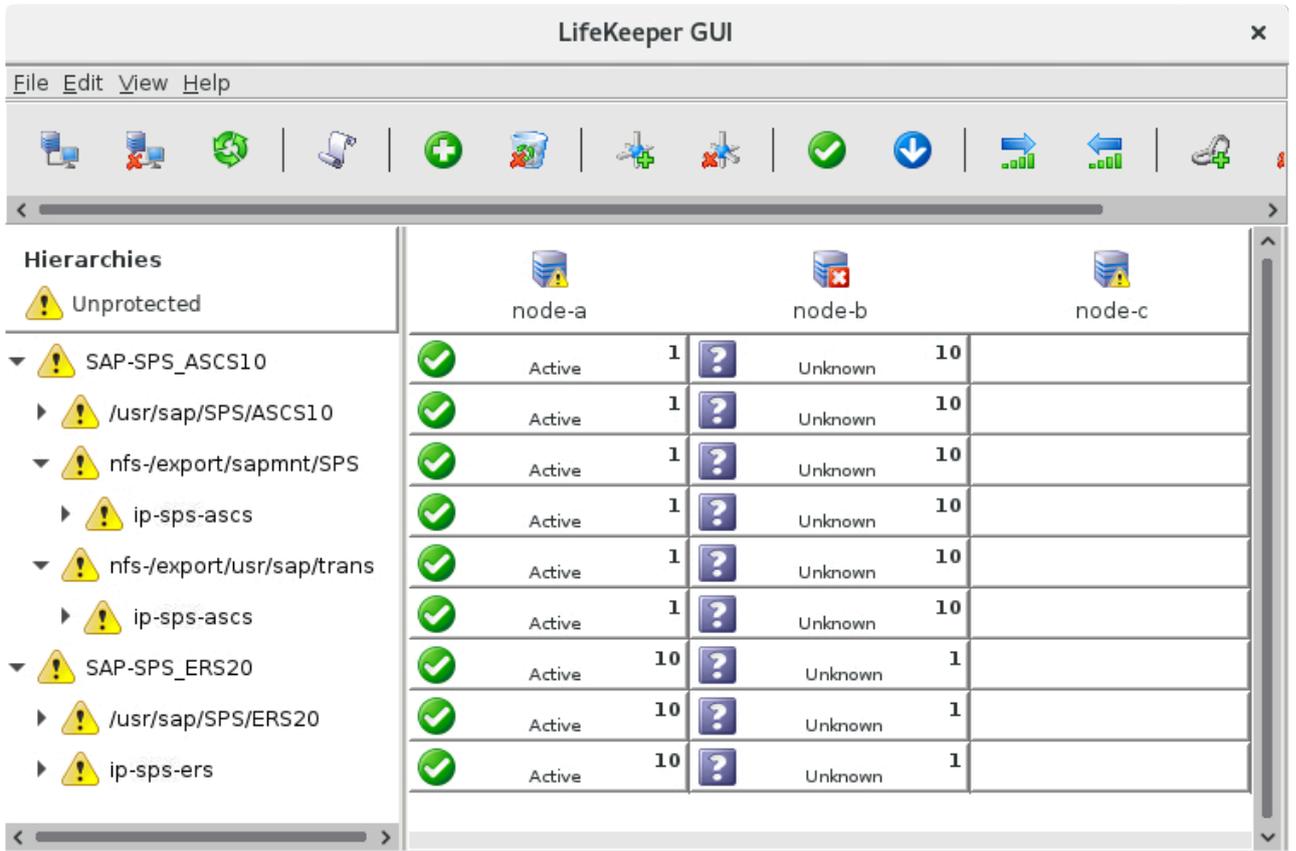


On Google Cloud, the LifeKeeper GUI should resemble the following image:



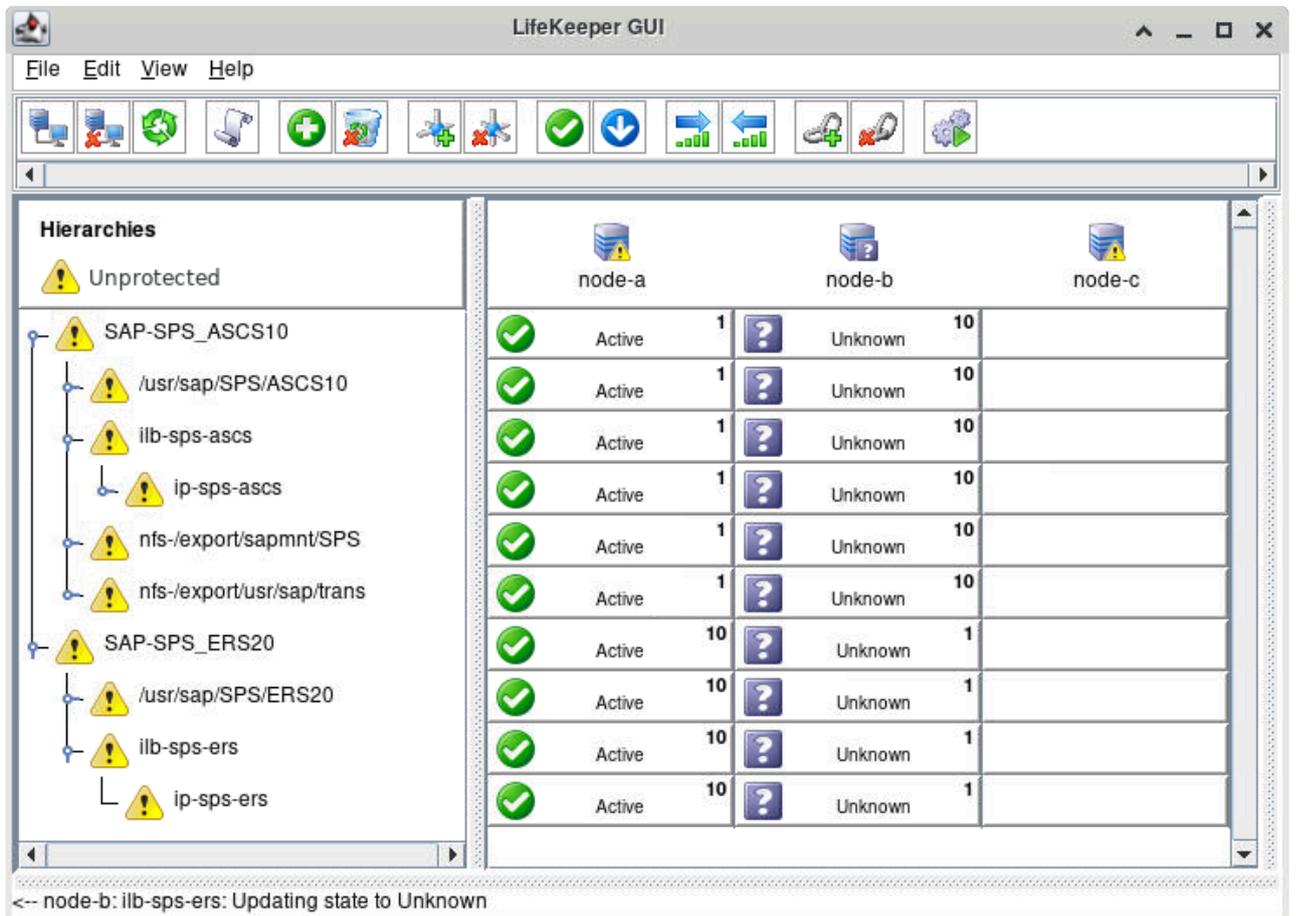
At this point, LifeKeeper will initiate automatic failover of the **SAP-SPS\_ASCS10** resource hierarchy back to node-a. The **SAP-SPS\_ASCS10** and **SAP-SPS\_ERS20** resource hierarchies will both be Active on node-a until node-b comes back online.

On AWS or Azure, the LifeKeeper GUI should resemble the following image:



<-- node-a: ec2-sps-asc: Updating state to Active

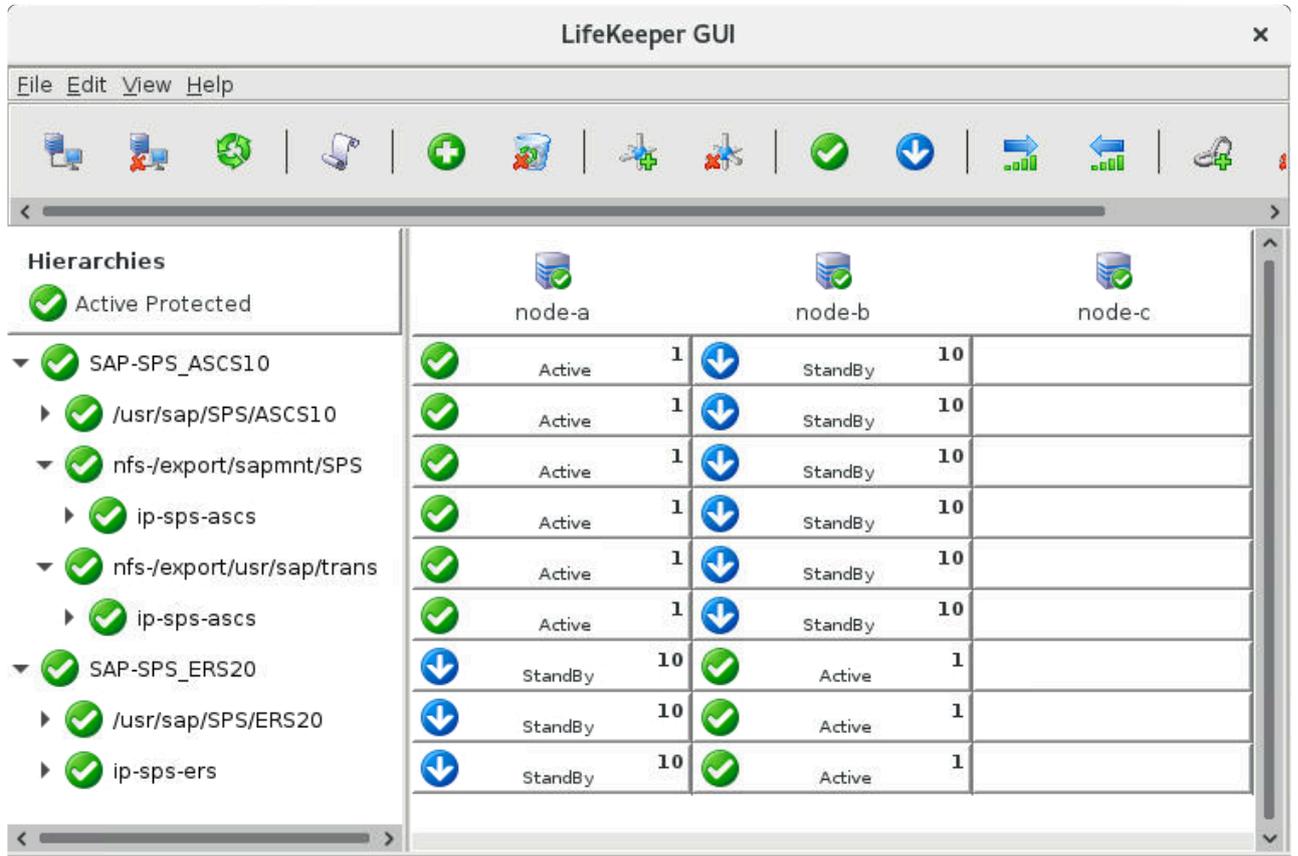
On Google Cloud, the LifeKeeper GUI should resemble the following image:



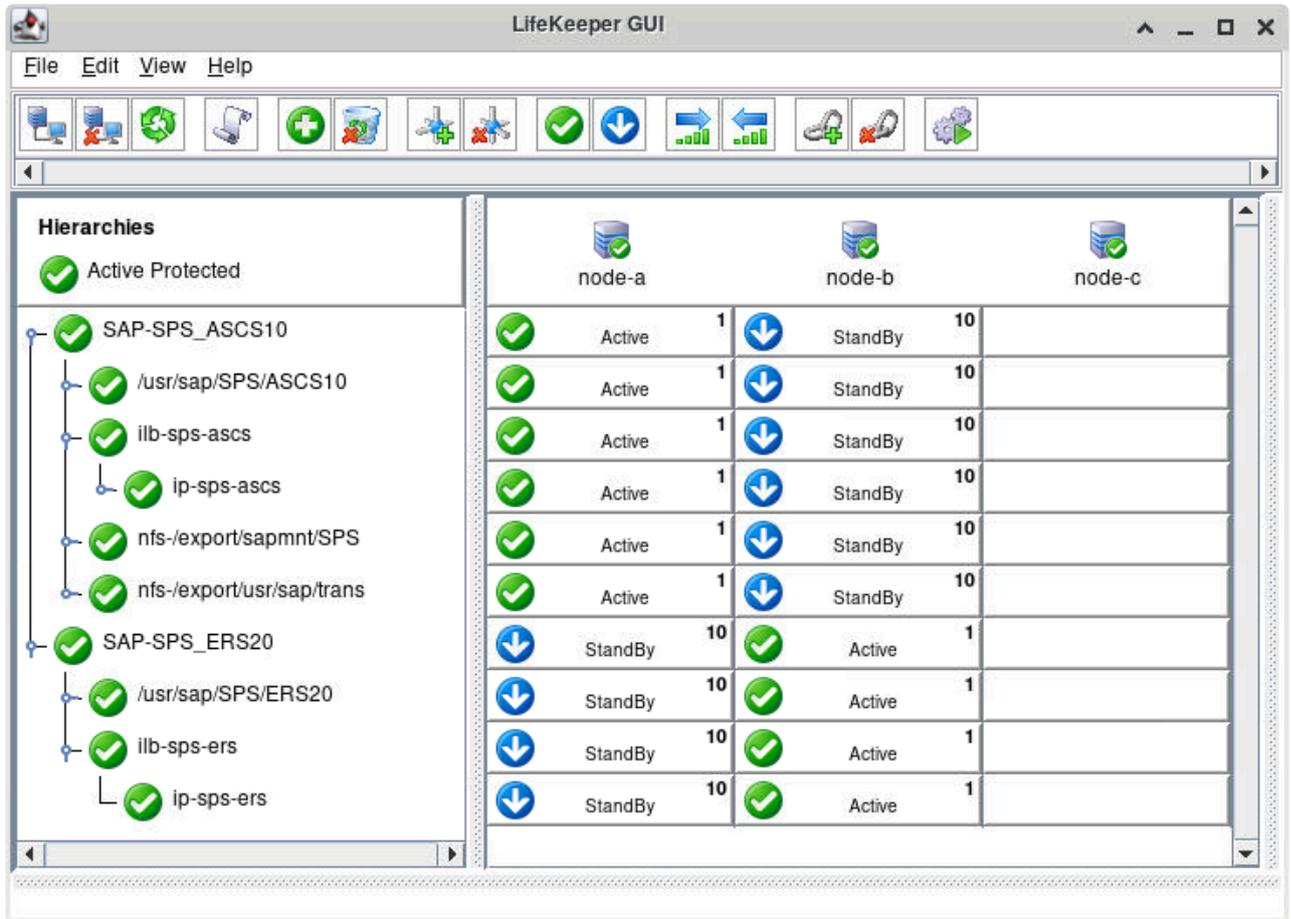
<-- node-b: ilb-sps-ers: Updating state to Unknown

Once node-b is back online, LifeKeeper will automatically relocate the **SAP-SPS\_ERS20** resource hierarchy back to node-b. This process may take several minutes to complete. Once this process is complete, the **SAP-SPS\_ASCS10** and **SAP-SPS\_ERS20** resource hierarchies will be back in-service on node-a and node-b, respectively.

On AWS or Azure, the LifeKeeper GUI should resemble the following image:



On Google Cloud, the LifeKeeper GUI should resemble the following image:



- Execute the sapcontrol commands given in steps 2 and 3 again to verify the expected state on each node.
- Execute the following command on node-a to release the 100 locks that were written in step 3:

```
[root@node-a ~]# su - spsadm -c "enq_admin --release_locks=100:X:DIAG::TAB:%u
pf=/usr/sap/SPS/SYS/profile/SPS_ASCS10_sps-asc"
Enqueue Server 2

2021-03-04 21:10:22; OK; 'Release Locks'; Response=36883 usec
=====
```

We have now verified the basic switchover and failover functionality of the ASCS and ERS resource hierarchies.

## 11.2.7.5.6. Protecting SAP HANA Resources

In this section we will deploy a highly-available SAP HANA 2.0 SPS04 cluster. Our configuration will use SAP HANA System Replication to replicate data from the active database host to the standby database host and will use the LifeKeeper SAP HANA Recovery Kit to handle monitoring, recovery, and failover the HDB database instance.

The details of the SAP HANA cluster nodes that will be created are given below. The instance provisioning described in this guide gives the **minimum** required specifications to run the applications in an evaluation context, and is not generally sufficient to handle productive workloads. Consult the documentation provided by SAP and your cloud provider for best practices when provisioning VM's in a production environment. For example, see [Hardware Requirements for SAP System Hosts](#).

The operating system and configuration used in the deployment must be supported by SAP, SIOS, and your cloud provider. Consult SAP's Product Availability Matrix (PAM) in order to determine which operating systems are supported for various versions of SAP HANA.

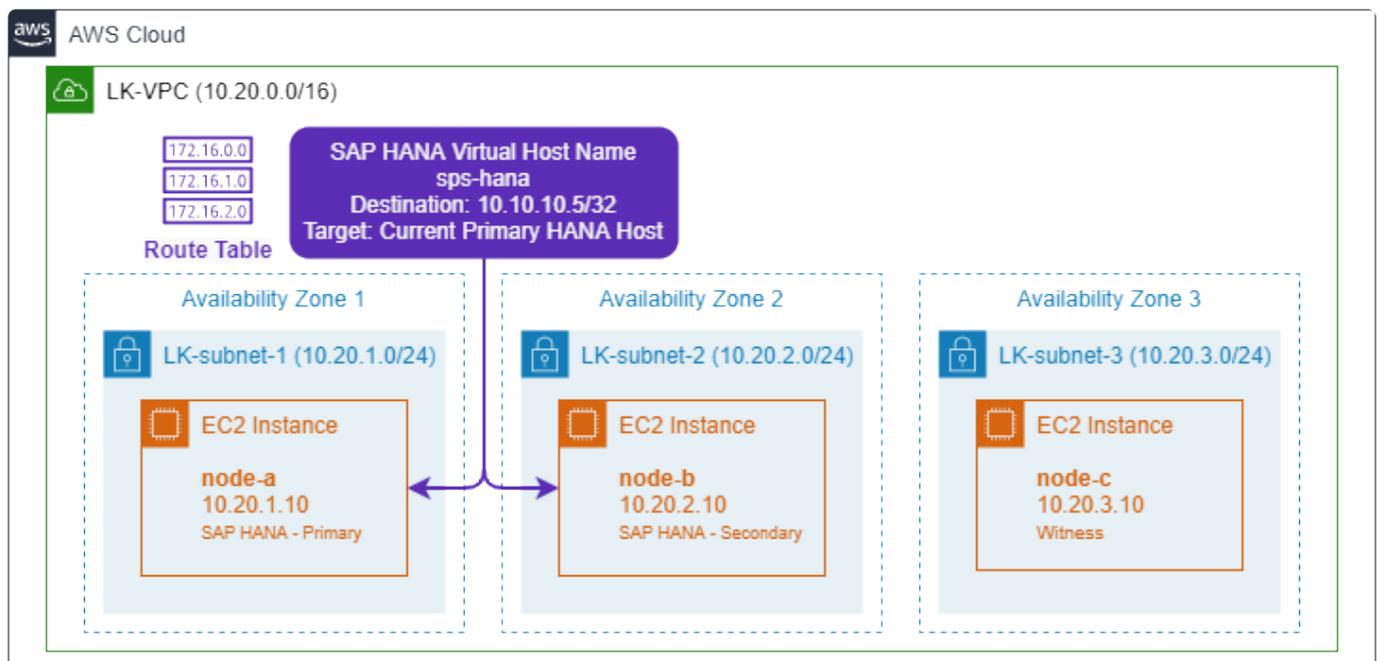
### Cluster Architecture for Evaluation Deployment

- node-a (Primary database host)
  - Private IP: 10.20.1.10
  - vCPU's: 8
  - Memory: 32GB
  - SAP HANA 2.0 SPS04 database instance HDB00 installed under SID 'SPS'
  - SAP HANA System Replication parameters:
    - Site name: SiteA
    - Replication mode: primary
    - Operation mode: primary
  - LifeKeeper for Linux installed with quorum/witness functionality enabled and the following recovery kits installed:
    - SAP HANA Recovery Kit
    - [AWS] Amazon EC2 Recovery Kit
    - [Azure, Google Cloud] Generic Application Recovery Kit for Load Balancer Health Checks (GenLB Recovery Kit)
  
- node-b (Standby database host)
  - Private IP: 10.20.2.10
  - vCPU's: 8
  - Memory: 32GB
  - SAP HANA 2.0 SPS04 database instance HDB00 installed under SID 'SPS'
  - SAP HANA System Replication parameters:
    - Site name: SiteB
    - Replication mode: sync
    - Operation mode: logreplay
  - LifeKeeper for Linux installed with quorum/witness functionality enabled and the following recovery kits installed:
    - SAP HANA Recovery Kit

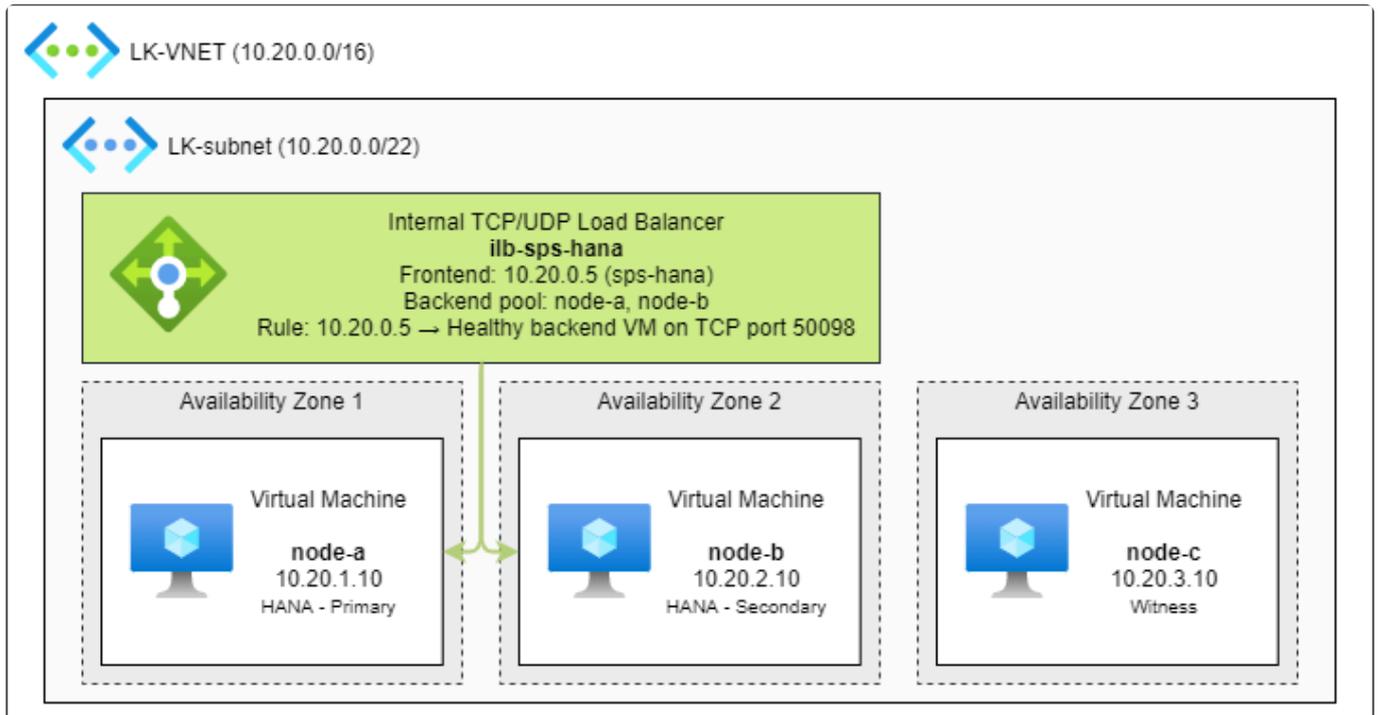
- [AWS] Amazon EC2 Recovery Kit
- [Azure, Google Cloud] Generic Application Recovery Kit for Load Balancer Health Checks (GenLB Recovery Kit)
  
- node-c (Quorum/Witness node)
  - Private IP: 10.20.3.10
  - vCPU's: 1
  - Memory: 2GB
  - LifeKeeper for Linux installed with quorum/witness functionality enabled
  
- SAP HANA virtual hostname: sps-hana

## Cloud-Specific Architecture Diagrams

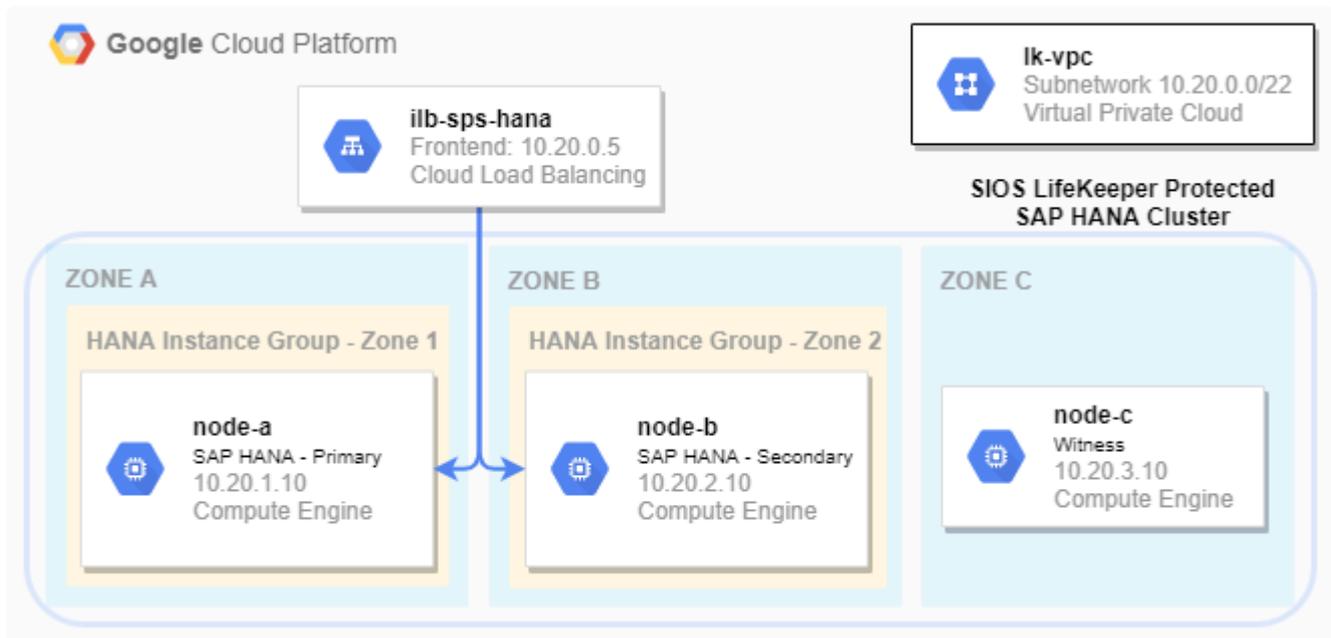
### SIOS Protected SAP HANA Cluster on Amazon Web Services (AWS)



## SIOS Protected SAP HANA Cluster on Microsoft Azure



## SIOS Protected SAP HANA Cluster on Google Cloud



## Prerequisites

In order to follow this guide, the following prerequisite steps must already be complete:

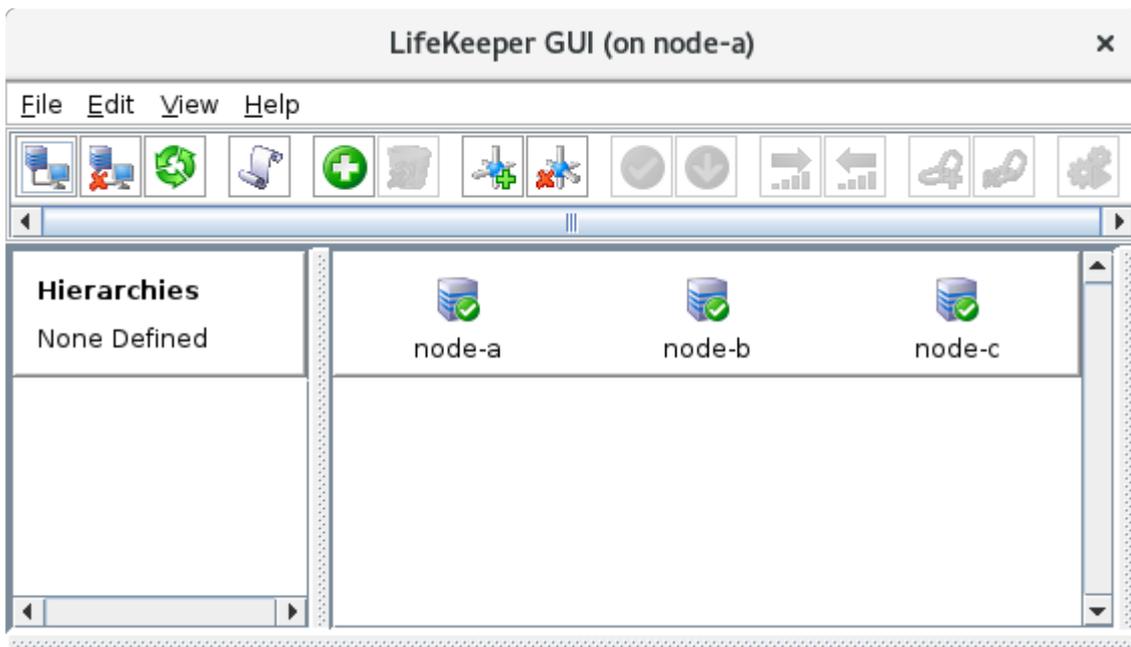
1. Three VM instances (node-a, node-b, and node-c) have been created using the networking conventions described earlier in the guide. Firewall rules are in place to allow inter-node communication as well as SSH connections. See [Configuring Network Components and Creating Instances](#) for details.

2. All three VM instances have been configured to run LifeKeeper for Linux. In particular, SELinux is disabled, the firewall is disabled, the /etc/hosts file on each node contains entries to resolve each node's hostname to its private IP, and the root user has the same password on each node. See [Configure Linux Nodes to Run LifeKeeper for Linux](#) for details.
3. LifeKeeper for Linux has been installed on all three nodes with quorum/witness functionality enabled and all required recovery kits installed. On node-c (the witness node), no additional recovery kits beyond the core LifeKeeper installation are required. All necessary SIOS licenses have been installed on each node. See [Install LifeKeeper for Linux](#) for details.

**Note:** Since the required recovery kits are installed only on node-a and node-b, all steps in this guide that are performed through the LifeKeeper GUI must be performed on either node-a or node-b.

4. LifeKeeper communication paths have been defined between all pairs of cluster nodes. Note that this requires creation of **three** bi-directional communication paths (node-a ↔ node-b, node-a ↔ node-c, node-b ↔ node-c). See [Login and Basic Configuration Tasks](#) for details.

After completing all of these tasks, the LifeKeeper GUI should resemble the following image.



<-- node-c: Adding app/res: scsi device

## 11.2.7.5.6.1. Create SAP HANA Primary Database Virtual IP

---

The virtual host name sps-hana will be configured to resolve to the current primary database host.

The recommended implementation of the virtual IP address varies by cloud platform. Please follow the steps provided in the appropriate section.

- [AWS – Create the SAP HANA Virtual IP](#)
- [Azure – Create the SAP HANA Primary Database Load Balancer](#)
- [Google Cloud – Create the SAP HANA Primary Database Load Balancer](#)

# 11.2.7.5.6.1.1. AWS – Create the SAP HANA Virtual IP

Before creating resources to protect the virtual IP addresses associated to the SAP HANA database, the following steps must be completed on node-a and node-b:

- [Install AWS CLI](#)
- [Assign Permission to Use EC2 Recovery Kit](#)
- [Disable PING Broadcasting](#)
- [AWS – Disable Source/Destination Checking](#)

We will now follow the process described in [Creating an AWS EC2 Resource \(RouteTable Scenario\)](#) to create resource hierarchies for the virtual IP address associated to the SAP HANA database.

Since the IP addresses used with the EC2 RouteTable resource type must be located outside of the VPC CIDR (which for LK-VPC is 10.20.0.0/16), we will use the following virtual host name and IP address:

Instance	Virtual Host Name	IP Address
HDB00	sps-hana	10.10.10.5

## Create the SAP HANA Virtual IP Resource Hierarchy

1. Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend a LifeKeeper IP resource (**ip-sps-hana**) to protect the SAP HANA virtual IP address on node-a and node-b. Notice that the IP resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The

 icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
IP Resource	10.10.10.5
Netmask	255.255.255.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-hana
<b>Pre-Extend Wizard</b>	
Target Server	node-b

Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.10.10.5 
Netmask	255.255.255.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-hana 

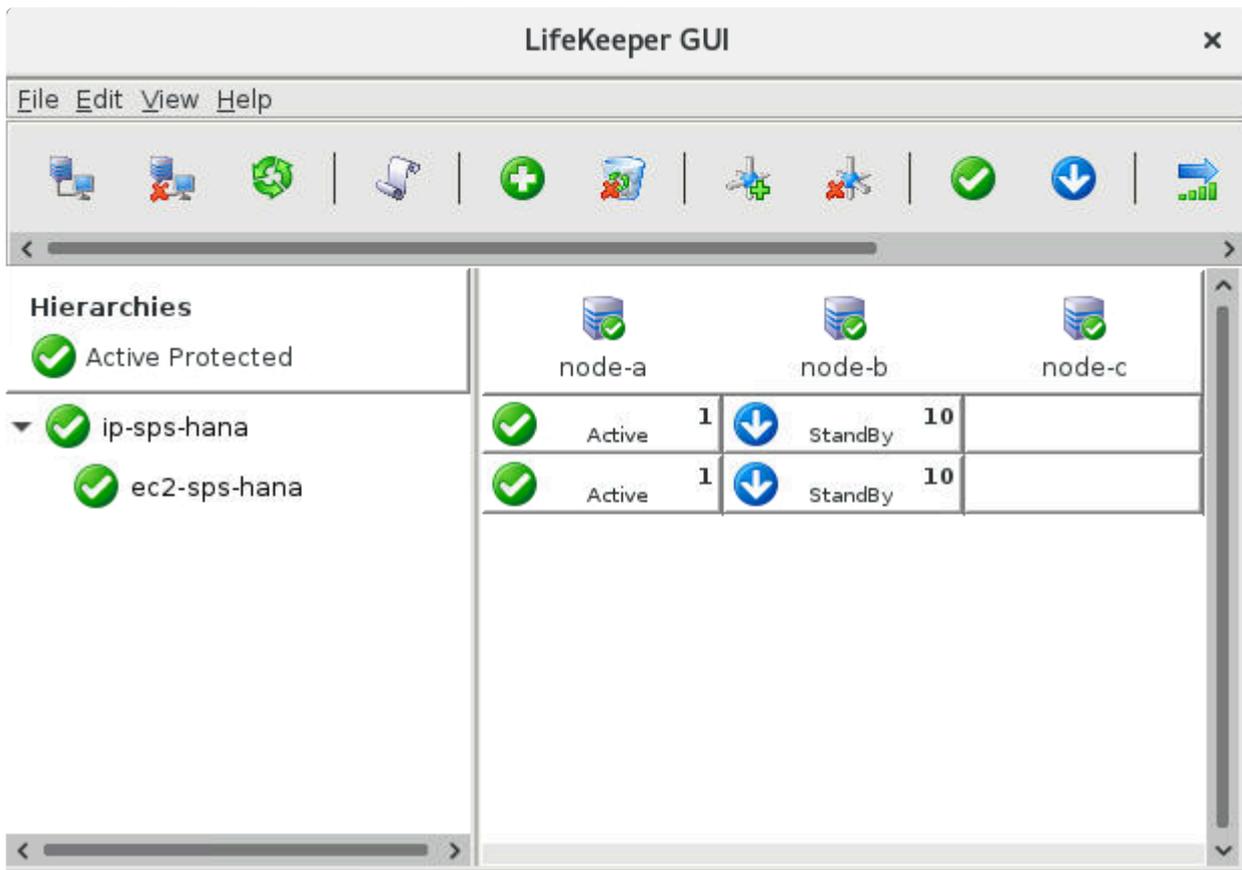
- Following the steps described in [Creating an AWS EC2 Resource \(RouteTable Scenario\)](#), use the following parameters to create and extend a LifeKeeper EC2 resource (**ec2-sps-hana**) to manage the backend manipulation of the route table in AWS when the IP resource is switched over. Notice that the EC2 resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
EC2 Resource type	RouteTable (Backend cluster) 
IP Resource	ip-sps-hana 
EC2 Resource Tag	ec2-sps-hana
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ec2 Resource Hierarchy Wizard</b>	

EC2 Resource Tag	ec2-sps-hana 
------------------	------------------------------------------------------------------------------------------------

 **Note:** If creation or extension of the EC2 resource fails, verify that (i) a route directing traffic from 10.10.10.5/32 to Node-A exists in the route table for LK-VPC in AWS, (ii) source/destination checks have been disabled on Node-A and Node-B, and (iii) that the priority values for the dependent IP resource **ip-sps-hana** are 1 on Node-A and 10 on Node-B, respectively.

Once the IP resource hierarchy has been successfully created, the LifeKeeper GUI should resemble the following image:



## Add Entry to /etc/hosts

Add the following entry to /etc/hosts on node-a and node-b to allow resolution of the virtual host name to its corresponding IP address:

```
10.10.10.5    sps-hana
```

# 11.2.7.5.6.1.2. Azure – Create the SAP HANA Primary Database Load Balancer

In Microsoft Azure, a TCP Internal Load Balancer is used to facilitate failover of the virtual IP associated with the primary HANA database host. The frontend of the load balancer is assigned an IP address from the subnet that it operates in. In this example, we will be using the following IP address for load balancer corresponding to the virtual host name:

Instance	Virtual Host Name	Internal TCP Load Balancer	Frontend IP Address
HDB00	sps-hana	ilb-sps-hana	10.20.0.5

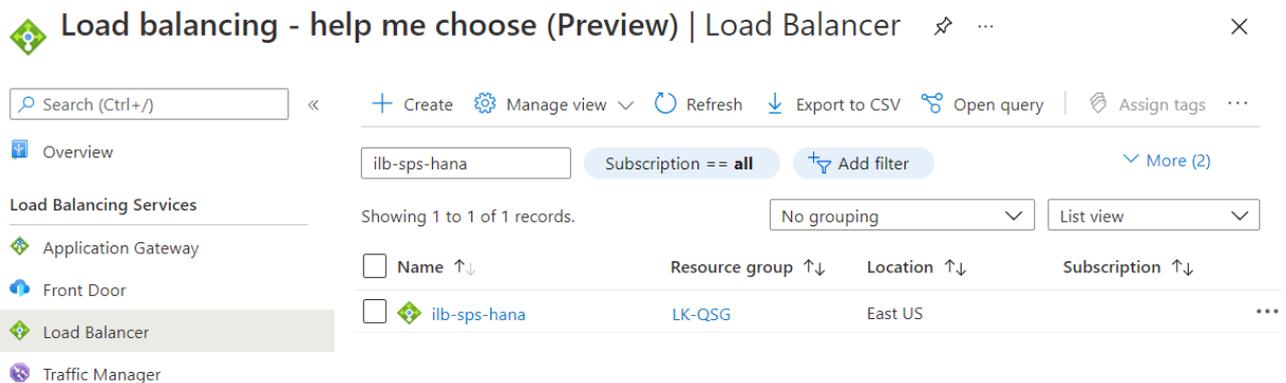
Following the steps provided in [Azure – Using an Internal Load Balancer](#) and [Responding to Load Balancer Health Checks](#), create and configure the following Azure and LifeKeeper resources.

1. Create a Load Balancer in Azure, **ilb-sps-hana**, with the following properties.

Name	ilb-sps-ascs-ers
<b>Create Load Balancer</b>	
Resource Group	LK-QSG
Name	ilb-sps-hana
Region	(same as the Virtual Machines)
Type	Internal
SKU	Standard (select Standard as the workload is distributed across Availability Zones)
Tier	Regional
<b>Frontend IP Configuration</b>	
Name	SPSHANAFrontEnd
Virtual Network	LK-VNET
Subnet	LK-subnet (10.20.0.0/22)
Assignment	Static
IP address	10.20.0.5
Availability zone	Zone-redundant
<b>Backend Pool</b>	
Name	backend-sps-hana
Backend Pool Configuration	NIC

IP Version	IPv4
Virtual machines	node-a, node-b
<b>Health Probe</b>	
Name	probe-sps-hana
Protocol	TCP
Port	50098 (This must agree with the corresponding LifeKeeper GenLB resource created later)
Interval	5 seconds
Unhealthy threshold	2 consecutive failures
<b>Load Balancing Rule</b>	
Name	ilb-rule-sps-hana
IP Version	IPv4
Frontend IP address	SPSHANAFrontEnd (10.20.0.5)
HA Ports	Click (allow forwarding to all ports for simplicity of evaluation deployment)
Backend pool	backend-sps-hana
Health probe	probe-sps-hana (TCP:50098)
Session persistence	None
Idle timeout	4 minutes
TCP reset	Disabled
Floating IP	Enabled

Once created, the **Load Balancing Services** → **Load Balancer** page in the Microsoft Azure Console will show the newly created load balancer.

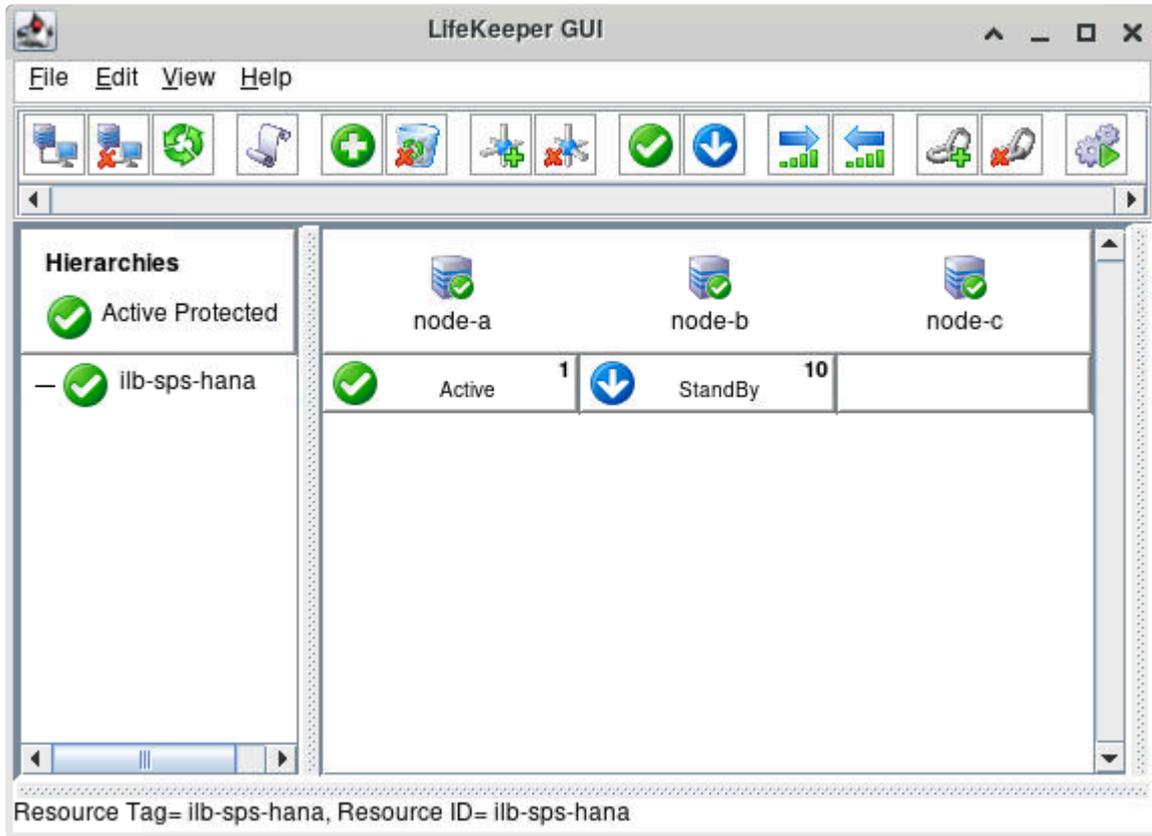


- Following the steps in [Responding to Load Balancer Health Checks](#), install the LifeKeeper Generic Application Recovery Kit for Load Balancer Health Checks (“GenLB Recovery Kit”) and create a LifeKeeper GenLB resource, **ilb-sps-hana**, with the following properties. Notice that **ilb-sps-hana**

is created on node-a and extended to node-b. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Restore Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/restore.pl
Remove Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/remove.pl
QuickCheck Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/quickCheck ( <b>Note:</b> Although the quickCheck script may be optional for some 'Generic Application' resource types, it is required for GenLB resources.)
Local Recovery Script	None (Empty)
Application Info	50098
Bring Resource In Service	Yes
Resource Tag	ilb-sps-hana
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Resource Tag	ilb-sps-hana 
Application Info	50098 

The resource will appear in the LifeKeeper GUI resource pane once it has been created and extended successfully.

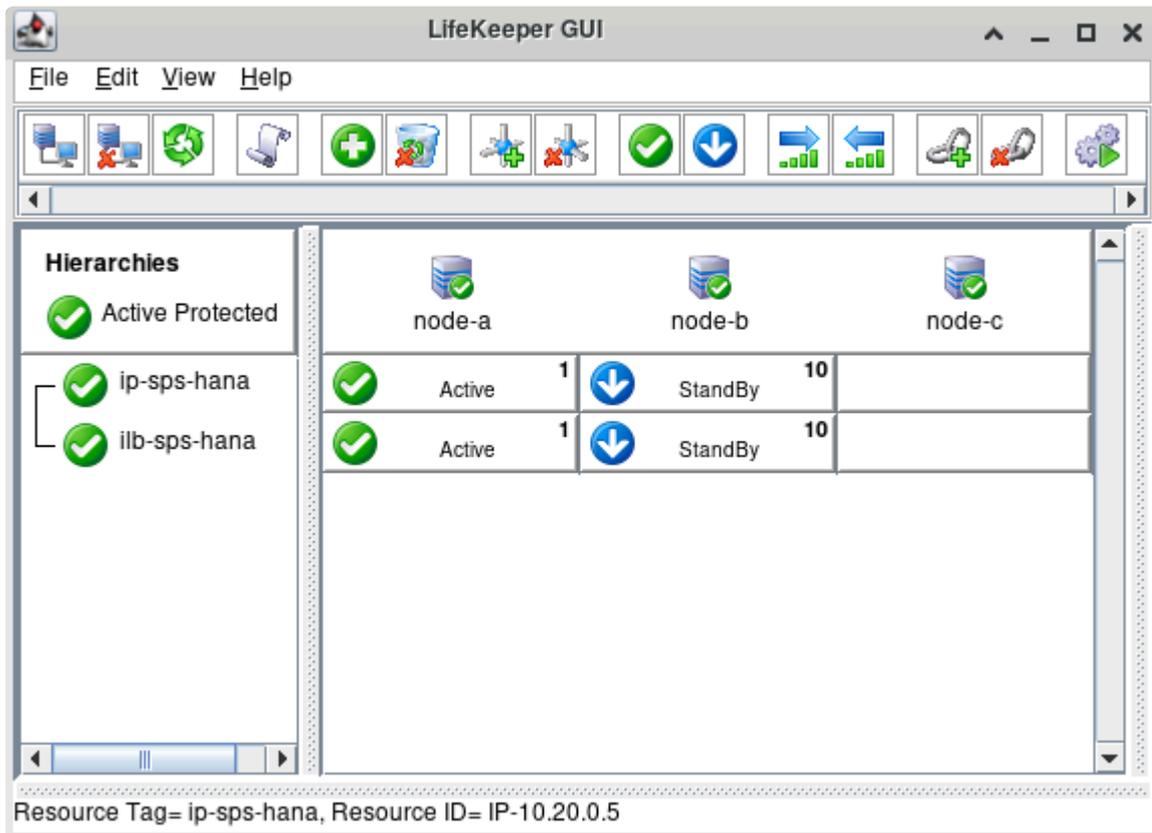


- Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend a LifeKeeper IP resource (**ip-sps-hana**) to protect the HANA database virtual IP address on node-a and node-b. Notice that the IP resource is being created on node-a and extended to node-b. Also note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
IP Resource	10.20.0.5
Netmask	255.255.252.0 
Network Interface	eth0 
IP Resource Tag	ip-sps-hana
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 

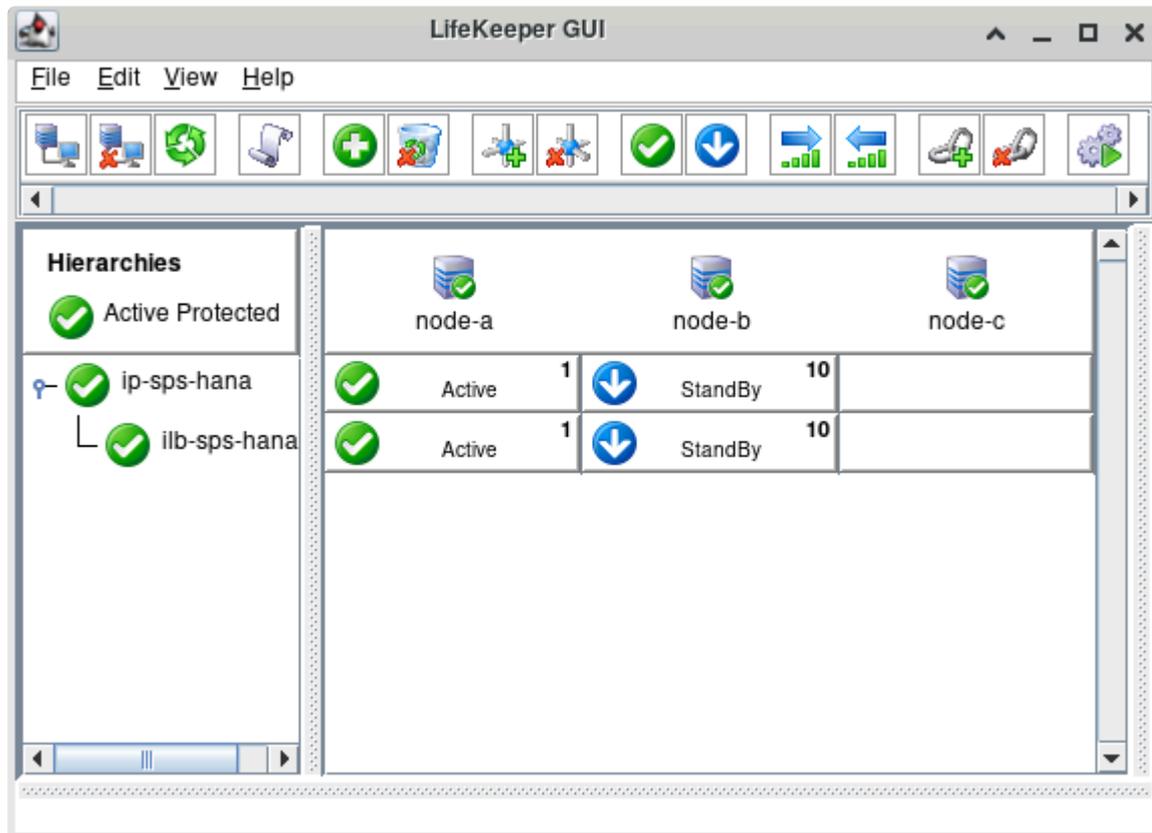
Template Priority	1 ✓
Target Priority	10 ✓
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.20.0.5 ✓
Netmask	255.255.252.0 ✓
Network Interface	eth0 ✓
IP Resource Tag	ip-sps-hana ✓

Once the IP resources have been created successfully, the LifeKeeper GUI should resemble the following image.



- Right-click the **ip-sps-hana** resource on node-a and click “Create Dependency...” Specify **ilb-sps-hana** as the Child Resource Tag and click **Create Dependency**.

Once the dependency has been created successfully, the LifeKeeper GUI should resemble the following image.



5. Add the following entry to `/etc/hosts` on node-a and node-b to allow resolution of the virtual host name to the frontend IP address of the corresponding load balancer:

```
10.20.0.5 sps-hana
```

6. Test switchover and failover of the GenLB resource hierarchy as described in the **Test GenLB Resource Switchover and Failover** section of [Responding to Load Balancer Health Checks](#). Correct any issues found or tune the parameters of the load balancer health check as required to achieve successful operation.

# 11.2.7.5.6.1.3. Google Cloud – Create the SAP HANA Primary Database Load Balancer

In Google Cloud, a TCP Internal Load Balancer is used to facilitate failover of the SAP HANA virtual IP. The frontend of the load balancer is assigned an ephemeral IP from the subnet that it operates in. In this example, we will be using the following IP address for load balancer corresponding to the sps-hana virtual hostname:

Instance	Virtual Host Name	Internal TCP Load Balancer	Frontend IP Address
HDB00	sps-hana	ilb-sps-hana	10.20.0.5

Following the steps provided in [Google Cloud – Using an Internal Load Balancer](#) and [Responding to Load Balancer Health Checks](#), create and configure the following Google Cloud and LifeKeeper resources.

1. Create two unmanaged instance groups: **ig-sps-hana-zone1**, containing node-a, and **ig-sps-hana-zone2**, containing node-b.

Instance groups

Instance groups are collections of VM instances that use load balancing and automated services, like autoscaling and autohealing. [Learn more](#)

Name	Instances	Template	Group type	Creation time	Recommendation	Autoscaling	Zone	In Use By
ig-sps-hana-zone1	1	-	Unmanaged	Mar 1, 2021, 9:36:49 AM UTC-05:00			us-east1-b	ilb-sps-hana
ig-sps-hana-zone2	1	-	Unmanaged	Mar 1, 2021, 9:43:34 AM UTC-05:00			us-east1-c	ilb-sps-hana

2. Create a TCP Internal Load Balancer, **ilb-sps-hana**, with the following properties. The icon indicates that the default option is chosen.

Name	ilb-sps-hana
<b>Backend configuration</b>	
Region	<Deployment region> (e.g., us-east1)
Network	lk-vpc
Backends	ig-sps-hana-zone1 ig-sps-hana-zone2
<b>Health Check</b>	
Name	hc-sps-hana
Protocol	TCP

Port	50098 (This must agree with the corresponding LifeKeeper GenLB resource created later)
Proxy protocol	NONE
Request	Leave empty
Response	Leave empty
Check interval	5 seconds
Timeout	5 seconds
Healthy threshold	2 consecutive successes
Unhealthy threshold	2 consecutive failures
Connection draining timeout	10 seconds
<b>Frontend configuration</b>	
Name	fe-sps-hana
Subnetwork	lk-subnet
Purpose	Non-shared
IP address	Ephemeral (Custom)
Custom ephemeral IP address	10.20.0.5
Ports	All
Global access	Disable
Service label	Leave empty

Once created, the **Network services** → **Load balancing** page in the Google Cloud Console will show the load balancer.

[Load balancers](#)   Backends   Frontends

Filter by name or protocol

<input type="checkbox"/> Name	Protocol	Region	Backends
<input type="checkbox"/> ilb-sps-hana	TCP (Internal)	us-east1	1 regional backend service (2 instance groups)

- Follow the steps described in the **Disable IP Forwarding** section of [Google Cloud – Using an Internal Load Balancer](#) to allow node-a and node-b to communicate through the frontend IP addresses of the load balancer. Reboot node-a and node-b for the changes to take effect.

Note that this step is not strictly required in order to connect to the primary database through the load balancer frontend IP address, but it will make it possible to test the database connection from within the cluster, as described in the [Test Switchover and Failover](#) section.

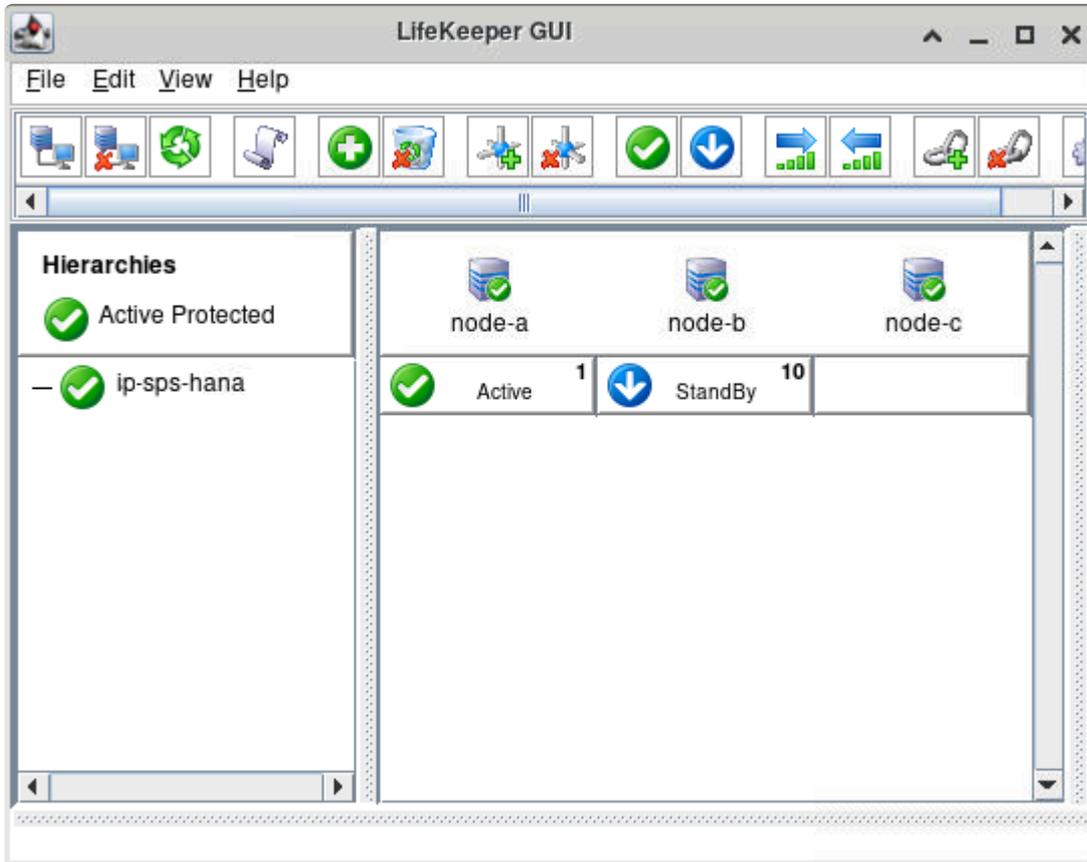
- Following the steps described in [Creating an IP Resource](#), use the following parameters to create and extend a LifeKeeper IP resources (**ip-sps-hana**), with net mask 255.255.255.255, to protect the SAP HANA virtual IP address on node-a and node-b. Notice that **ip-sps-hana** is created on node-a and extended to node-b. Also note that this IP resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

 The use of 255.255.255.255 as the value of the Netmask parameter (both during create and extend) is important. Be careful not to select “Accept Defaults” during the extend process as this will typically cause 255.255.255.0 to be incorrectly used as the network mask on the standby node, which may lead to unexpected networking issues.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
IP Resource	10.20.0.5
Netmask	255.255.255.255
Network Interface	eth0 
IP Resource Tag	ip-sps-hana
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend comm/ip Resource Hierarchy Wizard</b>	
IP Resource	10.20.0.5 
Netmask	255.255.255.255

Network Interface	eth0 
IP Resource Tag	ip-sps-hana 

Once the IP resource has been created successfully, the LifeKeeper GUI should resemble the following image.

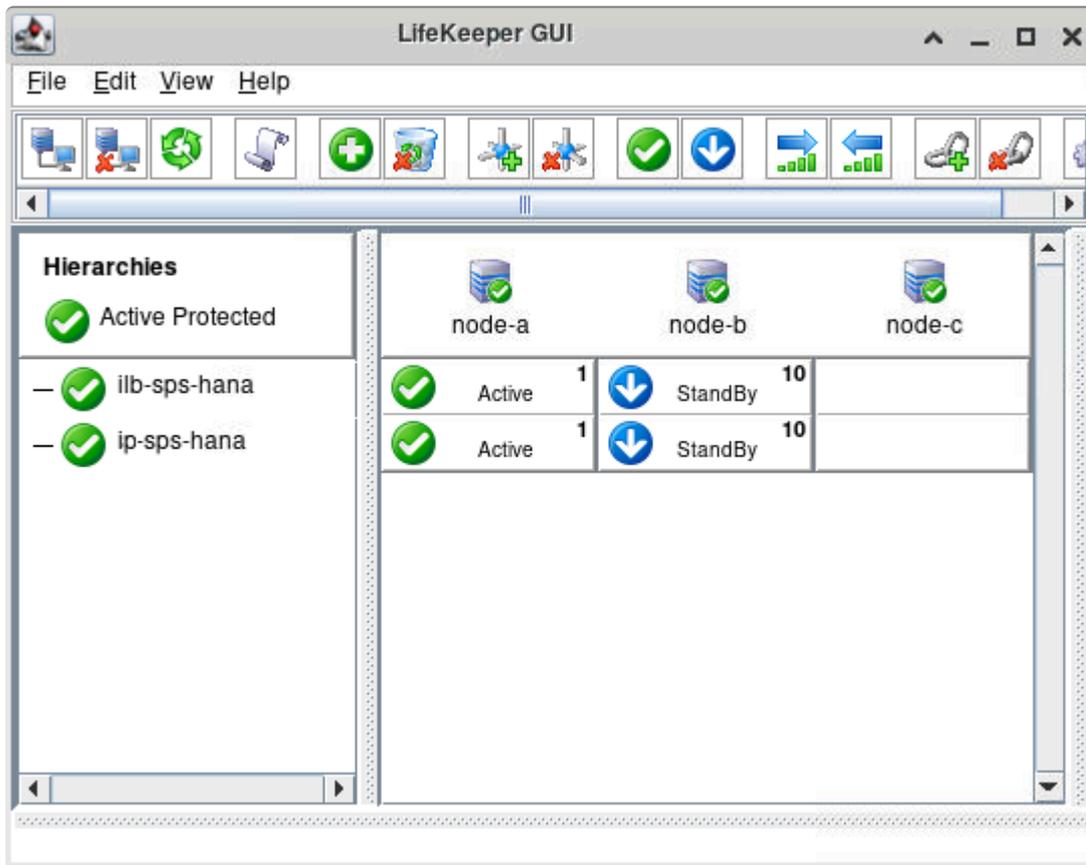


- Following the steps in [Responding to Load Balancer Health Checks](#), install the LifeKeeper Generic Application Recovery Kit for Load Balancer Health Checks (“GenLB Recovery Kit”) and create a LifeKeeper GenLB resource, **ilb-sps-hana**, with the following properties. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
Restore Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/restore.pl
Remove Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/remove.pl
QuickCheck Script	/opt/LifeKeeper/SIOS_Hotfixes/Gen-LB-PL-7172/quickCheck ( <b>Note:</b> Although the quickCheck script may be optional for some ‘Generic

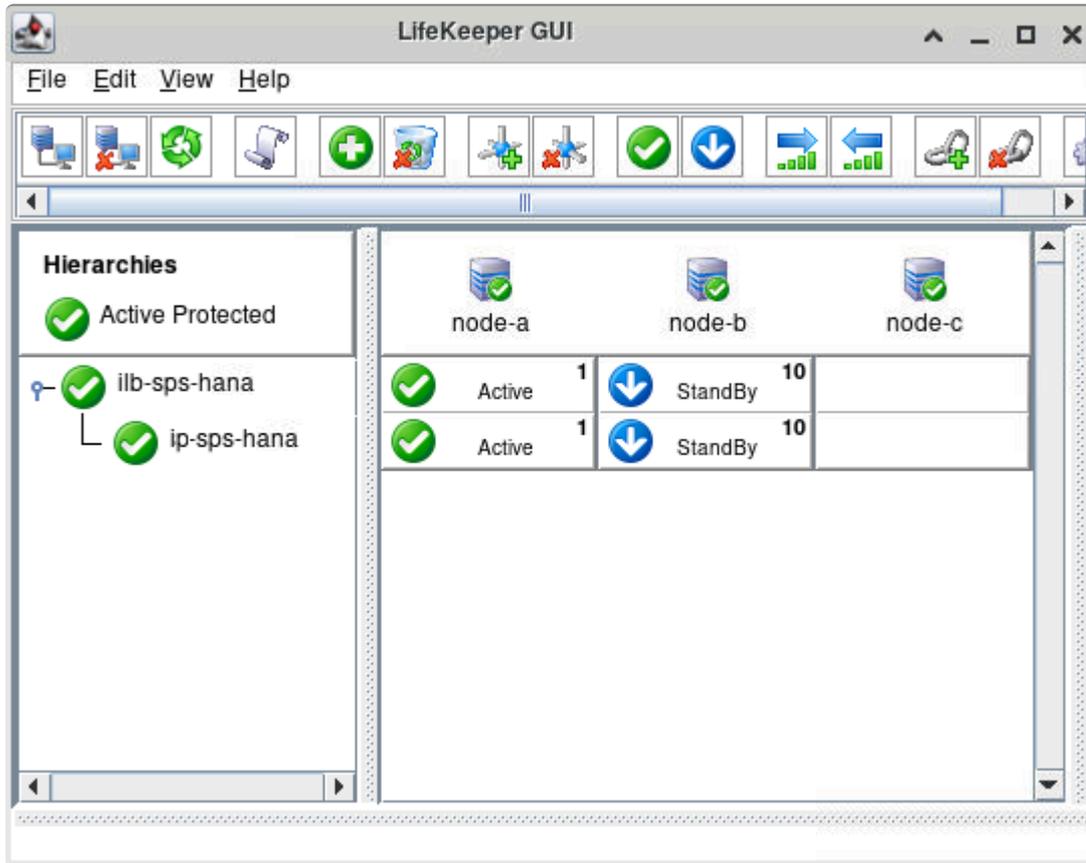
	Application' resource types, it is required for GenLB resources.)
Local Recovery Script	None (Empty)
Application Info	50098
Bring Resource In Service	Yes
Resource Tag	ilb-sps-hana
<b>Pre-Extend Wizard</b>	
Target Server	node-b
Switchback Type	intelligent 
Template Priority	1 
Target Priority	10 
<b>Extend gen/app Resource Hierarchy Wizard</b>	
Resource Tag	ilb-sps-hana 
Application Info	50098 

The resource will appear in the LifeKeeper GUI resource pane once it has been created and extended successfully.



- Right-click the **ilb-sps-hana** resource on node-a and click "Create Dependency..." Specify **ip-sps-hana** as the Child Resource Tag and click **Create Dependency**.

Once the dependency has been created, the LifeKeeper GUI should resemble the following image.



7. Add the following entry to /etc/hosts on node-a and node-b to allow resolution of the virtual hostname to the frontend IP address of the corresponding load balancer:

```
10.20.0.5 sps-hana
```

8. Test switchover and failover of the GenLB resources as described in the **Test GenLB Resource Switchover and Failover** section of [Responding to Load Balancer Health Checks](#). Correct any issues found or tune the parameters of the load balancer health checks as required to achieve successful operation.

## 11.2.7.5.6.2. Attach Disks for SAP HANA File Systems

1. Create and attach three additional disks to node-a and node-b to support the installation of SAP HANA. The device names used in this example may vary depending on your environment (e.g., /dev/xvdb instead of /dev/sdb). Note that these disk sizes are used for evaluation purposes only. Consult the relevant documentation from SAP, your cloud provider, and any third-party storage provider when provisioning resources in a production environment.

Device Name	Minimum Size	Mount Point
/dev/sdb	20GB	/hana/log
/dev/sdc	30GB	/hana/data
/dev/sdd	60GB	/hana/shared

 **Note:** Other configurations may be used when attaching storage to support the SAP HANA installation. For example, it may be convenient to use a single physical volume partitioned with multiple logical partitions via LVM.

2. Execute the following commands on both node-a and node-b to create a /dev/sdb1 partition with an xfs file system.

```
[root@node-a ~]# parted /dev/sdb --script mklabel gpt mkpart xfspart xfs 0% 10
0%

[root@node-a ~]# mkfs.xfs /dev/sdb1
meta-data=/dev/sdb1      isize=512    agcount=4, agsize=1310592 blks
           =             sectsz=4096   attr=2,   projid32bit=1
           =             crc=1        finobt=1, sparse=1, rmapbt=0
           =             reflink=1
data      =             bsize=4096   blocks=5242368, imaxpct=25
           =             sunit=0      swidth=0 blks
naming    =version 2     bsize=4096   ascii-ci=0,  ftype=1
log       =internal log bsize=4096   blocks=2560,  version=2
           =             sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none         extsz=4096   blocks=0,    rtextents=0

[root@node-a ~]# partprobe /dev/sdb1
```

Repeat these commands for /dev/sdc and /dev/sdd on both node-a and node-b.

3. Execute the following command to create the mount points for the SAP HANA file systems:

```
[root@node-a ~]# mkdir -p /hana/{log,data,shared}
```

4. Use the `blkid` command to view the UUID's for each partition.

```
[root@node-a ~]# blkid | sort
/dev/sdb1: UUID="f058a8de-4e17-4c1d-97a3-a2672aff58dd" TYPE="xfs" PARTLABEL="xfspart" PARTUUID="f780893d-bd61-4beb-89df-fb23ea04f16e"
/dev/sdc1: UUID="c0ab7218-ec24-4a6f-b883-37e505ec806c" TYPE="xfs" PARTLABEL="xfspart" PARTUUID="34c47333-7174-468c-b8f0-a3491cd0fe0c"
/dev/sdd1: UUID="057f4f63-8fb3-486b-82b5-49fbbd4914cc" TYPE="xfs" PARTLABEL="xfspart" PARTUUID="0c9cc7d3-9622-4b3f-8949-e744a02af466"
```

5. Referencing the UUID's found in step 3, add the following entries to `/etc/fstab` on `node-a` and `node-b` to allow the file systems to be mounted at boot. Replace the sample UUID's given in this example with the UUID's of the appropriate partitions on your systems.

```
UUID=f058a8de-4e17-4c1d-97a3-a2672aff58dd    /hana/log    xfs    default
s        0 0
UUID=c0ab7218-ec24-4a6f-b883-37e505ec806c    /hana/data    xfs    defau
lts        0 0
UUID=057f4f63-8fb3-486b-82b5-49fbbd4914cc    /hana/shared    xfs    defaul
ts        0 0
```

6. Execute the following commands on `node-a` and `node-b` to mount the SAP HANA file systems on both nodes:

```
mount /hana/log
mount /hana/data
mount /hana/shared
```

The file systems are now prepared for SAP HANA installation.

# 11.2.7.5.6.3. Install SAP HANA and Configure System Replication

## Install SAP HANA on node-a

1. Create a /sap-install directory on node-a which will contain the SAP HANA installation files.

```
mkdir /sap-install
```

 **Note:** If the root file system for your chosen instance type does not have sufficient disk space to store the required SAP HANA installation files once they are extracted, you may attach a temporary disk to each instance and mount it at /sap-install in order to complete installation of the SAP HANA software.

2. Download the SAP HANA installation files to /sap-install on node-a.
3. Extract the SAP HANA installation files. In this example we are installing SAP HANA 2.0 SPS04 Platform Edition.

```
[root@node-a ~]# unzip /sap-install/hana-2-platform-edition-sp4.zip -d /sap-install/
```

4. Run the hdblcm executable to begin the installation.

```
[root@node-a ~]# /sap-install/DATA_UNITS/HDB_LCM_LINUX_X86_64/hdblcm
```

5. Enter the following parameters to complete a basic installation of the SAP HANA Database Server.

Select other optional components during installation as required for your deployment. The  icon indicates that the default option is chosen.

Field	Value
Enter selected action index	1 (install)
Components to install	2 (HDB Server)
Enter Installation Path	/hana/shared 
Enter Local Host Name	node-a 
Do you want to add hosts to the system? (y/n)	n 
Enter SAP HANA System ID	SPS

Enter Instance Number	00 ✓
Enter Local Host Worker Group	default ✓
Select Usage / Enter Index	4 (custom) ✓
Enter Location of Data Volumes	/hana/data/SPS ✓
Enter Location of Log Volumes	/hana/log/SPS ✓
Restrict maximum memory allocation?	n ✓
Enter Certificate Host Name for Host 'node-a'	node-a ✓
Enter System Administrator (spsadm) Password	<spsadm User Password>
Confirm System Administrator (spsadm) Password	<spsadm User Password>
Enter System Administrator Home Directory	/usr/sap/SPS/home ✓
Enter System Administrator Login Shell	/bin/sh ✓
Enter System Administrator User ID	1001 (or default value) ✓
Enter ID of User Group (sapsys)	79 (or default value) ✓
Enter System Database User (SYSTEM) Password	<SYSTEM DB User Password>
Confirm System Database User (SYSTEM) Password	<SYSTEM DB User Password>
Restart system after machine reboot?	n ✓
Do you want to continue?	y ✓

**!** **Note:** If the installation of SAP HANA fails, verify that all required dependent packages (e.g., libatomic / libatomic1, compat-sap-c++, libtool-ltdl / libltdl7, etc.) are installed and reattempt installation.

6. Once the installation completes, execute the following command on node-a to verify that the HDB00 instance is running successfully:

```
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
04.03.2021 23:42:13
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2021 03 04 23:38:10, 0:04:03, 23804
```

```

hdbcompileserver, HDB Compileserver, GREEN, Running, 2021 03 04 23:38:41, 0:0
3:32, 24044
hdbnameserver, HDB Nameserver, GREEN, Running, 2021 03 04 23:38:11, 0:04:02, 2
3822
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2021 03 04 23:38:41, 0:03:3
2, 24047
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2021 03 04 23:39:24, 0:0
2:49, 24454
hdbindexserver, HDB Indexserver-SPS, GREEN, Running, 2021 03 04 23:38:41, 0:0
3:32, 24093
hdbxsengine, HDB XSEngine-SPS, GREEN, Running, 2021 03 04 23:38:41, 0:03:32, 2
4096

```

## Install SAP HANA on node-b

1. Repeat the installation steps from the previous section on node-b, replacing “node-a” by “node-b” wherever it appears.



**Important:** The parameters (including the user ID of the spsadm user and the group ID of the sapsys group) must be exactly the same in both installations.

2. Once the installation completes, execute the following command on node-b to verify that the HDB00 instance is running successfully.

```

[root@node-b ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
04.03.2021 23:44:09
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2021 03 04 23:38:16, 0:05:53, 10785
hdbcompileserver, HDB Compileserver, GREEN, Running, 2021 03 04 23:38:46, 0:0
5:23, 11111
hdbnameserver, HDB Nameserver, GREEN, Running, 2021 03 04 23:38:16, 0:05:53, 1
0804
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2021 03 04 23:38:46, 0:05:2
3, 11114
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2021 03 04 23:39:29, 0:0
4:40, 11438
hdbindexserver, HDB Indexserver-SPS, GREEN, Running, 2021 03 04 23:38:47, 0:0
5:22, 11159
hdbxsengine, HDB XSEngine-SPS, GREEN, Running, 2021 03 04 23:38:47, 0:05:22, 1
1162

```

## Back up the SYSTEMDB and SPS Databases on Both

## Nodes

Execute the following commands on both node-a and node-b to create backups of the SYSTEMDB database and the SPS tenant database, replacing <SYSTEM User Password> with the password for the SYSTEM database user:

```
# su - spsadm -c "hdbsql -i 00 -u SYSTEM -p <SYSTEM User Password> -d SystemD
B \"BACKUP DATA USING FILE ('/hana/shared/SPS/HDB00')\"
0 rows affected (overall time 7390.598 msec; server time 7388.081 msec)

# su - spsadm -c "hdbsql -i 00 -u SYSTEM -p <SYSTEM User Password> -d SPS \"BA
CKUP DATA USING FILE ('/hana/shared/SPS/HDB00')\"
0 rows affected (overall time 5104.380 msec; server time 5102.842 msec)
```

## Copy the PKI ssfs KEY and DAT Files

1. Execute the following command to stop the HDB00 instance on node-b.

```
[root@node-b ~]# su - spsadm -c "sapcontrol -nr 00 -function StopSystem HDB"
04.03.2021 23:55:10
StopSystem
OK
```

2. Verify that the HDB00 instance has been successfully stopped on node-b.

```
[root@node-b ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
04.03.2021 23:55:51
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GRAY, Stopped, , , 10785
```

3. Execute the following command to copy the /usr/sap/SPS/SYS/global/security/rsecssfs/data/SSFS\_SPS.DAT and /usr/sap/SPS/SYS/global/security/rsecssfs/key/SSFS\_SPS.KEY files from node-a to node-b.

```
[root@node-a ~]# scp -r /usr/sap/SPS/SYS/global/security/rsecssfs/ root@node-
b:/usr/sap/SPS/SYS/global/security/
SSFS_SPS.DAT          100% 2960      2.8MB/s   00:00
SSFS_SPS.KEY         100%  187      190.9KB/s  00:00
```

## Configure SAP HANA System Replication

1. Execute the following command on node-a to enable system replication using site name SiteA.

```
[root@node-a ~]# su - spsadm -c "hdbnsutil -sr_enable --name=SiteA"
nameserver is active, proceeding ...
successfully enabled system as system replication source site
done.
```

2. Execute the following command to register node-b as a secondary system replication site using site name SiteB.

```
[root@node-b ~]# su - spsadm -c "hdbnsutil -sr_register --remoteHost=node-a --
remoteInstance=00 --replicationMode=sync --operationMode=logreplay --name=Site
B"
adding site ...
nameserver node-a:30001 not responding.
collecting information ...
updating local ini files ...
done.
```

3. Execute the following command to start the HDB00 instance on node-b.

```
[root@node-b ~]# su - spsadm -c "sapcontrol -nr 00 -function StartSystem HDB"
05.03.2021 00:11:51
StartSystem
OK
```

4. Execute the following commands to verify that the HDB00 instance is running successfully on node-b and that the replication and operation modes are primary/primary and sync/logreplay on node-a and node-b, respectively.

```
[root@node-b ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
05.03.2021 00:13:36
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2021 03 05 00:11:52, 0:01:44, 15816
hdbcompileserver, HDB Compileserver, GREEN, Running, 2021 03 05 00:11:57, 0:0
1:39, 15940
hdbindexserver, HDB Indexserver-SPS, GREEN, Running, 2021 03 05 00:11:57, 0:0
1:39, 15997
hdbnameserver, HDB Nameserver, GREEN, Running, 2021 03 05 00:11:52, 0:01:44, 1
5835
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2021 03 05 00:11:57, 0:01:3
9, 15943
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2021 03 05 00:12:02, 0:0
1:34, 16217
hdbxsengine, HDB XSEngine-SPS, GREEN, Running, 2021 03 05 00:11:57, 0:01:39, 1
```

```

6000

[root@node-b ~]# su - spsadm -c "hdbnsutil -sr_state"
System Replication State
~~~~~

online: true

mode: sync
operation mode: logreplay
site id: 2
site name: SiteB

is source system: false
is secondary/consumer system: true
has secondaries/consumers attached: false
is a takeover active: false
active primary site: 1

primary masters: node-a

Host Mappings:
~~~~~

node-b -> [SiteB] node-b
node-b -> [SiteA] node-a

Site Mappings:
~~~~~
SiteA (primary/primary)
  |--SiteB (sync/logreplay)

Tier of SiteA: 1
Tier of SiteB: 2

Replication mode of SiteA: primary
Replication mode of SiteB: sync

Operation mode of SiteA: primary
Operation mode of SiteB: logreplay

Mapping: SiteA -> SiteB
done.

```

Now that the HDB00 instance is running on node-a and node-b and both nodes are registered in SAP HANA System Replication, we are ready to create the SAP HANA resource in LifeKeeper.

## 11.2.7.5.6.4. Create LifeKeeper SAP HANA Resource

---

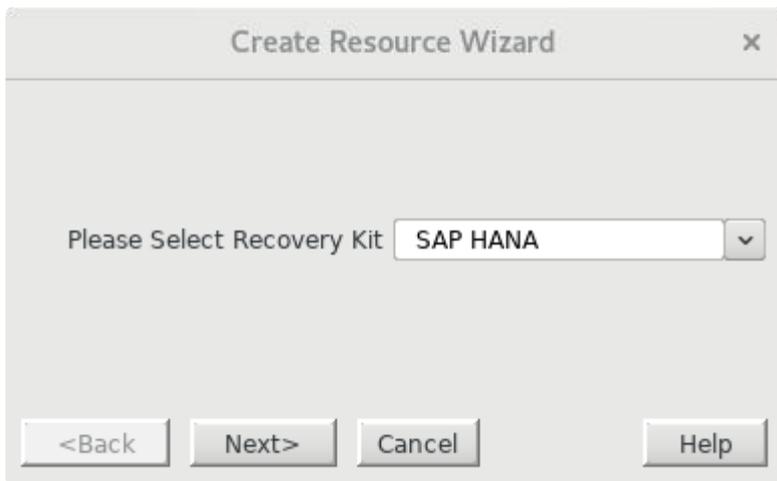
The steps required to create the LifeKeeper resources to protect the HDB00 instance vary by cloud platform. Please follow the steps provided in the section corresponding to your cloud platform:

- [AWS/Azure – Create LifeKeeper SAP HANA Resource](#)
- [Google Cloud – Create LifeKeeper SAP HANA Resource](#)

# 11.2.7.5.6.4.1. AWS/Azure – Create LifeKeeper SAP HANA Resource

\* **Note:** This section applies only to deployments on AWS or Microsoft Azure. For deployments on Google Cloud, see [Google Cloud – Create LifeKeeper SAP HANA Resource](#).

1. In the LifeKeeper GUI, click the  icon to open the **Create Resource Wizard**. Select the “SAP HANA” Recovery Kit.

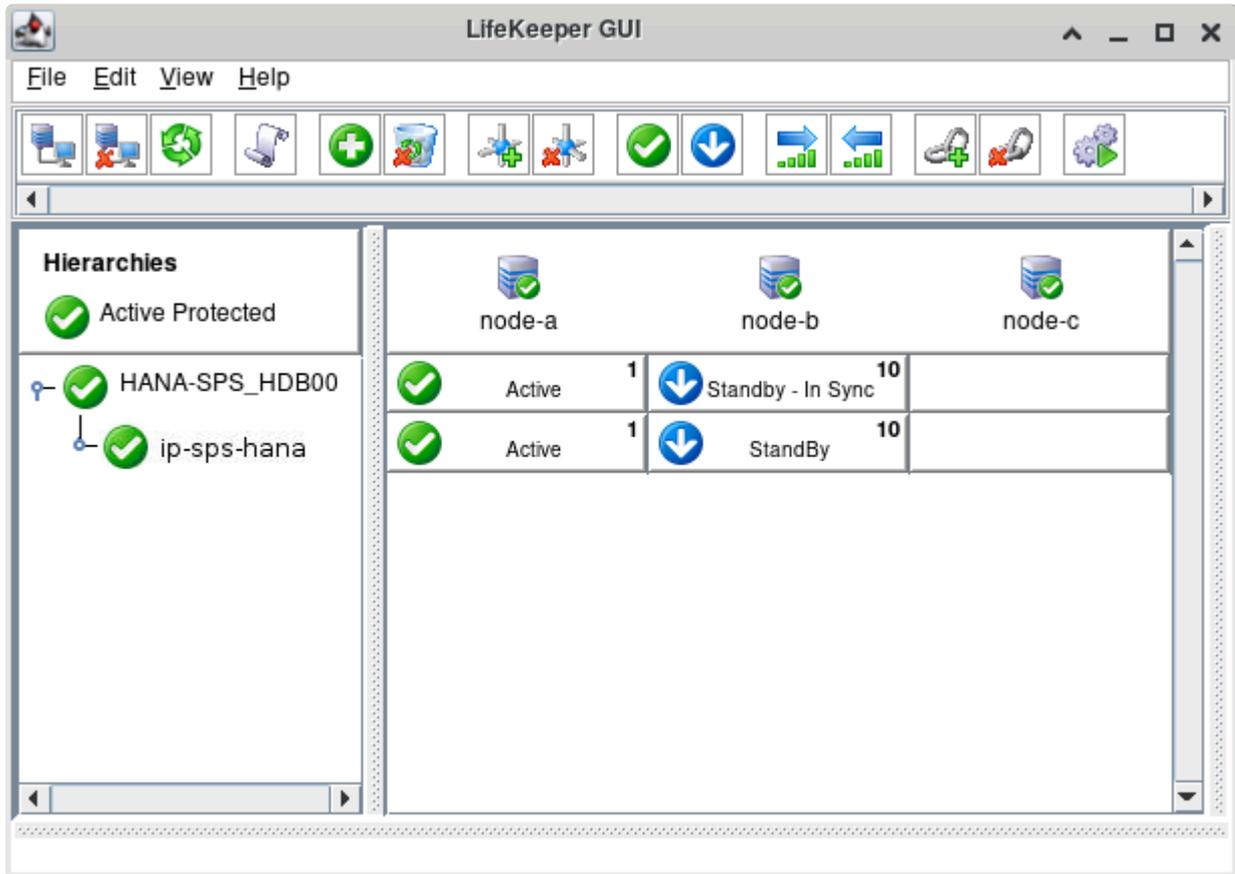


2. Enter the following values to create and extend a LifeKeeper resource (**HANA-SPS\_HDB00**) to protect the HDB00 database instance on node-a and node-b. Note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
SAP HANA SID	SPS 
SAP HANA Instance for SPS	HDB00 
IP child resource	ip-sps-hana
Enable/Disable Local Recovery	Enable 
HANA Tag	HANA-SPS_HDB00 

Pre-Extend Wizard	
Target Server	node-b
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
Extend database/hana Resource Hierarchy Wizard	
Root Tag	HANA-SPS_HDB00 ✓
Enable/Disable Local Recovery	Enable ✓

Once the **HANA-SPS\_HDB00** resource has been successfully created and extended to node-b, the LifeKeeper resource panel should look similar to the following.

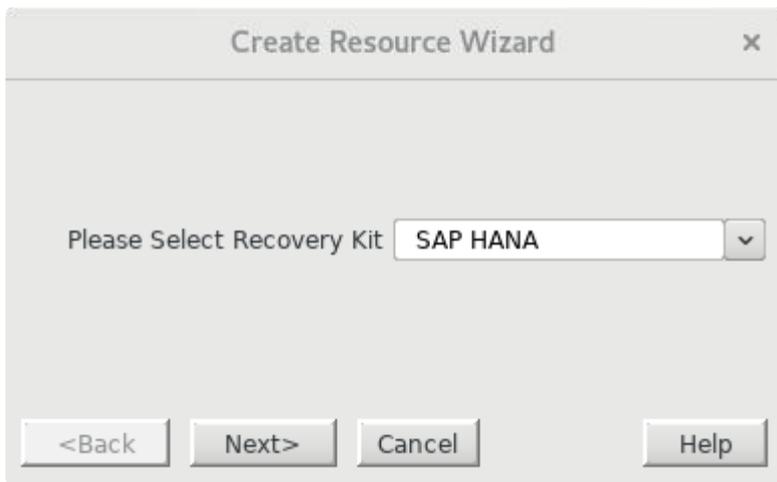


SAP HANA database instance HDB00 has now been successfully protected by LifeKeeper.

# 11.2.7.5.6.4.2. Google Cloud – Create LifeKeeper SAP HANA Resource

 **Note:** This section applies only to deployments on Google Cloud. For deployments on AWS or Microsoft Azure, see [AWS/Azure – Create LifeKeeper SAP HANA Resource](#).

1. In the LifeKeeper GUI, click the  icon to open the **Create Resource Wizard**. Select the “SAP HANA” Recovery Kit.

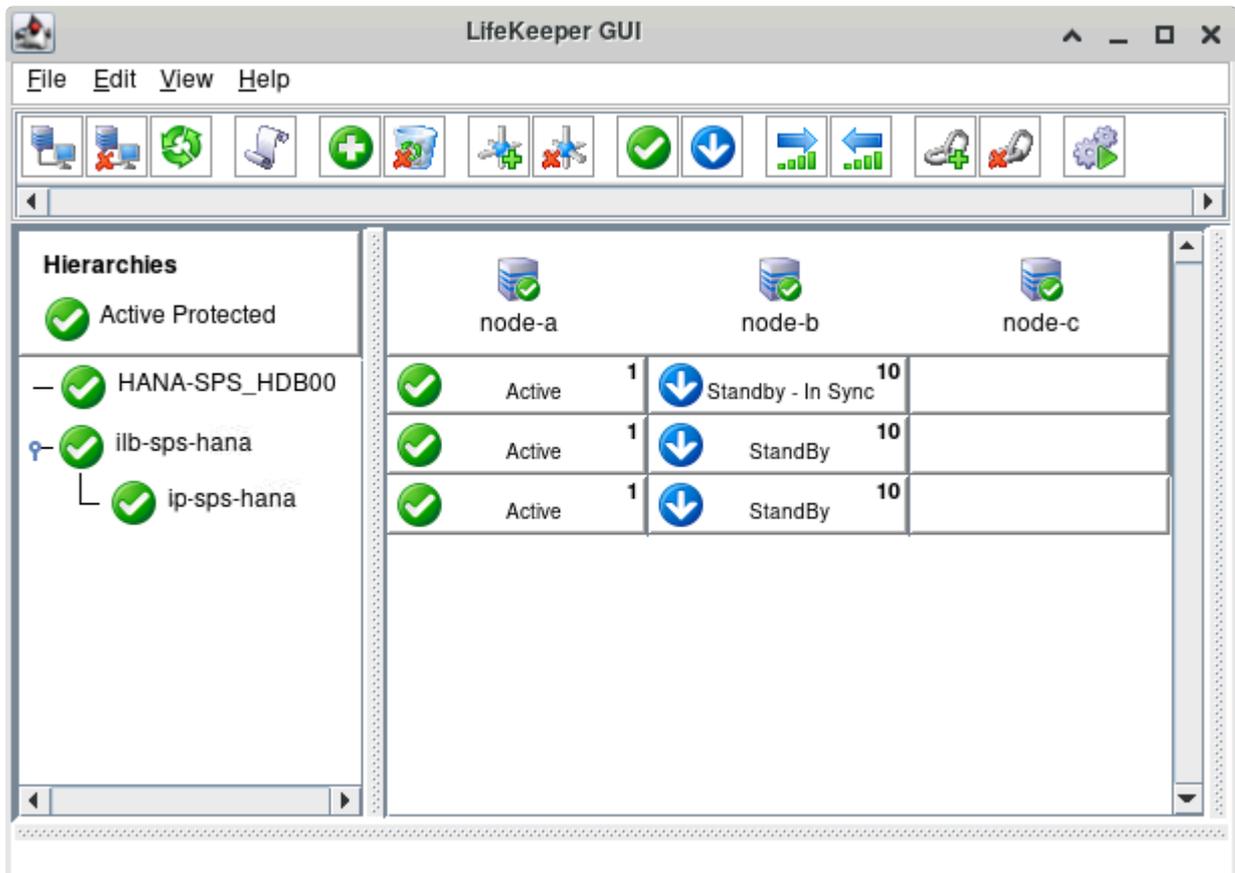


2. Enter the following values to create and extend a LifeKeeper resource (**HANA-SPS\_HDB00**) to protect the HDB00 database instance on node-a and node-b. Note that the resulting resource should not be extended to node-c, the witness node. The  icon indicates that the default option is chosen.

Field	Value
<b>Create Resource Wizard</b>	
Switchback Type	intelligent 
Server	node-a
SAP HANA SID	SPS 
SAP HANA Instance for SPS	HDB00 
IP child resource	none 
Enable/Disable Local Recovery	Enable 
HANA Tag	HANA-SPS_HDB00 

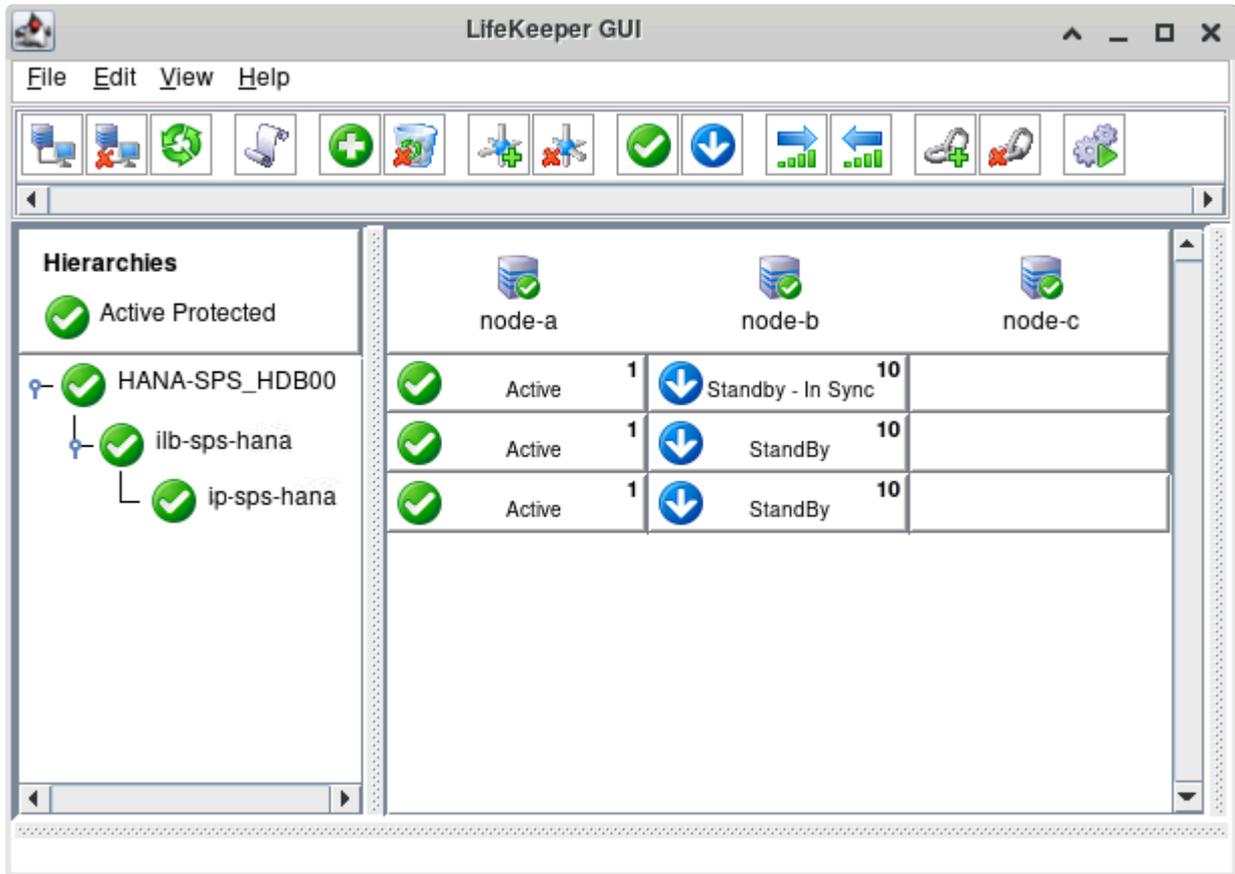
Pre-Extend Wizard	
Target Server	node-b
Switchback Type	intelligent ✓
Template Priority	1 ✓
Target Priority	10 ✓
Extend gen/app Resource Hierarchy Wizard	
Root Tag	HANA-SPS_HDB00 ✓
Enable/Disable Local Recovery	Enable ✓

Once the **HANA-SPS\_HDB00** resource has been successfully created and extended to node-b, the LifeKeeper resource panel should look similar to the following.



- Right-click the **HANA-SPS\_HDB00** resource on node-a and click "Create Dependency..." Specify **ilb-sps-hana** as the Child Resource Tag and click **Create Dependency**.

Once the dependency has been created, the LifeKeeper GUI should resemble the following image.

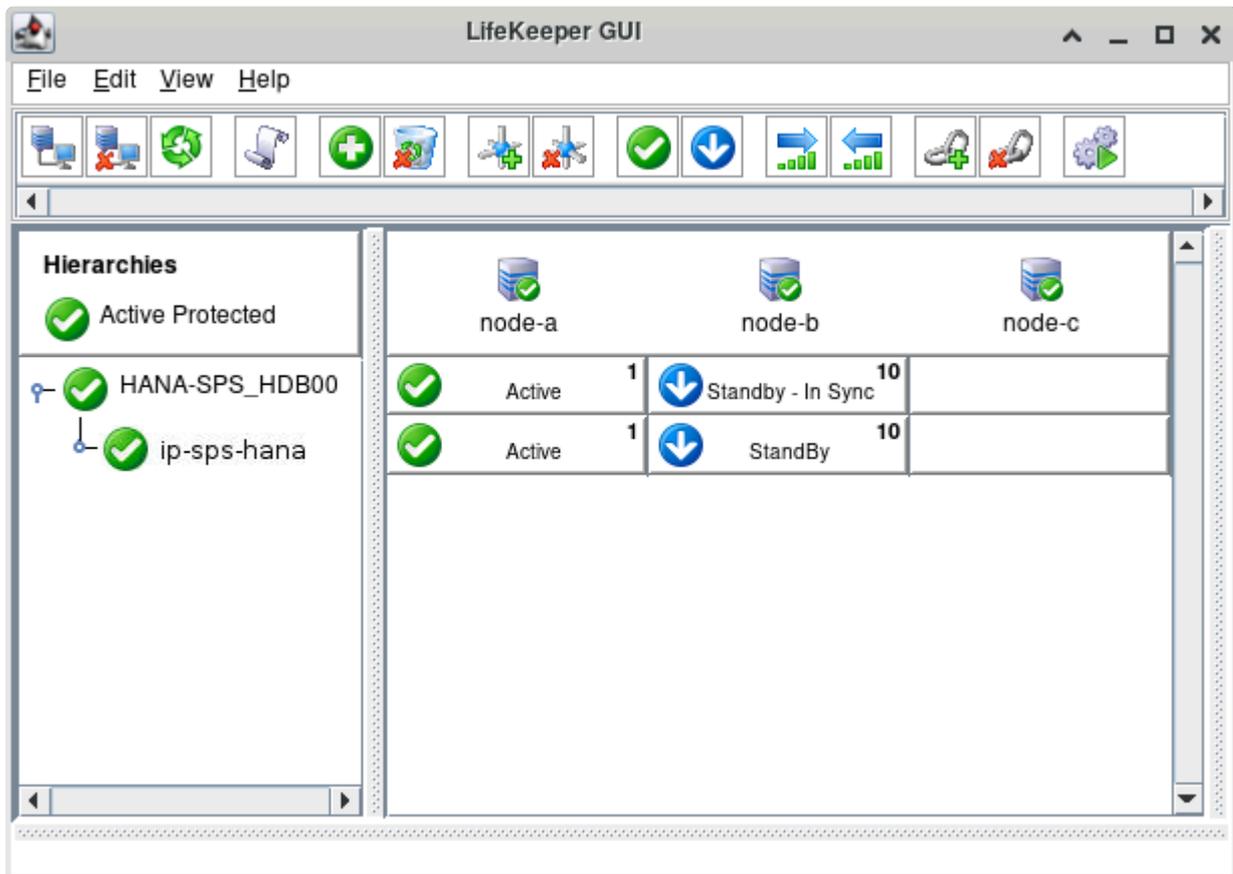


SAP HANA database instance HDB00 has now been successfully protected by LifeKeeper.

## 11.2.7.5.6.5. Test Switchover and Failover

In this section we will perform basic tests to verify the expected behavior of the **HANA-SPS\_HDB00** resource hierarchy on switchover and failover.

1. Verify that the **HANA-SPS\_HDB00** resource state is currently “Active” on node-a and “Standby – In-Sync” on node-b.



2. Execute the following commands on node-a and node-b to verify that the HDB00 instance is running successfully on both nodes, node-a and node-b are both registered appropriately in SAP HANA System Replication, and that a connection to the primary database can be established through the virtual hostname sps-hana. The replication and operation modes should be primary/primary and sync/logreplay on node-a and node-b, respectively.

\* In a Google Cloud deployment, the steps given in the **Disable IP Forwarding** section of [Google Cloud – Using an Internal Load Balancer](#) must be implemented in order to allow connections from node-b to the primary database instance on node-a using the virtual host name sps-hana.

```
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
04.03.2021 21:31:50
GetProcessList
OK
```

```

name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2021 03 04 19:44:27, 1:47:23, 19236
hdbcompileserver, HDB Compileserver, GREEN, Running, 2021 03 04 19:45:46, 1:46:04, 19618
hdbindexserver, HDB Indexserver-SPS, GREEN, Running, 2021 03 04 19:45:47, 1:46:03, 19666
hdbnameserver, HDB Nameserver, GREEN, Running, 2021 03 04 19:44:27, 1:47:23, 19254
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2021 03 04 19:45:46, 1:46:04, 19621
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2021 03 04 19:46:15, 1:45:35, 20241
hdbxsengine, HDB XSEngine-SPS, GREEN, Running, 2021 03 04 19:45:47, 1:46:03, 19669
    
```

```

[root@node-a ~]# su - spsadm -c "hdbnsutil -sr_state"
System Replication State
~~~~~

online: true

mode: primary
operation mode: primary
site id: 1
site name: SiteA

is source system: true
is secondary/consumer system: false
has secondaries/consumers attached: true
is a takeover active: false

Host Mappings:
~~~~~

node-a -> [SiteB] node-b
node-a -> [SiteA] node-a

Site Mappings:
~~~~~
SiteA (primary/primary)
  |--SiteB (sync/logreplay)

Tier of SiteA: 1
Tier of SiteB: 2
    
```

```

Replication mode of SiteA: primary
Replication mode of SiteB: sync

Operation mode of SiteA: primary
Operation mode of SiteB: logreplay

Mapping: SiteA -> SiteB
done.

```

```

[root@node-b ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
04.03.2021 21:38:10
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2021 03 04 19:48:50, 1:49:20, 13383
hdbcompileserver, HDB Compileserver, GREEN, Running, 2021 03 04 19:49:00, 1:49:10, 13747
hdbindexserver, HDB Indexserver-SPS, GREEN, Running, 2021 03 04 19:49:00, 1:49:10, 13794
hdbnameserver, HDB Nameserver, GREEN, Running, 2021 03 04 19:48:51, 1:49:19, 13401
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2021 03 04 19:49:00, 1:49:10, 13750
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2021 03 04 19:49:42, 1:48:28, 15167
hdbxsengine, HDB XSEngine-SPS, GREEN, Running, 2021 03 04 19:49:00, 1:49:10, 13797

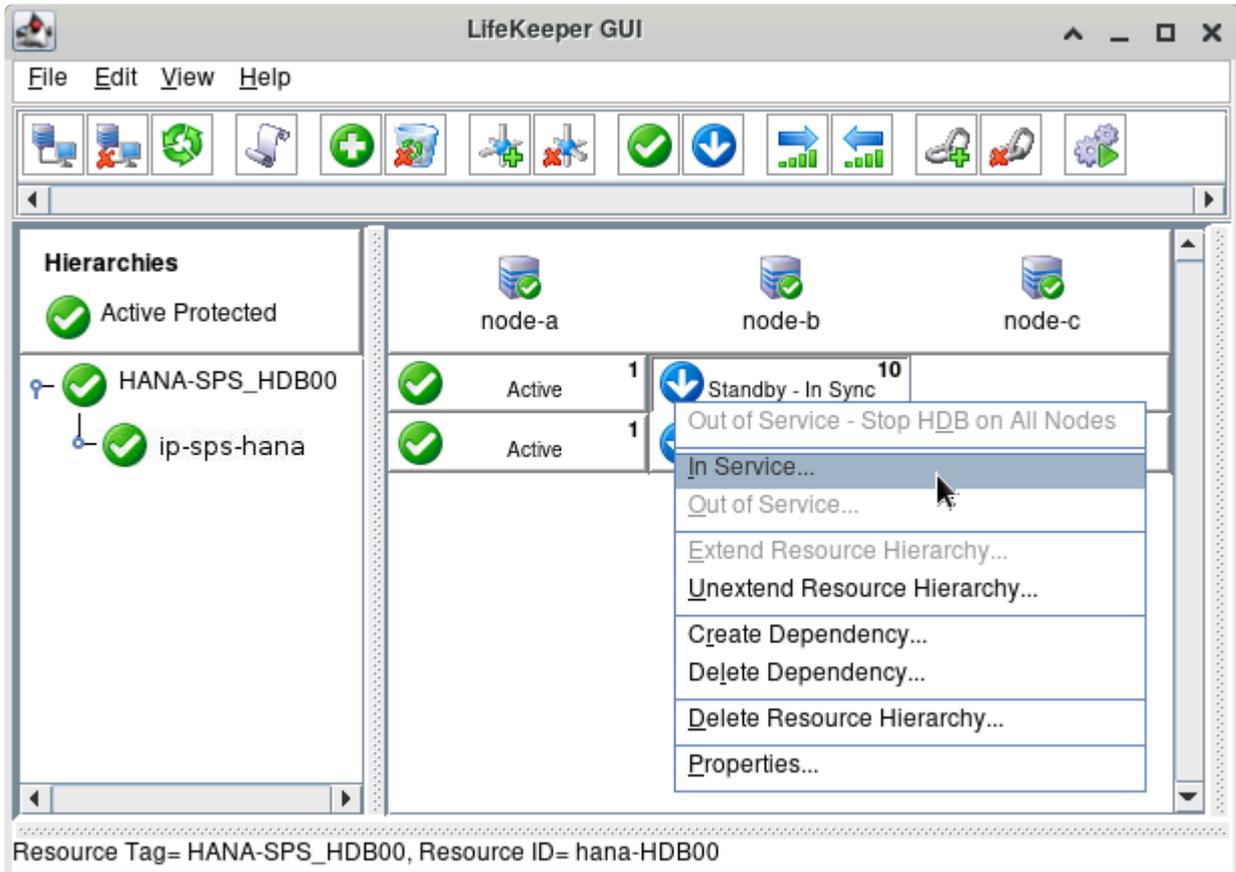
```

```

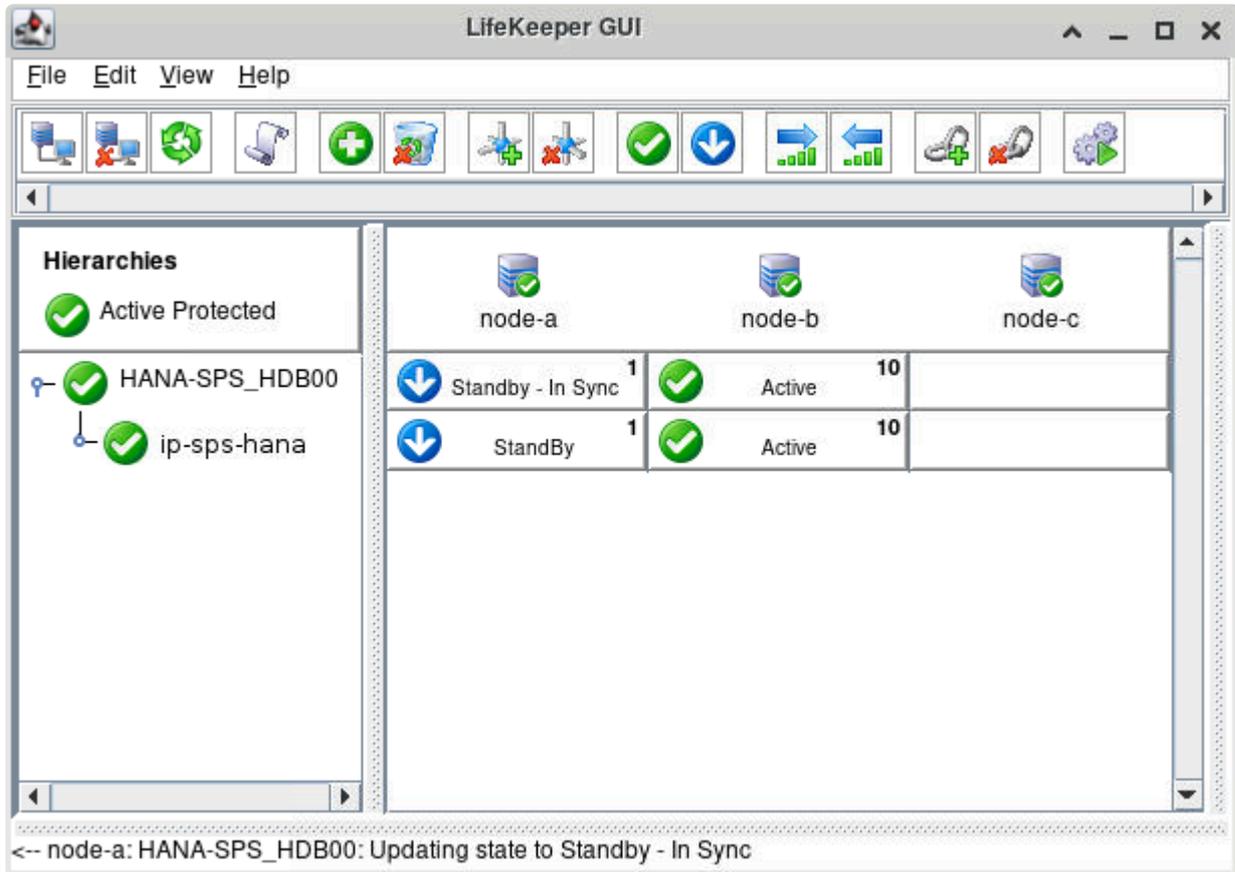
[root@node-b ~]# su - spsadm -c "hdbsql -n sps-hana -i 00 -u SYSTEM -p <SYSTEM user password> -d SPS '\s'"
host          : sps-hana:30013
sid           : SPS
dbname       : SPS
user         : SYSTEM
kernel version: 2.00.046.00.1581325702
SQLDBC version:          libSQLDBCHDB 2.04.182.1579711187
autocommit   : ON
locale       : en_US.UTF-8
input encoding: UTF8
sql port     : node-a:30015

```

3. Perform a switchover of the SAP HANA resource hierarchy by right-clicking the **HANA-SPS\_HDB00** resource on node-b and choosing the **In-Service...** operation. Click **In Service** to begin the switchover.



Allow time for the switchover to complete and for the resource states to switch to “Standby – In Sync” on node-a and “Active” on node-b. Note that it may take up to two minutes for the resource state to transition from “Standby – HDB Running” to “Standby – In-Sync” after the standby database instance is running.



- Execute the following commands on node-a and node-b to verify that the HDB00 instance is running successfully on both nodes, node-a and node-b are both registered appropriately in SAP HANA System Replication, and that a connection to the primary database can be established through the virtual hostname sps-hana. The replication and operation modes should now be sync/ logreplay and primary/primary on node-a and node-b, respectively.

```
[root@node-a ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
04.03.2021 21:50:20
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2021 03 04 21:43:51, 0:06:29, 156845
hdbcompileserver, HDB Compileserver, GREEN, Running, 2021 03 04 21:43:58, 0:06:22, 156970
hdbindexserver, HDB Indexserver-SPS, GREEN, Running, 2021 03 04 21:43:58, 0:06:22, 157101
hdbnameserver, HDB Nameserver, GREEN, Running, 2021 03 04 21:43:52, 0:06:28, 156864
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2021 03 04 21:43:58, 0:06:22, 156973
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2021 03 04 21:44:36, 0:05:44, 157429
hdbxsengine, HDB XSEngine-SPS, GREEN, Running, 2021 03 04 21:43:58, 0:06:22, 157104
```

```
[root@node-a ~]# su - spsadm -c "hdbssql -n sps-hana -i 00 -u SYSTEM -p <SYSTEM user password> -d SPS '\s'"
host          : sps-hana:30013
sid           : SPS
dbname       : SPS
user         : SYSTEM
kernel version: 2.00.046.00.1581325702
SQLDBC version:          libSQLDBC_HDB 2.04.182.1579711187
autocommit   : ON
locale       : en_US.UTF-8
input encoding: UTF8
sql port     : node-b:30015
```

```
[root@node-b ~]# su - spsadm -c "sapcontrol -nr 00 -function GetProcessList"
04.03.2021 21:52:36
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2021 03 04 19:48:50, 2:03:46, 13383
hdbcompileservers, HDB Compilerserver, GREEN, Running, 2021 03 04 19:49:00, 2:03:36, 13747
hdbindexserver, HDB Indexserver-SPS, GREEN, Running, 2021 03 04 19:49:00, 2:03:36, 13794
hdbnameserver, HDB Nameserver, GREEN, Running, 2021 03 04 19:48:51, 2:03:45, 13401
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2021 03 04 19:49:00, 2:03:36, 13750
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2021 03 04 19:49:42, 2:02:54, 15167
hdbxsengine, HDB XSEngine-SPS, GREEN, Running, 2021 03 04 19:49:00, 2:03:36, 13797
```

```
[root@node-b ~]# su - spsadm -c "hdbnsutil -sr_state"
System Replication State
~~~~~

online: true

mode: primary
operation mode: primary
site id: 2
site name: SiteB

is source system: true
is secondary/consumer system: false
```

```
has secondaries/consumers attached: true
is a takeover active: false

Host Mappings:
~~~~~

node-b -> [SiteB] node-b
node-b -> [SiteA] node-a

Site Mappings:
~~~~~
SiteB (primary/primary)
  |--SiteA (sync/logreplay)

Tier of SiteB: 1
Tier of SiteA: 2

Replication mode of SiteB: primary
Replication mode of SiteA: sync

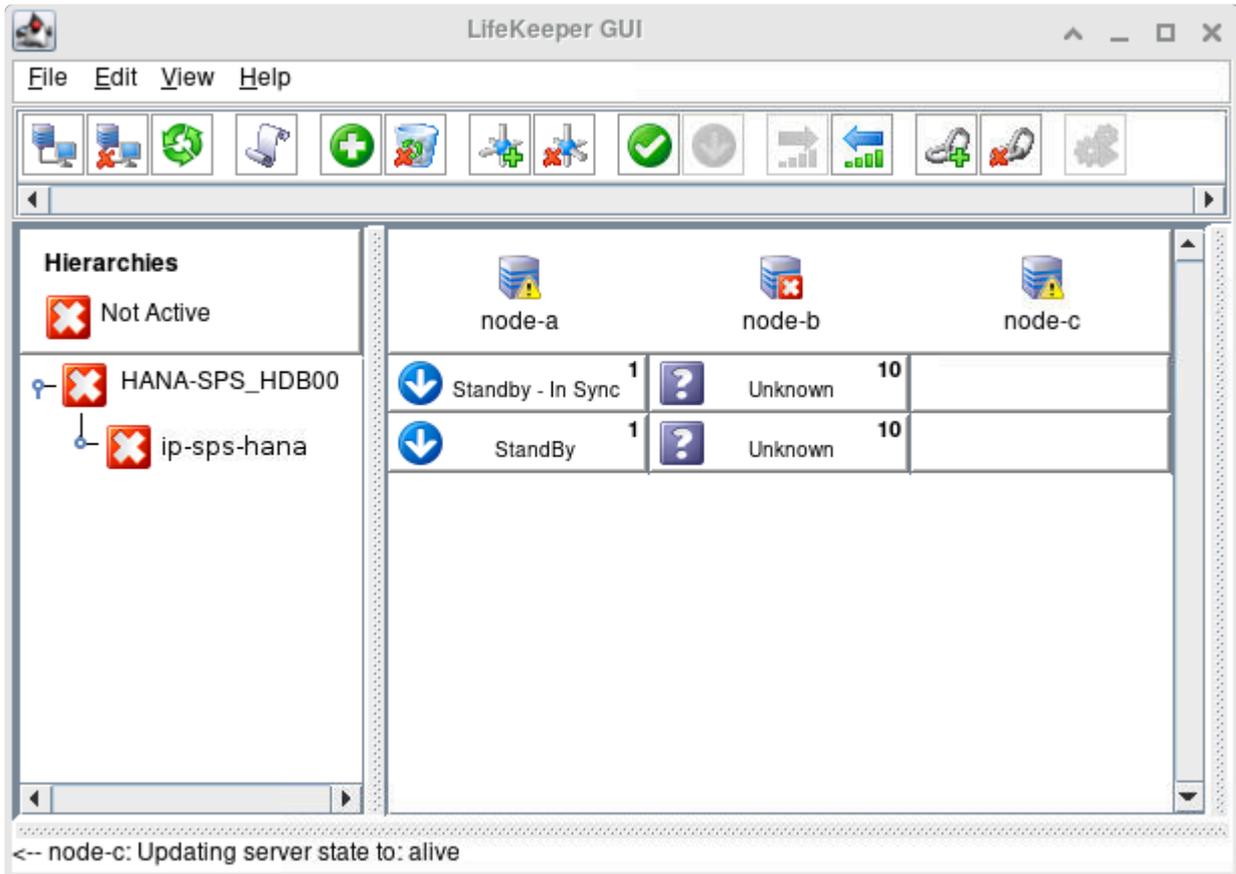
Operation mode of SiteB: primary
Operation mode of SiteA: logreplay

Mapping: SiteB -> SiteA
done.
```

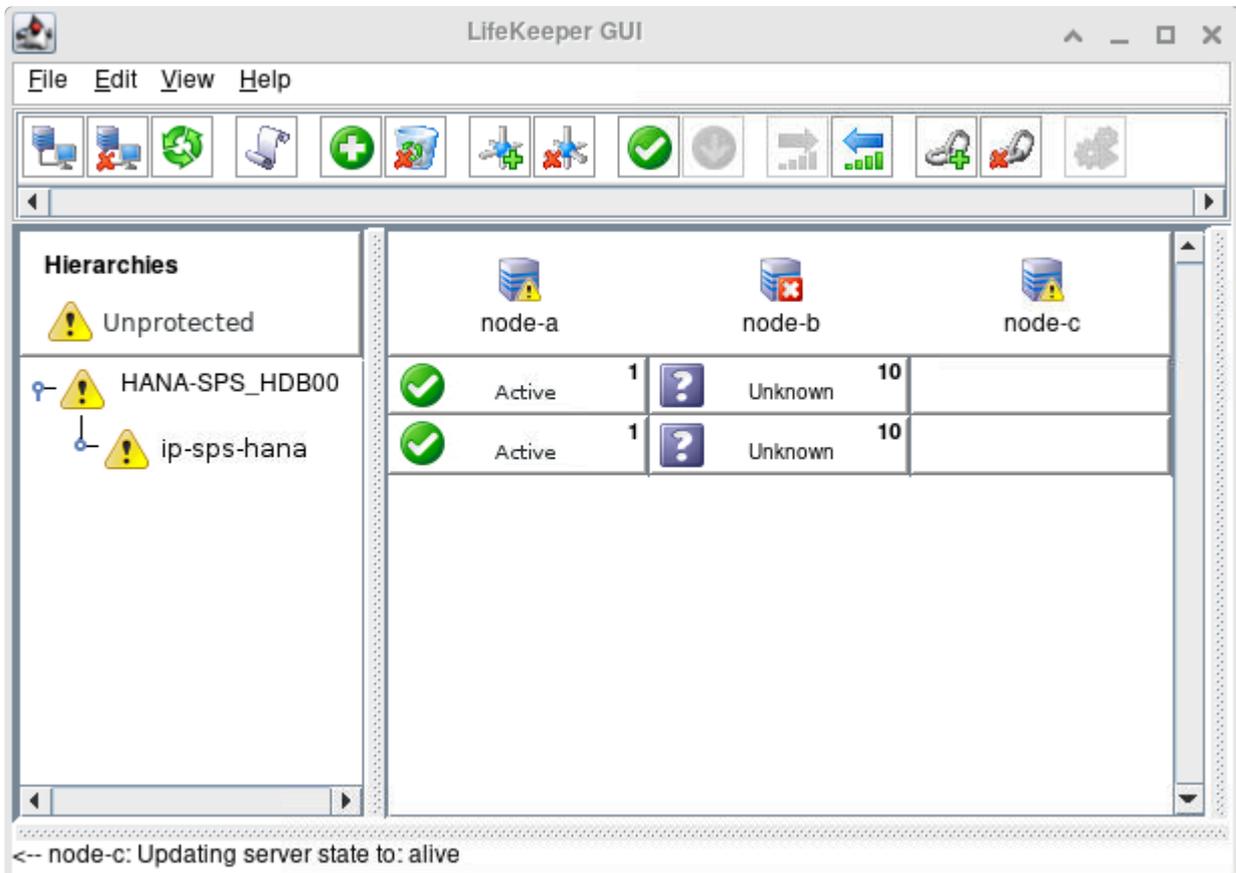
**5. Execute the following command to forcefully reboot node-b:**

```
[root@node-b ~]# echo b > /proc/sysrq-trigger
```

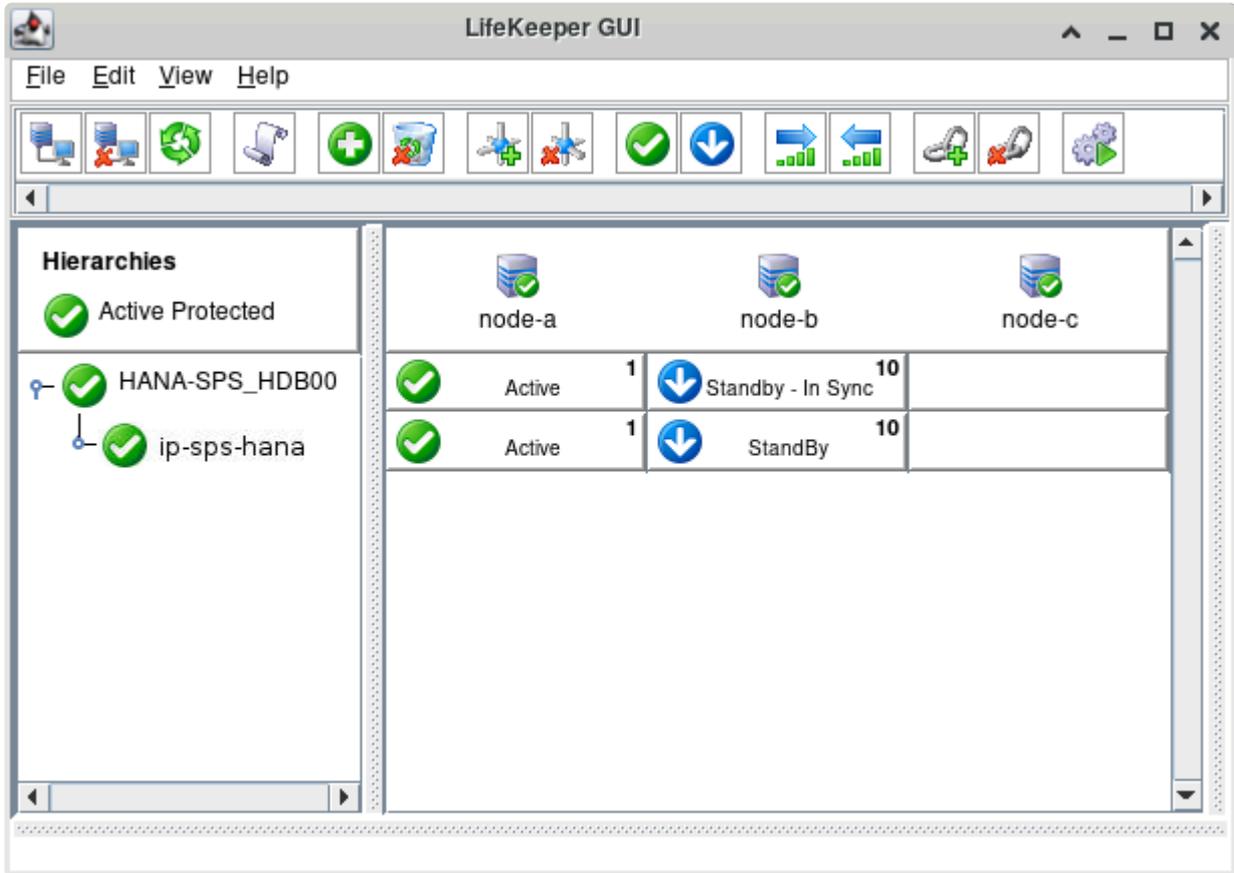
Once LifeKeeper has detected that node-b has been powered off, the status of node-b updates to “Unknown” in the LifeKeeper GUI.



At this point, LifeKeeper automatically initiates failover of the **HANA-SPS\_HDB00** resource hierarchy back to node-a.



Once node-b is back online, LifeKeeper will automatically re-register it as a secondary replication site in SAP HANA System Replication and restart the secondary database instance. This process may take several minutes to complete. Once this process is complete, the resource states will transition back to “Active” on node-a and “Standby – In-Sync” on node-b.



- Execute the commands given in step 2 again to verify that the database is functioning as expected on both nodes.

We have now verified the basic switchover and failover functionality of the SAP HANA resource hierarchy.

## 11.2.7.6. Common Tasks

---

The following topics contain common tasks across different resources.

- [How to Confirm if the Data Storage is Available on a Node](#)
- [Switchover the Data Storage to the Other Node](#)

# 11.2.7.6.1. How to Confirm if the Data Storage is Available on a Node

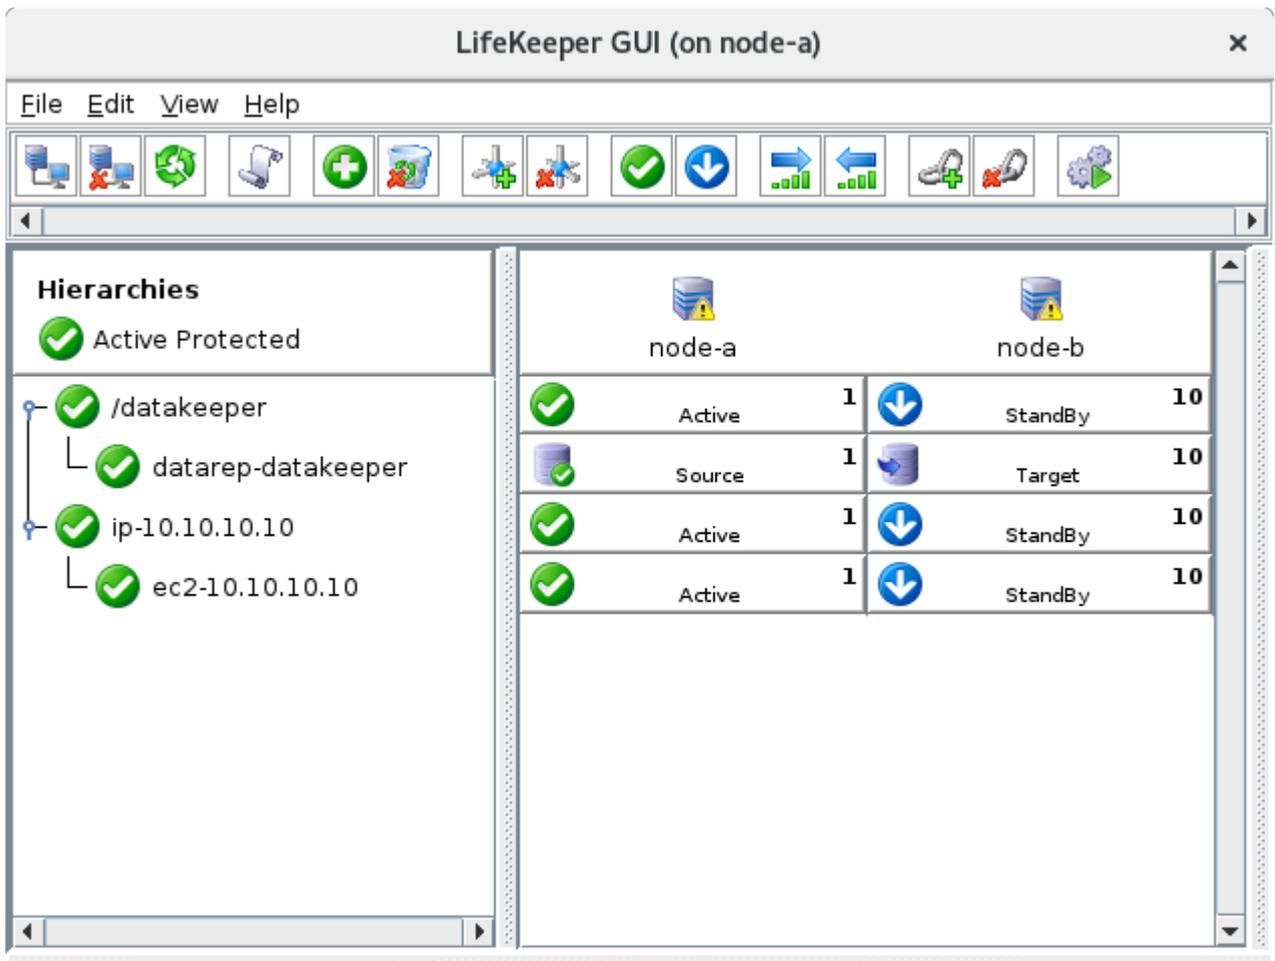
During these steps we will be switching data storage between nodes. Here are two options to check whether the data storage (per this guide, located at /datakeeper) is available in a node.

- View the output from the `df` command.

```

1 $ df
2 Filesystem      1K-blocks    Used Available Use% Mounted on
3 devtmpfs        916888         0   916888  0% /dev
4 tmpfs           7340032         0  7340032  0% /dev/shm
5 tmpfs           940184    17424   922760  2% /run
6 tmpfs           940184         0   940184  0% /sys/fs/cgroup
7 /dev/xvda2      31444972 17183608 14261364 55% /
8 tmpfs           188040         0   188040  0% /run/user/1000
9 /dev/md0        20960236 2510944 18449292 12% /datakeeper
    
```

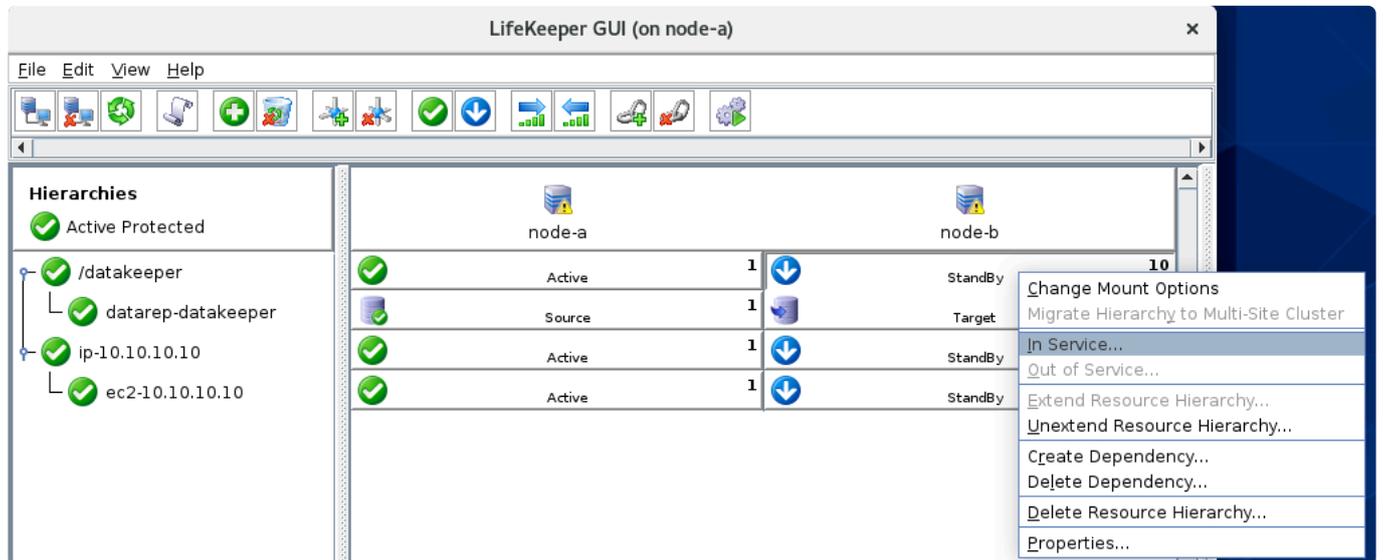
- Check the LifeKeeper GUI to see the status of the `datarep-datakeeper` resource. The node status should be “Source” and the other node should be labeled as “Target”.



<-- node-b: datarep-datakeeper: Updating state to Target

# 11.2.7.6.2. Switchover the Data Storage to the Other Node

To switchover the disk to the other node using the LifeKeeper GUI, select the `/datakeeper` resource on the standby node, then select "In Service...".



Once the `datarep-datakeeper` resource becomes active on the other node, the Source & Target labels are switched.

**LifeKeeper GUI (on node-a)** ✕

File Edit View Help

**Hierarchies**

- ✓ Active Protected
- ✓ /datakeeper
  - ✓ datarep-datakeeper
- ✓ ip-10.10.10.10
  - ✓ ec2-10.10.10.10

node-a		node-b	
↓	StandBy	1	✓ Active <b>10</b>
↓	Target	1	↓ Source <b>10</b>
✓	Active	1	↓ StandBy <b>10</b>
✓	Active	1	↓ StandBy <b>10</b>

<-- node-a: datarep-datakeeper: Updating state to Target

# 12. Quick Start Guides

---

[AWS Direct Connect Quick Start Guide](#)

[Microsoft Azure Quick Azure Guide](#)

[MySQL Cluster with Data Replication \(Shared Nothing Cluster\)](#)

[LifeKeeper for Linux in the AWS Cloud \(SAP\)](#)

[Connecting to a LifeKeeper Cluster using AWS Transit Gateway Quick Start Guide](#)

[AWS VPC Peering Connections Quick Start Guide](#)

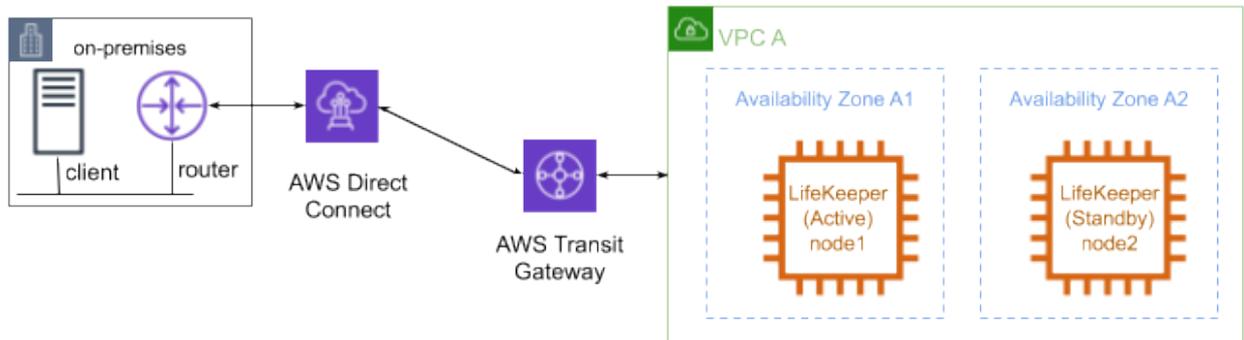
[PostgreSQL Cluster with Shared Storage \(iSCSI\)](#)

[Apache/MySQL Cluster Using Both Shared and Replicated Storage](#)

# 12.1. AWS Direct Connect Quick Start Guide

## Objective

With the release of AWS Transit Gateway, a route table scenario for the Recovery Kit for EC2 is now available with the configuration where the on-premises environment (on-premises in the figure below) using AWS Direct Connect is connected to the HA cluster nodes located in VPC (VPC A) via AWS Transit Gateway.



This document describes the requirements and basic operations for building connections from outside VPC with LifeKeeper for Linux v9.4.1.

This document does not cover the basic settings, operations, and technical details of LifeKeeper and Amazon Web Service (AWS). For terms, operations and technical information related to LifeKeeper and AWS, that are the prerequisites of this configuration, review the related documents and user websites.

\* **Note: Amazon Web Services, Powered by Amazon Web Services logo, AWS Amazon EC2, EC2, Amazon Elastic Compute Cloud, Amazon Virtual Private Cloud, Amazon Route 53 and Amazon VPC are trademarks of Amazon.com, Inc. or its affiliates in the United States and other countries.**

## 12.1.1. AWS Direct Connect Requirements

---

The following is a summary of requirements that should be met for an AWS environment and instances created on it.

### Requirements for AWS Environment

Create a base environment on AWS to provide services. The requirements for using this configuration are as follows.

#### Amazon Virtual Private Cloud (VPC)

- VPC needs to be configured in AWS.
- The subnet where the primary instance is located and the subnet where the standby instance is located must be created in different Availability Zones (AZ).

#### Amazon Elastic Compute Cloud (EC2)

- At least 2 instances are required.
- A primary instance and a standby instance need to be configured to start with different AZ for each.
- Instances are connected to Elastic Network Interface (ENI).
- Instances are required to satisfy LifeKeeper's installation requirements.
- The AWS Command Line Interface (AWS CLI) needs to be installed in each of EC2 the instances. For the details, please refer to [AWS Command Line Interface installation](#).
- You need to be able to access Amazon EC2 Web Services endpoint URL (EC2 URL) using https and Amazon EC2 metadata URL (<http://169.254.169.254/>) using http.

#### AWS Identity and Access Management (IAM)

In order for LifeKeeper to operate AWS, an IAM user or IAM role with the following access privilege is required. Configure an [EC2 IAM role](#) or configure [AWSCLI](#) appropriately so that it can be accessed from root user of the EC2 instance.

- ec2:DisassociateAddress
- ec2:DescribeAddresses
- ec2:AssociateAddress
- ec2:DescribeRouteTables

- ec2:ReplaceRoute

## AWS Transit Gateway

- The VPC with the cluster nodes and the on-premises environment where the clients are located must be connected via AWS Transit Gateway; not via Virtual Private Gateway.
- Enable the Default route table association and the Default route table propagation when creating AWS Transit Gateway.
- Connect VPC by creating Transit Gateway Attachment.
- Connect to AWS Direct Connect by selecting the created AWS Transit Gateway in the Gateway association configuration of Direct Connect Gateway. At this time, configure both the network address of the VPC where the cluster nodes are located and the virtual IP address in Allowed prefixes.

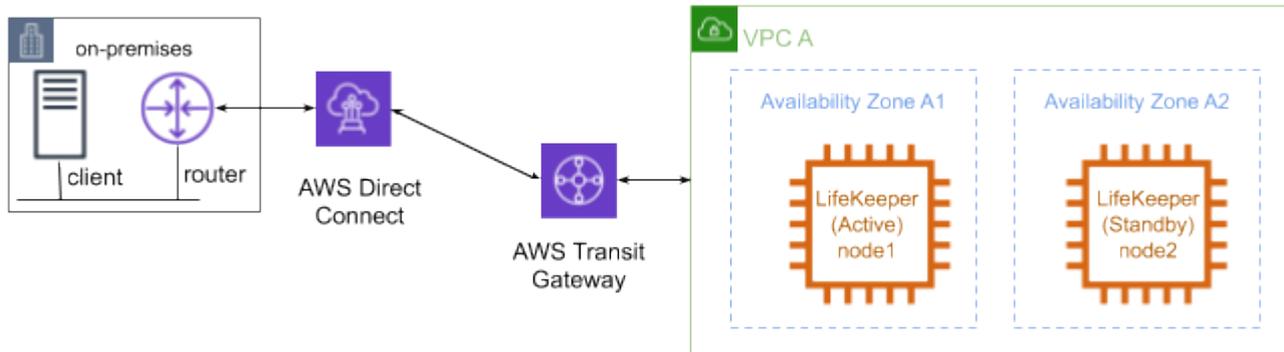
## LifeKeeper Software Requirements

You need to install the same version of LifeKeeper software and patches on each server. The Application Recovery Kit (ARK) required for this configuration is shown below. For the specific LifeKeeper requirements, please refer to: [LifeKeeper for Linux Technical Documentation](#) and [LifeKeeper for Linux Release Notes](#)

- LifeKeeper IP Recovery Kit
- LifeKeeper Recovery Kit for EC2

## 12.1.2. AWS Direct Connect Setup Procedure

This section describes the general procedure to setup the environment shown below.



## 12.1.2.1. AWS Direct Connect Preparations

---

Create an environment that meets the [AWS Direct Connect requirements](#). Install LifeKeeper on each instance and create a communication path between Node1 and Node2.

## 12.1.2.2. Creating Direct Connect Resources

---

### Creating an IP Resource

- Create a virtual IP resource. The IP resource address must be outside the CIDR block managed by the VPC.

### Creating an EC2 Resource

- Create EC2 resources. For the IP resource requested when creating resources, specify the resource created in “Create IP Resource” above. Specify the Route Table (Backend Cluster) as the EC2 resource type required when creating resources.

### Creating Resources for Protected Services

- Create resources for the services you want to protect. If an IP resource is required for resource creation, specify the resource created in “IP Resource Creation” above. Configure resource dependencies so that the resources of the protected service are the parent resources and the EC2 resources are the child resources.

## 12.1.2.3. Configuring a Route Table

---

Configure a route table as shown below.

- Add the route information to the on-premises environment network to the route table of the VPC or subnet where the cluster nodes are located.

Destination Address	Target
On-premises network address	Created Transit Gateway

- Add the route information to the Virtual IP address in the route table of the Transit Gateway.

Destination Address	Target
Virtual IP address	VPC where the cluster nodes are located

- Configure the routing information of clients and routers in the on-premises environment so that the destination of packets to the network address and virtual IP address of the VPC where the cluster nodes are located are the Direct Connect.

Once configured, make sure that the client can access the private address and virtual IP address of the cluster server.

## 12.1.3. Considerations for Settings and Operations in AWS Direct Connect

---

### Considering the Use of LifeKeeper I-O Fencing

Since the shared disk environment cannot be used in an AWS environment, you cannot use SCSI reservations to prevent a split-brain. IP resources may cause a split-brain as it uses the real IP resource with different IP addresses for each node.

For this reason, please consider the use of Quorum/Witness server or STONITH, an I/O fencing function of LifeKeeper to use this configuration safely.

Since you can implement I/O fencing separately without the Quorum server, if you use the TCP\_REMOTE setting in Quorum mode, it is easy to implement in the cloud environment. For more details, please refer to the following:

- [Quorum/Witness](#)
- [STONITH](#)

### AWS Direct Connect Known Issues and Troubleshooting

There are currently no Known Issues.

## 12.2. Microsoft Azure Quick Start Guide

This guide provides the steps to deploy LifeKeeper for virtual machines and increase their availability. It is based on actual work procedures to help you use LifeKeeper for Linux with Microsoft Azure (Azure) Resource Manager. It will guide you through the steps to configure the following examples of a cluster environment on an Azure VM.

- LifeKeeper cluster (2-node cluster) \*Multi-NIC configuration
- Data mirror type shared data area with DataKeeper replication
- IP resources, Oracle resources

 **Disclaimer:** This guide outlines a general LifeKeeper for Linux configuration in a Microsoft Azure environment. The setting values used in this guide are listed with an excerpt of the items required for explanation. Items that are not listed are assumed to be default values or arbitrary values. You can configure additional Azure infrastructure options by replacing the setting values described in this guide according to your actual requirements.

### Future Compatibility

The information contained in this document is based on the results of validation performed on Azure as of March 2020. We do not guarantee any changes in the specifications of Azure and LifeKeeper in the future. Please refer to the most recent documentation and make the appropriate settings if necessary.

### Restrictions

- Multi-NIC configurations are supported only on RHEL7 and CentOS7 where “Predictable Network Interface Name” can be used.
- RHEL5, CentOS5, RHEL6, CentOS6 only support single NIC configurations, not multi-NIC configurations.
- Active/Active configurations are not supported.
- Services that use multiple ports or configurations that provide multiple services are not supported.

### Terms

- **Internal Load Balancing or ILB** – This is a function that enables load balancing between virtual machines that reside inside Azure cloud services or virtual networks with limited regions.
- **[VIP protected with IP resources]** – A virtual IP address created with LifeKeeper.
- **[VIP configured by ILB]** – A virtual IP address of the Internal Load Balancing created in Azure.
- **[Azure private IP address]** – A private address used in Azure Virtual Network.

## 12.2.1. Microsoft Azure Overview

---

Azure is a public cloud service provided by Microsoft that allows users to use hardware, networks, storage (e.g. disks), server operating system (e.g. Windows Server and Linux), middleware (e.g. web servers and RDBMS), groupware and server applications, and application runtime environments (.Net Framework, etc.) on demand through network.

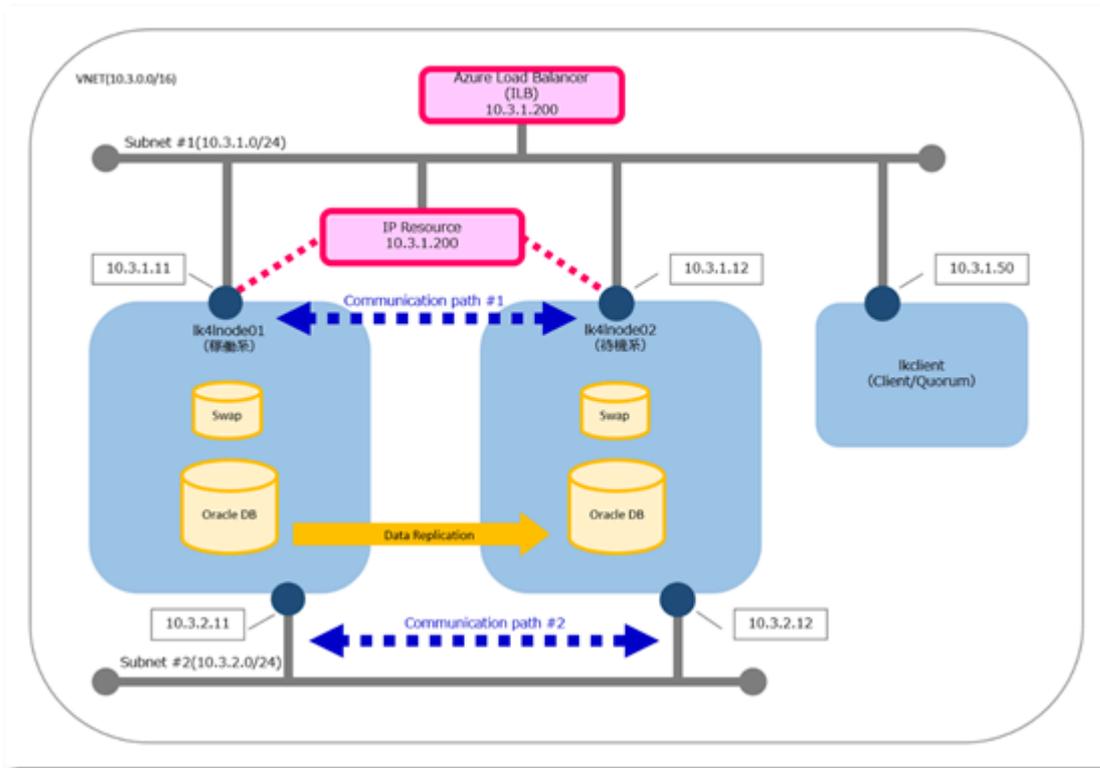
Azure enables a user to easily create, deploy and manage server applications over the Internet in Microsoft-managed data centers around the world. It allows you to use computing resources such as memory and disks when you need them without having to own servers or infrastructure. The pricing is different by resource and is based on an hourly pay-as-you-go system. The web browser-based management portal allows users to operate intuitively.

Azure offers services in four forms:

- Virtual Machines (VMs)
- Websites
- Mobile Services
- Cloud Services

# 12.2.2. Configurations

In this configuration, LifeKeeper is used to build an active/standby cluster as shown below. LifeKeeper is deployed on an Azure VM with a multi-NIC configuration.



## Server Configuration

<p><b>Virtual Machine size</b></p>	<p>Cluster node: Standard_A1 (1 core, 1.75 GB memory)</p> <p>Client/Witness Server is Standard B1s (1 vcpu number, 1 GiB memory)</p> <p>(0.25 core, 0.75GB memory)</p> <p>*Note 1</p>
<p><b>Data Disk</b></p>	<p>30GiB (for Oracle DB)</p> <p>10GiB (for swap area)</p> <p>(/dev/disk/cloud/azure_resource is not usable as a data disk)</p> <p>*Note 2)</p>

	*Note 1
<b>[VIP protected with IP resources]</b>	10.3.1.200
<b>[IP to set in ILB]</b>	10.3.1.200
<b>ILB forwarding port</b>	1521
<b>ILB load balancing target hosts and ports</b>	1521 port for each lk4lnode01/lk4lnode02
<b>[Azure private IP address]</b>	Cluster node (Active) 10.3.1.11 / 10.3.2.11 Cluster node (Standby) 10.3.1.12 / 10.3.2.12 Client and Witness server 10.3.1.50

\*Note 1: Prepare the instance size and virtual machine disks to satisfy Oracle installation requirements (1 GB of memory, Oracle database installation disk size, and swap area size).

\*Note 2: Most Azure VMs contain a temporary disk (/dev/disk/cloud/azure\_resource), which is not a managed Disk. The temporary disk provides short-term storage for applications and processes and is intended to only store data such as page or swap files. The temporary storage is automatically mounted on /mnt but in the examples below it is being mounted on /mnt/resource. The azure\_resource is often the /dev/sdb device node but can be a different device node depending on the configuration. This temporary disk is NOT suitable to be used as a LifeKeeper protected device such as storage used with DataKeeper. Refer to the file /mnt/DATALOSS\_WARNING\_README.txt for more information.

## Software Configuration

<b>OS</b>	RedHat Enterprise Linux 7.8 64bit
<b>LifeKeeper</b>	SIOS Protection Suite for Linux v9.5.2
<b>Oracle</b>	Oracle Database 19c (19.3) Enterprise Edition

## Monitored Items

<b>Monitored items</b>	VIP PROTECTED WITH IP RESOURCES / File system / DataReplication / Oracle DB / Oracle Listener
------------------------	-----------------------------------------------------------------------------------------------

## Network Configuration

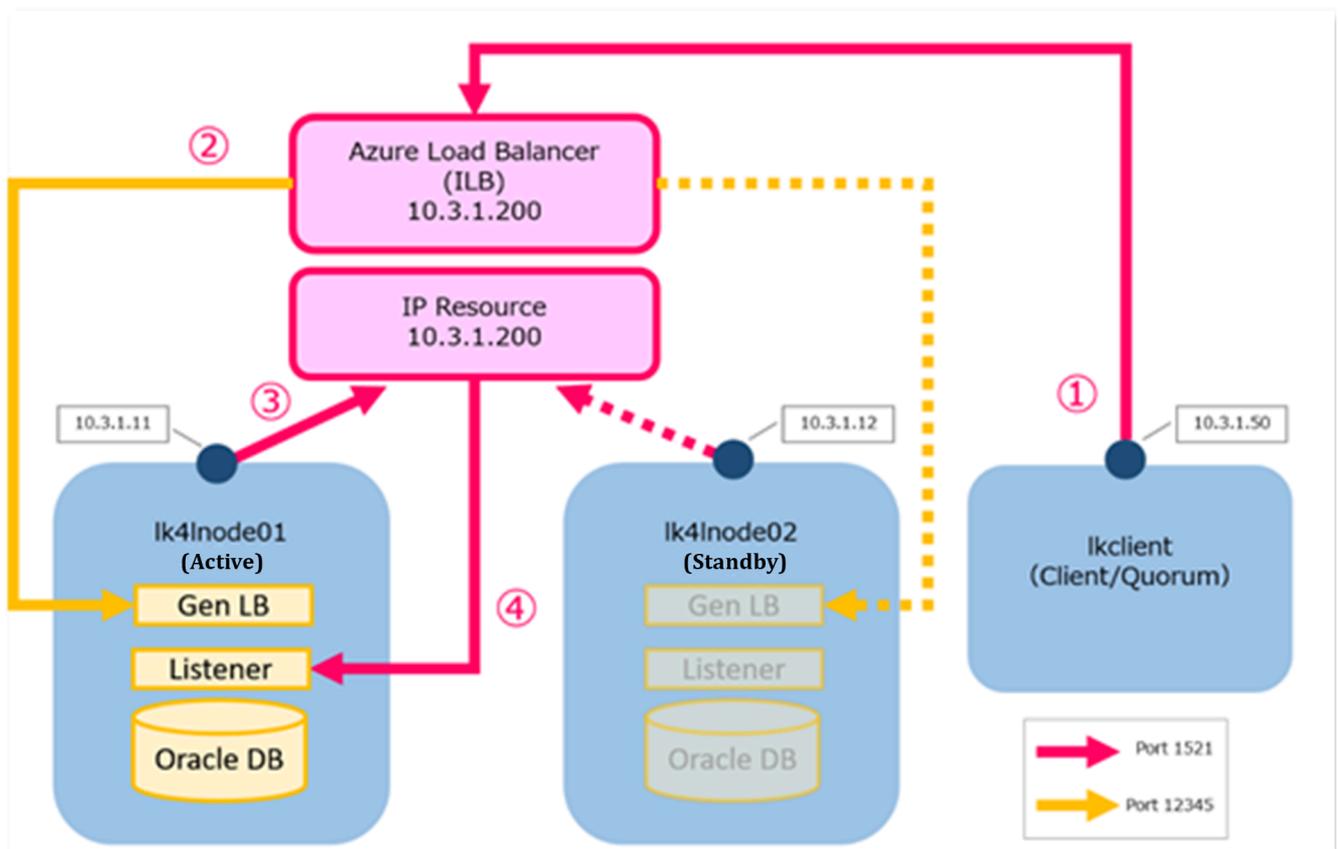
In Azure, a “virtual network” (VNET) is created to enable communication between virtual machines (VMs). The VNET enables communication between VMs within this subnet by specifying a subnet.

When configuring a network on a VM on Azure, it is common to create a VNET beforehand and specify this VNET for the VM. VNET allows you to provide a virtual private network (VPN) to the VM. Optionally, a VPN can be connected to the on-premises environment to enable hybrid or cross-premises solutions. VNET can control network topology, including DNS and IP address range configuration, through the management portal and Azure PowerShell.

## 12.2.3. LifeKeeper-Specific Configurations in Azure

### Internal Load Balancer (ILB)

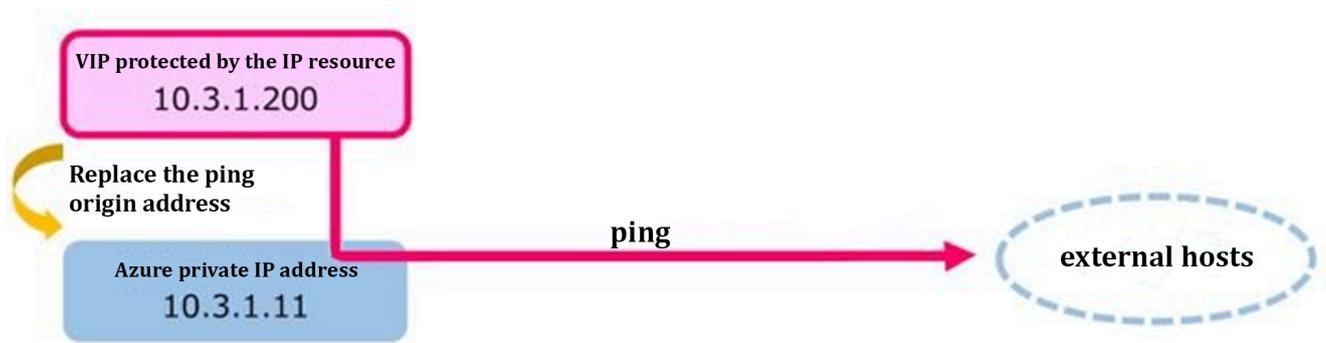
Azure is not capable of recognizing [VIP protected with IP resources] with the VNET. Because of this, network communication using [VIP protected with IP resources], which is usually assumed by LifeKeeper for Linux, cannot be performed. Therefore, LifeKeeper introduces ILB as follows with the [VIP set by ILB] set as a network communication path.



1. In order for the Oracle Client to connect to the Oracle Listener, start the connection to 10.3.1.200 (port 1521) via ILB.
2. In the ILB, 10.3.1.11 and 10.3.1.12 are registered in the load balancing destination (backend pool) with no port settings. Packets received by the ILB for 10.3.1.200 (port 1521) are forwarded to both the active and standby nodes.
3. Due to the LifeKeeper specifications, it is necessary to specify [VIP protected with IP resources] for the Oracle Listener resource and an IP resource for receiving needs to be created. Since the LifeKeeper IP resource (10.3.1.200) is in-service only on the active node, only the active node (10.3.1.11) receives requests.
4. As a result, the connection request from the Oracle Client is received by the active Oracle Listener (10.3.1.200: port 1521).

## Address Conversion

LifeKeeper uses PING to external hosts to monitor IP resources, but due to the Azure's specifications, PING with virtual IP as the source cannot be performed. However, by changing the packet source information as described below, PING can be enabled and network failures between external hosts can be detected.



1. From the IP resource, start ping with the VIP protected by the IP resource as an origin.
2. Replace the PING origin address from 10.3.1.200 (VIP protected by the IP resource) to 10.3.1.11 (Azure private IP address).
3. After replacing the origin, send the packet to the PING destination. If it is replaced with [Azure Private IP address], the response will be returned if the network is reachable.

## Quorum/Witness

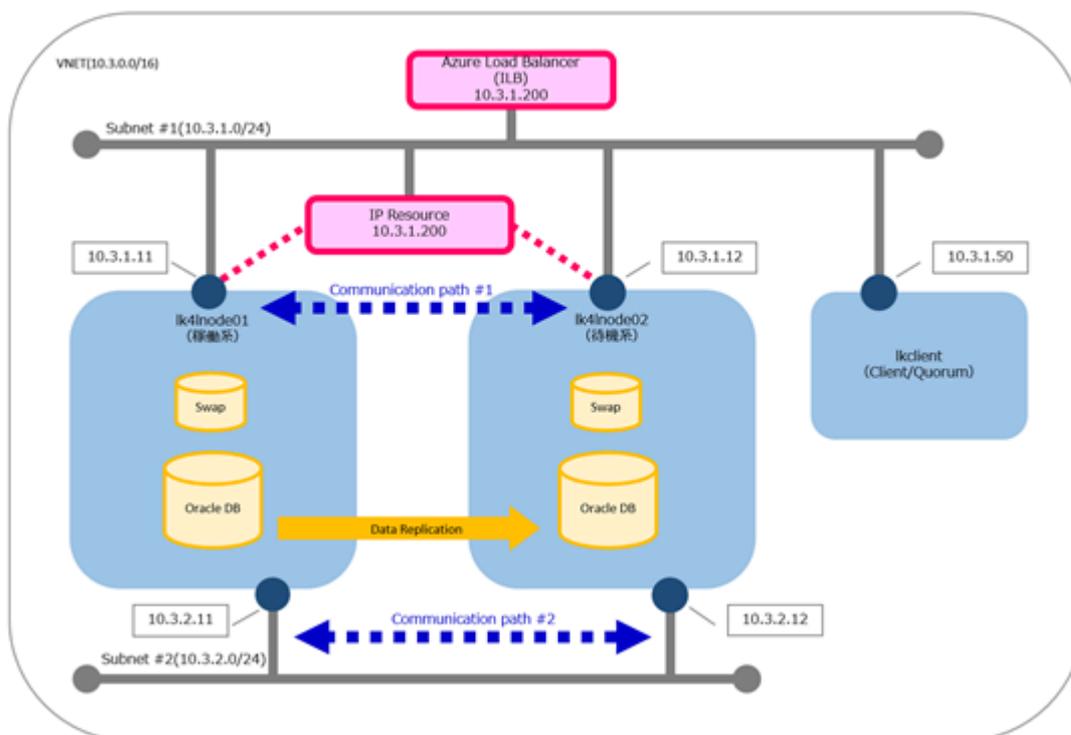
In an Azure environment, DataKeeper replication disks are used as an alternative to cluster shared disks. In a cluster that uses DataKeeper, when the network fails such that the communication path is completely disconnected, it may result in a split brain due to its mechanism. To avoid this, configure I/O Fencing using the Quorum/Witness functionality.

## 12.2.4. Building a Virtual Machine and Starting the OS

### Creating a Virtual Machine and Starting the OS

When using LifeKeeper with Azure it is necessary to make the following three settings specific to the Azure environment:

- Use internal load balancing (ILB) to enable connections to Oracle resources protected by LifeKeeper
- For communication using a virtual IP address, configure address conversion
- Setup Quorum to avoid split brain



## 12.2.4.1. Creating a Resource Group

1. Select **Azure Portal** > **Resource groups** > **Add** to create a resource group to be used in this environment. Enter the required values and click **Review + Create**.

[Home](#) > [Resource groups](#) >

### Create a resource group

**Basics**   Tags   Review + create

**Resource group** - A container that holds related resources for an Azure solution. The resource group can include all the resources for the solution, or only those resources that you want to manage as a group. You decide how you want to allocate resources to resource groups based on what makes the most sense for your organization. [Learn more](#) 

#### Project details

Subscription \* ⓘ

Visual Studio Premium with MSDN



Resource group \* ⓘ

#### Resource details

Region \* ⓘ

(US) East US

2. Review the details and click **Create**.

## 12.2.4.2. Creating a Virtual Network

Create a virtual network to use in this environment. Create one address space and two subnets, one of these subnets is for services and the other is for DataKeeper replication.

1. Select **Azure Portal > Virtual networks > Add** to create the virtual network and one subnet at a time for this environment.

Home > Virtual networks >

### Create virtual network

Basics IP Addresses Security Tags Review + create

Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. VNet enables many types of Azure resources, such as Azure Virtual Machines (VM), to securely communicate with each other, the internet, and on-premises networks. VNet is similar to a traditional network that you'd operate in your own data center, but brings with it additional benefits of Azure's infrastructure such as scale, availability, and isolation. [Learn more about virtual network](#)

**Project details**

Subscription \* ⓘ Visual Studio Premium with MSDN

Resource group \* ⓘ LK\_SPS\_Example  
[Create new](#)

**Instance details**

Name \*

Region \* (US) East US

Enter the required information and click **Create**.

Use the following values:

Item	Value to be entered or selected
Name	lk4l-vnet
Address space	10.3.0.0/16
<b>Subnet</b>	
Name	lk4l-nw01
Address range	10.3.1.0/24
Firewall	Disabled (default)

2. Select **Azure Portal > Virtual networks > (Virtual network name) > Subnets > Add** to add a second subnet for your virtual network.

Dashboard > Virtual networks > Ik4I-vnet-rtanaka - Subnets > Add subnet

### Add subnet

Ik4I-vnet-rtanaka

\* Name

\* Address range (CIDR block) ⓘ  
  
 10.3.0.0 - 10.3.0.255 (251 + 5 Azure reserved addresses)

---

Network security group  
 None >

---

Route table  
 None >

---

Service endpoints  
 Services ⓘ

---

Subnet delegation

**OK**

Enter the required information and click **OK**.

Use the following values:

Item	Value to be entered or selected
Name	Ik4I-nw02
Address range	10.3.2.0/24

## 12.2.4.3. Creating a Virtual Machine

---

Create a virtual machine to use in this environment.

The following three virtual machines are created.

- [Cluster Node \(Active\)](#)
- [Cluster Node \(Standby\)](#)
- [Client and Witness Server](#)

# 12.2.4.3.1. Creating a Cluster Node (Active)

1. Select **Azure Portal** > **Virtual machines** > **Add** and enter the basics settings.

Home > Virtual machines >

## Create a virtual machine

Subscription \* ⓘ Visual Studio Premium with MSDN

Resource group \* ⓘ LK\_SPS\_Example  
[Create new](#)

**Instance details**

Virtual machine name \* ⓘ lk4lnode01 ✓

Region \* ⓘ (US) East US

Availability options ⓘ Availability set

Availability set \* ⓘ No existing availability sets in current resource group and location.  
[Create new](#)

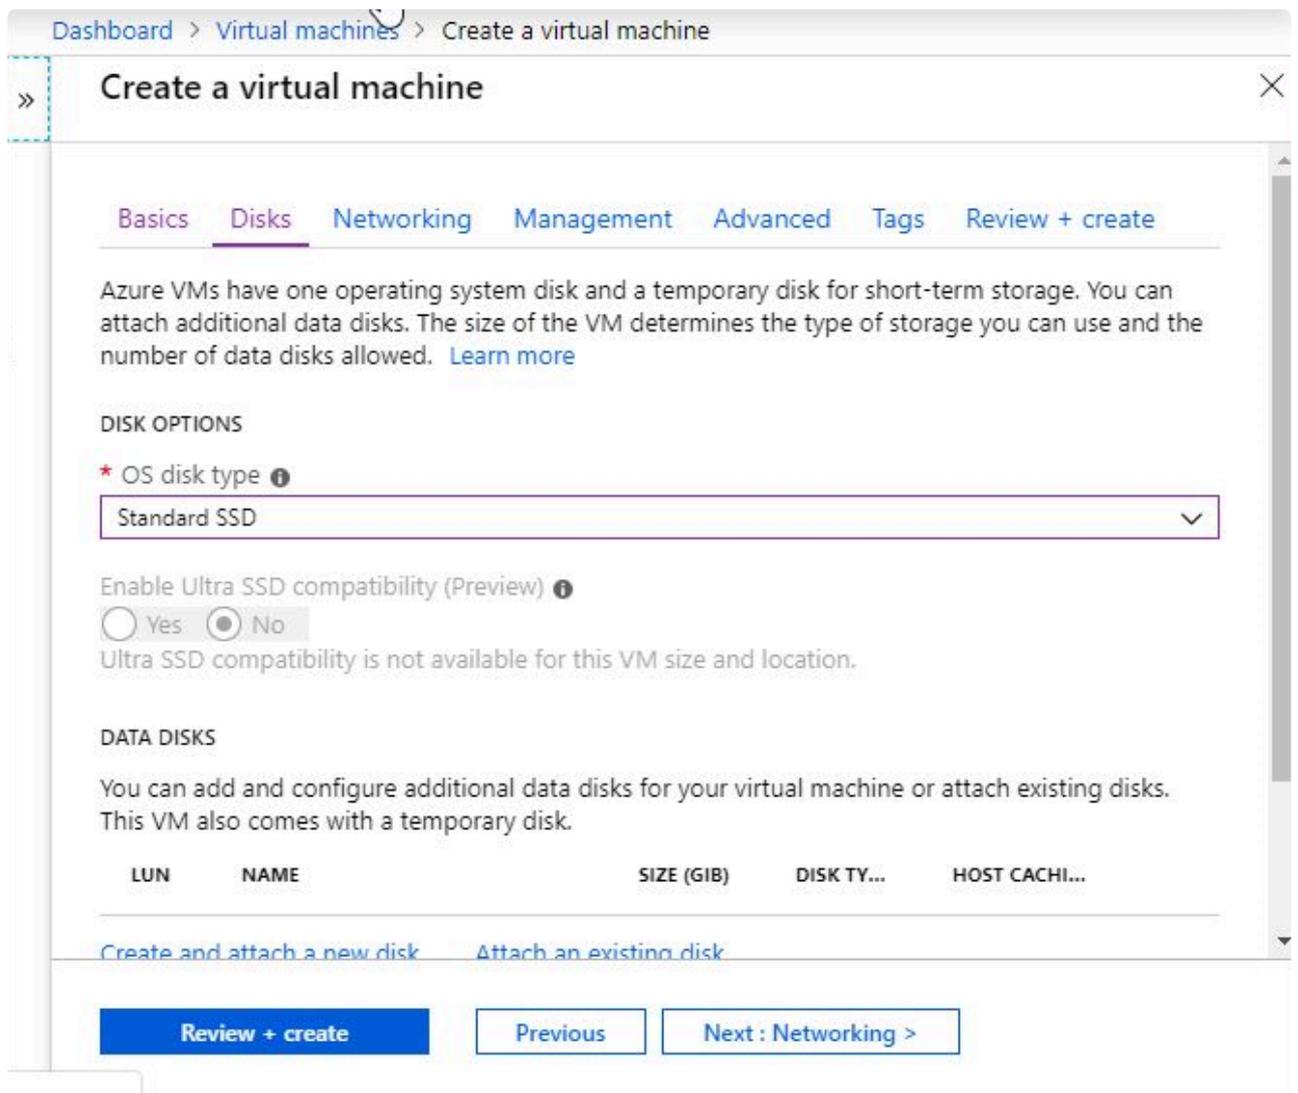
Use the following values:

Basics		
Item	Value to be entered or selected	Notes
<b>INSTANCE DETAILS</b>		
Virtual machine name	lk4lnode01	
Availability options	Availability set	*Note 1
	(create new)lk4lcluster	
Image	Red Hat Enterprise Linux 8.1	*Note 2
Size	Standard A1	*Note 3 *Note 4
<b>ADMINISTRATIVE ACCOUNT</b>		
Authentication type	Password	
Username	lkadmin	
Password	XXXXXXXX	
<b>INBOUND PORT RULES</b>		
Public inbound ports	Allow selected ports	
Select inbound ports	SSH	

- **Note 1:** **Availability Sets are a prerequisite for LifeKeeper clusters on Azure and must be set up.**
- **Note 2:** Select the OS version supported by LifeKeeper to be installed. Refer to the [Support Matrix](#) OS versions supported by LifeKeeper.
- **Note 3:** Be sure to select an instance size for the Availability Set. Availability Sets are not available for some instance sizes.
- **Note 4:** Be sure to select an instance size for the load balancing rules that can be applied. Load balancing rules are not available for some instance sizes.

2. Once the required values are entered click **Next:Disks**.
3. Configure the disk settings. In this case, a disk for Oracle DB (30GiB) and a disk for swap area (10GiB) are required in addition to the OS disk.

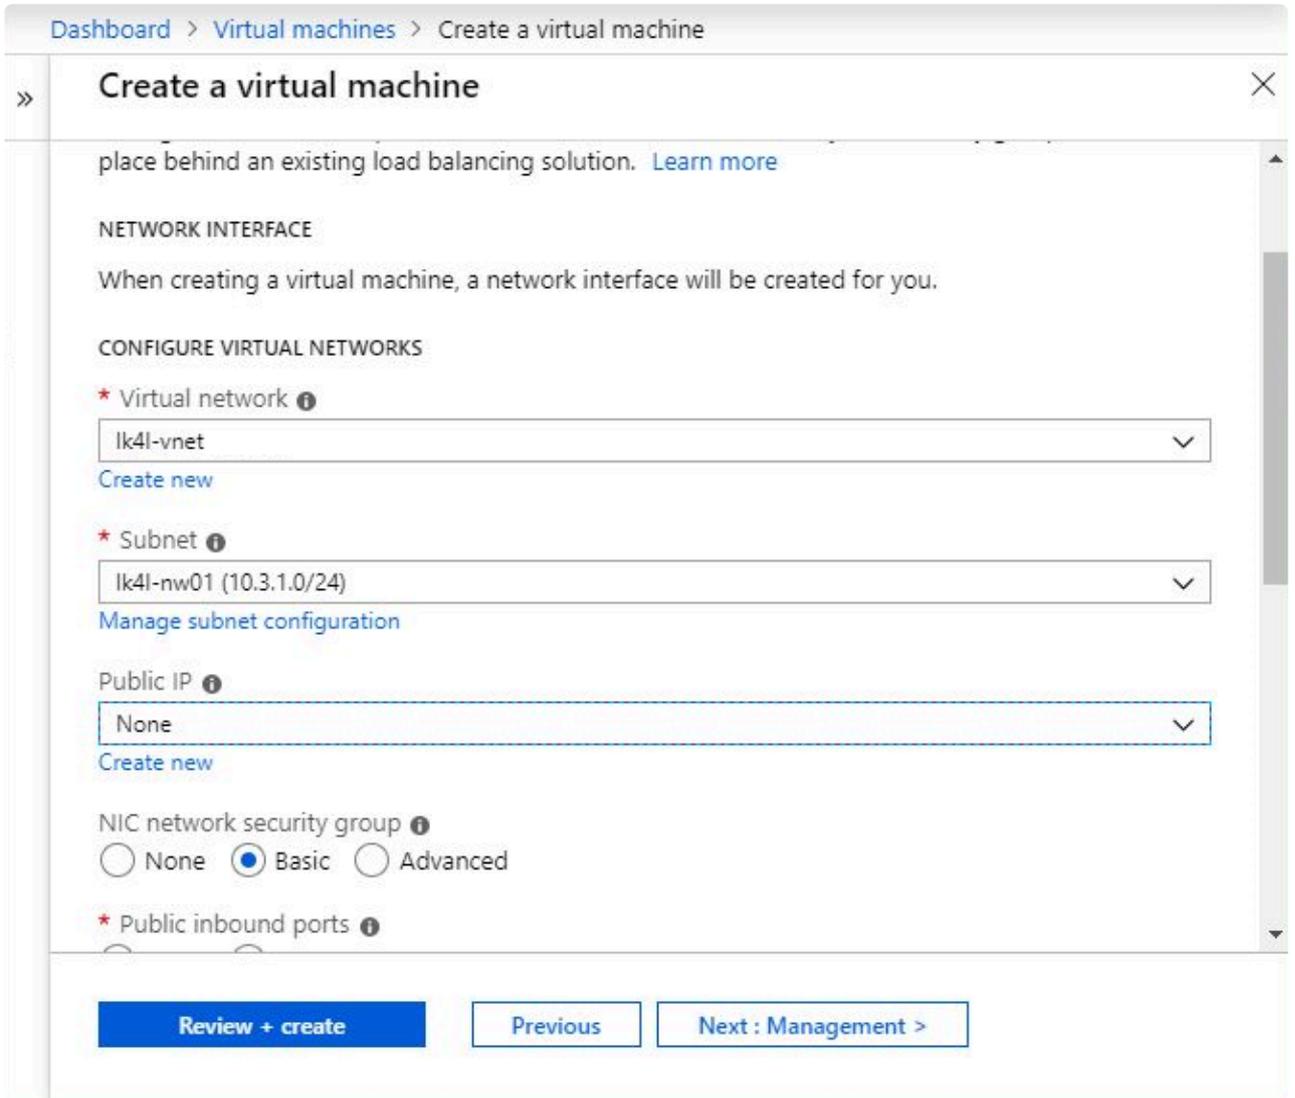
Click **Create and attach a new disk** to assign additional disks.



Use the following values:

Disks		
Item	Value to be entered or selected	Notes
<b>DISK OPTIONS</b>		
OS Disk type	Standard SSD	
<b>DATA DISKS</b>		
LUN	0	
NAME	(create and attach a new disk) Iklnode01_DataDisk_0	
Size(GB)	30	
DISK TYPE	Standard SSD	
HOST CACHING	None	*Note 1
<b>DATA DISKS</b>		
LUN	1	
NAME	(create and attach a new disk) Iklnode01_DataDisk_1	Any value
Size(GB)	10	
DISK TYPE	Standard SSD	
HOST CACHING	None	*Note1

- **Note 1:** In this case, the cache is disabled for database use and should be configured appropriately according to your requirements.
4. Once the required values are entered click **Next:Networking**.
  5. Enter the network settings.



At this time, configure the first subnet that you created in the previous step for the virtual machine. This setting attaches a single network interface. The second network interface is configured after the virtual machine is created.

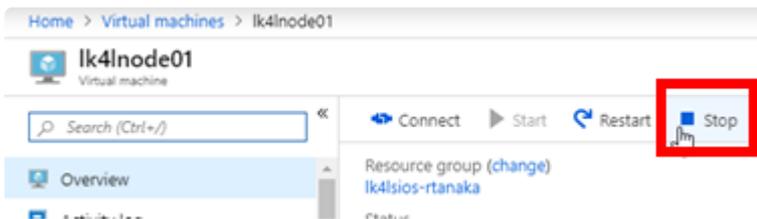
Use the following values:

Networking		
Item	Value to be entered or selected	Notes
Virtual network	lk4l-vnet	
Subnet	lk4l-nw01(10.3.1.0/24)	*Note 1
Public IP	(create new) Name:lk4lnode01pubip SKU:Basic Assignment:Static	*Note 2

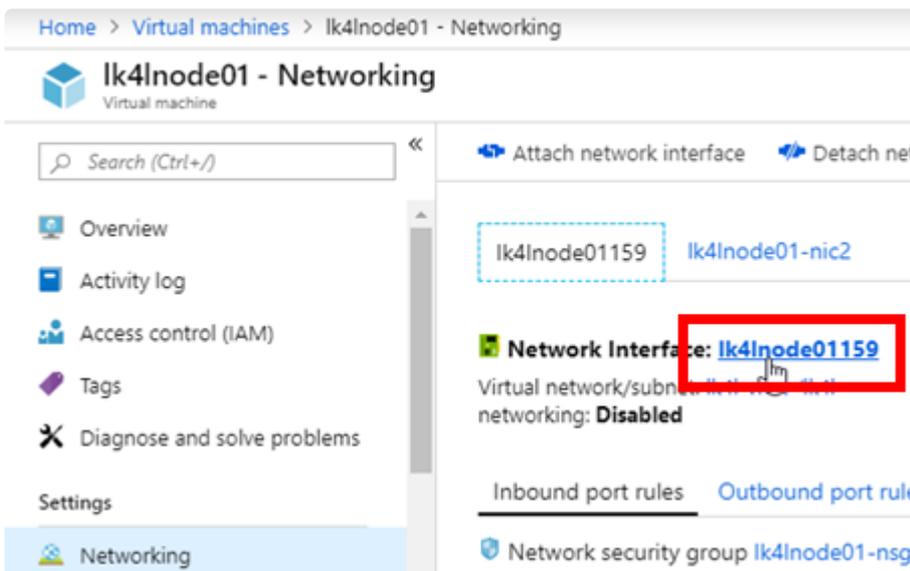
NIC network security group	Basic	
Public inbound ports	Allow selected ports	
Selected inbound ports	SSH	

- **Note 1:** Configure the first subnet that you created in the previous step for the virtual machine.
- **Note 2:** In this procedure the Public IP has been set to access via the Internet.

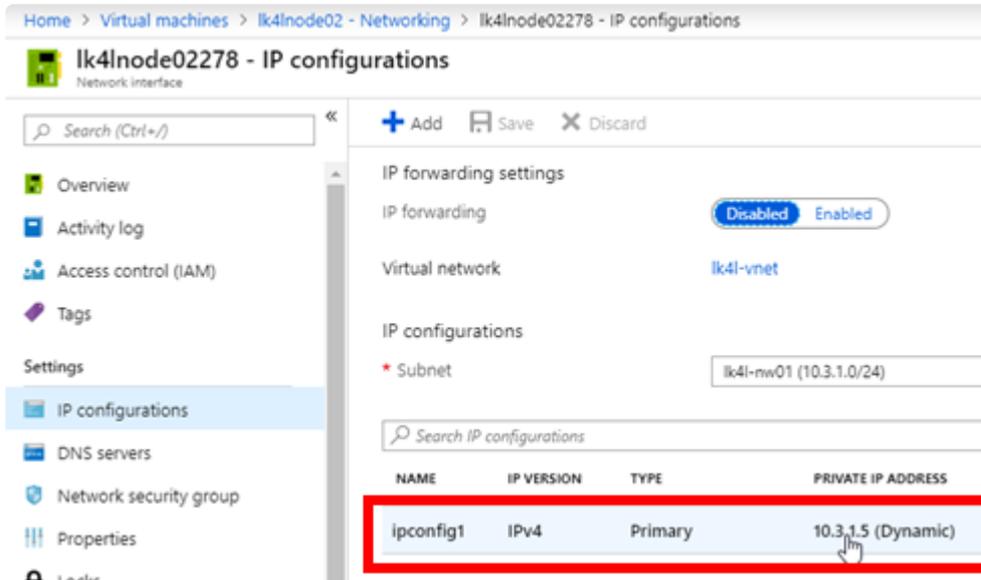
6. Configure administrative settings, advanced settings, and tag settings if necessary, then click **Review + Create**.
7. Review the details and click **Create**. When deployed with the above configuration, Red Hat Enterprise Linux starts with the installation completed.
8. Stop the virtual machine to change the network configuration.



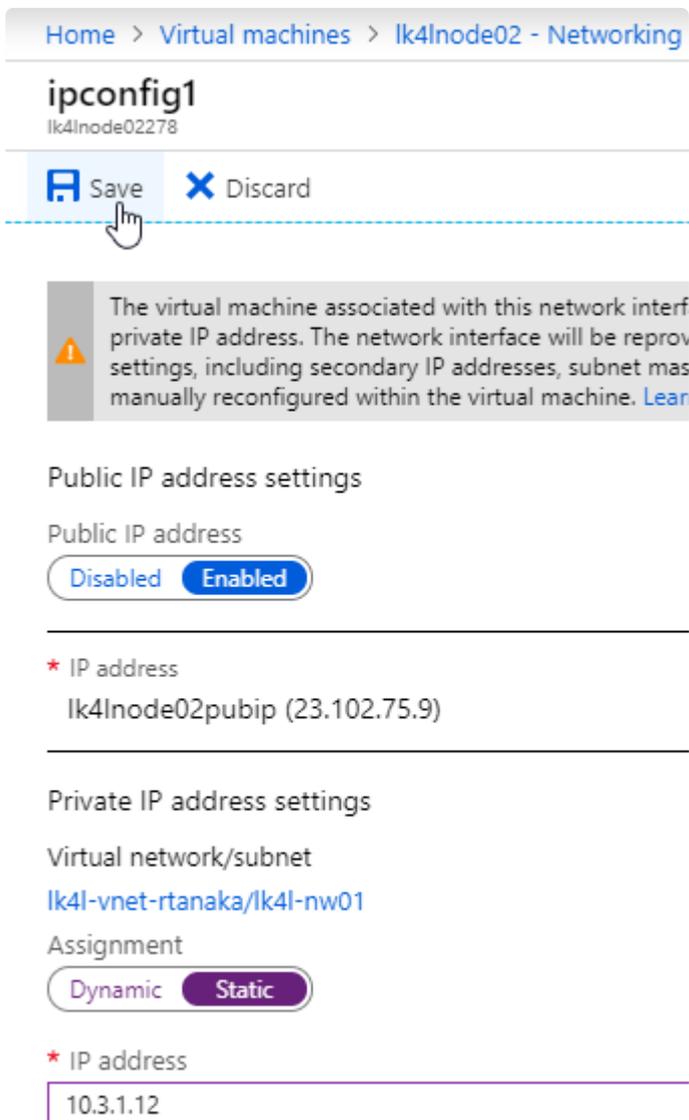
9. Change the first network interface to a static IP address. Click **Azure Portal > Virtual machines > (virtual machine name) > Networking > (first interface name)**.



10. Click **IP configurations > ipconfig1**. At this point the IP address is set to Dynamic.



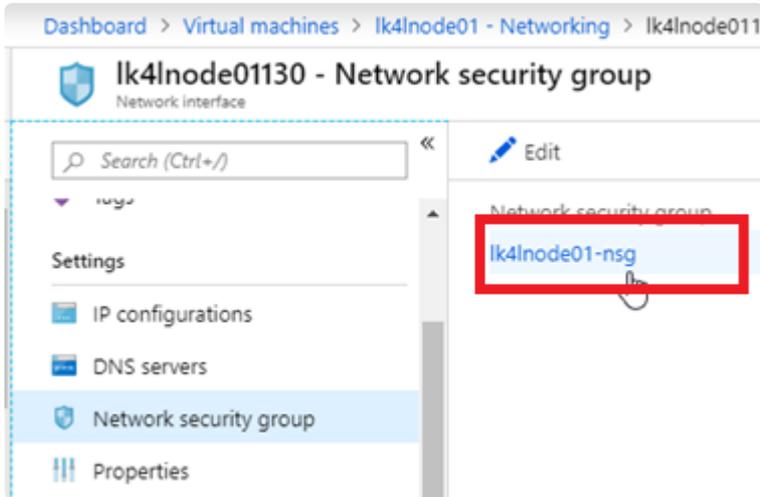
11. Change the IP address setting to static.



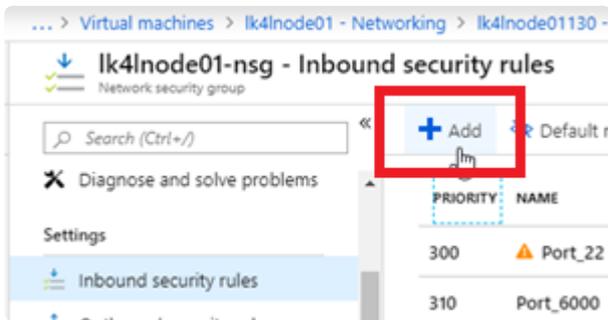
Use the following values:

Item	Value to be entered or selected	Notes
Assignment	Static	
IP address	10.3.1.11	

- Click **Save** to save the settings.
- Edit the network security group to allow the Oracle listener service to communicate using TCP Port 1521. Click **Azure Portal > Virtual machines > (virtual machine name) > Networking > (first interface name) > Network security group > (security group name)**.



- Click **Inbound security rules > Add**.



- Add inbound rules.

## Add inbound security rule

lk4lnode01-nsg

Basic

---

**\* Source** ⓘ

**\* Source port ranges** ⓘ

**\* Destination** ⓘ

**\* Destination port ranges** ⓘ

**\* Protocol**

Any
TCP
UDP

**\* Action**

Allow
Deny

**\* Priority** ⓘ

**\* Name**

---

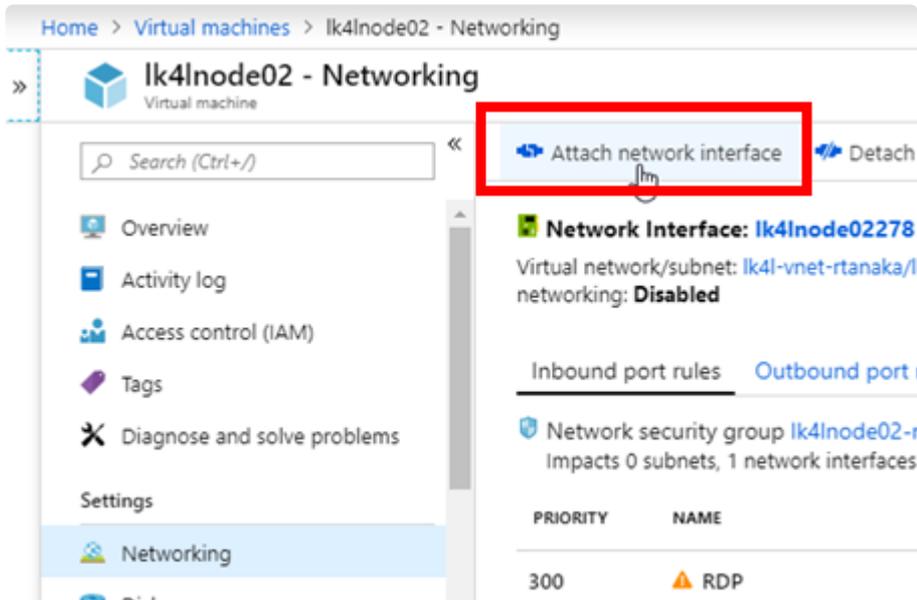
Add

Use the following values:

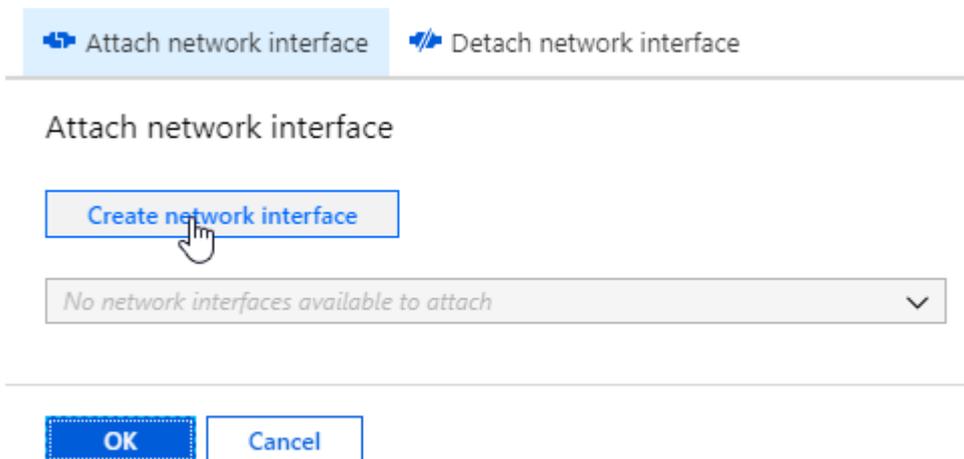
Item	Value to be entered or selected	Notes
Source	Any	
Source port ranges	•	
Destination	Any	
Destination port ranges	1521	*Note 1
Protocol	TCP	
Action	Allow	
Name	Port_1521	

- **Note 1:** Since the Oracle listener uses port1521, communication to the relevant port is permitted.

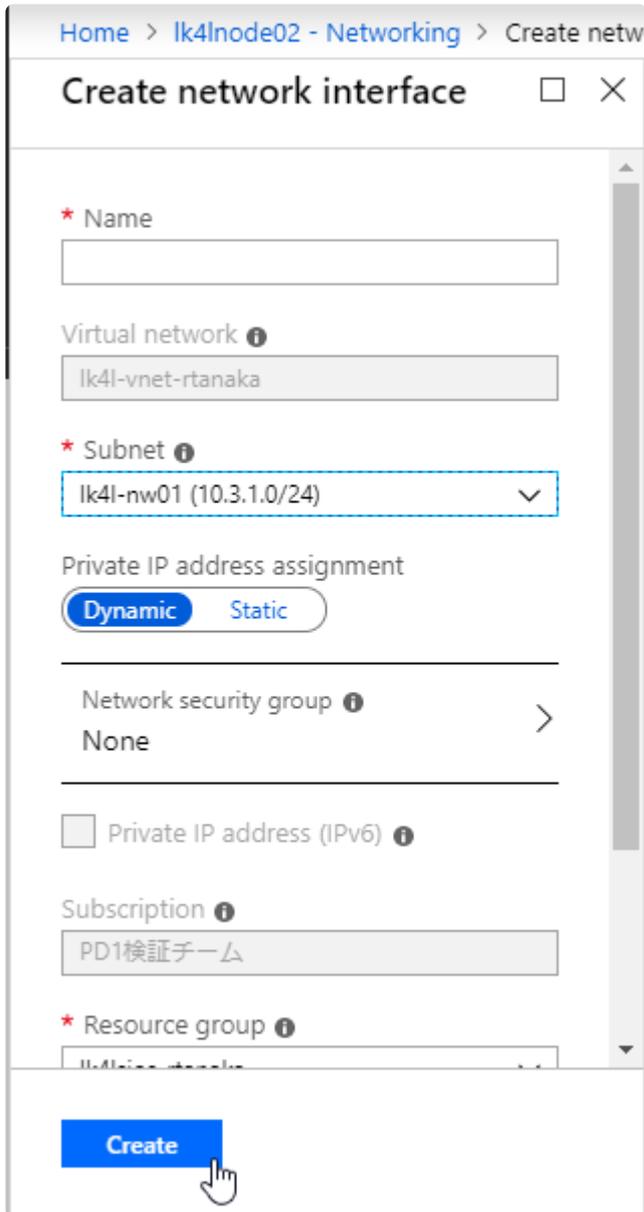
16. Click **Add** to save your settings.
17. Next, add a second network interface. Click **Azure Portal > Virtual machines > (Virtual machine name) > Networking > Attach network interface**.



18. Click **Create network interface**.



19. After entering the required values, click **Create**.



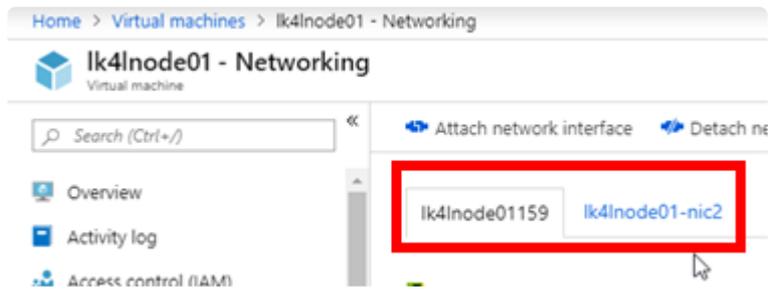
Use the following values:

Item	Value to be entered or selected	Notes
Name	lk4lnode02-nic2	
Subnet	10.3.2.0/24	*Note 1
Private IP address assignment	Static	
Private IP address	10.3.2.11	
Network security group	None	

- **Note 1:** Set the second subnet created in the previous step for the virtual machine.

20. Click **OK**.

21. Make sure that the two interfaces are attached.



22. Start the virtual machine.

## 12.2.4.3.2. Creating a Cluster Node (Standby)

1. Follow the same procedure to create a virtual machine that was used for the cluster nodes (standby node).

Use the following values for basic settings:

Basics		
Item	Value to be entered or selected	Notes
<b>INSTANCE DETAILS</b>		
Virtual machine name	lk4lnode02	

All other settings should be the same as the active node.

Use the following values for disk settings:

Disks		
Item	Value to be entered or selected	Notes
<b>DATA DISKS</b>		
LUN	0	
NAME	(create and attach a new disk) lklnode02_DataDisk_0	
<b>DATA DISKS</b>		
LUN	1	
NAME	(create and attach a new disk) lklnode02_DataDisk_1	Any value

All other settings should be the same as the active node.

Use the following values for network settings:

Networking		
Item	Value to be entered or selected	Notes
Public IP	(create new)	*Note 1

	Name:lk4Inode02pubip	
	SKU:Basic	
	Assignment:Static	

**Note 1:** In this procedure the Public IP has been set to access via the Internet.

All other settings should be the same as the active node.

2. Make administrative settings, advanced settings, and tagging if necessary, and then click **Review + Create**.
3. Review the details and click **Create**.
4. Since the network configuration will be changed after this, stop the virtual machine.
5. Change the first network interface to a static IP address. Click **Azure Portal > Virtual machines > (Virtual machine name) > Networking > (First interface name)**.
6. Click **IP configurations > ipconfig1**. At this point the IP address is set as Dynamic.
7. Set the IP address to Static.

Use the following values:

Item	Value to be entered or selected	Notes
Assignment	Static	
IP address	Static 10.3.1.12	

8. Click **Save** to save the settings.
9. Next, add a second network interface. Click **Azure Portal > Virtual machines > (Virtual machine name) > Networking > Attach network interface**.
10. Click **Create network interface**.
11. Click **Create** after entering the required values.

Use the following values:

Item	Value to be entered or selected	Notes
Name	lk4Inode02-nic2	
Subnet	10.3.2.0/24	*Note 1

Private IP address assignment	Static	
Private IP address	10.3.2.12	
Network security group	None	

- **Note 1:** Set the second subnet created in the previous step for the virtual machine.

12. Click **OK**.
13. Make sure that the two interfaces are attached.
14. Start the virtual machine.

## 12.2.4.3.3. Creating a Client and Witness Server

1. Follow the same procedure that was used for the cluster nodes (standby node).

Use the following values for basic settings:

Basics		
Item	Value to be entered or selected	Notes
<b>INSTANCE DETAILS</b>		
Virtual machine name	lk4lclient	
Availability options	No Infrastructure redundancy required	*Note 1
Image	Red Hat Enterprise Linux 7.5	
Size	Basic AO	
<b>INBOUND PORT RULES</b>		
Public inbound parts	Allow selected ports	
Select inbound ports	SSH	

- **Note 1:** It is not necessary to set the Availability Set except for the cluster nodes. If the client is used for Quorum you should consider including it in the Availability Set for high availability.

After entering the required values, click **Next: Disks**.

2. Enter the disk settings.

A disk is not added at this time. Click **Next: Networking**.

3. Enter the network settings.

Use the following values:

Networking		
Item	Value to be entered or selected	Notes
Virtual network	lk4l-vnet	
Subnet	lk4l-nw01(10.3.1.0/24)	*Note 1
Public IP	(create new)	*Note 2

	Name:lk4lclientpubip SKU:Basic Assignment:Static	
NIC network security group	Basic	
Public inbound ports	Allow selected ports	
Selected inbound ports	SSH	

- **Note 1:** Set the first subnet created in the previous step for the virtual machine.
- **Note 2:** In this procedure, the Public IP is set because it is accessed via the Internet.

4. Make administrative settings, advanced settings, and tagging if necessary, and then click **Review + Create**.
5. Review the details and click **Create**.

After the virtual machine is started, change the first network interface to a static IP address. Click **Azure Portal > Virtual machines > (virtual machine name) > (Networking) > (first interface name)**.

6. Click **IP configurations > ipconfig1**. At this point the IP address is set to Dynamic.
7. Change the setting of the IP address to Static.

Use the following values:

Item	Value to be entered or selected	Notes
Assignment	Static	
IP address	10.3.1.50	

8. Click **Save** to save the settings.

Now three virtual machines have been created.

Home > Virtual machines

### Virtual machines

既定のディレクトリ

[+](#) Add   [🕒](#) Reservations   [☰](#) Edit columns   [🔄](#) Refresh   [🏷️](#) Assign tags

**Subscriptions:**

3 items

<input type="checkbox"/>	NAME ↑↓	TYPE ↑↓	STATUS	RESOURCE
<input type="checkbox"/>	lk4lclient	Virtual machine	Running	lk4lsio
<input type="checkbox"/>	lk4lnode01	Virtual machine	Running	lk4lsio
<input type="checkbox"/>	lk4lnode02	Virtual machine	Running	lk4lsio

## 12.2.4.4. Creating a Load Balancer

On Azure the virtual IP created by LifeKeeper cannot be used for communication. A load balancer is required for the virtual IP to work.

There are two types of Azure load balancers:

- External load balancer
- Internal load balancer

Here we will use the internal load balancer. The IP address of this internal load balancer is also used as the virtual IP address that will be set later. It will become the access point of the protected services.

1. Select **Azure Portal > Load balancers > Add** to create a load balancer used in this environment. Enter the required values and click **Review + Create**.

Use the following values:

Item	Value to be entered or selected	Notes
<b>INSTANCE DETAILS</b>		
Name	lk4lsiosilb	
Type	Internal	*Note 1
SKU	Basic	
<b>PUBLIC IP ADDRESS</b>		
Virtual network	lk4l-vnet	
Subnet	lk4l-nw01(10.3.1.0/24)	
IP address assignment	Static	
Private IP address	10.3.1.200	*Note 2

- **Note 1:** Be sure to select "Internal" to create an internal load balancer.
- **Note 2:** This value is also used as the value of the virtual IP address that will be set later. It becomes the access point of the protected services.

2. Review the details and click **Create**.
3. Next, add two virtual machines for the created cluster to the backend pool of the internal load balancer. Select **Azure Portal > Load balancers > (load balancer name) > Backend pools > Add**.

lk4l All resource gr

1 items

<input checked="" type="checkbox"/>	NAME ↑↓	RESOURC
<input checked="" type="checkbox"/>	 <a href="#">lk4lsiosilb</a>	lk4lsios

Home > Load balancers > lk4lsiosilb - Backend pools

### lk4lsiosilb - Backend pools

Load balancer

Search (Ctrl+/) << **+ Add** Refresh

- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems

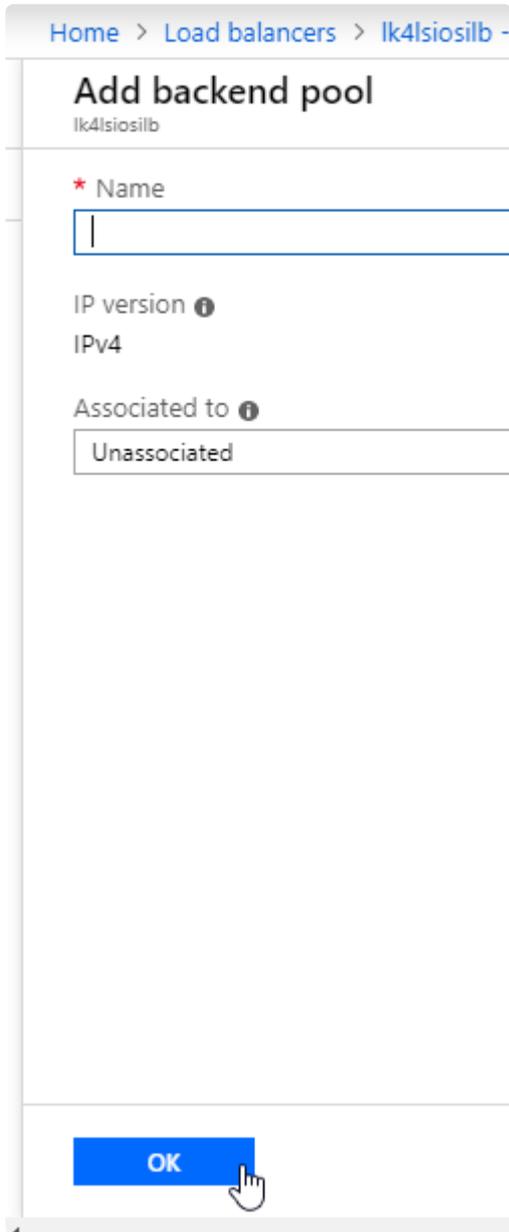
Settings

- Frontend IP configuration
- Backend pools**

Search backend address

VIRTUAL MACHINE

No results.



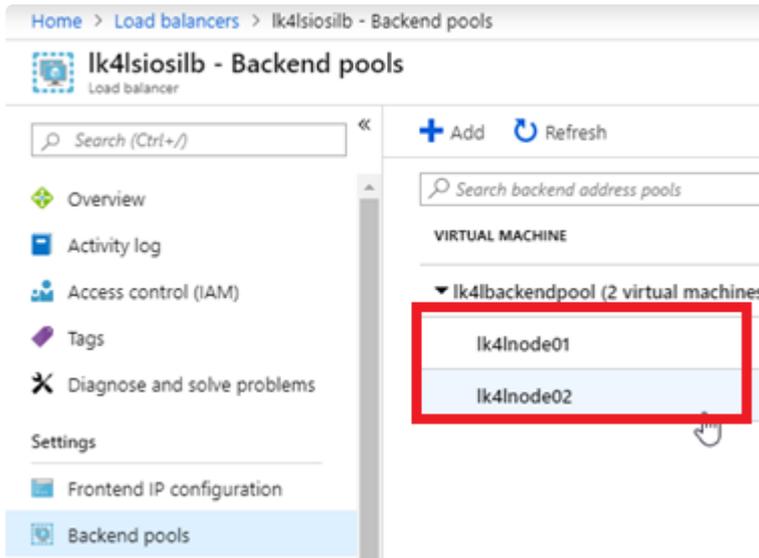
Use the following values:

Item	Value to be entered or selected	Notes
Name	lk4lbackendpool	
Associated to	Availability set	
Availability set	lk4lcluster	*Note 1
<b>Target #1</b>		
Target virtual machine	lk4lnode01	
Network IP configuration	ipconfig1(10.3.1.11)	*Note 2
<b>Target #2</b>		
Target virtual machine	lk4lnode02	
Network IP configuration	ipconfig1(10.3.1.12)	*Note 2

- **Note 1:** Specify the created Availability Set.
- **Note 2:** Associate the backend pool with the primary interface of the virtual machine in the created Availability set.

Enter the required values and click **OK**.

4. Verify that the backend pool has been created.



5. Configure the probe. Click **Azure Portal > Load balancers > (load balancer name) > Health Probes > Add**.

Dashboard > Load balancers > lk4lsiosilb - Health

### Add health probe

lk4lsiosilb

\* Name

IP version  
IPv4

Protocol ⓘ

\* Port ⓘ

\* Interval ⓘ

\* Unhealthy threshold ⓘ

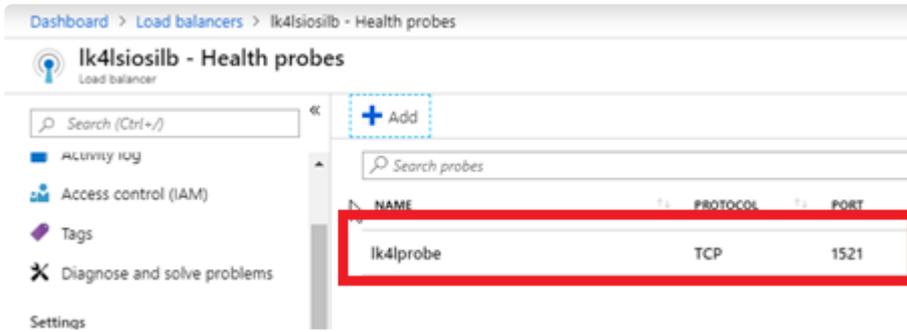
**OK**

Use the following values:

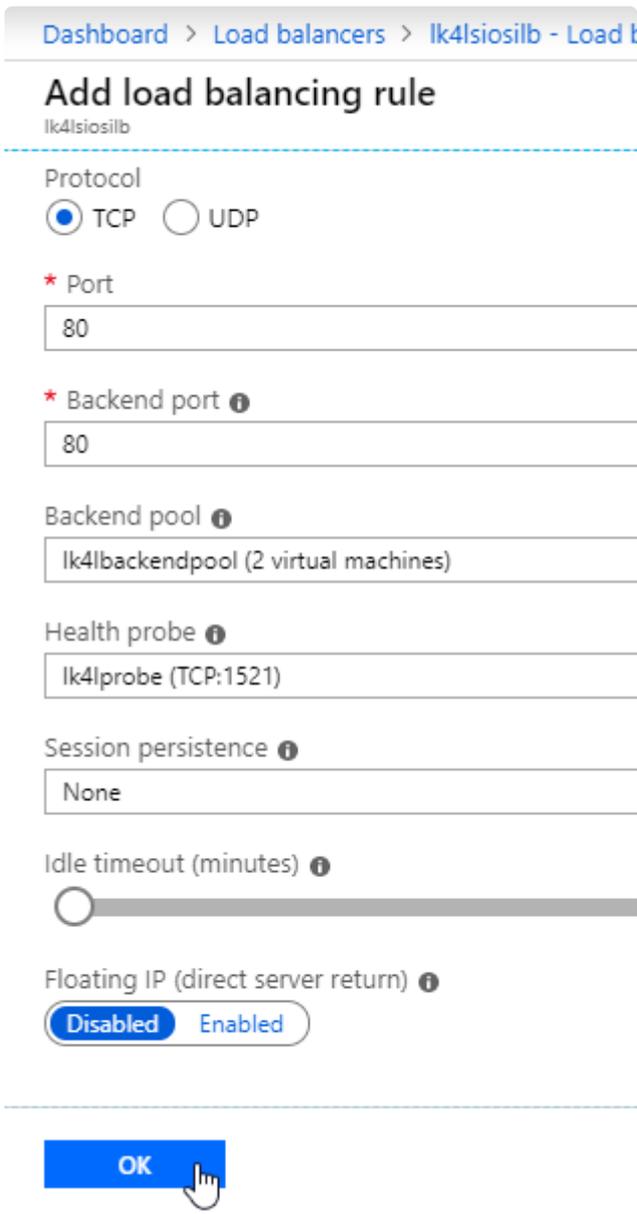
Item	Value to be entered or selected	Notes
Name	lk4lprobeport	
Protocol	TCP	*Note 1
Port	1521	*Note 1

- **Note 1:** For the probe, specify the protocol and port used by the Oracle listener.

6. Enter the required values and click **OK**.
7. Verify that the probe has been created.



- Configure the load balancing rules. Click **Azure Portal > Load balancers > (load balancer name) > Load balancing rules > Add**.



Use the following values:

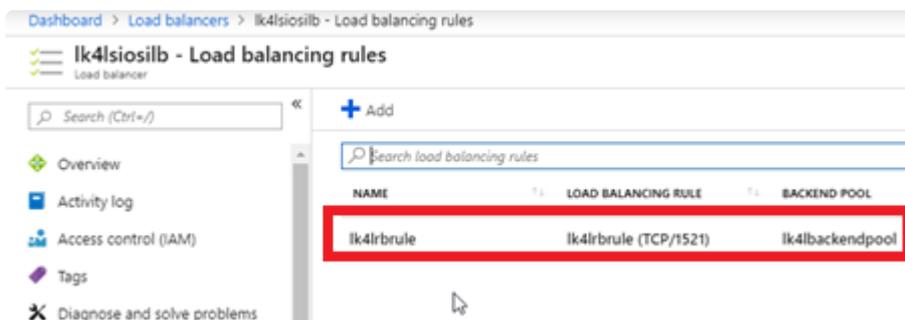
Item	Value to be entered or selected	Notes
Name	Ik4lrbrule	

Protocol	TCP	
Port	1521	*Note 1
Backend port	1521	*Note 1
Backend pool	lk4lbackendpool	
Health probe	lk4lprobe	
Floating IP address	Enabled	*Note 2

- **Note 1:** For the port for load balancing, specify the port used by the Oracle listener.
- **Note 2:** Since the communication from the client to the cluster must be a packet communication to an IP resource (VIP), the Floating IP address must be enabled.

9. Enter the required values and click **OK**.

10. Verify that the load balancing rules are created.



Now the load balancer has been configured.

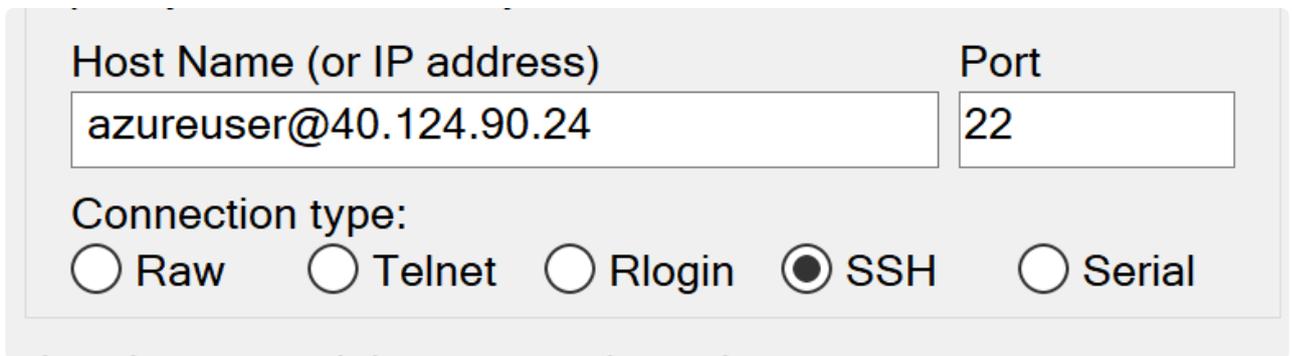
## 12.2.4.5. Configuring the OS

Prior to installing LifeKeeper, configure the OS for the three virtual machines that you have created.

### Logging In to a Virtual Machine

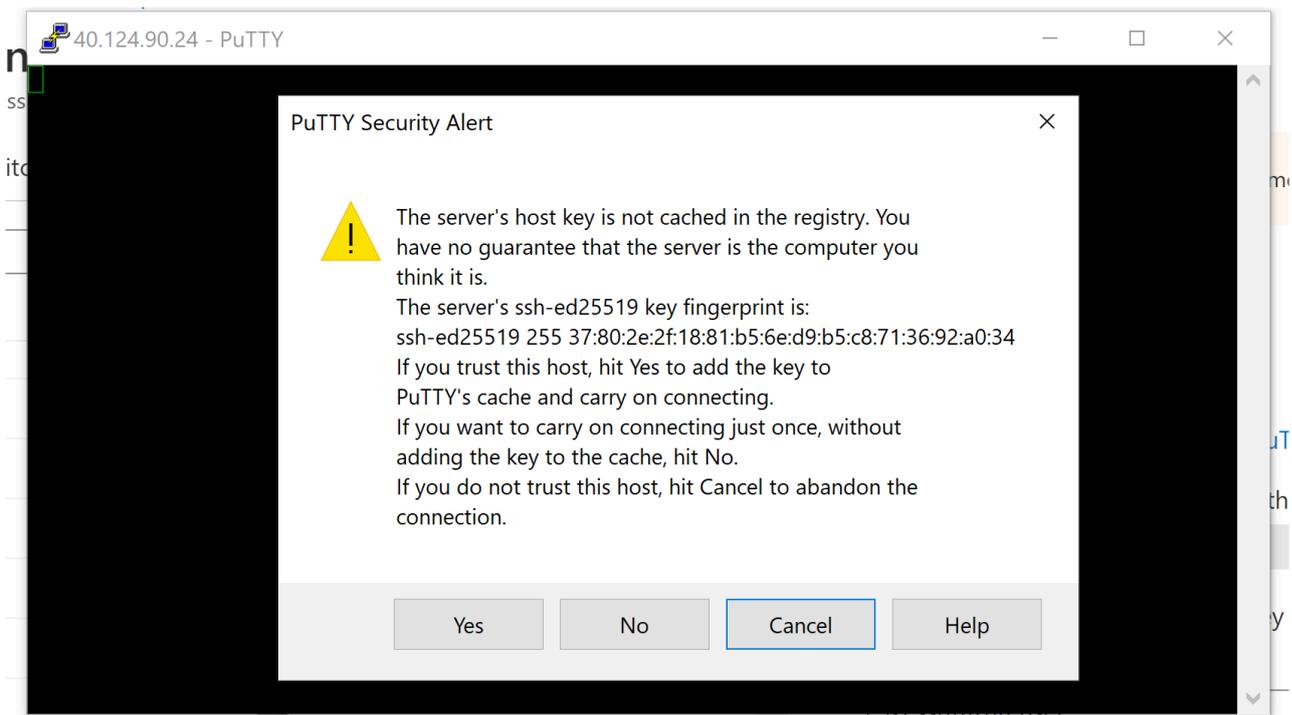
Perform the following steps on **both cluster nodes and clients**.

1. Select **Azure Portal > Virtual machines > Overview > Connect** and review the information required for the connection.
2. Using terminal software (Putty, Tera Term, etc.) on the local PC, SSH to the public IP of the virtual machine with the account created when creating the virtual machine.



The screenshot shows the PuTTY configuration dialog box. The 'Host Name (or IP address)' field contains 'azureuser@40.124.90.24' and the 'Port' field contains '22'. Under 'Connection type:', the 'SSH' radio button is selected, while 'Raw', 'Telnet', 'Rlogin', and 'Serial' are unselected.

3. Log in to the virtual machine. (The screenshot below displays what you'll see when you login to PuTTY for the **first time**. Click **Yes** to Proceed.)



4. Set it up for subsequent use by root. Change to root privileges.

```
$ sudo su -
```

- Set the root password.

```
# passwd
```

- Allow root login.

```
# cp /etc/ssh/sshd_config /etc/ssh/sshd_config.org
# vi /etc/ssh/sshd_config
< omitted >
PermitRootLogin yes //Uncomment
< omitted >
#systemctl restart sshd
```

- Launch another window with the terminal software and confirm that you can log in as root (**do not close the terminal window that you have been using**).
- After successfully logging in, perform all subsequent tasks as root.

## Fixing NIC Addresses

The NIC address is fixed to prevent it from being updated when the Azure infrastructure is updated. Make the following settings on **both cluster nodes**.

- Display the NIC information of the virtual machine and take a note of each interface name and MAC address.

```
#ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.3.1.11 netmask 255.255.255.0 broadcast 10.3.1.255
    inet6 fe80::20d:3aff:fe50:42bf prefixlen 64 scopeid 0x20<link>
    ether 00:0d:3a:50:42:bf txqueuelen 1000 (Ethernet)
<omitted>
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.3.2.11 netmask 255.255.255.0 broadcast 10.3.2.255
    inet6 fe80::e3c9:ec16:62a:bbd0 prefixlen 64 scopeid 0x20<link>
    ether 00:0d:3a:51:fc:01 txqueuelen 1000 (Ethernet)
<omitted>
```

- By default, there is only a network configuration file for eth0, so create a network configuration file for eth1.

```
#cp /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-eth1
```

- Edit the eth0 and eth1 configuration files. For HWADDR, use the value you took a note in the previous step.

```
#vi /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
ONBOOT=yes
<omitted>
HWADDR=00:0d:3a:50:42:bf //Add

#vi /etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE=eth1 //change
ONBOOT=yes
<omitted>
HWADDR=00:0d:3a:51:fc:01 //Add
```

## SELinux

Disable SELinux. Make the following settings on **both cluster nodes**.

1. Open the following file.

```
# vi /etc/selinux/config
```

2. Change the following parameters from “enforcing” to “disabled” and save.

```
SELINUX=disabled
```

3. Reboot the virtual machine for the settings to take effect.
4. Confirm that SELinux is disabled.

```
# getenforce
Disabled
```

## Firewall

Since LifeKeeper uses a specific port, refer to [Running LifeKeeper with a Firewall](#) to change the OS firewall configuration.

Disable the firewall. Configure the following settings on **both cluster nodes**.

- Execute the following commands.

```
# systemctl stop firewalld.service
# systemctl disable firewalld.service
```

## Address Conversion

Configure the following settings at the same time.

- To enable communication to the VIP, the destination address of the packet for port 1521 where the load is balanced by ILB is converted to [VIP protected with IP resources].
- Convert the source address of the ICMP packet to [Private IP address of Azure] so that the monitoring process of IP resources can be performed successfully.

Make the following settings on **both cluster nodes**.

1. Install the IP tables package.

```
# yum -y install iptables-services
```

2. Enable IP tables.

```
#systemctl start iptables.service
#systemctl enable iptables.service
```

3. Make sure the IP tables are enabled.

```
#systemctl list-unit-files | grep iptables
iptables.service                enabled
```

4. Configure the IP tables.

Execute the following command on the active node (lk4lnode1):

```
# iptables -t nat -A PREROUTING -p tcp --dport 1521 -j DNAT --to-destination
10.3.1.200:1521
# iptables -t nat -A POSTROUTING -p icmp -s 10.3.1.200 -j SNAT --to-source 10.3.1.11
# service iptables save
```

Execute the following command on the standby node (lk4lnode02):

```
# iptables -t nat -A PREROUTING -p tcp --dport 1521 -j DNAT --to-destination
10.3.1.200:1521
# iptables -t nat -A POSTROUTING -p icmp -s 10.3.1.200 -j SNAT --to-source 10.3.1.12
# service iptables save
```

5. Confirm that the ip table settings have been added. The following is an example of the active node.

```
#cat /etc/sysconfig/iptables
<omitted>
:POSTROUTING ACCEPT [16:960]
-A PREROUTING -p tcp -m tcp --dport 1521 -j DNAT --to-destination 10.3.1.200:1521
-A POSTROUTING -s 10.3.1.200/32 -p icmp -j SNAT --to-source 10.3.1.11
<omitted>
```

6. Reboot the virtual machine for the settings to take effect.

## GUI Connection Settings

By default, Azure virtual machines do not have a GUI environment installed. We need to set up an X11 environment on the cluster nodes and a GUI environment on the client to use the LifeKeeper GUI.

Make the following settings on **both cluster nodes**.

7. Install the GUI package. Execute the following command.

```
# yum -y groupinstall "X11"
```

8. After installing the package, execute the following command.

```
# systemctl set-default graphical.target
```

9. Next, set the permission for intra-subnet communication. Edit the following file and add an entry before [REJECT].

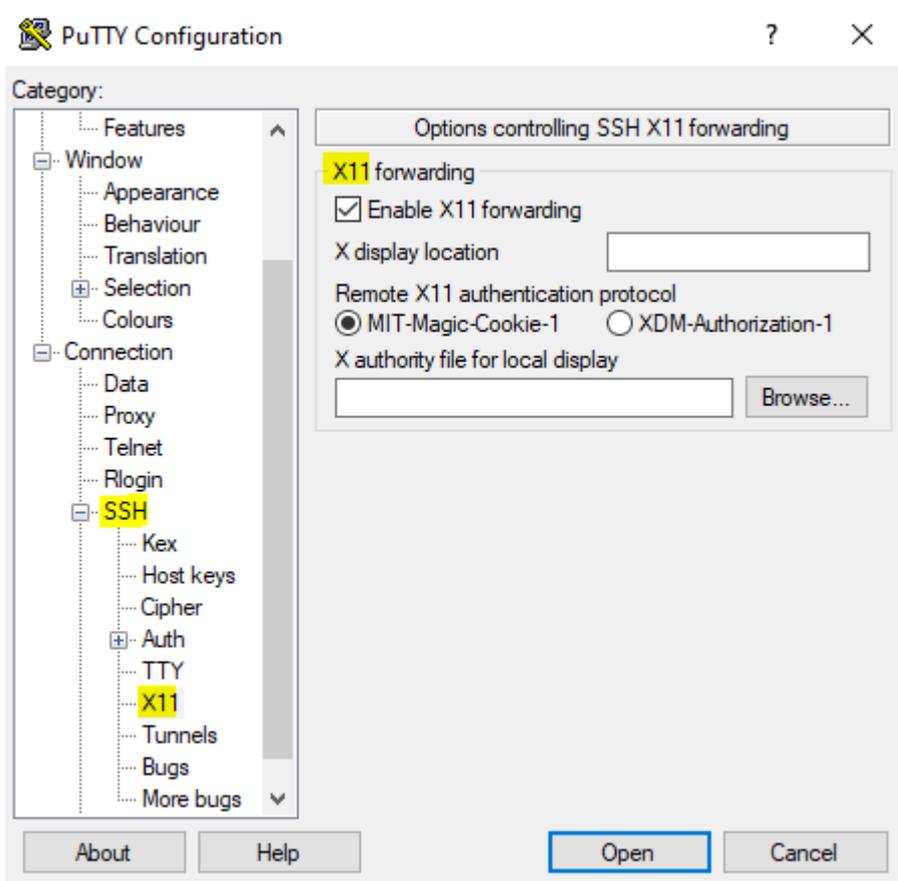
```
# vi /etc/sysconfig/iptables
<omitted>
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [4710:957955]
-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT
-A INPUT -s 10.3.1.0/24 -p tcp -j ACCEPT //Add
-A INPUT -s 10.3.2.0/24 -p tcp -j ACCEPT //Add
-A INPUT -j REJECT --reject-with icmp-host-prohibited
<omitted>
```

10. Reboot the virtual machine for the settings to take effect.

```
# reboot
```

11. Configure the local PC. Install the X Server software (such as Xming) on your local PC and start it.
12. Next, configure the X 1 1 forwarding in the local PC terminal software (such as Putty or TeraTerm).

The setting method varies depending on the terminal software you are using. The following is an example of the configuration in PuTTY.



## Name Resolution

Configure the following settings on **both cluster nodes and clients**.

1. Register the information of the cluster node (active) and cluster node (standby) clients and three Witness servers in */etc/hosts*.

```
#vi /etc/hosts
10.3.1.11 lk4lnode01 //Add
10.3.1.12 lk4lnode02 //Add
10.3.1.50 lkclient //Add
```

2. Verify that each of them can communicate with each other with host name.

```
[root@lk4lnode01 ~]# ping lk4lnode02 -c 4  
[root@lk4lnode01 ~]# ping lk4lclient -c 4
```

```
[root@lk4lnode02 ~]# ping lk4lnode01 -c 4  
[root@lk4lnode02 ~]# ping lk4lclient -c 4
```

```
[root@lk4lclient ~]# ping lk4lnode01 -c 4  
[root@lk4lclient ~]# ping lk4lnode02 -c 4
```

## Checking the Kernel Version

Check the kernel version of RHEL.

```
# uname -r  
3.10.0-862.11.6.el7.x86_64  
# cat /etc/redhat-release  
Red Hat Enterprise Linux Server release 7.5 (Maipo)
```

For each distribution that supports LifeKeeper, refer to the [LifeKeeper for Linux Support Matrix](#).

Now you are ready to install LifeKeeper.

## 12.2.5. Building an HA Cluster with LifeKeeper

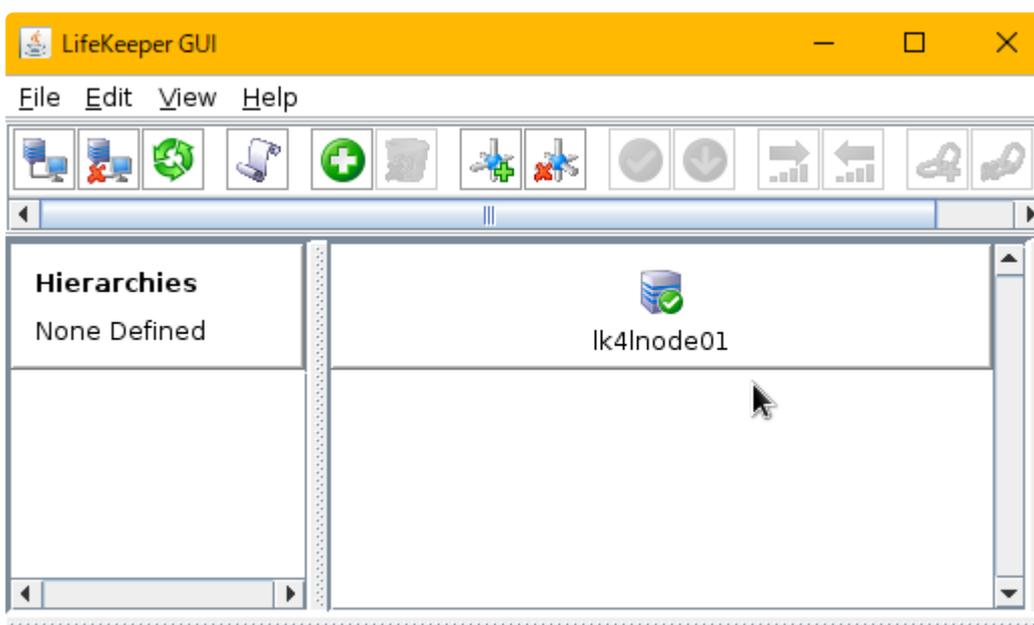
Install LifeKeeper on the virtual machine OS and build an HA cluster.

### Installing and Launching LifeKeeper

1. Upload the LifeKeeper installation image and license file on the virtual machine using the scp command etc. in advance.
2. Install LifeKeeper on **both cluster nodes**. For installation instructions, refer to the [LifeKeeper for Linux Installation Guide](#).
  - Select Use Quorum / Witness Functions.
  - For LifeKeeper Authentication, enter root.
  - In the Recovery Kit Selection Menu, select LifeKeeper Oracle RDBMS Recovery Kit and DataKeeper for Linux.
  - Select LifeKeeper Startup After Install.
3. After installing, add the LifeKeeper command path to your shell environment variables.

```
# vi /root/.bash_profile
PATH=$PATH:/opt/LifeKeeper/bin MANPATH=$MANPATH:/opt/LifeKeeper/man
export PATH MANPATH
# source /root/.bash_profile
```

4. Start the LifeKeeper GUI.



```
<-- lk4lnode01: Adding app/res: scsi netraid
```

## 12.2.5.1. Creating a Communication Path

---

LifeKeeper performs live monitoring of the nodes by checking the heartbeat response through all the configured communication paths. If the heartbeat communication does not receive a predetermined number of responses, it is determined that the communication path is disconnected, and the status is set to DEAD. If the status of all communication paths is DEAD, LifeKeeper determines it as a node failure and initiates a failover.

Since two or more communication paths between servers are recommended with LifeKeeper, configure them as follows.

Type	Priority	Active	Standby
TCP	1	10.3.1.11	10.3.1.12
TCP	2	10.3.2.11	10.3.2.12

Refer to [Creating a Communication Path](#) for details.

## 12.2.5.2. Configuring Quorum/Witness

---

Configure Quorum/Witness on the client to prevent split brain. For Quorum / Witness settings, please refer to the following URLs and choose the mode that meets your requirements.

[Quorum/Witness](#)

[Quorum Parameters List](#)

## 12.2.5.3. Disable Broadcast Ping

---

In an Azure virtual machine environment, the response to Broadcast Ping cannot be obtained immediately after creating the IP resource. Therefore, disable the alive monitoring with Broadcast Ping on both cluster nodes.

Edit the following file.

```
# vi /etc/default/LifeKeeper  
<omitted>  
NOBCASTPING=1 //Change from an initial value 0
```

## 12.2.5.4. Creating IP Resources

Create an IP resource with the LifeKeeper GUI.

1. Select **Create Resource Hierarchy** from the LifeKeeper GUI administration screen.



2. Follow the Create Resource wizard. Use the following values.

Item	Value to be entered or selected	Notes
Select Recovery Kit	IP	
Switchback Type	Intelligent	
IP Resource	10.3.1.200	*Note 1
Netmask	255.255.255.0	
Network Interface	eth0	
IP Resource Tag	ip-10.3.1.200	

- **Note 1:** Use the same value as the configured Azure load balancer frontend IP.

When extending, use the following values.

Item	Value to be entered or selected	Notes
Target Server	lk4lnode02	
Switchback Type	Intelligent	
Template Priority	1	
Target Priority	10	
IP Resource	10.3.1.200	*Note 1
Netmask	255.255.255.0	
Network Interface	eth0	
IP Resource Tag	ip-10.3.1.200	

- **Note 1:** Use the same value as the configured Azure load balancer frontend IP.

3. IP resource is created.

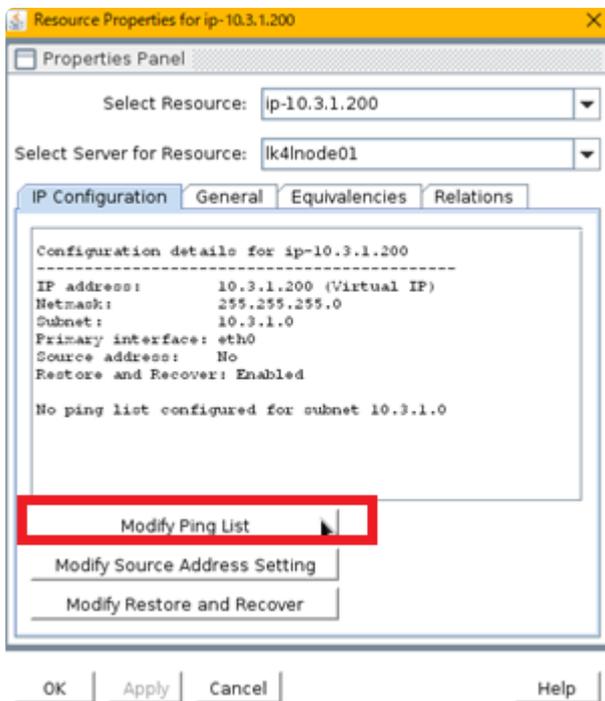


### [Creating an IP Resource Hierarchy](#)

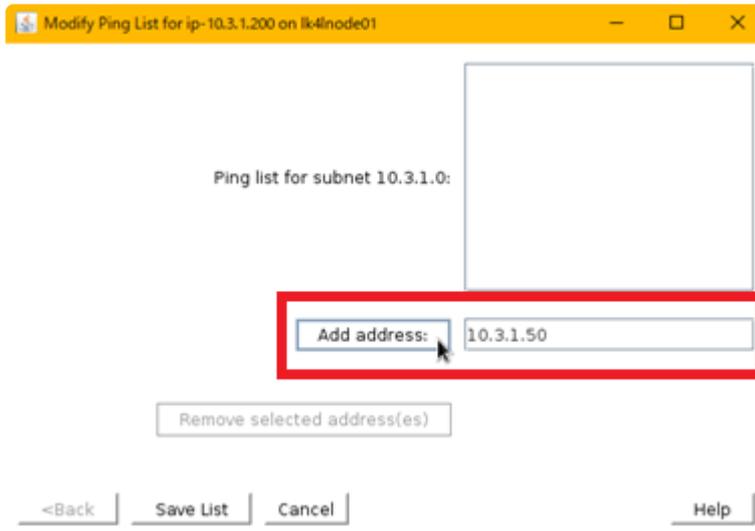
In this environment, since the monitoring by Broadcast Ping is disabled, we will perform live monitoring using Unicast Ping.

### Example

Open the IP resource property screen and click **Modify Ping List**.



Enter the IP address to set in the Ping List and click **Add address**.



Click **Save List** to save the settings.

Verify that you can get a ping response from the Ping List address you configured.

```
# ping -I 10.3.1.200 10.3.1.50
```

## 12.2.5.5. Creating a Data Replication Resource Hierarchy

---

This document describes how to create a data replication resource.

### Configuring a Disk

First, configure the disk to be used by the data replication resource.

1. Create a directory for an oracle database mount point on both nodes of the cluster.

```
# mkdir /mnt/ORA
```

2. On both nodes, verify that the disk (the disk specified when creating the instance) is attached.

```
#parted -l
(omitted)
/dev/sdc: unrecognised disk label
Model: Msft Virtual Disk (scsi)
Disk /dev/sdc: 33.8GB
(omitted)
Model: Msft Virtual Disk (scsi)
Disk /dev/sdd: 11.8GB
(omitted)
```

3. Create a disk partition for Oracle on both nodes.

```
#parted /dev/sdc
GNU Parted 3.1
Using /dev/sdc
..... (omitted)
```

4. Check the created partition.

```
# parted -l
(omitted)
Model: Msft Virtual Disk (scsi)
Disk /dev/sdc: 32.2GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  32.2GB             ora

Model: Msft Virtual Disk (scsi)
Disk /dev/sdd: 10.7GB
Sector size (logical/physical): 512B/512B
Partition Table: loop
Disk Flags:

Number  Start   End     Size    File system  Flags
  1      0.00B  10.7GB  linux-swap(v1)
```

5. Create a mount point for the file system for Oracle on both cluster nodes.

```
# mkdir /mnt/ORA
```

6. Next, configure DataKeeper to recognize the device.

When performing data replication using DataKeeper for Linux on Azure, since the ID used to identify the storage with the standard configurations cannot be obtained, you need to create a GUID Partition (GPT) and assign a unique ID to the partition or use an LVM. Refer to [Troubleshooting](#) for more information.

Now the disk is configured.

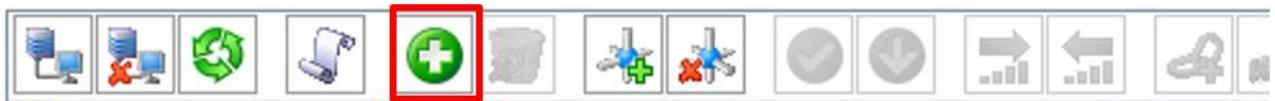
## Creating a Data Replication Resource Hierarchy

Next, create a Data Replication resource with the LifeKeeper GUI.

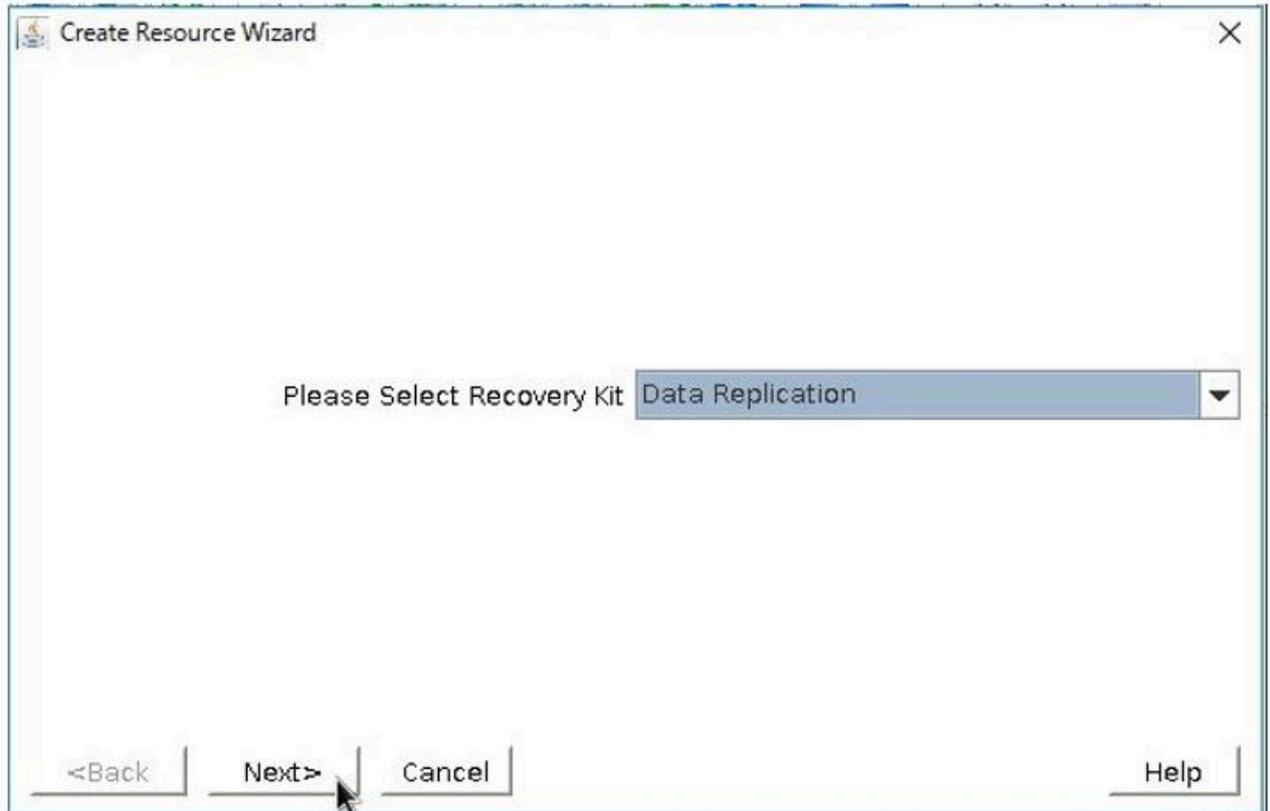
1. Start the LifeKeeper GUI with the following command on the primary node.

```
# /opt/LifeKeeper/bin/lkGUIapp &
```

2. Click the **Create Resource Hierarchy** icon to start creating a resource.



3. The Create Resource Wizard appears. Select **Data Replication** for the Recovery Kit from the pull-down menu and click **Next**.

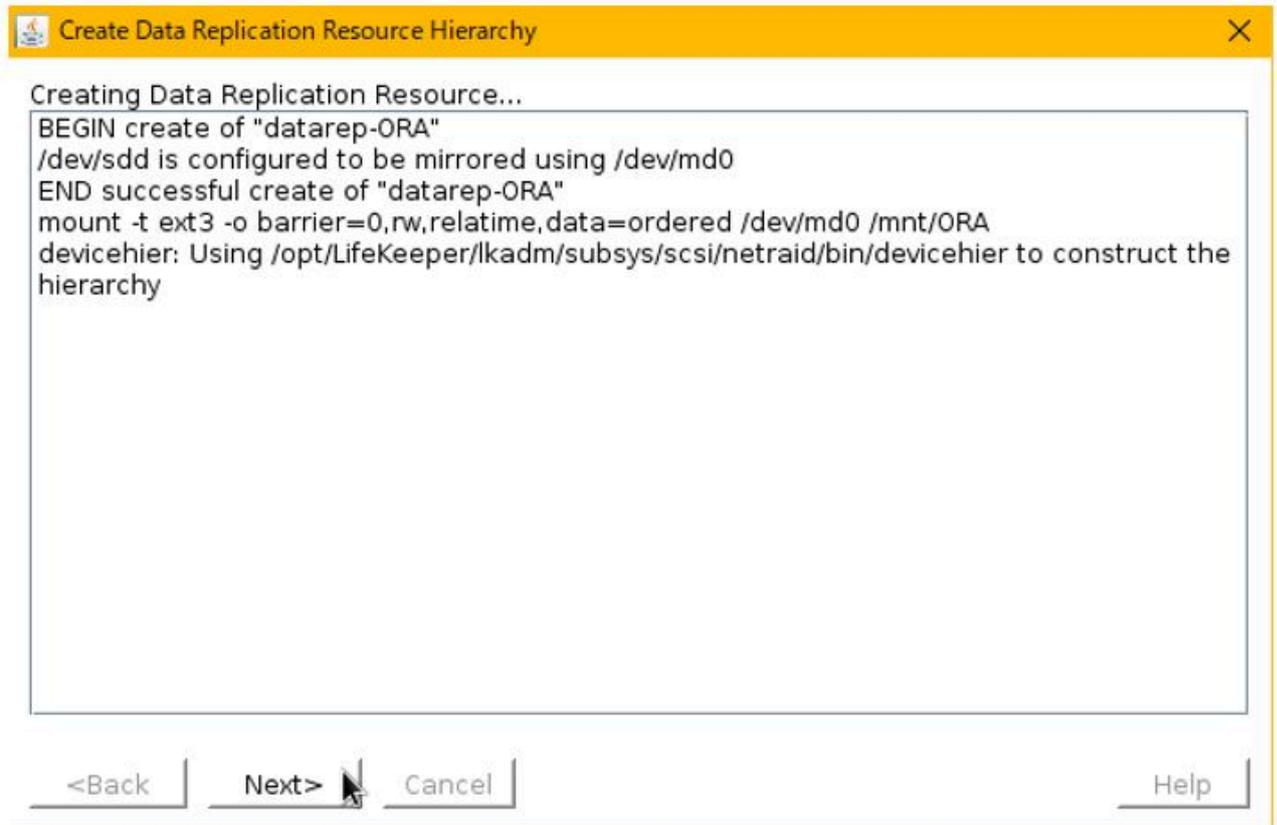


4. In the Create Resource Wizard, enter the following values.

Item	Value to be entered or selected	Notes
Switchback Type	Intelligent	
Server	lk4lnode01	
Hierarchy Type	Replicate New Filesystem	
Source Disk	/dev/sdc1	*Note 1
New Mount Point	/mnt/ORA	*Note 2
New Filesystem Type	xfs	
Data Replication Resource Tag	datarep-ORA	
File System Resource Tag	/mnt/ORA	
Bitmap File	/opt/LifeKeeper/bitmap_mnt_ORA	*Note 3
Enable Asynchronous Replication?	no	

- **Note 1:** After this, “**ATTENTION! <Device name> is not shareable with any other server**” is displayed, you can ignore it because it is always displayed when creating a replication.
- **Note 2:** Specify the mount point that has been created.
- **Note 3:** Although the default path is used this time, consider specifying a faster private area to improve performance.

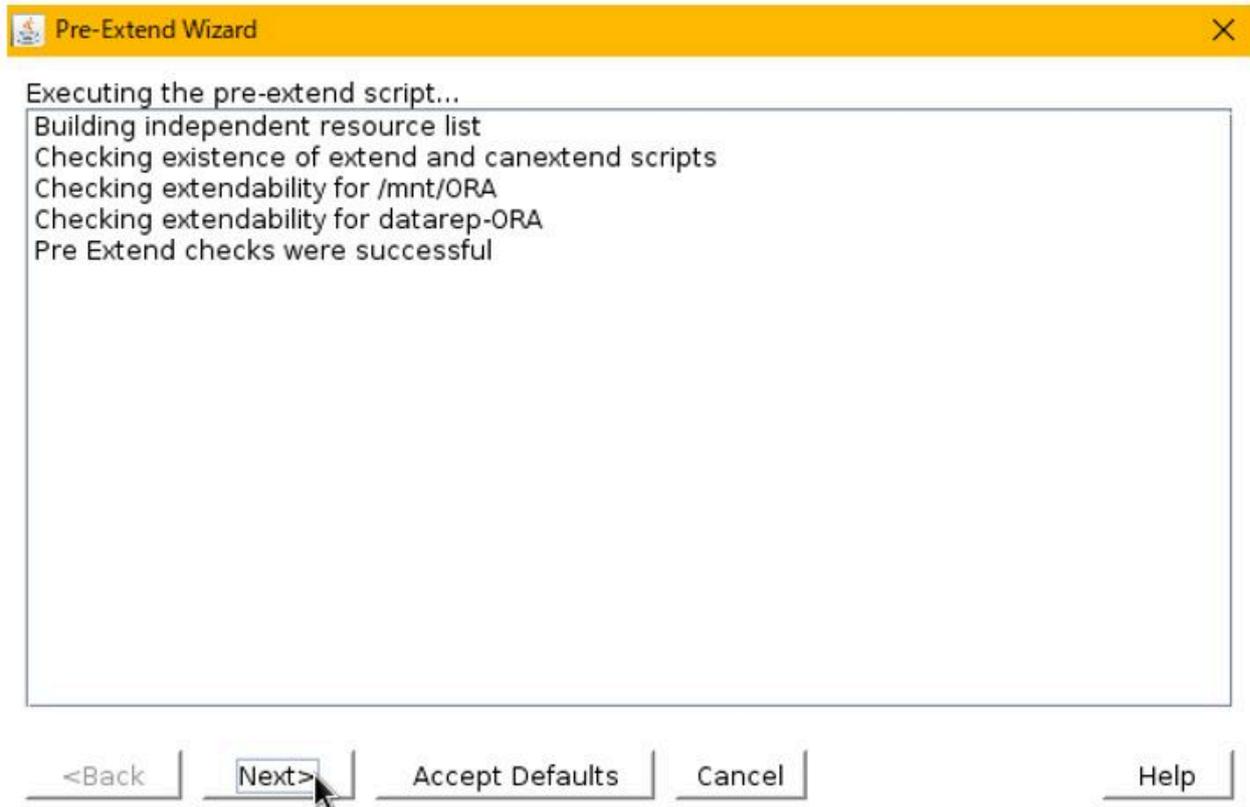
5. Creation of the data replication resource will begin.
6. “End of successful Create of...” is displayed on the screen to confirm it is successful. Click **Next** to go to the Pre-Extend Wizard.



7. Enter the following values.

Item	Value to be entered or selected
Target Server	lk4lnode02
Switchback Type	Intelligent
Template Priority	1
Target Priority	10

8. Pre-Extend is started. It is successful if “Pre Extend checks were successful” is displayed. Click **Next**.

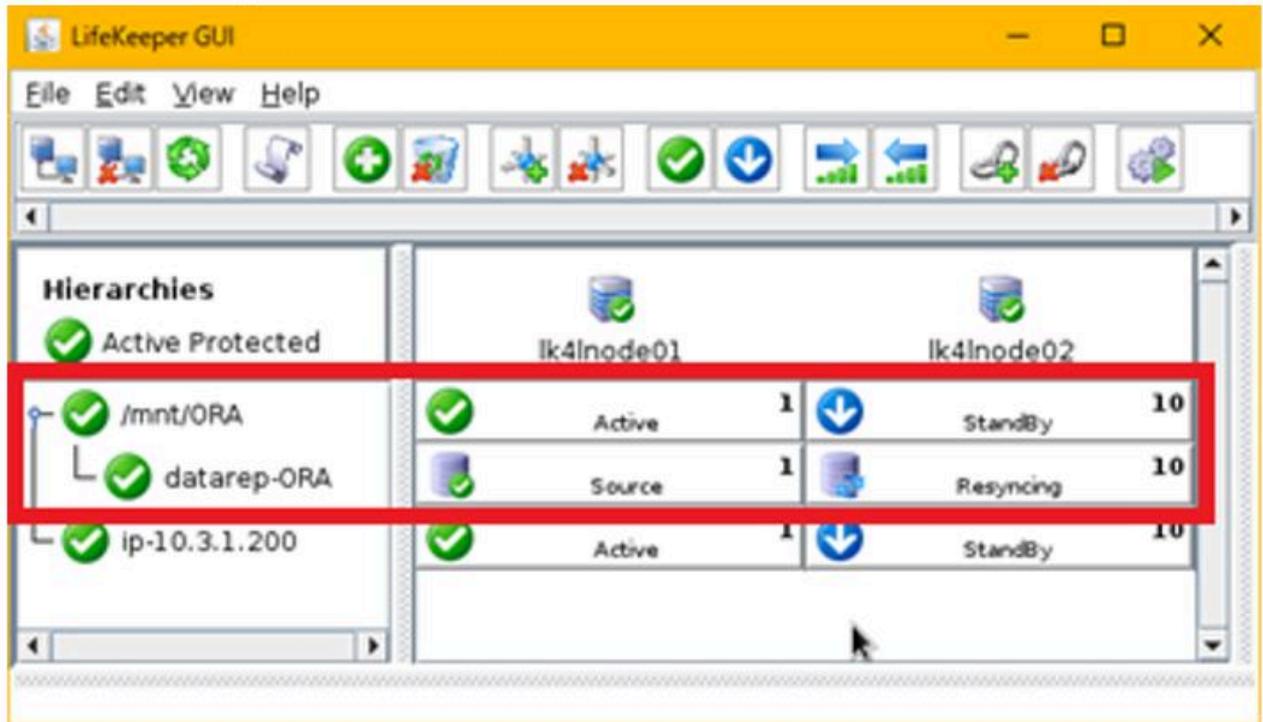


9. In the Extend Resource Wizard, enter the following values.

Item	Value to be entered or selected	Notes
Mount Point	/mnt/ORA	
Root Tag	/mnt/ORA	
Target Disk	/dev/sdc1	
Data Replication Resource Tag	datarep-ORA	
Bitmap File	/opt/LifeKeeper/bitmap_mnt_ORA	
Replication Path	10.3.2.11/10.3.2.12	*Note 1

\*Note1: Specify the communication path on the second subnet.

10. Extending starts. If “Hierarchy successful extended” is displayed, it is successful. Click **Finish**.
11. Click **Done** to exit the wizard.
12. Data replication resource is created.



\*Data replication resource is not created for the swap area.

Refer to [Creating a DataKeeper Resource Hierarchy](#) for details.

## 12.2.5.6. Creating an Oracle Resource Hierarchy

Set up the environment for installing Oracle Database 12c. For details on the prerequisites for installing Oracle, please refer to the Oracle documentation.

Make the following settings on **both cluster nodes**.

1. Create a swap area.

At least 3GB of swap space is required to install Oracle. Configure the swap area using the 10GB of disk (`/dev/sdd`) assigned to the instance.

```
# mkswap /dev/sdd
# swapon /dev/sdd
# free
      total        used        free      shared  buff/cache   available
Mem:    1706084        189248        1348300          9184        168536        1300564
Swap:   12582904            0        12582904
# blkid | grep /dev/sdd
/dev/sdd: UUID="268e56bd-c64f-4798-9197-e411a497ba11" TYPE="swap" //write down the UUID
# vi /etc/fstab
# /etc/fstab
# Created by anaconda on Fri Mar 23 17:41:14 2018
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=50a9826b-3a50-44d0-ad12-28f2056e9927 /        xfs        defaults        0 0
UUID=268e56bd-c64f-4798-9197-e411a497ba11 swap swap defaults 0 0 ←追加
```

2. Register the RHEL option repository. Add a repository containing the required package "compat-libstdc++-33".

```
# yum-config-manager --enable rhui-REGION-rhel-server-extras rhui-REGION-rhel-server-optional
```

3. Install the package.

```
# yum -y install binutils compat-libcap1 compat-libstdc++-33 compat-libstdc++-33.i686
gcc gcc-c++ glibc glibc.i686 glibc-devel glibc-devel.i686 ksh libgcc libgcc.i686 libstdc++
libstdc++.i686 libstdc++-devel libstdc++-devel.i686 libaio libaio.i686 libaio-devel libaio-
devel.i686 libXext libXext.i686 libX11 libX11.i686 libxcb libxcb.i686 libXi libXi.i686 make
sysstat
```

4. Edit the kernel parameters.

**Note:** Symbols and blanks may be changed to different characters due to Word and Exel specifications and an error may occur when you execute the command using copy and paste. Do not add a blank at the end. The command may fail because the variable is recognized as a string instead of a number.

```
# MEMTOTAL=$(free -b | sed -n '2p' | awk '{print $2}')
# SHMMAX=$(expr $MEMTOTAL / 2)
# SHMMNI=4096
# PAGESIZE=$(getconf PAGE_SIZE)
# cat >> /etc/sysctl.conf << EOF
>fs.aio-max-nr = 1048576
>fs.file-max = 6815744
>kernel.shmmax = $SHMMAX
>kernel.shmall = `expr ¥( $SHMMAX / $PAGESIZE ¥) ¥* ¥( $SHMMNI / 16 ¥)`
>kernel.shmmni = $SHMMNI
>kernel.sem = 250 32000 100 128
>net.ipv4.ip_local_port_range = 9000 65500
>net.core.rmem_default = 262144
>net.core.rmem_max = 4194304
>net.core.wmem_default = 262144
>net.core.wmem_max = 1048576
>EOF
```

5. Apply to the kernel.

```
# sysctl -p
```

6. Create a dedicated user group for Oracle and configure the system environment.

```
# i=54321; for group in oinstall dba backupdba oper dgdba kmdba; do groupadd -g $i
$group; i=`expr $i + 1` ;done
# useradd -u 1200 -g oinstall -G dba,oper,backupdba,dgdba,kmdba -d /home/oracle
```

7. Set the password for the Oracle user.

```
# passwd oracle
Changing password for user oracle.
New password:
Retype new password:
Passwd: all authentication tokens updated successfully.
```

8. Verify that `/mnt/ORA` is protected by DK and mounted on the primary node. (This will only be done on the primary node)

```
# df
Filesystem      1K-blocks    Used Available Use% Mounted on
<<中 omitted>>
/dev/md0        31864752    7396580   22842880   25% /mnt/ORA
<<後 omitted>>
#lcdstatus -e
lk4lnode02  LSNR.LISTENER  dblistener-11079  ISP      1  lk4lnode01
lk4lnode02  /mnt/ORA       /mnt/ORA          ISP      1  lk4lnode01
lk4lnode02  datarep-ORA   /dev/sdc          ISP      1  lk4lnode01
lk4lnode02  ip-10.3.1.200  ISP              1  lk4lnode01
<<後 omitted>>
```

9. Change directory creation, access rights, and owner. (Perform this only on the primary node)

```
# mkdir -p /mnt/ORA/app/oracle
# chown -R oracle:oinstall /mnt/ORA/app
# chmod -R 775 /mnt/ORA
```

10. Add it in */etc/pam.d/login* around line 14.

```
# vi /etc/pam.d/login
session    required    pam_selinux.so open
session    required    pam_namespace.so
session    required    pam_limits.so //Add
session    optional   pam_keyinit.so force revoke
```

11. Add it to the last line of */etc/security/limits.conf*.

```
# vi /etc/security/limits.conf
oracle soft nproc 2047 //Add
oracle hard nproc 16384 //Add
oracle soft nofile 1024 //Add
oracle hard nofile 65536 //Add
oracle soft stack 10240 //Add
oracle hard stack 32768 //Add
```

12. Set up the environment for Oracle. The work is done with the “oracle” user created in the previous step.

```
# su oracle
$ vi ~/.bash_profile
// Add the following to the last line
umask 022
export ORACLE_BASE=/mnt/ORA/app/oracle
export ORACLE_HOME=/mnt/ORA/app/oracle/product/12.1.0/dbhome_1
export PATH=$PATH:$ORACLE_HOME/bin
export ORACLE_SID=<SID name>
export NLS_LANG=Japanese_Japan.AL32UTF8

$source ~/.bash_profile
```

Write down the ORACLE\_SID you set up here for use in later steps. In this example it is "lkoracle".

13. Copy the root xauth settings to the oracle user. (This is done only on the primary node)

```
# xauth list
# su oracle
$ xauth add <DisplayName> <ProtocolName> <Hexkey>
```

14. Verify that the xauth output for the oracle user is the same as root.

```
$ xauth list
```

15. Create a directory for installation. (This is done only on the primary node)

```
$ mkdir /home/oracle/tmp
```

Now you are ready to install.

## 12.2.5.6.1. Installing Oracle DB

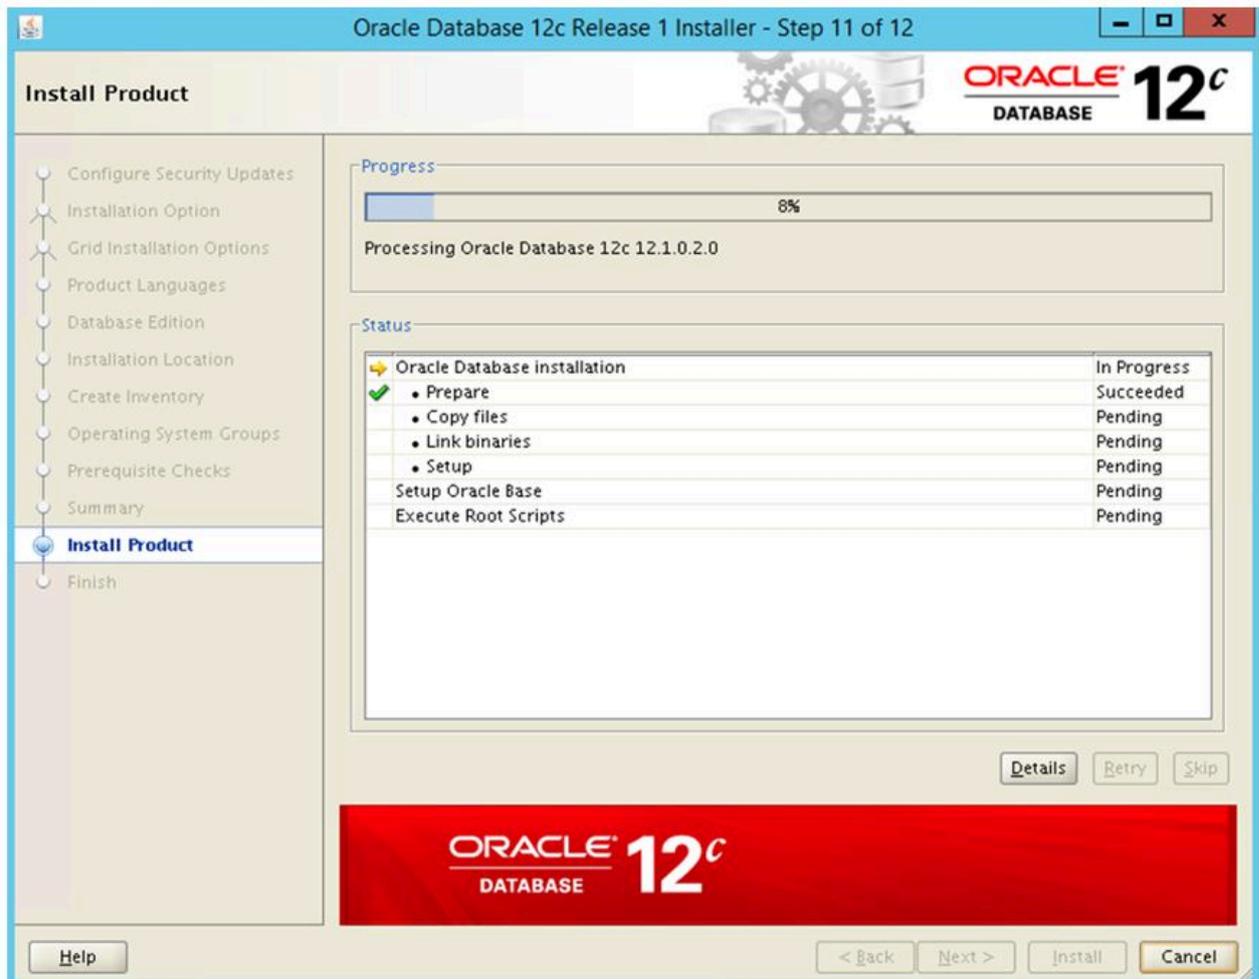
---

Perform the following steps on the cluster node (primary). For Oracle installation, please refer to the Oracle documentation.

1. Download the Oracle 12c installation package and unzip it to the appropriate directory under */home/oracle*.
2. Switch to the oracle user with su oracle
3. A directory named “database” is created under the unzipped directory. The directory will launch the installer, which is located:

```
./runInstaller &
```

4. Configure Security Update. Enter the registered email address and support password for Oracle and click **Next**. This is not mandatory.
5. Installation option. Select **Install Database Software only** and click **Next**.
6. Grid Installation options. Select **Single Instance Database** installation and click **Next**.
7. Product Language. Select your **preferred language** and click **Next**.
8. Database Edition. Select **Enterprise Edition** and click **Next**.
9. Installation Location. If it is  
Oracle base = /mnt/ORA/app/oracle  
Software Location = /mnt/ORA/app/oracle/product/12.1.0/dbhome\_1  
click **Next**.
10. Create Inventory. Use the **default value** and click **Next**.
11. Product list checks. Use the **default value** and click **Next**.
12. Operating System Group. Use the **default value** and click **Next**.
13. Summary. If there is no problem, run **Install**.
14. Installation is started.



15. The Execute Configuration scripts window will open. Follow the instructions and run the script as root.

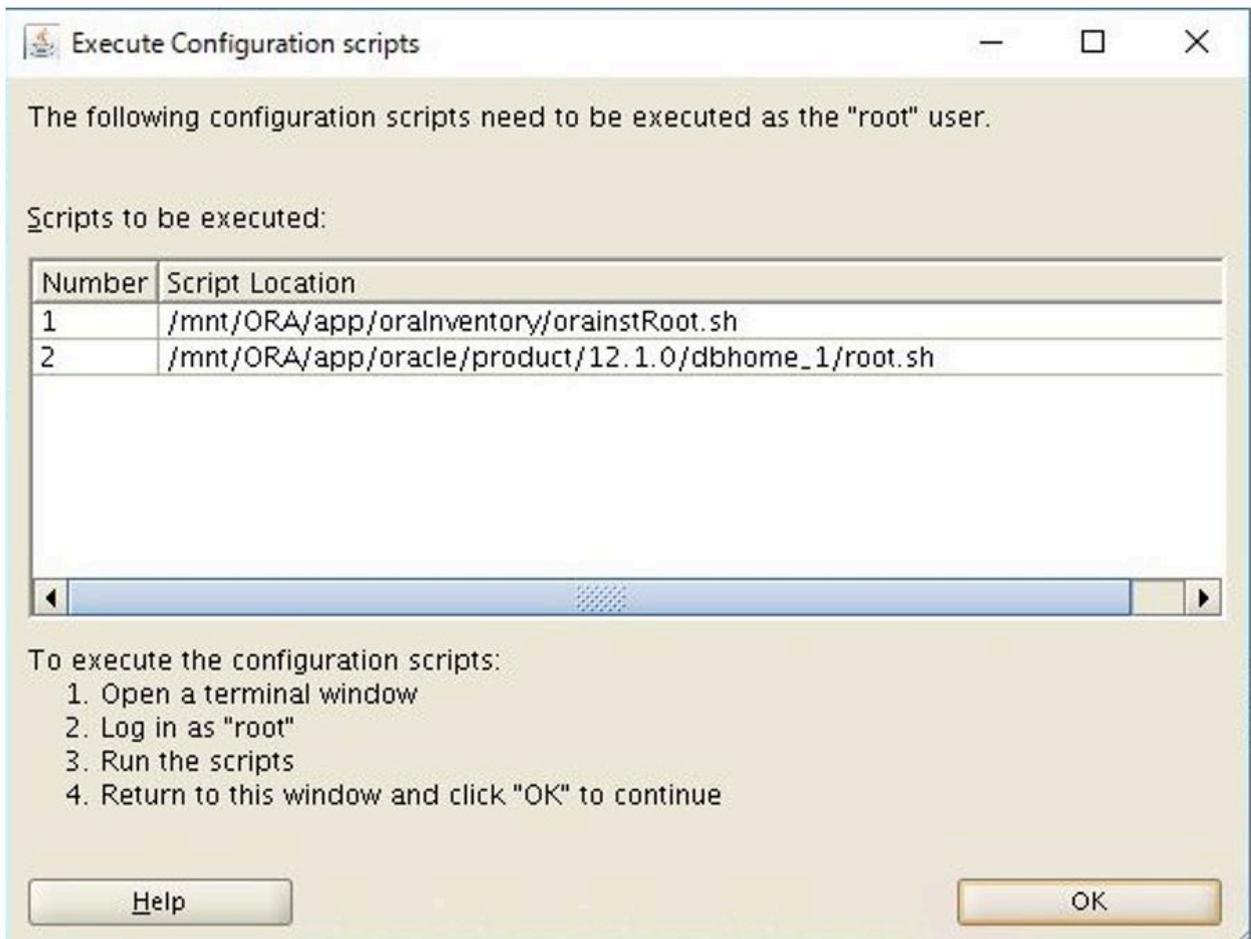
```
# /mnt/ORA/app/oraInventory/orainstRoot.sh
Changing permissions of /mnt/ORA/app/oraInventory.
Adding read,write permissions for group.
Removing read,write,execute permissions for world.

Changing groupname of /mnt/ORA/app/oraInventory to oinstall.
The execution of the script is complete.
# /mnt/ORA/app/oracle/product/12.1.0/dbhome_1/root.sh
Performing root user operation.

The following environment variables are set as:
  ORACLE_OWNER= oracle
  ORACLE_HOME= /mnt/ORA/app/oracle/product/12.1.0/dbhome_1

Enter the full pathname of the local bin directory: [/usr/local/bin]:
  Copying dbhome to /usr/local/bin ...
  Copying oraenv to /usr/local/bin ...
  Copying coraenv to /usr/local/bin ...

Creating /etc/oratab file...
Entries will be added to the /etc/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root script.
Now product-specific root actions will be performed.
```



When the script is completed, click **OK**.

16. Click **Close** to complete the installation.

## 12.2.5.6.2. Configuring a Listener

Configure a listener. On the cluster node (primary), make the following settings:

1. As the oracle user, execute the following command.

```
$ /mnt/ORA/app/oracle/product/12.1.0/dbhome_1/bin/netca &
```

2. Select **Listener Configuration** and click **Next**.



3. Select **Add** and click **Next**.



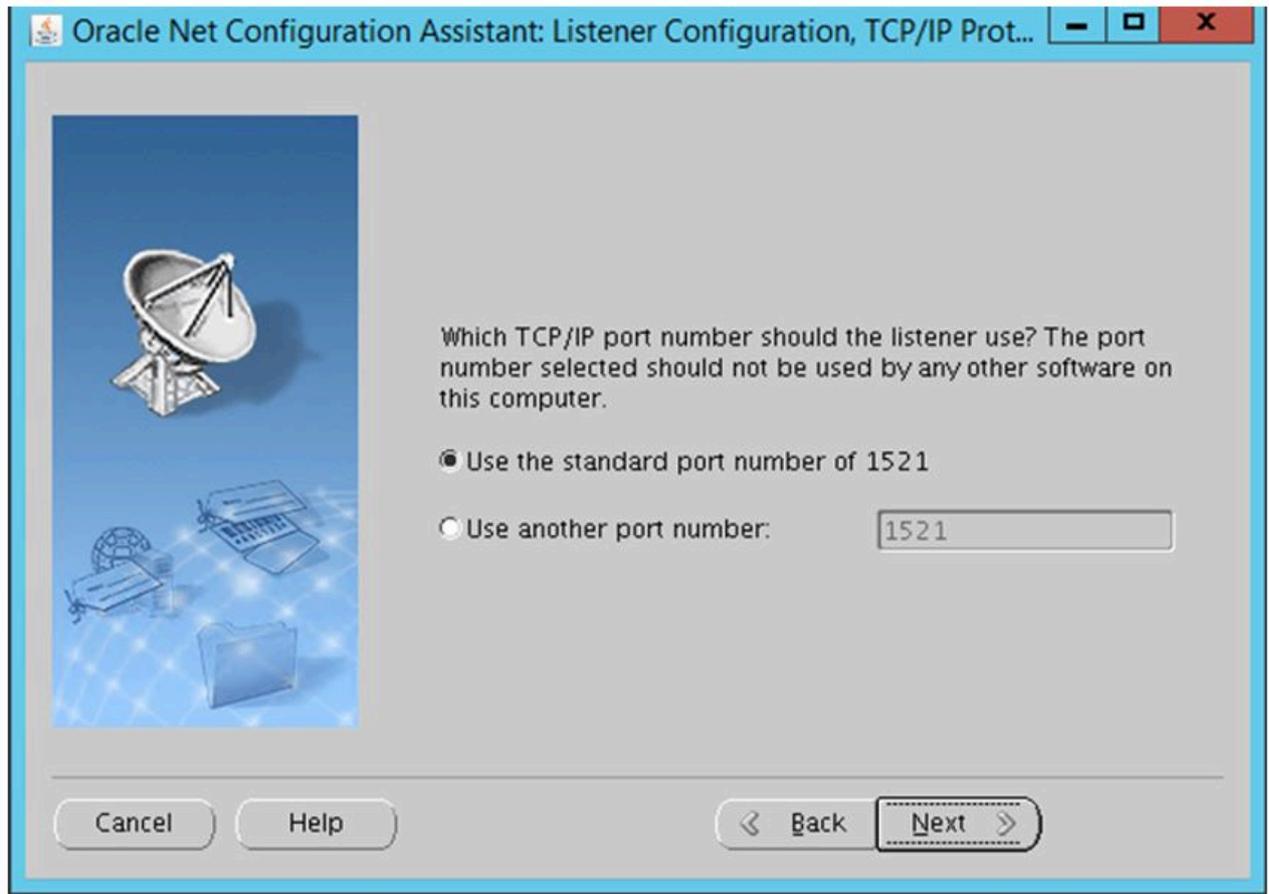
- Specify a name for the Listener and click **Next**.



5. Click **Next**.



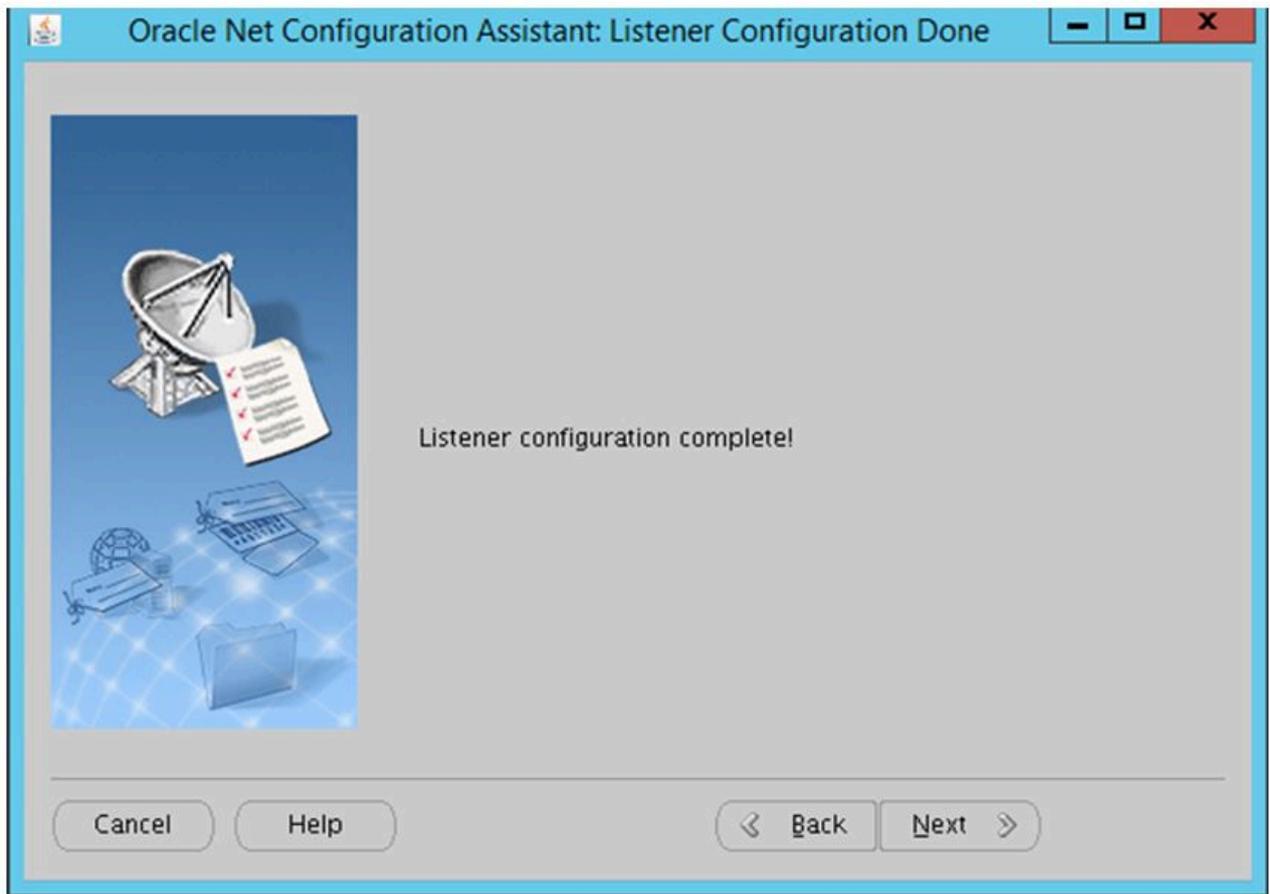
6. Select the default **port 1521** and click **Next**.



7. Complete the Listener configuration and click **Next**.



- The Listener configuration is completed. Click the [X] at the top right of the dialog to complete.



The Listener configuration is now completed.

## 12.2.5.6.3. Creating the DB

---

Configure the following settings on the cluster node (primary).

1. Execute the following command as an oracle user:

```
$ /mnt/ORA/app/oracle/product/12.1.0/dbhome_1/bin/dbca &
```

2. Database operation. Select **Create Database** and click **Next**.
3. Creation Mode. Select **Advanced Mode** and click **Next**.
4. Database Template. Select **General Purpose** or **Transaction Processing** and click **Next**.
5. Database Identification. Set the **database name** and **SID** as follows:  
Global Database Name = lkoracle  
SID = lkoracle
7. Management Options. Use the **default value** and click **Next**.
8. Database Credentials. Select **Use the Same Administrative Password for All Accounts** to register the password.
9. Network Configuration. Since the Listener has been created, use the **default value** and click **Next**.
10. Storage Locations. Use the **default value** and click **Next**.
11. Database Options. Use the **default value** and click **Next**.
12. Initialize Parameters. Use the **default value** and click **Next**.
13. Creation Options. Use the **default value** and click **Next**.
14. If the settings are correct, click **Finish** to start creating the database.
15. The database is created.
16. Next, modify */etc/oratab*. Note that this setting will be undone in a later step.

```
# vi /etc/oratab  
oracle:/mnt/ORA/app/oracle/product/12.1.0/dbhome_1:Y ←change N to Y
```

17. Copy */etc/oratab* from the primary node to the secondary node.

The installation of Oracle DB is now completed.

\*Installation of Oracle DB on the secondary node is not necessary.

## 12.2.5.6.4. Configuring Oracle

The configuration file must be modified before creating the Listener resource since the Listener is being accessed by the virtual IP address.

On the cluster node (primary), configure the following settings.

1. Modify `$ORACLE_HOME/network/admin/listener.ora`.

```
$vi $ORACLE_HOME/network/admin/listener.ora
<<中 omitted>>
  SID_LIST_LISTENER =
    (SID_LIST =
      (SID_DESC =
        (SID_NAME = LKORACLE)
      )
    )

  LISTENER =
    (DESCRIPTION_LIST =
      (DESCRIPTION =
        (ADDRESS = (PROTOCOL = TCP)(HOST = 10.3.1.200)(PORT = 1521))
      )
    )
  )
```

2. Modify `$ORACLE_HOME/network/admin/tnsnames.ora`.

```
$vi $ORACLE_HOME/network/admin/tnsnames.ora
LISTENER_LKORACLE =
  (ADDRESS = (PROTOCOL = TCP)(HOST = 10.3.1.200)(PORT = 1521))

LKORACLE =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = 10.3.1.200)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = lkoracle)
    )
  )
)
```

3. Create a password file to be used for login via the Listener. Refer to the official Oracle documentation for the format of the password file.

```
$orapwd file=orapw|koracle force=y ignorecase=y password=XXXXXXXX
```

4. Check the connection to the database via the Listener.

```
$ sqlplus sys/XXXXXXXX@10.3.1.200:1521/lkoracle as sysdba
```

```
SQL*Plus: Release 12.1.0.2.0 Production on 木 5月 23 07:52:41 2019
```

```
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

```
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production  
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing options  
に接続されました。
```

```
SQL> select instance_name from v$instance ;
```

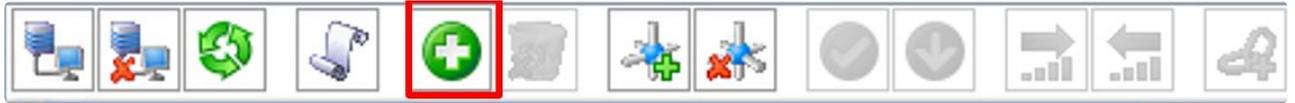
```
INSTANCE_NAME
```

```
-----  
lkoracle
```

# 12.2.5.6.5. Creating an Oracle Database Listener Resource Hierarchy

First, create an Oracle Database Listener resource.

1. Click the **Create Resource Hierarchy** icon to start creating a resource.



2. The Create Resource Wizard appears. Select **Oracle Database Listener** for the Recovery Kit from the pull-down menu and click **Next**.

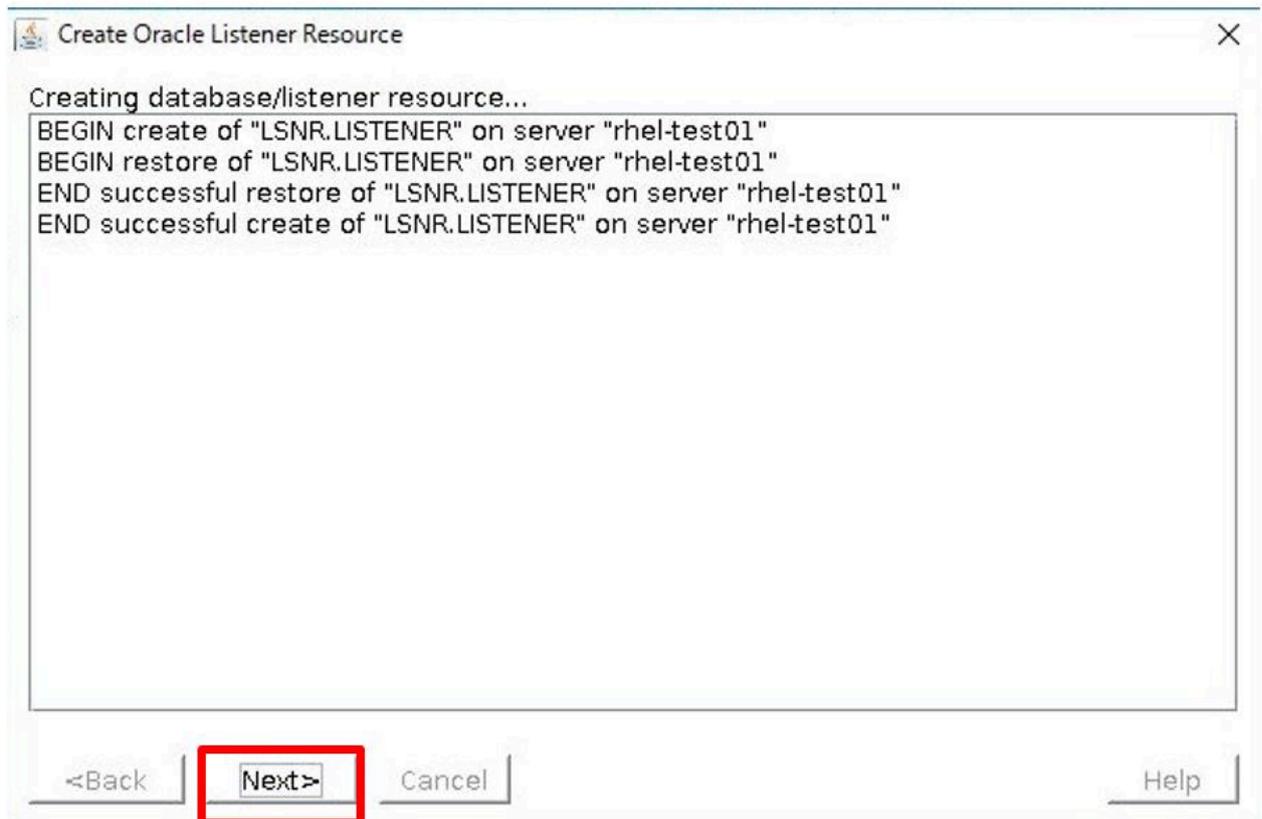


3. In the Create Resource wizard, enter the following values.

No.	Item	Value to be entered or selected
1	Switchback Type	Intelligent
2	Server	lk4lnode01
3	Listener Configuration File Path	/mnt/ORA/app/oracle/product/12.1.0/dbhome_1/network/admin/listener.ora
4	Listener Name(s)	LISTENER
5	Listener Executable(s)	/mnt/ORA/app/oracle/product/12.1.0/dbhome_1/bin
6	Listener Protection Level	Full Control(Start,Stop,Monitor,& Recover)

7	Listener Recovery Level	Standard, (On)
8	IP Address Name	ip-10.3.1.200
9	Listener Tag	LSNR.LISTENER

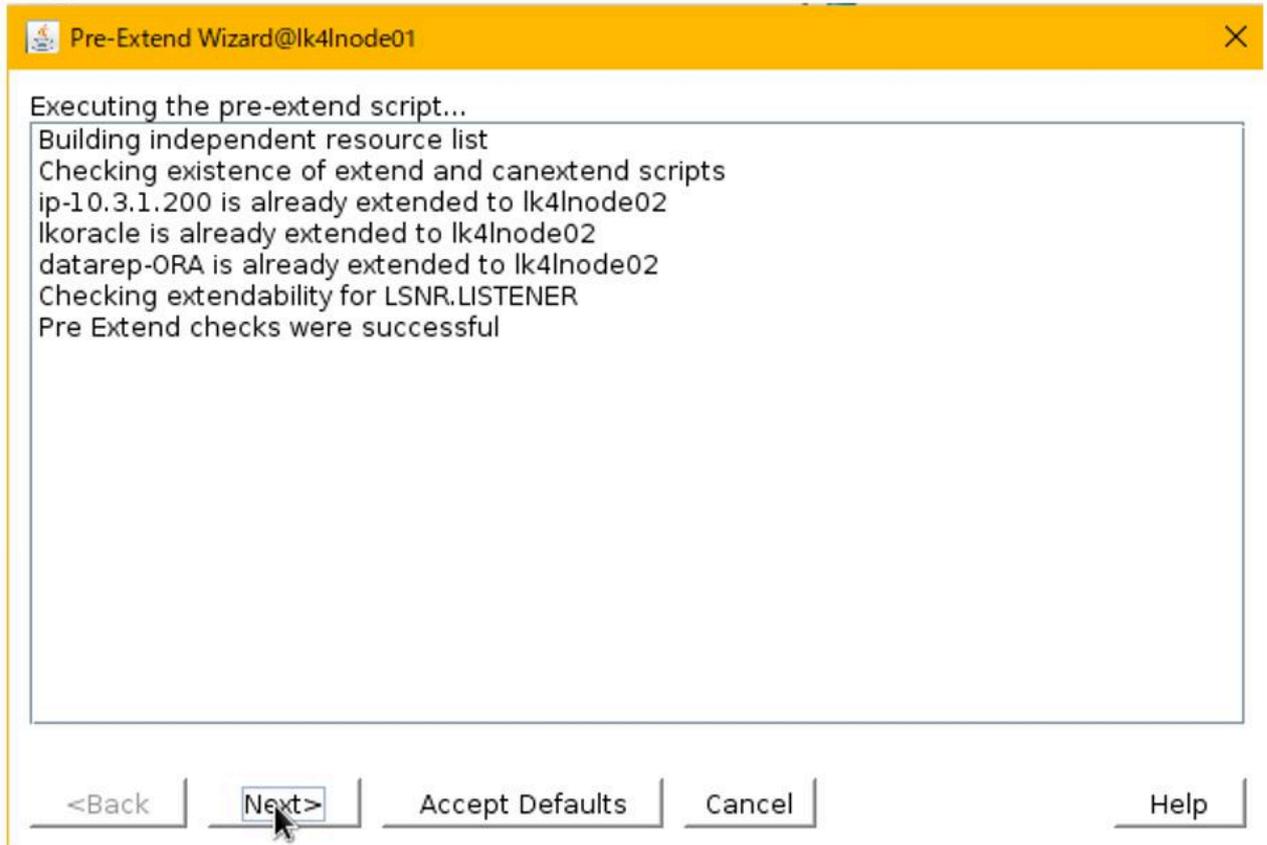
- Creation of the Oracle Database Listener resource begins.
- It is successful when “End successful create of ...” is displayed. Click **Next** to proceed to the Pre-Extend wizard.



- Enter the following values.

No.	Item	Value to be entered or selected
1	Target Server	lk4lnode02
2	Switchback Type	Intelligent
3	Template Priority	1
4	Target Priority	10

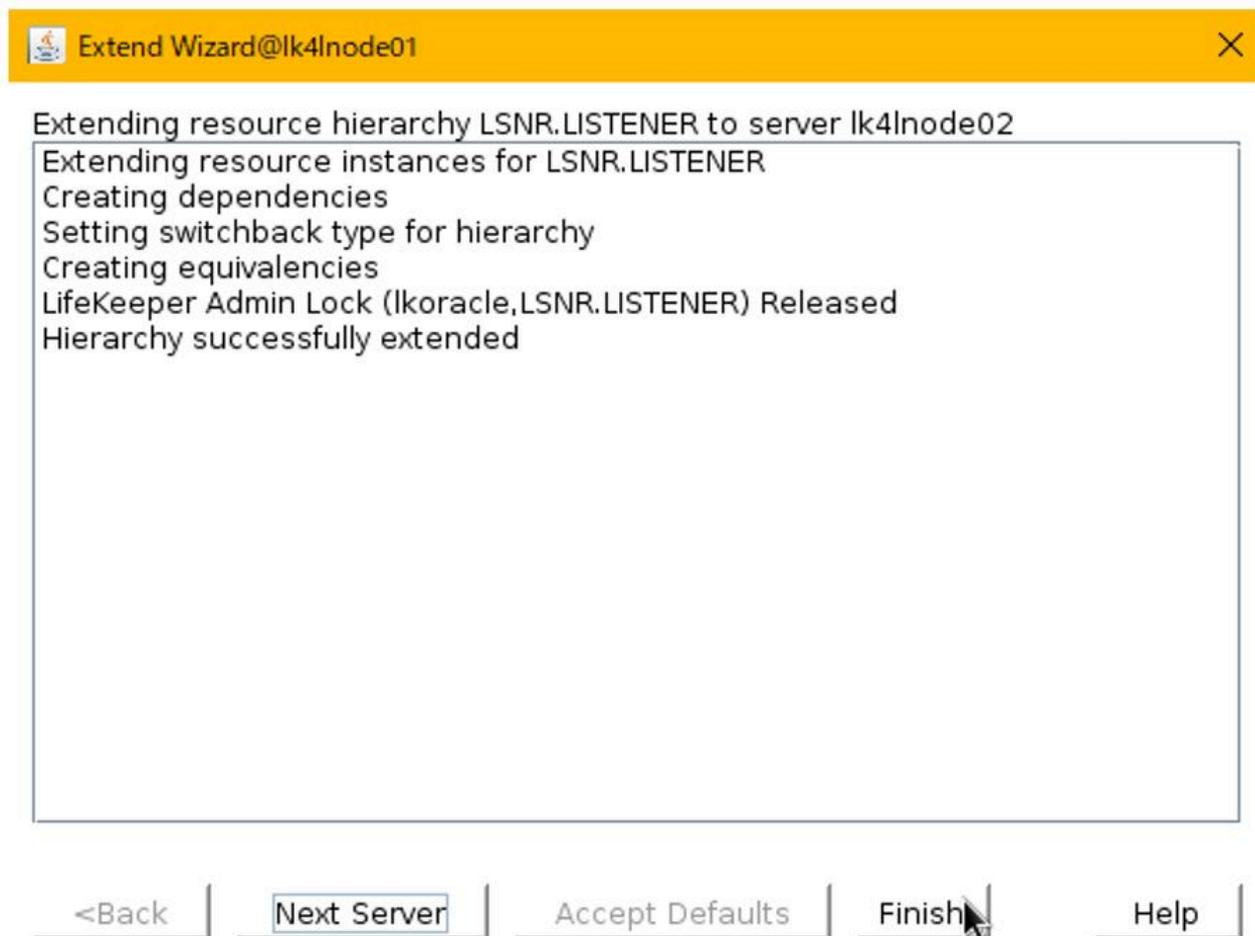
- Pre-Extend is started. It is successful when “Pre Extend checks were successful” is displayed. Click **Next**.



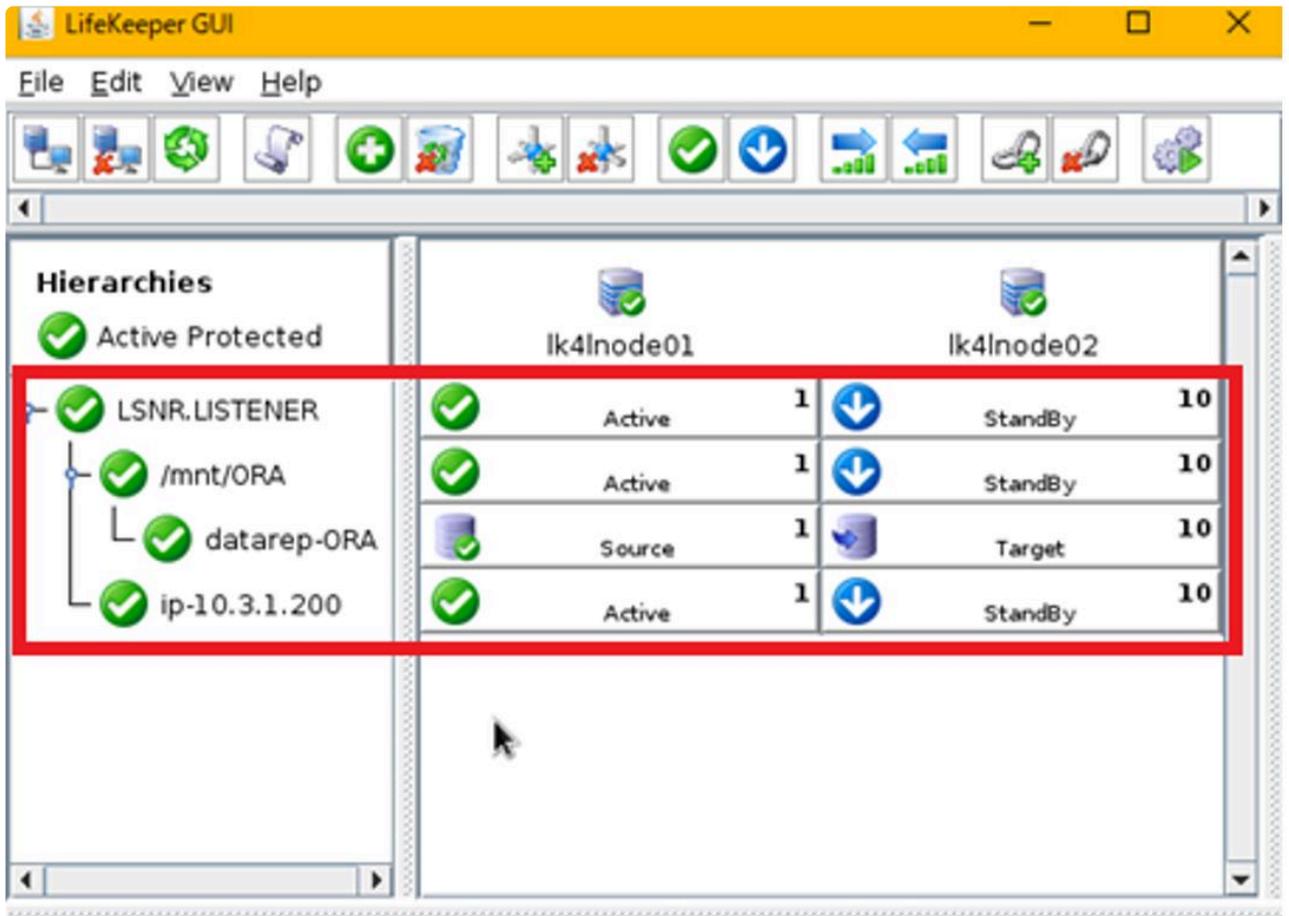
8. In the Extend Oracle Listener Resource Wizard, enter the following values.

No.	Item	Value to be entered or selected
1	Enter the Listener Configuration File Path	/mnt/ORA/app/oracle/product/12.1.0/dbhome_1/network/admin/listener.ora
2	Enter the Path to the Listener Executable(s)	/mnt/ORA/app/oracle/product/12.1.0/dbhome_1/bin
3	Listener Tag	LSNR.LISTENER

9. Extend is started. It is successful when “Hierarchy successfully extended” is displayed. Click **Finish**.



10. Press **Done** to exit the wizard.
11. Oracle Database Listener resource is created.



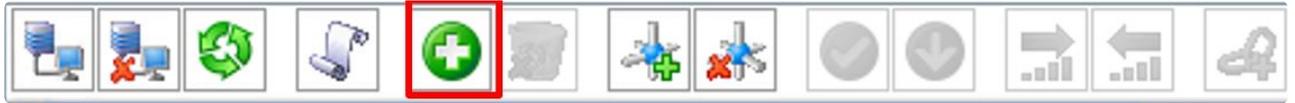
<-- lk4lnode02: datarep-ORA: Updating state to Target

Refer to [Creating an Oracle Resource Hierarchy](#) for more information.

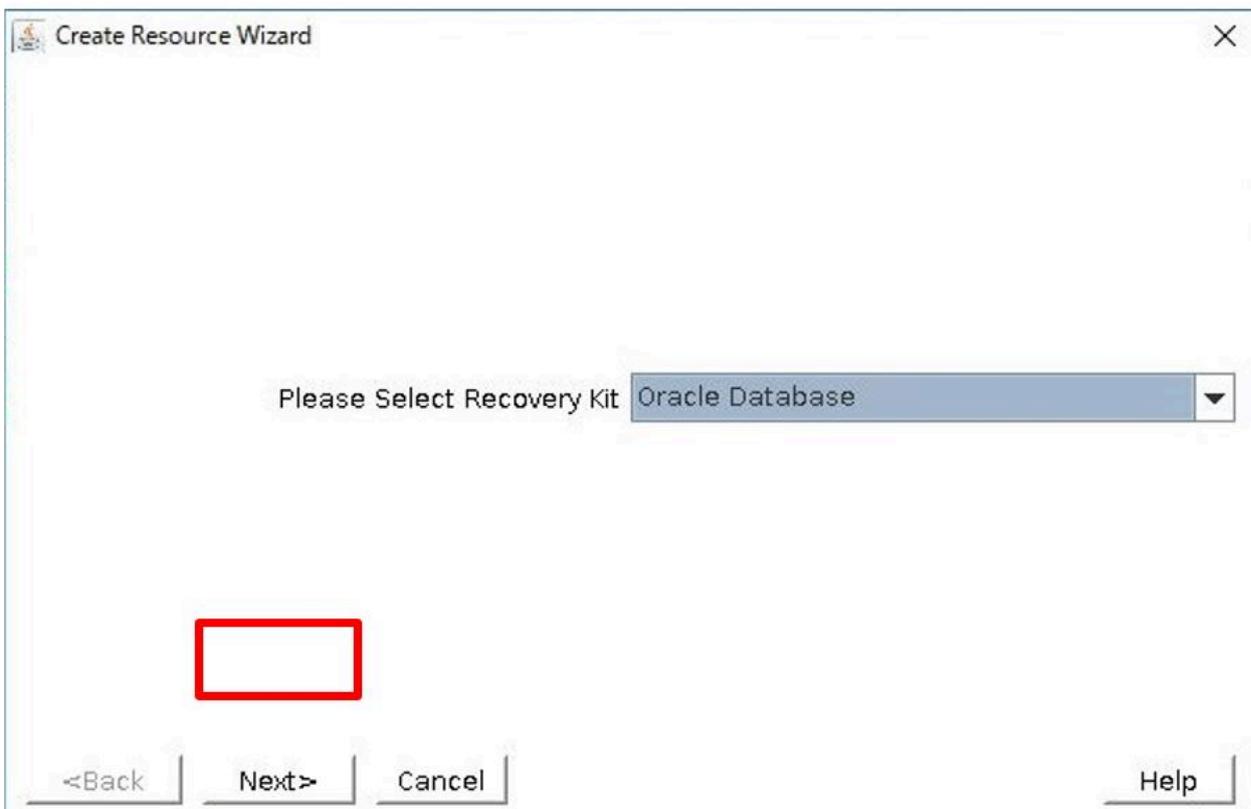
# 12.2.5.6.6. Creating an Oracle Resource Hierarchy

Next, create a resource for the Oracle database.

1. Click the Create Resource Hierarchy icon to start creating a resource.



2. The Create Resource Wizard appears. Select **Oracle Database** for the Recovery Kit from the pull-down menu and click **Next**.

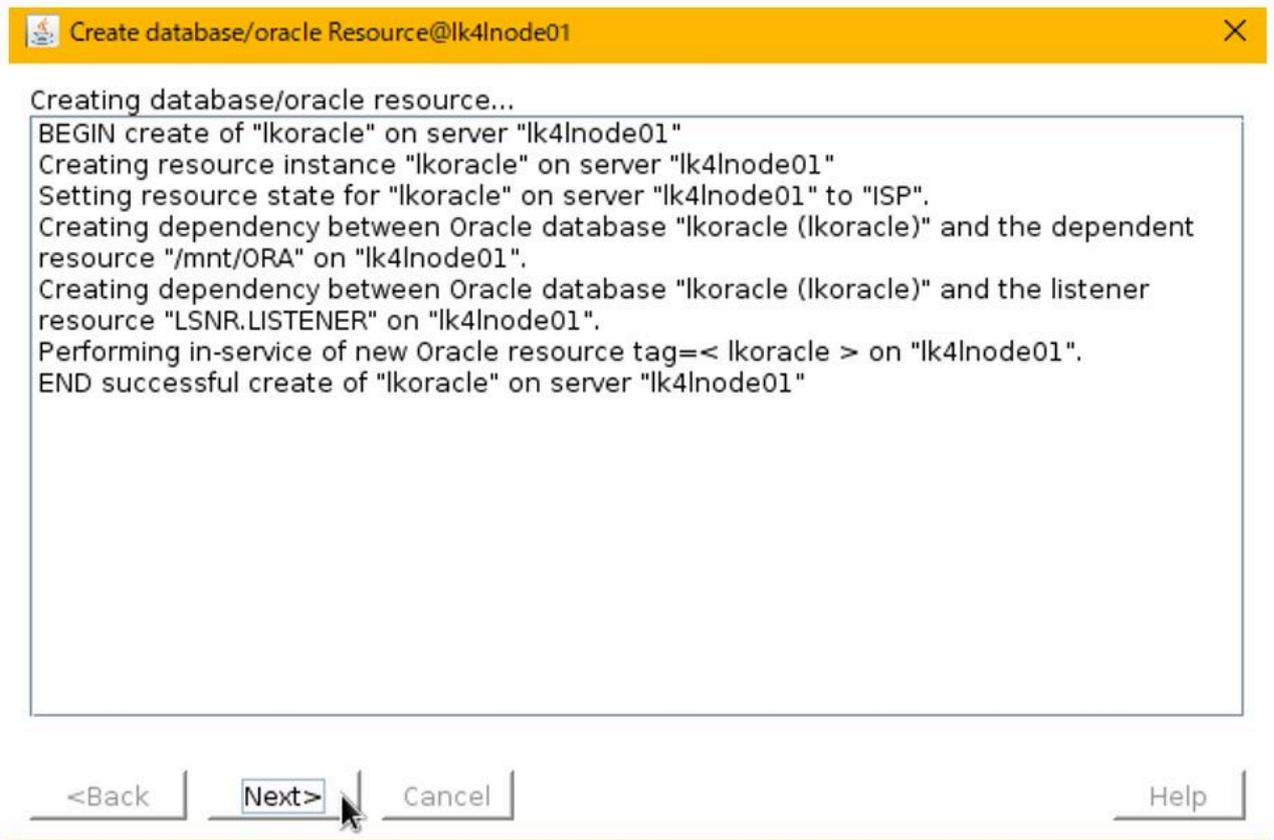


3. In the Create Resource Wizard, enter the following values.

No.	Item	Value to be entered or selected
1	Switchback Type	Intelligent
2	Server	lk4lnode01
3	Oracle _ SID for Database	lkoracle
4	User Name	oracle
5	Password	<< Password specified when creating Oracle DB >>

6	Select the Oracle Listener	LSNR.LISTENER
7	Database Tag	lkoracle

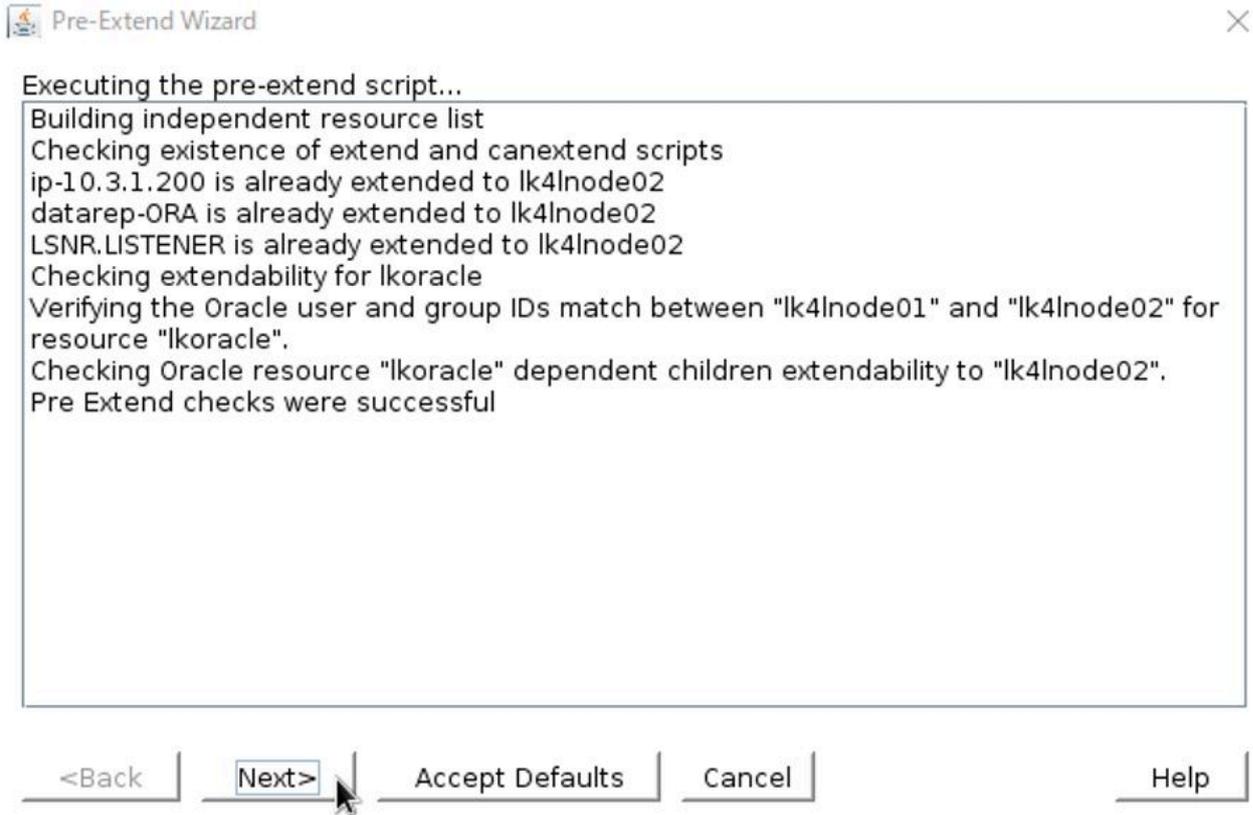
- The Oracle Database resource creation starts.
- It is successful when “End successful Create of ...” is displayed. Click **Next** to proceed to the Pre-Extend wizard.



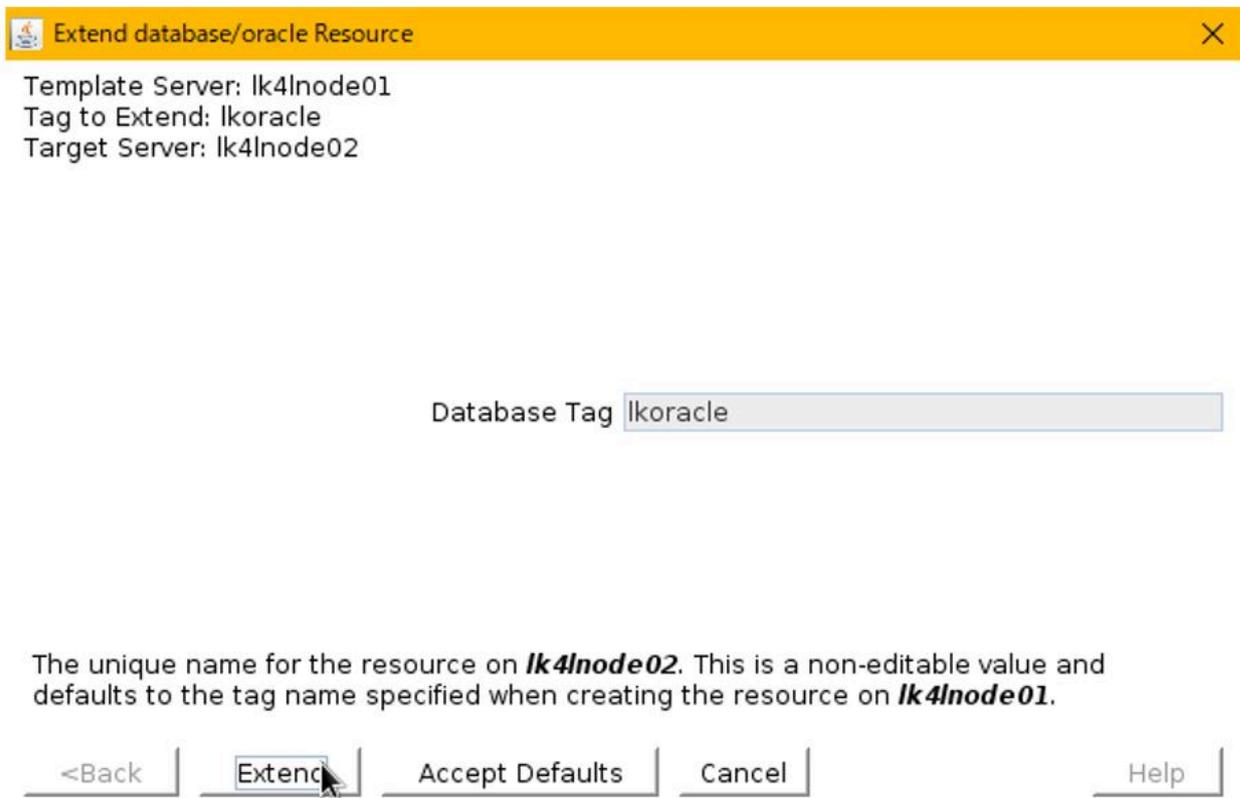
- Enter the following values.

No.	Item	Value to be entered or selected
1	Target Server	lk4lnode02
2	Switchback Type	Intelligent
3	Template Priority	1
4	Target Priority	10

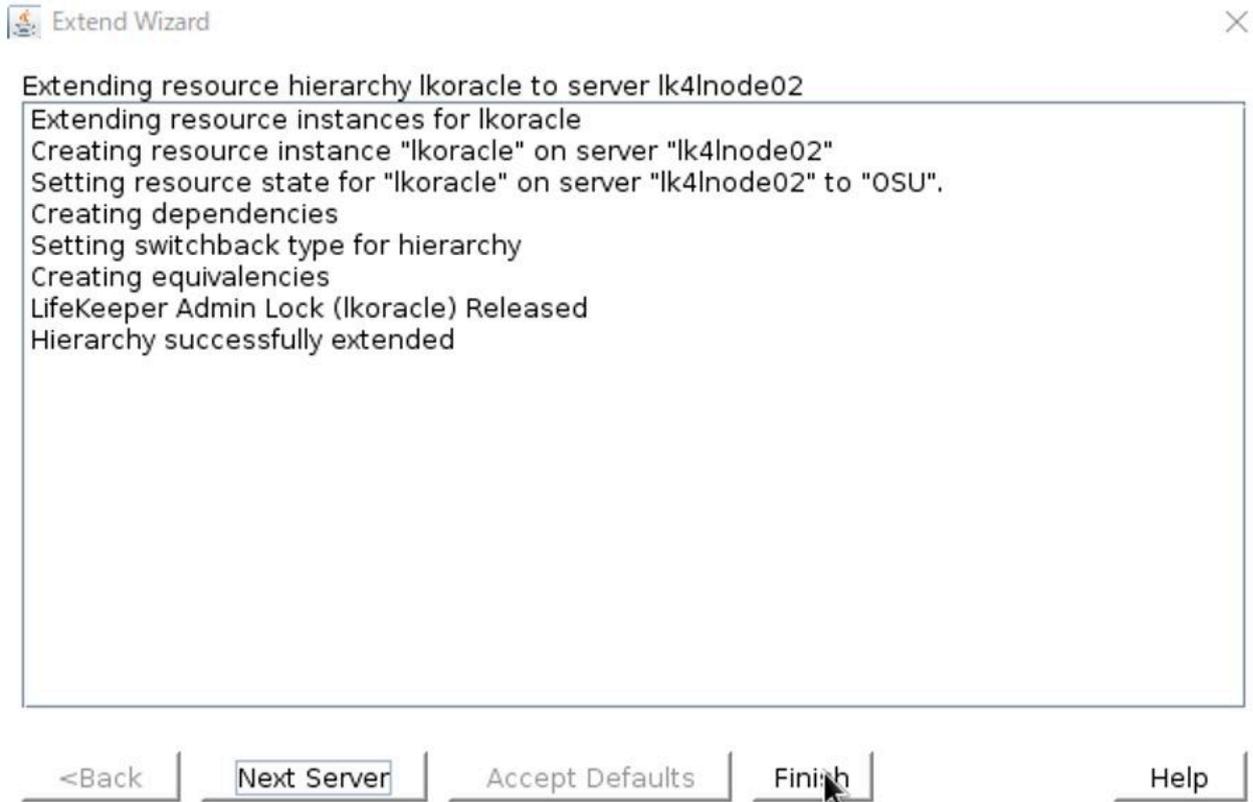
- The Pre-Extend is started. It is successful when “Pre Extend checks were successful” is displayed. Click **Next** to proceed.



8. When the Extend database/oracle Resource wizard appears, click **Extend**.

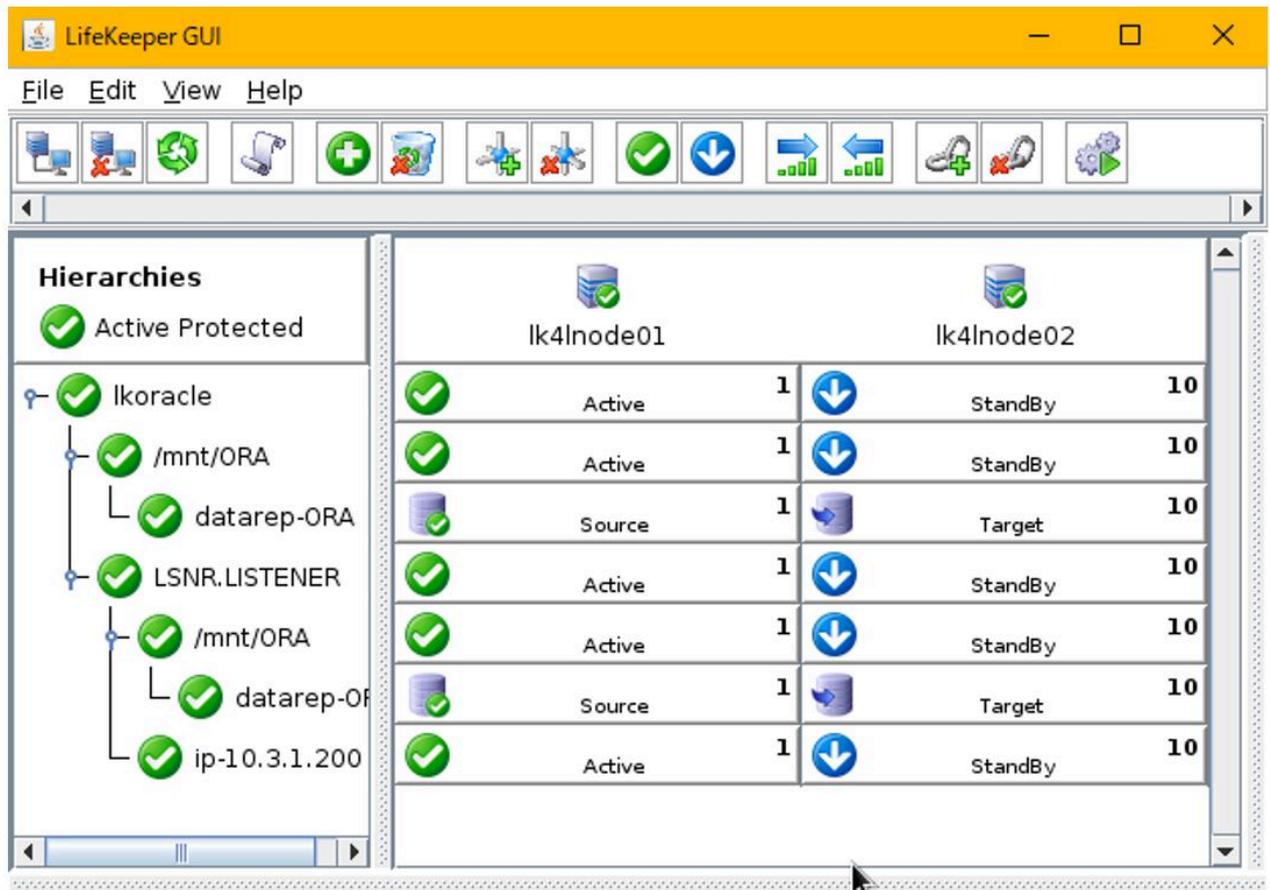


9. Extend is started. It is successful when “Hierarchy successfully extended” is displayed. Click **Finish**.



10. Click **Done** to exit the wizard.
11. The Oracle Database resource is created.

Below is the created resource tree.



12. Undo the modifications in *etc/oratab*.

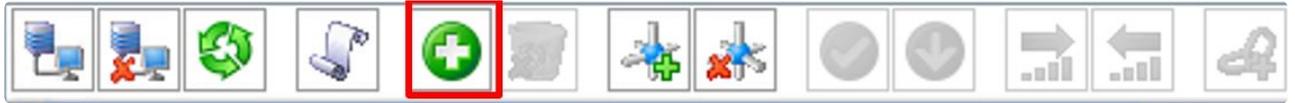
```
# vi /etc/oratab
oracle:/mnt/ORA/app/oracle/product/12.1.0/dbhome_1:N ← change Y to N
```

Refer to [Creating an Oracle Resource Hierarchy](#) for more information.

# 12.2.5.6.7. Creating an Oracle Pluggable Database Resource Hierarchy

Next, create a resource for the Oracle Pluggable database.

1. Click on the [Create Resource Hierarchy] icon to begin resource creation.



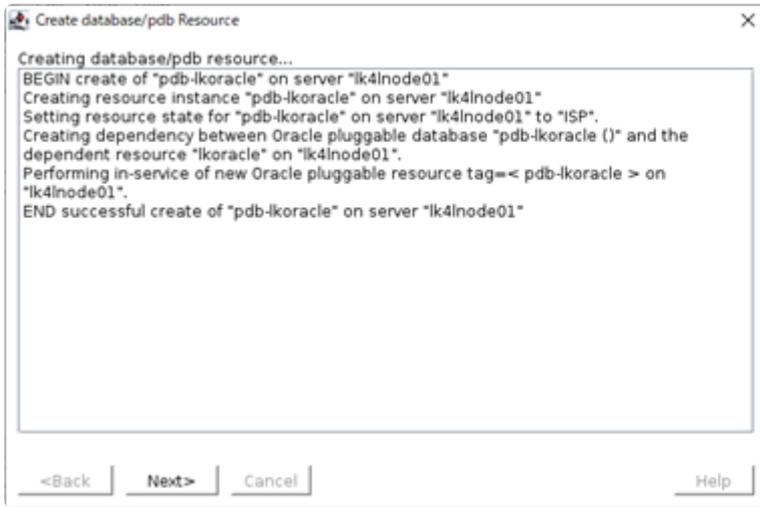
2. The Create Resource wizard will open. Select [Oracle Pluggable Database] for the Recovery Kit in the dropdown menu and click [Next].



3. In the Create Resource wizard, enter the following values.

No	Entry	Value to be entered or selected
1	Switchback Type	Intelligent
2	Server	lk4lnode01
3	Oracle_ SID for Database	lkoracle
4	Select the Oracle PDBs	LKPDB
5	Database Tag	pdb-lkoracle

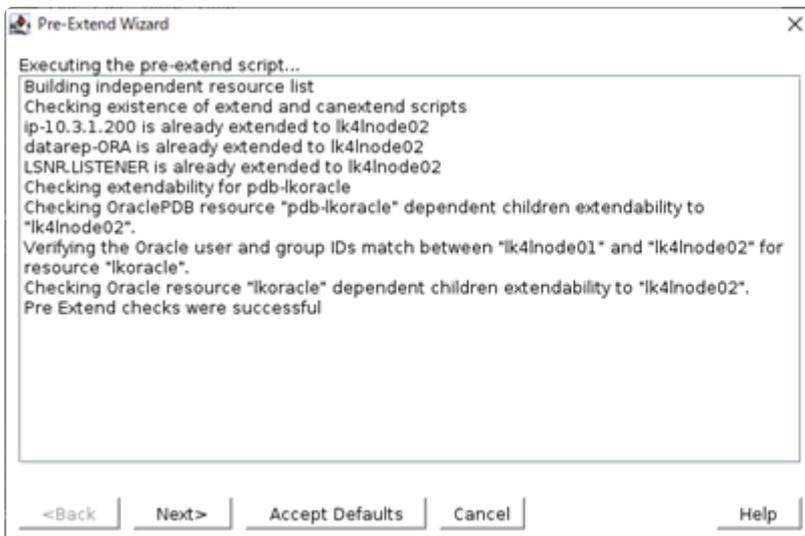
4. Creation of the Oracle Pluggable Database resource will begin.
5. If the message “End of successful Create of...” is displayed, then it is successful. Click [Next] to move to the Pre-Extend Wizard.



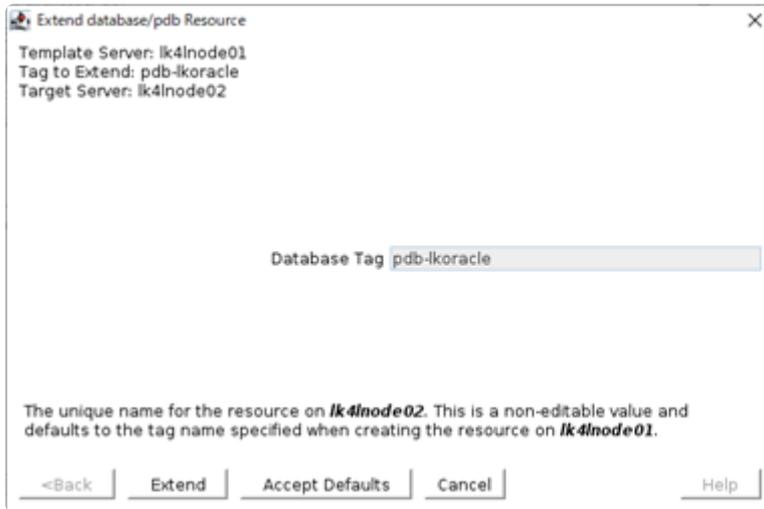
6. Enter the following values.

No	Entry	Value to be entered or selected
1	Target Server	ik4lnode02
2	Switchback Type	Intelligent
3	Template Priority	1
4	Target Priority	10

7. Pre-Extend will begin. If "Pre Extend checks were successful" is displayed, it is successful. Click [Next] to continue.



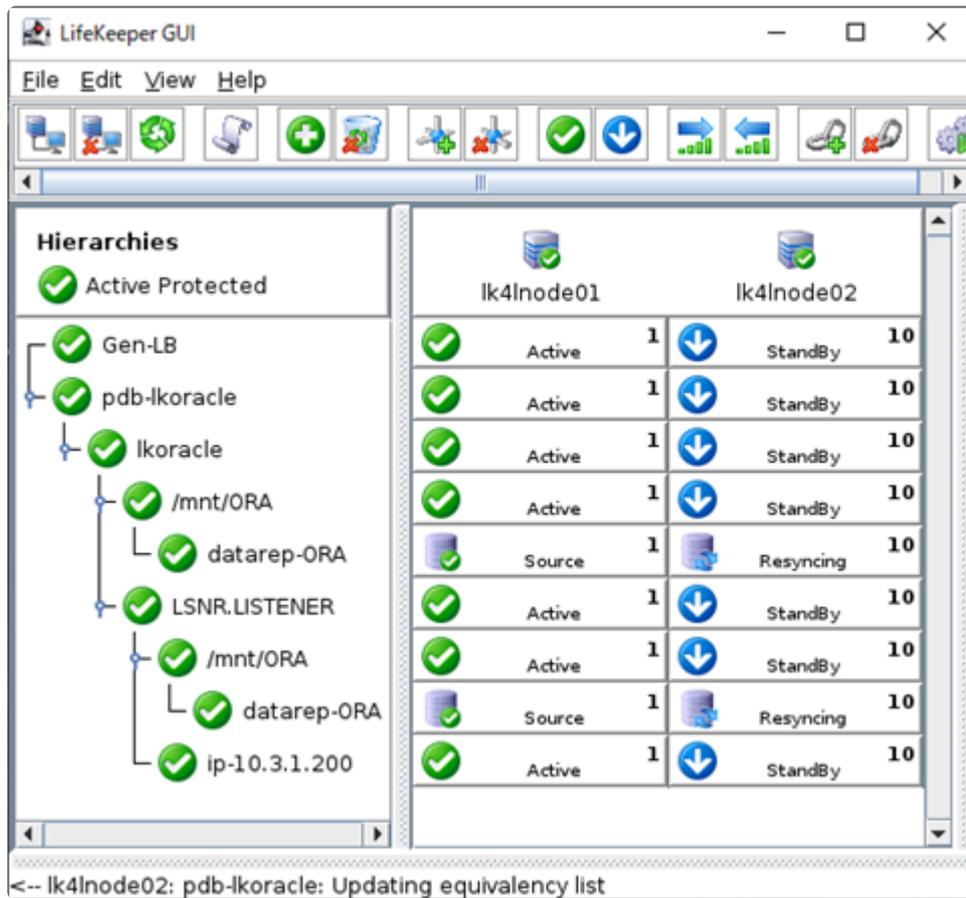
8. When the Extend database/pdb Resource wizard appears, click [Extend].



9. Extend is started. If the message “Hierarchy successful extended” is displayed, it is successful. Click [Finish].



10. Press [Done] to exit the wizard.
11. The Oracle Database resource has been created.  
The completed resource tree.



Refer to [Configuring a Pluggable Database with Oracle Multitenant](#) for more information.

# 12.2.5.6.7.1. Setting Resource Dependencies

Create a dependency between the created GenLB resource and the Oracle Pluggable Database resource.

- Using the LifeKeeper GUI, click on the [Create Dependency] icon to start creating the dependency.



- In the Create Dependency wizard, enter the following values.

Entry	Value
Server	lk4lnode01
Parent Resource Tag	Gen-LB
Child Resource Tag	pdb-lkoracle

- Dependency creation begins and is successful when you see “The dependency creation was successful.” Click [Done].



- The completed resource tree.

LifeKeeper GUI

File Edit View Help

**Hierarchies**  
 Active Protected

- Gen-LB
  - pdb-lkoracle
    - lkoracle
      - /mnt/ORA
        - datarep-ORA
      - LSNR.LISTENER
        - /mnt/ORA
          - datarep-ORA
        - ip-10.3.1.200

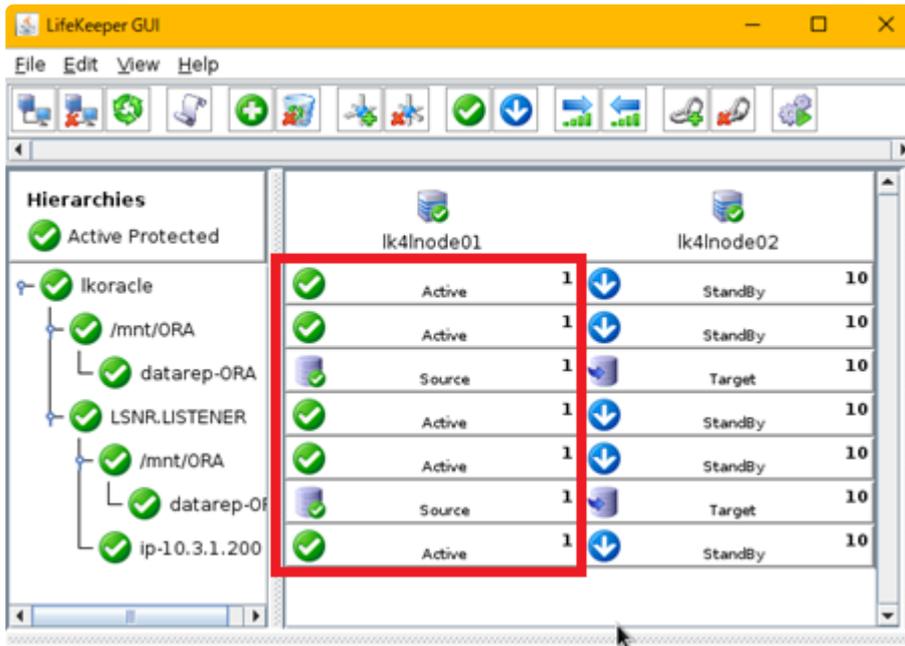
lk4lnode01		lk4lnode02	
Active	1	StandBy	10
Active	1	StandBy	10
Active	1	StandBy	10
Active	1	StandBy	10
Source	1	Resyncing	10
Active	1	StandBy	10
Active	1	StandBy	10
Source	1	Resyncing	10
Active	1	StandBy	10

<-- lk4lnode02: pdb-lkoracle: Adding parent: Gen-LB

# 12.2.5.7. Connectivity Check

Check the connection to the database from the remote client.

1. **Install** the Oracle client software on the client.
2. Ensure that the resource is active on the primary node.



3. Connect to the database via the listener from the remote client using the virtual IP address for connection.

```
# hostname
lk4client
# sqlplus sys/XXXXXXXX@10.3.1.200:1521/ikoracle as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Sat May 25 17:21:32 2019
Version 19.3.0.0.0

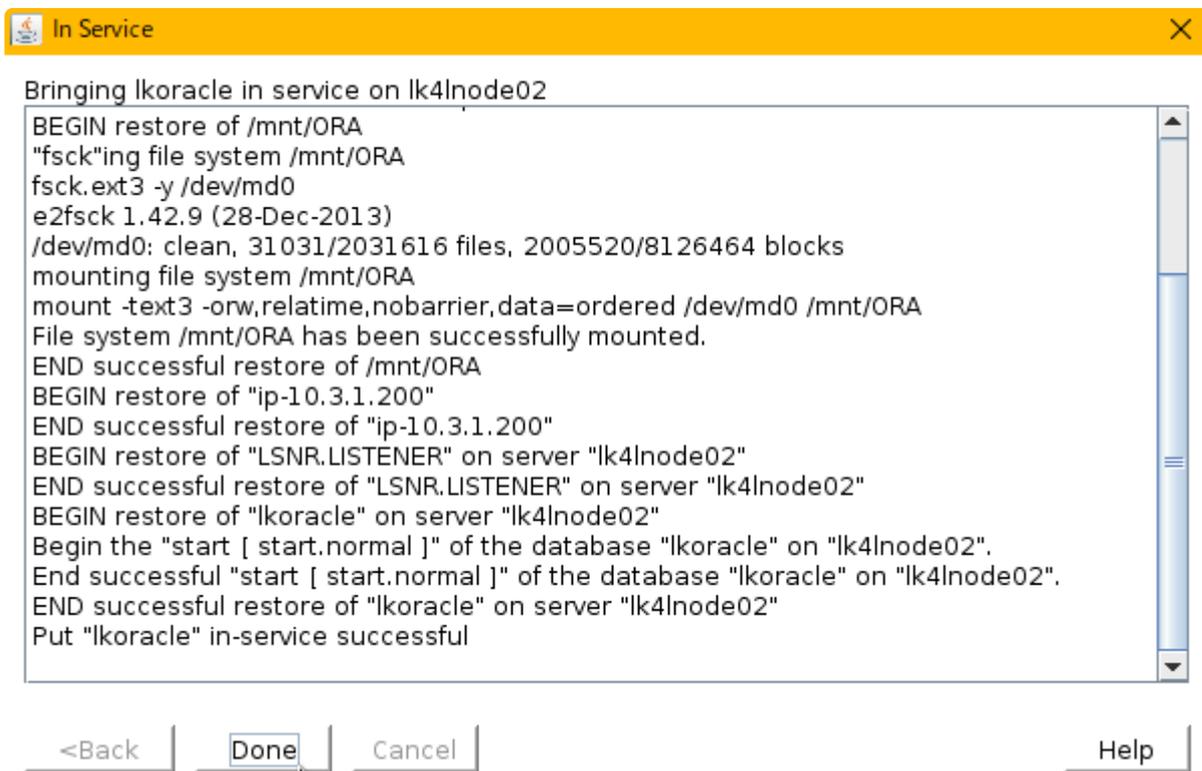
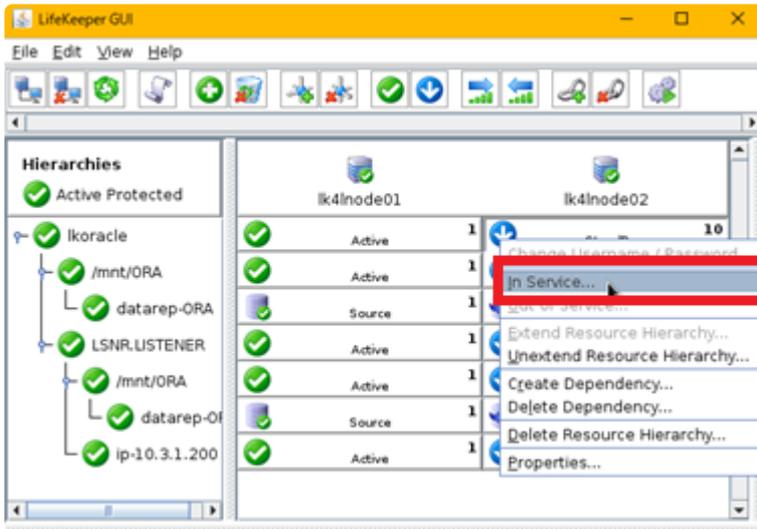
Copyright (c) 1982, 2019, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing options

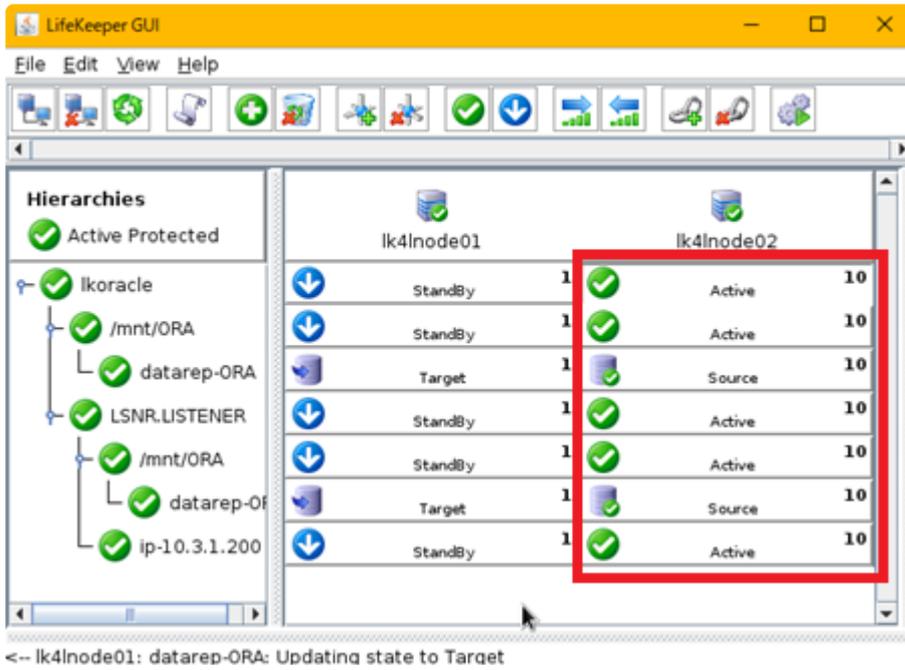
SQL> select instance_name from v$instance ;

INSTANCE_NAME
-----
Ikoracle
```

4. Switch over the Oracle Database resource from the LifeKeeper GUI. Right-click the **Oracle Database** resource and select **In Service**.



5. The switchover is completed.



6. Connect from the remote client via the listener as follows (access with the cluster host name).

```
# hostname
lk4client
# sqlplus sys/XXXXXXXX@10.3.1.200:1521/lkoracle as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Sat May 25 17:21:32 2019
Version 19.3.0.0.0

Copyright (c) 1982, 2019, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing options

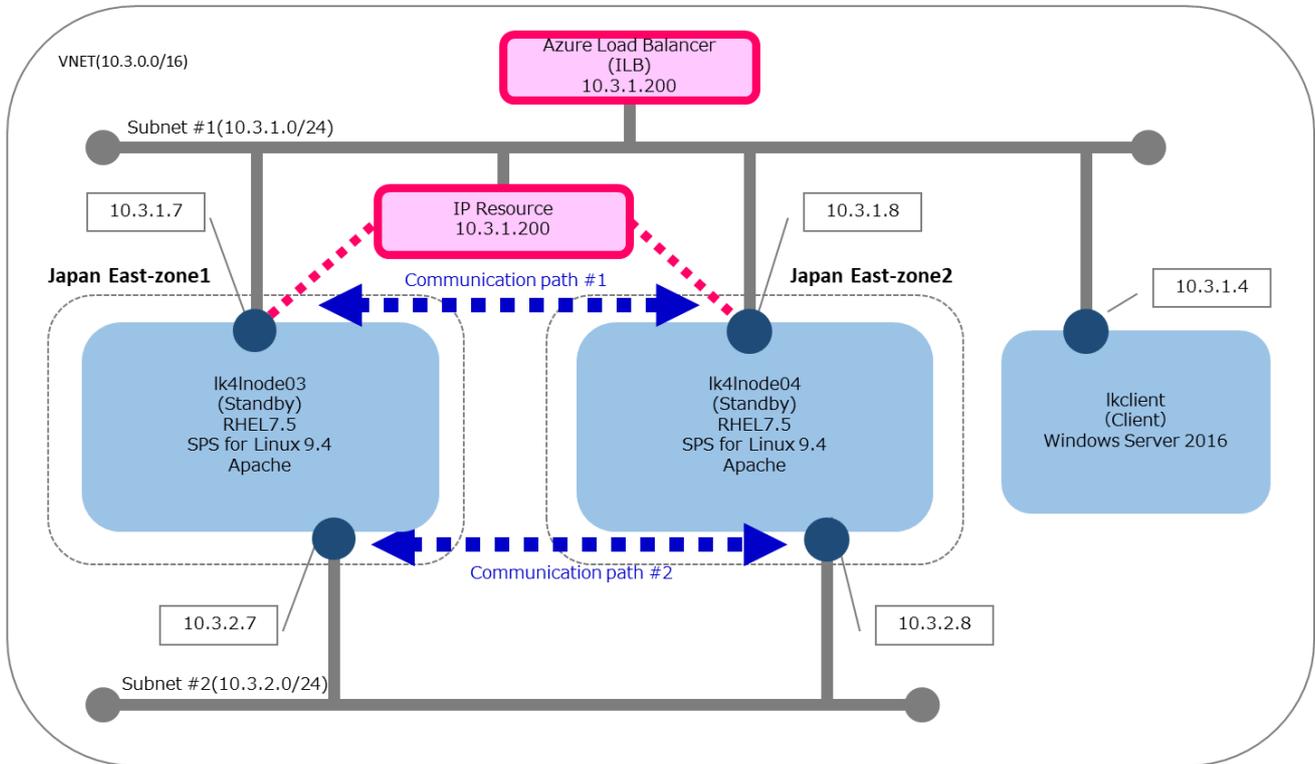
SQL> select instance_name from v$instance ;

INSTANCE_NAME
-----
lkoracle
```

7. It was confirmed that the same IP address (virtual IP address) can be used to connect to Oracle before and after switching.

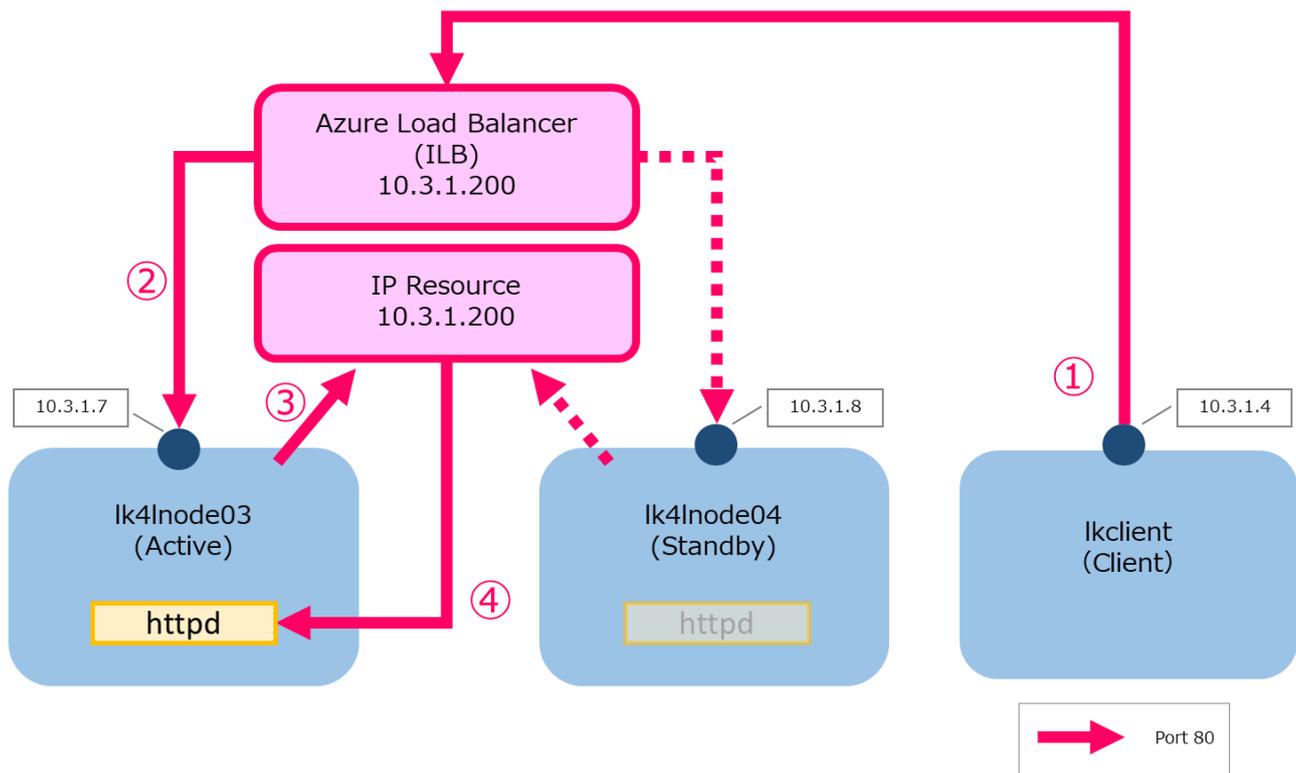
# 12.2.6. Availability Zone (High Availability Zone)

## Configuration for the verification



The application to be protected is Apache and we checked the connection from the client PC with a web browser. The active server (Ik4lnode03) and the standby server (Ik4lnode04) belong to different zones. The http request from the client (Ikclient) is redirected to the active server virtual IP (10.3.1.200) via the ILB probing at 80.

## Overview of client redirection with the ILB



The outline of connecting steps from the client PC is as follows.

1. In order for the client to connect to httpd, start the connection to 10.3.1.200 (port 80) via the ILB.
2. In the ILB 10.3.1.7 (active node: Ik4Inode03) and 10.3.1.8 (standby node: Ik4Inode04) are registered in the load balancing destination (backend pool) without port setting. The packets received by the ILB for 10.3.1.200 (port 80) are forwarded to both the active and standby nodes.
3. Due to LifeKeeper specifications, it is necessary to specify [VIP protected with IP resources] for Apache resources. Create an IP resource (10.3.1.200) for receiving. Since the LifeKeeper IP resource (10.3.1.200) is in service only on the active node, only the active node (10.3.1.7) receives requests.
4. As a result, httpd (10.3.1.200: port80) of the active node receives the connection request from the client.

## 12.2.6.1. Azure Configuration

---

To use the Availability Zone, configure it when you create the VM as follows.

1. For the Region, select (Asia Pacific) Japan East.\*1
2. Select Availability Zone in Availability options.
3. In Availability Zone, select 1, 2, or 3. The active node and the standby node should be in different zones: active (Zone 1) and standby (Zone 2).

\*1 If you select a region such as Japan West where only the Availability Set is available, you cannot select "Availability Zone" in Availability options.

Since the protected application used for verification is Apache, set port 80 to probe in the ILB settings. Other configurations are the same as the one using the Availability Set.

The environment used to verify the configuration using the Availability Zone is as follows.

1. Network, load balancer

### **[Network]**

VNet: 10.3.0.0/16

Subnet:

Subnet1: 10.3.1.0/24

Subnet2: 10.3.2.0/24

### **[Load Balancer]**

Location: Japan East

SKU: Standard

Availability Zone: Zone redundant

FrontendIP:

Type: Private

Subnet: 10.3.1.0/24

IP Address: 10.3.1.200(Static)

BackendPool :

lk4Inode03

lk4Inode04

Health Probe:

Protocol: TCP

Port: 22

Load Balancing Rule:

Port: 80

Backend Port: 80

Inbound NAT rules :

0 rules

Session persistence: None

Floating IP(Direct Server Return): Enable

2. Virtual server

**[Active Node]**

Name: lk4lnode03

Operating system: Linux (RHEL 7.5)

Size: Standard A1 v2 (1 vcpus, 2 GiB memory)

Network Interface:

Private IP: 10.3.1.7

Public IP: 52.156.45.126

Installed Software:

httpd: httpd-2.4.6-80.el7\_5.1.x86\_64

iptables: iptables-1.4.21-24.1.el7\_5.x86\_64

LifeKeeper: LifeKeeper for Linux v9.5.0 Build 108

IP tables:

```
...  
-A PREROUTING -p tcp -m tcp --dport 80 -j DNAT --to-destination 10.3.1.200:80  
-A POSTROUTING -s 10.3.1.200/32 -p icmp -j SNAT --to-source 10.3.1.8  
...  
:OUTPUT ACCEPT [4188:806503]  
-A INPUT -s 10.3.0.0/16 -p tcp -j ACCEPT  
...
```

httpd.conf:

```
...  
Listen 10.3.1.200:80...  
...  
ServerName 10.3.1.200:80  
...
```

### [Standby Node]

Name: lk4lnode04

Network Interface:

Private IP: 10.3.1.8

Public IP: None

IP tables:

```
...  
-A PREROUTING -p tcp -m tcp --dport 80 -j DNAT --to-destination 10.3.1.200:80  
-A POSTROUTING -s 10.3.1.200/32 -p icmp -j SNAT --to-source 10.3.1.7  
...  
:OUTPUT ACCEPT [4188:806503]  
-A INPUT -s 10.3.0.0/16 -p tcp -j ACCEPT  
...
```

Same as the active node for others.

### [Client Node]

Name: lkclient

Operating System: Windows (Windows Server 2012 Datacenter)

Size: Basic A1 (1 vcpus, 1.75 GiB memory)

Network Interface:

Private IP: 10.3.1.4

Public IP: 104.41.171.59

Installed Software:

Google Chrome : Version 80.0.3987.149 (Official Build) (64-bit)

Xming : 6.9.0.31

Tera Team : 4.102

WinSCP : 5.17.2



The screenshot shows two windows on a Windows Server 2012 desktop. The left window is the LifeKeeper GUI, displaying a hierarchy of protected resources and a table of nodes. The right window is a web browser showing a test page for the Azure Multi-AZ environment.

**LifeKeeper GUI Hierarchy:**

- Active Protected
  - apache-etc:htpd
  - /mnt/ik
  - ip-10.3.1.200

**LifeKeeper GUI Node Table:**

Node	State	Priority	Role	Weight
lk4inode03	Active	1	StandBy	10
	Active	1	StandBy	10
	Active	1	StandBy	10
lk4inode04	Active	1	StandBy	10
	Active	1	StandBy	10
	Active	1	StandBy	10

**Test Page Content:**

### Azure Multi-AZ Environment Test Page Ik4Inode03

This page is used to test the proper operation of the LifeKeeper for Linux after it has been installed. If you can read this page, it means that the LifeKeeper for Linux on Azure Multi-AZ Environment is working properly.

**About LifeKeeper for Linux :**

The LifeKeeper product includes fault detection and recovery software that provides high availability for file systems, network addresses, applications and processes running on Linux. LifeKeeper supports the configuration and switchover of a given application across multiple servers. The servers on which the application is configured are assigned priorities to determine the sequence in which the application will move from server to server in the event of multiple failures.

LifeKeeper for Linux provides switchover protection for a range of system resources. Automatic recovery is supported for the following resource types:

- \*Processes and Applications
- \*Shared Storage Devices (Including VMWare virtual hard disks)
- \*Network Attached Storage Devices
- \*File Systems (ext3, ext4, vxfs, xfs and nfs) Note: btrfs is not currently supported by the SIOS Protection Suite for Linux.
- \*Communication Resources (TCP/IP)
- \*Database Applications (Oracle, MySQL, DB2, SAP MaxDB, postgresSQL, EnterpriseDB Postgres Plus Advanced Server, Sybase)

For information on SIOS technology, please visit the [SIOS Technology, Inc. website](#). The documentation for LifeKeeper for Linux is [available on the SIOS Technology, Inc. website](#).

**About Single Server Protection for Linux:**

LifeKeeper Single Server Protection allows for application monitoring in a single node configuration (i.e., no cluster requirements or restraints). Single node environments may be physical or virtual (vSphere, KVM, Amazon EC2). LifeKeeper Single Server Protection is built on the proven and stable architecture of SIOS LifeKeeper. LifeKeeper Single Server Protection provides superior application monitoring and can perform recovery of failed applications and system infrastructure items (e.g., NFS share, IP address, File System). If an application cannot be recovered for some reason, LifeKeeper Single Server Protection will initiate a restart of the node via a system reboot or via a VMWare HA restart for VMWare virtual machines configured for VM and Application Monitoring.

**Log Message:** ik4inode04: datarep-ik: Updating state to Target

**Windows Server 2012**

## 12.2.7. References and Acknowledgements

---

### SIOS Technology:

[SIOS Technical Documentation](#)

[LifeKeeper for Linux 9.5.2 Technical Documentation](#)

[Oracle Recovery Kit Administration Guide](#)

### External websites:

Microsoft Azure

<https://azure.microsoft.com/en-us/>

Microsoft Azure Linux Virtual Machine Document

<https://docs.microsoft.com/en-us/azure/virtual-machines/linux/>

Microsoft Azure Blog

<https://azure.microsoft.com/en-us/blog/>

High Availability for a file share using WSFC, ILB and 3rd-party Software SIOS DataKeeper (used as a reference in ILB settings)

<http://azure.microsoft.com/blog/2014/11/11/high-availability-for-a-file-share-using-wsfc-ilb-and-3rd-party-software-sios-datakeeper/>

### Acknowledgements

Microsoft Japan approved the use of Microsoft Azure for the verification. We appreciate their great support in building and configuring the environment and creating the test cases.

Microsoft Azure Website

<https://azure.microsoft.com/en-us/>

## 12.3. MySQL Cluster with Data Replication ("Shared Nothing" Cluster)

---

### Objective

This document is intended to aid you in installing, configuring and using the LifeKeeper for Linux evaluation product to make MySQL highly available. If MySQL is not already installed, please allocate some time to install it on your servers. Once this task has been completed, you may install and configure LifeKeeper for Linux.

There are five phases in this process:

- Prepare to Install
- Configure Storage
- Install and Configure MySQL
- Install LifeKeeper for Linux
- Configure your LifeKeeper Cluster
- Test Your Environment

## 12.3.1. Terms to Know

---

The following terms are used throughout this document and, while some may be familiar to you, it may be helpful to review how SIOS defines and uses these terms.

### Network Communication Terms

**Crossover cable** – A cable used to directly connect computing devices together, instead of being connected to a network switch, hub or router. This cable creates an isolated, private network to allow cluster-related and data replication traffic to flow between systems.

### Types of LifeKeeper Servers

**Server** – A computer system dedicated to running software application programs.

**Active Server** – This is the server where the resource hierarchy is currently running (IN SERVICE).

**Standby Server** – This is the server where the resource hierarchy is defined, but is not currently running. This server is available to bring the resource hierarchy into service should something happen to the resource hierarchy on the Active Server.

**Primary Server** – This is the server in a LifeKeeper configuration with the highest priority for a given resource hierarchy. It is the server that provides services for the resource hierarchy under normal circumstances.

**Secondary Server** – This is the server in a LifeKeeper configuration with the 2nd highest priority for a given resource hierarchy.

**Source Server** – In a LifeKeeper cluster, using data replication, this is the Active Server. It is where the resource hierarchy is currently running and the replicated partition (Source Partition) is accessible for writes.

**Target Server** – In a LifeKeeper cluster, using data replication, this is the Standby Server. The replicated partition (Target Partition) is updated with writes from the Source Partition by the SIOS Data Replication system. This partition should not be accessed/modified manually.

### SIOS Data Replication Terms

**Replication** – Transferring data from one partition to another via a sector-by-sector copy. During replication, the target partition should not be accessed or modified assuring your data integrity.

**Synchronous** – A replication scheme in which the data is confirmed written and valid on the target before the write operation occurs on the source disk through a series of information exchanges. Synchronous mirrors should only be implemented on high speed (100Mbps+) networks due to the network overhead involved.

**Asynchronous** – A replication scheme in which the data is released for writing on the source immediately and is sent to the target(s) simultaneously for writing as fast as the data can get there and can be written on them.

**Rate of Change** – A measure of the amount of data which is changing over a set period of time.

**Compression** – An algorithm which is optionally implemented to reduce the amount of traffic between source and target nodes. Nine levels of compression are offered. Compression is turned off by default.

**Throttling** – An optionally implemented mechanism to limit the bandwidth used for replication.

## LifeKeeper Product Terms

**Communications Path** – A mechanism supporting communication between nodes in a LifeKeeper cluster. SIOS highly recommends implementing multiple communication paths between all servers in the cluster to eliminate a single point of failure.

**Heartbeat** – A periodic message exchanged between nodes in a LifeKeeper cluster that provides server health monitoring. A heartbeat message is one type of inter-node cluster communication sent over a communications path.

**Split Brain** – A situation in which all communications paths between cluster members fail, but all servers remain up and running. In this situation, both systems believe the other has failed and both believe they should keep or bring resources into service.

**Failover** – The unplanned migration of a resource hierarchy to the Standby Server because of a system or resource failure on the Active Server.

**Switchover** – The planned migration of a resource hierarchy from the Active Server to the Standby Server.

**Switchback** – The setting that governs the recovery behavior of the server where the resource was in service when it failed. If the setting is “intelligent”, the server acts as a possible backup for the given resource. If the setting is “Automatic”, the server actively attempts to re-acquire the resource without further notice.

**Resource** – A system asset that can be protected by LifeKeeper. Resources can be used to represent disk partitions, virtual IP addresses, applications, etc.

**Extend a Resource** – Create or define an already configured LifeKeeper resource onto another server in the cluster and build an equivalency relationship that prevents the resource from coming in service on both systems simultaneously.

**Resource Hierarchy** – A grouping of resources, in a predetermined order, from high to low. This may also be referred to as simply a Hierarchy.

**Shared Storage** – One or more logical disk partitions that are physically attached to all nodes in a cluster. LifeKeeper ensures that the volume is only accessible by one server at a time. This is formally

called I/O fencing.

**Data Replication (Disk Mirroring)** – The replication of logical disk partitions to separate physical hard disks in real time to ensure continuous availability, currency and accuracy of data.

**Source** – The partition on the source server used for replication. The “gold” copy of the data.

**Target** – The partition on the target server used for replication.

**Switchable IP Address** – A unique IP address that may be moved between systems in the cluster. Client systems connect to this address and the system where the virtual IP resource is active will respond to requests.

## 12.3.2. The Evaluation Process – MySQL Cluster

---

SIOS strongly recommends performing your evaluation of LifeKeeper for Linux in a test lab environment. SIOS is not responsible and cannot provide support for evaluation software installed in a production environment.

All questions during the evaluation period should be directed to [evalsupport@us.sios.com](mailto:evalsupport@us.sios.com) or your local Pre-Sales Engineering contact. Pre-sales support will contact you by the next business day to answer questions. Once you are a licensed customer on software maintenance, you will have access to 24 × 7 post-sales technical support.

 **Important:** Your evaluation license is valid for a limited period of time from the day you receive the SIOS product evaluation package and licenses via email from the SIOS sales team.

## 12.3.3. Prepare to Install

---

### Hardware Requirements

#### Primary and Secondary Servers

- Systems must meet the minimum requirements for the Linux distribution to be used during the evaluation.
- 512MB RAM minimum; 1GB RAM recommended.
- 2GB of available hard disk space recommended.
- Multiple Network Interface Cards (NIC's) are recommended.
- Configure one or more additional partitions to be used for data replication. On the primary server, these will become the source partitions. On the secondary server(s), these will become the target partitions. In this evaluation example we will be replicating MySQL data (which will be a partition mounted at /var/lib/mysql)
- For replicated partitions, a target partition's size must equal to or larger than the size of its source partition.
- The system ( / ) and boot (/boot) partitions are not eligible for replication.

 **Note:** You may use more than one partition for replicated data, allowing for separation of multiple database files and/or log directories for performance reasons.

#### Client

This system is not required but is recommended for testing the cluster environment.

- A standard linux terminal running the MySQL client can be used to test the configuration.

### Software Requirements

#### Primary Server and Secondary Server

- Linux Distribution x86\_64, AMD 64:
  - RedHat Enterprise Linux 5 (5.4+ recommended) or 6.x
  - CentOS Linux 5 (5.4+ recommended) or 6.x
  - Oracle Enterprise Linux 5 (5.4+ recommended), 6.3, 6.4 (**RedHat Compatibility Kernel**)

**Only)**

- SuSE Linux Enterprise Server 10 or 11 (11 recommended)
- See [Linux Release Notes](#) for a full list of supported Operating Systems
- Current patches / security updates are recommended
- Satisfied dependencies; especially if the Linux installation package selection was base/minimal you will need to refer to the dependencies documentation at [Linux Dependencies](#)
- It is recommended that IPTables is disabled
  - # /etc/init.d/iptables off
  - # chkconfig iptables off
  - See [Running LifeKeeper With a Firewall](#) for information regarding the ports LifeKeeper for Linux uses.
- Disable SELinux :
  - Edit /etc/selinux/config
  - Set SELINUX=disabled (note: permissive mode is also acceptable)
- Check the configuration of your /etc/hosts file
  - localhost.localdomain and localhost are the only entries that can be on 127.0.0.1
  - Create a separate entry for your hostname with a static address
- GUI Authentication with PAM
  - LifeKeeper for Linux now leverages the Pluggable Authentication Module (PAM) provided in the Linux Standard Base (LSB).
  - Users are identified and authenticated against the system's PAM configuration. Privilege levels are determined from group membership as provided through PAM.
  - In order to access the GUI, a user must be a member in one of the three LifeKeeper groups: lkadmin, lkoper or lkguest.
  - See the following URL for more information on this topic:

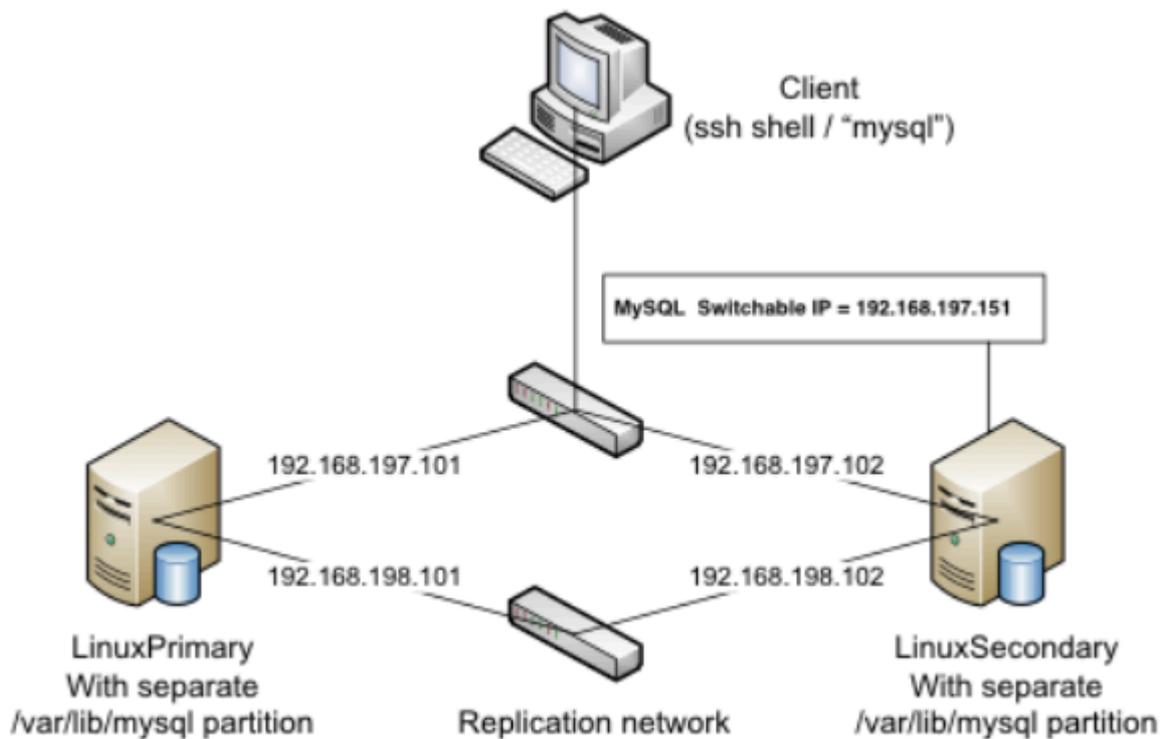
- See [Configuring GUI Users](#) for more information.

## Network Requirements

For your evaluation, we recommend configuring your machines similarly to the following example.

LinuxPrimary and LinuxSecondary are multi-homed, between two LAN segments (the second NIC in each server could even be connected via a cross-over cable if a second physical network is not available). The second NIC is optional in this configuration, but highly recommended in production environments to avoid a single point of failure.

In this example evaluation scenario we will be leveraging local, replicated storage.



**Network Configuration Example**

## Primary Server and Secondary Servers

- Configure the Host file with entries for all LifeKeeper protected servers. This is typically /etc/hosts.

Example:

```
192.168.197.101    LinuxPrimary
192.168.197.102    LinuxSecondary
```

- See your Network Administrator to obtain an unused IP Address to be used as the switchable IP Address. This switchable IP Address will be created later in the configuration process.
- Public Network connection(s) configured with:

- Static IP address
  - Correct subnet mask
  - Correct gateway address
  - Correct DNS server address(es)
- 
- Private Network connection(s) configured with:
    - Static IP address (on a different subnet from the public network)
    - Correct network mask
    - No gateway IP address
    - No DNS server addresses

## Client

Must be able to communicate on the same subnet/network as the servers Public interface addresses. In our example, this is the 192.168.197.0/24 network.

## 12.3.4. Configure Storage

---

### Before You Begin

#### Ensure the following:

- You have an extra disk/partition on both servers that can be used for data replication. A target volume's size must be equal to or larger than the size of its source disk/partition.

### Partition local storage for use with SIOS Data Replication

#### Primary Server

On your Primary server, perform the following actions:

1. Identify an existing free, unused disk partition to use as our Apache repository. Alternatively, create a new partition. Use the “gdisk” utility to partition your disk appropriately. In this example /dev/sdb is an unused disk where we will create a single partition
  - a. `gdisk /dev/sdb`
  - b. Press “n” to create a new partition
  - c. This example uses a new disk, so we will use all default values (Partition 1, entire disk and Linux filesystem partition type) Hit Enter four times to confirm these parameters.
  - d. Press “w” to write the partition table
  - e. Press “Y” to confirm to overwrite existing partitions

#### Example

```
[root@LinuxPrimary ~]# gdisk /dev/sdb
```

```
GPT fdisk (gdisk) version 0.8.10
```

```
Partition table scan:
```

```
MBR: protective
```

```
BSD: not present
```

```
APM: not present
```

```
GPT: present
```

```
Found valid GPT with protective MBR; using GPT.
```

```
Command (? for help): n
```

```
Partition number (1-128, default 1): <enter>
```

First sector (34-2047, default = 34) or {+-}size{KMGTP}: **<enter>**  
 Last sector (34-2047, default = 2047) or {+-}size{KMGTP}: **<enter>**  
 Current type is 'Linux filesystem'  
 Hex code or GUID (L to show codes, Enter = 8300): **<enter>**  
 Changed type of partition to 'Linux filesystem'

Command (? for help): **w**

Final checks complete. About to write GPT data. THIS WILL OVERWRITE EXISTING PARTITIONS!!

Do you want to proceed? (Y/N): **Y**  
 OK; writing new GUID partition table (GPT) to /dev/sdb.  
 Warning: The kernel is still using the old partition table.  
 The new table will be used at the next reboot.  
 The operation has completed successfully.

[root@LinuxPrimary ~]#

2. Format the newly created disk partition

```
# mkfs.ext3 /dev/sdb1
```

3. Mount the partition temporarily at /mnt

```
# mount /dev/sdb1 /mnt
```

4. Move any existing data from /var/lib/mysql/ into this new disk partition (assumes a default MySQL configuration)

```
# cd /var/lib/mysql
# mv * /mnt
```

5. Remount /dev/sdb1 at /var/lib/mysql

```
# cd /root
# umount /mnt
# mount /dev/sdb1 /var/lib/mysql
```

6. Note: there is no need to add this partition to /etc/fstab. Lifekeeper will take care of mounting this automatically.

### Result:

```
[root@LinuxPrimary ~]# df /var/lib/mysql
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/sdb1 253855 11083 229666 5% /var/lib/mysql<
```

## Secondary Server

7. On your Secondary server, only perform Step #1 above, where you partition the disk. The size of the Target disk/partition needs to be the same size, or greater, than our Source disk/partition.

## 12.3.5. Install, Configure and Start MySQL

### Primary Server

On your Primary server, perform the following actions:

1. Install both the “mysql” and “mysql-server” rpm packages if they do not exist on your system.

Apply any required dependencies as well

```
# yum install mysql mysql-server
```

2. Verify that your local disk partition is still mounted at /var/lib/mysql via the “df” command

3. If this is a fresh MySQL install, initialize a sample MySQL database:

```
# /usr/bin/mysql_install_db --datadir="/var/lib/mysql" --user=mysql
```

4. Ensure that all files in your MySQL data directory (/var/lib/mysql) have correct permissions and ownership

```
# chown -R mysql:mysql /var/lib/mysql
```

```
# chmod 755 /var/lib/mysql
```

5. Finally, manually start the MySQL daemon from the command line.

 **Do Not** start it via the “service” command, or the /etc/init.d/ scripts

```
# mysqld_safe --user=root --socket=/var/lib/mysql/mysql.sock --port=3306 --datadir=/var/lib/mysql --log &
```

6. Verify MySQL is running by connecting with the mysql client:

```
[root@LinuxPrimary mysql]# mysql
```

```
Welcome to the MySQL monitor. Commands end with ; or \g.
```

```
Your MySQL connection id is 2
```

```
Server version: 5.0.77-log Source distribution
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.
```

```
mysql> exit
```

Bye

```
[root@LinuxPrimary mysql]#
```

7. Update the root password for your mysql configuration. In this example we set the MySQL root password to "SteelEye"

```
# echo "update user set Password=PASSWORD where User='root'; flush privileges" |  
mysql mysql
```

8. Verify your new password:

```
# mysql mysql -u root -p
```

(Enter "SteelEye" as the password)

```
#exit
```

9. Create a MySQL configuration file. We will place this in the same shared directory (/var/lib/mysql/my.cnf)

```
# vi /var/lib/mysql/my.cnf
```

### Example

```
# cat /var/lib/mysql/my.cnf
```

```
[mysqld]
```

```
datadir=/var/lib/mysql
```

```
socket=/var/lib/mysql/mysql.sock
```

```
pid-file=/var/lib/mysql/mysqld.pid
```

```
user=root
```

```
port=3306
```

```
# Default to using old password format for compatibility with mysql 3.x
```

```
# clients (those using the mysqlclient10 compatibility package).
```

```
old_passwords=1
```

```
# Disabling symbolic-links is recommended to prevent assorted security risks;
```

```
# to do so, uncomment this line:
```

```
# symbolic-links=0

[mysqld_safe]

log-error=/var/log/mysql.log

pid-file=/var/run/mysql/mysql.pid

[client]

user=root

password=SteelEye
```

10. Delete the original MySQL configuration file, located in /etc

```
# rm /etc/my.cnf
```

## Secondary Server

On your Secondary Server:

1. Install both the “mysql” and “mysql-server” rpm packages if they do not exist on your system.

Apply any required dependencies as well

```
# yum install mysql mysql-server
```

2. Ensure that all files in your MySQL data directory (/var/lib/mysql) have correct permissions and ownership

```
# chown -R mysql:mysql /var/lib/mysql
```

```
# chmod 755 /var/lib/mysql
```

3. There is no need to perform any of the additional steps taken on the Primary Server

## 12.3.6. Install LifeKeeper for Linux – MySQL Cluster

---

For ease of installation, SIOS has provided the LifeKeeper for Linux with an installation script. Towards the end of the script, the desired Application Recovery Kits (ARKs) should be selected for installation. The software will be installed to the following locations:

SPS for Linux Component	Install Location
LifeKeeper Software	/opt/LifeKeeper
LifeKeeper Config File	/etc/default/LifeKeeper

Perform the following actions on **both** Primary and Secondary server.

### Download Software

1. Open the LifeKeeper for Linux evaluation email you received from SIOS.
2. Download the LifeKeeper for Linux Software from the link provided in your email. It is generally easiest to use “wget” to recursively download all files. Example:
  - a. # cd /root
  - b. # wget -r <URL>
  - c. After you have successfully downloaded the software you will have downloaded contents similar to the follow directory listing:

```
[root@LinuxPrimary ~]# ls -l <directory> total 63680
-rw-r--r-- 1 root root 23163 May 30 14:03 EULA.pdf
-rw-r--r-- 1 root root 536 May 30 14:03 readme.txt
-rw-r--r-- 1 root root 65179648 May 30 14:03 sps.img
```

3. Download your Evaluation license key from the link specified in your evaluation email. Save the license file to an easy to remember location on both servers.

### Run the LifeKeeper Installer Script

1. Loopback mount the sps.img file previously downloaded, which is an ISO9660 image file. Run the “setup” script inside:
  - a. # mount -o loop sps.img /mnt

b. # cd /mnt

c. # ./setup

2. During this procedure, you will hit Enter in most cases to accept default values and continue to the next screen. Note the following exceptions:
  - a. On the screen titled “High AvailabilityNFS” you may select “n” as in this particular eval guide we will not be creating a highly available NFS server cluster configuration.
  - b. If you have plans to create a highly available NFS service, adjust your response accordingly.
3. Towards the end of the setup script, you can choose to install a trial license key now, or later. We will install the license key in the next step, so you can safely select “n” at this point.
4. In the final screen of the “setup” select the DataKeeper from the list displayed on the screen.
5. The following RPMs should be installed:
  - a. steeleye-lkSQL-<version>.noarch.rpm
  - b. steeleye-lkDR-<version>.noarch.rpm
6. Un-mount the Distribution Enabling disk image:
  - a. # cd /root
  - b. # umount /mnt

## Install the Evaluation License Keys

The last phase of the setup process installs the licensing keys. You must install the evaluation license key file (“.lic”) that you downloaded with your evaluation software before starting the LifeKeeper for Linux.

1. To install your trial license key, run the “lkkeyins” command on both Primary and Secondary Server. This command is located at /opt/LifeKeeper/bin/lkkeyins . Example:

```
# /opt/LifeKeeper/bin/lkkeyins <path_to_license/<filename>.lic
```

2. Validate your license keys were installed via the /opt/LifeKeeper/bin/lklicmgr command

```
# /opt/LifeKeeper/bin/lklicmgr
License File: 20101230.lic
```

Product	Type	Expiry
LifeKeeper for Linux	Eval	27 Mar 2013 (87 days)

SIOS Data Replication ARK	Eval	27 Mar 2013 (87 days)
MySQL Recovery Kit	Eval	27 Mar 2013 (87 days)

## Start the LifeKeeper for Linux

1. Start

```
# /opt/LifeKeeper/bin/lkstart
```

## 12.3.7. Configure the Cluster

---

### Primary Server

Complete the following steps on the primary server to configure the cluster:

- Create TCP Communication (Comm) Path(s)
- Verify the Communication (Comm) Path(s)

Before you begin, SIOS recommends at least two TCP communications paths between each server within the cluster to each remote server for heartbeat redundancy.

 **Important** Supported configurations require that you define redundant comm. paths, so that the failure of a single communication line will not cause a split brain where resource hierarchies may come in-service on multiple servers simultaneously.

### Access the LikeKeeper GUI

The LifeKeeper Graphical User Interface (GUI) is a Java based application that can be run as a native Linux application.

The LifeKeeper GUI is based on Java RMI with callbacks. Hostnames must be resolvable or you may receive a Java 115 or 116 error.

1. Verify that both short and fully qualified hostnames of all cluster nodes resolve to the proper locations

```
# ping LinuxPrimary
```

```
# ping LinuxPrimary.domain.com
```

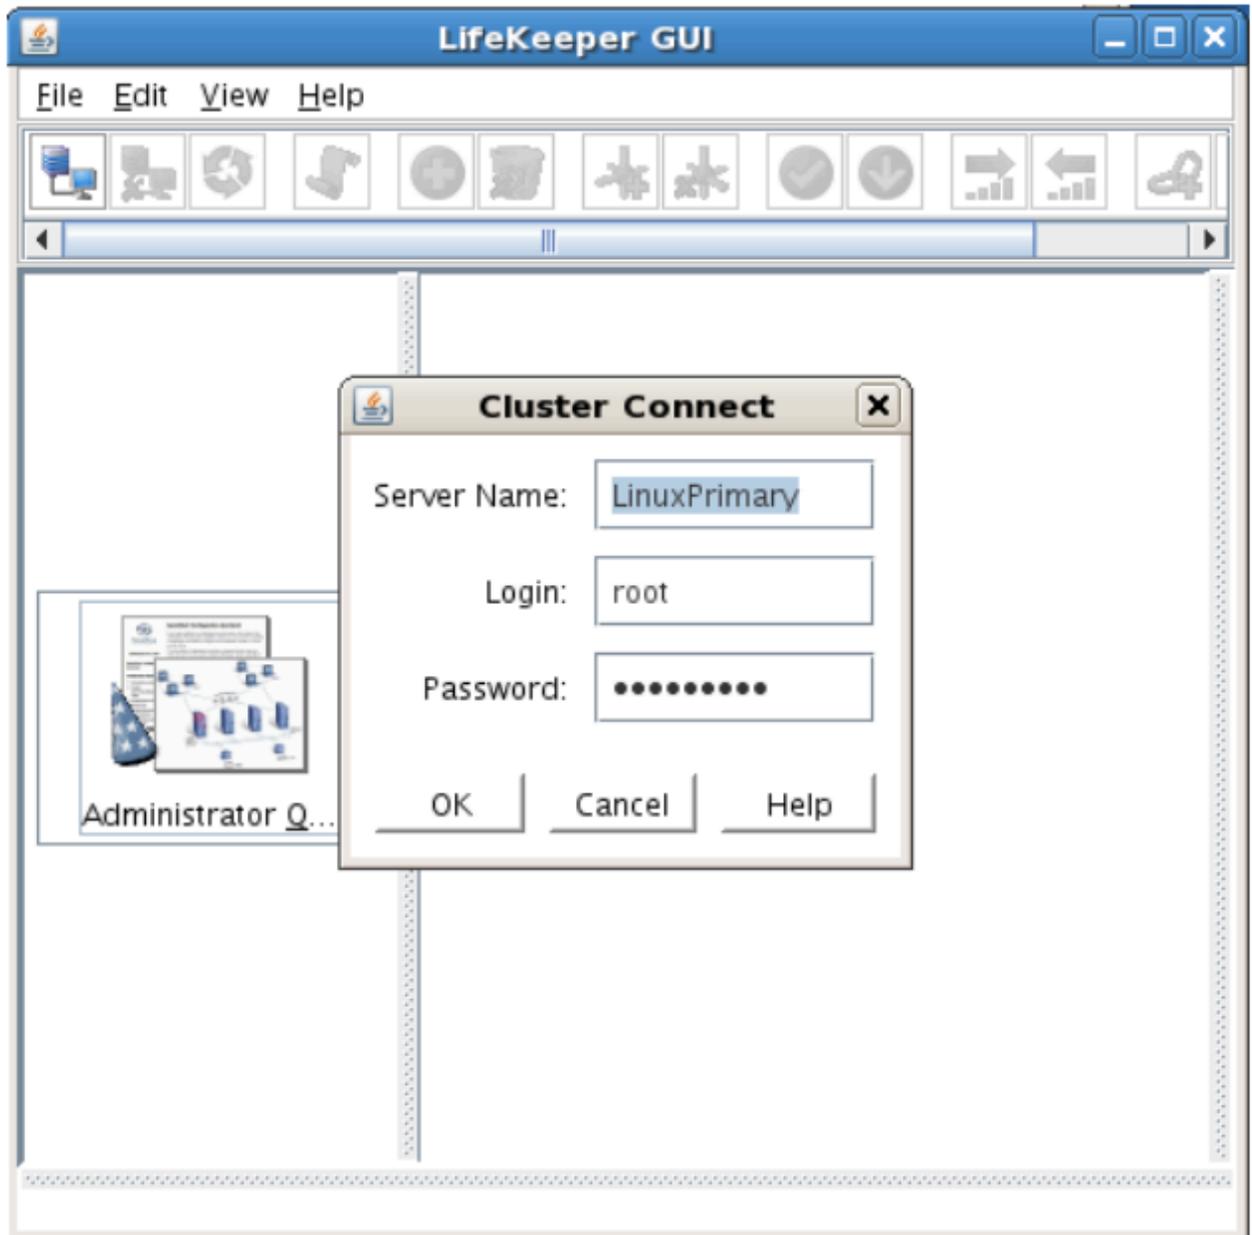
```
# ping LinuxSecondary
```

```
# ping LinuxSecondary.domain.com
```

2. To start the LifeKeeper Linux GUI Application:

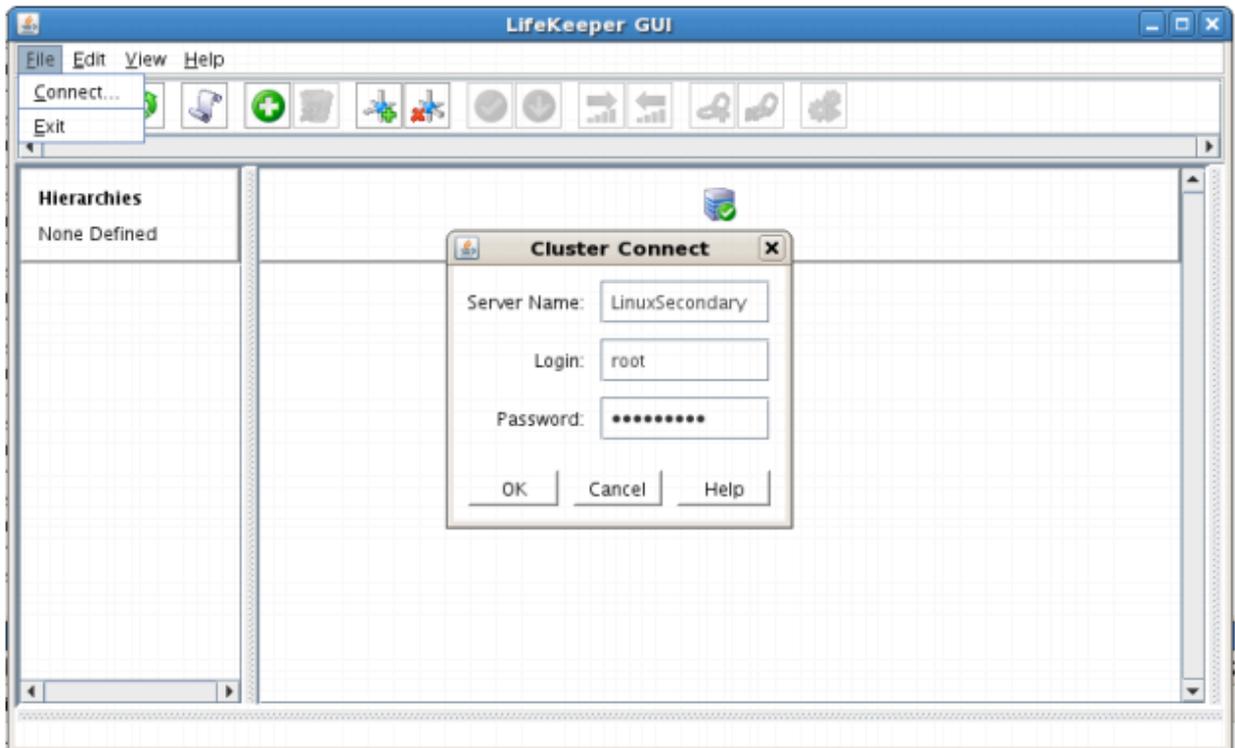
```
/opt/LifeKeeper/bin/lkGUIapp &
```

3. Enter the name of the server you wish to connect to (this field will be populated with the name of the server you are on, if you are running the GUI from a server with LifeKeeper installed) along with your root credentials and click OK.



## Create Communication (Comm) Paths

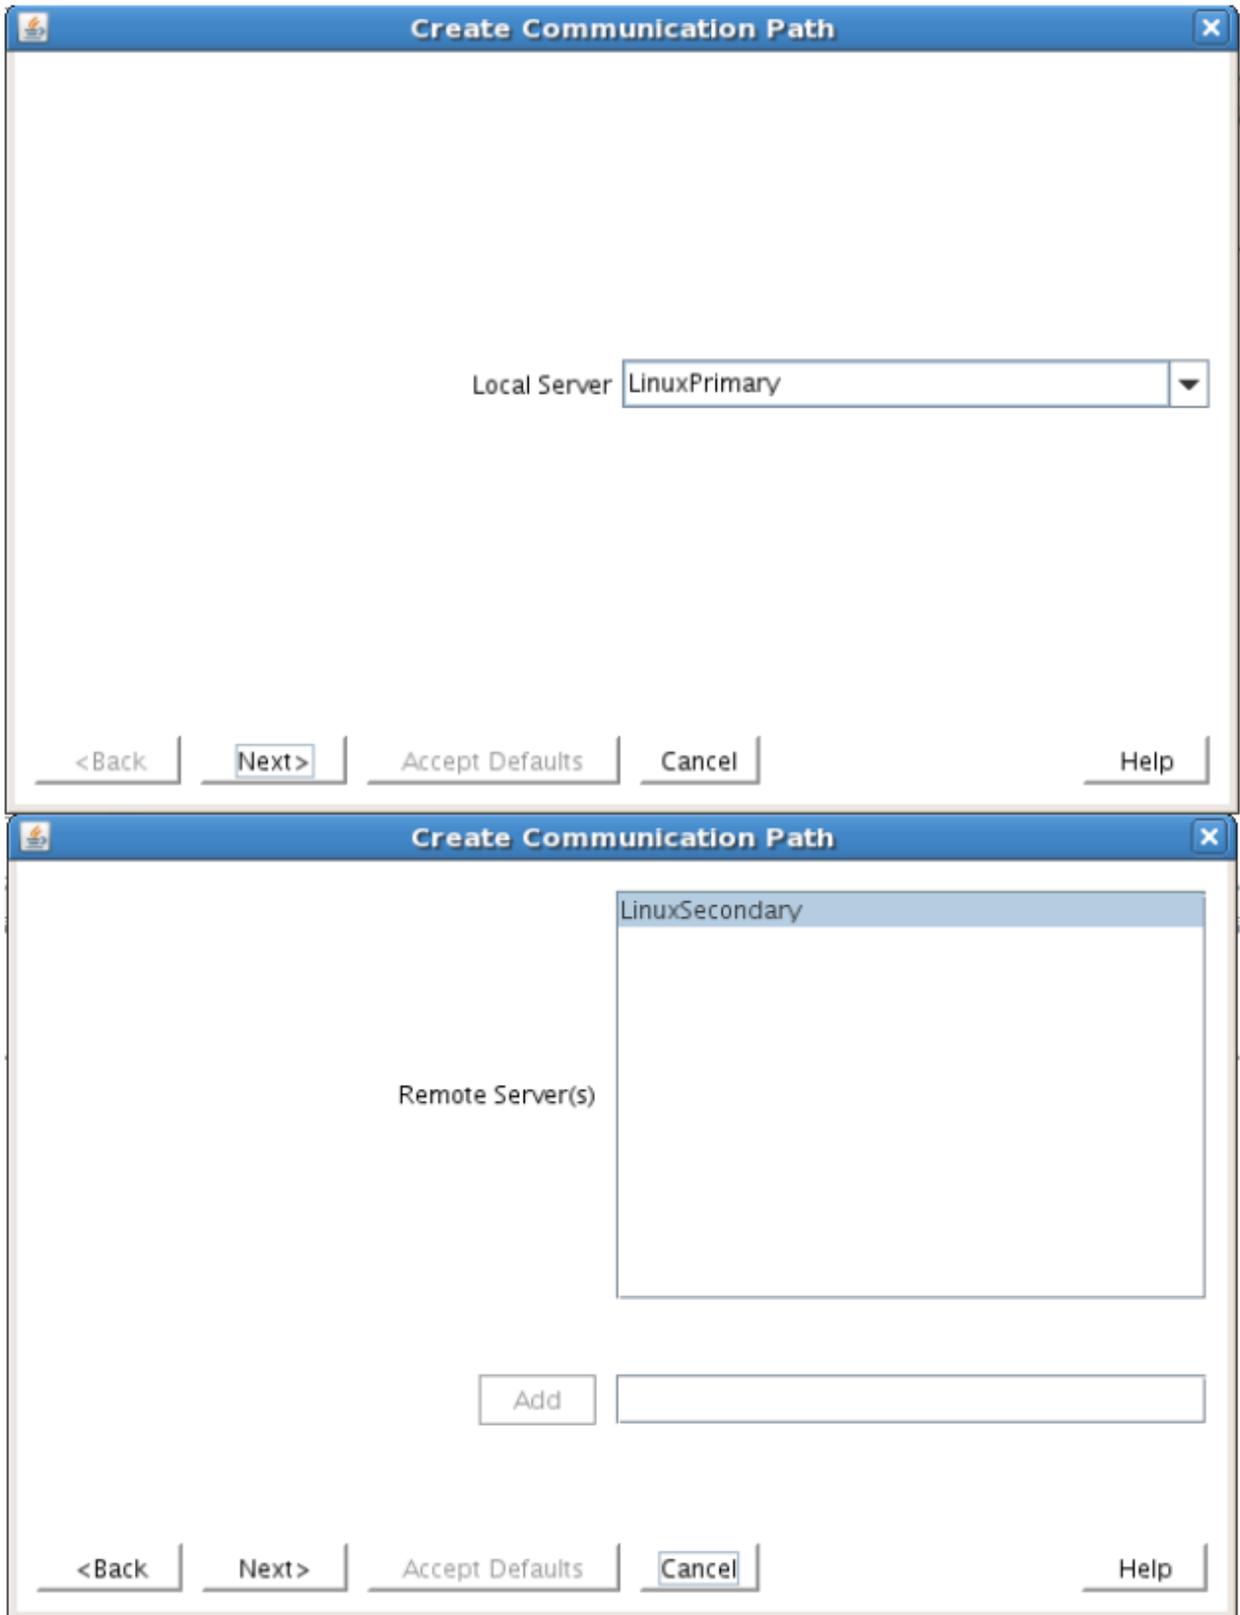
4. Within the LifeKeeper GUI, from the File menu, select Connect. Enter the name of your Secondary server, login and password when the Cluster Connect window displays.



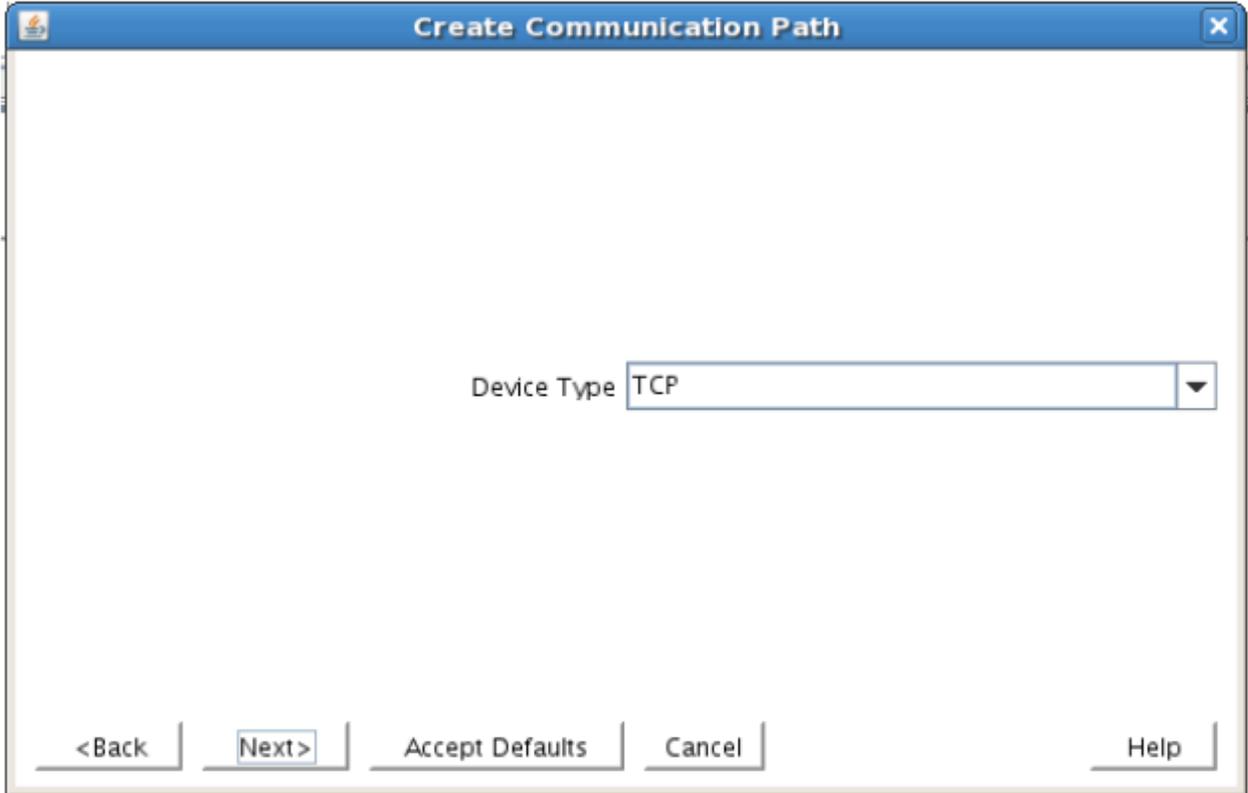
5. Within the LifeKeeper GUI, click the Create Comm Path button on the toolbar. You can also right click one of the servers and click Create Comm Path from the pop-up menu as well.



6. Select your Local and Remote Server(s) from the list box. If a server does not appear in the list box, you may enter it by typing its name and clicking the Add Server button. When using the Add Server procedure, you must make sure that the computer names for both network interfaces on the servers respond correctly when you ping them (from all of the partner server(s)) using the **ping -a IP ADDRESS** syntax. If they do not, this must be corrected prior to continuing. Click Next.

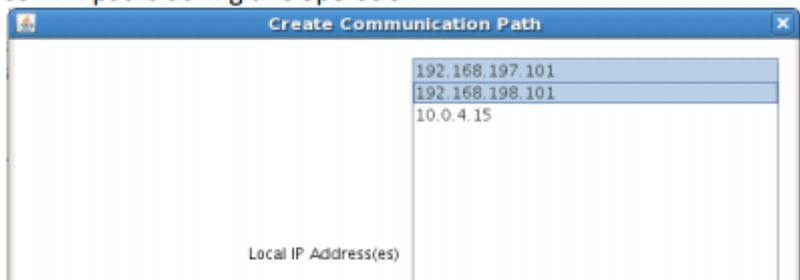


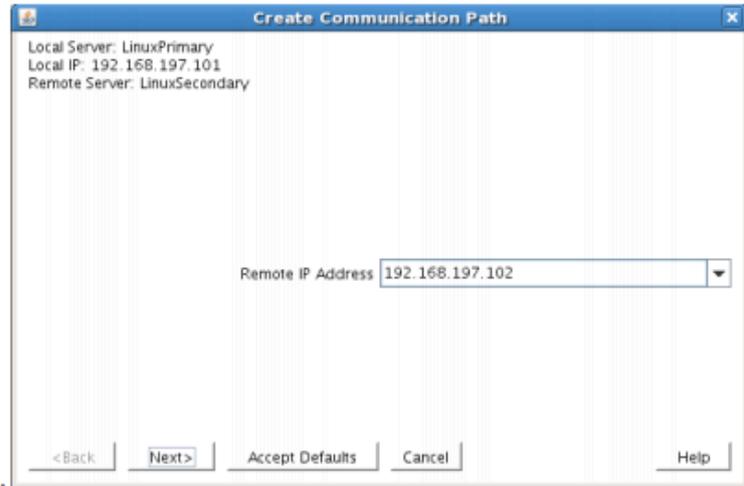
7. Select TCP for Device Type and Click Next.



- 8. Provide all the required information and click Next for the following series of dialog boxes. For each field in the dialog box you can click Help for further information or refer to the table below for an explanation or recommendation

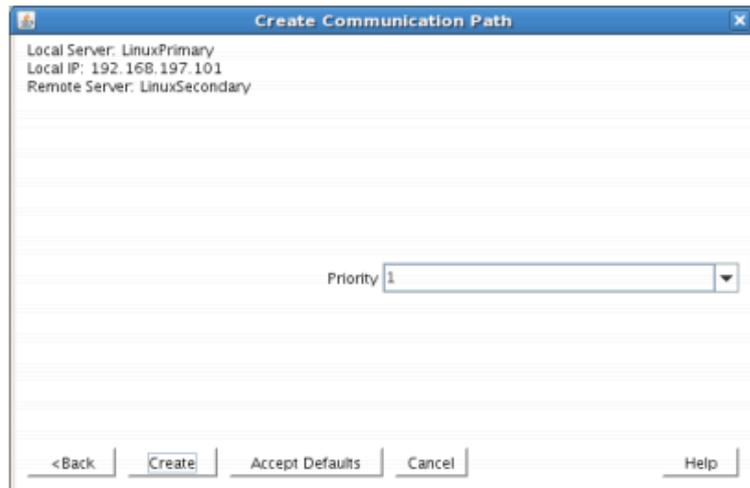
Field	Tips
<b>For TCP/IP Comm Path...</b>	
Local IP Address	Choose the IP address to be used by the local server for this comm path. Select both interfaces so that the wizard creates multiple comm. paths during this operation
Remote IP Address	Choose the IP address to be used by the remote server for this comm path



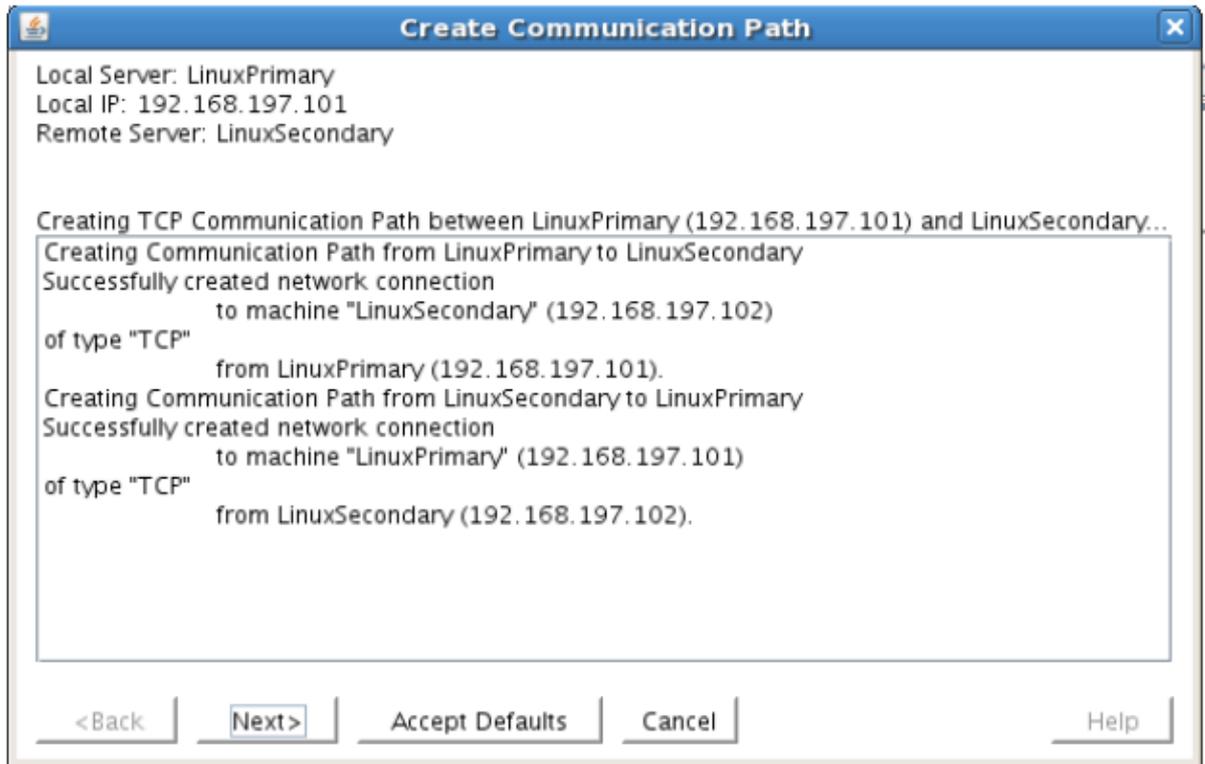


Enter the priority for the comm path on the local server. The priority will be used to determine the order that the comm paths between two servers will be used. Priority 1 is the highest; 99 is the lowest.

Priority



9. After entering data in all the required fields, select Create. A message will display indicating the network communication path is successfully created. Click Next.



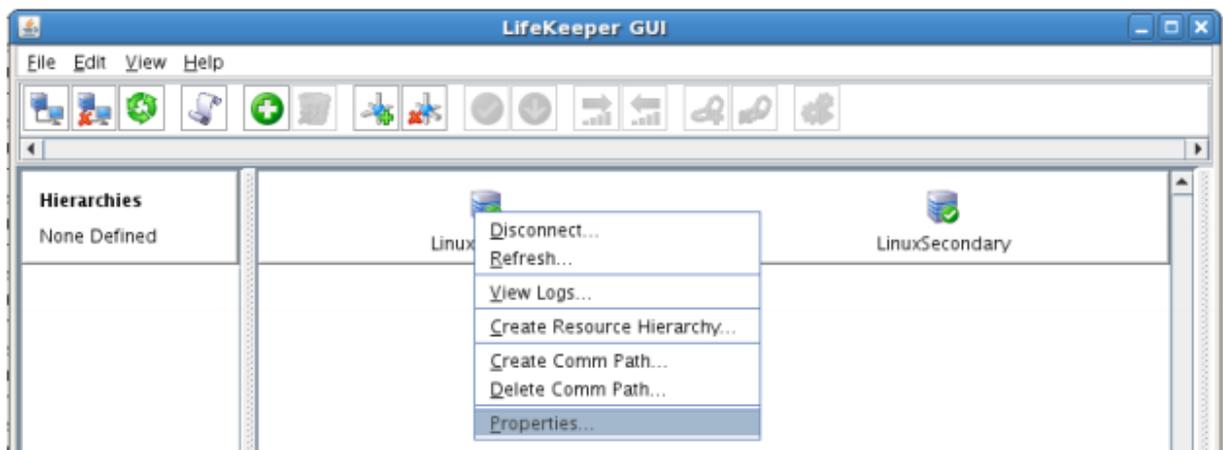
If you selected multiple Local IP Addresses or multiple Remote Servers and the Device Type was set to TCP, then the procedure will return you to the setup wizard the next Comm Path.

- 10. Select Done in the last dialog box.

Repeat this process until you have defined all the communication paths you plan to use. SIOS strongly recommends that you define at least two communication paths for redundancy.

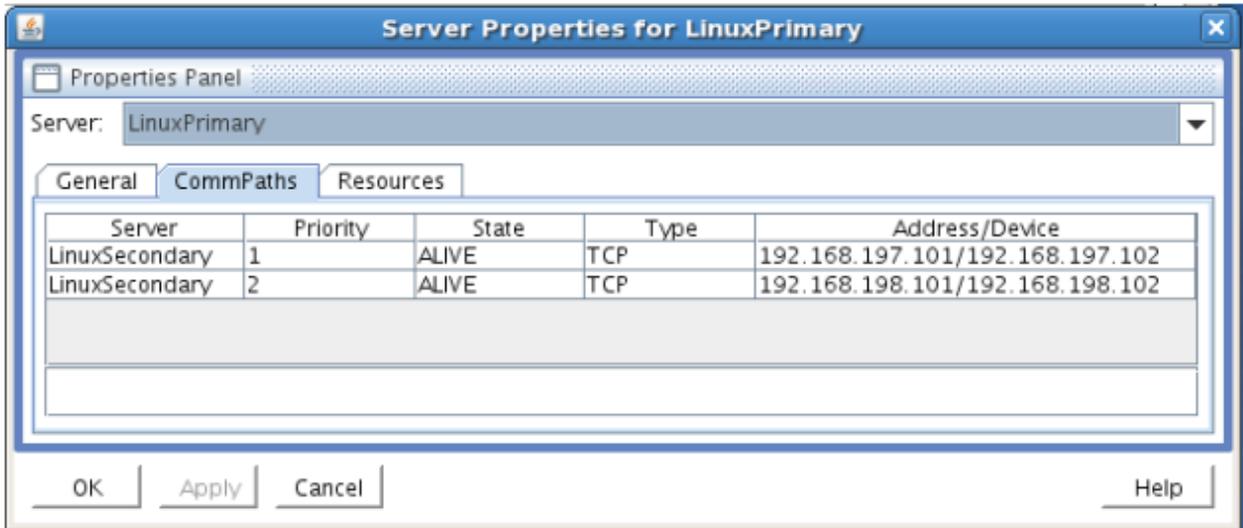
### Verify the Communications Paths

- 1. Verify that the communications paths are configured properly by viewing the Server Properties dialog box. From the LK GUI, select Edit, Server, Properties and then the Comm Paths tab.



- 2. Note the State displayed is ALIVE. You can also check the server icon in the right, main pane of

the GUI. If only one comm path has been created, the server icon shows a yellow warning icon on the server icon, indicating that one comm. path is ALIVE, but there is no redundant comm path. The server icon will display a green heartbeat checkmark when there are at least two comm paths configured and ALIVE.



## Create the LifeKeeper Hierarchy

### Create and Extend an IP Resource

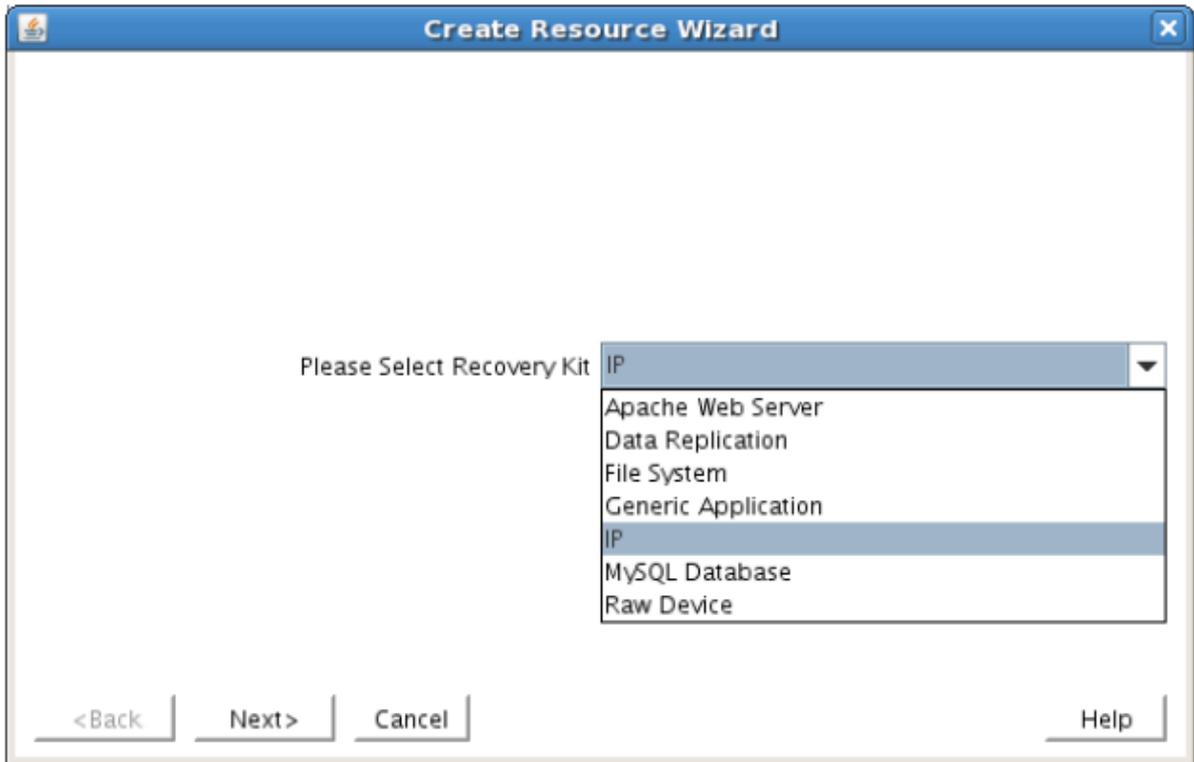
In LifeKeeper, create an IP resource and extend it to the secondary server by completing the following steps. This Virtual IP will have the ability to move between cluster nodes along the application that depends on it.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

2. Select IP Address and click Next.



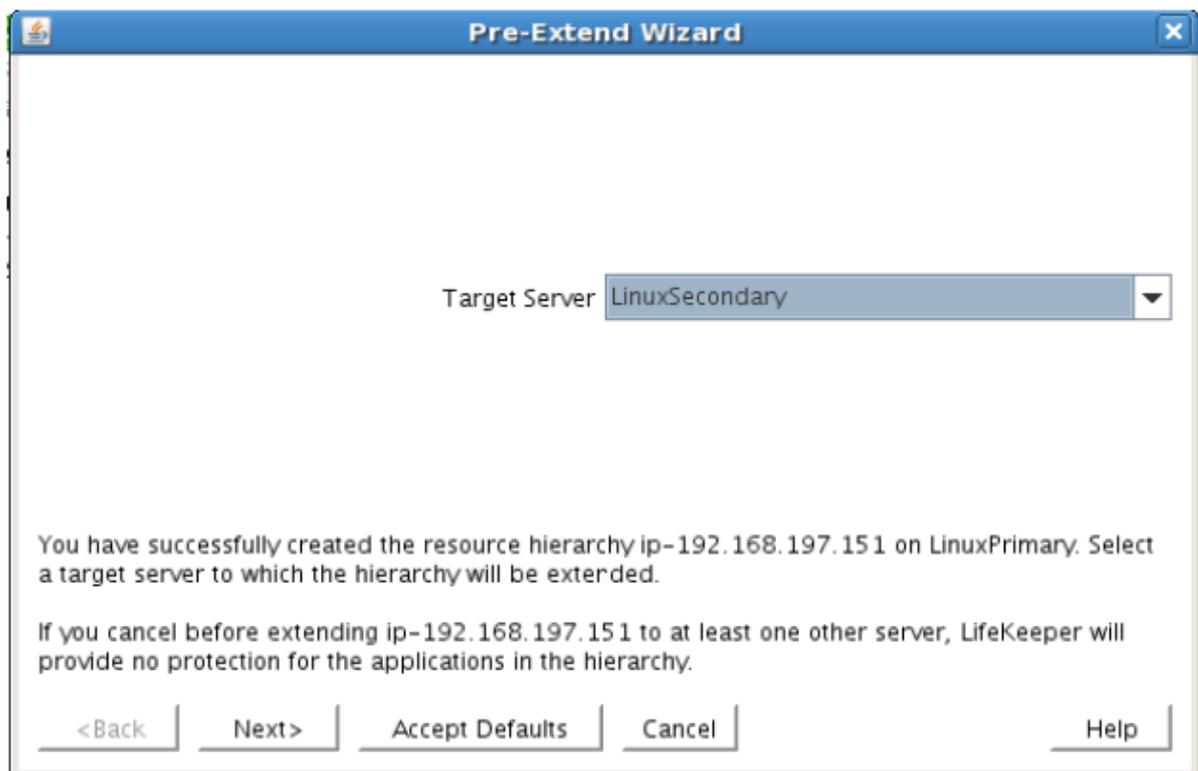
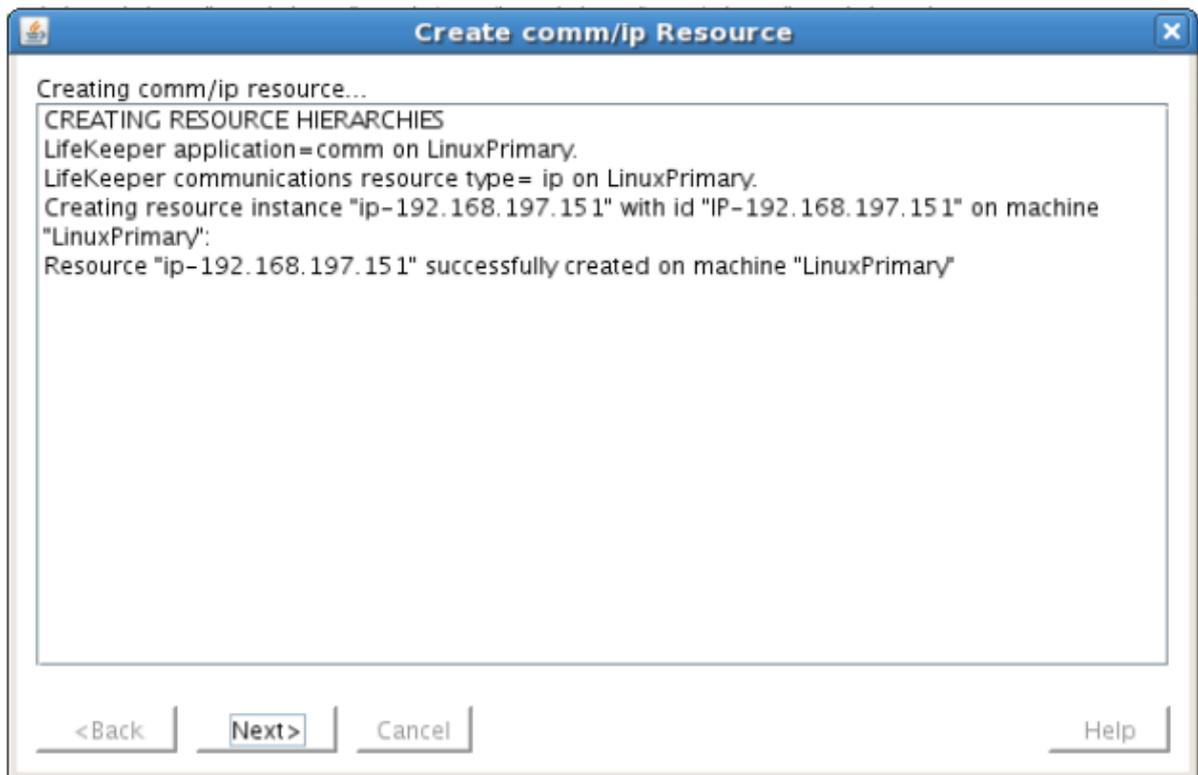
3. Enter the appropriate information for your configuration. The table below contains a list of the fields that display and additional information to assist you as you complete this procedure. Recommended values are also show below. You can also click the Help button for further information. Press Next to continue after entering the required information.

## IP Creation Field Definitions

Field	Tips
Resource Type	Select IP Address as the resource type and click Next.
Switchback Type	Select Intelligent and click Next.
Server	Select the Server where the IP resource will be created. Select your Primary server and click Next.
IP Resource	<p>Enter the virtual IP information and click Next</p> <p>Example <b>192.168.167.151</b></p> <p><b>Note</b> This is an IP address that is not currently in use anywhere on your network. This is the address that all clients will use to connect to the protected resources.</p>

Netmask	<p>The IP subnet mask that your TCP/IP resource will use on the target server. Any standard netmask for the class of the specific TCP/IP resource address is valid.</p> <p>In our sample configuration 255.255.255.0 is used for a subnet mask on both networks.</p> <p><b>Note:</b> The subnet mask you choose, combined with the IP address, determines the subnet that will be used by the TCP/IP resource and should be consistent with the network configuration.</p>
Network Connection	<p>This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. Select the correct NIC and click Next.</p>
IP Resource Tag	<p>Accept the default value and click Next. This value only affects how the IP is displayed in the GUI. The IP resource will be created on our Primary server.</p>

- LifeKeeper will create and validate your resource. After receiving the message that the resource has been created successfully, click Next when the following dialog box appears so that you can complete the process of Extending the IP Resource to our Secondary server, below.



Extending the IP resource will start automatically after you have finished creating an IP address resource if you clicked Next in the dialog box displayed above. You can also start this from an existing IP address resource by right clicking on the active resource and selecting Extend Resource Hierarchy.

Refer to the table below to complete the Extend IP Resource procedure.

Field	Recommended Entries or Notes
Switchback Type	Leave as "intelligent" and click Next
Template Priority	Leave as default (1)
Target Priority	Leave as default (10)
Network Interface	This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. The correct physical NIC should be selected by default. Please verify and then click Next
IP Resource Tag	Leave as default.
Target Restore Mode	Select Enable and click Next.
Target Local Recovery	Select Yes to enable Local Recovery for the SQL resource on the Target server.
Backup Priority	Accept the default value.

- After receiving the message Hierarchy extend operations completed, click Finish and then click Done
- Your IP resource (192.168.197.151) is now fully protected and has the ability to "float" between cluster nodes as needed. Looking at the LifeKeeper GUI you will notice that the IP resource is Active on the Primary cluster node and Standby on the Secondary cluster node



## Create a Mirror and Begin Data Replication

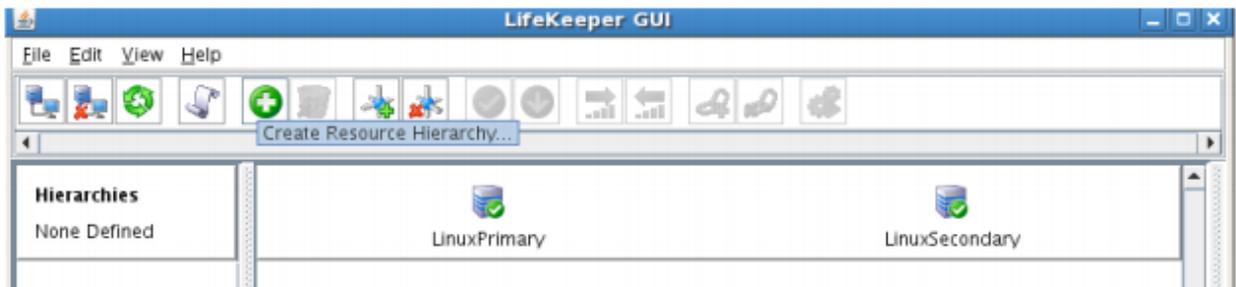
In this section we will setup and configure the Data Replication resource, which be used to synchronize

our MySQL's data between cluster nodes. The data we will replicate resides in the /var/lib/mysql partition on our Primary cluster node

Please note:

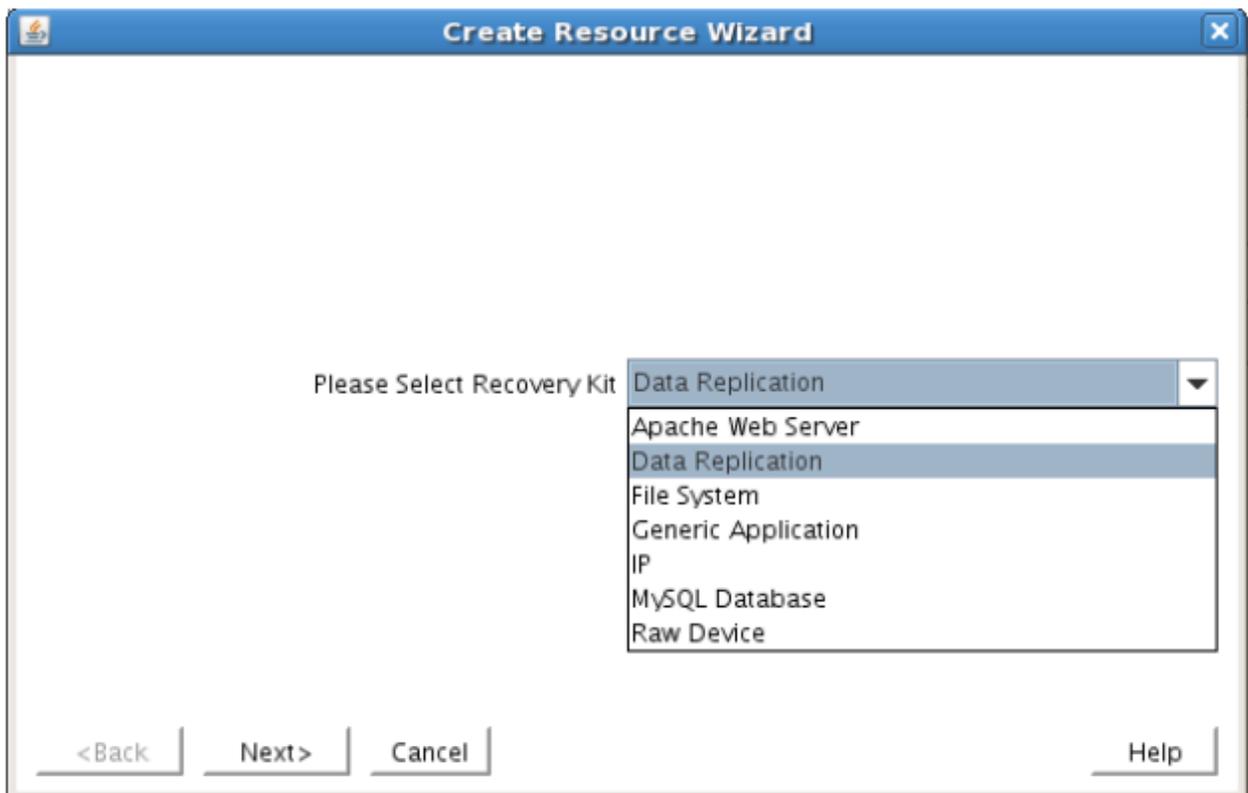
- The source volume to be replicated must be mounted on the Primary server
- The target volume, which will received replicated data, must **NOT** be mounted on the Secondary server.
- The target volume's size must equal to or larger than the size of its source volume.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

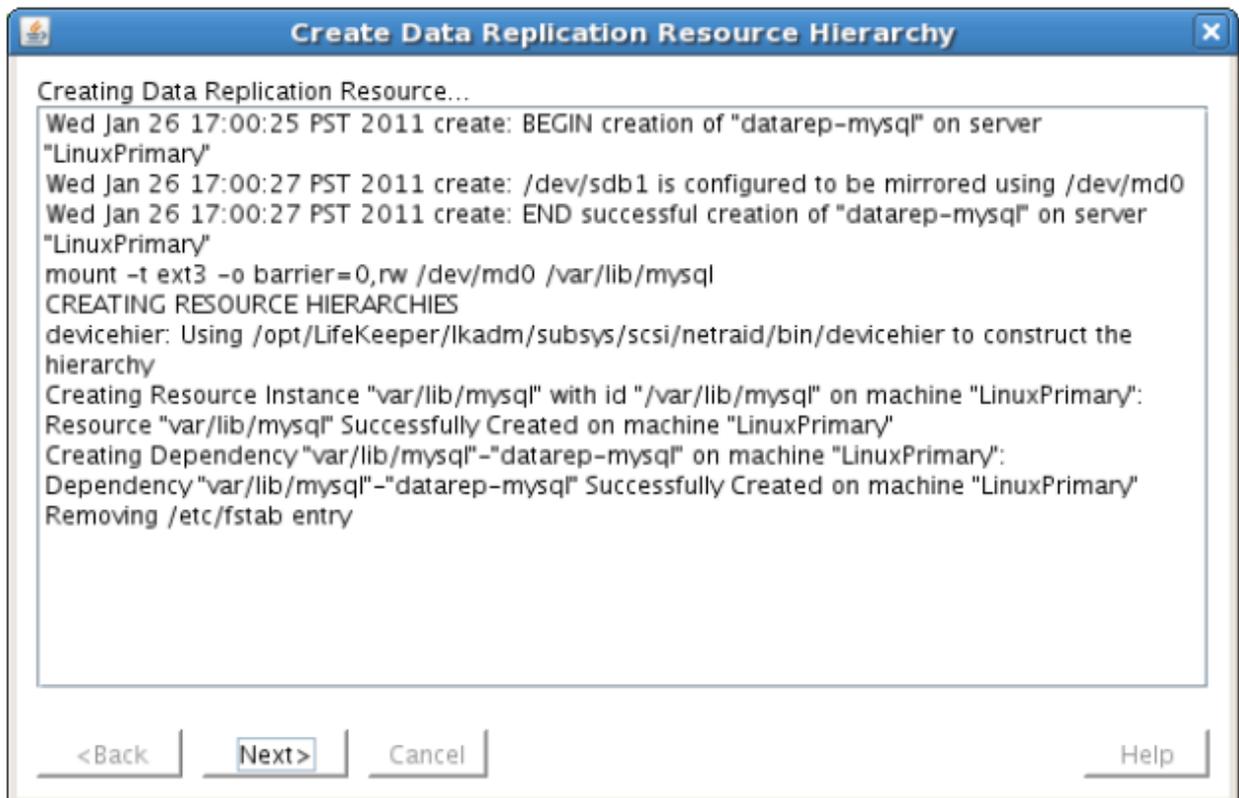
2. Select Data Replication and click Next.



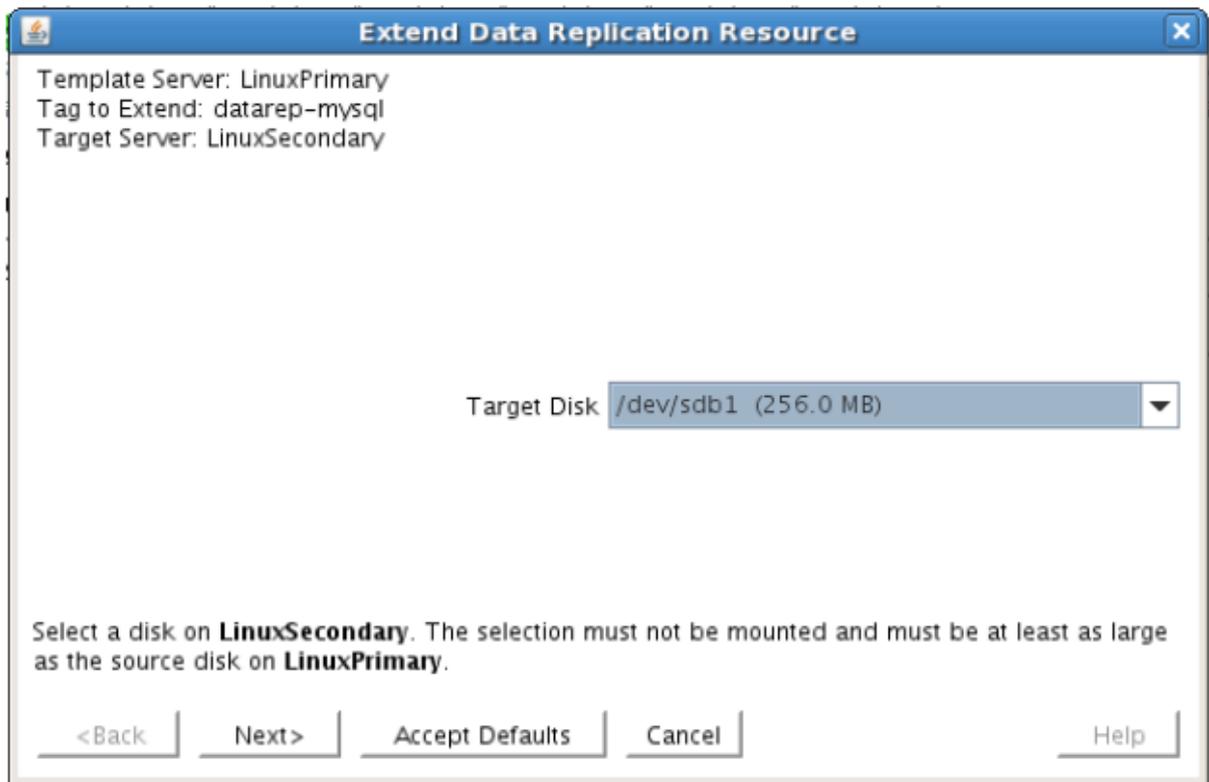
3. Follow the Data Replication wizard, and enter the following values:

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node, i.e. Mirror Source)
Hierarchy Type	Select: <b>“Replicate Existing Filesystem”</b>
Existing Mount Point	At this step you will select the mounted partition to replicate. In our example, select <code>“/var/lib/mysql”</code>
Data Replication Resource Tag	Leave as default
File System Resource Tag	Leave as default
Bitmap File	Leave as default (Note: if using high speed SSD storage you will want to create a small partition and use it for bitmap placement, i.e. <code>/bitmaps</code> )
Enable Asynchronous Replication	Leave as default (Yes)

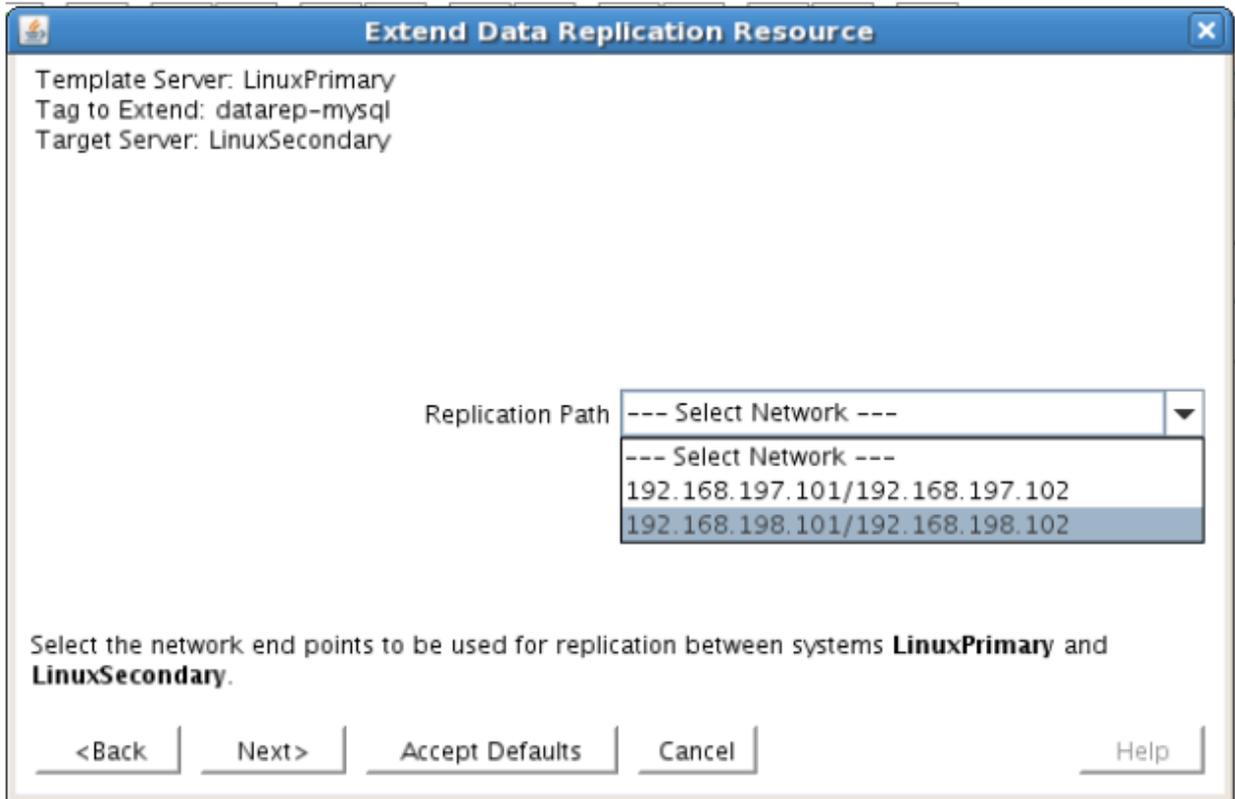
4. Click Next to begin creation of the Data Replication resource hierarchy. Status will be displayed in the GUI as follows:



5. Click Next to begin the process to Extend the Data Replication Resource. Select all default settings. When it asks for the target disk, select a free partition on your Target server which is the same size (or greater) than the Source Volume we are replicating. This partition should NOT be mounted on the Target system.



6. Continue through the wizard, and you will be prompted to select the network you would like replication to take place over. In general, it's a best practice to separate your user/application and your replication traffic. In our example setup we will replicate over our backend network, 192.168.198.X



- 7. Click Next and continue through the wizard. Once completed, your resource hierarchy will look as follows



### Create the MySQL Resource Hierarchy

Create a MySQL resource to protect the MySQL database and make it high available between cluster nodes.

**✳ Important** At this point, MySQL needs to be already running on the Primary Server. It should NOT be running on the Secondary Server. See section titled “Install, Configure, and Start MySQL” above to review the process to configure and start MySQL as needed.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.
2. Select **MySQL Database** and click **Next**.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node)
Location of my.cnf	Enter "/var/lib/mysql". Note that earlier in the MySQL
executables	configuration process we created a my.cnf file in this directory
Location of MySQL	Leave as default (/usr/bin) since we are using a standard MySQL install/configuration in this example
Database tag	Leave as default

4. Select Create to define the MySQL resource hierarchy on the Primary Server
5. Click Next to Extend the File System Resource to the Secondary Server
6. In the Extend Wizard, select "Accept Defaults"
7. As a result the MySQL resource is now protected on both cluster nodes. Click Finish to exit the Extend wizard.
8. Note: LifeKeeper will automatically identify that the MySQL resource has a dependency on the FileSystem (Data Replication) resource (/var/lib/mysql). The Filesystem Resource will appear underneath the MySQL resource in the GUI
9. Your resource hierarchy should look as follows:



### Create the MySQL IP Address Dependency

In this step will define an additional dependency: that MySQL depends on a Virtual IP (192.168.197.151) so that the IP address follows the MySQL database as it moves.

1. From the LifeKeeper GUI toolbar, right-click on the “mysql” resource
2. Select “Create Dependency” from the right-click context menu
3. In the Child Resource Tag dropdown menu, select “ip-192.168.197.151”
4. Click Next
5. Click Create Dependency
6. Click Done
7. The Virtual IP address resource (192.168.197.151) will now appear underneath the MySQL resource in the LifeKeeper user interface. This ensures that resources move together, and are started/stopped in the proper order.
8. Your resource hierarchy should look as follows



At this point in the Evaluation, we have fully protected MySQL, and its dependent resources: IP addresses, and replicated storage.

## 12.3.8. Test Your Environment

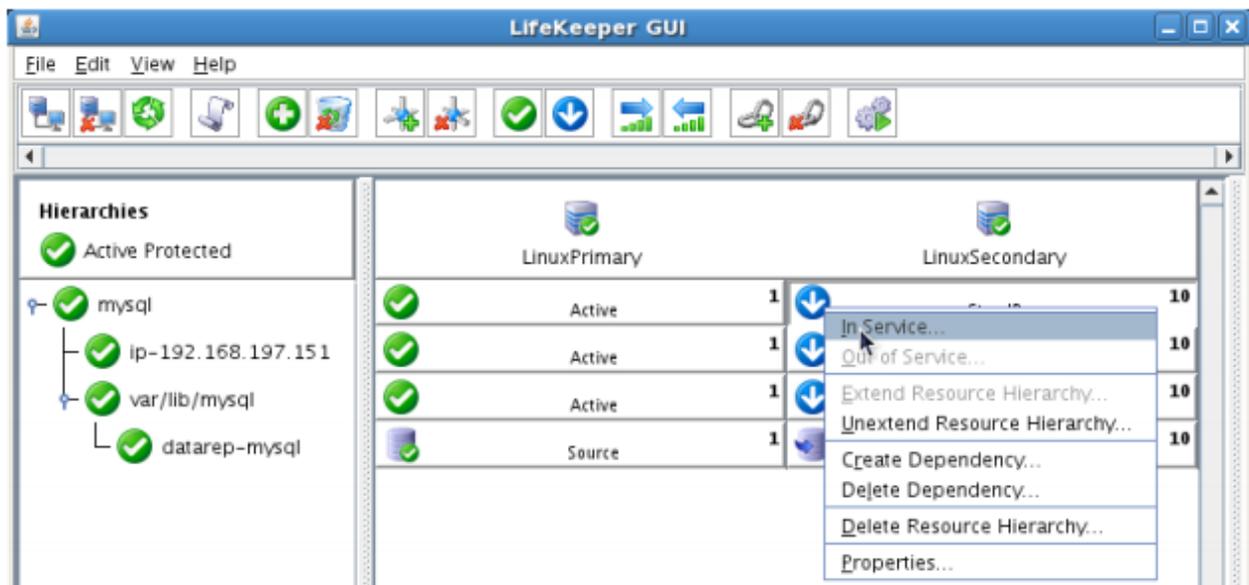
The following test scenarios have been included to guide you as you get started evaluating LifeKeeper for Linux. **Before beginning these tests, make sure the data replication resources are in the mirroring state.**

- \* **Note:** For these test examples, the Primary Server is referred to as LINUXPRIMARY. The Backup or Secondary Server is referred to as LINUXSECONDARY.

### Manual Switchover of the MySQL Hierarchy to Secondary Server

#### Procedure:

- From the LifeKeeper GUI, right click on the MySQL resource on the Secondary Server (LINUXSECONDARY) and choose IN SERVICE.
- Click "In Service" in the window that pops up



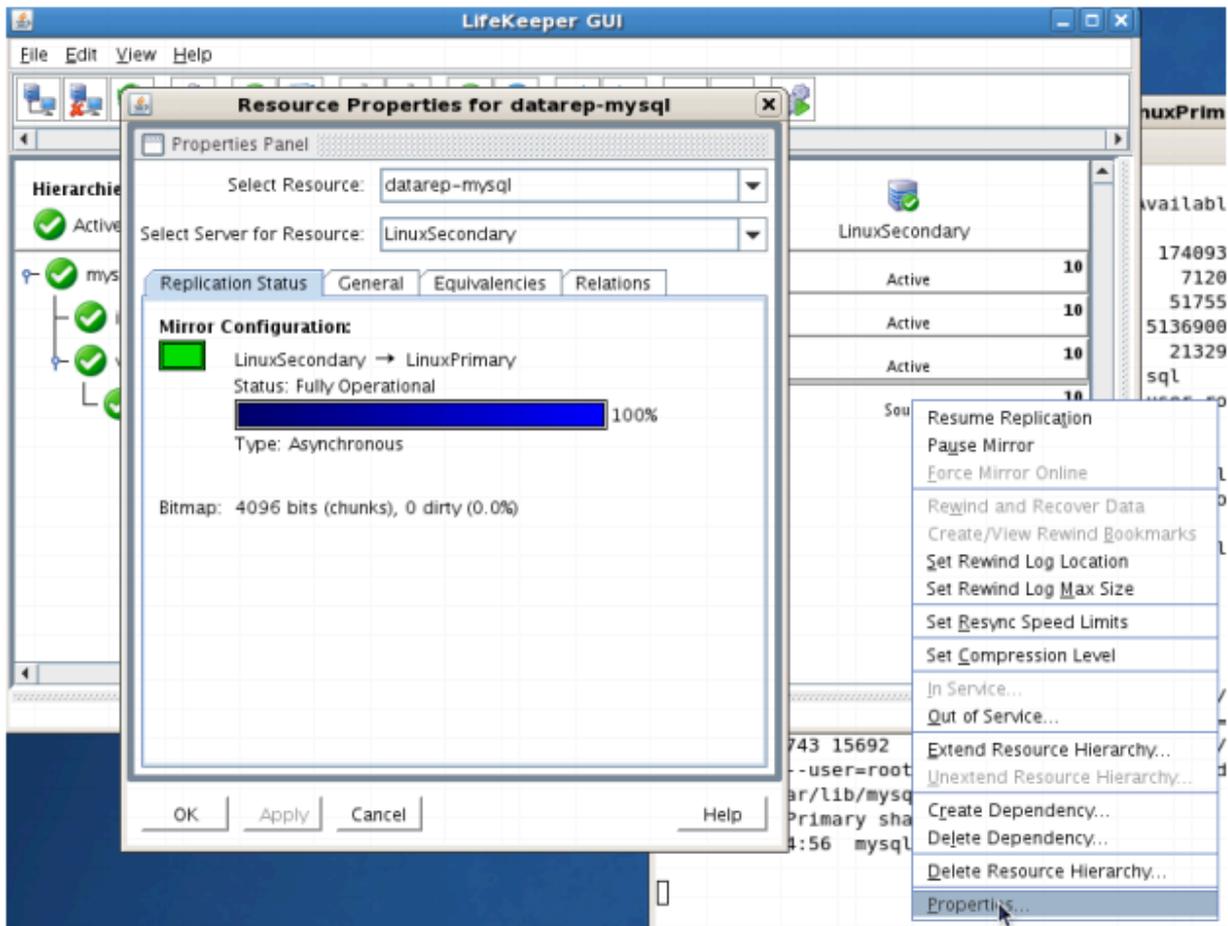
#### Expected Result:

- Beginning with the MySQL resource, all resources will be removed from service on the Active Server (LINUXPRIMARY).
- Beginning with the dependent resources (IP and Replicated Volume), all resources will be brought in service on LINUXSECONDARY.
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXSECONDARY -> LINUXPRIMARY
- At this point, all resources are now active on LINUXSECONDARY.



**Tests/Verification:**

- Using the LifeKeeper GUI, verify that the MySQL and dependent resources are active on LINUXSECONDARY.
- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the “datarep-mysql” resource and select Properties



- Run “ifconfig -a” on LINUXSECONDARY to validate that the IP Address 192.168.197.151 is active on LINUXSECONDARY
- Run “df -h” to verify that the /var/lib/mysql replicated filesystem is mounted as an “md” device

(example: /dev/md0”) on LINUXSECONDARY

- Verify the MySQL services are running on LINUXSECONDARY by running “ps -ef | grep -i mysql”
- On LINUXSECONDARY run the following command to verify client connectivity to the MySQL database:
  - # mysql -S /var/lib/mysql/mysql.sock -u root -p
  - (enter password “SteelEye”)

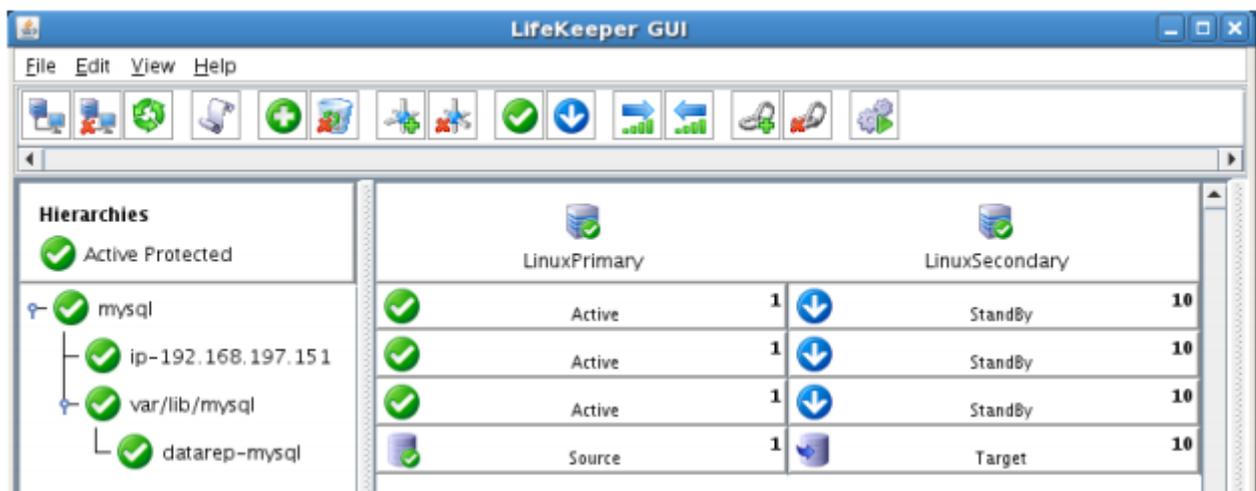
## Manual Switchover of the MySQL Hierarchy back to Primary Server

### Procedure:

- From the LifeKeeper GUI, right click on the MySQL resource on the Primary Server (LINUXPRIMARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up

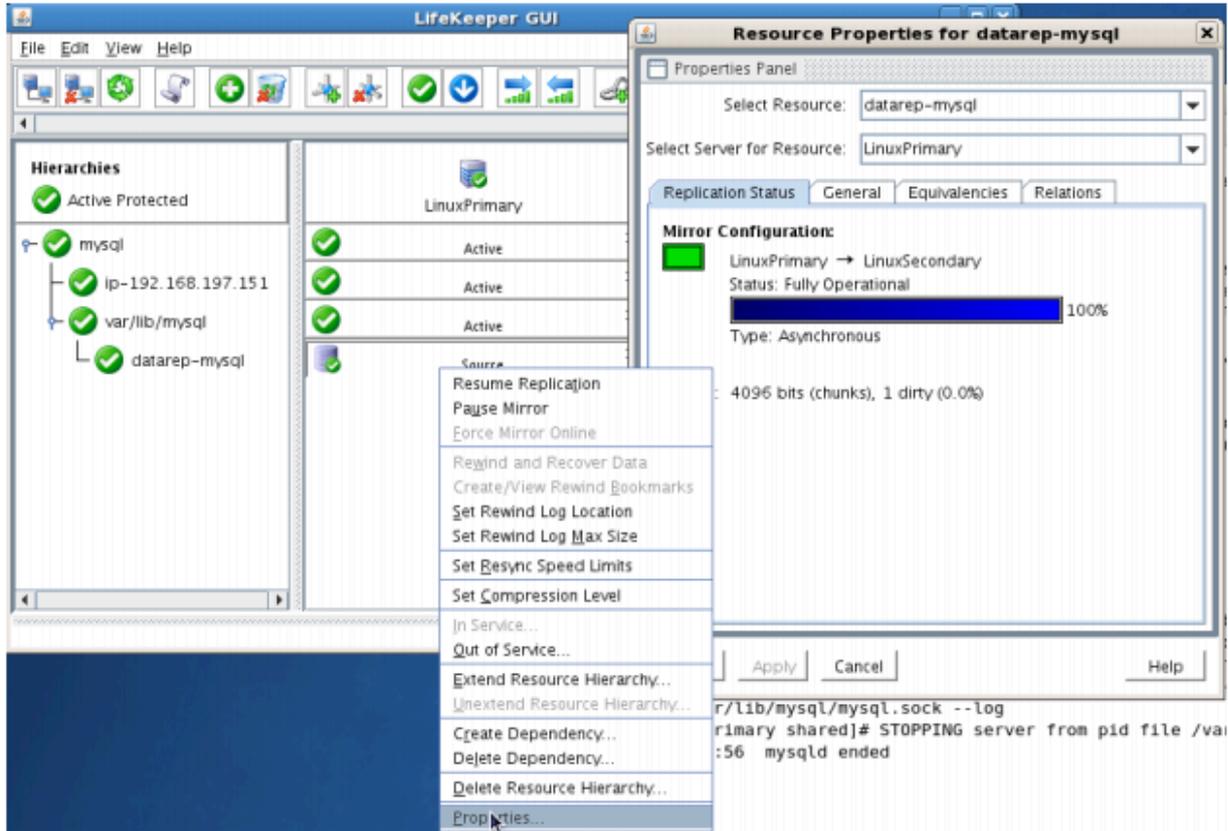
### Expected Result:

- Beginning with the MySQL resource, all resources will be removed from service on the Active Server (LINUXSECONDARY).
- Beginning with the dependent resources (IP and Replicated Volume), all resources will be brought in service on LINUXPRIMARY
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXPRIMARY -> LINUXSECONDARY



### Tests/Verification:

- Using the LifeKeeper GUI, verify that the MySQL and dependent resources are active on LINUXPRIMARY.
- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the “datarep-mysql” resource and select Properties



- Run “ifconfig –a” on LINUXPRIMARY to validate that the IP Address 192.168.197.151 is active on LINUXPRIMARY
- Run “df –h” to verify that the /var/lib/mysql replicated filesystem is mounted as an “md” device (example: /dev/md0”) on LINUXPRIMARY
- Verify the MySQL services are running on LINUXPRIMARY by running “ps –ef | grep –i mysql”
- On LINUXPRIMARY run the following command to verify client connectivity to the MySQL database:
  - # mysql –S /var/lib/mysql/mysql.sock –u root –p
  - (enter password “SteelEye”)

## Simulate a network failure on the Primary Server by failing the IP resource

**!** **IMPORTANT:** Only perform this test if you have more than one communications path configured.

If you perform this test and have only one communications path configured, your system will enter a split-brain scenario as described in the LifeKeeper Administration Guide found [here](#). Refer to this document for more information or contact SIOS presales technical support for assistance in resolving this condition.

**Procedure:**

- On LINUXPRIMARY, pull the network cable attached to the NIC that the virtual IP address is configured on

**Expected Result:**

- The IP Resource should fail first.
- The entire hierarchy should failover to LINUXSECONDARY

**Tests/Verification:**

- Check the LifeKeeper Log to verify the IP resource failed – “/opt/LifeKeeper/bin/lk\_log log”
- Using the LifeKeeper GUI, verify the MySQL and Apache resource hierarchies fail over successfully to LINUXSECONDARY



## Hard failover of the resource from the Secondary Server back to the Primary Server

**Procedure:**

- Pull the power cord on LINUXSECONDARY, as this is the server with all resources currently In Service.

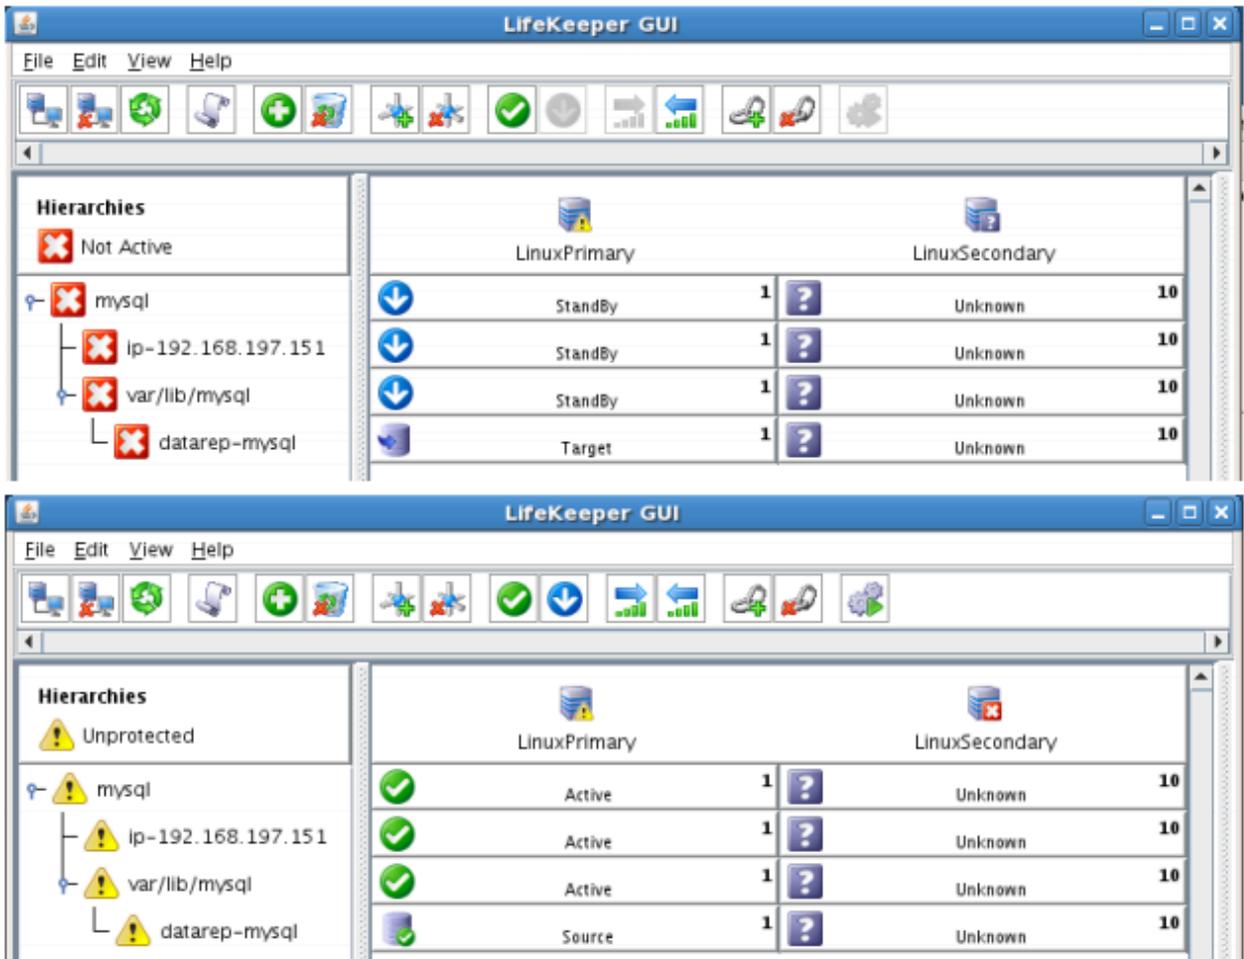
**Expected Result:**

- After failure has been detected, beginning with the dependent resources (IP and Volume), all resources will be brought in service on LINUXPRIMARY.

**Tests/Verification:**

- Using the LifeKeeper GUI, verify the mirror has reversed and is in a Resync Pending state waiting for LINUXSECONDARY to come back on line.

- Verify the Apache and MySQL Server services are running on LINUXPRIMARY.
- Verify that the client can still connect to the Webserver and database running on LINUXPRIMARY.
- Verify you can write data to the replicated volume, /var/lib/mysql on LINUXPRIMARY.



## Bring Failed Server back on line

### Procedure:

- Plug the power cord back into LINUXSECONDARY and boot it up.

### Expected Result:

- Using the LifeKeeper GUI, verify that LINUXSECONDARY is coming back up and has become the Standby Server.

### Tests/Verification:

- Verify the mirror performs a quick partial resync and moves to the Mirroring state
- Verify the Apache and MySQL Hierarchy are in service on LINUXPRIMARY and standby on LINUXSECONDARY.



## Verify Local Recovery of MySQL Server

### Procedure:

- Kill the MySQL processes via the command line:
  - # ps -ef | grep sql
  - # killall mysqld mysqld\_safe
  - run “ps -ef | grep sql once again to verify that the processes no longer exist

### Expected Result: (Assumes Local Recovery for MySQL resource is set to YES)

- The MySQL Server service should stop.
- The MySQL quickcheck process will automatically restart the MySQL Server Service when it runs periodically.
- No failure of MySQL should occur.

### Tests/Verification:

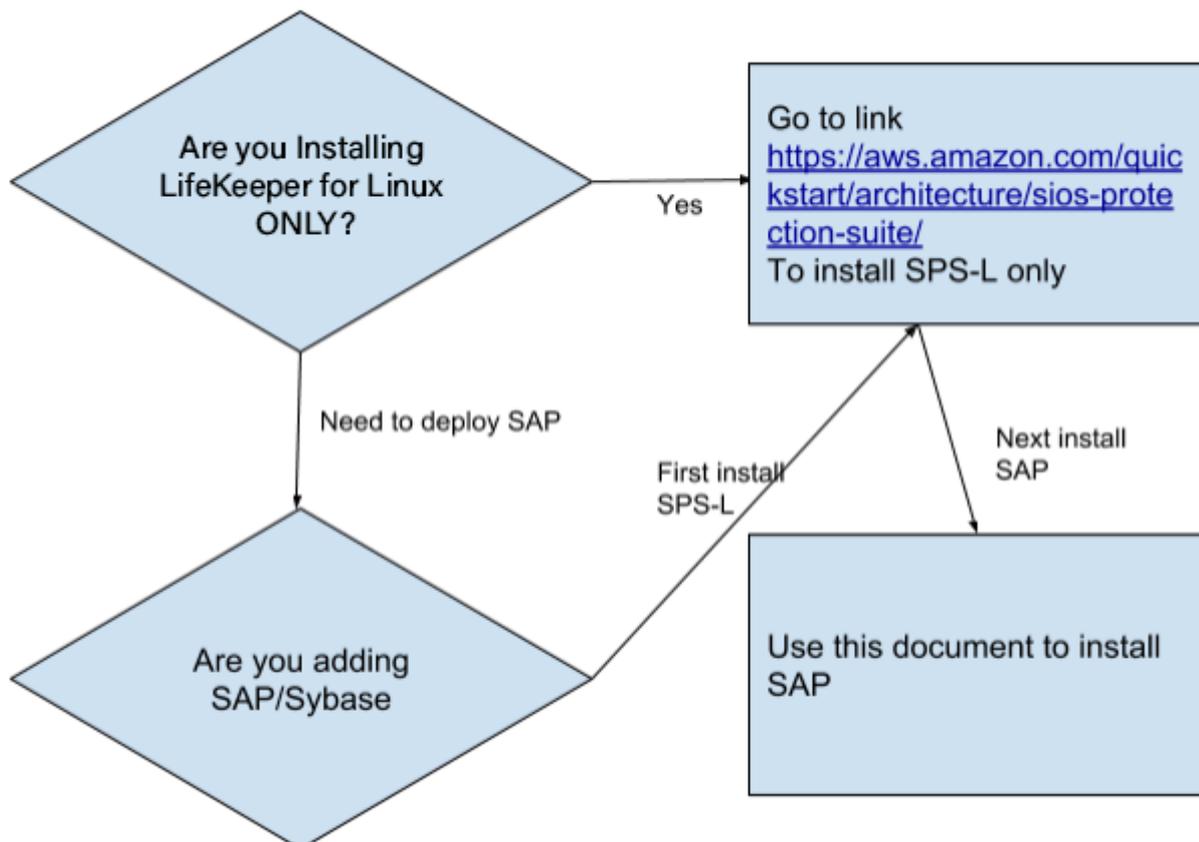
- Execute “ps -ef | grep sql” once again to verify that the mysql processes have been restored locally on LINUXPRIMARY.
- Verify connectivity to the MySQL database by running:
  - ↑
  - mysql -S /var/lib/mysql/mysql.sock -u root -p
  - (Enter password “SteelEye”)
- If you inspect the LifeKeeper logs, you will see information indicating that LifeKeeper detected the failure of the MySQL service and recovered it locally. Run /opt/LifeKeeper/bin/lk\_log log for more information.

# 12.4. LifeKeeper for Linux in the AWS Cloud (SAP)

## Overview

This document will guide the user during LifeKeeper for Linux installation.

Follow the quick decision matrix to understand how to install LifeKeeper for Linux for SAP environment.



**Note:** The link above is provided [here](#) so it can be copied.

## 12.4.1. Additional Steps to Configure SAP on LifeKeeper

---

Follow the steps below to configure SAP on LifeKeeper for Linux.

### Step 1. Test the Deployment

To connect to the LifeKeeper for Linux nodes, you need to connect to Windows jumpbox. To connect to a Windows machine, you need to connect to remote desktop terminal.

In AWS console, select the windows jumpbox node that was created, click on **Actions** and click on **Connect**. You can now download remote desktop program to connect. You will also need to decrypt the password that will need to be used to login to the machine.

Once you are connected to Windows machine, we suggest you download Putty and VNC Viewer. Download them from these sites.

- **Putty** – [www.putty.org](http://www.putty.org)
- **VNC Viewer** – <https://www.realvnc.com/en/connect/download/viewer/>

You can now use Putty to connect to the private IP address of each node, as well as VNC Viewer to connect to the node using the same private IP address. Note that the nodes are not accessible outside the windows jumpbox, but the nodes should be able to access the internet using the NAT gateway. (**Note:** If there are issues with the NAT gateway, make sure to check the security group rules/main route).

Once you have connected to one of the nodes, you can su to root using the password you created in the template earlier, and run the program vncserver. This will allow you to connect using VNC Viewer to that node in a graphical interface.

*Code snippet for installing VNC Server*

Run as root the command vncserver with the following options:

Enter password and repeat for confirmation  
Set Read-Only password to **No**

```
Optionally edit /root/.vnc/config and add  
securitytypes=none  
vncserver -kill:1
```

Access to VNC is ipv4:5901 where 5901 is the port number specified.

Right click on the desktop and click on **Open Terminal**, and enter the command **/opt/LifeKeeper/bin/lkGUIapp**, that will connect to the LifeKeeper GUI. Login using root and password setup previously. You will see the 2 nodes connected.

Now that you have reached this point, basic LifeKeeper 2 node is setup. Proceed with SAP installation and protection of SAP services using LifeKeeper.

## Step 2. Configure Virtual IP

Now that SAP has been setup on the node, you can continue to setup LifeKeeper protecting SAP services and filesystems.

### *Amazon AWS Elastic Compute Cloud (EC2) setup*

The AWS command line interface (cli) needs to be installed on each node. For details, please refer to [AWS Command Line Interface Installation](#). All the EC2 instances must be able to access Amazon EC2 services endpoints using the protocols HTTP and HTTPS. In order to obtain metadata of Amazon EC2 instance, it is necessary to have an access to IP address 169.254.169.254 using the HTTP protocol.

In order for LifeKeeper to operate AWS, an IAM user or IAM role with the following access privilege is required. Please configure an [EC2 IAM role or configure AWS CLI](#) appropriately so that it can be accessed from root user of the EC2 instance.

- ec2:DisassociateAddress
- ec2:DescribeAddresses
- ec2:AssociateAddress
- ec2:DescribeRouteTables
- ec2:ReplaceRoute

### *Create the virtual IP resource*

Determine the IP address. The IP address should be an IP address outside the CIDR block range of the current IP of the nodes. The IP address should be placed in the VPC route table for the node.

Note in the following diagram we placed the ip address of 10.1.0.10/32 and associated it to one of the nodes, using the eni- network adapter.

Buttons: **Create Route Table**, Delete Route Table, Set As Main Table

Search: Search Route Tables and their X

<input type="checkbox"/>	Name	Route Table ID	Explicitly Associat	Main	VPC
<input type="checkbox"/>	Private subnet 1A	rtb-0f25600d73f059...	0 Subnets	No	vpc-003113c78
<input checked="" type="checkbox"/>		rtb-0259a7dda97e3...	2 Subnets	Yes	vpc-003113c78
<input type="checkbox"/>	Private subnet 2A	rtb-0d48d96bfa35f3...	0 Subnets	No	vpc-003113c78
<input type="checkbox"/>	Public Subnets	rtb-09d48b85b3b18...	2 Subnets	No	vpc-003113c78

rtb-0259a7dda97e3fbd8

Buttons: Summary, **Routes**, Subnet Associations, Route Propagation, Tags

Buttons: Cancel, **Save**

View: All rules

Destination	Target	Status	Propagated	Remo
10.0.0.0/16	local	Active	No	
<input type="text" value="0.0.0.0/0"/>	<input type="text" value="nat-082059721765b9ac9"/>	Active	No	
<input type="text" value="10.1.0.10/32"/>	<input type="text" value="eni-012388209b2453af5"/>	Active	No	

Edit /etc/default/LifeKeeper and set NOBCASTPING=1 to disable broadcast ping before continuing.

Click the **green plus** icon to create a new resource:

File Edit View Help

Please Select Recovery Kit

Follow the wizard to create the IP resource with these selections:

**Select Recovery Kit: IP**

**Switchback Type: Intelligent**

**IP Resource: 10.1.0.10**

**Netmask: 255.255.255.0**

**Network Interface: eth0**

**IP Resource Tag: ip-10.1.0.10**

Extend the IP resource with these selections:

**Switchback Type: Intelligent**

**Template Priority: 1**

**Target Priority: 10**

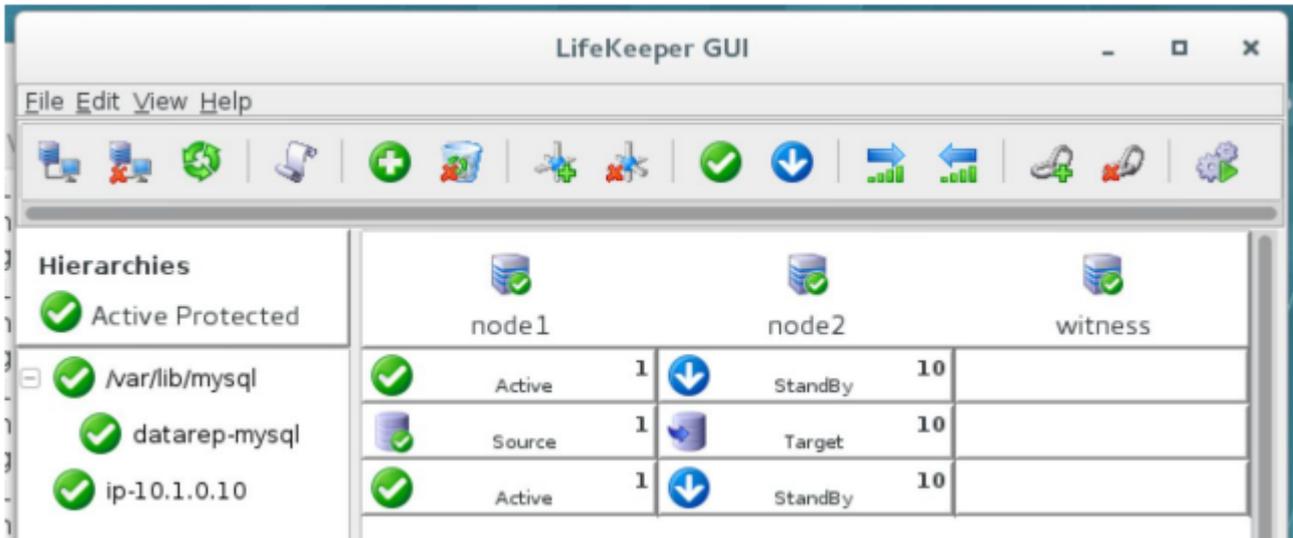
**IP Resource: 10.1.0.10**

**Netmask: 255.255.255.0**

**Network Interface: eth0**

**IP Resource Tag: ip-10.1.0.10**

The cluster will now look like this, with both Mirror and IP resources created:



### Step 3. Setup SAP

Download the SAP software and setup on the node. You can access the SAP marketplace to download SAP software on each node.

There are a number of choices to setup SAP. The decision to implement one would depend on various factors, such as cost, experience and RAS (Reliability, Availability and Serviceability) factors.

- [ASCS without NFS](#)
- [ASCS + ERS with NFS](#)

Each configuration has advantages and disadvantages. We recommend that you work with SAP experts at your site, or you engage with SIOS Professional Services to determine the best fit for your environment.

**Note:** Future documentation will detail installation for HANA, all-in-one, using EFS (AWS Elastic File System) and Cloudwatch. There are also planned automated installation quick start scripts and using SAP Landscape Manager (LaMa) to manage the installation.

## 12.4.2. ASCS without NFS

---

Here are the steps to setup ASCS without NFS:

- [General Setup Steps for ASCS without NFS](#)
- [Installing SAP](#)
- [Creating the SAP Resource Hierarchy](#)

## 12.4.2.1. General Setup Steps for ASCS without NFS

---

1. Create Virtual IP, done in earlier steps
2. Create an EC2 resource and create as dependency for virtual IP, done in earlier steps
3. Install SAP on node 1 on “virtual hostname” based on “virtual IP”
4. Stopsap on node1
5. Use the LifeKeeper GUI to “In-service” the virtual IP to node 2, and Install SAP on node 2 on “virtual hostname” based on “virtual IP”
6. Stopsap on node 2 and modify profile files on both nodes (see below)
7. Use the LifeKeeper GUI to “In-service” the virtual IP back to node 1
8. Create replication resource for the mount points needed for SAP, done in earlier steps, as advised by SAP consultants
9. Startsap on node 1 and ensure SAP is working properly
10. In /etc/default/LifeKeeper add the follow entries to the end on both nodes:

```
SAP_EXPERTMODE=1
```

```
SAP_NFS_CHECK_IGNORE=1
```

```
SAP_DB_CHECK_IGNORE=1
```

11. Re-run the LifeKeeper setup program to add the SAP Recovery Kit

```
./setup -k
```

Select the recovery kit for SAP from the menu of available recovery kits using the arrow keys and pressing the <spacebar> to select, press <enter> to continue and complete the installation.

12. Create SAP resources following the [SAP Recovery Kit guide](#)

## 12.4.2.2. Installing SAP

---

1. ASCS should be installed based on “virtual hostname” based on “virtual IP”, which should have been added to host files during earlier installation steps. Please be sure to do so if they have not already been done prior to installing SAP.
- When installing SAP (specifically ASCS or SCS) you need to specify the `SAPINST_USE_HOSTNAME=vip` option. This is not required for ERS. (**Note:** Document the `SAPINST_USE_HOSTNAME` virtual IP address as it will be used later during creation of the SAP resources in LifeKeeper.)

**Note:** Specify `sapinst SAPINST_USE_HOSTNAME=vip` where **vip** is the virtual IP that will float between the nodes.

Run `./sapinst SAPINST_USE_HOSTNAME= {hostname}`

- In seven phases, the **Core Services** should be created and started. If permission errors occur on `jdbconnect.jar`, go to `/sapmnt/STC/exe/uc/linuxx86_64` and make that directory as well as file `jdbconnect.jar` writable (`chmod 777 ---`).

ASCS profiles should be pointing to local mount point containing `/usr/sap`, `sapmnt` or any other necessary for SAP files in your environment.

The ASCS and ERS instance profiles must be modified in order to prevent the enqueue server and enqueue replicator processes from being automatically restarted by the `sapstart` utility. After updating the instance profiles, SAP Start Service for each of these instances must be restarted in order for the changes to take effect. Follow the steps provided in [Modify ASCS and ERS Instance Profile Settings](#), then return to this page to continue the setup process.

2. Sapstop SAP on node 1
3. Using the LifeKeeper GUI, right click on the IP address resource created for ASCS, select “In-Service” and select node 2 to switch the IP onto node 2.
4. Repeat step 1 to install SAP onto node 2 and ensure that it’s able to run correctly
5. Sapstop SAP on node 2
6. Using the LifeKeeper GUI, right click on the IP address resource created for ASCS, select “In-Service” and select node1 to switch the IP back onto node 1.
7. Sapstart SAP on node 1 and ensure that it’s able to run correctly

## 12.4.2.3. Creating the SAP Resource Hierarchy

---

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

A dialog box will appear with a drop-down list box with all recognized recovery kits installed within the cluster. Select **SAP** from the drop-down listing.



Please Select Recovery Kit

Click **Next**

When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful should you encounter an error that might require you to correct previously entered information.

If you click **Cancel** at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the SAP instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either **intelligent** or **automatic**. **Intelligent switchback** requires administrative intervention to switch the instance back to the primary/original server. **Automatic switchback** means the switchback will occur as soon as the primary server comes back on line and re-establishes LifeKeeper communication paths.



Switchback Type

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box.

Click **Next**.

3. Select the Server where you want to place the SAP PAS, ASCS or SCS (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop-down list box.

Server

4. Select the **SAP SID**. This is the system identifier of the SAP PAS, ASCS or SCS system being protected.

SAP SID

Click **Next**

5. Select the SAP Instance Name (ex. ASCS<No.>) (Core Instance first) for the SID being protected.

SAP Instance for PRS

Click **Next**

**Note:** Additional screens may appear related to customization of Protection and Recovery Levels.

6. Select the **IP Child Resource**. This is typically either the Virtual Host IP address noted during SAP installation (SAPINST\_USE\_HOSTNAME) or the IP address needed for failover.

IP child resource

ip-jamie
ip-mutt-104.57
ip-jeff-104.58
none

7. Select or enter the **SAP Tag**. This is a tag name that LifeKeeper gives to the SAP hierarchy. You can select the default or enter your own tag name. The default tag is SAP-<SID>\_<ID>.

SAP Tag

When you click **Create**, the **Create SAP Resource Wizard** will create your SAP resource.

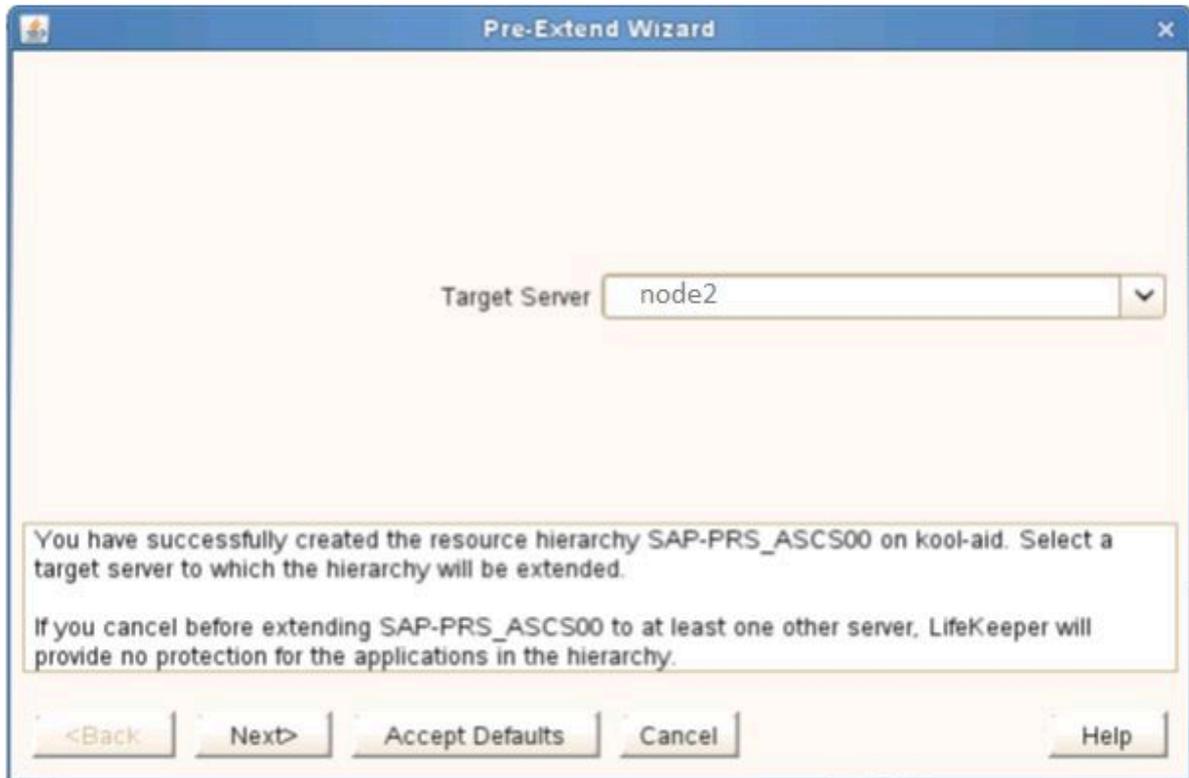
- At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your SAP resource hierarchy. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created. There may also be errors or messages output from the SAP startup scripts that are displayed in the information box.

```
Creating appsuite/sap resource...
Additional information is available in the LifeKeeper and system logs
Wed Aug 24 14:23:01 EDT 2011 restore: 112025: All processes for SAP SID "PRS" and
Instance "ASCS00" are "running" on "kool-aid".
Additional information is available in the LifeKeeper and system logs
Wed Aug 24 14:23:01 EDT 2011 restore: 112003: The SAP Instance "ASCS00" and all required
processes were started successfully during the "restore" on server "kool-aid".
Additional information is available in the LifeKeeper and system logs.
Wed Aug 24 14:23:01 EDT 2011 restore: END successful restore of "SAP-PRS_ASCS00" on
server "kool-aid"
Wed Aug 24 14:23:01 EDT 2011 create: END successful create of "SAP-PRS_ASCS00" on
server "kool-aid"
```

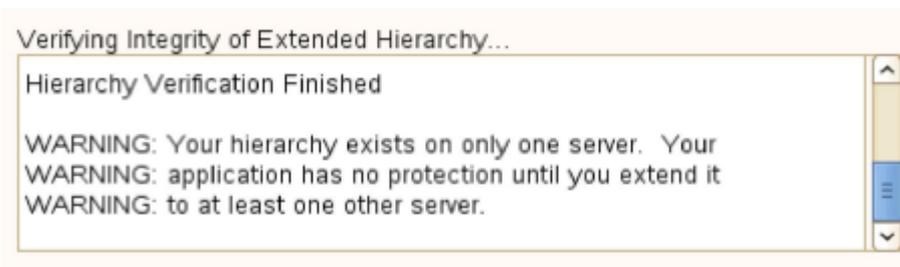
Click **Next**

- Another information box will appear explaining that you have successfully created an SAP resource hierarchy, and you must **Extend** that hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

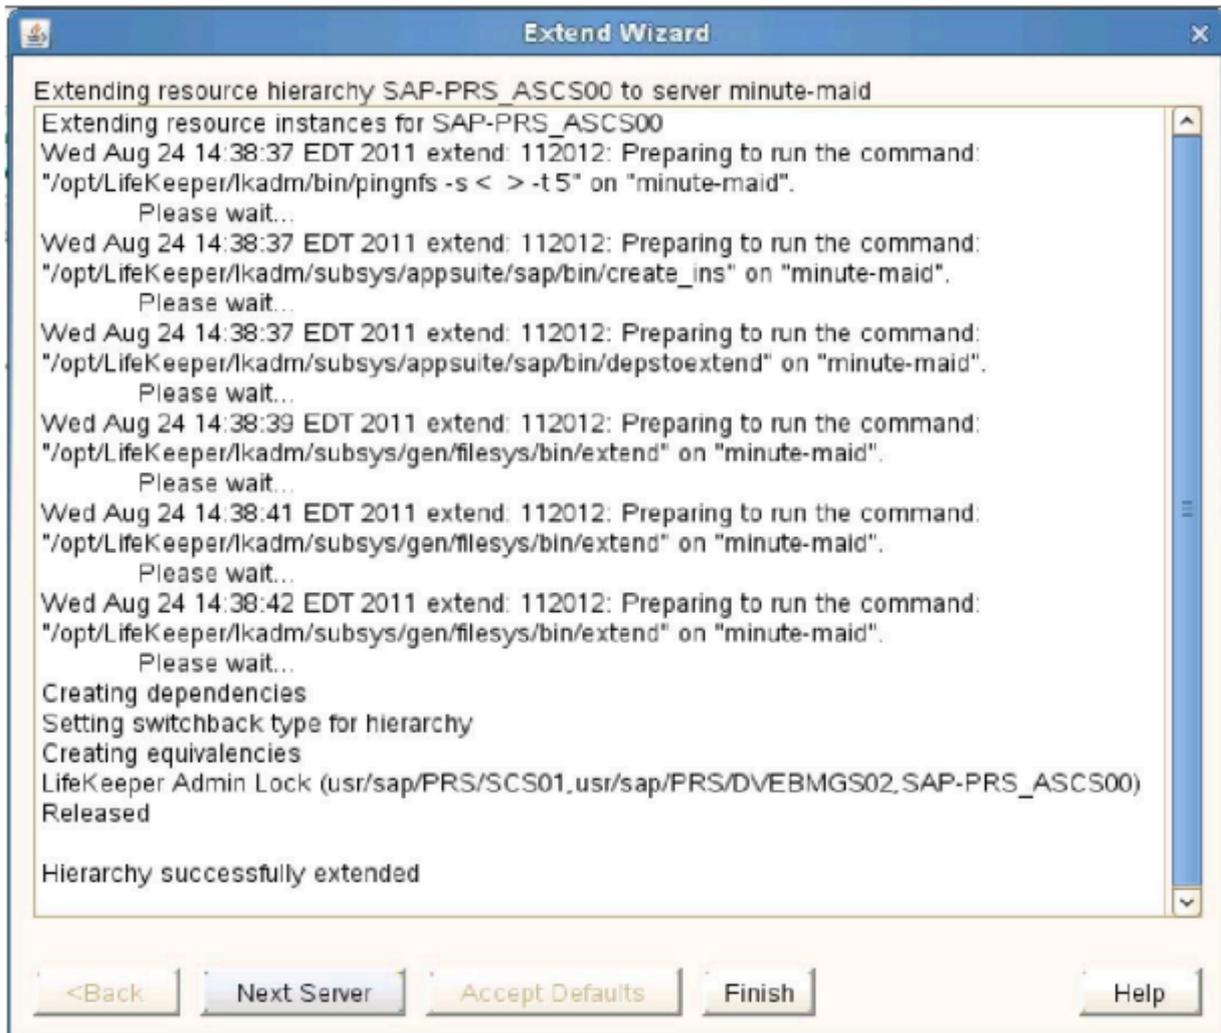
When you click **Next**, LifeKeeper will launch the **Pre-Extend Wizard** that is explained later in this section.



If you click **Cancel** now, a dialog box will appear warning you that you will need to come back and extend your SAP resource hierarchy to another server at some other time to put it under LifeKeeper protection.



10. The **Extend Wizard** dialog will appear stating **Hierarchy successfully extended**. Click **Finish**.



The **Hierarchy Integrity Verification** dialog appears. Once Hierarchy Verification finishes, click **Done** to exit the **Create Resource Hierarchy** menu selection.

### Hierarchy with the Core as the Top Level



While LifeKeeper can be used to protect the PAS and AAS servers, most customers would simply use them as independent standby servers with no additional HA on them. This guide does not cover their protection steps but you can refer to our [SAP Recovery Kit](#) documentation for details and steps.

## 12.4.3. ASCS + ERS with NFS

---

Here are the steps to setup ASCS + ERS with NFS:

- [General Setup Steps](#)
- [Installing SAP](#)
- [Setting up NFS](#)
- [Creating an NFS Resource Hierarchy](#)
- [Creating the SAP Resource Hierarchy](#)
- [Create the ERS Resource](#)

## 12.4.3.1. General Setup Steps

---

1. Create Virtual IP, done in earlier steps on node1, extend, done in earlier steps
2. Create EC2 resource and create as dependency for virtual IP, done in earlier steps
3. Install SAP on node1 on “virtual hostname” based on “virtual IP”
4. Stopsap on node 1
5. Use the LifeKeeper GUI to “In-service” the virtual IP to node 2, and Install SAP on node 2 on “virtual hostname” based on “virtual IP”
6. Stopsap on node 2 and modify profile files on both nodes (see below)
7. Use the LifeKeeper GUI to “In-service” the virtual IP back to node 1
8. Create replication resource for the mount points needed for SAP, done in earlier steps, as advised by SAP consultants
9. Startsap on node1 and ensure SAP is working properly
10. In /etc/default/LifeKeeper on both nodes add the follow entries to the end:

```
SAP_EXPERTMODE=1
```

```
SAP_NFS_CHECK_IGNORE=1
```

```
SAP_DB_CHECK_IGNORE=1
```

11. Re-run the LifeKeeper setup program to add the SAP recovery kit

Mount the sps.img file (downloaded as per earlier steps) using the following command:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

Where PATH is the path to the image

IMAGE\_NAME is the name of the image

MOUNT\_POINT is the path to mount location

Change to the sps.img mounted directory and type the following: ./setup -k

You will now be shown a menu of recovery kits available. Select the recovery kit for SAP by using the arrow keys and pressing the <spacebar> to select, press <enter> to continue and complete the installation.

12. Setup NFS servers
13. Copy file systems onto the SAP server and create replication resources on the file systems for redundancy and failover
14. Create NFS resources following the [NFS Recovery Kit guide](#)

Simplified steps are given [here](#)

15. Create SAP resources following the [SAP Recovery Kit guide](#)

Simplified steps are given [here](#)

## 12.4.3.2. Installing SAP

---

1. ASCS and ERS should be installed based on “virtual hostname” based on “virtual IP”, which should have been added to hosts files during earlier installation steps. Be sure to do so if they have not already been done prior to installing SAP.

- When installing SAP (specifically ASCS or SCS) you need to specify the `SAPINST_USE_HOSTNAME=vip` option. This is not required for ERS. (**Note:** Document the `SAPINST_USE_HOSTNAME` virtual IP address as it will be used later during creation of the SAP resources in LifeKeeper.)

**Note:** Specify `sapinst SAPINST_USE_HOSTNAME=vip` where **vip** is the virtual IP that will float between the nodes.

Run `./sapinst SAPINST_USE_HOSTNAME= {hostname}`

- In seven phases, the **Core Services** should be created and started. If permission errors occur on `jdbccconnect.jar`, go to `/sapmnt/STC/exe/uc/linuxx86_64` and make that directory as well as file `jdbccconnect.jar` writable (`chmod 777 ---`).

Enqueue replication should be configured and checked working based on SAP documentation and best practices.

2. ASCS and ERS profiles should be pointing to local mount point containing `/usr/sap`, `sapmnt` or any other necessary for SAP files in your environment, the actual files will be moved onto NFS mount points after it is installed and configured.

The ASCS and ERS instance profiles must be modified in order to prevent the enqueue server and enqueue replicator processes from being automatically restarted by the `sapstart` utility. After updating the instance profiles, SAP Start Service for each of these instances must be restarted in order for the changes to take effect. Follow the steps provided in [Modify ASCS and ERS Instance Profile Settings](#), then return to this page to continue the setup process.

3. Sapstop SAP on node 1
4. Using the LifeKeeper GUI, right click on the IP address resource created for ASCS, select “In-Service” and select node2 to switch the IP onto node 2.
5. Repeat step 1 to install SAP onto node 2 and ensure that it’s able to run correctly
6. Sapstop SAP on node 2
7. Using the LifeKeeper GUI, right click on the IP address resource created for ASCS, select “In-Service” and select node1 to switch the IP back onto node 1.
8. Sapstart SAP on node 1 and ensure that it’s able to run correctly

## 12.4.3.3. Setting up NFS

NFS server should have been installed on both cluster nodes prior to installation of SIOS as a prerequisite.

Create the NFS exports based on the SAP’s requirements in your SAP design. Below are examples that may be use as a guide but not a representation of your SAP environment.

LifeKeeper maintains NFS share information using inodes; therefore, every NFS share is required to have a unique inode. Since every file system root directory has the same inode, NFS shares must be at least one directory level down from root in order to be protected by LifeKeeper. For example, referring to the information above, if the `/usr/sap/trans` directory is NFS shared on the SAP server, the `/trans` directory is created on the shared storage device which would require mounting the shared storage device as `/usr/sap`. It is not necessarily desirable, however, to place all files under `/usr/sap` on shared storage which would be required with this arrangement. To circumvent this problem, it is recommended that you create an `/exports` directory tree for mounting all shared file systems containing directories that are NFS shared and then create a soft link between the SAP directories and the `/exports` directories, or alternately, locally NFS mount the NFS shared directory. (**Note:** The name of the directory that we refer to as `/exports` can vary according to user preference; however, for simplicity, we will refer to this directory as `/exports` throughout this documentation.) For example, the following directories and links/ mounts for our example on the SAP Primary Server would be:

For the <code>/usr/sap/trans</code> share	
Directory	Notes
<code>/trans</code>	created on share file system and shared through NFS
<code>/exports/usr/sap</code>	mounted to <code>/</code> (on shared file system)
<code>/user/sap/trans</code>	soft linked to <code>/exports/usr/sap/trans</code>

The following directories and links for the `<sapmnt>/<SAPSID>` share would be:

For the <code>&lt;sapmnt&gt;/&lt;SAPSID&gt;</code> share	
Directory	Notes
<code>/&lt;SAPSID&gt;</code>	created on shared file systems and shared through NFS
<code>/exports/sapmnt</code>	mounted to <code>/</code> (on shared file system)
<code>&lt;sapmnt&gt;/&lt;SAPSID&gt;</code>	NFS mounted to <code>&lt;virtual SAP server&gt;:/exports/sapmnt/&lt;SAPSID&gt;</code>

### Local NFS Mounts

The recommended directory structure for SAP in a LifeKeeper environment requires a locally mounted NFS share for one or more SAP system directories. If the NFS export point for any of the locally mounted NFS shares becomes unavailable, the system may hang while waiting for the export point to become available again. Many system operations will not work correctly, including a system reboot. You

should be aware that the NFS server for the SAP cluster should be protected by LifeKeeper and should not be manually taken out of service while local mount points exist.

### Location of <INST> Directories

Since the /usr/sap/<SAPSID> path is not NFS shared, it can be mounted to the root directory of the file system. The /usr/sap/<SAPSID> path contains the SYS subdirectory and an <INST> subdirectory for each SAP instance that can run on the server. For certain configurations, there may only be one <INST> directory, so it is acceptable for it to be located under /usr/sap/<SAPSID> on the shared file system. For other configurations, however, the backup server may also contain a local AS instance whose <INST> directory should not be on a shared file system since it will not always be available. To solve this problem, it is recommended that for certain configurations, the PAS's, ASCS's or SCS's /usr/sap/<SAPSID>/<INST>, /usr/sap/<SAPSID>/<ASCS-INST> or /usr/sap/<SAPSID>/<SCS-INST> directories should be mounted to the shared file system instead of /usr/sap/<SAPSID> and the /usr/sap/<SAPSID>/SYS and /usr/sap/<SAPSID>/<AS-INST> for the AS should be located on the local server.

For example, the following directories and mount points should be created for the ABAP+Java Configuration

Directory	Notes
usr/sap/<SAPSID>/DVEBMS<No.>	mounted to / (Replicated non-NFS file system )
usr/sap/<SAPSID>/SCS<No.>	mounted to / (Replicated non-NFS file system )
usr/sap/<SAPSID>/ERS<No.> (for SCS instance)	should be locally mounted on all cluster nodes or mounted from a NAS share (should not be mounted on shared storage)
usr/sap/<SAPSID>/ASCS<Instance No. >	mounted to / (Replicated from non-NFS file system)
usr/sap/<SAPSID>/ERS<No.> (for ASCS instance)	should be locally mounted on all cluster nodes or mounted from a NAS share (should not be mounted on shared storage)
usr/sap/<SAPSID>AS<Instance No. >	created for AS backup server

### Mount NFS and Move File Systems

After mount points has been created for the main SAP file systems, mount them accordingly (required).

At this point, stop all SAP services before proceeding with these steps.

```
mount /dev/sap/sapmnt /exports/sapmnt
```

```
mount /dev/sap/saptrans /exports/saptrans
```

## Move Data to NFS

1. Edit `/etc/exports` and insert the mount points for SAP's main directories.

```
/exports/sapmnt * (rw, sync, no_root_squash)
```

```
/exports/saptrans * (rw, sync, no_root_squash)
```

## Example NFS Export



Replace each occurrence of `<nfsvip>` with the appropriate virtual IP that will be used for that NFS share. Depending on the design of your SAP system, different virtual IP's may be used to share different filesystems.

```
# more /etc/exports
```

```
/exports/sapmnt 10.2.0.69(rw, sync, all_squash, anonuid=0, anongid=1001)
```

```
/exports/sapmnt 10.2.0.11(rw, sync, all_squash, anonuid=0, anongid=1001)
```

```
/exports/usr/sap/<instance name>/ASCS01 10.2.0.69(rw, sync, all_squash, anonuid=0, anongid=1001)
```

```
/exports/sap/<instance name>/ASCS01
10.2.0.11(rw, sync, all_squash, anonuid=0, anongid=1001)
```

```
# more /etc/fstab
```

```
#
```

```
# /etc/fstab
```

```
# Created by anaconda on Mon Nov 9 20:20:10 2015
```

```
#
```

```
# Accessible filesystems, by reference, are maintained under '/dev/disk'
```

```
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
```

```
#
```

```
UUID=367df610-4210-4a5a-8c8d-51ddf499fc17 / xfs defaults 0 0
```

```
/dev/xvdb swap swap defaults 0 0
```

```
/dev/xvdc /tmp xfs nodev,nosuid,noexec,relatime 0 0
```

```
/dev/xvdp1 /var xfs defaults 0 0
```

```
/dev/xvdp2 /var/log xfs defaults 0 0
```

```
/dev/xvdp3 /var/log/audit xfs defaults 0 0
```

```
/dev/xvdp4 /home xfs defaults,nodev 0 0
```

```
/tmp /var/tmp none bind,nodev,nosuid 0 0
```

```
/dev/xvdj /usr/sap xfs defaults 0 0
```

```
/dev/xvdg /exports/usr/sap/P4G/ASCS01 xfs defaults 0 0
```

```
/dev/xvdh /usr/sap/P4G/D00 xfs defaults 0 0
```

```
/dev/xvdi /sapcd xfs defaults 0 0
```

```
/dev/xvdk /exports/sapmnt xfs defaults 0 0
```

```
<nfsvip>:/exports/usr/sap/P4G/ASCS01 /usr/sap/<instance name>/ASCS01 nfs
nfsvers=3,proto=udp,rw,sync,bg 0 0
```

```
<nfsvip>:/exports/sapmnt /sapmnt nfs nfsvers=3,proto=udp,rw,sync,bg 0 0
```

```
<nfsvip>:/exports/usr/sap/P4G/ASCS01 /usr/sap/PG4/ASCS01 nfs nfsvers=3,proto=udp,rw,sync,bg 0 0
(Note: This ERS entry will only be present if using an ERSv2 configuration with a shared ERS
filesystem.)
```

2. Start the NFS server using the `systemctl start nfs-server.service` command. If the NFS server is already active, you may need to do an `exportfs -va` to export those mount points.
3. On both node1 & 2, execute the following mount commands (**note the usage of udp; this is important for failover and recovery**), ensuring you are able to mount the NFS shares.

```
mount {virtual ip}:/exports/sapmnt/<PG4> /sapmnt/<PG4> -o rw,sync,bg,udp
```

```
mount {virtual ip}:/exports/saptrans /usr/sap/trans -o rw,sync,bg,udp
```

4. From node 1, copy the necessary file systems from the `/usr/sap` and `/sapmnt` or any other required files into the NFS mount points, mounted from the NFS servers onto node 1.

5. Log in to SAP and start SAP (after su to stcadm).

```
startsap sap{No.}
```

6. Make sure all processes have started.

```
ps -ef | grep en.sap (2 processes)
```

```
ps -ef | grep ms.sap (2 processes)
```

```
ps -ef | grep dw.sap (17 processes)
```

“SAP Logon” or “SAP GUI for Windows” is an SAP supplied Windows client the Windows client. The program can be downloaded from the SAP download site. The virtual IP address may be used as the “Application Server” on the **Properties** page. This ensures that a connection to the primary machine where the virtual ip resides is active.

7. If not already done, create the Data Replication Cluster resource on the NFS shares mount points to replicate the data from node1 to node2.

## Reduce Switchover Times using SAP and NFSv4/TCP

In SAP/NFS environment the following changes can be made to reduce switchover times with SAP and NFSv4/TCP:

**Take all LK NFS resources out of service**, then execute:

```
# systemctl stop nfs
# echo 10 > /proc/fs/nfsd/nfsv4gracetime
# echo 10 > /proc/fs/nfsd/nfsv4leasetime
# systemctl start nfs
```

**!** **Note:** LEASE time must be equal to GRACE time, so set both. With this change, remove of NFS resources now takes only 10 seconds, rather than the default of 90.

To set the grace and lease times persistently, for example:

**On RHEL 8, execute the following:**

```
# nfsconf --set nfsd grace-time 10
# nfsconf --set nfsd lease-time 10
# systemctl restart nfs-server
```

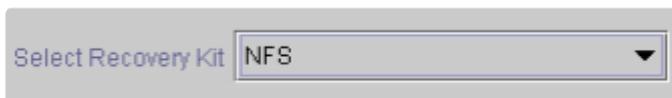
## 12.4.3.4. Creating a Resource Hierarchy

To create a resource instance from the primary server, you should complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

To change a selection already entered or if an error message is encountered during any step in the creation of your NFS resource hierarchy, use the **Back** button to change your selection or make corrections (assuming the **Back** button is enabled).

A dialog box will appear with a drop-down menu listing all recognized Recovery Kits installed within the cluster. Select **NFS** from the drop-down menu

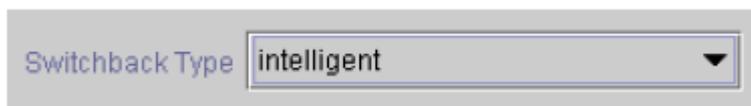


Select Recovery Kit

Click **Next** to proceed to the next dialog box.

**Note:** If you click the **Cancel** button at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the NFS instance will be switched back to the primary server when it comes back into service after a failover to the backup server. Choose either **Intelligent** or **Automatic**. **Intelligent switchback** requires administrative intervention to switch the instance back to the primary/original server. **Automatic switchback** means the switchback will occur as soon as the primary server comes back on line and is under LifeKeeper protection



Switchback Type

The switchback type can be changed later, if desired, from the **General** tab of the **Resource Properties** dialog box.

3. Select the **Server** where you want to create the NFS resource (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop-down menu



Server

Click **Next** to proceed to the next dialog box.

4. The **Export Point** dialog displays a drop-down list of export points for NFS file systems that meet the following criteria:

- The export point has been exported by NFS.
- The export point is on a shared drive.
- If the underlying file system is LifeKeeper-protected, it must be in service and have the highest priority on the server selected on the **Server** dialog.
- The exported point with fsid=0 and its sub directory export point are not listed.

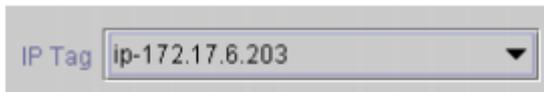
Select the NFS export point to be protected from the drop-down list.



The screenshot shows a dialog box with a label "Export Point" on the left. To its right is a dropdown menu with a light gray background and a dark border. The text "/export/public" is displayed inside the dropdown, and a small downward-pointing arrow is visible on the right side of the dropdown box.

Click **Next** to proceed to the next dialog box.

5. The **IP Tag** dialog displays a drop-down list of tags corresponding to virtual IP addresses currently under LifeKeeper protection and in service on the server where the NFS resource is being created. Select the **tag** for the virtual IP address used by clients to access the protected NFS file system.



The screenshot shows a dialog box with a label "IP Tag" on the left. To its right is a dropdown menu with a light gray background and a dark border. The text "ip-172.17.6.203" is displayed inside the dropdown, and a small downward-pointing arrow is visible on the right side of the dropdown box.

**Note:** At this point, LifeKeeper will check to ensure that there is a protected IP resource available.

**Note:** If you are using other LifeKeeper Recovery Kits that have virtual IP address dependencies, you should create a different virtual IP address for the NFS resource. Otherwise, if the virtual IP resource fails over to a backup server, all of the resources that depend on that IP resource will fail over at the same time.

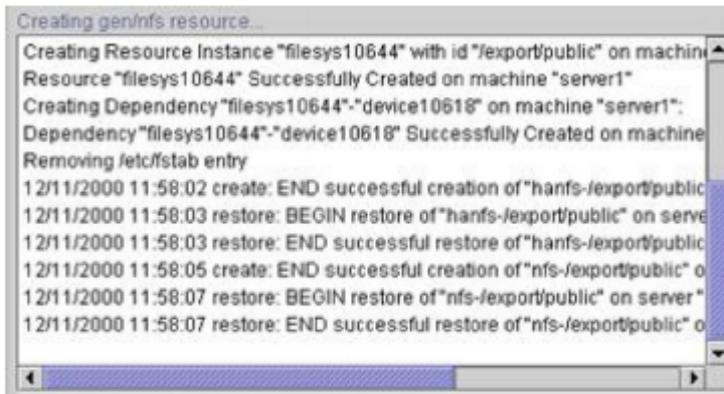
Click **Next** to proceed to the next dialog box

6. Select or enter the **NFS Tag**. This is a tag name given to the NFS hierarchy. You can select the default or enter your own tag name.



The screenshot shows a dialog box with a label "NFS Tag" on the left. To its right is a text input field with a light gray background and a dark border. The text "nfs-/export/public" is entered into the field.

When you click the **Create** button, the **Create Resource Wizard** will create your NFS resource.



When you click **Next**, LifeKeeper will launch the **Pre-Extend Wizard** that is discussed in Extending Your Hierarchy.

**Note:** The NFS resource hierarchy should be created successfully at this point. However, error messages may be encountered indicating that the new NFS instance has failed to start correctly. Note that the new NFS hierarchy must be started (In Service) before it can be extended to another system. If startup fails, you may pause at this point and correct the problem based on the error message displayed. If the errors are not correctable, you will only be given the choice to cancel which cancels the resource create.

Bring the new hierarchy In Service before proceeding with extending your hierarchy.

\*\*\* Repeat the steps above to create an additional resource hierarchy for each NFS share.

## 12.4.3.5. Creating the SAP Resource Hierarchy

---

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

A dialog box will appear with a drop-down list box with all recognized recovery kits installed within the cluster. Select **SAP** from the drop-down listing.



Please Select Recovery Kit

Click **Next**.

When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful should you encounter an error that might require you to correct previously entered information.

If you click **Cancel** at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the SAP instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either **intelligent** or **automatic**. **Intelligent switchback** requires administrative intervention to switch the instance back to the primary/original server. **Automatic switchback** means the switchback will occur as soon as the primary server comes back on line and re-establishes LifeKeeper communication paths.



Switchback Type

The switchback type can be changed later from the **General** tab of the **Resource Properties** dialog box.

Click **Next**

3. Select the Server where you want to place the SAP PAS, ASCS or SCS (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop-down list box.



Server

4. Select the **SAP SID**. This is the system identifier of the SAP PAS, ASCS or SCS system being protected.

A screenshot of a web form showing a dropdown menu for 'SAP SID'. The selected value is 'PRS'. The dropdown arrow is visible on the right side of the field.

Click **Next**.

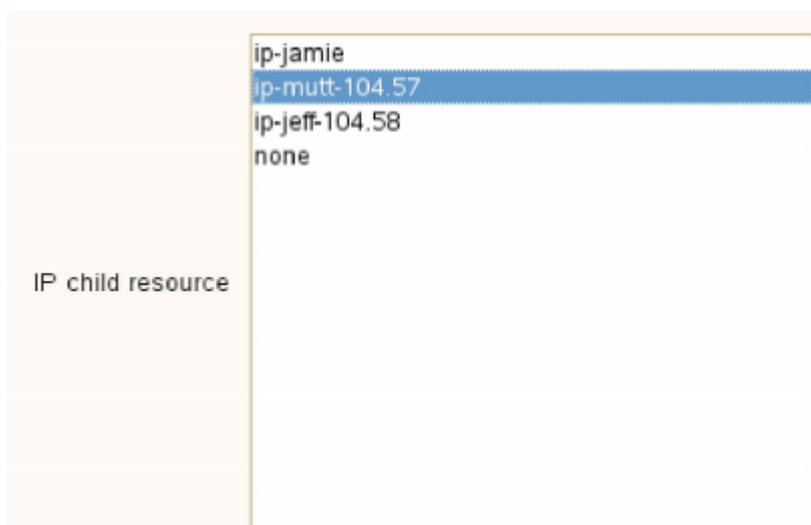
5. Select the SAP Instance Name (ex. ASCS<No.>) (Core Instance first) for the SID being protected.

A screenshot of a web form showing a dropdown menu for 'SAP Instance for PRS'. The selected value is 'ASCS00'. The dropdown arrow is visible on the right side of the field.

Click **Next**.

**Note:** Additional screens may appear related to customization of Protection and Recovery Levels.

6. Select the **IP Child Resource**. This is typically either the Virtual Host IP address noted during SAP installation (SAPINST\_USE\_HOSTNAME) or the IP address needed for failover.

A screenshot of a web form showing a list of IP child resources. The list is titled 'IP child resource' and contains four items: 'ip-jamie', 'ip-mutt-104.57', 'ip-jeff-104.58', and 'none'. The item 'ip-mutt-104.57' is highlighted with a blue background.

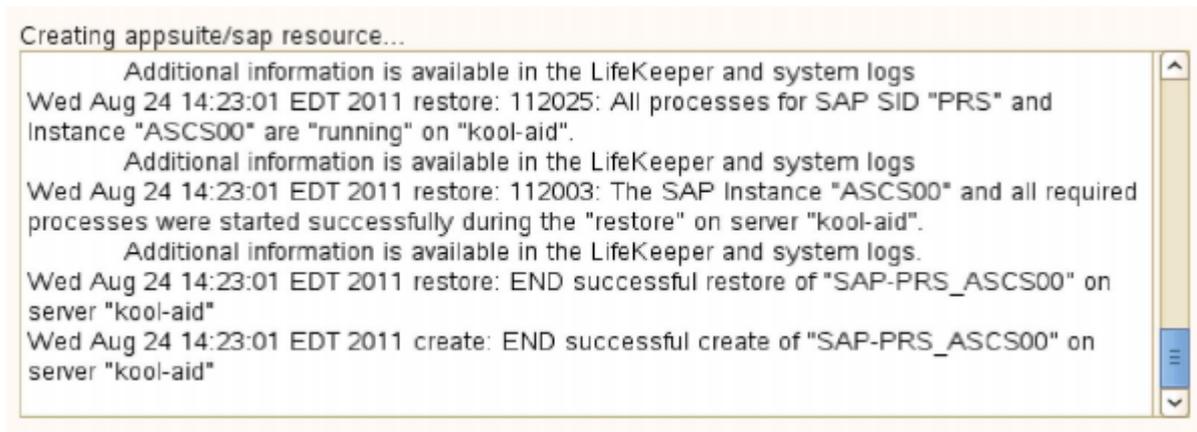
7. Select or enter the **SAP Tag**. This is a tag name that LifeKeeper gives to the SAP hierarchy. You can select the default or enter your own tag name. The default tag is *SAP-<SID>\_<ID>*.

A screenshot of a web form showing an input field for 'SAP Tag'. The text 'SAP-PRS\_ASCS00' is entered into the field.

When you click **Create**, the **Create SAP Resource Wizard** will create your SAP resource.

8. At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your SAP resource hierarchy. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created. There

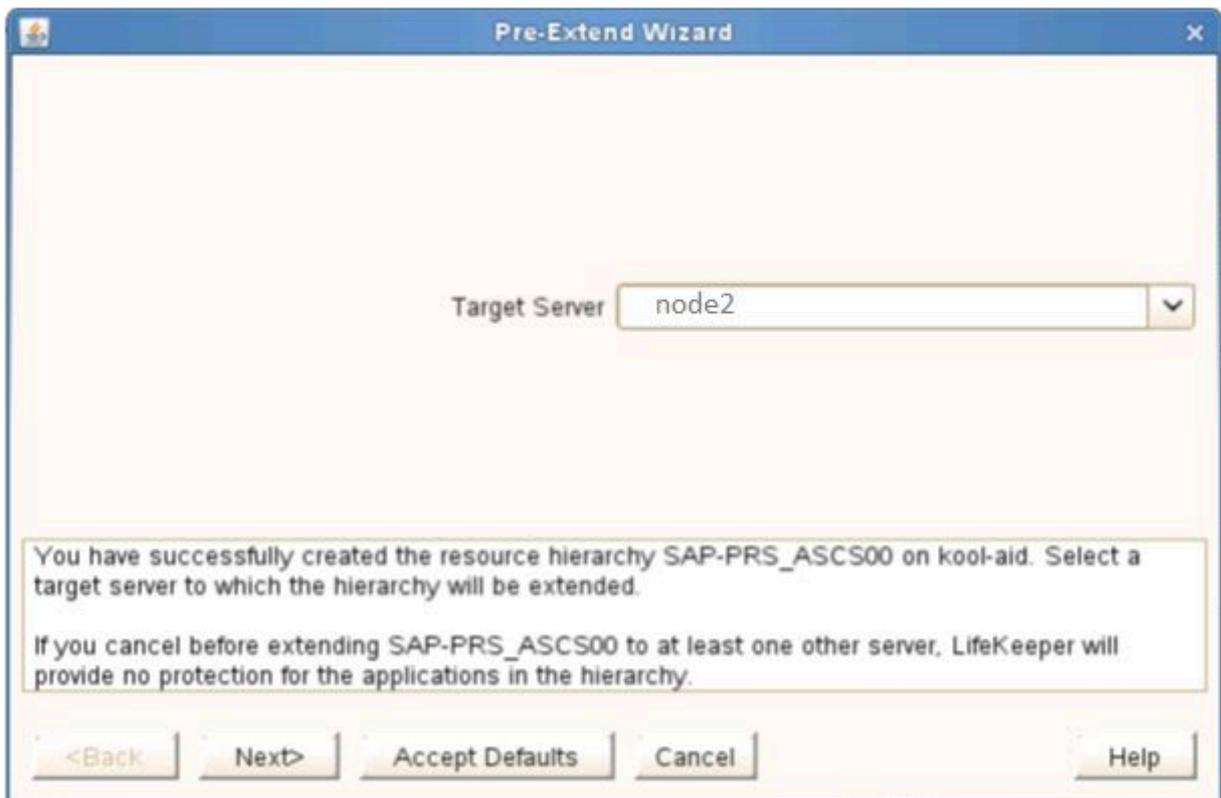
may also be errors or messages output from the SAP startup scripts that are displayed in the information box.



Click **Next**.

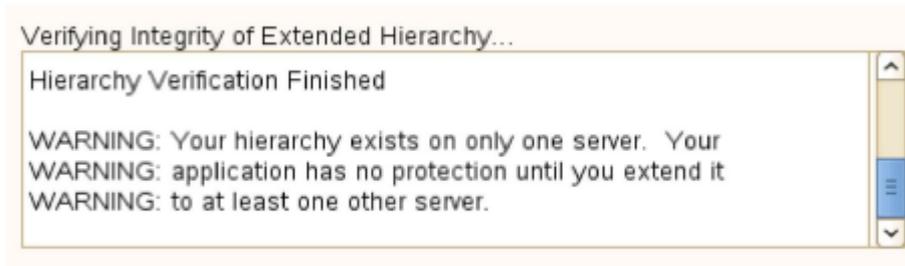
- 9. Another information box will appear explaining that you have successfully created an SAP resource hierarchy, and you must **Extend** that hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

When you click **Next**, LifeKeeper will launch the **Pre-Extend Wizard** that is explained later in this section

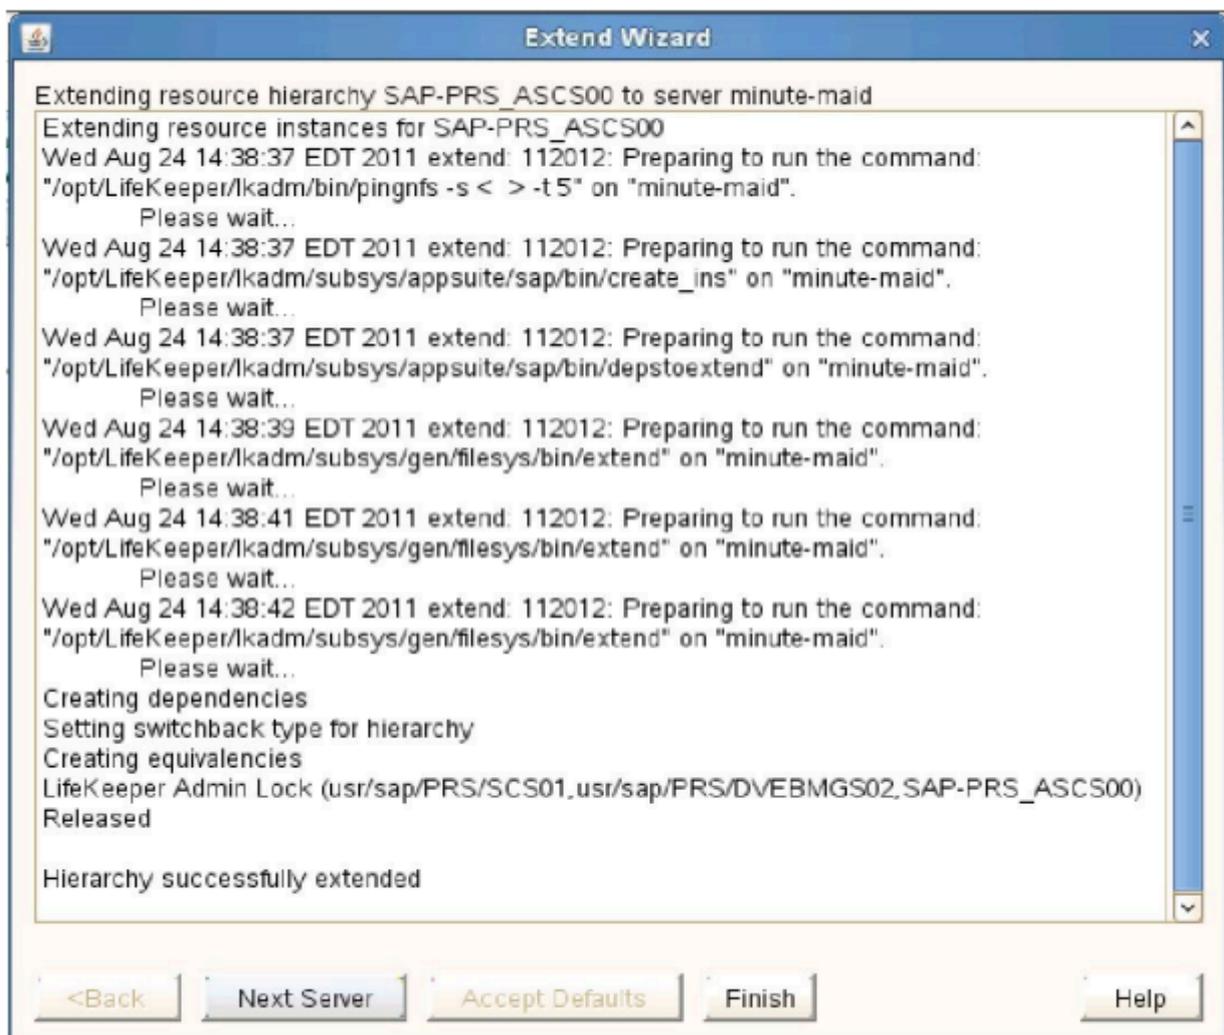


If you click **Cancel** now, a dialog box will appear warning you that you will need to come back and extend your SAP resource hierarchy to another server at some other time to put it under

LifeKeeper protection.



10. The **Extend Wizard** dialog will appear stating **Hierarchy successfully extended**. Click **Finish**.



The **Hierarchy Integrity Verification** dialog appears. Once Hierarchy Verification finishes, click **Done** to exit the **Create Resource Hierarchy** menu selection.

## Hierarchy with the Core as the Top Level



## 12.4.3.6. Create the ERS Resources

---

The ERS resource provides additional protection against a single point of failure of a Core Instance (Central Services Instance) or enqueue server process. When a Core Instance (Central Services Instance) fails and is restarted, it will retrieve the current status of the lock table and transactions. The result is that, in the event of the enqueue server failure, no transactions or updates are lost and the service for the SAP system continues.

Perform the following steps to create this ERS Resource.

1. For this same SAP SID, repeat the above steps to create the ERS Resource selecting your **ERS instance** when prompted.
2. You will then be prompted to select **Dependent Instances**. Select the **Core Resource** that was created above, and then click **Next**.
3. Follow prompts to **extend resource hierarchy**.
4. Once **Hierarchy Successfully Extended** displays, select **Finish**.
5. Select **Done**

**Note:** The Enqueue Replication Server (ERS) resource will be in-service (ISP) on the primary node in your cluster. However, the architecture and function of the ERS requires that the actual processes for the instance run on the backup node. This allows the standby server to hold a complete copy of the lock table information for the primary server and primary enqueue server instance. When the primary server running the enqueue server fails, it will be restarted by LifeKeeper on the backup server on which the ERS process is currently running. The lock table (replication table) stored on the ERS is transferred to the enqueue server process being recovered and the new lock table is created from it. Once this process is complete, the active replication server is then deactivated (it closes the connection to the enqueue server and deletes the replication table). LifeKeeper will then restart the ERS processes on the new current backup node (formerly the primary) which has been inactive until now. Once the ERS process becomes active, it connects to the enqueue server and creates a replication table. For more information on the ERS process and SAP architecture features, visit <http://help.sap.com> and search for **Enqueue Replication Service**.

## Hierarchy with ERS as Top Level



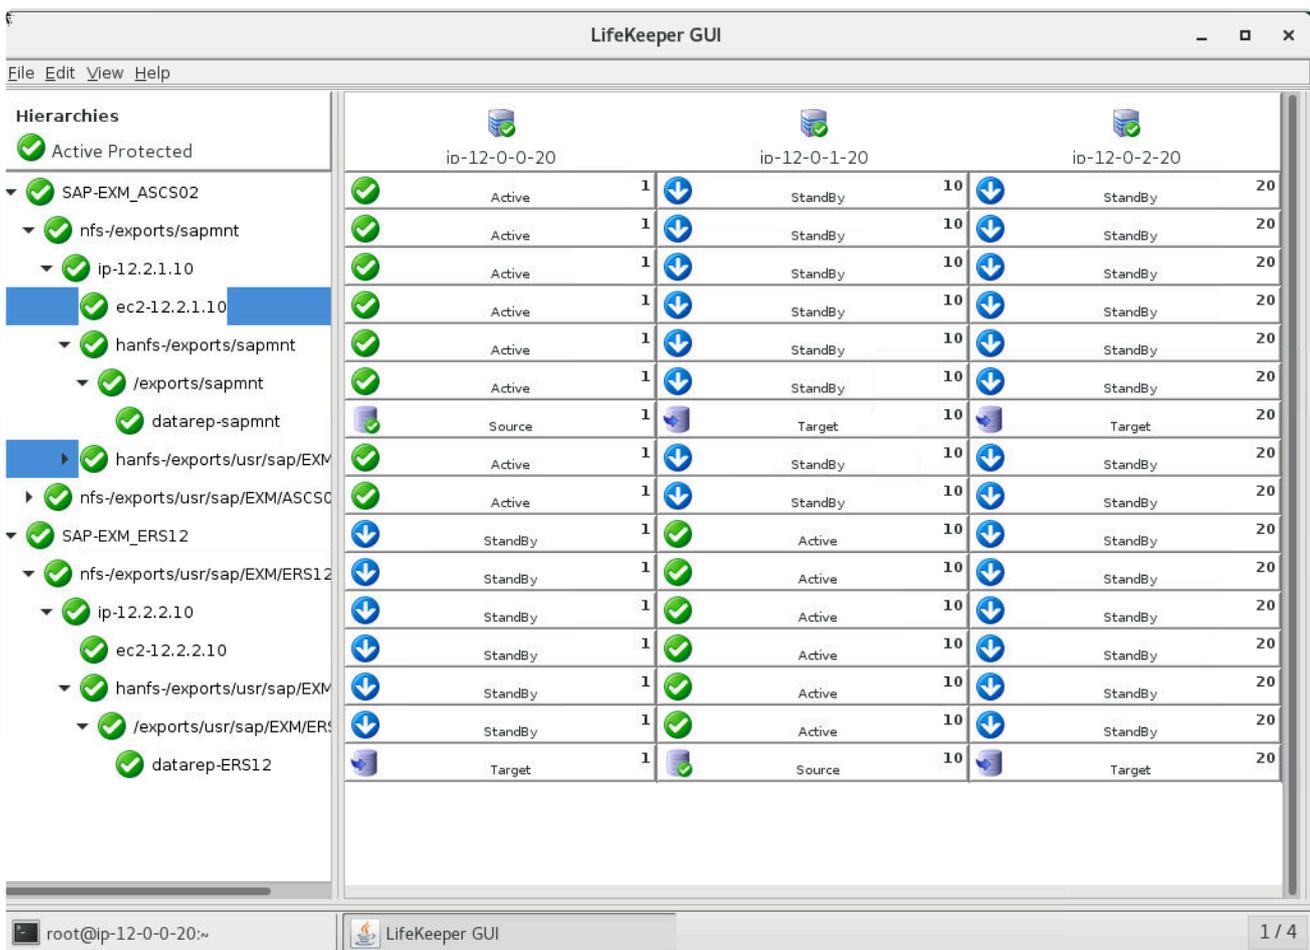
While LifeKeeper can be used to protect the PAS and AAS servers, most customers would simply use them as independent standby servers with no additional HA on them. This guide does not cover their protection steps but you can refer to our [SAP Recovery Kit documentation](#) for details and steps.

# 12.4.3.7. Enforcing ASCS/ERS Avoidance Behavior When Using ENSA2/ERSv2

ERSv2 is intended to be active (in-service) on a node in the cluster where the ASCS resource is not active. The ERS quickCheck will automatically transfer the ERS hierarchy if ERS and ASCS are active on the same node and another node is available. To avoid getting into the situation where ASCS and ERS are both active (in-service) on the same node after a switchover or failover, a gen/app terminal leaf resource can be created to automatically route the in-service to a node where the corresponding resource hierarchy is not active (in-service). To facilitate creating this terminal leaf node a new utility is provided, /opt/LifeKeeper/bin/create\_terminal\_leaf (1M).

To create the avoidance terminal leaf the utility takes two parameters, the ASCS and ERS hierarchy root resource tags. The two hierarchies should be fully extended to all nodes in the cluster and in-service on a node in the cluster. It does not require that the hierarchies be in-service on the node where the utility is run as long as the utility is run on a node in the cluster. The terminal leaf node will be named "avoid\_<tag>" where tag is the appropriate root node to be avoided. The terminal leaf node will be attached as a child dependency on each branch of the hierarchy.

For example, in a configuration with SAP-EXM\_ASCS02 and SAP-EXM\_ERS12 as the root nodes:



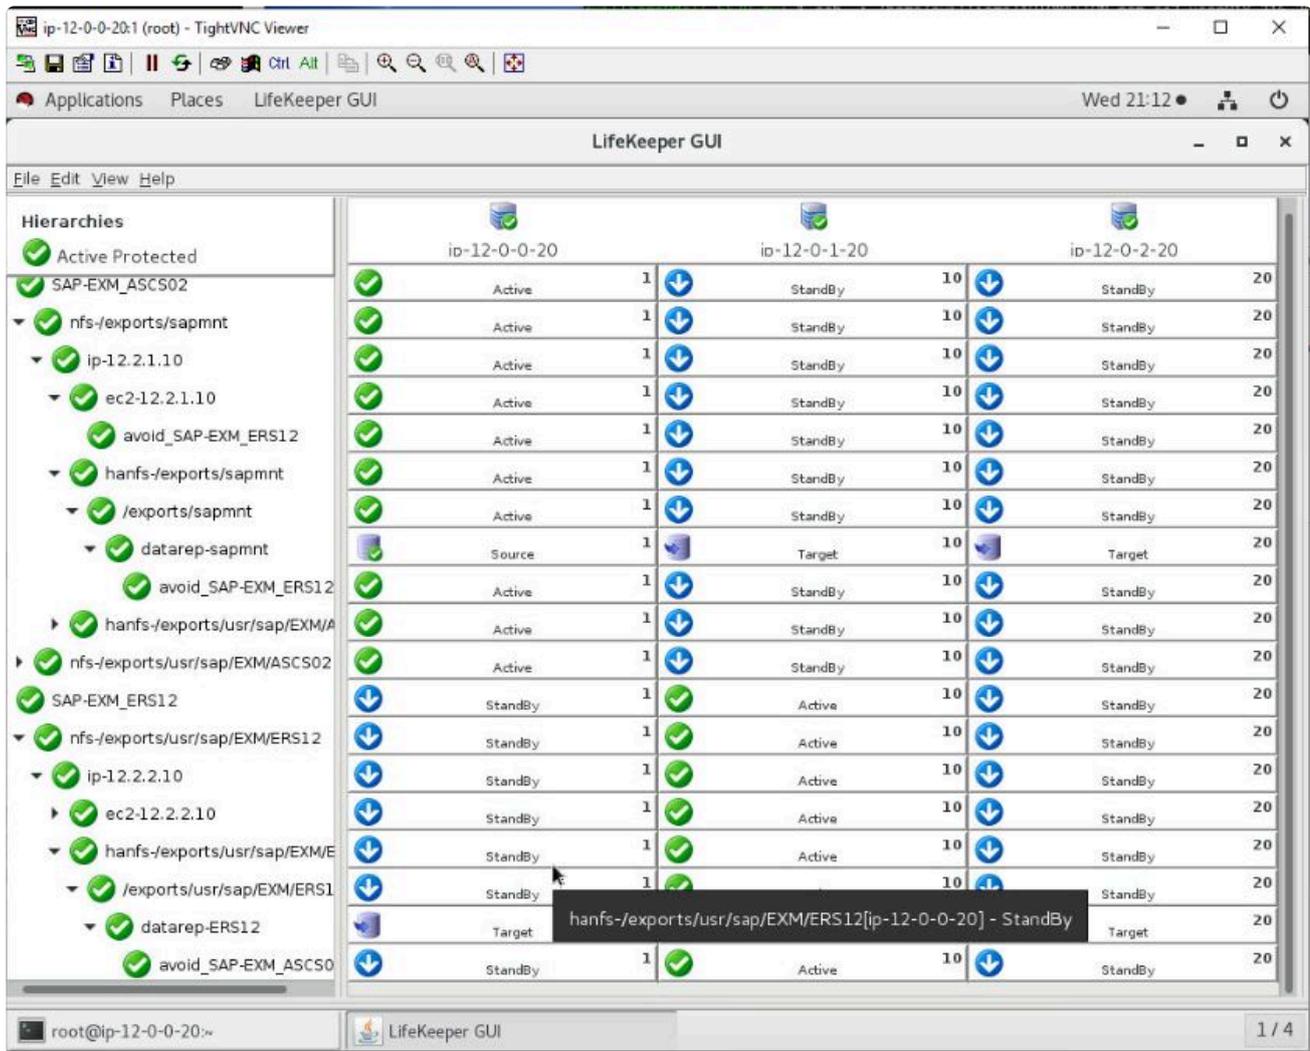
The appropriate terminal leaf nodes are created by running /opt/LifeKeeper/bin/create\_terminal\_leaf SAP-EXM\_ASCS02 SAP-EXM\_ERS12

```

[root@ip-12-0-2-20 ~]# /opt/LifeKeeper/bin/create_terminal_leaf SAP-EXM_ASCS02 SAP-EXM_ERS12
Create avoidance terminal leaf resource for root hierarchy 'SAP-EXM_ASCS02' and 'SAP-EXM_ERS12'.
creapphier: WARNING No quickCheck script specified
creapphier: WARNING No local recovery script specified
BEGIN create of "avoid_SAP-EXM_ERS12"
creating resource "avoid_SAP-EXM_ERS12"
resource "avoid_SAP-EXM_ERS12" successfully created
restoring resource "avoid_SAP-EXM_ERS12"
BEGIN restore of "avoid_SAP-EXM_ERS12"
Attempting to avoid resource avoid_SAP-EXM_ASCS02. Since the resource was not found on ip-12-0-0-20, restore of avoid_SAP-EXM_ERS12 is successful on ip-12-0-0-20.
END successful restore of "avoid_SAP-EXM_ERS12"
resource "avoid_SAP-EXM_ERS12" restored
END successful create of "avoid_SAP-EXM_ERS12"
Extending resource instances for avoid_SAP-EXM_ERS12
Creating dependencies
Setting switchback type for hierarchy
Creating equivalencies
LifeKeeper Admin Lock (avoid_SAP-EXM_ERS12) Released
Hierarchy successfully extended
BEGIN extend of "avoid_SAP-EXM_ERS12"
END successful extend of "avoid_SAP-EXM_ERS12"
Extending resource instances for avoid_SAP-EXM_ERS12
Creating dependencies
Creating equivalencies
LifeKeeper Admin Lock (avoid_SAP-EXM_ERS12) Released
Hierarchy successfully extended
BEGIN extend of "avoid_SAP-EXM_ERS12"
END successful extend of "avoid_SAP-EXM_ERS12"
Create dependency on 'ip-12-0-0-20' for 'avoid_SAP-EXM_ERS12' to 'ec2-12.2.1.10'.
Create dependency on 'ip-12-0-1-20' for 'avoid_SAP-EXM_ERS12' to 'ec2-12.2.1.10'.
Create dependency on 'ip-12-0-2-20' for 'avoid_SAP-EXM_ERS12' to 'ec2-12.2.1.10'.
Create dependency on 'ip-12-0-0-20' for 'avoid_SAP-EXM_ERS12' to 'datarep-ASCS02'.
Create dependency on 'ip-12-0-1-20' for 'avoid_SAP-EXM_ERS12' to 'datarep-ASCS02'.
Create dependency on 'ip-12-0-2-20' for 'avoid_SAP-EXM_ERS12' to 'datarep-ASCS02'.
Create dependency on 'ip-12-0-0-20' for 'avoid_SAP-EXM_ERS12' to 'datarep-sapmnt'.
Create dependency on 'ip-12-0-1-20' for 'avoid_SAP-EXM_ERS12' to 'datarep-sapmnt'.
Create dependency on 'ip-12-0-2-20' for 'avoid_SAP-EXM_ERS12' to 'datarep-sapmnt'.
creapphier: WARNING No quickCheck script specified
creapphier: WARNING No local recovery script specified
BEGIN create of "avoid_SAP-EXM_ASCS02"
creating resource "avoid_SAP-EXM_ASCS02"
resource "avoid_SAP-EXM_ASCS02" successfully created
restoring resource "avoid_SAP-EXM_ASCS02"
BEGIN restore of "avoid_SAP-EXM_ASCS02"
Attempting to avoid resource avoid_SAP-EXM_ERS12. Since the resource is not ISP on ip-12-0-1-20, restore of avoid_SAP-EXM_ASCS02 is successful on ip-12-0-1-20.
END successful restore of "avoid_SAP-EXM_ASCS02"
resource "avoid_SAP-EXM_ASCS02" restored
END successful create of "avoid_SAP-EXM_ASCS02"
Extending resource instances for avoid_SAP-EXM_ASCS02
Creating dependencies
Creating equivalencies
LifeKeeper Admin Lock (avoid_SAP-EXM_ASCS02) Released
Hierarchy successfully extended
BEGIN extend of "avoid_SAP-EXM_ASCS02"
END successful extend of "avoid_SAP-EXM_ASCS02"
Extending resource instances for avoid_SAP-EXM_ASCS02
Creating dependencies
Creating equivalencies
LifeKeeper Admin Lock (avoid_SAP-EXM_ASCS02) Released
Hierarchy successfully extended
BEGIN extend of "avoid_SAP-EXM_ASCS02"
END successful extend of "avoid_SAP-EXM_ASCS02"
Create dependency on 'ip-12-0-0-20' for 'avoid_SAP-EXM_ASCS02' to 'ec2-12.2.2.10'.
Create dependency on 'ip-12-0-1-20' for 'avoid_SAP-EXM_ASCS02' to 'ec2-12.2.2.10'.
Create dependency on 'ip-12-0-2-20' for 'avoid_SAP-EXM_ASCS02' to 'ec2-12.2.2.10'.
Create dependency on 'ip-12-0-0-20' for 'avoid_SAP-EXM_ASCS02' to 'datarep-ERS12'.
Create dependency on 'ip-12-0-1-20' for 'avoid_SAP-EXM_ASCS02' to 'datarep-ERS12'.
Create dependency on 'ip-12-0-2-20' for 'avoid_SAP-EXM_ASCS02' to 'datarep-ERS12'.
[root@ip-12-0-2-20 ~]#

```

The terminal leaf node has now been attached:



The avoid\_SAP-EXM\_ERS12 resource will not allow the SAP-EXM\_ASCS02 hierarchy to come in-service on a node if the avoid\_SAP-EXM\_ASCS02 resource (in the SAP-EXM\_ERS12 hierarchy) is in-service on that node and there is another viable node available in the cluster. A node is **NOT** a viable option when:

1. The node is not responding.
2. LifeKeeper is not running on the node.
3. A local recovery has failed on the node. This is determined by checking the output of `/opt/LifeKeeper/bin/flg_list` for the flag `!volatile!recover_fail_<tag>`.

The avoidance leaf can be disabled on a particular system by creating the flag:

1. "ignore\_avoidance\_leaf" – will disable the avoidance leaf checking for any resource, aka the avoidance leaf will come in-service at all times.
2. "ignore\_<tag>" – will disable the particular so that it will always come in-service but other avoidance leaves will still avoid in-service.

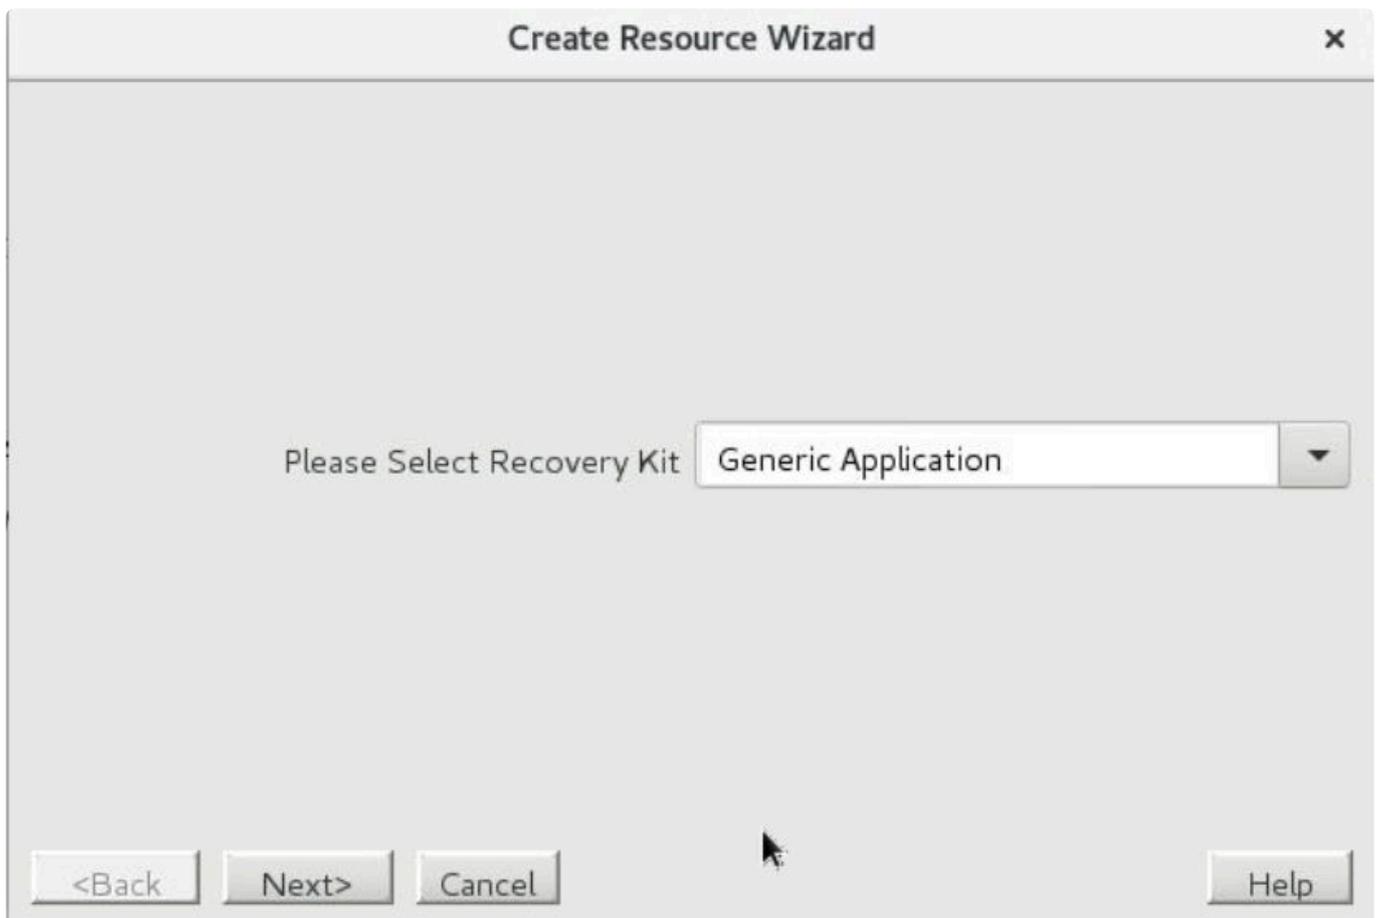
**NOTE:** The flags will not affect the SAP quickCheck from migrating ERS when it detects it is running on ASCS. The flag "sap\_no\_ers\_relocation\_<tag>" will disable quickCheck from relocating ERS where <tag> is the ERS resource tag.

## Creating Avoidance Terminal Leaf Node Using the GUI

The avoidance terminal leaf node is a gen/app resource that can be created using the GUI. The restore script for the avoidance terminal leaf is '/opt/LifeKeeper/lkadm/bin/avoid\_restore'. The remove script should be '/bin/true'. There is no quickCheck script. The info field should be the name of the tag to avoid.

For example, there are two resources, app1 and app2, that you want to be on different nodes when possible. You can create two gen/app resources, "avoid\_app1" and "avoid\_app2". The 'info' field for avoid\_app1 would have 'avoid\_app2'. The 'info' field for avoid\_app2 would have 'avoid\_app1'. The 'avoid\_app2' is a dependent child resource to 'app1' and 'avoid\_app1' is a child resource to 'app2'.

**Note:** The tag name is not required to be 'avoid\_<tag>' but this makes it clear what the resource is doing.



**Create gen/app Resource** ✕

Restore Script

Enter the pathname for the shell script or object program which starts the application. The **restore** script is responsible for bringing a protected application resource in-service. The **restore** script should not impact an active resource application when invoked.

Valid characters allowed in the script pathname are letters, digits, and the following special characters:  
- \_ ! . /

A copy of this script or program will be saved under:  
**/opt/LifeKeeper/subsys/gen/resources/app/actions**  
Whenever this resource is extended to a new server, the copy will be passed to that server.

**Create gen/app Resource** ✕

Remove Script

Enter the pathname for the shell script or object program which stops the application. The **remove** script is responsible for stopping a protected application resource and putting it in the out-of-service state.

Valid characters allowed in the script pathname are letters, digits, and the following special characters:  
- \_ ! . /

A copy of this script or program will be saved under:  
**/opt/LifeKeeper/subsys/gen/resources/app/actions**  
Whenever this resource is extended to a new server, the copy will be passed to that server.

### Create gen/app Resource ✕

QuickCheck Script [optional]

Enter the pathname for the shell script or object program which monitors the application. The **quickCheck** script is called periodically, and is responsible for performing a health check of the protected application.

The **quickCheck** script is optional. If one is not provided it will always be assumed that the application is in an OK state.

Valid characters allowed in the script pathname are letters, digits, and the following special characters:  
- \_ ! . /

A copy of this script or program will be saved under:  
**/opt/LifeKeeper/subsys/gen/resources/app/actions**  
Whenever this resource is extended to a new server, the copy will be passed to that server.

### Create gen/app Resource ✕

Application Info [optional]

Enter any optional data for the application resource instance that may be needed by the **restore** and **remove** scripts.

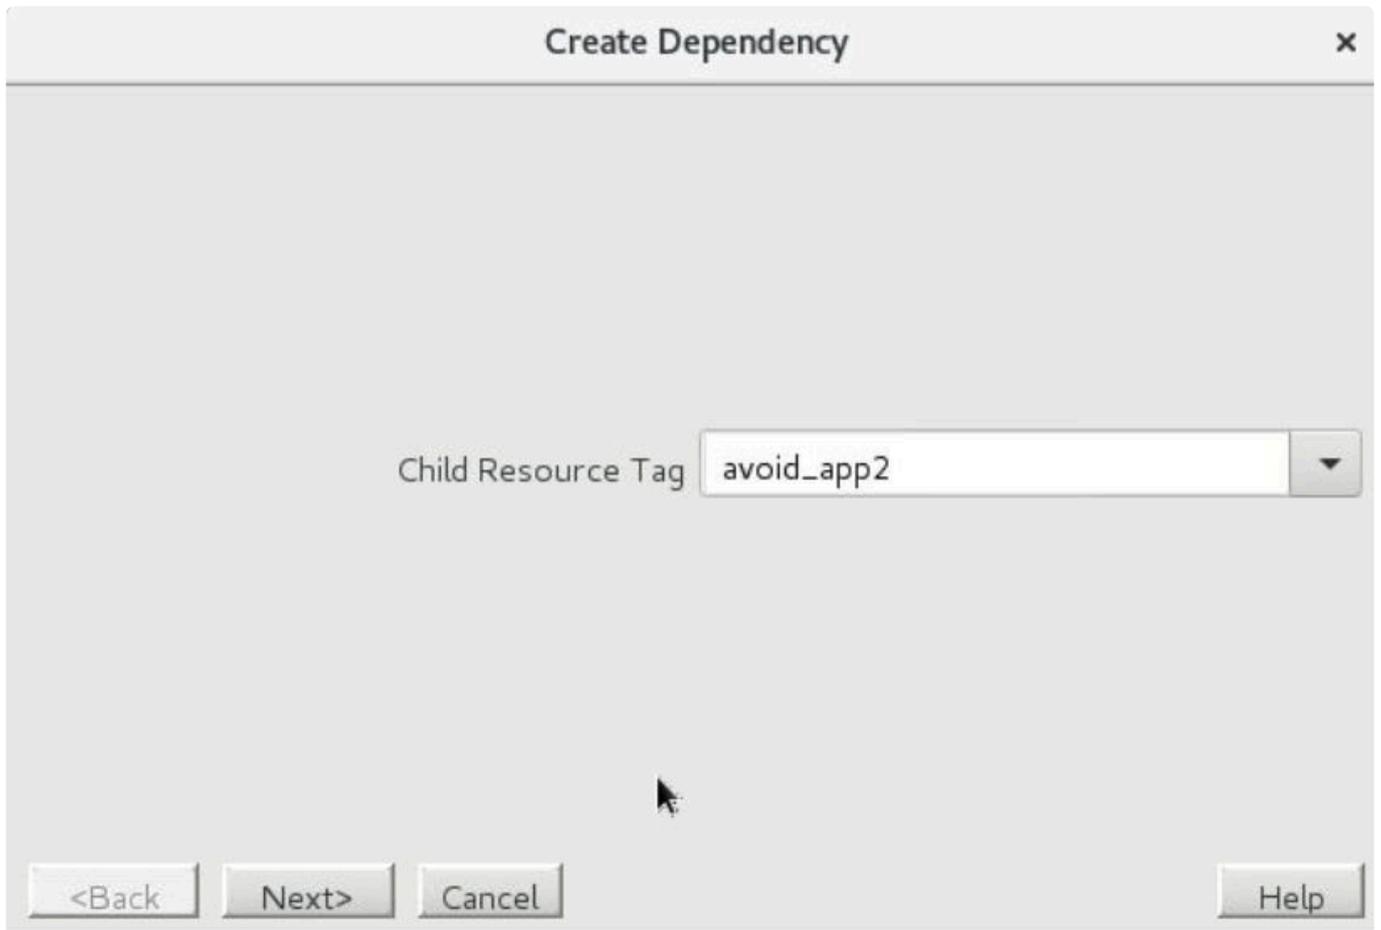
The valid characters allowed for the data field are letters, digits, and the following special characters:  
- \_ . / = **[space]**



After creating 'avoid\_app2' extend it to all nodes with the same priorities that app1 has.

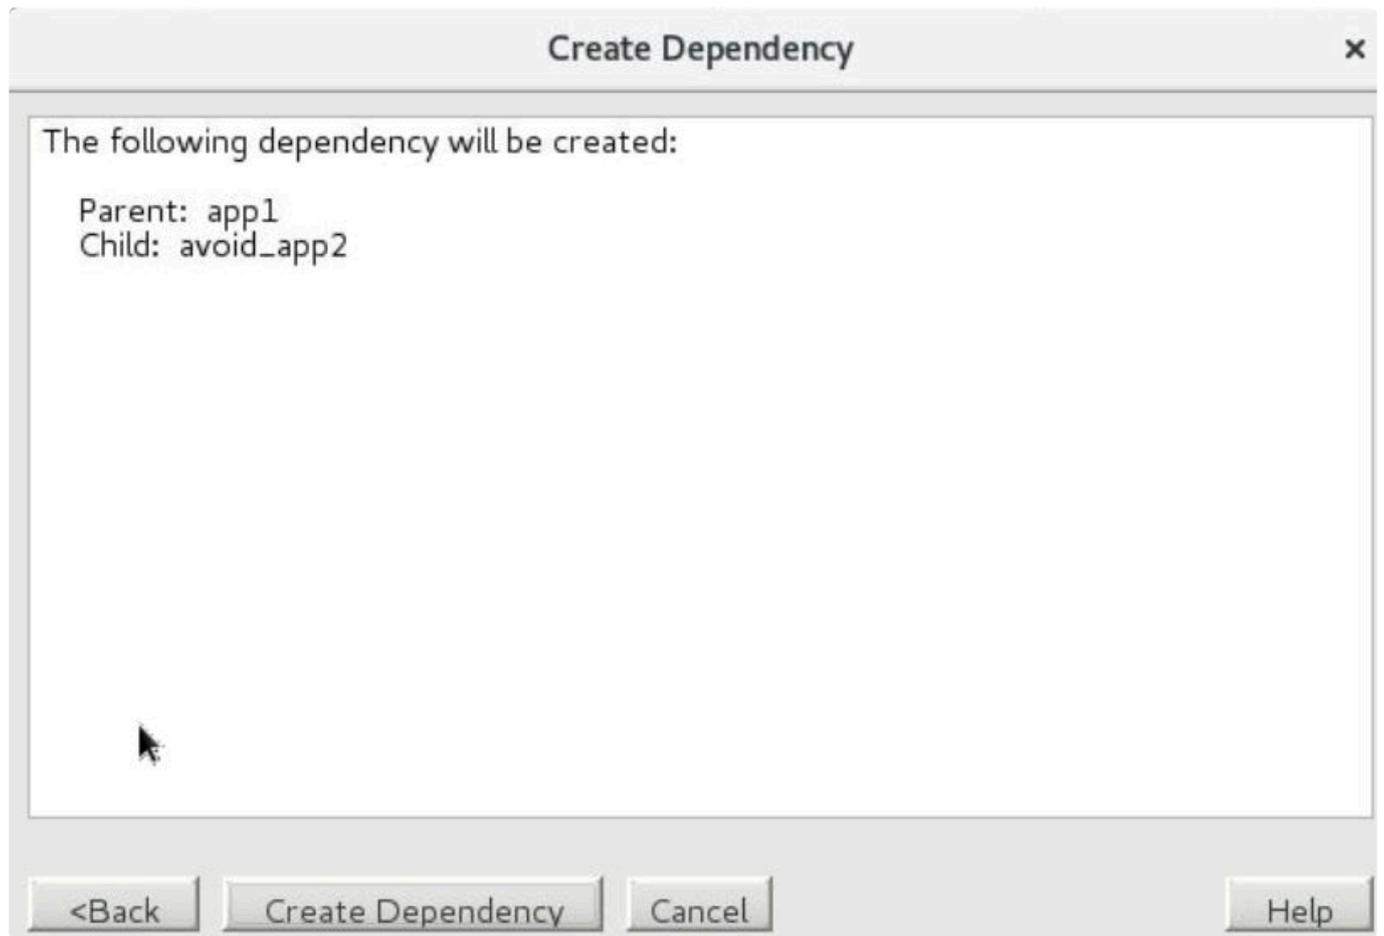


Then select the 'app1' resource and create a child dependency with 'avoid\_app2'.



After creating the avoid\_app1 resource similarly the hierarchy will look like:

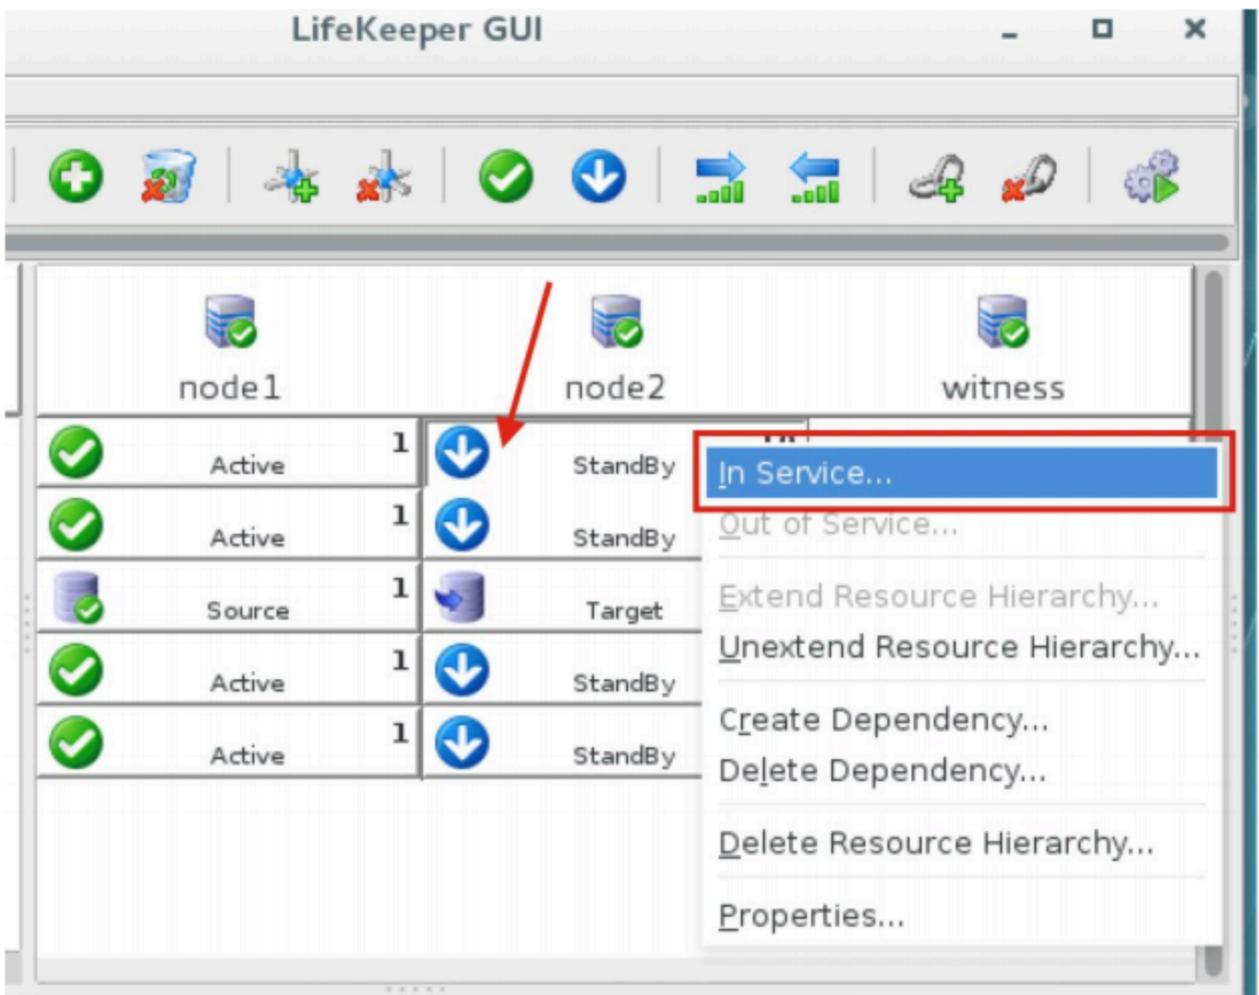




## 12.4.4. Switchover and Failover Testing

Steps below are for testing the switchover and failover of a SIOS cluster for SAP. Open “SAP Logon” or “SAP GUI for Windows”, which is an SAP supplied Windows client the Windows client. The program can be downloaded from the SAP download site. The virtual IP address may be used as the “Application Server” on the **Properties** page. This ensures that a connection to the primary machine where the virtual ip resides is active.

- Using the LifeKeeper GUI, failover from Node1 -> Node2. Right click on the top resource in the cluster underneath node2, and select “In Service...”. This demonstrates that node 2 is able to take over from node1 during a failure



After switchover has completed, check on the SAP GUI or reconnect where necessary and examine that SAP is still running normally.

You may also check if SAP processes are running in the OS.

- Using the LifeKeeper GUI, failover from Node2 -> Node1. Right click on the top resource in the cluster underneath node2, and select “In Service...”. This demonstrates that node 1 is able to take over from node2 during a failure

After switchover has completed, check on the SAP GUI or reconnect where necessary and examine that SAP is still running normally.

You may also check if SAP processes are running in the OS.

3. On the command line interface on node 1 (active node), issue a halt as provided by your system administrator to perform an immediate clean shutdown of the OS.

After failover has completed, use the LifeKeeper GUI on node 2 to check visually that the services are failed over normally.

Check on the SAP GUI or reconnect where necessary and examine that SAP is still running normally.

You may also check if SAP processes are running in the OS.

Turn on node 1 again and use the LifeKeeper GUI on node 2 to check visually that the services are on node 1 becomes standby, and replication is started.

**Note:** Before attempting to do any more switchover or failover testing, ensure that the data replication resources have already completed their synchronization and are in sync.

4. Repeat step 2 or step 1 as necessary to switchover back to node 1, or to perform another crash testing on node 2.

## 12.4.4.1. Additional Resources

---

### AWS services

- Amazon EC2  
<https://aws.amazon.com/documentation/ec2/>
- AWS CloudFormation  
<https://aws.amazon.com/documentation/cloudformation/>
- Amazon VPC  
<https://aws.amazon.com/documentation/vpc/>

### LifeKeeper for Linux

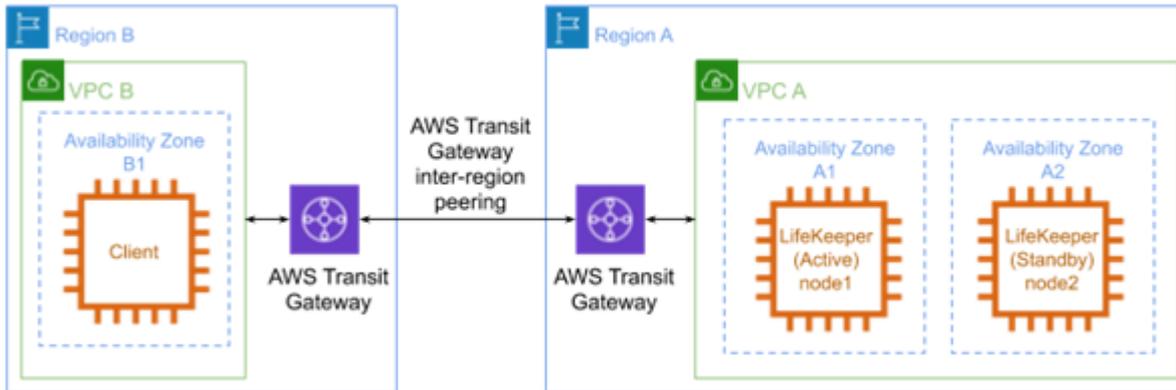
- Step-By-Step: How to configure a Linux failover cluster in Amazon EC2 without shared storage  
<http://www.linuxclustering.net/2016/03/21/step-by-step-how-to-configure-a-linux-failover-cluster-in-amazon-ec2-without-shared-storage-amazon-aws-sanless-cluster/>

### Quick Start reference deployments

- AWS Quick Start home page  
<https://aws.amazon.com/quickstart/>

## 12.5. Connecting to a LifeKeeper Cluster using AWS Transit Gateway Quick Start Guide

With the release of AWS Transit Gateway and AWS Transit Gateway inter-region peering, the Recovery Kit for EC2 route table scenario is now available for configurations where a client in a VPC (VPC B in the figure below) connects to an HA cluster located in a different region and VPC (VPC A in the figure below).



This document describes the requirements and basic operations for building a configuration where a client connects to a LifeKeeper for Linux HA cluster in another region.

This document does not explain the basic settings, operations or technical details of LifeKeeper or Amazon Web Service (AWS). For terms, operations and technical information related to LifeKeeper and AWS required for this configuration, review the related documents and user websites.

**Note:** AWS Transit Gateway inter-region peering is available only in the Eastern U.S. (N. Virginia, Ohio), Western U.S. (Oregon) and Europe (Ireland, Frankfurt) as of February 2020. If you deploy the server or client in another region, the configuration described in this document cannot be used. If you place your server or client in a region where AWS Transit Gateway inter-region peering is not available, consider using the Route53 Recovery Kit to update DNS A records (corresponding IP address to host names) registered in “Route53” of the AWS DNS service.

**Note:** This document is for configurations where cluster nodes are located within a single VPC. Route table scenarios cannot be used with configurations where cluster nodes are located across multiple regions or multiple VPCs.

**Note:** Amazon Web Services, Powered by Amazon Web Services logo, AWS, Amazon EC2, EC2, Amazon Elastic Compute Cloud, Amazon Virtual Private Cloud, AWS Direct Connect, AWS Identity and Access Management, AWS Transit Gateway, AWS Transit Gateway inter-region peering and Amazon VPC are trademarks of Amazon.com, Inc. or

its affiliates in the United States and other countries.

## 12.5.1. AWS VPC Peering Connections Requirements

---

The following requirements should be met when using this configuration. Below is a summary of requirements for the AWS environment and instances created on it.

### Requirements for AWS environment

Create a base environment on AWS to provide services. The requirements for using this configuration are as follows.

### Amazon Virtual Private Cloud (VPC)

- A VPC needs to be configured in AWS.
- The VPC where the client is located must be configured in a different region from the VPC where the cluster nodes are located.
- Create a subnet for the primary instance and a subnet for the standby instance in the VPC where the cluster nodes reside. The subnets must be created in different Availability Zones (AZ).
- The security groups for the subnets in the VPC containing the cluster nodes must be configured to allow incoming traffic from the subnet in the VPC containing the client, and vice-versa.

### Amazon Elastic Compute Cloud (EC2)

- At least 2 instances are required.
- A primary instance and a standby instance need to be configured in different AZs from each other.
- Cluster node instances are connected to an Elastic Network Interface (ENI).
- Cluster node instances must satisfy LifeKeeper installation requirements.
- The AWS Command Line Interface (AWS CLI) must be installed on all of the cluster node instances. Refer to [Installing the AWS CLI](#) for more details. The path to the AWS CLI executable files must be appended to the PATH parameter in the LifeKeeper defaults file `/etc/default/LifeKeeper` if it is not already present there.
- The cluster nodes need to be able to access the Amazon EC2 web service endpoint URL (see [https://docs.aws.amazon.com/general/latest/gr/rande.html#ec2\\_region](https://docs.aws.amazon.com/general/latest/gr/rande.html#ec2_region)) using https and the Amazon EC2 metadata URL (`http://169.254.169.254/`) using http.

### AWS Identity and Access Management (IAM)

In order for LifeKeeper to operate in AWS, an IAM user or IAM role with the following access privilege is required. Please configure [EC2 IAM role](#) or configure [AWS CLI](#) appropriately so that it can be accessed

by the root user of the EC2 instance.

- ec2:DisassociateAddress
- ec2:DescribeAddresses
- ec2:AssociateAddress
- ec2:DescribeRouteTables
- ec2:ReplaceRoute

## AWS Transit Gateway

- The VPC with cluster nodes and VPCs with clients should not be directly connected to each other with AWS Inter-Region VPC Peering. Instead, create an AWS Transit Gateway in each region and connect the AWS Transit Gateways with AWS Transit Gateway inter-region peering.
- Enable the default route table association and the default route table propagation when creating each AWS Transit Gateway.
- Create a Transit Gateway Attachment in each region to connect each AWS Transit Gateway to its corresponding VPC.
- An AWS Transit Gateway inter-region peering connection between AWS Transit Gateways should be enabled by creating a Transit Gateway Attachment. Note that this step requires manual confirmation in the target region before the Transit Gateway Attachment will actually be created by AWS.

## 12.5.1.1. LifeKeeper Software Requirements for AWS Environment

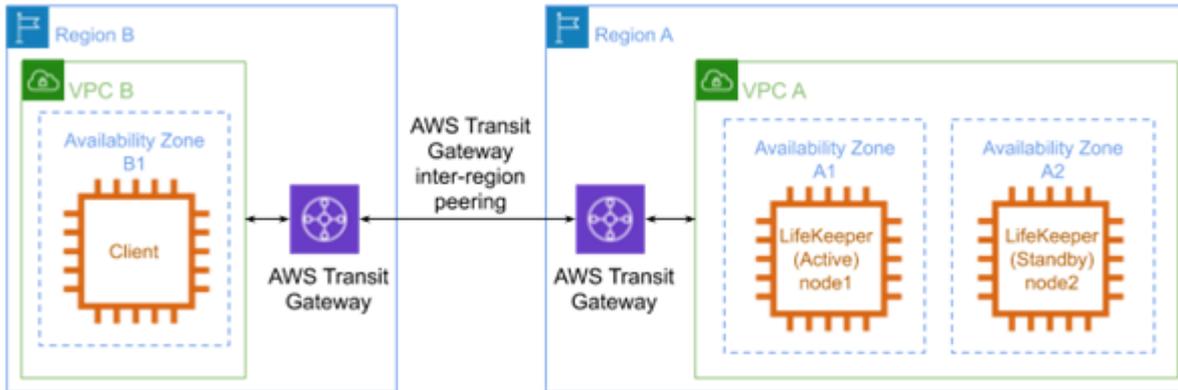
---

Install the same version of LifeKeeper software and patches on each server in the high-availability cluster. The Application Recovery Kits (ARK) required for this configuration are shown below. For the specific LifeKeeper requirements, please refer to: [LifeKeeper for Linux Technical Documentation](#) and [LifeKeeper for Linux Release Notes](#).

- [LifeKeeper IP Recovery Kit](#)
- [LifeKeeper EC2 Recovery Kit](#)

# 12.5.2. AWS VPC Peering Setup Procedure

This section describes the general procedure to set up the environment shown in the figure below.



## Preparations

- Create an environment that satisfies the [AWS VPC Peering Connections Requirements](#).
- Install LifeKeeper on each instance and create a communication path between node 1 and node 2.

## Creating an IP Resource

- Create a Virtual IP resource on node 1 and extend it to node 2. The IP resource address must be outside of the CIDR block managed by VPC A.

## Creating an EC2 Resource

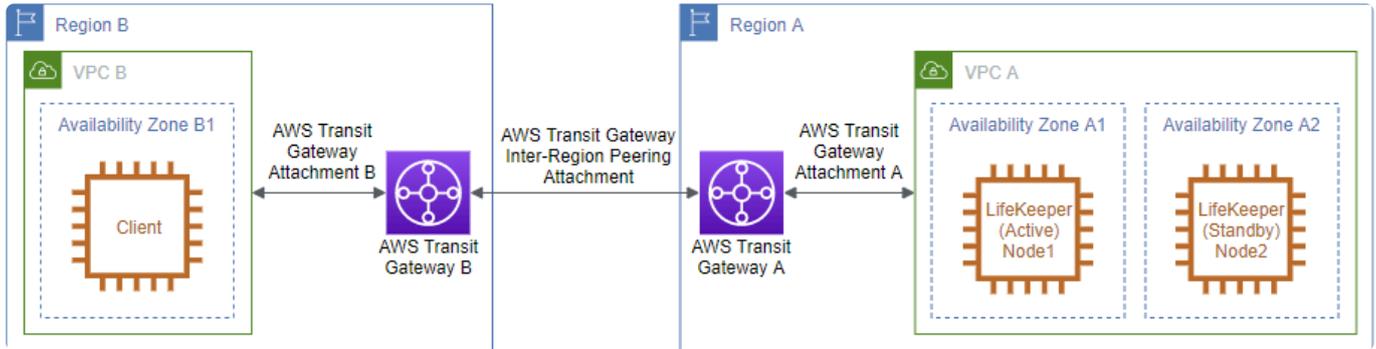
Create an EC2 resource on node 1 and extend it to node 2. For the IP resource requested when creating a resource, specify the resource created in “Creating an IP Resource”. Specify “Route Table (Backend Cluster)” for the EC2 resource type when creating the resource.

## Creating Resources for Protected Services

Create a resource for the service or application you want to protect. If an IP resource is required when creating a resource, specify the resource created in “Creating an IP Resource”. If necessary, configure resource dependencies so that the service/application is the parent resource and the EC2 resource is the child resource.

# 12.5.3. Configuring the Route Table

The AWS environment should be configured as in the following diagram:



Add the following routes to the route table for VPC B or the subnet that contains the client instance:

Destination Address	Target
VPC A CIDR Block	AWS Transit Gateway B
Virtual IP Address	AWS Transit Gateway B

Add the following routes to the route table for AWS Transit Gateway B:

CIDR	Choose Attachment
VPC A CIDR Block	AWS Transit Gateway Inter-Region Peering Attachment
Virtual IP address	AWS Transit Gateway Inter-Region Peering Attachment

Add the following routes to the route table for AWS Transit Gateway A:

CIDR	Choose Attachment
VPC B CIDR Block	AWS Transit Gateway Inter-Region Peering Attachment
Virtual IP address	AWS Transit Gateway Attachment A

Add the following route to the route table for VPC A or the subnets that contain the LifeKeeper instances:

Destination Address	Target
VPC B CIDR Block	AWS Transit Gateway A

Once configured, make sure that the client can access the private address of each server in the high-availability cluster as well as the virtual IP address.

## 12.5.4. Considerations for Settings and Operations in AWS VPC Peering

---

[Considering the Use of LifeKeeper I-O Fencing – AWS VPC Peering](#)

## 12.5.4.1. Considering the Use of LifeKeeper I-O Fencing – AWS VPC Peering

---

Since an AWS environment does not support shared disk configurations, SCSI reservations cannot be used to prevent split brain scenarios. For this reason, consider using the Quorum/Witness Server or STONITH, LifeKeeper's I/O fencing functionality, to operate more safely with this configuration.

Quorum fencing functionality can be easily configured in cloud environments by using the TCP\_REMOTE Quorum mode, instead of setting up a separate Quorum server. For details, please see the URL below

[Quorum/Witness](#)

[STONITH](#)

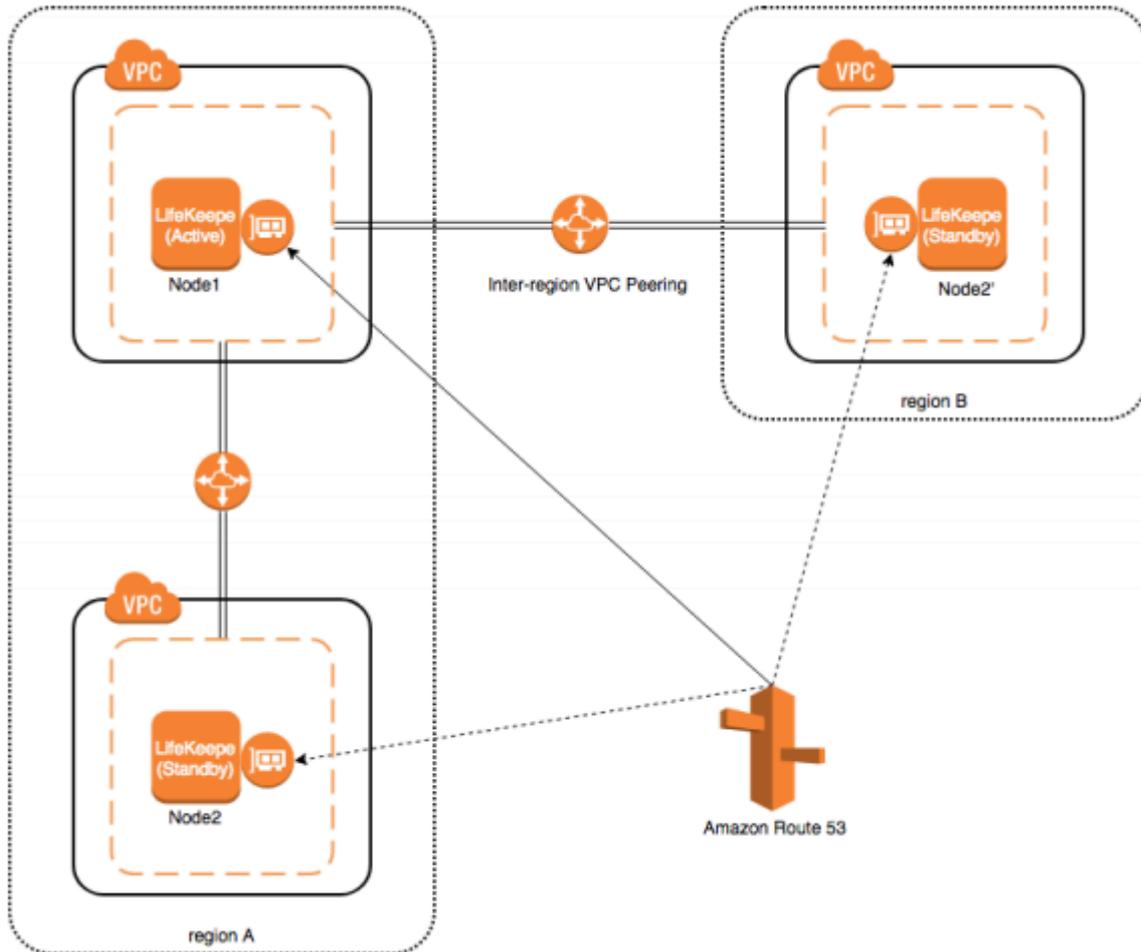
## 12.5.5. AWS Direct Connect Known Issues and Troubleshooting

---

There are currently no Known Issues.

## 12.6. Connecting to a LifeKeeper Cluster using AWS VPC Peering Quick Start Guide

### Objective



This document describes the requirements and basic operations for building connections among VPCs with LifeKeeper for Linux for v9.5.0.

You can also build HA clusters in the AWS environment using the existing Recovery Kit for EC 2; however, you cannot connect from your on-premises environment with AWS Direct Connect due to the problems described below.

Recovery Kit for EC2 provides two functions: “Route Tables Scenario” and “Elastic IP Scenario.”

“**Route Tables Scenario**” manages VPC route tables are configured to be routed to an active IP resources. An address of IP resource should be outside CIDR block which is managed within the VPC. However, the address should be the one within the VPC CIDR block in order to connect from other VPC via VPC Peering Connection. With this route table scenario, you cannot connect to the VPC from the on-premises environment.

“**Elastic IP Scenario**” can be used where the access from the Internet is available since the elastic IP address is a public address. An access from the on-premises environment is enabled through the Internet. In this case, you can access to HA cluster nodes on VPC without VPC Peering Connection.

For above reasons, Recovery kit for EC2 does not support an access to VPC from other VPC using VPC Peering Connection. If you need to access to HA cluster nodes on the VPC via VPC Peering Connection, please use the configuration provided in this document.

 It is also now possible to use AWS Transit Gateways for inter-region peering. Refer to [Connecting to a LifeKeeper Cluster using AWS Transit Gateway Quick Start Guide](#) for information.

Please note that this document does not describe the basic settings, operations, and technical details of LifeKeeper and Amazon Web Service (AWS). For terms, operations and technical information related to LifeKeeper and AWS, that are the prerequisites of this configuration, please read related documents and user websites beforehand.

 **Note:** “Amazon Web Services,” “Powered by Amazon Web Services” logo, “AWS,” “Amazon EC2,” “EC2,” “Amazon Elastic Compute Cloud,” “Amazon Virtual Private Cloud,” “Amazon Route 53” and “Amazon VPC” is trademarks of Amazon.com, Inc. or its affiliates in the United States and other countries.

## 12.6.1. Connecting to a LifeKeeper Cluster using AWS Requirements

---

Some requirements should be met when using this configuration. Below is a summary of requirements for the AWS environment and instances created on it.

### Requirements for AWS environment

Create a base environment on AWS to provide services. The requirements for using this configuration are as follows.

### Amazon Virtual Private Cloud (VPC)

- VPC needs to be configured in AWS.
- Need to create more than two subnets in different Availability Zones (AZ) or in different VPCs.

### Amazon Elastic Compute Cloud (EC2)

- At least 2 instances are required.
- A primary instance and a standby instance need to be configured to start with different AZ or different VPC for each.
- Instances are connected to Elastic Network Interface (ENI).
- Instances are required to satisfy LifeKeeper's installation requirements.
- AWS Command Line Interface (AWS CLI) needs to be installed in each of EC2 the instances. For the details, please refer to [AWS Command Line Interface Installation](#).
- Instances need to have an access to route53.amazonaws.com with HTTPS protocol. Please configure EC2 and the OS properly

### AWS Identity and Access Management (IAM)

In order for LifeKeeper to operate AWS, IAM user or IAM role with the following access privilege is required. Please configure [EC2 IAM role](#) or configure [AWS CLI](#) appropriately so that it can be accessed from root user of the EC2 instance.

- route53:GetChange
- route53:ListHostedZones
- route53:ChangeResourceRecordSets
- route53:ListResourceRecordSets

## Amazon Route 53

- You need to register your domain name on Amazon Route 53 to use the service. This is required to create a Route53 resource.

## 12.6.1.1. Peering Requirements for Connecting to a LifeKeeper Cluster using AWS

---

You need to install the same version of LifeKeeper software and patches on each server. The Application Recovery Kit (ARK) required for this configuration is shown below. For the specific LifeKeeper requirements, please refer to: [LifeKeeper for Linux Technical Documentation](#) and [LifeKeeper for Linux Release Notes](#).

- LifeKeeper IP Recovery Kit
- LifeKeeper Route53 Recovery Kit

## 12.6.1.2. Other AWS VPC Requirements

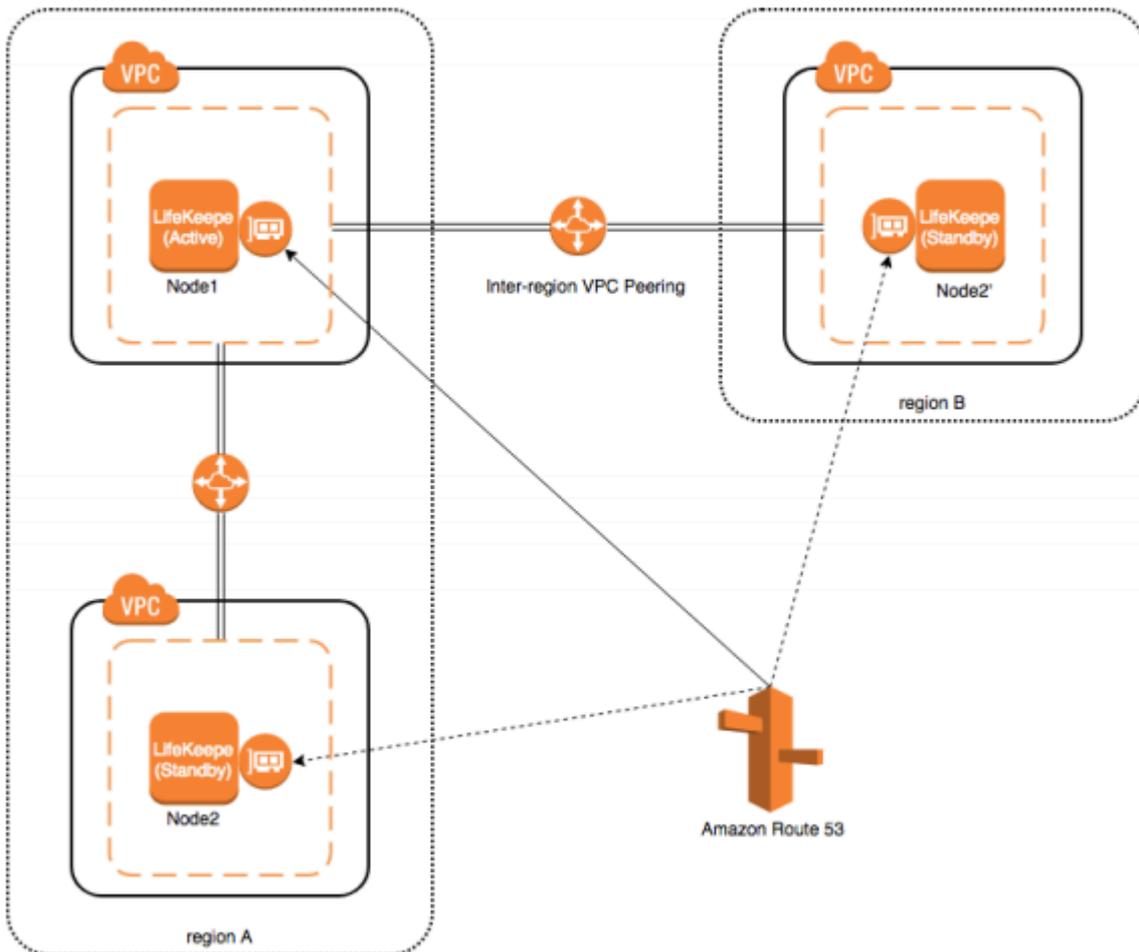
---

Requirements for using this service from other VPCs are as follows:

- Clients using the service should be able to resolve names of the hosts that are protected by Route53 resources.
- Clients using the service should access with the host name protected by Route53 resource.

## 12.6.2. Setup Procedure for Connecting to a LifeKeeper Cluster using AWS

In this section, a general procedure to setup the environment shown as the figure below



### Preparations

Create an environment that satisfies [Requirements](#). Please install LifeKeeper on each instance and create a communication path between Node1 and Node2 (or Node2'). Please confirm that you can access from other VPC environments to the ENI's real IP address connected to Node1/Node2 (or Node2').

### Creating IP Resource

Create an IP resource: not a virtual IP resource but a real IP resource (**Note:** resource for a primary IP address configured for Network Interface). Please specify IP address "0.0.0.0" when creating a resource. Also, specify IP address "0.0.0.0" for an extension target node when extending. Refer to the [IP Recovery Kit Administration Guide](#) for more information.

## Creating Route53 Resource

Create Route53 resource. Please specify the IP resource created in [Creating IP Resources](#) if required when creating Route53 resource.

## Creating Resources for Protected Services

Create resources for protected services. Please specify the IP resource created in [Creating IP Resources](#) if required when creating resources. Also, please create a resource dependency to enable the resources of the services protected by the parent resource and the child resource to become Route53 resources.

## 12.6.3. Related LifeKeeper Resources for AWS VPC Peering

---

### Route53 Resource

#### Summary

When switchover occurs, it is necessary to update Amazon Route 53 DNS information in order to continue to secure the connection to the service. This feature is provided in Route53 resources. When the status of Route53 resource becomes "In Service," the IP address of the IP resource with a dependency is registered in the corresponding DNS A record using API.

### IP Resource

#### Summary

IP resource is a resource generated with using IP Recovery Kit included in the LifeKeeper Core product. In order to support this configuration, it is now possible to generate IP resource (real IP resource) with a real IP address. This allows you to use real IP addresses as a LifeKeeper resource.

Please do not use the real IP resource except for this configuration.

For more information, please refer to: [IP Recovery Kit Technical Documentation](#)

## 12.6.4. Connecting to a LifeKeeper Cluster using AWS Settings and Operations Considerations

---

Please see the following topics for additional considerations when connecting to an AWS cluster using VPC peering and AWS Route53:

[Considering the Use of LifeKeeper I-O Fencing when Connecting to a LifeKeeper Cluster using AWS](#)

[Connecting to LifeKeeper Using AWS Requirements](#)

## 12.6.4.1. Considering the Use of LifeKeeper I-O Fencing when Connecting to a LifeKeeper Cluster using AWS

---

Since the shared disk environment cannot be used in AWS environment, you cannot use SCSI reservations to prevent a split-brain. Also, IP resource may cause the split-brain as it uses the real IP resource with different IP addresses for each node.

For this reason, please consider the use of Quorum/Witness server or STONITH, an I/O fencing function of LifeKeeper to use this configuration safely.

Especially, because you can implement I/O fencing function separately without the Quorum server if you use the TCP\_REMOTE setting in Quorum mode, it is easy to be implemented in the cloud environment. For more details, please refer to the following URLs:

[Quorum/Witness](#)

[STONITH](#)

# 12.7. PostgreSQL Cluster with Shared Storage (ISCSI)

---

## Objective

This document is intended to aid you in installing, configuring and using the LifeKeeper for Linux evaluation product, to make PostgreSQL highly available. If PostgreSQL is not already installed, please allocate some time to install it on your servers. Once this task has been completed, you may install and configure LifeKeeper for Linux.

There are five phases in this process:

- Phase 1 – Prepare to Install
- Phase 2 – Configure Storage
- Phase 3 – Install and Configure PostgreSQL
- Phase 4 – Install LifeKeeper for Linux
- Phase 5 – Configure your LifeKeeper Cluster
- Phase 6 – Test Your Environment

## 12.7.1. Terms to Know – PostgreSQL

---

The following terms are used throughout this document and, while some may be familiar to you, it may be helpful to review how SIOS defines and uses these terms.

### Network Communication Terms

**Crossover cable** – A cable used to directly connect computing devices together, instead of being connected to a network switch, hub or router. This cable creates an isolated, private network to allow cluster-related and data replication traffic to flow between systems.

### Types of LifeKeeper Servers

**Server** – A computer system dedicated to running software application programs.

**Active Server** – This is the server where the resource hierarchy is currently running (IN SERVICE).

**Standby Server** – This is the server where the resource hierarchy is defined, but is not currently running. This server is available to bring the resource hierarchy into service should something happen to the resource hierarchy on the Active Server.

**Primary Server** – This is the server in a LifeKeeper configuration with the highest priority for a given resource hierarchy. It is the server that provides services for the resource hierarchy under normal circumstances.

**Secondary Server** – This is the server in a LifeKeeper configuration with the 2nd highest priority for a given resource hierarchy.

**Source Server** – In a LifeKeeper cluster, using data replication, this is the Active Server. It is where the resource hierarchy is currently running and the replicated partition (Source Partition) is accessible for writes.

**Target Server** – In a LifeKeeper cluster, using data replication, this is the Standby Server. The replicated partition (Target Partition) is updated with writes from the Source Partition by the SIOS Data Replication system. This partition should not be accessed/modified manually.

### SIOS Data Replication Terms

**Replication** – Transferring data from one partition to another via a sector-by-sector copy. During replication, the target partition should not be accessed or modified assuring your data integrity.

**Synchronous** – A replication scheme in which the data is confirmed written and valid on the target before the write operation occurs on the source disk through a series of information exchanges. Synchronous mirrors should only be implemented on high speed (100Mbps+) networks due to the network overhead involved.

**Asynchronous** – A replication scheme in which the data is released for writing on the source immediately and is sent to the target(s) simultaneously for writing as fast as the data can get there and can be written on them.

**Rate of Change** – A measure of the amount of data which is changing over a set period of time.

**Compression** – An algorithm which is optionally implemented to reduce the amount of traffic between source and target nodes. Nine levels of compression are offered. Compression is turned off by default.

**Throttling** – An optionally implemented mechanism to limit the bandwidth used for replication.

## LifeKeeper Product Terms

**Communications Path** – A mechanism supporting communication between nodes in a LifeKeeper cluster. SIOS highly recommends implementing multiple communication paths between all servers in the cluster to eliminate a single point of failure.

**Heartbeat** – A periodic message exchanged between nodes in a LifeKeeper cluster that provides server health monitoring. A heartbeat message is one type of inter-node cluster communication sent over a communications path.

**Split Brain** – A situation in which all communications paths between cluster members fail, but all servers remain up and running. In this situation, both systems believe the other has failed and both believe they should keep or bring resources into service.

**Failover** – The unplanned migration of a resource hierarchy to the Standby Server because of a system or resource failure on the Active Server.

**Switchover** – The planned migration of a resource hierarchy from the Active Server to the Standby Server.

**Switchback** – The setting that governs the recovery behavior of the server where the resource was in service when it failed. If the setting is “intelligent”, the server acts as a possible backup for the given resource. If the setting is “Automatic”, the server actively attempts to re-acquire the resource without further notice.

**Resource** – A system asset that can be protected by LifeKeeper. Resources can be used to represent disk partitions, virtual IP addresses, applications, etc.

**Extend a Resource** – Create or define an already configured LifeKeeper resource onto another server in the cluster and build an equivalency relationship that prevents the resource from coming in service on both systems simultaneously.

**Resource Hierarchy** – A grouping of resources, in a predetermined order, from high to low. This may also be referred to as simply a Hierarchy.

**Shared Storage** – One or more logical disk partitions that are physically attached to all nodes in a cluster. LifeKeeper ensures that the volume is only accessible by one server at a time. This is formally

called I/O fencing.

**Data Replication (Disk Mirroring)** – The replication of logical disk partitions to separate physical hard disks in real time to ensure continuous availability, currency and accuracy of data.

**Source** – The partition on the source server used for replication. The “gold” copy of the data.

**Target** – The partition on the target server used for replication.

**Switchable IP Address** – A unique IP address that may be moved between systems in the cluster. Client systems connect to this address and the system where the virtual IP resource is active will respond to requests.

## 12.7.2. The Evaluation Process – PostgreSQL

---

SIOS strongly recommends performing your evaluation of LifeKeeper for Linux in a test lab environment. SIOS is not responsible and cannot provide support for evaluation software installed in a production environment.

All questions during the evaluation period should be directed to [evalsupport@us.sios.com](mailto:evalsupport@us.sios.com) or your local Pre-Sales Engineering contact. Pre-sales support will contact you by the next business day to answer questions. Once you are a licensed customer on software maintenance, you will have access to 24 × 7 post-sales technical support.

 **Important** Your evaluation license is valid for a limited period of time from the day you receive the SIOS product evaluation package and licenses via email from the SIOS sales team.

## 12.7.3. Prepare to Install – PostgreSQL

### Hardware Requirements

#### Primary and Secondary Servers

- Systems must meet the minimum requirements for the Linux distribution to be used during the evaluation.
- 512MB RAM minimum; 1GB RAM recommended.
- 2GB of available hard disk space recommended.
- Multiple Network Interface Cards (NIC's) are recommended.
- Configure one or more additional disk partitions to be used for data replication. On the primary server, these will become the source partitions. On the secondary server(s), these will become the target partitions. In this evaluation example we will be replicating MySQL data (which will be a partition mounted at /var/lib/mysql)
- For replicated partitions, a target partition's size must equal to or larger than the size of its source partition.
- The system ( / ) and boot (/boot) partitions are not eligible for replication.

 **Note:** You may use more than one partition for replicated data, allowing for separation of multiple database files and/or log directories for performance reasons.

#### Client

This system is not required but is recommended for testing the cluster environment.

- A standard linux terminal running the MySQL client can be used to test the configuration.

### Software Requirements

#### Primary Server and Secondary Server

- Linux Distribution x86\_64, AMD 64:
  - RedHat Enterprise Linux 5 (5.4+ recommended) or 6.x
  - CentOS Linux 5 (5.4+ recommended) or 6.x
  - Oracle Enterprise Linux 5 (5.4+ recommended), 6.3, 6.4
    - RedHat Compatibility Kernel Only
  - SuSE Linux Enterprise Server 10 or 11 (11 recommended)
  - See [Linux Release Notes](#) for a full list of supported Operating Systems
- Current patches / security updates are recommended.
- Satisfied dependencies; especially if the Linux installation package selection was base/minimal you will need to refer to the dependencies documentation at [Linux Dependencies](#)

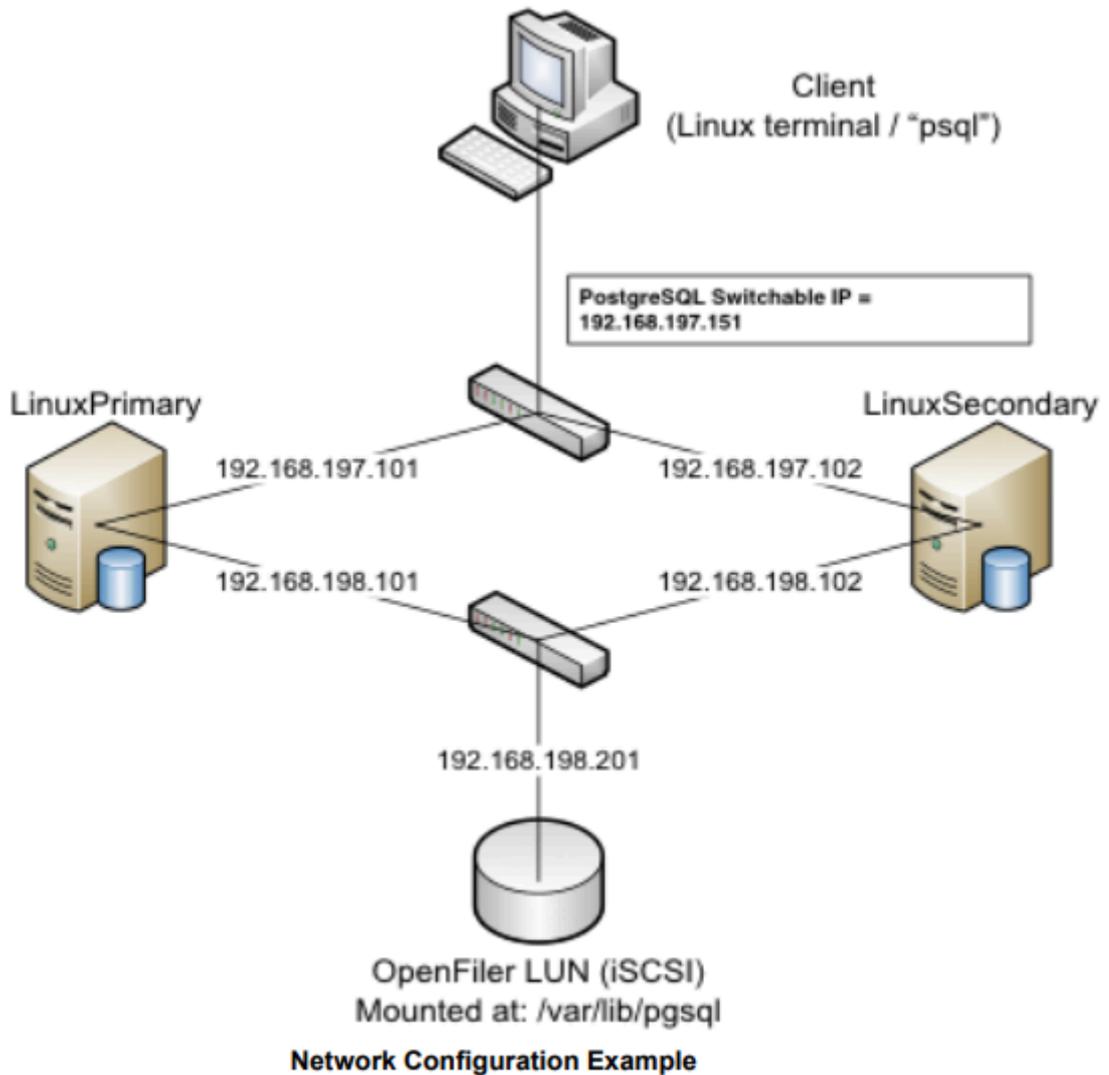
- Its recommended that IPtables is disabled
  - # /etc/init.d/iptables off
  - # chkconfig iptables off
  - See [here](#) for information regarding the ports LifeKeeper for Linux uses.
  
- Disable SELinux :
  - Edit /etc/selinux/config
  - Set SELINUX=disabled (note: permissive mode is also acceptable)
  
- Check the configuration of your /etc/hosts file
  - localhost.localdomain and localhost are the only entries that can be on 127.0.0.1
  - Create a separate entry for your hostname with a static address
  
- GUI Authentication with PAM
  - LifeKeeper for Linux now leverages the Pluggable Authentication Module (PAM) provided in the Linux Standard Base (LSB).
  - Users are identified and authenticated against the system's PAM configuration. Privilege levels are determined from group membership as provided through PAM.
  - In order to access the GUI, a user must be a member in one of the three LifeKeeper groups: **lkadmin, lkoper or lkguest**.
  - See [Configuring GUI Users](#) for more information.

## Network Requirements

For your evaluation, we recommend configuring your machines similarly to the following example.

LinuxPrimary and LinuxSecondary are multi- in production environments to avoid a single point of failure.

In this example evaluation scenario we will be leveraging Shared (iSCSI) Storage with our PostgreSQL database. OpenFiler is a storage appliance server that will serve an iSCSI target to LinuxPrimary and LinuxSecondary.



## Primary Server and Secondary Servers

- Configure the Host file with entries for all LifeKeeper protected servers. This is typically /etc/hosts.

Example:

```
192.168.197.101 LinuxPrimary
```

```
192.168.197.102 LinuxSecondary
```

- See your Network Administrator to obtain an unused IP Address to be used as the switchable IP Address. This switchable IP Address will be created later in the configuration process.
- Public Network connection(s) configured with:
  - Static IP address
  - Correct subnet mask
  - Correct gateway address

- Correct DNS server address(es)
  
- Private Network connection(s) configured with:
  - Static IP address (on a different subnet from the public network)
  
  - Correct network mask
  
  - No gateway IP address
  
  - No DNS server addresses

## Client

Must be able to communicate on the same subnet/network as the servers Public interface addresses. In our example, this is the 192.168.197.0/24 network.

## 12.7.4. Configure Storage – PostgreSQL

---

### Before you Begin

Ensure the following:

- If planning to use replicated storage, have an extra volume/partition on both servers that can be used for data replication. A target volume's size must equal to or larger than the size of its source volume.
- If planning to use shared storage, as in this example, ensure the Shared storage is configured and accessible to your cluster nodes. This can either be Fiber Channel SAN, iSCSI, NAS, etc. In this example we will review configuration of an iSCSI target for use as our PostgreSQL database storage repository.

### Configure iSCSI initiator, discover and login to iSCSI target

This Evaluation guide will not cover how to setup an iSCSI Target Server. It is assumed that the shared storage already exists in your environment. If you don't have shared storage and wish to configure it, a simple solution is to use OpenFiler (<http://www.openfiler.com/>), an Open Source storage management appliance, which can be run on physical hardware or as a virtual machine.

On both Primary and Secondary servers, perform the following functions:

1. If not already installed, ensure that the **iscsi-initiator-utils** rpm package is installed:

```
# yum install iscsi-initiator-utils
```

2. Start the iscsid service and enable it to automatically start when the system boots

```
# service iscsid start
```

```
# chkconfig iscsid on
```

3. Configure the iscsi service to automatically start, which logs into iSCSI targets needed at system start up.

```
# chkconfig iscsi on
```

4. Use the iscsiadm command to discover all available targets on the network storage server (OpenFiler)

```
# iscsiadm -m discovery -t sendtargets -p <name or IP of iSCSI server>
```

#### Example

```
[root@LinuxPrimary init.d]# iscsiadm -m discovery -t sendtargets -p 192.168.198.201  
iqn.2006-01.com.openfiler:tsn.postgres
```

## 5. Manually Login to the iSCSI Target

```
# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.postgres -p 192.168.198.201 -- login
```

## 6. Configure Automatic Login

```
# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.postgres -p 192.168.198.201 --op  
update -n node.startup -v automatic
```

## 7. Use the “gdisk” command to format your iSCSI LUN, if needed

```
# gdisk /dev/sdc
```

## 8. Create a filesystem on your new iSCSI LUN Partition, sdc1

```
# mkfs.ext3 /dev/sdc1
```

## 9. Mount your iSCSI LUN at /var/lib/pgsql (assuming a default postgres configuration). If data already exists in this directory, make sure to move it into the shared iSCSI LUN

```
# mount mount /dev/sdc1 /var/lib/pgsql
```

## 10. At this point you now have an iSCSI shared LUN, /dev/sdc1, mounted at /var/lib/pgsql. Our disk layout now look as follows (example):

### Example

```
[root@LinuxPrimary postgresql]# df  
Filesystem 1K-blocks Used Available Use% Mounted on  
/dev/sda2 25967432 3683016 1976400 66% /  
/dev/sda1 101086 24659 71208 26% /boot  
tmpfs 517552 0 517552 0% /dev/shm  
/dev/sdc1 966644 38944 878596 5% /var/lib/pgsql
```

## 12.7.5. Install, Configure and Start PostgreSQL

---

### Primary Server

On your Primary server, perform the following actions:

1. Install both the “postgresql-server” and “postgresql” rpm packages if they do not exist on your system. Apply any required dependencies as well

```
# yum install postgresql postgresql-server
```

2. Verify that your Shared iSCSI LUN is still mounted at /var/lib/pgsql via the “df” command

3. If this is a fresh PostgreSQL install, initialize a sample PostgreSQL database:

```
# su – postgres
```

```
# initdb --pgdata=/var/lib/pgsql/data
```

4. Ensure that all files in your PostgreSQL data directory (/var/lib/pgsql) have correct permissions and ownership

```
# chown -R postgres:postgres /var/lib/pgsql
```

```
# chmod 755 /var/lib/pgsql
```

5. Finally, manually start the PostgreSQL daemon from the command line. Note: **Do Not** start it via the “service” command, or the /etc/init.d/ scripts

```
# su – postgres
```

```
# pg_ctl start -D /var/lib/pgsql/data -l /var/lib/pgsql/pgstartup.log -o “-p 5432” -w
```

6. Verify PostgreSQL is running by connecting with the psql client (ensure you are still running as the “postgres” linux user):

```
-bash-3.2$ psql
```

```
Welcome to psql 8.1.22, the PostgreSQL interactive terminal.
```

```
Type: \copyright for distribution terms
```

```
\h for help with SQL commands
```

```
\? for help with psql commands
```

\g or terminate with semicolon to execute query

\q to quit

```
postgres=# \q
```

```
-bash-3.2$
```

## Secondary Server

On your Secondary Server:

1. install both the “postgresql” and “postgresql-server” rpm packages if they do not exist on your system. Apply any required dependencies as well

```
# yum install postgresql postgresql-server
```

2. Ensure that the PostgreSQL data directory (/var/lib/pgsql) has correct permissions and ownership

```
# chown -R postgres:postgres /var/lib/pgsql
```

```
# chmod 755 /var/lib/pgsql
```

3. There is no need to perform any of the additional steps taken on the Primary Server

## 12.7.6. Install LifeKeeper for Linux – PostgreSQL

---

For ease of installation, SIOS has provided the LifeKeeper for Linux with an installation script. Towards the end of the script, the desired Application Recovery Kits (ARKs) should be selected for installation. The software will be installed to the following locations:

SPS for Linux Component	Install Location
LifeKeeper Software	/opt/LifeKeeper
LifeKeeper Config File	/etc/default/LifeKeeper

Perform the following actions on **both** Primary and Secondary server.

### Download Software

1. Open the LifeKeeper evaluation email you received from SIOS.
2. Download the LifeKeeper for Linux Software from the link provided in your email. It is generally easiest to use “wget” to recursively download all files. Example:

a. # cd /root

b. # wget -r <URL>

c. After successful download you will have downloaded contents similar to the follow directory listing:

```
[root@LinuxPrimary ~]# ls -l <directory> total 63680
-rw-r--r-- 1 root root 23163 May 30 14:03 EULA.pdf
-rw-r--r-- 1 root root 536 May 30 14:03 readme.txt
-rw-r--r-- 1 root root 65179648 May 30 14:03 sps.img
```

3. Download your Evaluation license key from the link specified in your evaluation email. Save the license file to an easy to remember location on both servers.

### Run the LifeKeeper Installer Script

1. Loopback mount the sps.img file previously downloaded, which is an ISO9660 image file. Run the “setup” script inside:

```
# mount -o loop sps.img /mnt
```

```
# cd /mnt
```

```
# ./setup
```

2. During this procedure, you will hit Enter in most cases to accept default values and continue to the next screen. Note the following exceptions:
  - a. On the screen titled “High AvailabilityNFS” you may select “n” as in this particular eval guide we will not be creating a highly available NFS server cluster configuration.
  - b. If you have plans to create a highly available NFS service, adjust your response accordingly.
3. Towards the end of the setup script, you can choose to install a trial license key now, or later. We will install the license key in the next step, so you can safely select “n” at this point
4. In the final screen of the “setup” select the DataKeeper from the list displayed on the screen.
5. The following RPMs should be installed:
  - a. steeleye-lkDR-<version>.noarch.rpm
6. Un-mount the Distribution Enabling disk image:

```
# cd /root
```

```
# umount /mnt
```

## Install the Evaluation License Keys

The last phase of the setup process installs the licensing keys. You must install the evaluation license key file (“.lic”) that you downloaded with your evaluation software before starting the LifeKeeper for Linux.

1. To install your trial license key, run the “lkkeyins” command on both Primary and Secondary Server. This command is located at /opt/LifeKeeper/bin/lkkeyins . Example:

```
# /opt/LifeKeeper/bin/lkkeyins <path_to_license/<filename>.lic
```

2. Validate your license keys were installed via the /opt/LifeKeeper/bin/lklicmgr command

```
# /opt/LifeKeeper/bin/lklicmgr
```

License File: 20101230.lic

Product	Type	Expiry
LifeKeeper for Linux	Eval	27 Mar 2013 (87 days)

SIOS Data Replication ARK	Eval	27 Mar 2013 (87 days)
PostgreSQL Recovery Kit	Eval	27 Mar 2013 (87 days)

## Start the LifeKeeper for Linux

1. Start:

```
# /opt/LifeKeeper/bin/lkstart
```

## 12.7.7. Configure the Cluster – PostgreSQL

### Primary Server

Complete the following steps on the primary server to configure the cluster:

- Create TCP Communication (Comm) Path(s)
- Verify the Communication (Comm) Path(s)

Before you begin, SIOS recommends at least two TCP communications paths between each server within the cluster to each remote server for heartbeat redundancy.

 **Important** Supported configurations require that you define redundant comm. paths, so that the failure of a single communication line will not cause a split brain where resource hierarchies may come in-service on multiple servers simultaneously.

### Access the LikeKeeper GUI

The LifeKeeper Graphical User Interface (GUI) is a Java based application that can be run as a native Linux application.

The LifeKeeper GUI is based on Java RMI with callbacks. Hostnames must be resolvable or you may receive a Java 115 or 116 error.

1. Verify that both short and fully qualified hostnames of all cluster nodes resolve to the proper locations

```
# ping LinuxPrimary
```

```
# ping LinuxPrimary.domain.com
```

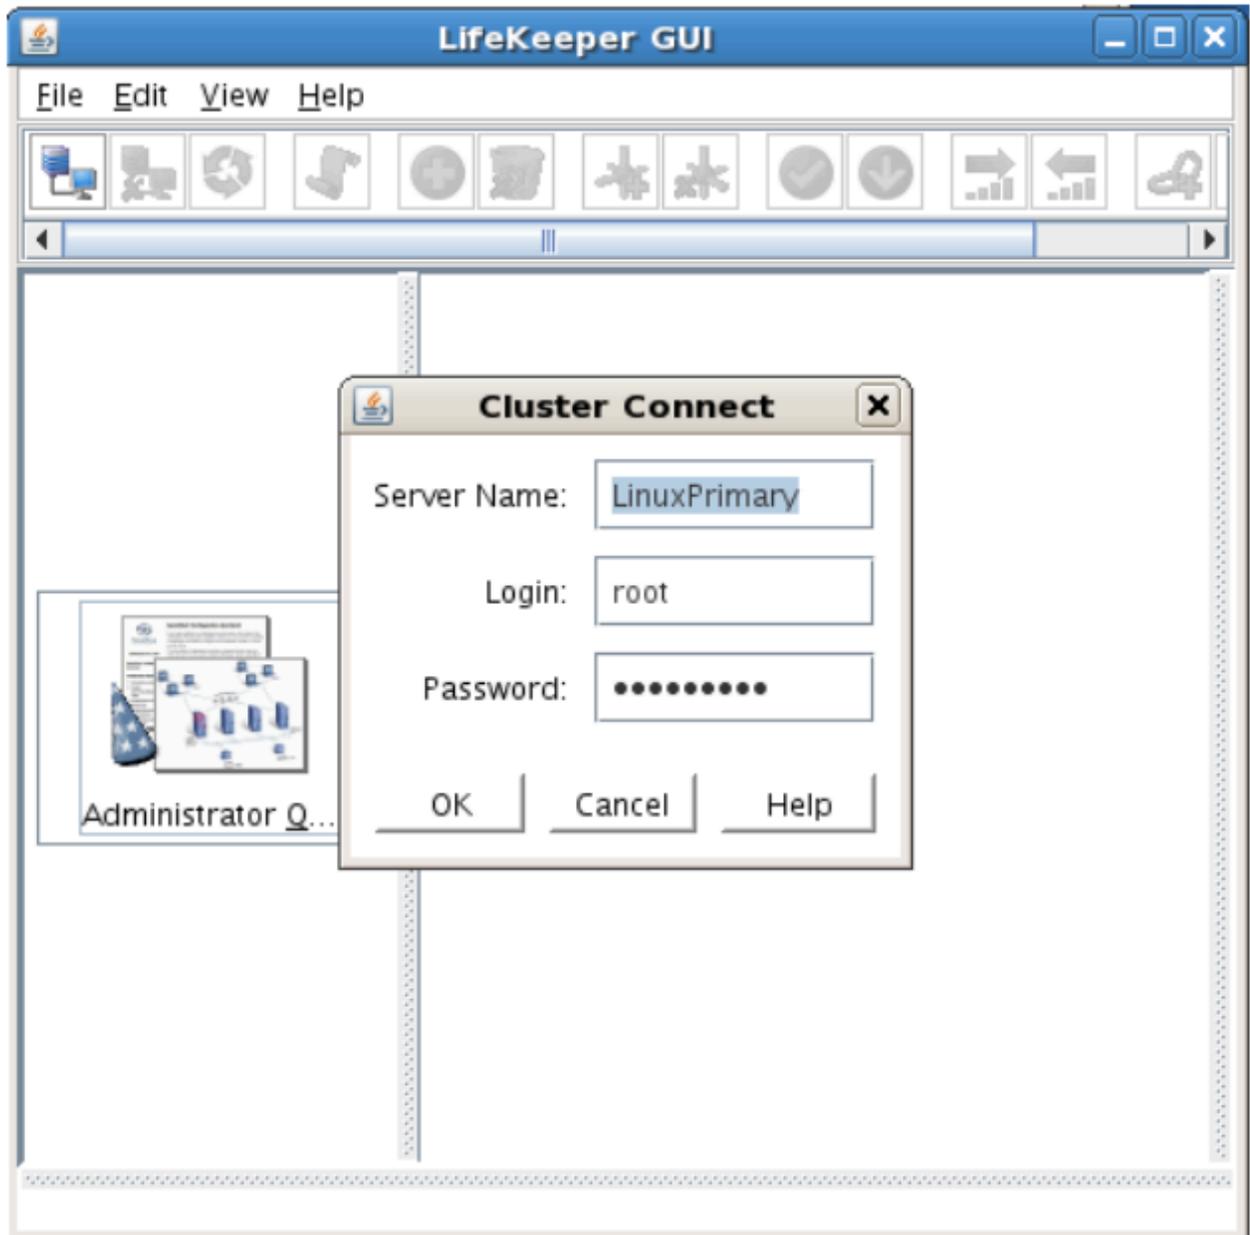
```
# ping LinuxSecondary
```

```
# ping LinuxSecondary.domain.com
```

2. To start the LifeKeeper Linux GUI Application:

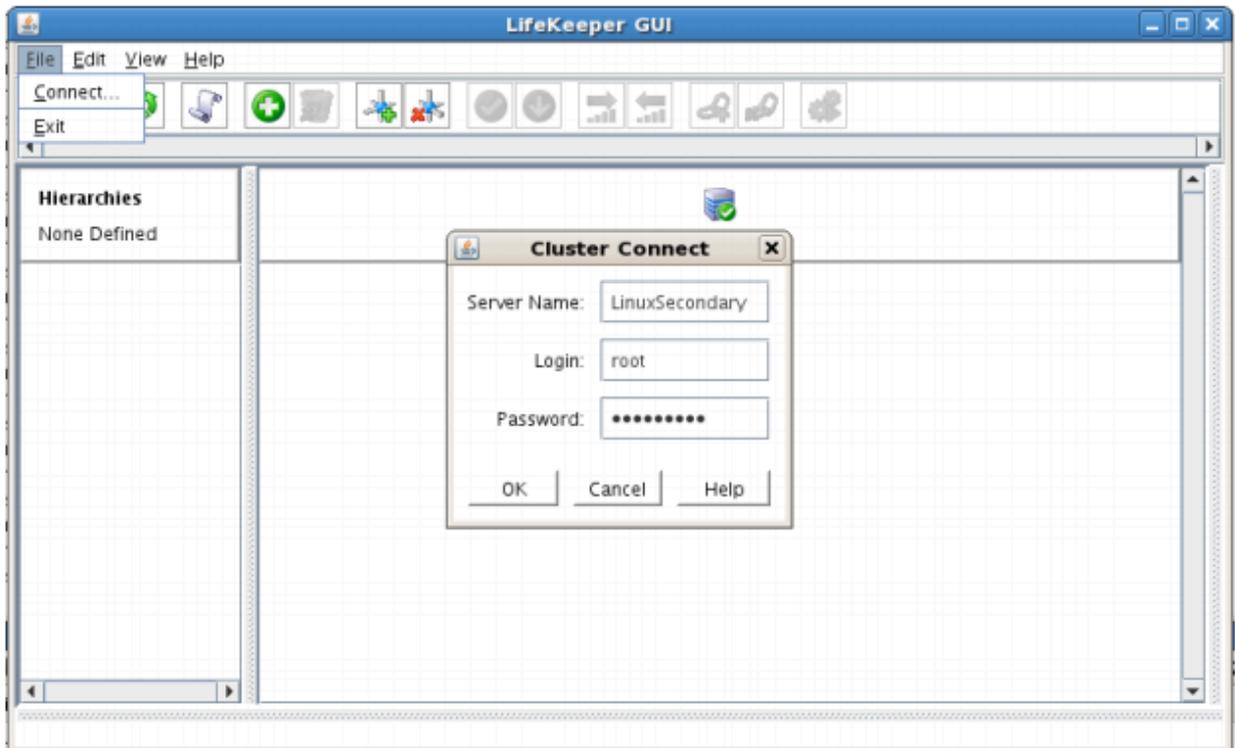
```
a. /opt/LifeKeeper/bin/lkGUlapp &
```

3. Enter the name of the server you wish to connect to (this field will be populated with the name of the server you are on, if you are running the GUI from a server with LifeKeeper installed) along with your root credentials and click OK.

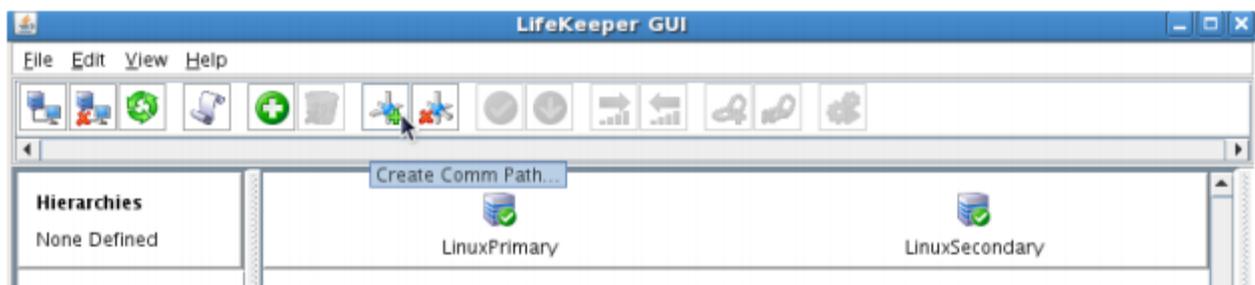


## Create Communication (Comm) Paths

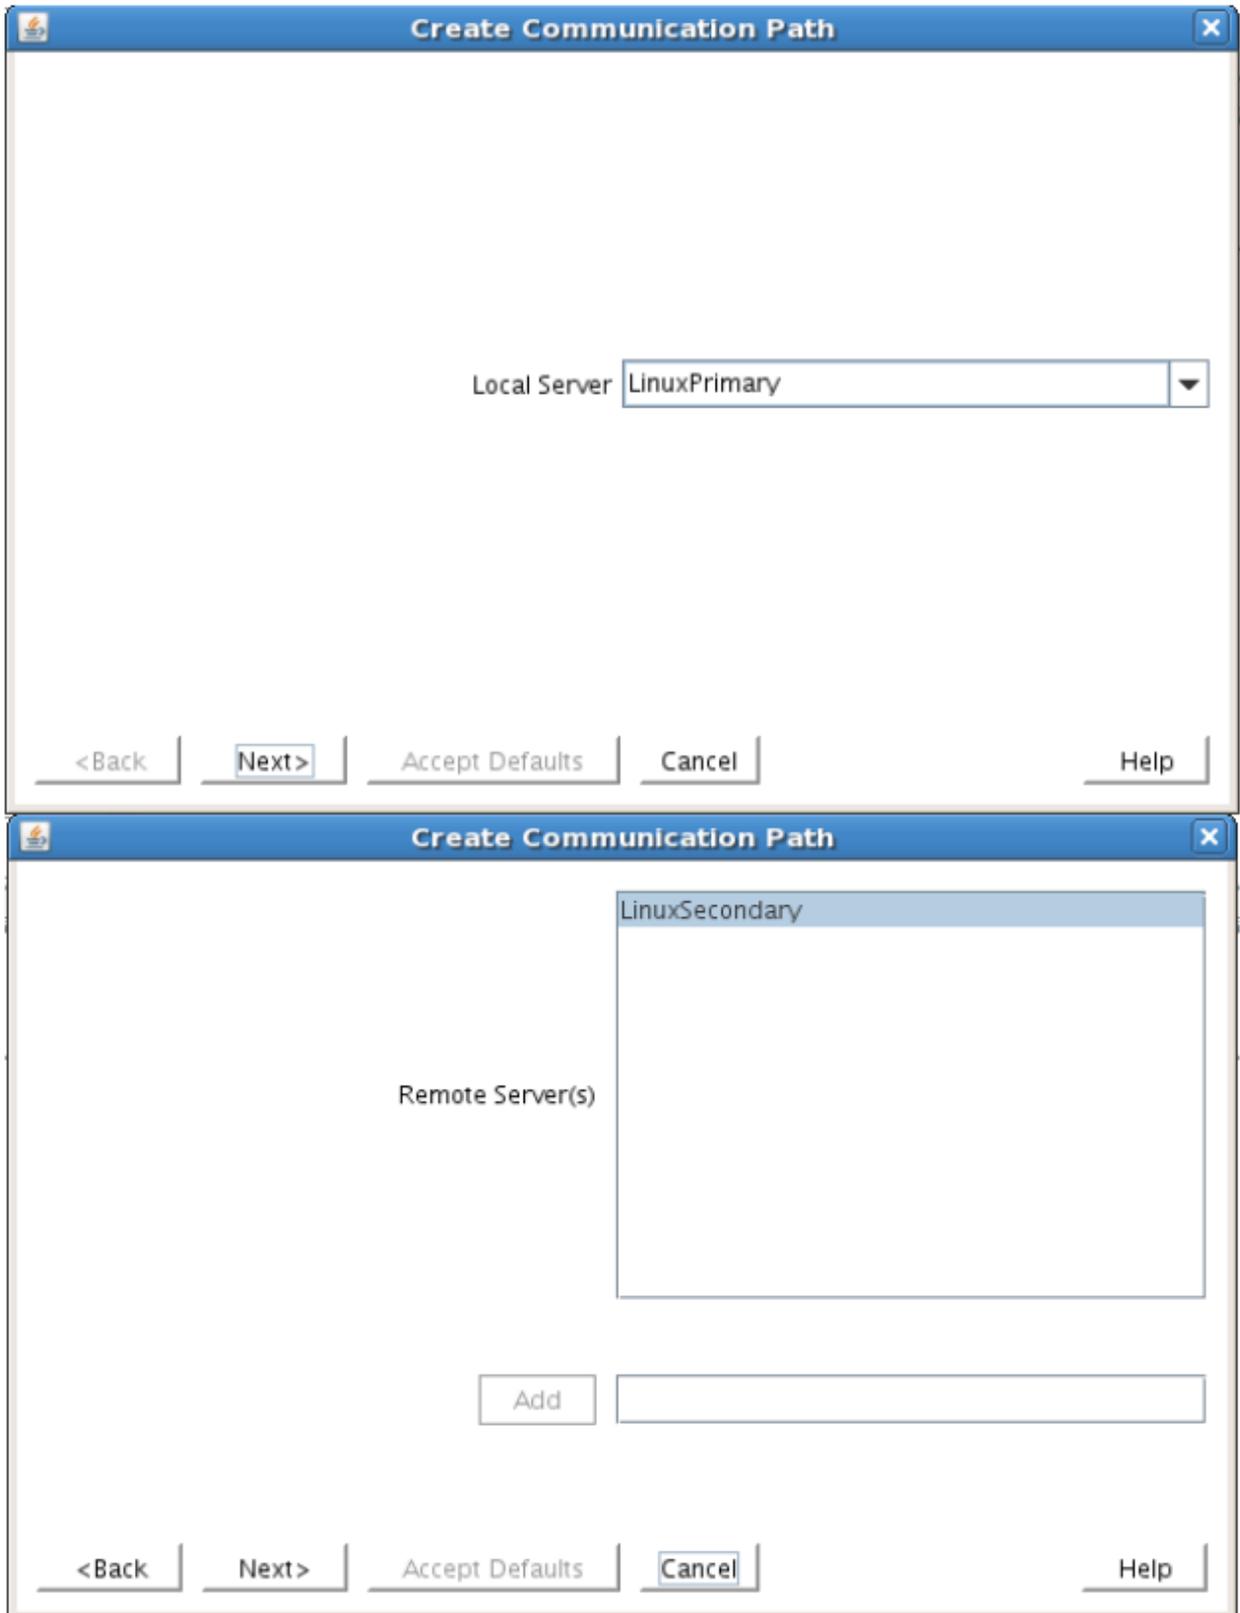
4. Within the LifeKeeper GUI, from the File menu, select Connect. Enter the name of your Secondary server, login and password when the Cluster Connect window displays.



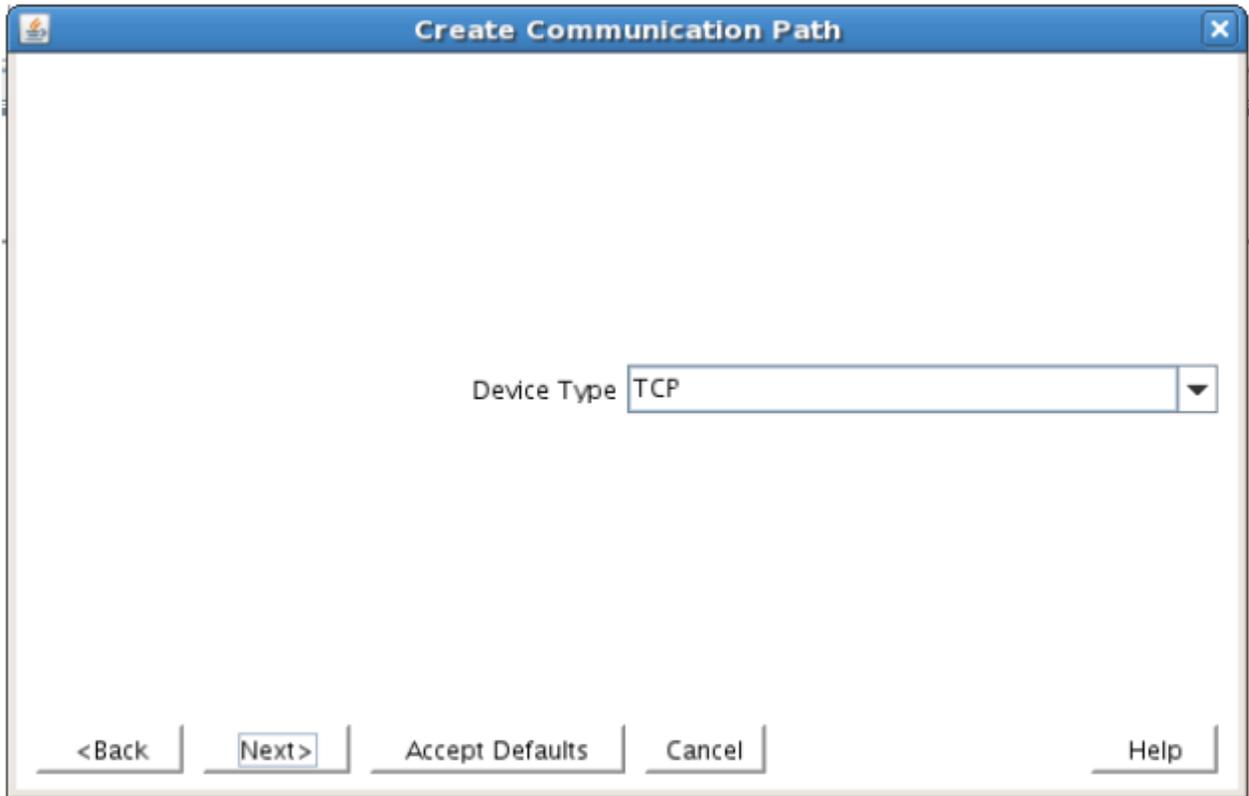
5. Within the LifeKeeper GUI, click the Create Comm Path button on the toolbar. You can also right click one of the servers and click Create Comm Path from the pop-up menu as well.



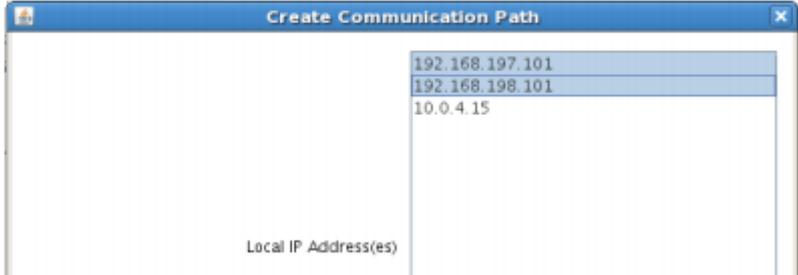
6. Select your Local and Remote Server(s) from the list box. If a server does not appear in the list box, you may enter it by typing its name and clicking the Add Server button. When using the Add Server procedure, you must make sure that the computer names for both network interfaces on the servers respond correctly when you ping them (from all of the partner server(s)) using the **ping -a IP ADDRESS** syntax. If they do not, this must be corrected prior to continuing. Click Next.

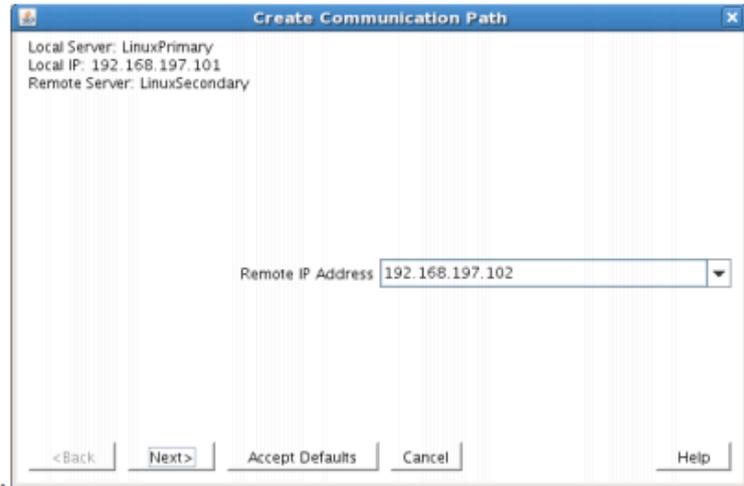


7. Select TCP for Device Type and Click Next.



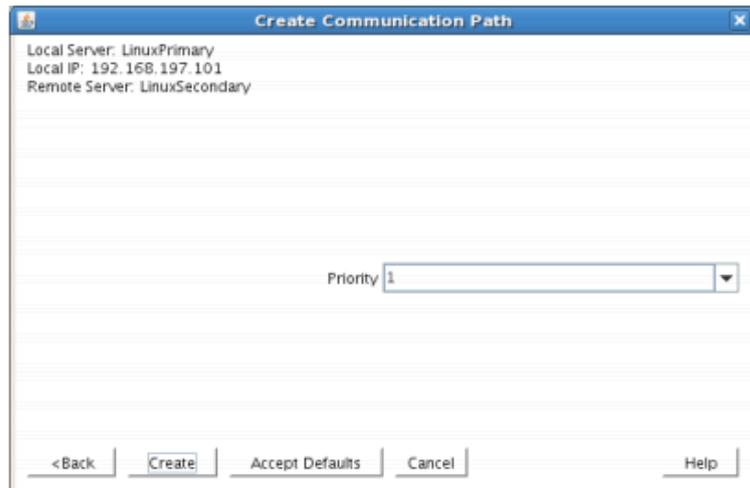
- 8. Provide all the required information and click Next for the following series of dialog boxes. For each field in the dialog box you can click Help for further information or refer to the table below for an explanation or recommendation

Field	Tips
<b>For TCP/IP Comm Path...</b>	Choose the IP address to be used by the local server for this comm path. Select both interfaces so that the wizard creates multiple comm. paths during this operation
Local IP Address	
Remote IP Address	Choose the IP address to be used by the remote server for this comm path

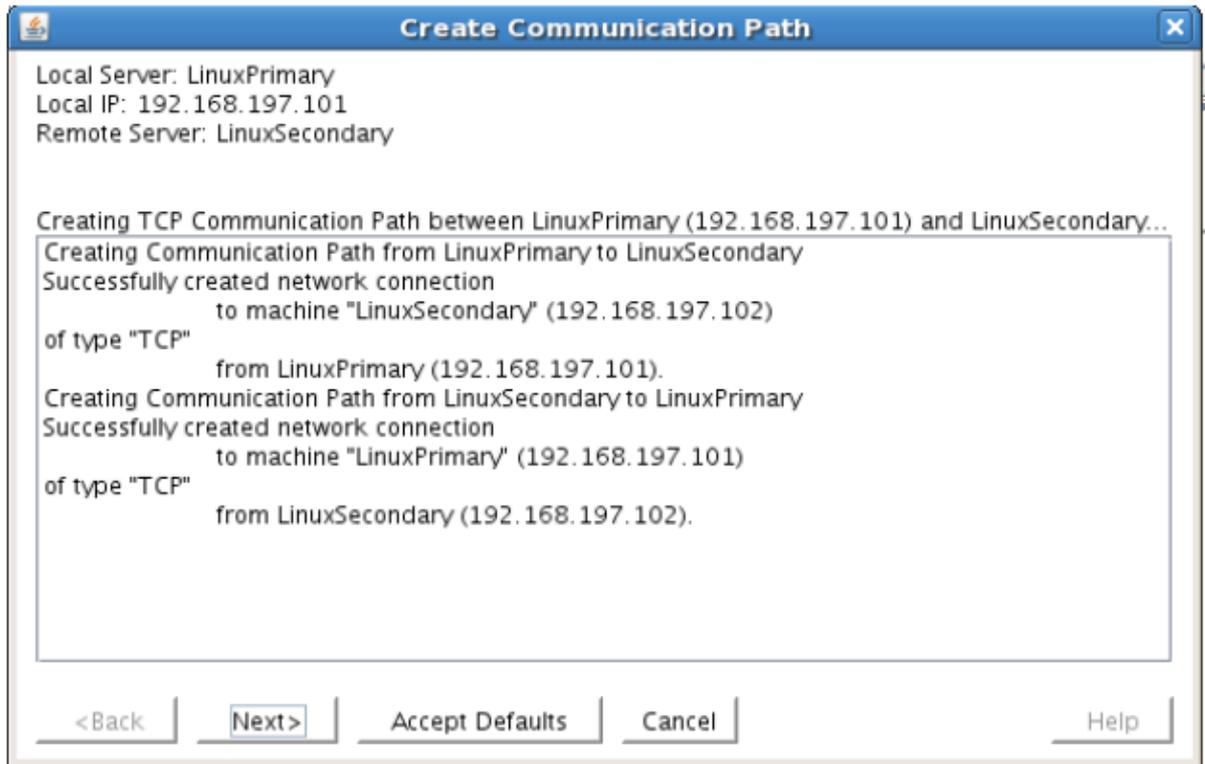


Enter the priority for the comm path on the local server. The priority will be used to determine the order that the comm paths between two servers will be used. Priority 1 is the highest; 99 is the lowest.

Priority



9. After entering data in all the required fields, select Create. A message will display indicating the network communication path is successfully created. Click Next.



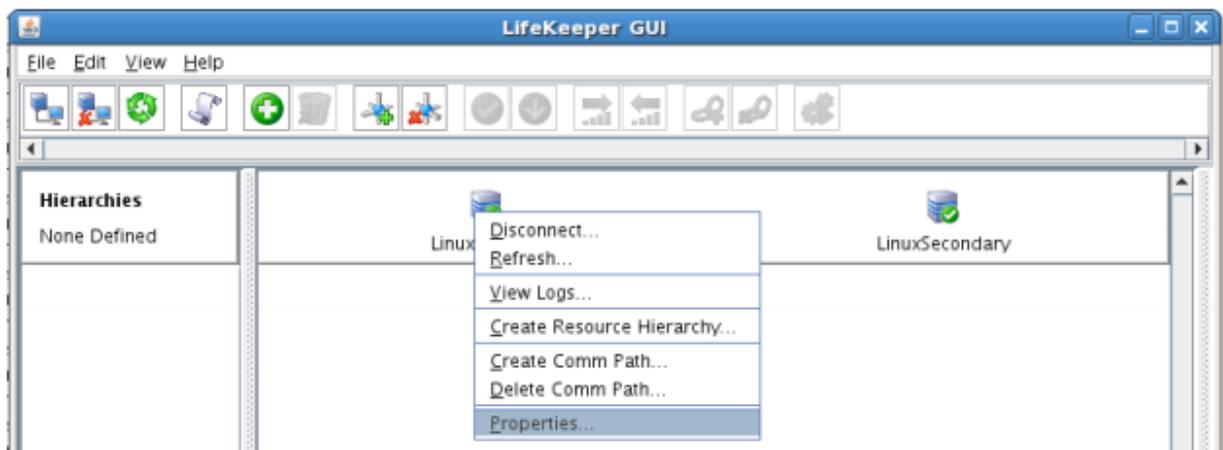
If you selected multiple Local IP Addresses or multiple Remote Servers and the Device Type was set to TCP, then the procedure will return you to the setup wizard the next Comm Path.

10. Select Done in the last dialog box.

Repeat this process until you have defined all the communication paths you plan to use. SIOS strongly recommends that you define at least two communication paths for redundancy.

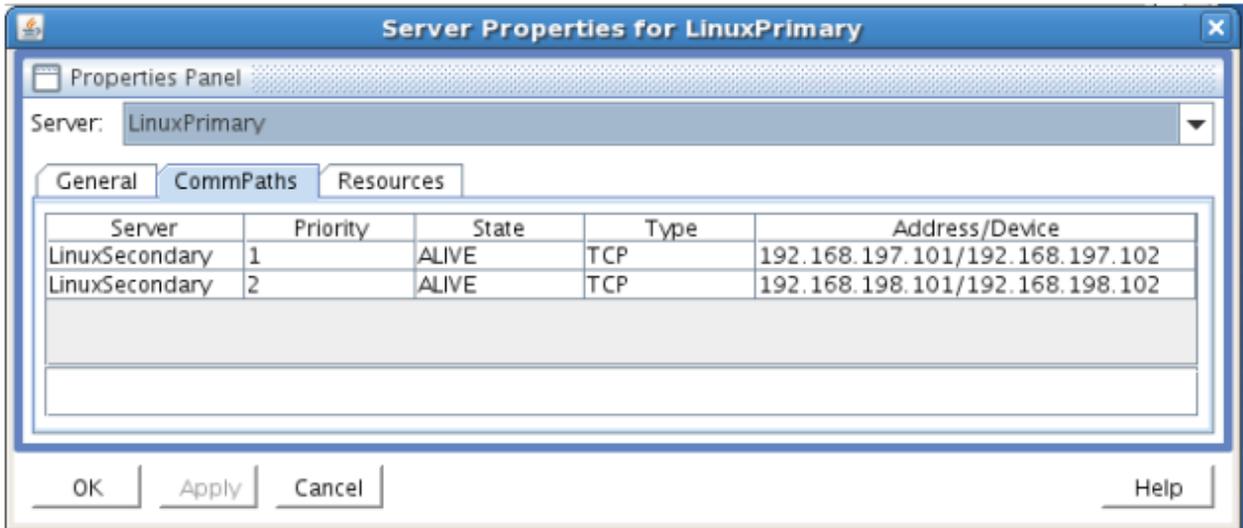
### Verify the Communications Paths

1. Verify that the communications paths are configured properly by viewing the Server Properties dialog box. From the LK GUI, select Edit, Server, Properties and then the Comm Paths tab.



2. Note the State displayed is ALIVE. You can also check the server icon in the right, main pane of

the GUI. If only one comm path has been created, the server icon shows a yellow warning icon on the server icon, indicating that one comm. path is ALIVE, but there is no redundant comm path. The server icon will display a green heartbeat checkmark when there are at least two comm paths configured and ALIVE.

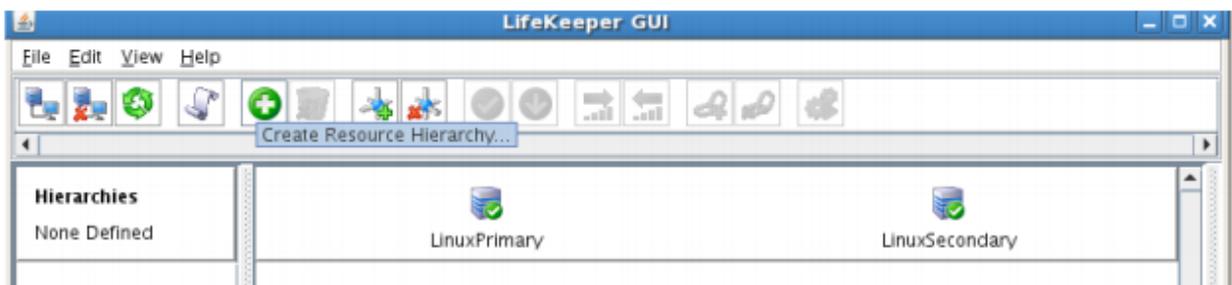


## Create the LifeKeeper Hierarchy

### Create and Extend an IP Resource

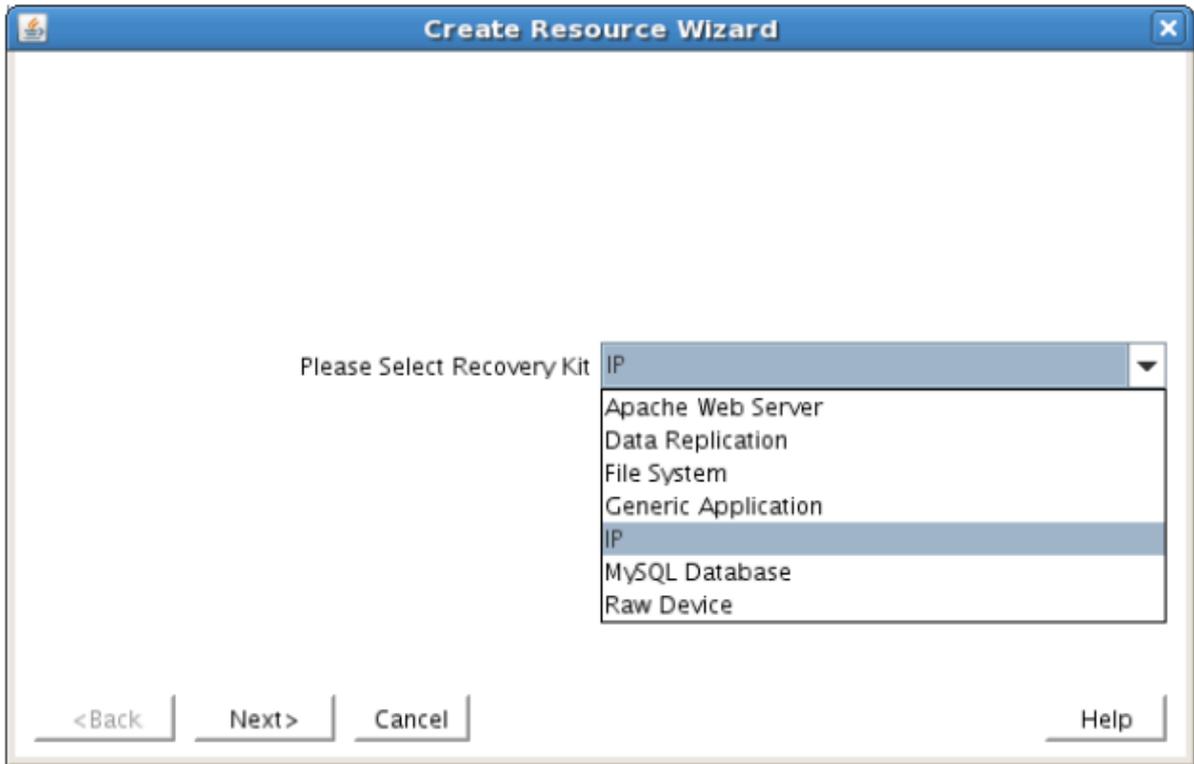
In LifeKeeper, create an IP resource and extend it to the secondary server by completing the following steps. This Virtual IP will have the ability to move between cluster nodes along the application that depends on it.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

2. Select IP Address and click Next.



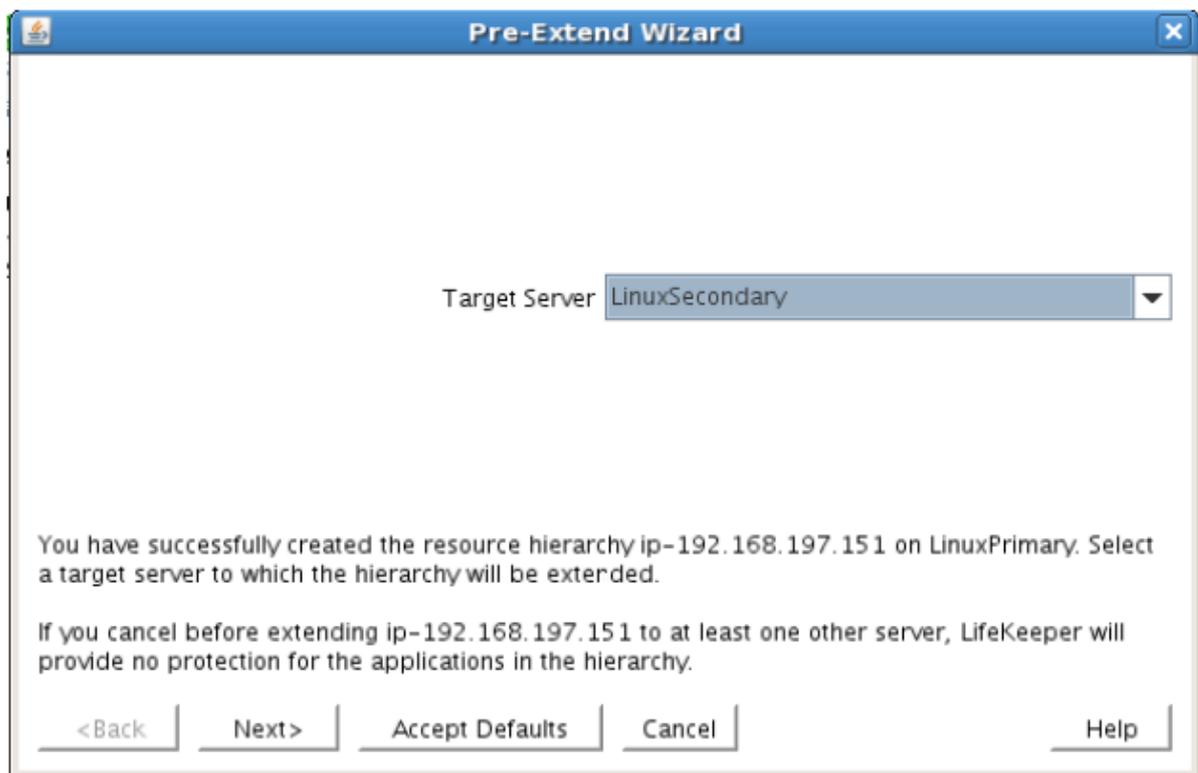
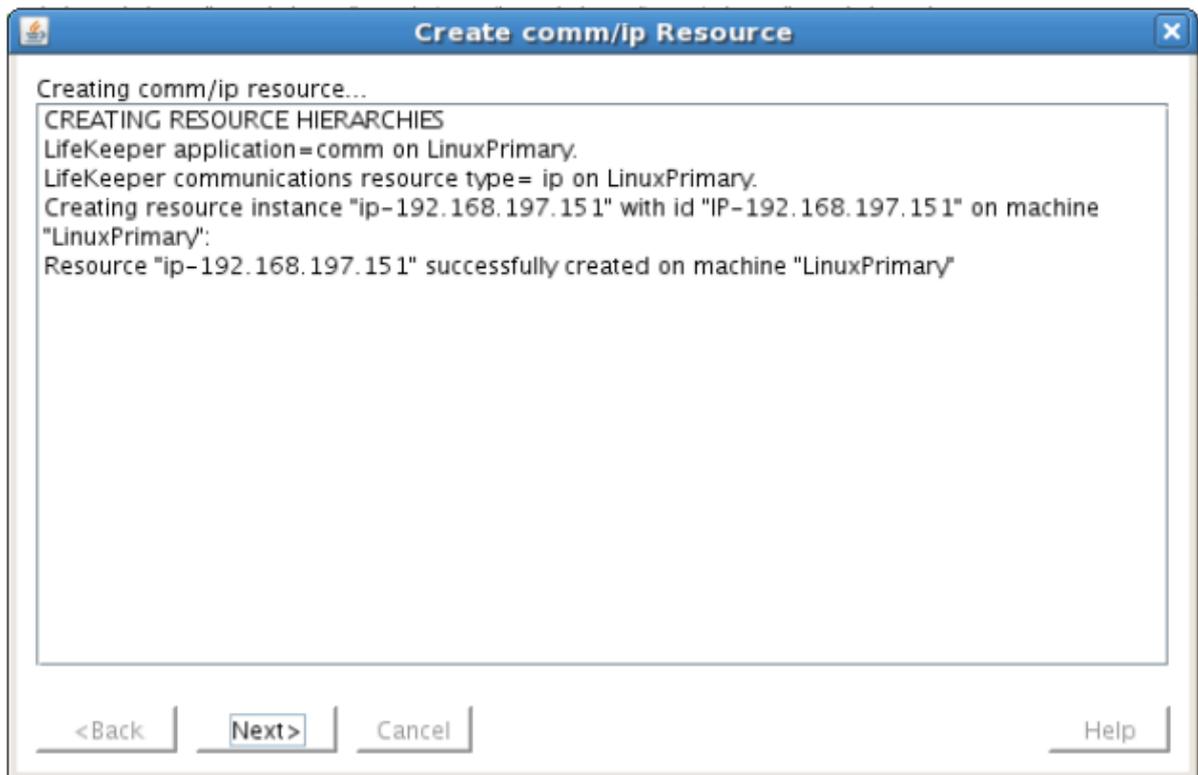
3. Enter the appropriate information for your configuration. The table below contains a list of the fields that display and additional information to assist you as you complete this procedure. Recommended values are also show below. You can also click the Help button for further information. Press Next to continue after entering the required information.

## IP Creation Field Definitions

Field	Tips
Resource Type	Select IP Address as the resource type and click Next.
Switchback Type	Select Intelligent and click Next.
Server	Select the Server where the IP resource will be created. Select your Primary server and click Next.
IP Resource	<p>Enter the virtual IP information and click Next</p> <p>Example <b>192.168.167.151</b></p> <p><b>Note</b> This is an IP address that is not currently in use anywhere on your network. This is the address that all clients will use to connect to the protected resources.</p>

Netmask	<p>The IP subnet mask that your TCP/IP resource will use on the target server. Any standard netmask for the class of the specific TCP/IP resource address is valid.</p> <p>In our sample configuration 255.255.255.0 is used for a subnet mask on both networks.</p> <p><b>Note:</b> The subnet mask you choose, combined with the IP address, determines the subnet that will be used by the TCP/IP resource and should be consistent with the network configuration.</p>
Network Connection	<p>This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. Select the correct NIC and click Next.</p>
IP Resource Tag	<p>Accept the default value and click Next. This value only affects how the IP is displayed in the GUI. The IP resource will be created on our Primary server.</p>

- LifeKeeper will create and validate your resource. After receiving the message that the resource has been created successfully, click Next when the following dialog box appears so that you can complete the process of Extending the IP Resource to our Secondary server, below.



Extending the IP resource will start automatically after you have finished creating an IP address resource if you clicked Next in the dialog box displayed above. You can also start this from an existing IP address resource by right clicking on the active resource and selecting Extend Resource Hierarchy.

Refer to the table below to complete the Extend IP Resource procedure.

Field	Recommended Entries or Notes
Switchback Type	Leave as "intelligent" and click Next
Template Priority	Leave as default (1)
Target Priority	Leave as default (10)
Network Interface	This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. The correct physical NIC should be selected by default. Please verify and then click Next
IP Resource Tag	Leave as default.
Target Restore Mode	Select Enable and click Next.
Target Local Recovery	Select Yes to enable Local Recovery for the SQL resource on the Target server.
Backup Priority	Accept the default value.

- After receiving the message Hierarchy extend operations completed, click Finish and then click Done
- Your IP resource (192.168.197.151) is now fully protected and has the ability to "float" between cluster nodes as needed. Looking at the LifeKeeper GUI you will notice that the IP resource is Active on the Primary cluster node and Standby on the Secondary cluster node



## Create the Shared Filesystem Resource Hierarchy

Create a Filesystem resource to protect the shared iSCSI filesystem and make it high available between

cluster nodes. LifeKeeper leverages SCSI Persistent Group Reservations (PGR) to lock the LUN, ensuring that only the active cluster node for the storage resource can access it.

**\* Important** At this point, the shared iSCSI LUN needs to already be mounted on the Primary Server. It should NOT be mounted on the Secondary Server. See section titled “Configure iSCSI initiator, discover and login to iSCSI target” above to review the steps involved.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.
2. Select File System and click Next.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node)
Mount Point	Select /var/lib/pgsql . Note that LifeKeeper scans the system for LUNS that are sharable between cluster nodes. The list of possible shared LUNS is presented automatically in this step of the wizard.

4. Select Create Instance to define this resource hierarchy on the Primary Server
5. Click Next to Extend the File System Resource to the Secondary Server
6. In the Extend Wizard, select “Accept Defaults”
7. As a result the File System resource is now protected on both cluster nodes. Click Finish to exit the Extend wizard.
8. Your resource hierarchy should look as follows:



## Create the PostgreSQL Resource Hierarchy

Create a PostgreSQL resource to protect the PostgreSQL database and make it high available between cluster nodes.

**\* Important** At this point, PostgreSQL needs to be already running on the Primary Server. It should NOT be running on the Secondary Server. See section titled “Install, Configure, and Start PostgreSQL” above to review the process to configure and start PostgreSQL as needed.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.
2. Select PostSQL Database and click Next.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node)
PostgreSQL Executable	Leave as default (/usr/bin) since we are using a standard PostgreSQL
Location	install/configuration in this example. This field is used to specify the directory path containing the PostgreSQL executables.
PostgreSQL Client	Leave as default (/usr/bin/psql) . This field is used to specify the directory
Executable Location	path containing the PostgreSQL executable psql.
PostgreSQL Administration	Leave as default (/usr/bin/pg_ctl). This field is used to specify the
Executable Location	directory path containing the PostgreSQL executable pg_ctl.
PostgreSQL Data Directory	/var/lib/pgsql/data. This field is used to specify the location of the PostgreSQL data directory (datadir) that will be placed under LifeKeeper protection. The specified directory must exist and reside on a shared or replicated file system.
PostgreSQL Port 5432.	This field is used to specify the TCP/IP port number on which the postmaster daemon is listening for connections from client applications
PostgreSQL Socket Path	Leave as default (/tmp/.s.PGSQL.5432 ) . This field is used to specify the full path to the Unix- domain socket on which the postmaster daemon is listening for connections from client applications.
Enter Database Administrator User	Enter "postgres" . This field is used to specify a PostgreSQL Database Administrator User name for the specified database instance with connection and administrator privileges for the instance
PostgreSQL Logfile	Leave as default (/tmp/pgsql-5432.lk.log) . This field is used to specify the log file path that will be used for the PostgreSQL log file.
Database tag	Leave as default

4. Select Create to define the PostgreSQL resource hierarchy on the Primary Server
5. Click Next to Extend the File System Resource to the Secondary Server
6. In the Extend Wizard, select "Accept Defaults"
7. As a result the PostgreSQL resource is now protected on both cluster nodes. Click Finish to exit the Extend wizard.
8. Note: LifeKeeper will automatically identify that the PostgreSQL resource has a dependency on the FileSystem resource (/var/lib/pgsql). The FileSystem Resource will appear underneath the PostgreSQL resource in the GUI

9. Your resource hierarchy should look as follows:



### Create the PostgreSQL IP Address Dependency

In this step will define an additional dependency: that PostgreSQL depends on a Virtual IP (192.168.197.151) so that the IP address follows the PostgreSQL database should it move.

1. From the LifeKeeper GUI toolbar, right-click on the “pgsql-5432” resource
2. Select “Create Dependency” from the right-click context menu
3. In the Child Resource Tag dropdown menu, select “ip-192.168.197.151”
4. Click Next
5. Click Create Dependency
6. Click Done
7. The Virtual IP address resource (192.168.197.151) will now appear underneath the PostgreSQL resource in the LifeKeeper user interface. This ensures that resources move together, and are started/stopped in the proper order.
8. Your resource hierarchy should look as follows



At this point in the Evaluation, we have fully protected PostgreSQL, and its dependent resources: IP addresses, and Shared Storage.

## 12.7.8. Test Your Environment – PostgreSQL

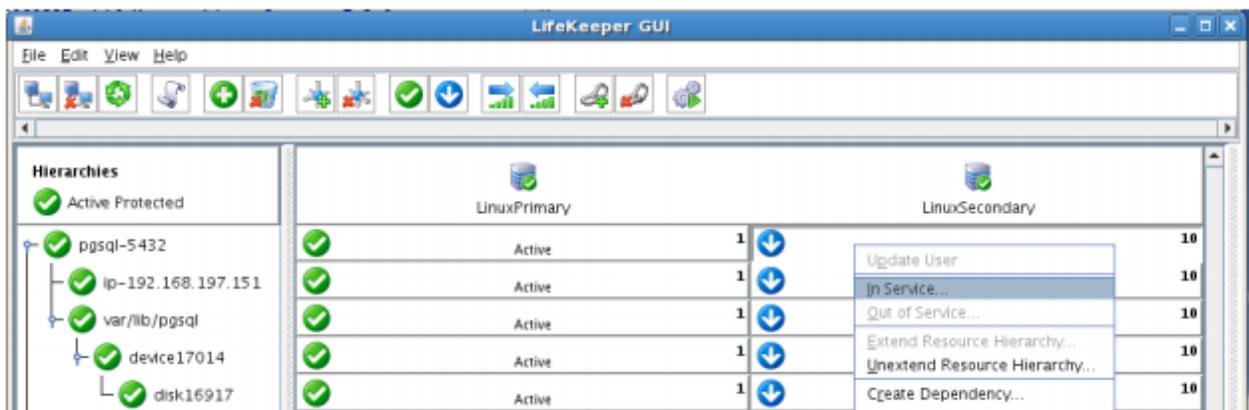
The following test scenarios have been included to guide you as you get started evaluating LifeKeeper for Linux. **Before beginning these tests, make sure the data replication resources are in the mirroring state.**

Note: For these test examples, the Primary Server is referred to as LINUXPRIMARY. The Backup or Secondary Server is referred to as LINUXSECONDARY.

### Manual Switchover of the PostgreSQL Hierarchy to Secondary Server

#### Procedure:

- From the LifeKeeper GUI, right click on the PostgreSQL resource on the Secondary Server (LINUXSECONDARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up



#### Expected Result:

- Beginning with the PostgreSQL resource, all resources will be removed from service on the Active Server (LINUXPRIMARY).
- Beginning with the dependent resources (IP and Shared Volume), all resources will be brought in service on LINUXSECONDARY

#### Tests/Verification:

- Using the LifeKeeper GUI, verify that the PostgreSQL and dependent resources are active on LINUXSECONDARY.
- Run “ifconfig -a” on LINUXSECONDARY to validate that the IP Address 192.168.197.151 is active on LINUXSECONDARY
- Run “df -h” to verify that the /var/lib/pgsql shared iSCSI filesystem is mounted on LINUXSECONDARY
- Verify the PostgreSQL services are running on LINUXSECONDARY by running “ps -ef | grep -i

postgres”

- On LINUXSECONDARY run the following command to verify client connectivity to the PostgreSQL database:
  - # su – postgres
  - # psql
- Verify that the SCSI reservation has properly locked the iSCSI LUN to the currently active cluster node. On LINUXPRIMARY, run “mount /dev/sdc1 /var/lib/pgsql”. This should FAIL because LINUXPRIMARY does not own the SCSI reservation on this LUN.

## Manual Switchover of the PostgreSQL Hierarchy back to Primary Server

### Procedure:

- From the LifeKeeper GUI, right click on the PostgreSQL resource on the Primary Server (LINUXPRIMARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up

### Expected Result

- Beginning with the PostgreSQL resource, all resources will be removed from service on the Active Server (LINUXSECONDARY).
- Beginning with the dependent resources (IP and Shared Volume), all resources will be brought in service on LINUXPRIMARY

pgsql-5432	Active	1	StandBy	10
ip-192.168.197.151	Active	1	StandBy	10
var/lib/pgsql	Active	1	StandBy	10
device17014	Active	1	StandBy	10
disk16917	Active	1	StandBy	10

### Tests/Verification:

- Using the LifeKeeper GUI, verify that the PostgreSQL and dependent resources are active on LINUXPRIMARY.
- Run “ifconfig –a” on LINUXPRIMARY to validate that the IP Address 192.168.197.151 is active on LINUXPRIMARY
- Run “df –h” to verify that the /var/lib/pgsql shared iSCSI filesystem is mounted on LINUXPRIMARY
- Verify the PostgreSQL services are running on LINUXPRIMARY by running “ps –ef | grep –i postgres”
- On LINUXPRIMARY run the following command to verify client connectivity to the PostgreSQL database:
  - # su – postgres

- #35; psql

- Verify that the SCSI reservation has properly locked the iSCSI LUN to the currently active cluster node. On LINUXSECONDARY, run “mount /dev/sdc1 /var/lib/pgsql”. This should FAIL because LINUXSECONDARY does not own the SCSI reservation on this LUN.

## Simulate a network failure on the Primary Server by failing the IP resource

**✿ IMPORTANT NOTE:** Only perform this test if you have more than one communications path configured.

If you perform this test and have only one communications path configured, your system will enter a split-brain scenario as described in the LifeKeeper Administration Guide found here. Refer to this document for more information or contact SIOS presales technical support for assistance in resolving this condition.

### Procedure

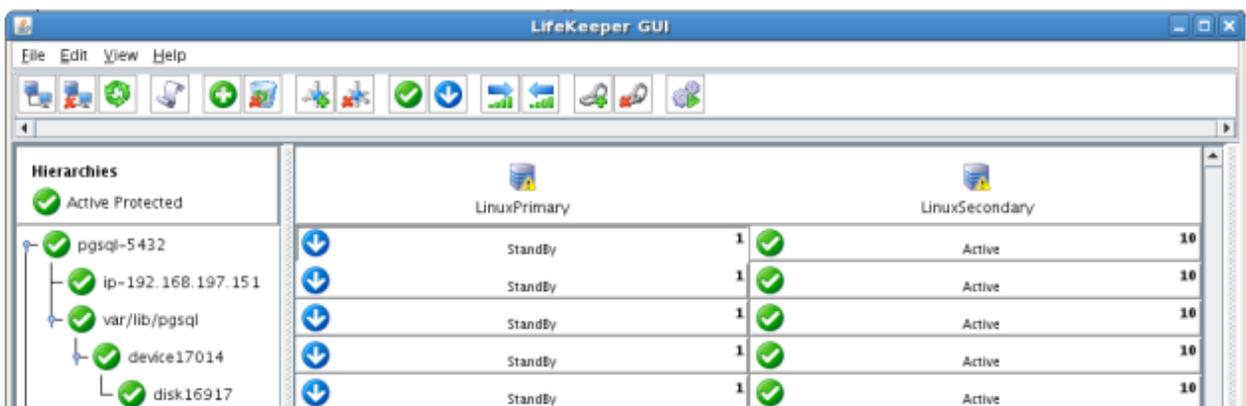
- On LINUXPRIMARY, pull the network cable attached to the NIC that the virtual IP address is configured on

### Expected Result:

- The IP Resource should fail first.
- The entire hierarchy should failover to LINUXSECONDARY

### Tests/Verification:

- Check the LifeKeeper Log to verify the IP resource failed – “/opt/LifeKeeper/bin/lk\_log log”
- Using the LifeKeeper GUI, verify the PostgreSQL resource hierarchy fails over successfully to LINUXSECONDARY
- After this test has been completed, re-connect the network cable on LINUXPRIMARY



## Hard failover of the resource from the Secondary Server back to the Primary Server

### Procedure:

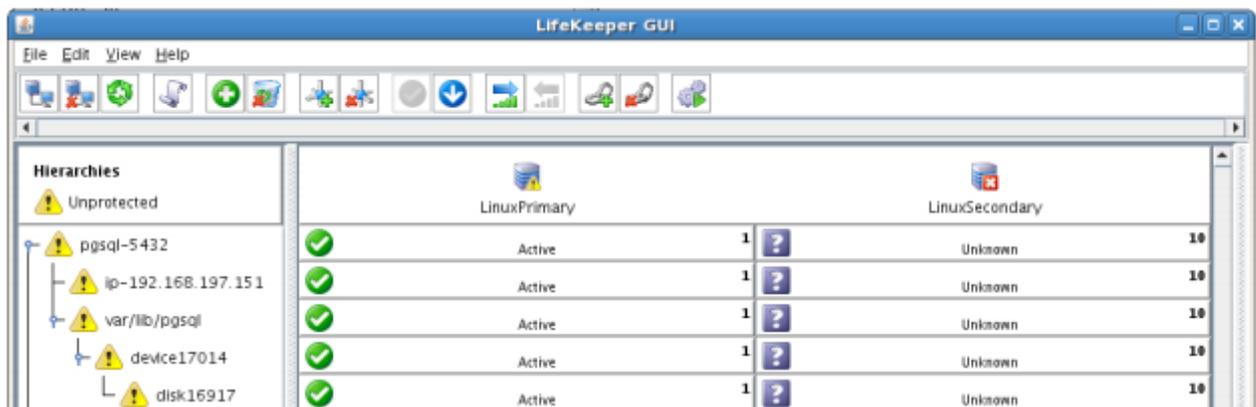
- Pull the power cord on LINUXSECONDARY, as this is the server with all resources currently In Service.

### Expected Result:

- After failure has been detected, beginning with the dependent resources (IP and Volume), all resources will be brought in service on LINUXPRIMARY.

### Tests/Verification:

- Using the LifeKeeper GUI, verify the mirror has reversed and is in a Resync Pending state waiting for LINUXSECONDARY to come back on line.
- Verify the PostgreSQL Server services are running on LINUXPRIMARY.



## Bring Failed Server back on line

### Procedure:

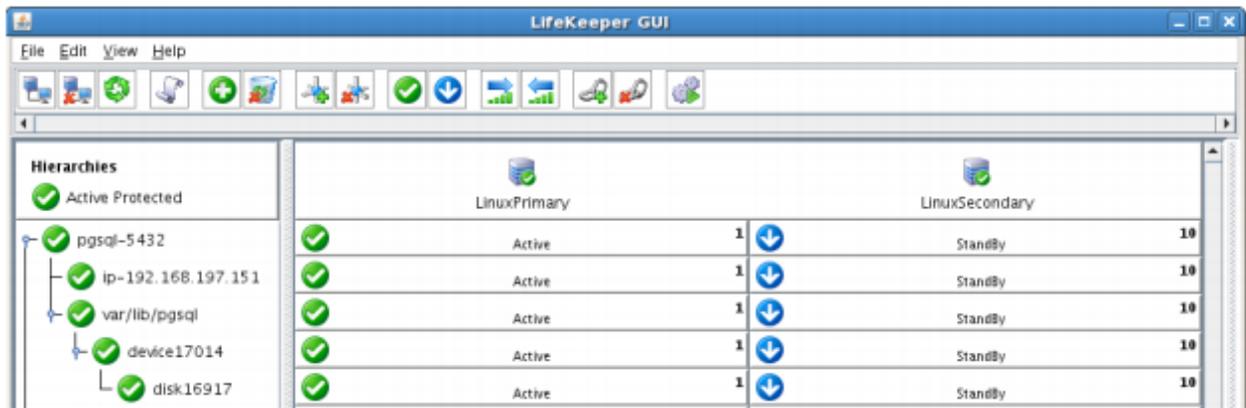
- Plug the power cord back into LINUXSECONDARY and boot it up.

### Expected Result:

- Using the LifeKeeper GUI, verify that LINUXSECONDARY is coming back up and has become the Standby Server.

### Tests/Verification:

- Verify the PostgreSQL Hierarchy is in service on LINUXPRIMARY and standby on LINUXSECONDARY.



## Verify Local Recovery of PostgreSQL Server

### Procedure:

- Kill the PostgreSQL processes via the command line:
- # ps -ef | grep postgres
- # (kill -9 the PIDs returned)
- run “ps -ef | grep postgres once again to verify that the processes no longer exist

### Expected Result: (Assumes Local Recovery for SQL resource is set to YES)

- The PostgreSQL Server service should stop.
- The PostgreSQL quickcheck process will automatically restart the PostgreSQL Server Service when it runs periodically.
- No failure of PostgreSQL should occur.

### Tests/Verification:

- Execute “ps -ef | grep postgres” once again to verify that the postgresql processes have been restored locally on LINUXPRIMARY.
- Verify connectivity to the postgresql database by running:
  - ↑
  - # su - postgres
  - # psql
- If you inspect the LifeKeeper logs, you will see information indicating that LifeKeeper detected the failure of the PostgreSQL service and recovered it locally. Run /opt/LifeKeeper/bin/lk\_log log for more information.

# 12.8. Apache/MySQL Cluster Using Both Shared and Replicated Storage

---

## Objective

This document is intended to aid you in installing, configuring and using the LifeKeeper for Linux evaluation product to make Apache and MySQL highly available. If Apache and MySQL are not already installed, please allocate some time to install it on your servers. Once this task has been completed, you may install and configure LifeKeeper for Linux.

There are five phases in this process:

- Phase 1 – Prepare to Install
- Phase 2 – Configure Storage
- Phase 3 – Install and Configure Apache/PHP
- Phase 4 – Install and Configure MySQL
- Phase 5 – Install LifeKeeper for Linux
- Phase 6 – Configure your LifeKeeper Cluster
- Phase 7 – Test Your Environment

## 12.8.1. Terms to Know – Apache

---

The following terms are used throughout this document and while some may be familiar to you, it may be helpful to review how SIOS defines and uses these terms.

### Network Communication Terms

**Crossover cable** – A cable used to directly connect computing devices together instead of being connected to a network switch, hub or router. This cable creates an isolated, private network to allow cluster-related and data replication traffic to flow between systems.

### Types of LifeKeeper Servers

**Server** – A computer system dedicated to running software application programs.

**Active Server** – This is the server where the resource hierarchy is currently running (IN SERVICE).

**Standby Server** – This is the server where the resource hierarchy is defined, but is not currently running. This server is available to bring the resource hierarchy into service should something happen to the resource hierarchy on the Active Server.

**Primary Server** – This is the server in a LifeKeeper configuration with the highest priority for a given resource hierarchy. It is the server that provides services for the resource hierarchy under normal circumstances.

**Secondary Server** – This is the server in a LifeKeeper configuration with the 2nd highest priority for a given resource hierarchy.

**Source Server** – This is the server in a LifeKeeper cluster that is using data replication (Active Server). It is where the resource hierarchy is currently running and the replicated partition (Source Partition) is accessible for writes.

**Target Server** – This is the server in a LifeKeeper cluster using data replication (Standby Server). The replicated partition (Target Partition) is updated with writes from the Source Partition by the SIOS Data Replication system. This partition should not be accessed/modified manually.

### SIOS Data Replication Terms

**Replication** – Transferring data from one partition to another via a sector-by-sector copy. During replication, the target partition should not be accessed or modified assuring your data integrity.

**Synchronous** – A replication scheme in which the data is confirmed written and valid on the target before the write operation occurs on the source disk through a series of information exchanges. Synchronous mirrors should only be implemented on high speed (100Mbps+) networks due to the network overhead involved.

**Asynchronous** – A replication scheme in which the data is released for writing on the source immediately and is sent to the target(s) simultaneously for writing as fast as the data can get there and can be written on them.

**Rate of Change** – A measure of the amount of data which is changing over a set period of time.

**Compression** – An algorithm which is optionally implemented to reduce the amount of traffic between source and target nodes. Nine levels of compression are offered. Compression is turned off by default.

**Throttling** – An optionally implemented mechanism to limit the bandwidth used for replication.

## LifeKeeper Product Terms

**Communications Path** – A mechanism supporting communication between nodes in a LifeKeeper cluster. SIOS highly recommends implementing multiple communication paths between all servers in the cluster to eliminate a single point of failure.

**Heartbeat** – A periodic message exchanged between nodes in a LifeKeeper cluster that provides server health monitoring. A heartbeat message is one type of inter-node cluster communication sent over a communications path.

**Split Brain** – A situation in which all communications paths between cluster members fail, but all servers remain up and running. In this situation, both systems believe the other has failed and both believe they should keep or bring resources into service.

**Failover** – The unplanned migration of a resource hierarchy to the Standby Server because of a system or resource failure on the Active Server.

**Switchover** – The planned migration of a resource hierarchy from the Active Server to the Standby Server.

**Switchback** – The setting that governs the recovery behavior of the server where the resource was in service when it failed. If the setting is “intelligent”, the server acts as a possible backup for the given resource. If the setting is “Automatic”, the server actively attempts to re-acquire the resource without further notice.

**Resource** – A system asset that can be protected by LifeKeeper. Resources can be used to represent disk partitions, virtual IP addresses, applications, etc.

**Extend a Resource** – Create or define an already configured LifeKeeper resource onto another server in the cluster and build an equivalency relationship that prevents the resource from coming in service on both systems simultaneously.

**Resource Hierarchy** – A grouping of resources, in a predetermined order, from high to low. This may also be referred to as simply a Hierarchy.

**Shared Storage** – One or more logical disk partitions that are physically attached to all nodes in a cluster. LifeKeeper ensures that the volume is only accessible by one server at a time. This is formally

called I/O fencing.

**Data Replication (Disk Mirroring)** – The replication of logical disk partitions to separate physical hard disks in real time to ensure continuous availability, currency and accuracy of data.

**Source** – The partition on the source server used for replication. The “gold” copy of the data.

**Target** – The partition on the target server used for replication.

**Switchable IP Address** – A unique IP address that may be moved between systems in the cluster. Client systems connect to this address and the system where the virtual IP resource is active will respond to requests.

## 12.8.2. The Evaluation Process – Apache

---

SIOS strongly recommends performing your evaluation of LifeKeeper for Linux in a test lab environment. SIOS is not responsible and cannot provide support for evaluation software installed in a production environment.

All questions during the evaluation period should be directed to [evalsupport@us.sios.com](mailto:evalsupport@us.sios.com) or your local Pre-Sales Engineering contact. Pre-sales support will contact you by the next business day to answer questions. Once you are a licensed customer on software maintenance, you will have access to 24 × 7 post-sales technical support.

 **Important** Your evaluation license is valid for a limited period of time from the day you receive the SIOS product evaluation package and licenses via email from the SIOS sales team.

## 12.8.3. Prepare to Install – Apache

---

### Hardware Requirements

#### Primary and Secondary Servers

- Systems must meet the minimum requirements for the Linux distribution to be used during the evaluation.
- 512MB RAM minimum; 1GB RAM recommended.
- 2GB of available hard disk space recommended.
- Multiple Network Interface Cards (NIC's) are recommended.
- Configure one or more additional disk partitions to be used for data replication. On the primary server, these will become the source partitions. On the secondary server(s), these will become the target partitions. In this evaluation example we will be replicating MySQL data (which will be a partition mounted at /var/lib/mysql)
- For replicated partitions, a target partition's size must equal to or larger than the size of its source partition.
- The system ( / ) and boot (/boot) partitions are not eligible for replication.

**Note:** You may use more than one partition for replicated data, allowing for separation of multiple database files and/or log directories for performance reasons.

#### Client

This system is not required but is recommended for testing the cluster environment.

- A standard linux terminal running the MySQL client can be used to test the configuration.

### Software Requirements

#### Primary Server and Secondary Server

- Linux Distribution x86\_64, AMD 64:
  - RedHat Enterprise Linux 5 (5.4+ recommended) or 6.x
  - CentOS Linux 5 (5.4+ recommended) or 6.x
  - Oracle Enterprise Linux 5 (5.4+ recommended), 6.3, 6.4
    - RedHat Compatibility Kernel Only
  - SuSE Linux Enterprise Server 10 or 11 (11 recommended)
  - See [Linux Release Notes](#) for a full list of supported Operating Systems

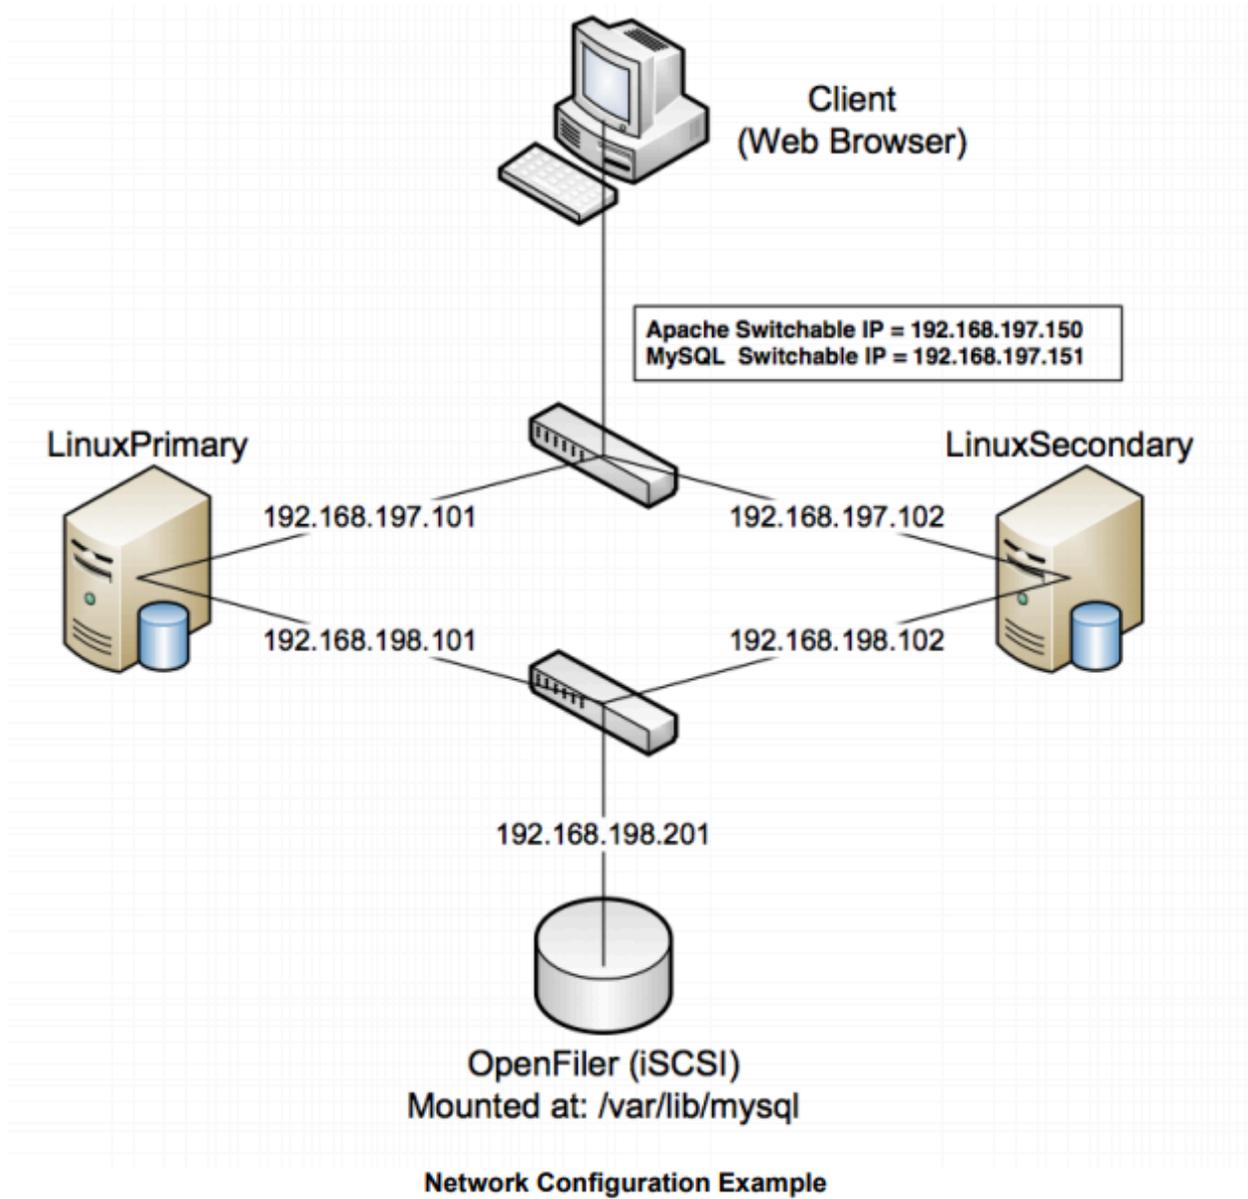
- Current patches / security updates are recommended.
- Satisfied dependencies; especially if the Linux installation package selection was base/minimal you will need to refer to the dependencies documentation at [Linux Dependencies](#)
- Its recommended that IPtables is disabled
  - # /etc/init.d/iptables off
  - # chkconfig iptables off
  - See [here](#) for information regarding the ports LifeKeeper for Linux uses.
- Disable SELinux :
  - Edit /etc/selinux/config
  - Set SELINUX=disabled (note: permissive mode is also acceptable)
- Check the configuration of your /etc/hosts file
  - localhost.localdomain and localhost are the only entries that can be on 127.0.0.1
  - Create a separate entry for your hostname with a static address
- GUI Authentication with PAM
  - LifeKeeper for Linux now leverages the Pluggable Authentication Module (PAM) provided in the Linux Standard Base (LSB).
    - Users are identified and authenticated against the system's PAM configuration. Privilege levels are determined from group membership as provided through PAM.
    - In order to access the GUI, a user must be a member in one of the three LifeKeeper groups: lkadmin, lkoper or lkguest.
    - See the following URL for more information on this topic:
      - [Configuring GUI Users](#)

## Network Requirements

For your evaluation, we recommend configuring your machines similarly to the following example.

LinuxPrimary and LinuxSecondary are multi-homed, between two LAN segments (the second NIC in each server could even be connected via a cross-over cable if a second physical network is not available). The second NIC is optional in this configuration, but highly recommended in production environments to avoid a single point of failure.

In this example evaluation scenario we will be leveraging both local, replicated storage (with the Apache configuration) as well as Shared Storage (iSCSI, for the MySQL configuration). OpenFiler is a storage appliance server that will serve an iSCSI target to LinuxPrimary and LinuxSecondary.



## Primary Server and Secondary Servers

- Configure the Host file with entries for all LifeKeeper protected servers. This is typically `/etc/hosts`.

Example:

192.168.197.101 LinuxPrimary

192.168.197.102 LinuxSecondary

- See your Network Administrator to obtain an unused IP Address to be used as the switchable IP Address. This switchable IP Address will be created later in the configuration process.
- Public Network connection(s) configured with:
  - Static IP address
  - Correct subnet mask
  - Correct gateway address
  - Correct DNS server address(es)
  
- Private Network connection(s) configured with:
  - Static IP address (on a different subnet from the public network)
  - Correct network mask
  - No gateway IP address
  - No DNS server addresses

## Client

Must be able to communicate on the same subnet/network as the servers Public interface addresses. In our example, this is the 192.168.197.0/24 network.

## 12.8.4. Configure Storage – Apache

---

### Before you Begin

Ensure the following:

- You have an extra disk/partition on both servers that can be used for data replication. A target volume's size must equal to or larger than the size of its source disk/partition.
- Shared storage is available. This can either be Fiber Channel SAN, iSCSI, NAS, etc. In this example we will review configuration of an iSCSI target for use as our MySQL database storage repository.

### Partition local storage for use with SIOS DataKeeper for Linux

#### Primary Server

On your Primary server, perform the following actions:

1. Identify an existing free, unused disk partition to use as our Apache repository. Alternatively, create a new partition. Use the "gdisk" utility to partition your disk appropriately. In this example /dev/sdb is an unused disk where we will create a single partition
  - a. `gdisk /dev/sdb`
  - b. Press "n" to create a new partition
  - c. Press "p" to create a primary partition
  - d. This example uses a new disk, so we will use all default values (Partition 1, entire disk) Hit Enter twice to confirm these parameters
  - e. Press "w" to write the partition table and exit gdisks

#### Example

```
[root@LinuxPrimary ~]# gdisk /dev/sdb
```

Command (m for help): **n**

Command action

e extended

p primary partition (1-4)

**p**

Partition number (1-4): **1**

First cylinder (1-256, default 1): **<enter>**

Using default value 1

Last cylinder or +size or +sizeM or +sizeK (1-256, default 256): **<enter>**

Using default value 256

Command (m for help): **w**

The partition table has been altered! Calling ioctl() to re-read partition table.

Syncing disks.

```
[root@LinuxPrimary ~]#
```

```
[root@LinuxPrimary ~]# df /var/www
```

```
Filesystem 1K-blocks Used Available Use% Mounted on
```

```
/dev/sdb1 253855 11083 229666 5% /var/www
```

```
[root@LinuxPrimary ~]# ls /var/www
```

```
cgi-bin error html icons lost+found manual usage
```

2. Format the newly created disk partition

```
# mkfs.ext3 /dev/sdb1
```

3. Mount the partition temporarily at /mnt

```
# mount /dev/sdb1 /mnt
```

4. Move any existing data from /var/www/ into this new disk partition (assumes a default apache configuration)

```
# cd /var/www
```

```
# mv * /mnt
```

5. Remount /dev/sdb1 at /var/www

```
# cd /root
```

```
# umount /mnt
```

```
# mount /dev/sdb1 /var/www
```

- Note: there is no need to add this partition to `/etc/fstab`. Lifekeeper will take care of mounting this automatically.

## Result

```
[root@LinuxPrimary ~]# df /var/www
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/sdb1 253855 11083 229666 5% /var/www
[root@LinuxPrimary ~]# ls /var/www
cgi-bin error html icons lost+found manual usage
```

## Secondary Server

- On your Secondary server, only perform Step #1 above, where you partition the disk. The size of the Target volume needs to be the same size, or greater, than our Source volume.

# Configure iSCSI initiator, discover and login to iSCSI target

This Evaluation guide will not cover how to setup an iSCSI Target Server. It is assumed that the shared storage already exists in your environment. If you don't have shared storage and wish to configure it, a simple solution is to use OpenFiler (<http://www.openfiler.com/>), an Open Source storage management appliance, which can be run on physical hardware or as a virtual machine.

On both Primary and Secondary servers, perform the following functions:

- If not already installed, ensure that the **iscsi-initiator-utils** rpm package is installed:

```
# yum install iscsi-initiator-utils
```

- Start the `iscsid` service and enable it to automatically start when the system boots

```
# service iscsid start
```

```
# chkconfig iscsid on
```

- Configure the `iscsi` service to automatically start, which logs into iSCSI targets needed at system start up.

```
# chkconfig iscsi on
```

- Use the `iscsiadm` command to discover all available targets on the network storage server (OpenFiler)

```
# iscsiadm -m discovery -t sendtargets -p <name or IP of iSCSI server>
```

**Example**

```
[root@LinuxPrimary init.d]# iscsiadm -m discovery -t sendtargets -p 192.168.198.201  
iqn.2006-01.com.openfiler:tsn.mysql
```

5. Manually Login to the iSCSI Target

```
# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.mysql -p 192.168.198.201 --login
```

6. Configure Automatic Login

```
# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.mysql -p 192.168.198.201 --op  
update -n node.startup -v automatic
```

7. Use the “gdisk” command to format your iSCSI LUN, if needed

```
# gdisk /dev/sdc
```

8. Create a filesystem on your new iSCSI LUN Partition, sdc1

```
# mkfs.ext3 /dev/sdc1
```

9. Mount your iSCSI LUN at /var/lib/mysql (assuming a default mysql configuration). If data already exists in this directory, make sure to move it into the shared iSCSI LUN

```
# mount /dev/sdc1 /var/lib/mysql
```

10. At this point you now have a local partition, /dev/sdb1 mounted at /var/www and an iSCSI shared LUN, /dev/sdc1, mounted at /var/lib/mysql. Our disk layout now look as follows (example):

**Example**

```
[root@LinuxPrimary mysql]# df  
Filesystem 1K-blocks Used Available Use% Mounted on  
/dev/sda2 25967432 3683016 1976400 66% /  
/dev/sda1 101086 24659 71208 26% /boot  
tmpfs 517552 0 517552 0% /dev/shm  
/dev/sdb1 253855 11132 229617 5% /var/www  
/dev/sdc1 966644 38944 878596 5% /var/lib/mysql
```

## 12.8.5. Install and Configure Apache and PHP

---

### Before you Begin

Ensure the following:

1. If you are not familiar with installing and configuring Apache, please refer to the Apache documentation.
2. Apache should not be running at the time you attempt to protect it with LifeKeeper.

 **Important:** If Apache and PHP are already installed, skip to step #2 in the section below. You will need to verify the Apache configuration and/or relocate the location of the website data.

### Primary Server

1. Install Apache.
  - a. This example assumes that you are running a RHEL or CentOS 5.X based Linux distribution. For other Linux distros, please refer to your apache documentation for syntax differences.
  - b. If the Apache and PHP packages are not already installed, from the command line, , run: “yum install httpd php”
  - c. A number of dependencies will most likely be discovered. Install those as well.
2. Apache should be configured so that it will not automatically start when the server boots. LifeKeeper will control the start/stop of the webserver once its protected
  - a. Check the status of the webserver: “/etc/init.d/httpd status
  - b. If running, please stop of: /etc/init.d/httpd stop
  - c. Disable automatic startup: chkconfig httpd off
3. Apache Configuration. In this example, we will assume default Apache settings and directory locations, specifically:
  - a. Validate the following settings in /etc/httpd/conf/httpd.conf:
    - i. ServerRoot “/etc/httpd”
    - ii. DocumentRoot “/var/www/html”

- iii. Listen 192.168.197.150:80 (note 192.168.197.150 is the Apache Switchable IP we will configure later in the LifeKeeper user interface)
  - b. Edit the Listen parameter in the /etc/httpd/conf.d/ssl.conf configuration file
    - i. Listen 192.168.197.150:443
4. Create a sample Index page:
  - a. vi /var/www/html/index.php
  - b. Insert the following single line of code into this file:  

```
<? phpinfo(); ?>
```

## Secondary Server

1. Install Apache/PHP exactly as you did on the primary server, making all of the same configuration changes.
2. There is no need to perform Step #4 in which you create a sample index page. All data in /var/www will be replicated from LinuxPrimary to LinuxSecondary
3. Stop all Apache services on the secondary server

## 12.8.6. Install, Configure, and Start MySQL – Apache

---

### Primary Server

On your Primary server, perform the following actions:

1. Install both the “mysql” and “mysql-server” rpm packages if they do not exist on your system.

Apply any required dependencies as well

```
# yum install mysql mysql-server
```

2. Verify that your Shared iSCSI LUN is still mounted at /var/lib/mysql via the “df” command
3. If this is a fresh MySQL install, initialize a sample MySQL database:

```
# /usr/bin/mysql_install_db --datadir="/var/lib/mysql" --user=mysql
```

4. Ensure that all files in your MySQL data directory (/var/lib/mysql) have correct permissions and ownership

```
# chown -R mysql:mysql /var/lib/mysql
```

```
# chmod 755 /var/lib/mysql
```

5. Finally, manually start the MySQL daemon from the command line. Note: **Do Not** start it via the “service” command, or the /etc/init.d/ scripts

```
# mysqld_safe --user=root --socket=/var/lib/mysql/mysql.sock --port=3306 --  
datadir=/var/lib/mysql --log &
```

6. Verify MySQL is running by connecting with the mysql client:

```
[root@LinuxPrimary mysql]# mysql
```

```
Welcome to the MySQL monitor. Commands end with ; or \g.
```

```
Your MySQL connection id is 2
```

```
Server version: 5.0.77-log Source distribution
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.
```

```
mysql> exit
```

Bye

```
[root@LinuxPrimary mysql]#
```

7. Update the root password for your mysql configuration. In this example we set the MySQL root password to "SteelEye"

```
# echo "update user set Password=PASSWORD where User='root'; flush
privileges" | mysql mysql
```

8. Verify your new password:

```
# mysql mysql -u root -p

(Enter "SteelEye" as the password)

#exit
```

9. Create a MySQL configuration file. We will place this in the same shared directory (/var/lib/mysql/my.cnf)

```
# vi /var/lib/mysql/my.cnf
```

### Example

```
# cat /var/lib/mysql/my.cnf
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
pid-file=/var/lib/mysql/mysqld.pid
user=root
port=3306
# Default to using old password format for compatibility with mysql 3.x
# clients (those using the mysqlclient10 compatibility package).
old_passwords=1
# Disabling symbolic-links is recommended to prevent assorted security risks;
# to do so, uncomment this line:
# symbolic-links=0
[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid
[client]
user=root
password=SteelEye
```

10. Delete the original MySQL configuration file, located in /etc

```
# rm /etc/my.cnf
```

## Secondary Server

On your Secondary Server:

1. Install both the “mysql” and “mysql-server” rpm packages if they do not exist on your system.

Apply any required dependencies as well

```
# yum install mysql mysql-server
```

2. Ensure that all files in your MySQL data directory (/var/lib/mysql) have correct permissions and ownership

```
# chown -R mysql:mysql /var/lib/mysql
```

```
# chmod 755 /var/lib/mysql
```

3. There is no need to perform any of the additional steps taken on the Primary Server

## 12.8.7. Install LifeKeeper for Linux – Apache

For ease of installation, SIOS has provided the LifeKeeper for Linux with an installation script. Towards the end of the script, the desired Application Recovery Kits (ARKs) should be selected for installation. The software will be installed to the following locations:

SPS for Linux Component	Install Location
LifeKeeper Software	/opt/LifeKeeper
LifeKeeper Config File	/etc/default/LifeKeeper

Perform the following actions on **both** Primary and Secondary server.

### Download Software

1. Open the LifeKeeper evaluation email you received from SIOS.
2. Download the LifeKeeper Software from the link provided in your email. It is generally easiest to use “wget” to recursively download all files. Example:

a. # cd /root

b. # wget -r <URL>

c. After successful download you will have downloaded contents similar to the follow directory listing:

```
[root@LinuxPrimary ~]# ls -l <directory> total 63680
-rw-r--r-- 1 root root 23163 May 30 14:03 EULA.pdf
-rw-r--r-- 1 root root 536 May 30 14:03 readme.txt
-rw-r--r-- 1 root root 65179648 May 30 14:03 sps.img
```

3. Download your Evaluation license key from the link specified in your evaluation email. Save the license file to an easy to remember location on both servers.

### Run the LifeKeeper Installer Script

1. Loopback mount the sps.img file previously downloaded, which is an ISO9660 image file. Run the “setup” script inside:

```
# mount -o loop sps.img /mnt
```

```
# cd /mnt
```

```
# ./setup
```

2. During this procedure, you will hit Enter in most cases to accept default values and continue to the next screen. Note the following exceptions:
  - a. On the screen titled “High AvailabilityNFS” you may select “n” as in this particular eval guide we will not be creating a highly available NFS server cluster configuration.
  - b. If you have plans to create a highly available NFS service, adjust your response accordingly.
3. Towards the end of the setup script, you can choose to install a trial license key now, or later. We will install the license key in the next step, so you can safely select “n” at this point
4. In the final screen of the “setup” select the DataKeeper from the list displayed on the screen.
5. The following RPMs should be installed:
  - a. steeleye-lkAPA-<version>.noarch.rpm
  - b. steeleye-lkSQL-<version>.noarch.rpm
  - c. steeleye-lkDR-<version>.noarch.rpm
6. Un-mount the Distribution Enabling disk image:

```
# cd /root
```

```
# umount /mnt
```

## Install the Evaluation License Keys

The last phase of the setup process installs the licensing keys. You must install the evaluation license key file (“.lic”) that you downloaded with your evaluation software before starting the LifeKeeper for Linux.

1. To install your trial license key, run the “lkkeyins” command on both Primary and Secondary Server. This command is located at /opt/LifeKeeper/bin/lkkeyins . Example:

```
# /opt/LifeKeeper/bin/lkkeyins <path_to_license/<filename>.lic
```

2. Validate your license keys were installed via the /opt/LifeKeeper/bin/lklicmgr command

```
# /opt/LifeKeeper/bin/lklicmgr
```

License File: 20101230.lic

Product	Type	Expiry

LifeKeeper for Linux	Eval	27 Mar 2013 (87 days)
Apache Recovery Kit	Eval	27 Mar 2013 (87 days)
SIOS Data Replication ARK	Eval	27 Mar 2013 (87 days)
MySQL Recovery Kit	Eval	27 Mar 2013 (87 days)

## Start the LifeKeeper for Linux

1. Start:

```
# /opt/LifeKeeper/bin/lkstart
```

## 12.8.8. Configure the Cluster – Apache

### Primary Server

Complete the following steps on the primary server to configure the cluster:

- Create TCP Communication (Comm) Path(s)
- Verify the Communication (Comm) Path(s)

Before you begin, SIOS recommends at least two TCP communications paths between each server within the cluster to each remote server for heartbeat redundancy.

 **Important:** Supported configurations require that you define redundant comm. paths, so that the failure of a single communication line will not cause a split brain where resource hierarchies may come in-service on multiple servers simultaneously.

### Access the LikeKeeper GUI

The LifeKeeper Graphical User Interface (GUI) is a Java based application that can be run as a native Linux application.

The LifeKeeper GUI is based on Java RMI with callbacks. Hostnames must be resolvable or you may receive a Java 115 or 116 error.

1. Verify that both short and fully qualified hostnames of all cluster nodes resolve to the proper locations

```
# ping LinuxPrimary
```

```
# ping LinuxPrimary.domain.com
```

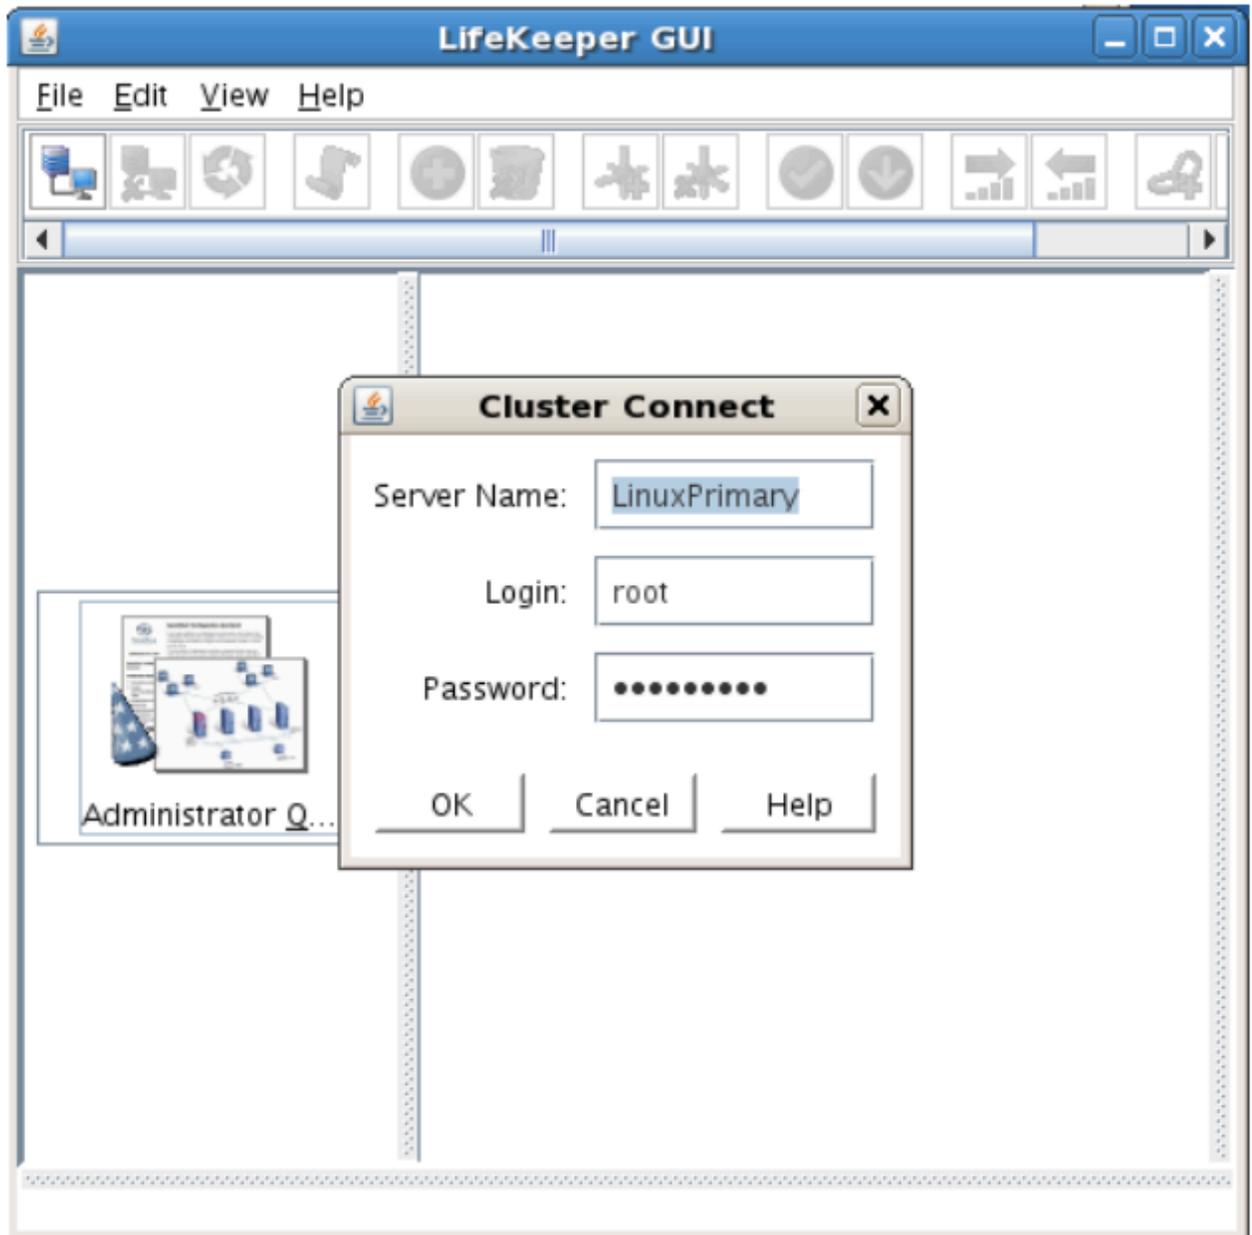
```
# ping LinuxSecondary
```

```
# ping LinuxSecondary.domain.com
```

2. To start the LifeKeeper Linux GUI Application:

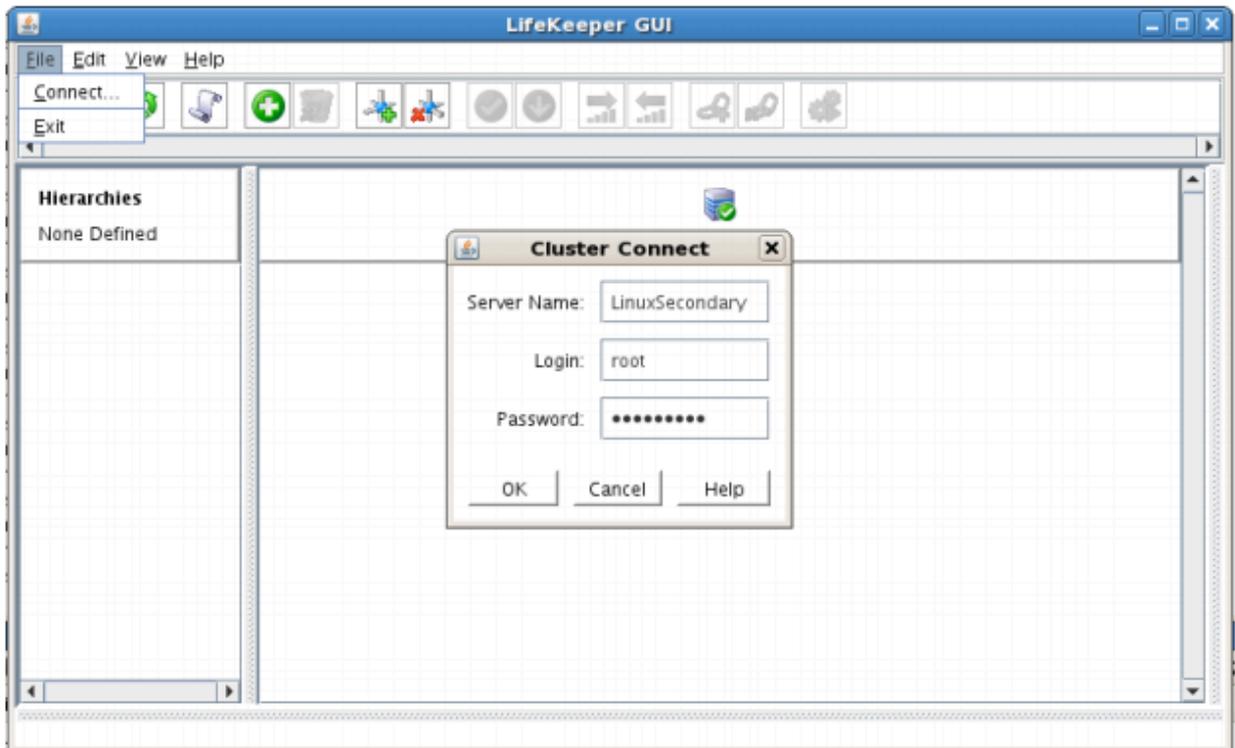
```
a. /opt/LifeKeeper/bin/lkGUIapp &
```

3. Enter the name of the server you wish to connect to (this field will be populated with the name of the server you are on, if you are running the GUI from a server with LifeKeeper installed) along with your root credentials and click OK.



## Create Communication (Comm) Paths

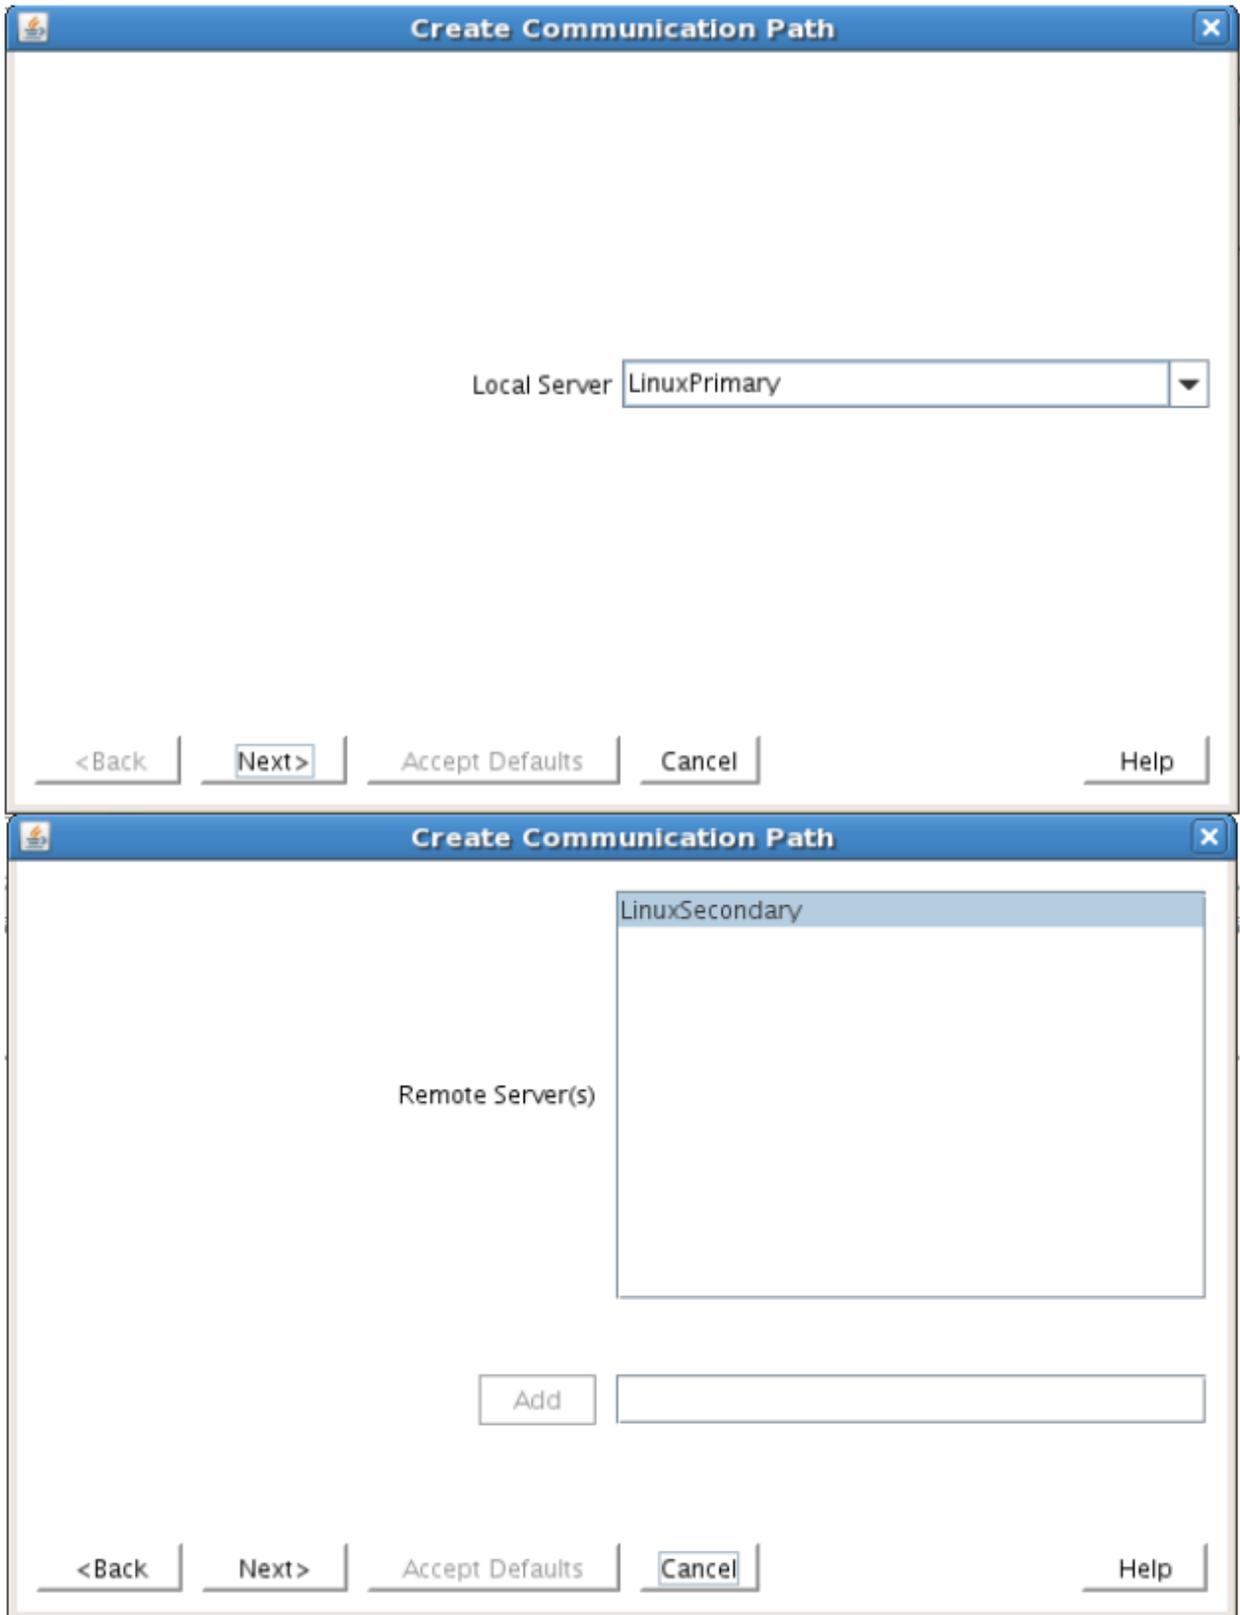
4. Within the LifeKeeper GUI, from the File menu, select Connect. Enter the name of your Secondary server, login and password when the Cluster Connect window displays.



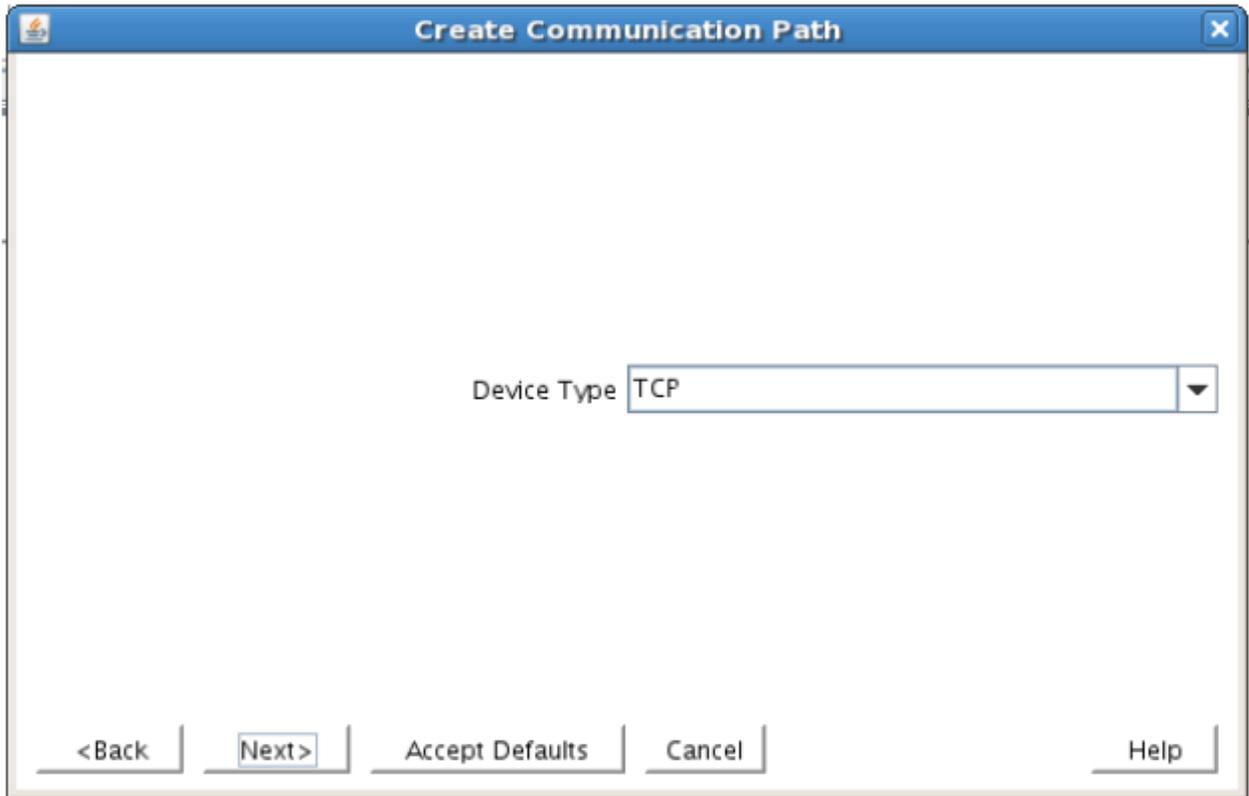
5. Within the LifeKeeper GUI, click the Create Comm Path button on the toolbar. You can also right click one of the servers and click Create Comm Path from the pop-up menu as well.



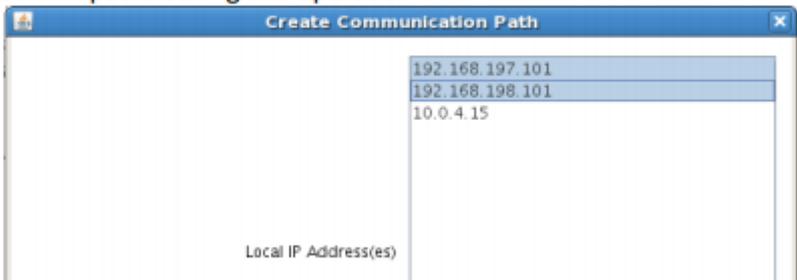
6. Select your Local and Remote Server(s) from the list box. If a server does not appear in the list box, you may enter it by typing its name and clicking the Add Server button. When using the Add Server procedure, you must make sure that the computer names for both network interfaces on the servers respond correctly when you ping them (from all of the partner server(s)) using the **ping -a IP ADDRESS** syntax. If they do not, this must be corrected prior to continuing. Click Next.

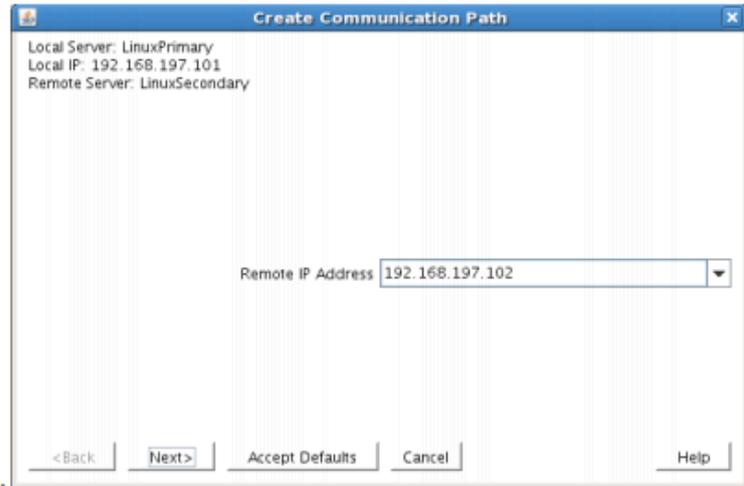


7. Select TCP for Device Type and Click Next.



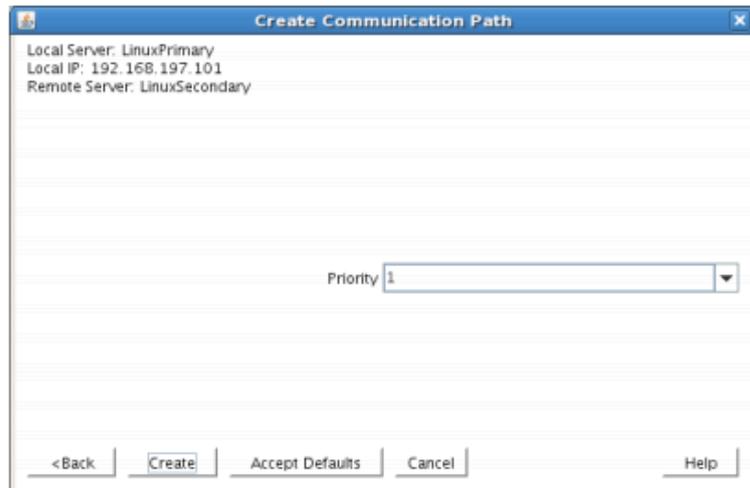
- 8. Provide all the required information and click Next for the following series of dialog boxes. For each field in the dialog box you can click Help for further information or refer to the table below for an explanation or recommendation

Field	Tips
<b>For TCP/IP Comm Path...</b>	Choose the IP address to be used by the local server for this comm path. Select both interfaces so that the wizard creates multiple comm. paths during this operation
Local IP Address	
Remote IP Address	Choose the IP address to be used by the remote server for this comm path

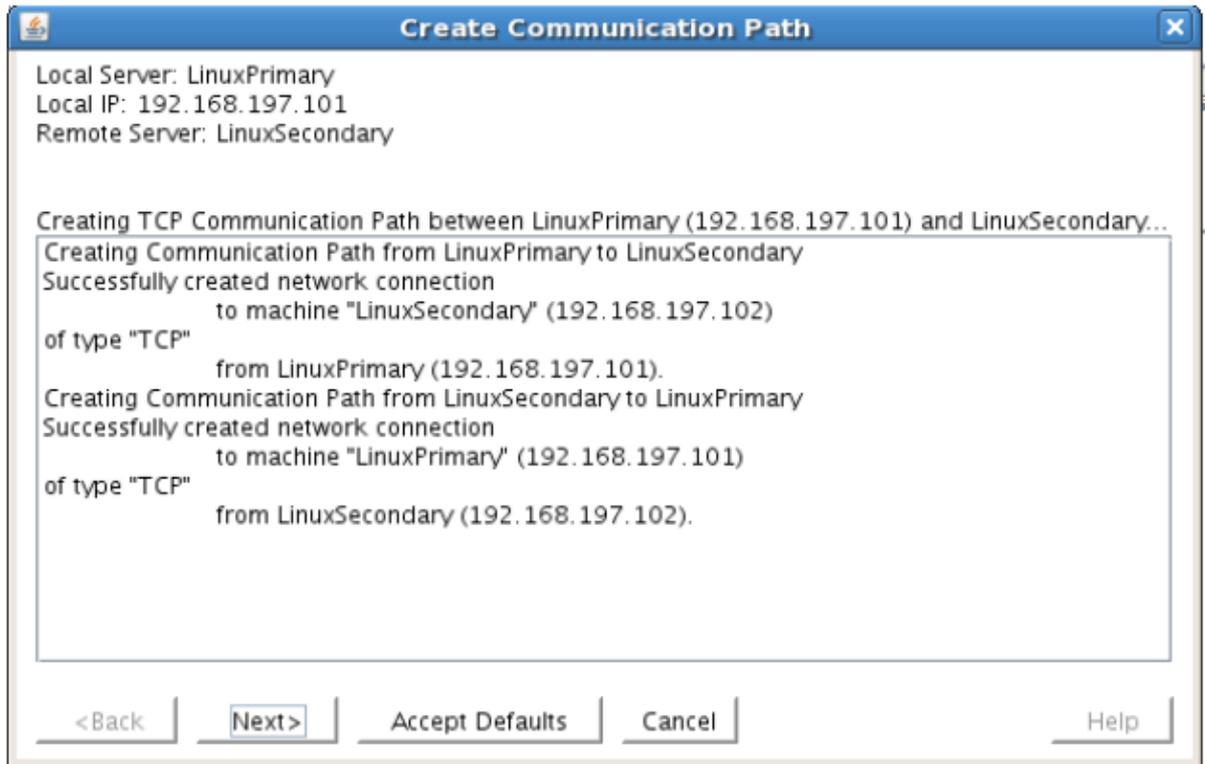


Enter the priority for the comm path on the local server. The priority will be used to determine the order that the comm paths between two servers will be used. Priority 1 is the highest; 99 is the lowest.

Priority



9. After entering data in all the required fields, select Create. A message will display indicating the network communication path is successfully created. Click Next.



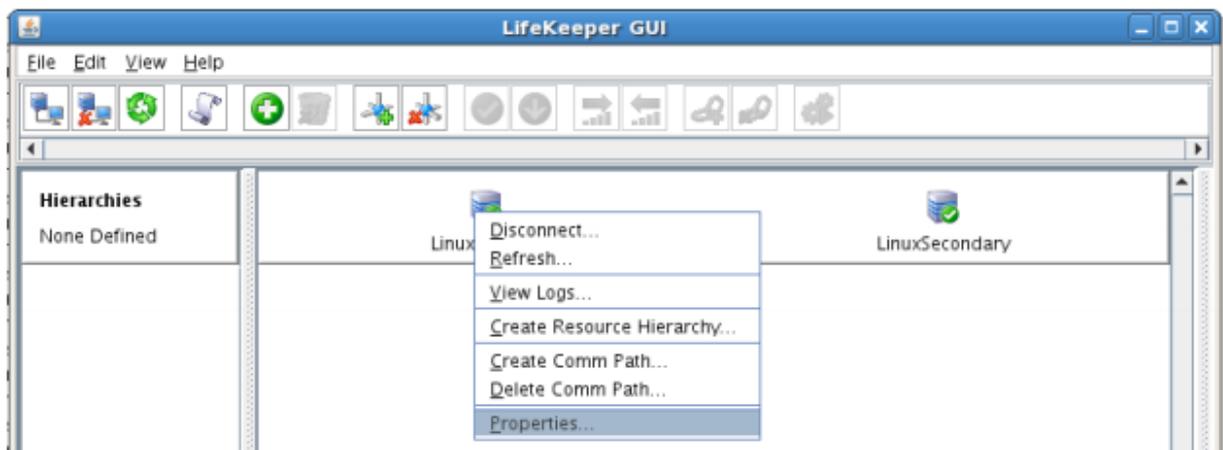
If you selected multiple Local IP Addresses or multiple Remote Servers and the Device Type was set to TCP, then the procedure will return you to the setup wizard the next Comm Path.

10. Select Done in the last dialog box.

Repeat this process until you have defined all the communication paths you plan to use. SIOS strongly recommends that you define at least two communication paths for redundancy.

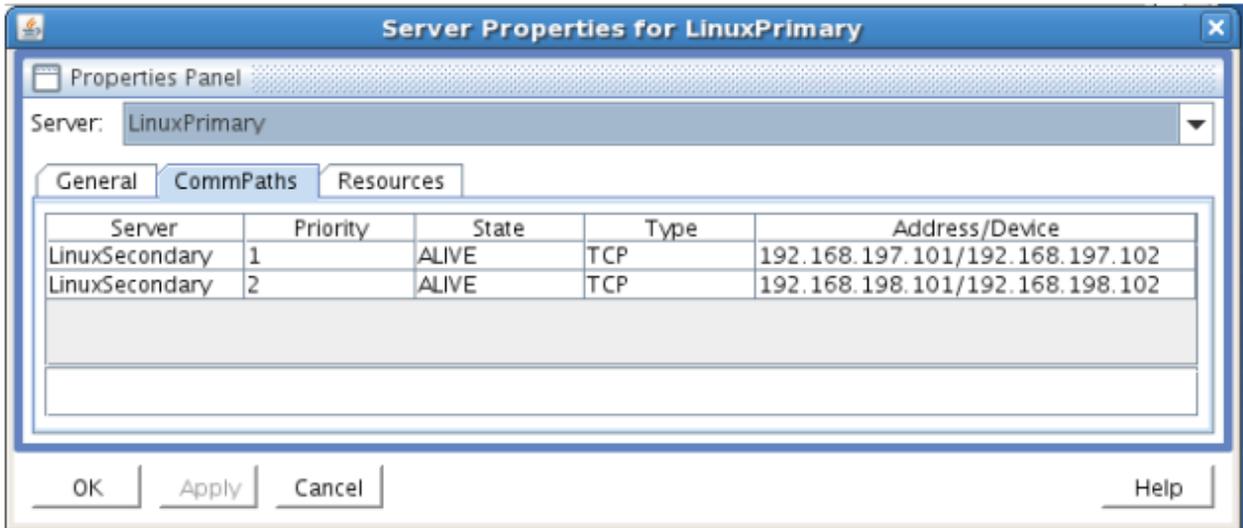
## Verify the Communications Paths

1. Verify that the communications paths are configured properly by viewing the Server Properties dialog box. From the LK GUI, select Edit, Server, Properties and then the Comm Paths tab.



2. Note the State displayed is ALIVE. You can also check the server icon in the right, main pane of

the GUI. If only one comm path has been created, the server icon shows a yellow warning icon on the server icon, indicating that one comm. path is ALIVE, but there is no redundant comm path. The server icon will display a green heartbeat checkmark when there are at least two comm paths configured and ALIVE.



## Create the LifeKeeper Hierarchy

### Create and Extend an IP Resource

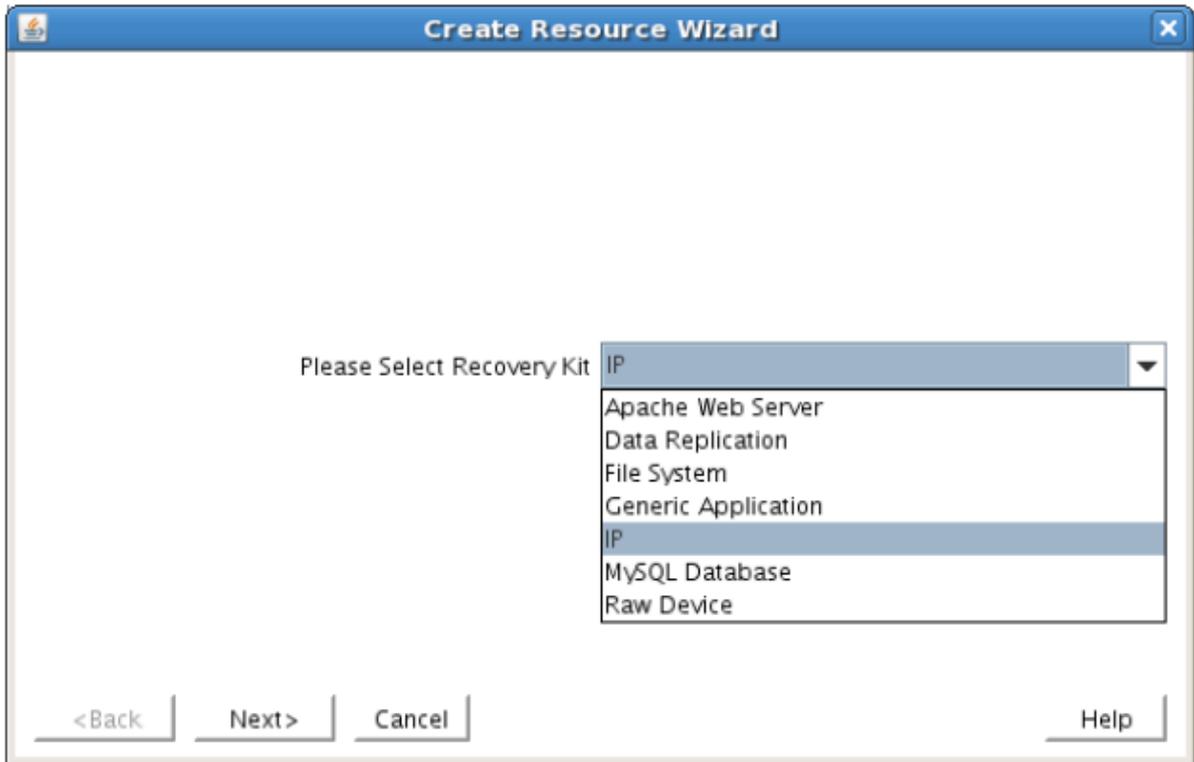
In LifeKeeper, create an IP resource and extend it to the secondary server by completing the following steps. This Virtual IP will have the ability to move between cluster nodes along the application that depends on it.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

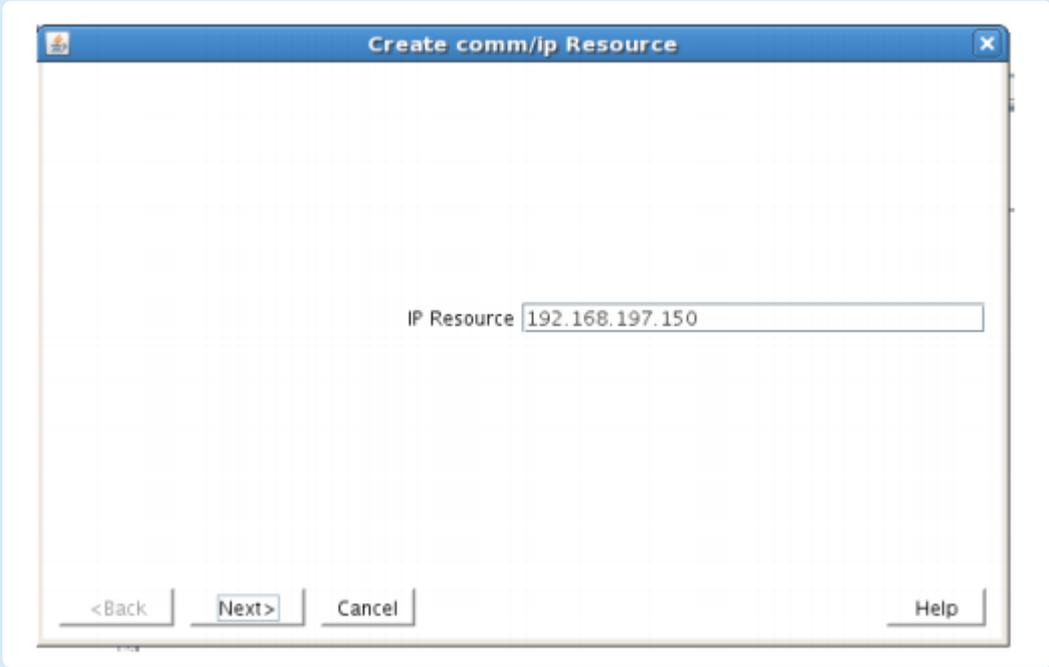
2. Select IP Address and click Next.

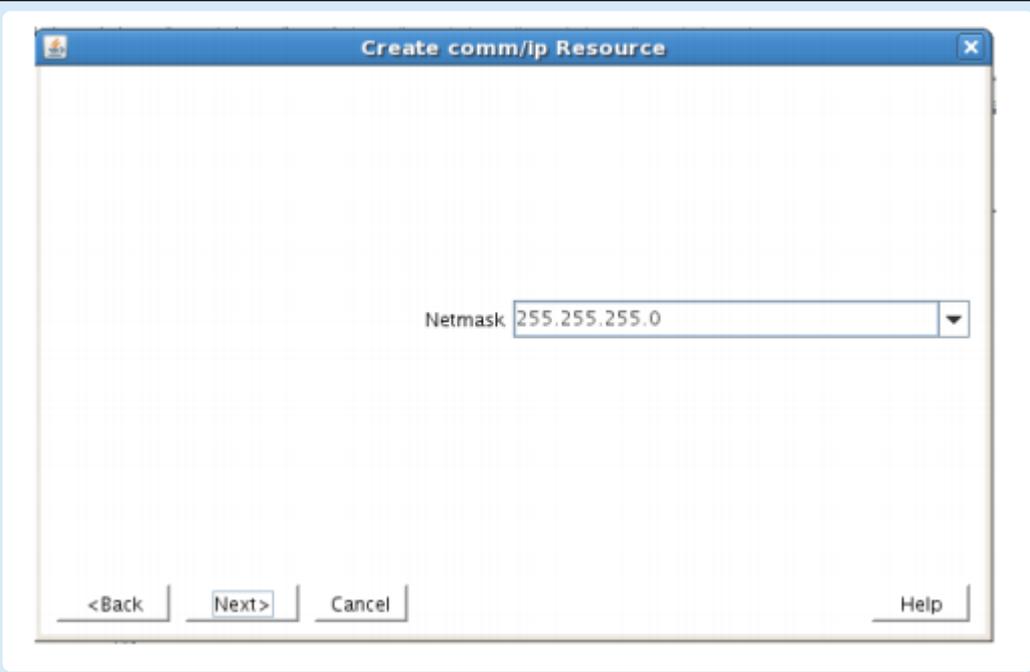


3. Enter the appropriate information for your configuration. The table below contains a list of the fields that display and additional information to assist you as you complete this procedure. Recommended values are also show below. You can also click the Help button for further information. Press Next to continue after entering the required information.

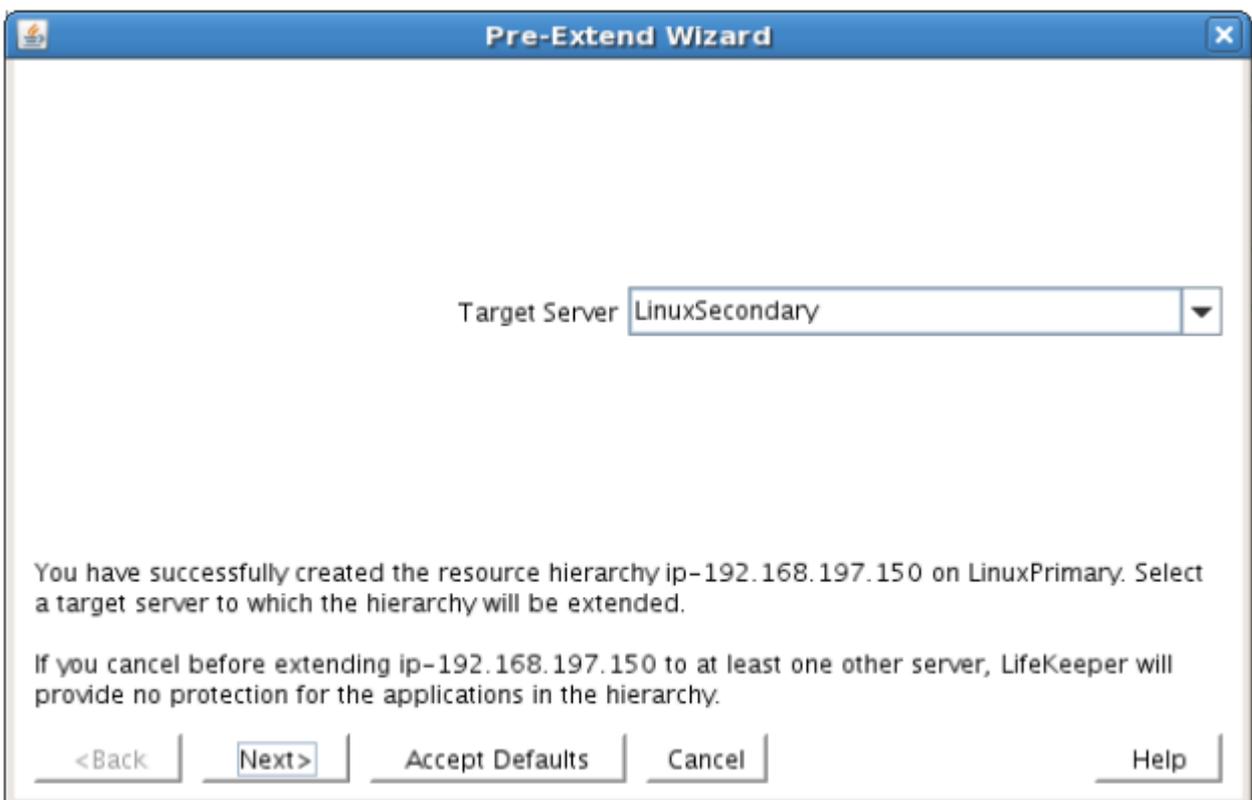
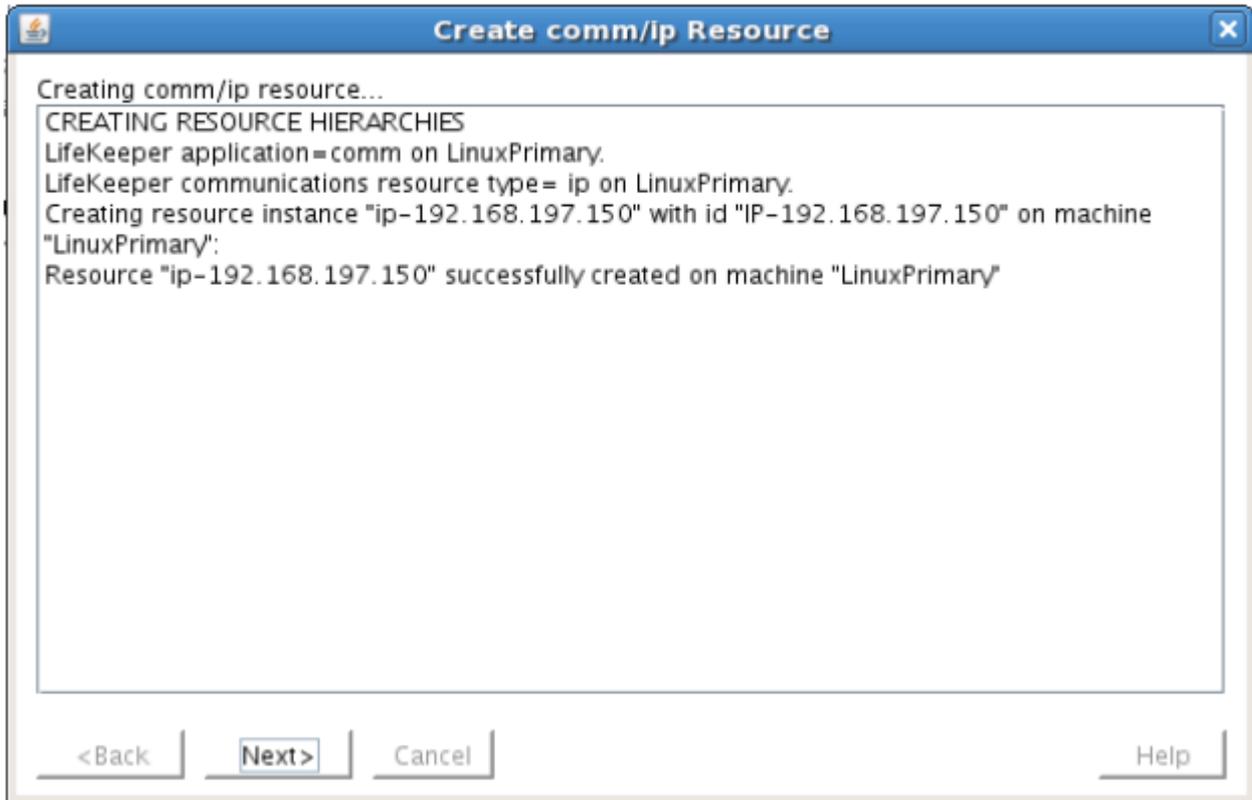
## IP Creation Field Definitions

Field	Tips
Resource Type	Select IP Address as the resource type and click Next.
Switchback Type	Select Intelligent and click Next.
Server	Select the Server where the IP resource will be created. Select your Primary server and click Next.
IP Resource	<p>Enter the virtual IP information and click Next</p> <p>Example <b>192.168.167.151</b></p> <p><b>Note</b> This is an IP address that is not currently in use anywhere on your network. This is the address that all clients will use to connect to the protected resources.</p> <p>In this configuration example, we will be protecting two (2) virtual IPs.</p> <p>First we will protect 192.168.197.150, which our Apache webserver will</p>

	<p>use.</p>  <p>The second time through this wizard we will protect 192.168.197.151, which will be used by MySQL</p>
<p>Netmask</p>	<p>The IP subnet mask that your TCP/IP resource will use on the target server. Any standard netmask for the class of the specific TCP/IP resource address is valid</p> <p>In our sample configuration 255.255.255.0 is used for a subnet mask on both networks.</p> <p><b>Note:</b> The subnet mask you choose, combined with the IP address, determines the subnet that will be used by the TCP/IP resource and should be consistent with the network configuration.</p>

	
<p>Network Connection</p>	<p>This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. Select the correct NIC and click Next.</p>
<p>IP Resource Tag</p>	<p>Accept the default value and click Next. This value only affects how the IP is displayed in the GUI. The IP resource will be created on our Primary server.</p>

- LifeKeeper will create and validate your resource. After receiving the message that the resource has been created successfully, click Next when the following dialog box appears so that you can complete the process of Extending the IP Resource to our Secondary server, below.



Extending the IP resource will start automatically after you have finished creating an IP address resource if you clicked Next in the dialog box displayed above. You can also start this from an existing IP address resource by right clicking on the active resource and selecting Extend Resource Hierarchy.

Refer to the table below to complete the Extend IP Resource procedure.

Field	Recommended Entries or Notes
Switchback Type	Leave as "intelligent" and click Next
Template Priority	Leave as default (1)
Target Priority	Leave as default (10)
Network Interface	This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. The correct physical NIC should be selected by default. Please verify and then click Next
IP Resource Tag	Leave as default.
Target Restore Mode	Select Enable and click Next.
Target Local Recovery	Select Yes to enable Local Recovery for the SQL resource on the Target server.
Backup Priority	Accept the default value.

- After receiving the message Hierarchy extend operations completed, click Finish and then click Done
- Your IP resource (192.168.197.151) is now fully protected and has the ability to "float" between cluster nodes as needed. Looking at the LifeKeeper GUI you will notice that the IP resource is Active on the Primary cluster node and Standby on the Secondary cluster node



### Create a Second IP Resource

Repeat the procedure above to protect a 2nd IP resource.

This second time, protect 192.168.197.151, which is the IP address our MySQL database will later use.

As a result, your LifeKeeper GUI will display as follows, with both IP resources Active and protected on the Primary cluster node:



### Create a Mirror and Begin Data Replication

In this section we will setup and configure the Data Replication resource, which be used to synchronize our Apache Webserver’s data between cluster nodes. The data we will replicate resides in the /var/www partition on our Primary cluster node

Please note:

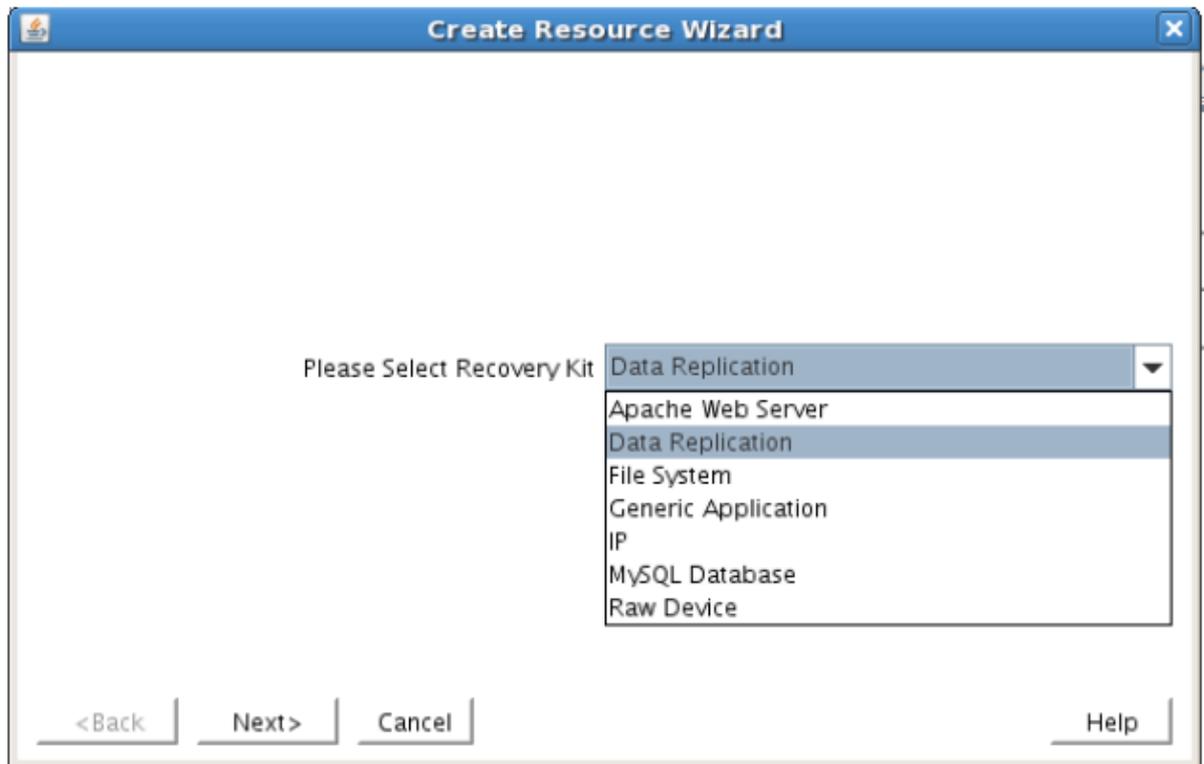
- The source volume to be replicated must be mounted on the Primary server
- The target volume, which will received replicated data, must NOT be mounted on the Secondaryserver.
- The target volume’s size must equal to or larger than the size of its source volume.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

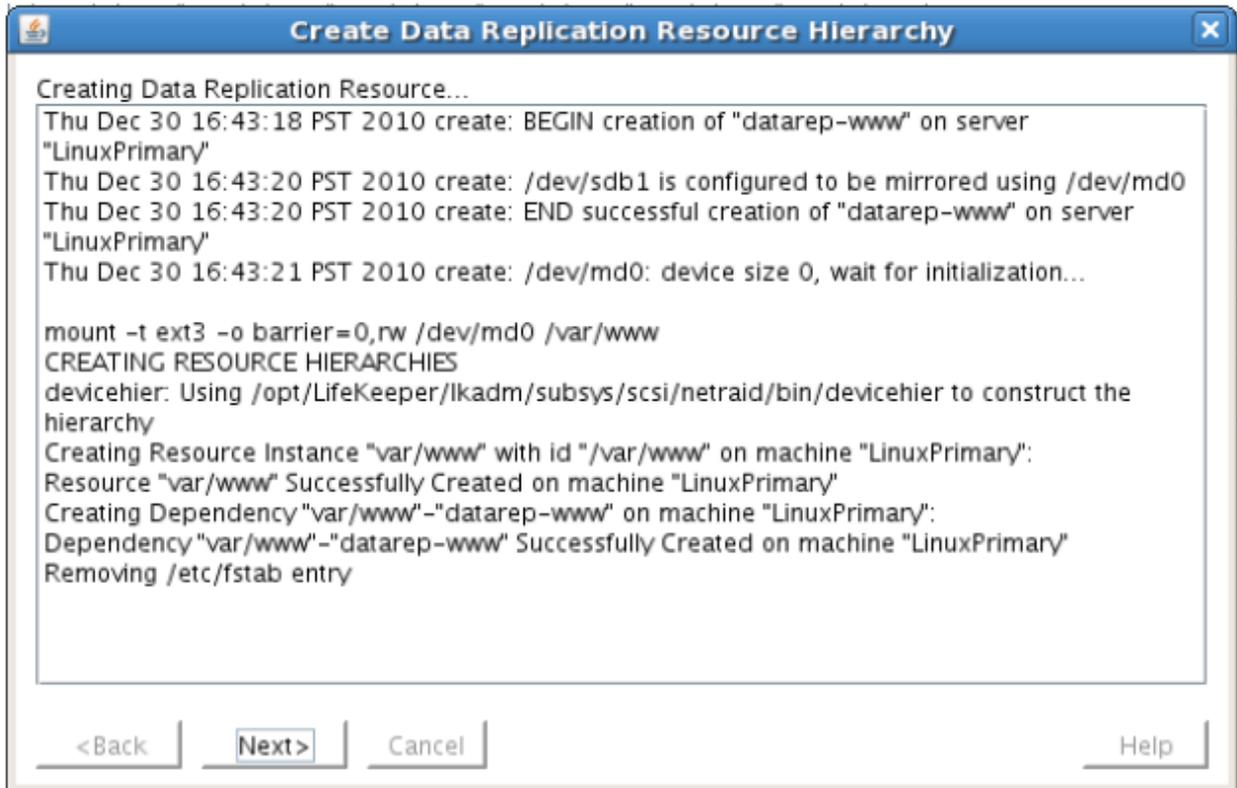
2. Select Data Replication and click Next.



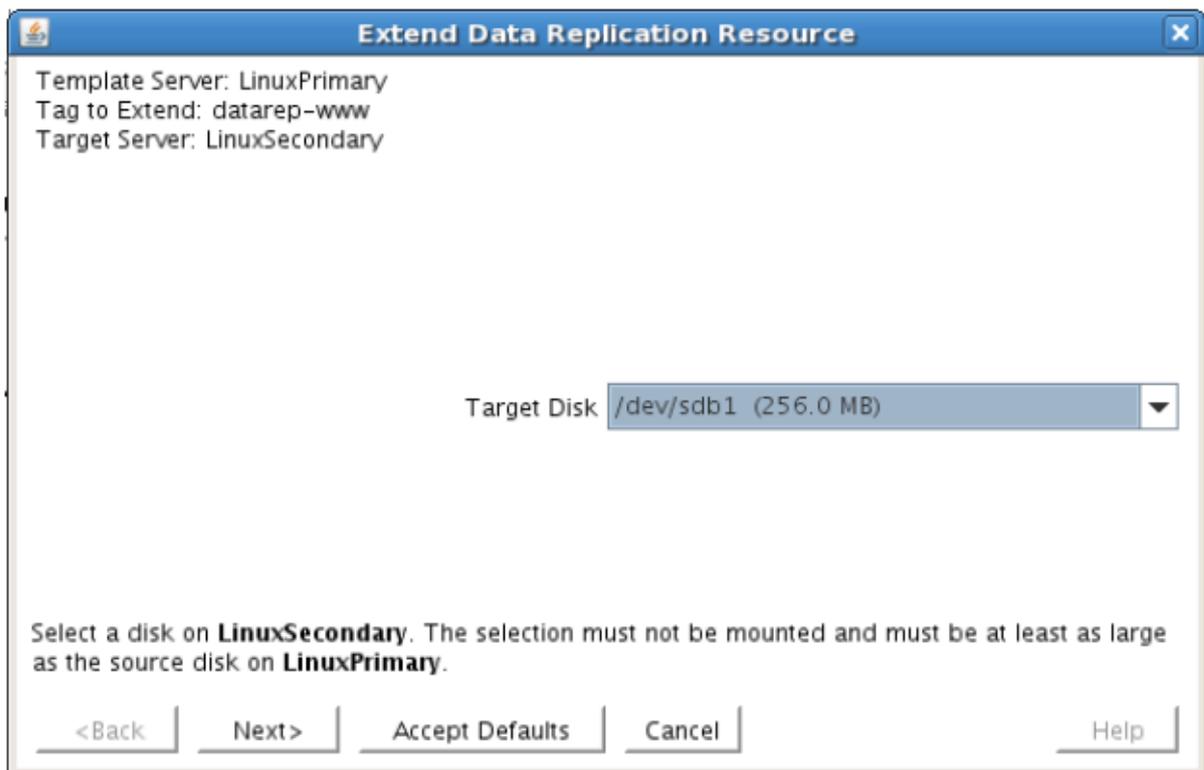
3. Follow the Data Replication wizard, and enter the following values:

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node, i.e. Mirror Source)
Hierarchy Type	Select: <b>“Replicate Existing Filesystem”</b>
Existing Mount Point	At this step you will select the mounted partition to replicate. In our example, select <code>“/var/lib/mysql”</code>
Data Replication Resource Tag	Leave as default
File System Resource Tag	Leave as default
Bitmap File	Leave as default (Note: if using high speed SSD storage you will want to create a small partition and use it for bitmap placement, i.e. <code>/bitmaps</code> )
Enable Asynchronous Replication	Leave as default (Yes)

4. Click Next to begin creation of the Data Replication resource hierarchy. Status will be displayed in the GUI as follows:

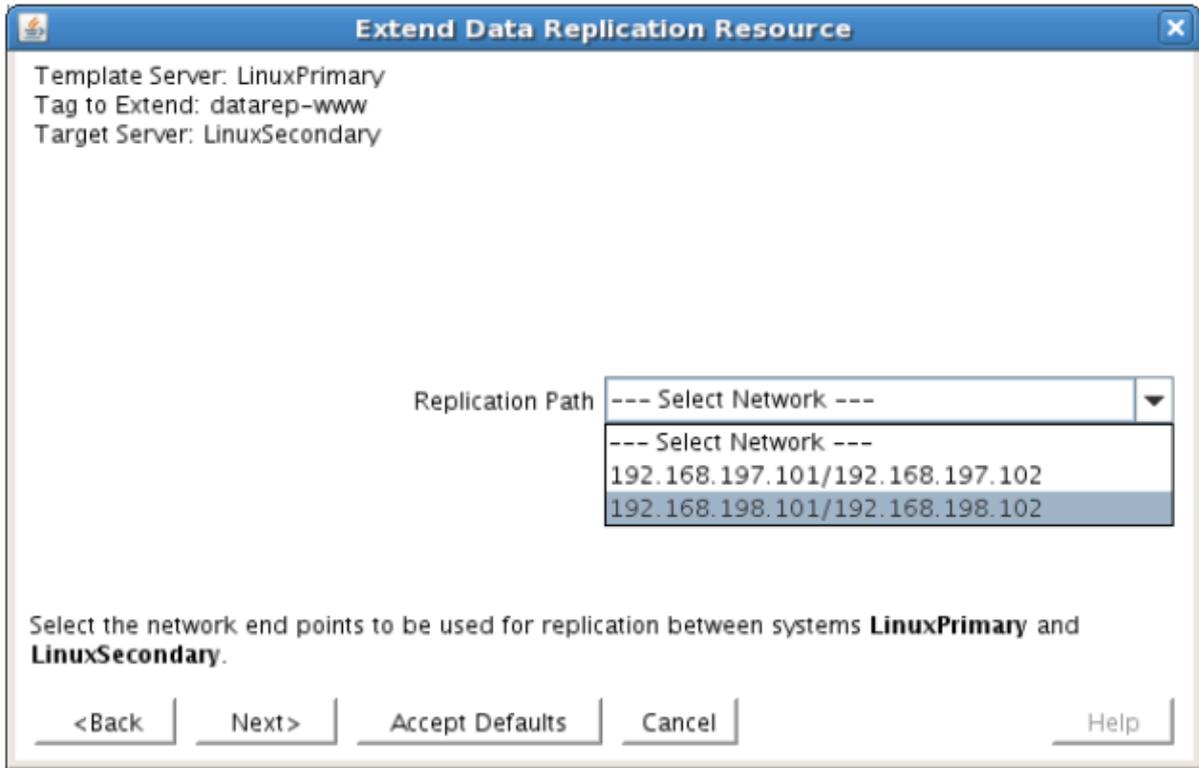


5. Click Next to begin the process to Extend the Data Replication Resource. Select all default settings. When it asks for the target disk, select a free partition on your Target server which is the same size (or greater) than the Source Volume we are replicating. This partition should NOT be mounted on the Target system.



6. Continue through the wizard, and you will be prompted to select the network you would like

replication to take place over. In general, it's a best practice to separate your user/application and your replication traffic. In our example setup we will replicate over our backend network, 192.168.198.X



- 7. Click Next and continue through the wizard. Once completed, your resource hierarchy will look as follows

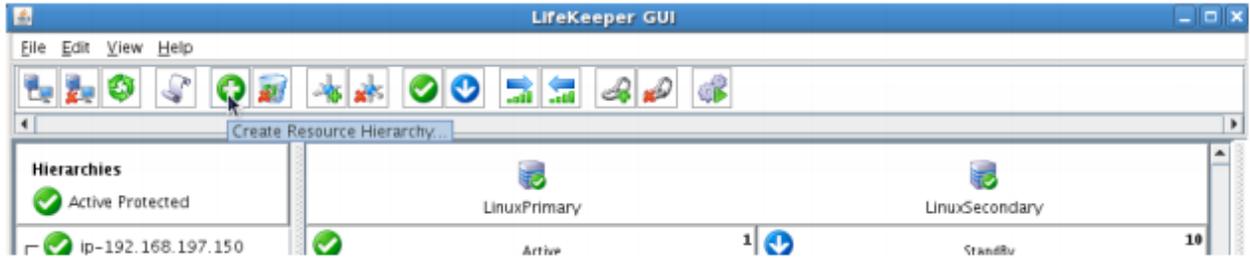


### Create the Apache Hierarchy

In this section we will create an Apache resource hierarchy on the primary server and extend it to the backup server. This step will create a dependency on the IP resource created in previous the step.

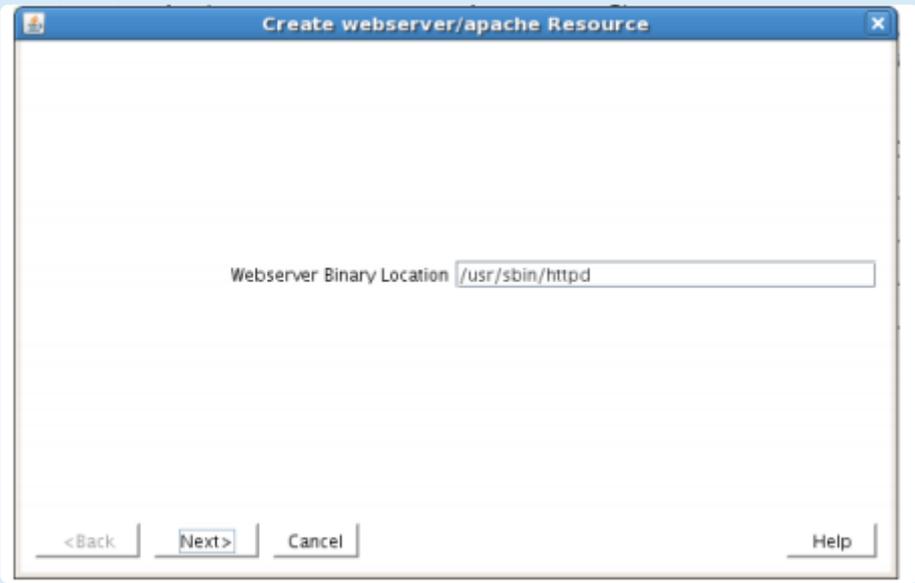
**Important:** The Apache web server should not be running at this time.

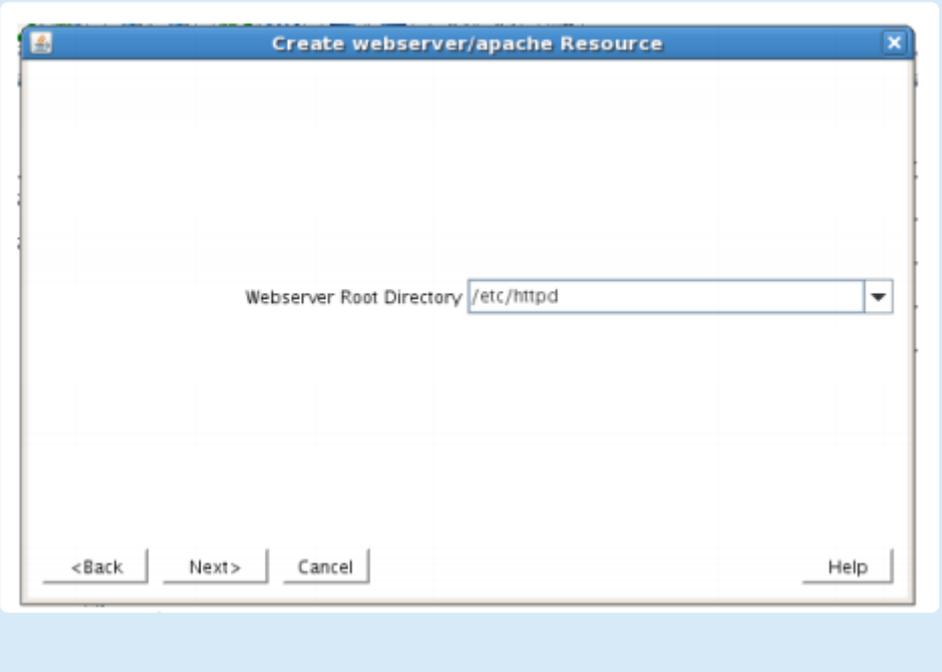
1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



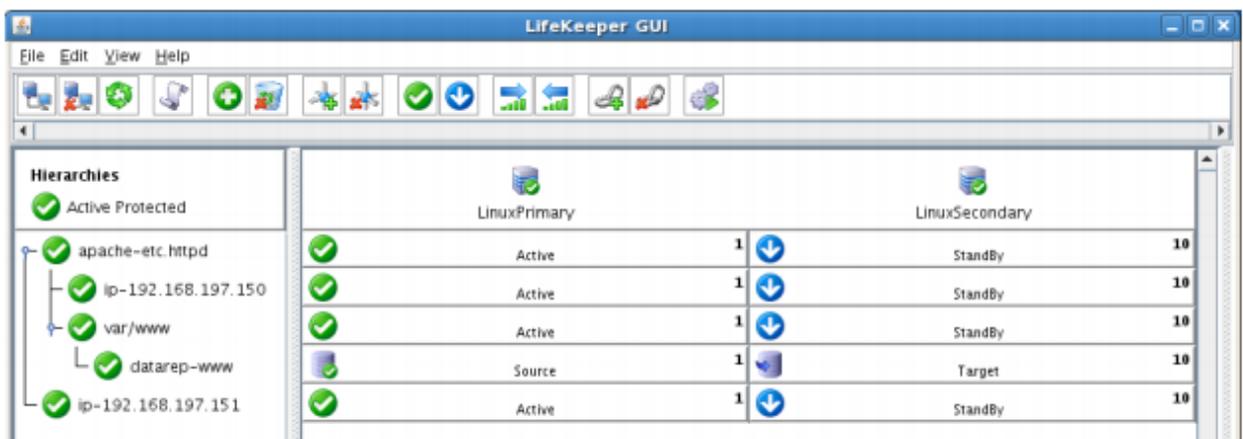
The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

2. Select Apache Web Server and click Next.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node, i.e. Mirror Source)
Webserver Binary Location	/usr/sbin/httpd (assumes a standard apache config) 
Webserver Root Directory	/etc/httpd (assumes a standard apache config)

	
Root Tag	Leave as default

4. Click “Create” to begin resource hierarchy creation on the primary server. Once complete, click “Next” to extend this resource to the secondary server.
5. During the Extend Resource wizard, leave all settings as default.
6. Note: during the resource creation process, LifeKeeper extracted the existing configuration of the existing Apache webserver, and identified that it depends on the IP resource (192.168.197.150) and the Data Replication resource (/var/www) that were created in previous steps. These resources now appear underneath the newly created Apache resource, to indicate the dependency relationship.



## Create the Shared Filesystem Resource Hierarchy

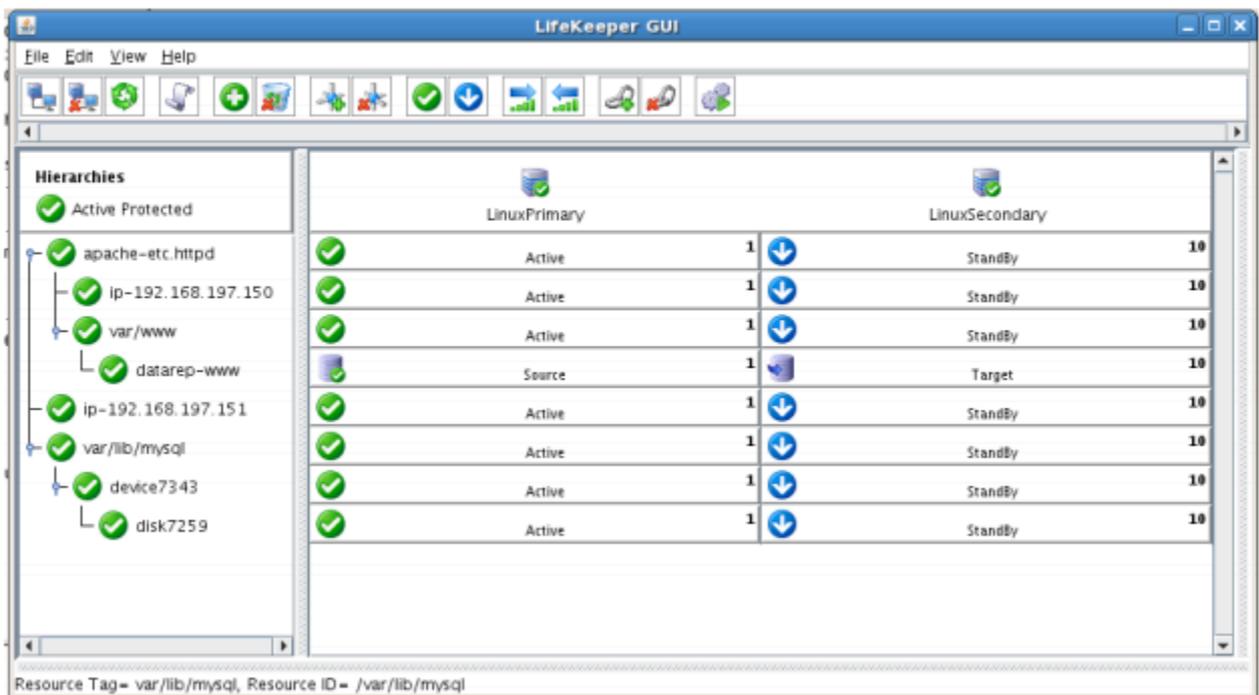
Create a Filesystem resource to protect the shared iSCSI filesystem and make it high available between cluster nodes. LifeKeeper for Linux leverages SCSI Persistent Group Reservations (PGR) to lock the LUN, ensuring that only the active cluster node for the storage resource can access it.

**❁ Important:** At this point, the shared iSCSI LUN needs to already be mounted on the Primary Server. It should NOT be mounted on the Secondary Server. See section titled “Configure iSCSI initiator, discover and login to iSCSI target” above to review the steps involved.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.
2. Select File System and click Next.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node)
Mount Point	Select /var/lib/mysql. Note that LifeKeeper scans the system for LUNS that are sharable between cluster nodes. The list of possible shared LUNS is presented automatically in this step of the wizard.

4. Select Create Instance to define this resource hierarchy on the Primary Server
5. Click Next to Extend the File System Resource to the Secondary Server
6. In the Extend Wizard, select “Accept Defaults”
7. As a result the File System resource is now protected on both cluster nodes. Click Finish to exit the Extend wizard.
8. Your resource hierarchy should look as follows:



## Create the MySQL Resource Hierarchy

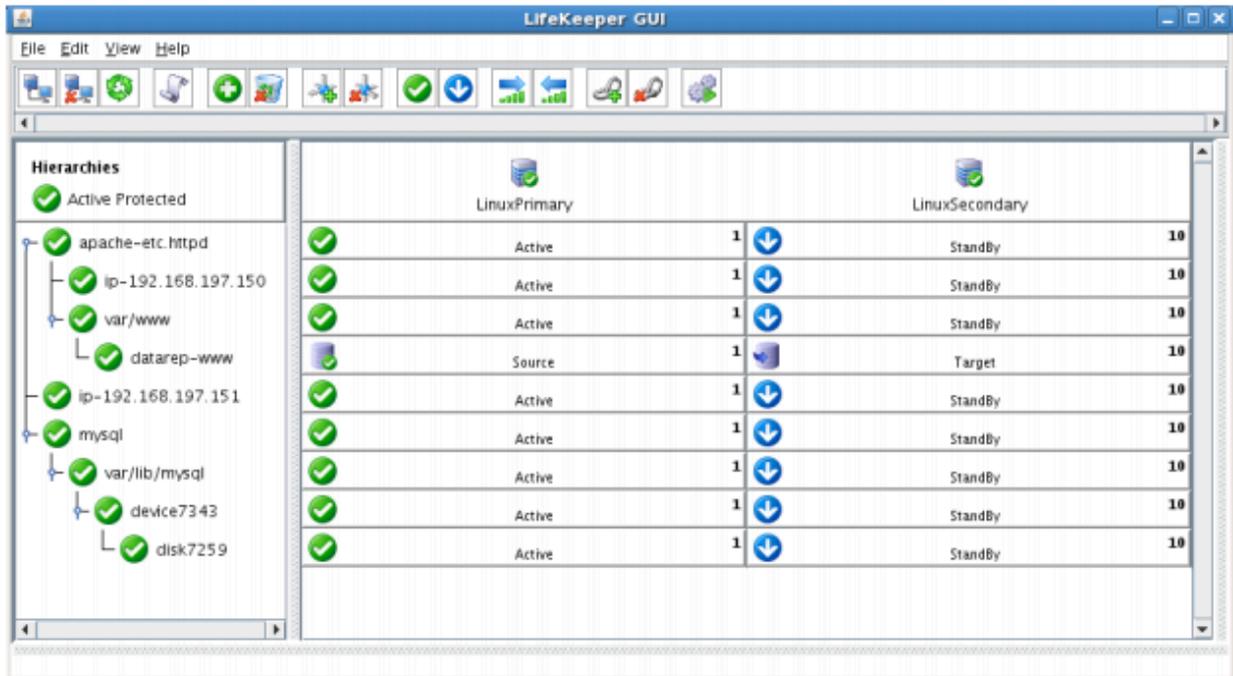
Create a MySQL resource to protect the MySQL database and make it high available between cluster nodes.

 **Important:** At this point, MySQL needs to be already running on the Primary Server. It should NOT be running on the Secondary Server. See section titled “Install, Configure, and Start MySQL” above to review the process to configure and start MySQL as needed.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.
2. Select **MySQL Database** and click Next.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node)
Location of my.cnf	Enter “/var/lib/mysql”. Note that earlier in the MySQL configuration process we created a my.cnf file in this directory.
Location of MySQL executables	Leave as default (/usr/bin) since we are using a standard MySQL install/ configuration in this example
Database tag	Leave as default

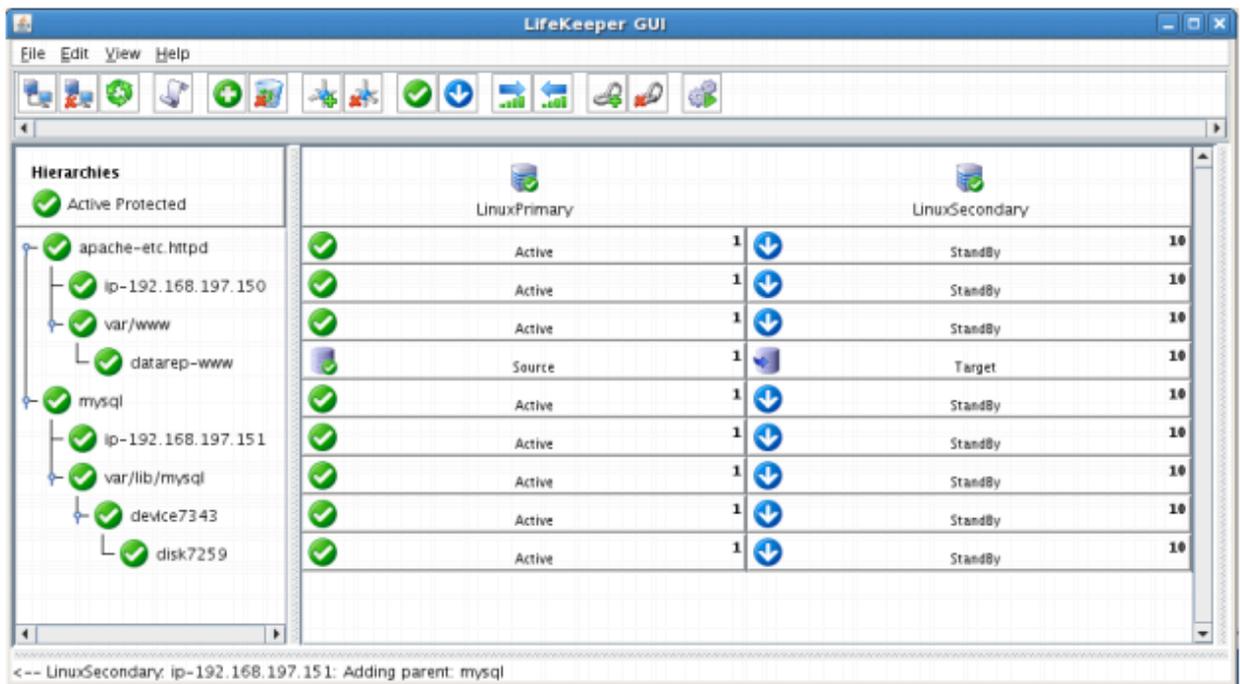
4. Select Create to define the MySQL resource hierarchy on the Primary Server
5. Click Next to Extend the File System Resource to the Secondary Server
6. In the Extend Wizard, select “Accept Defaults”
7. As a result the MySQL resource is now protected on both cluster nodes. Click Finish to exit the Extend wizard.
8. Note: LifeKeeper will automatically identify that the MySQL resource has a dependency on the FileSystem resource (/var/lib/mysql). The Filesystem Resource will appear underneath the MySQL resource in the GUI
9. Your resource hierarchy should look as follows:



## Create the MySQL IP Address Dependency

In this step will define an additional dependency: that MySQL depends on a Virtual IP (192.168.197.151) so that the IP address follows the MySQL database should it move. This IP (.151) is the IP the webserver will use to access the MySQL database.

1. From the LifeKeeper GUI toolbar, right-click on the “mysql” resource
2. Select “Create Dependency” from the right-click context menu
3. In the Child Resource Tag dropdown menu, select “ip-192.168.197.151”
4. Click Next
5. Click Create Dependency
6. Click Done
7. The Virtual IP address resource (192.168.197.151) will now appear underneath the MySQL resource in the LifeKeeper user interface. This ensures that resources move together, and are started/stopped in the proper order.
8. Your resource hierarchy should look as follows



At this point in the Evaluation, we have fully protected Apache, MySQL, and their dependent resources: IP addresses, and Storage, both shared and replicated.

# 12.8.9. Test Your Environment – Apache

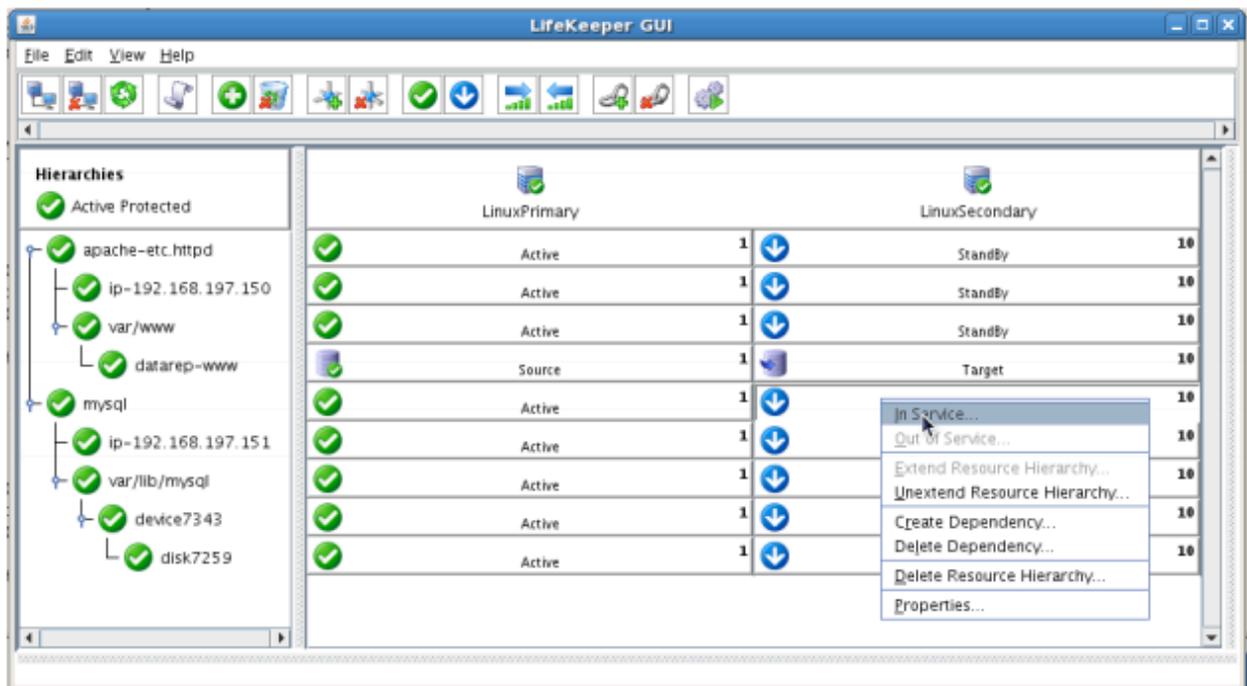
The following test scenarios have been included to guide you as you get started evaluating LifeKeeper for Linux. **Before beginning these tests, make sure the data replication resources are in the mirroring state.**

**Note:** For these test examples, the Primary Server is referred to as LINUXPRIMARY. The Backup or Secondary Server is referred to as LINUXSECONDARY.

## Manual Switchover of the MySQL Hierarchy to Secondary Server

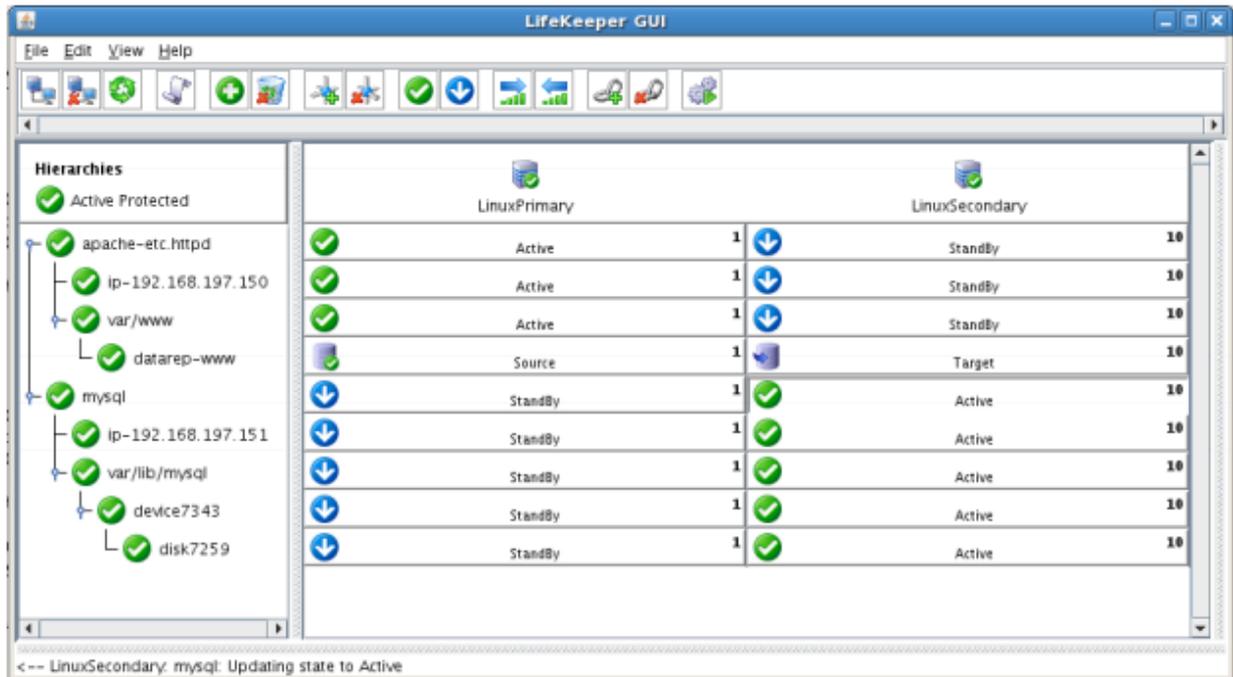
### Procedure:

- From the LifeKeeper GUI, right click on the MySQL resource on the Secondary Server (LINUXSECONDARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up



### Expected Result:

- Beginning with the MySQL resource, all resources will be removed from service on the Active Server (LINUXPRIMARY).
- Beginning with the dependent resources (IP and Shared Volume), all resources will be brought in service on LINUXSECONDARY
- At this point, we now have an “Active/Active” cluster because both cluster nodes are actively running resources.



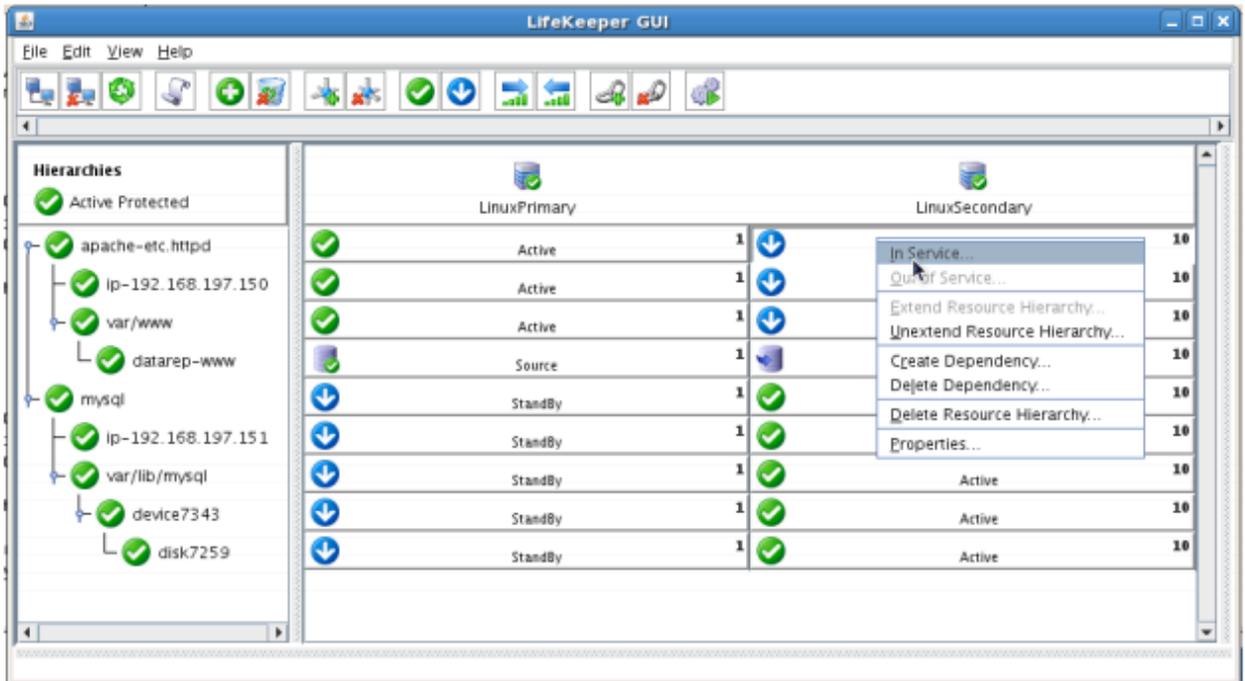
### Tests/Verification:

- Using the LifeKeeper GUI, verify that the MySQL and dependent resources are active on LINUXSECONDARY.
- Run “ifconfig -a” on LINUXSECONDARY to validate that the IP Address 192.168.197.151 is active on LINUXSECONDARY
- Run “df -h” to verify that the /var/lib/mysql shared iSCSI filesystem is mounted on LINUXSECONDARY
- Verify the MySQL services are running on LINUXSECONDARY by running “ps -ef | grep -i mysql”
- On LINUXSECONDARY run the following command to verify client connectivity to the MySQL database:
  - # mysql -S /var/lib/mysql/mysql.sock -u root -p
  - (enter password “SteelEye”)
- Verify that the SCSI reservation has properly locked the iSCSI LUN to the currently active cluster node. On LINUXPRIMARY, run “mount /dev/sdc1 /var/lib/mysql”. This should FAIL because LINUXPRIMARY does not own the SCSI reservation on this LUN.

## Manual Switchover of the Apache Hierarchy to Secondary Server

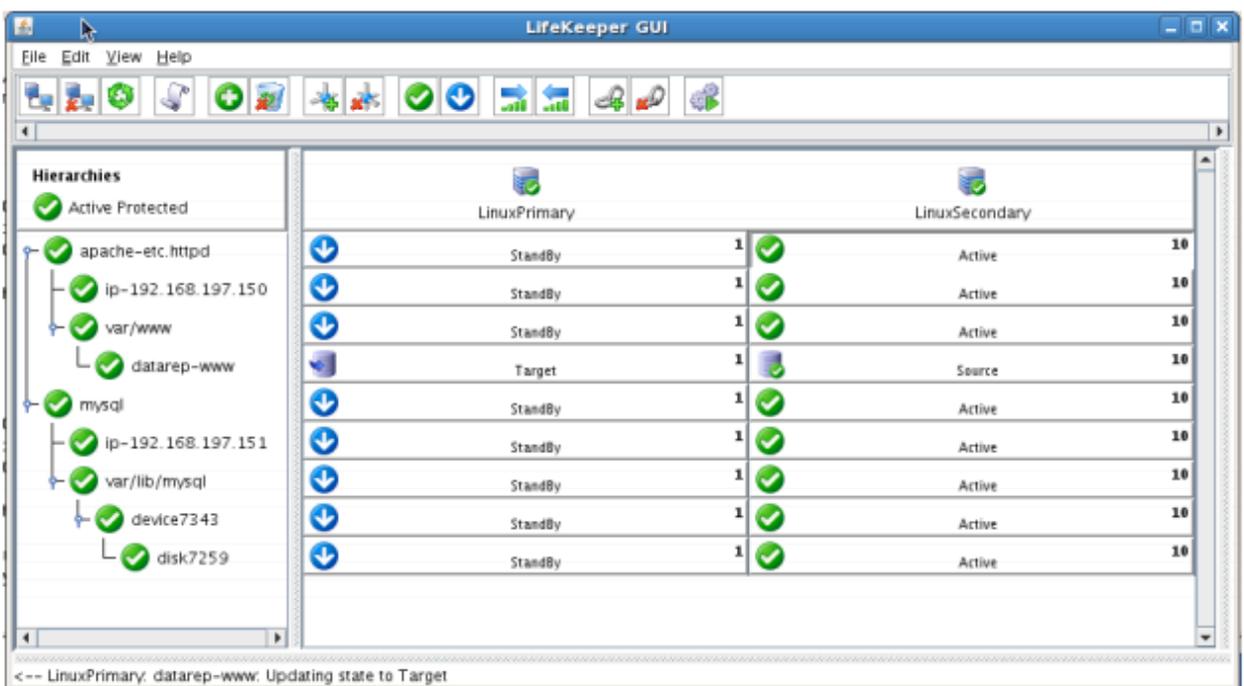
### Procedure:

- From the LifeKeeper GUI, right click on the Apache resource on the Secondary Server (LINUXSECONDARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up



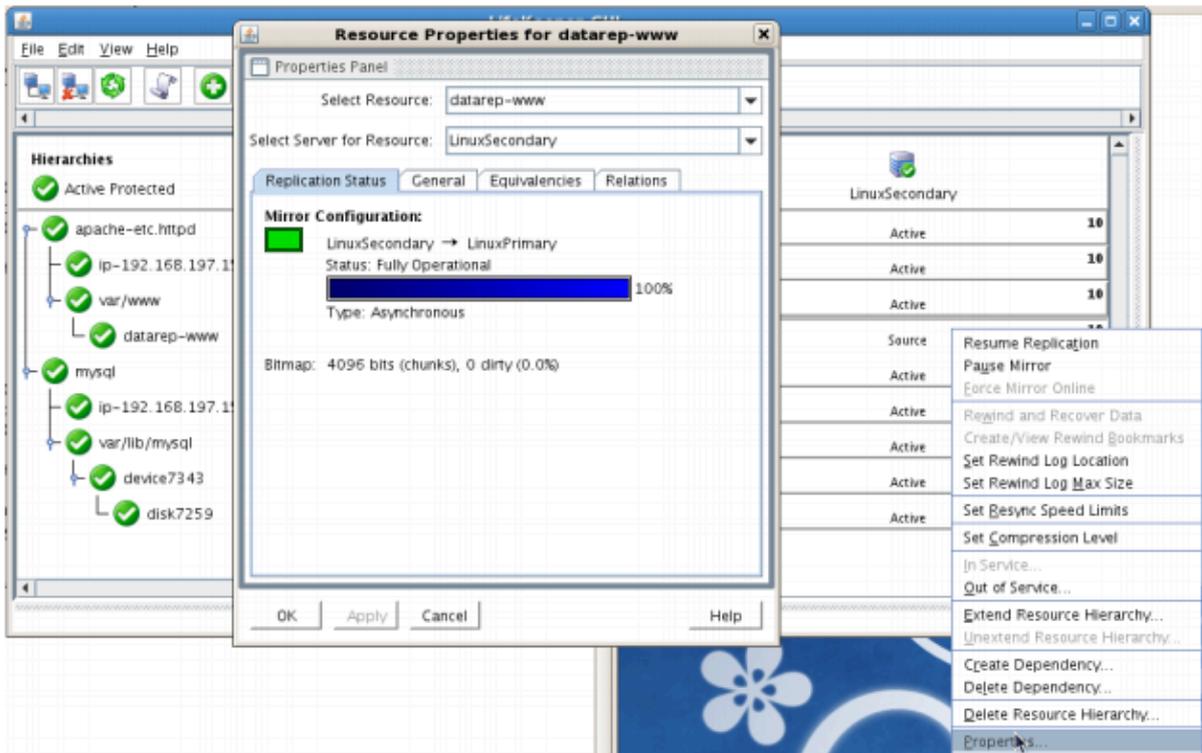
**Expected Result:**

- Beginning with the Apache resource, all resources will be removed from service on the Active Server (LINUXPRIMARY).
- Beginning with the dependent resources (IP and Replicated Volume), all resources will be brought in service on LINUXSECONDARY
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXSECONDARY -> LINUXPRIMARY
- At this point, we now back to an “Active/Passive” cluster because all services are now actively running on LINUXSECONDARY



**Tests/Verification:**

- Using the LifeKeeper GUI, verify that the Apache and dependent resources are active on LINUXSECONDARY.
- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the “datarep-www” resource and select Properties



- Run “ifconfig -a” on LINUXPRIMARY to validate that the IP Address 192.168.197.150 is active on LINUXPRIMARY
- Run “df -h” to verify that the /var/www replicated filesystem is mounted as an “md” device (example: /dev/md0) on LINUXPRIMARY
- Verify the Apache services are running on LINUXPRIMARY by running “ps -ef | grep -i httpd”
- Open a Web Browser to <http://192.168.197.150> and verify that it can successfully connect. The PHPInfo output should indicate that the system name is “LinuxPrimary

**PHP Version 5.1.6** 

<b>System</b>	Linux LinuxPrimary 2.6.18-194.26.1.el5 #1 SMP Tue Nov 9 12:54:40 EST 2010 i686
<b>Build Date</b>	Nov 29 2010 16:49:03
<b>Configure</b>	'./configure' '--build=i686-redhat-linux-gnu' '--host=i686-redhat-linux-gnu' '--target=i386-

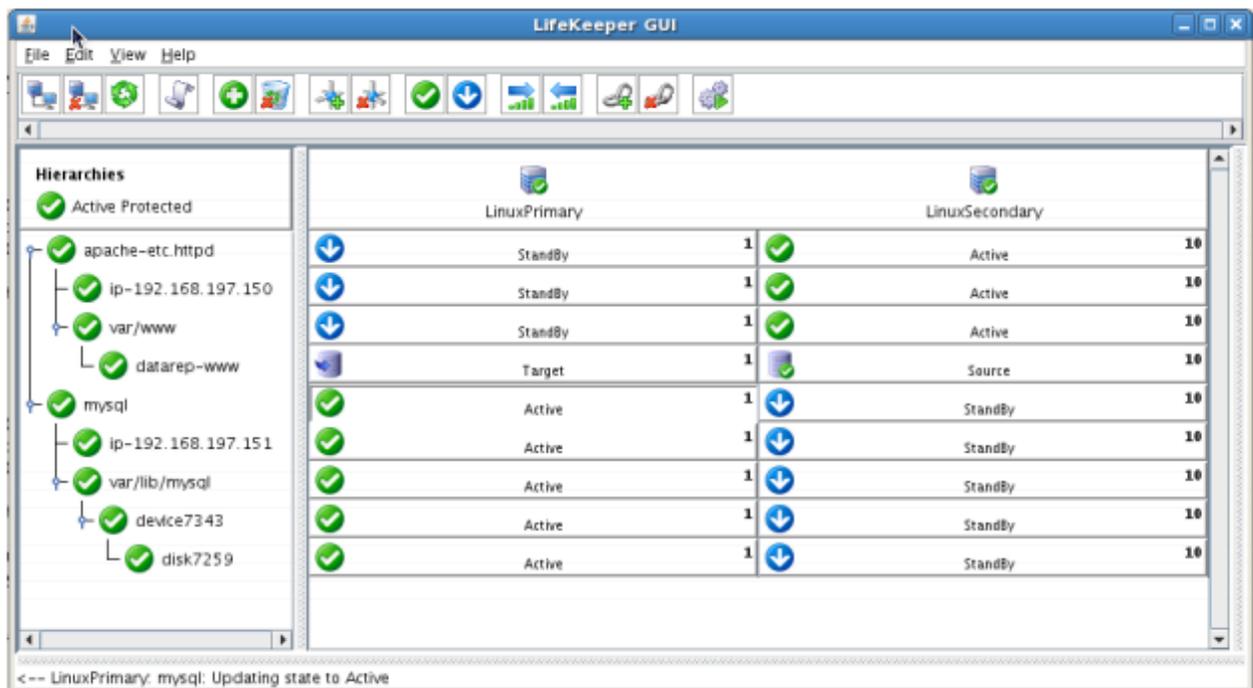
## Manual Switchover of the MySQL Hierarchy back to Primary Server

### Procedure:

- From the LifeKeeper GUI, right click on the MySQL resource on the Primary Server (LINUXPRIMARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up

### Expected Result:

- Beginning with the MySQL resource, all resources will be removed from service on the Active Server (LINUXSECONDARY).
- Beginning with the dependent resources (IP and Shared Volume), all resources will be brought in service on LINUXPRIMARY



### Tests/Verification:

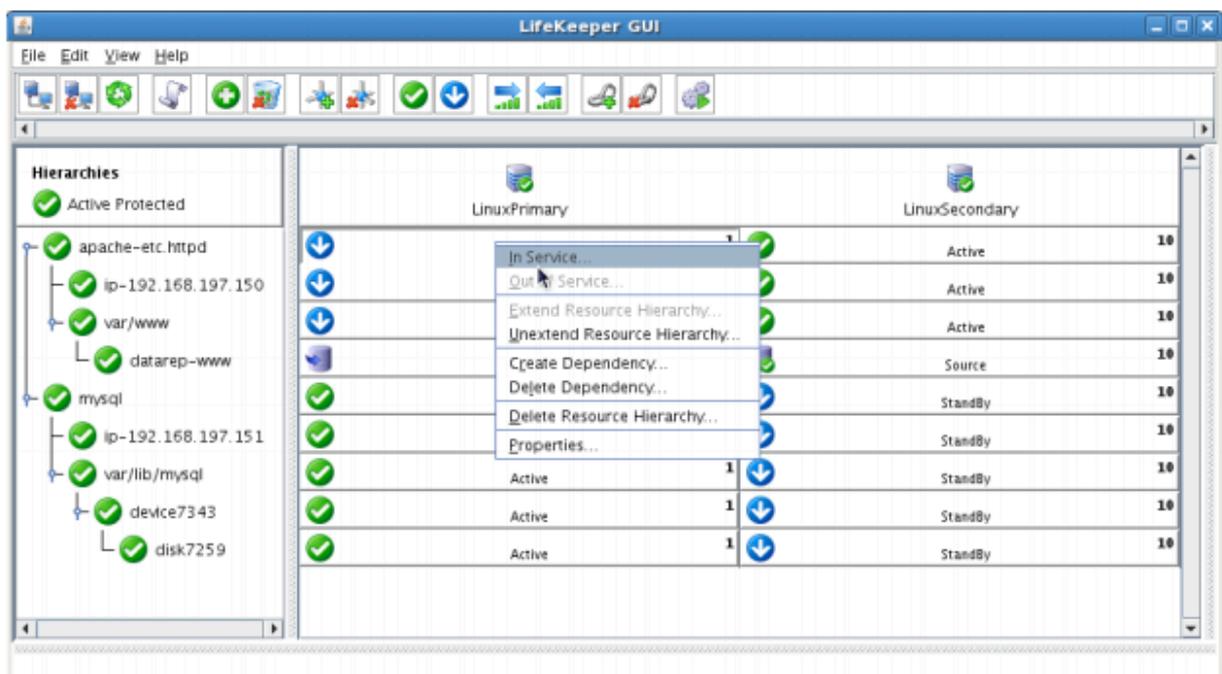
- Using the LifeKeeper GUI, verify that the MySQL and dependent resources are active on LINUXPRIMARY.
- Run “ifconfig -a” on LINUXPRIMARY to validate that the IP Address 192.168.197.151 is active on LINUXPRIMARY
- Run “df -h” to verify that the /var/lib/mysql shared iSCSI filesystem is mounted on LINUXPRIMARY
- Verify the MySQL services are running on LINUXPRIMARY by running “ps -ef | grep -i mysql”
- On LINUXPRIMARY run the following command to verify client connectivity to the MySQL database:
  - # mysql -S /var/lib/mysql/mysql.sock -u root -p

- (enter password “SteelEye”)
- Verify that the SCSI reservation has properly locked the iSCSI LUN to the currently active cluster node. On LINUXSECONDARY, run “mount /dev/sdc1 /var/lib/mysql”. This should FAIL because LINUXSECONDARY does not own the SCSI reservation on this LUN.

## Manual Switchover of the Apache Hierarchy back to the Primary Server

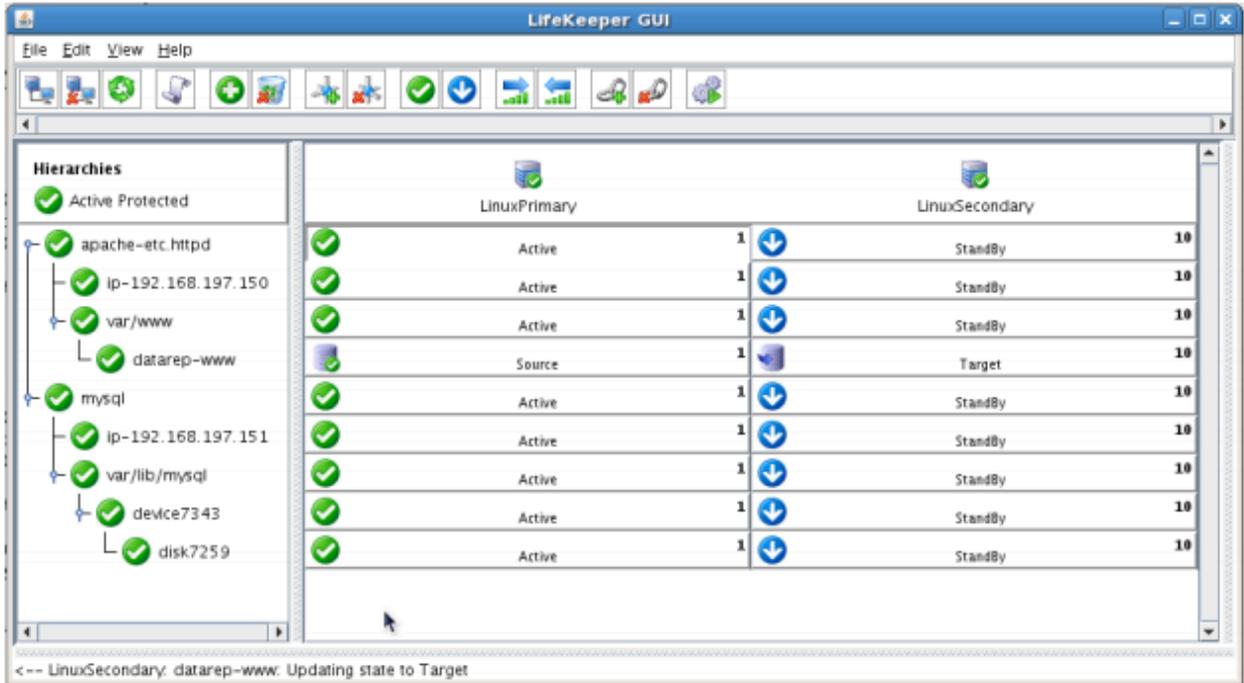
### Procedure:

- From the LifeKeeper GUI, right click on the Apache resource on the Primary Server (LINUXPRIMARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up



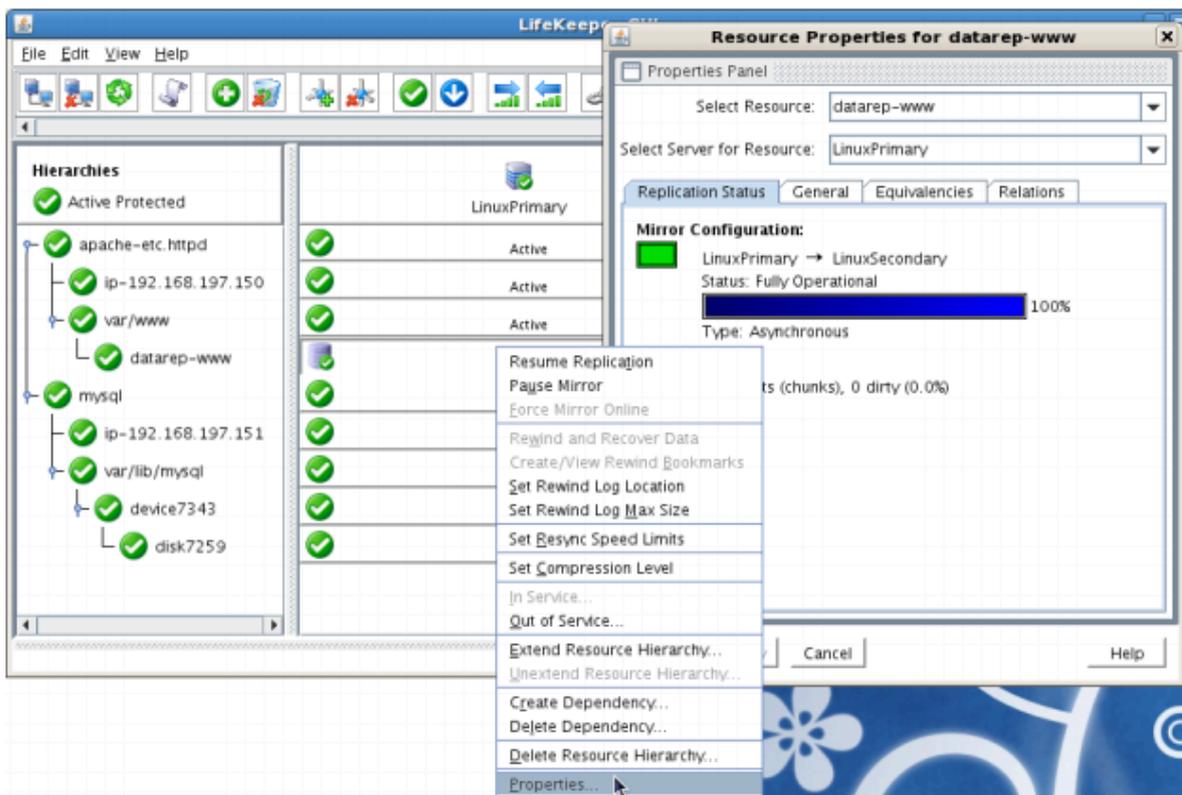
### Expected Result:

- Beginning with the Apache resource, all resources will be removed from service on the Active Server (LINUXSECONDARY).
- Beginning with the dependent resources (IP and Replicated Volume), all resources will be brought in service on LINUXPRIMARY
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXPRIMARY -> LINUXSECONDARY



**Tests/Verification:**

- Using the LifeKeeper GUI, verify that the Apache and dependent resources are active on LINUXSECONDARY.
- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the “datarep-www” resource and select Properties



- Run “ifconfig -a” on LINUXPRIMARY to validate that the IP Address 192.168.197.150 is active on

## LINUXPRIMARY

- Run “df –h” to verify that the /var/www replicated filesystem is mounted as an “md” device (example: /dev/md0”) on LINUXPRIMARY
- Verify the Apache services are running on LINUXPRIMARY by running “ps –ef | grep –i httpd”
- Open a Web Browser to <http://192.168.197.150> and verify that it can successfully connect. The PHPInfo output should indicate that the system name is “LinuxPrimary



<b>System</b>	Linux LinuxPrimary 2.6.18-194.26.1.el5 #1 SMP Tue Nov 9 12:54:40 EST 2010 i686
<b>Build Date</b>	Nov 29 2010 16:49:03
<b>Configure</b>	'./configure' '--build=i686-redhat-linux-gnu' '--host=i686-redhat-linux-gnu' '--target=i386-

## Simulate a network failure on the Primary Server by failing the IP resource

**!** **IMPORTANT:** Only perform this test if you have more than one communications path configured.

If you perform this test and have only one communications path configured, your system will enter a split-brain scenario as described in the LifeKeeper Administration Guide found [here](#). Refer to this document for more information or contact SIOS presales technical support for assistance in resolving this condition.

### Procedure:

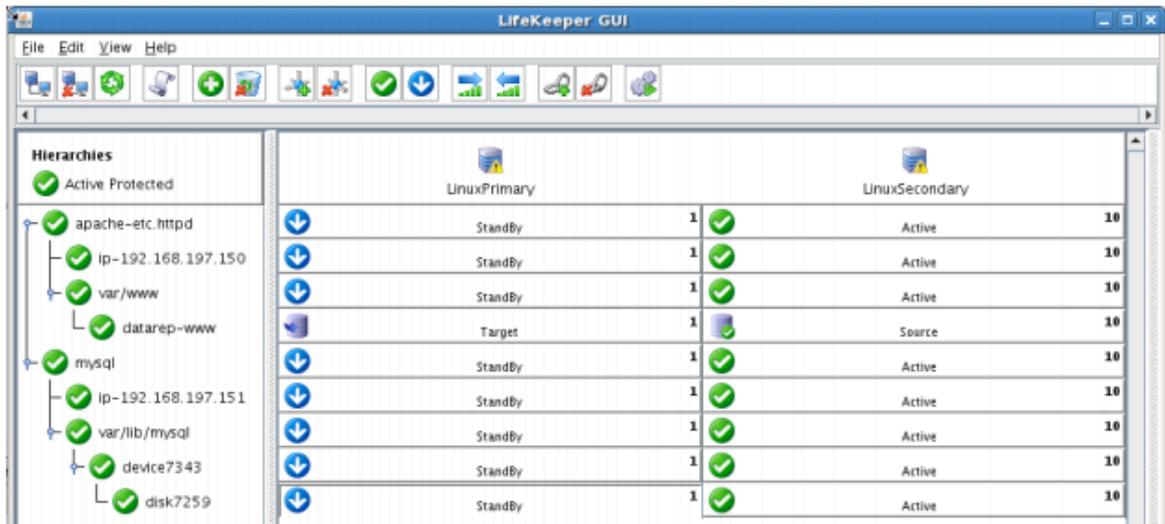
- On LINUXPRIMARY, pull the network cable attached to the NIC that the virtual IP address is configured on

### Expected Result:

- The IP Resource should fail first.
- The entire hierarchy should failover to LINUXSECONDARY

### Tests/Verification:

- Check the LifeKeeper Log to verify the IP resource failed – “/opt/LifeKeeper/bin/lk\_log log”
- Using the LifeKeeper GUI, verify the MySQL and Apache resource hierarchies fail over successfully to LINUXSECONDARY



## Hard failover of the resource from the Secondary Server back to the Primary Server

### Procedure:

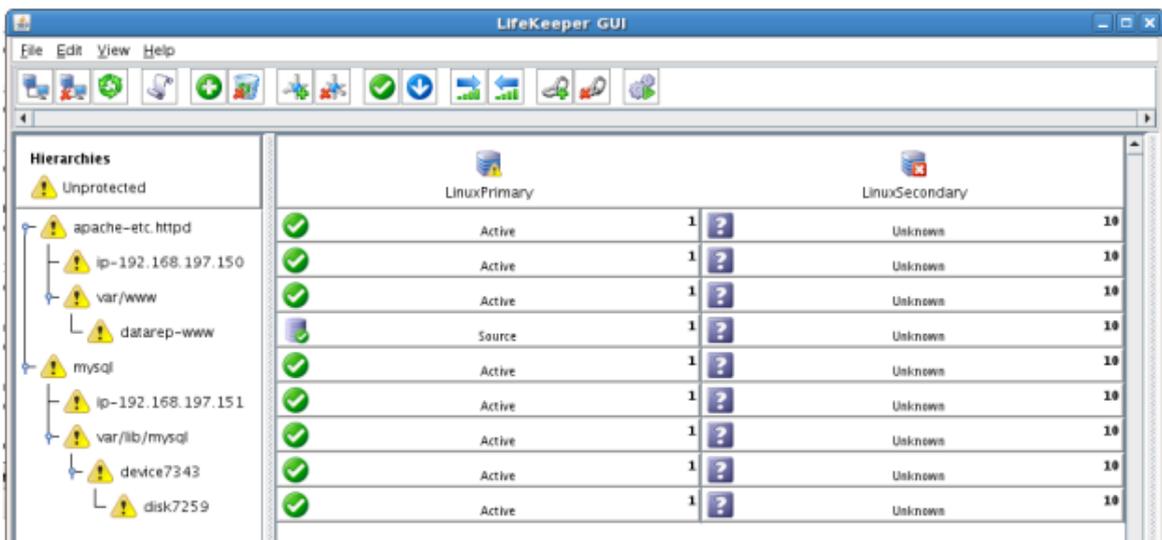
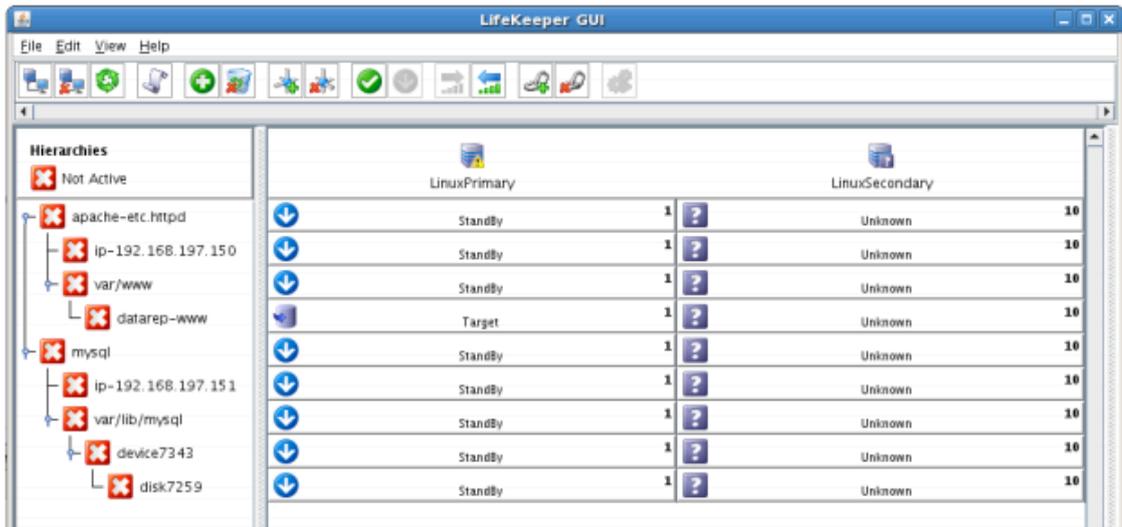
- Pull the power cord on LINUXSECONDARY, as this is the server with all resources currently In Service.

### Expected Result:

- After failure has been detected, beginning with the dependent resources (IP and Volume), all resources will be brought in service on LINUXPRIMARY.

### Tests/Verification:

- Using the LifeKeeper GUI, verify the mirror has reversed and is in a Resync Pending state waiting for LINUXSECONDARY to come back on line.
- Verify the PostgreSQL Server services are running on LINUXPRIMARY.
- Verify that the client can still connect to the Webserver and database running on LINUXPRIMARY.
- Verify you can write data to the replicated volume, /var/www on LINUXPRIMARY.



## Bring Failed Server back on line

### Procedure:

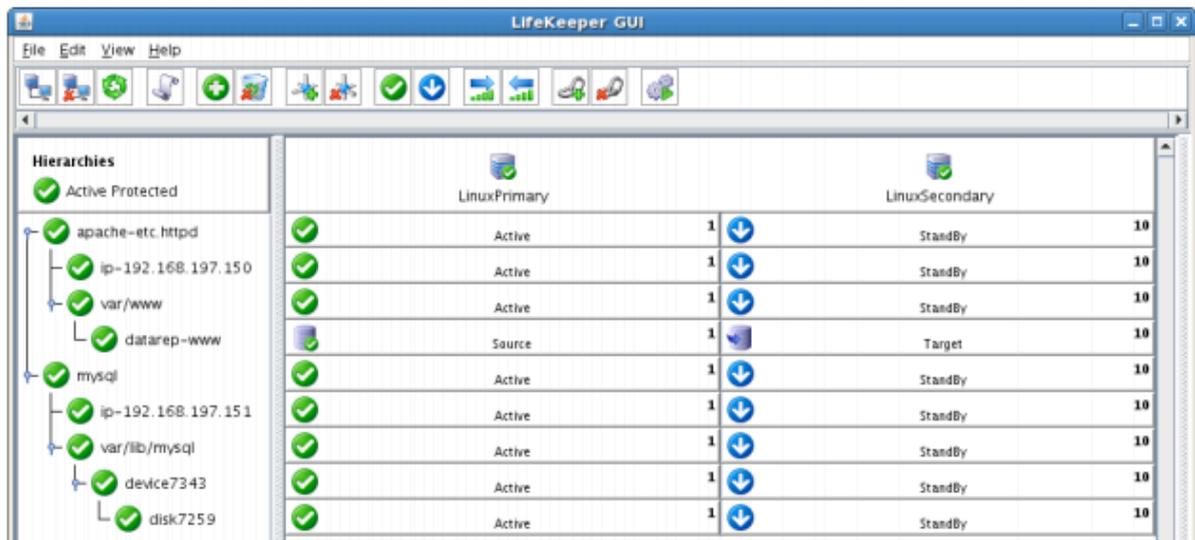
- Plug the power cord back into LINUXSECONDARY and boot it up.

### Expected Result:

- Using the LifeKeeper GUI, verify that LINUXSECONDARY is coming back up and has become the Standby Server.

### Tests/Verification:

- Verify the mirror performs a quick partial resync and moves to the Mirroring state
- Verify the Apache and MySQL Hierarchy are in service on LINUXPRIMARY and standby on LINUXSECONDARY.



## Verify Local Recovery of MySQL Server

### Procedure:

- Kill the PostgreSQL processes via the command line:
  - # ps -ef | grep sql
  - # killall mysqld mysqld\_safe
  - run “ps -ef | grep sql once again to verify that the processes no longer exist

### Expected Result: (Assumes Local Recovery for SQL resource is set to YES)

- The MySQL Server service should stop.
- The MySQL quickcheck process will automatically restart the MySQL Server Service when it runs periodically.
- No failure of MySQL should occur.

### Tests/Verification:

- Execute “ps -ef | grep sql” once again to verify that the mysql processes have been restored locally on LINUXPRIMARY.
- Verify connectivity to the MySQL database by running:
  - # mysql -S /var/lib/mysql/mysql.sock -u root -p
  - (Enter password “SteelEye”)
- If you inspect the LifeKeeper logs, you will see information indicating that LifeKeeper detected the failure of the MySQL service and recovered it locally. Run /opt/LifeKeeper/bin/lk\_log log for more information.

# 13. LifeKeeper Single Server Protection

---

[LifeKeeper Single Server Protection Release Notes](#)

[LifeKeeper Single Server Protection for Linux Installation Guide](#)

# 13.1. LifeKeeper Single Server Protection for Linux Release Notes

---

## Version 9.6.1

Released April 20, 2022

### Important!!

**Read This Document Before Attempting To Install Or Use This Product!**

**This document contains last minute information that must be considered before, during and after installation.**

## Introduction

This release notes document is written for the person who installs, configures and/or administers the LifeKeeper Single Server Protection for Linux product. The document contains important information not detailed in the formal LifeKeeper Single Server Protection documentation set such as system requirements, new features and links to product restrictions and troubleshooting hints and tips. It is important that you review this document before installing and configuring LifeKeeper Single Server Protection software.

## LifeKeeper Single Server Protection Product Description

LifeKeeper Single Server Protection allows for application monitoring in a single node configuration (i.e., no cluster requirements or restraints). Single node environments may be physical or virtual (vSphere, KVM, Amazon EC2). LifeKeeper Single Server Protection is built on the proven and stable architecture of SIOS LifeKeeper. LifeKeeper Single Server Protection provides superior application monitoring and can perform recovery of failed applications and system infrastructure items (e.g., NFS share, IP address, File System). If an application cannot be recovered for some reason, LifeKeeper Single Server Protection will initiate a restart of the node via a system reboot or via a VMware HA restart for VMware virtual machines configured for VM and Application Monitoring.

## Components

LifeKeeper Single Server Protection Software is bundled and runs on 64-bit systems (x86\_64, AMD64) and is comprised of the following components:

Package	Package Name	Protected Applications
<a href="#">LifeKeeper Core</a>	steeleye- lk-9.6.1-7412.x86_64.rpm	The LifeKeeper package provides recovery software for failures associated with core system components such as memory, CPUs, the operating system, the SCSI disk

		subsystem and file systems.
<a href="#">LifeKeeper GUI</a>	steeleye- lkGUI-9.6.1-7412.x86_64.rpm	The LifeKeeper GUI package provides a graphical user interface for LifeKeeper administration and status monitoring.
<a href="#">LifeKeeper IP Recovery Kit</a>	steeleye- lkIP-9.6.1-7412.noarch.rpm	The SSP IP Recovery Kit provides recovery software for automatic switchover of IP addresses.
<a href="#">Quick Service Protection</a>	steeleye- lkQSP-9.6.1-7412.noarch.rpm	The Quick Service Protection Recovery Kit provides a simple disaster recovery function for various services.
LifeKeeper Man Pages	steeleye- lkMAN-9.6.1-7412.noarch.rpm	The LifeKeeper Man Page package provides reference manual pages for the SSP product.

## LifeKeeper Single Server Protection Optional Software

Package	Package Name	Protected Applications
<a href="#">LifeKeeper Apache Web Server Recovery Kit</a>	steeleye- lkAPA-9.6.1-7412.noarch.rpm	Apache Web Server v2.4
<a href="#">LifeKeeper SAP MaxDB Recovery Kit</a>	steeleye- lkSAPDB-9.6.1-7412.noarch.rpm	SAP MaxDB v7.9
<a href="#">LifeKeeper DB2 Recovery Kit</a>	steeleye- lkDB2-9.6.1-7412.noarch.rpm	IBM DB2 Universal Database v10.5, v11.1 and v11.5 IBM DB2 Enterprise Server Edition (ESE) v10.5, v11.1 and v11.5 IBM DB2 Workgroup Server Edition (WSE) v10.5, v11.1 and v11.5 IBM DB2 Express Edition v10.5, v11.1 and v11.5
<a href="#">LifeKeeper Oracle Recovery Kit</a>	steeleye- lkORA-9.6.1-7412.noarch.rpm	Oracle Database Enterprise Edition v12c, v12c R2, v18c and v19c (excluding ASM and pluggable database), v21c (excluding ASM and pluggable database features)

		Oracle Database Standard Edition 2 (SE2) v12c, v12c R2, v18c and v19c (excluding ASM and pluggable database), v21c (excluding ASM and pluggable database features)
<a href="#">LifeKeeper MySQL Recovery Kit</a>	steeleye- lkSQL-9.6.1-7412.noarch.rpm	MySQL and MySQL Enterprise v5.7 and v8.0  MariaDB v10.3, v10.4, and v10.5
<a href="#">LifeKeeper PostgreSQL Recovery Kit</a>	steeleye- lkPGSQL-9.6.1-7412.noarch.rpm	PostgreSQL v9.6, v10, v11, v12, v13 and v14  EnterpriseDB Postgres Plus Advanced Server/EDB Postgres Advanced Server v9.6, v10.0, v11.0, v12.0, v13.0 and v14.0
<a href="#">LifeKeeper Sybase ASE Recovery Kit</a>	steeleye- lkSYBASE-9.6.1-7412.noarch.rpm	Sybase ASE 15.7 and 16.0
<a href="#">LifeKeeper Postfix Recovery Kit</a>	steeleye- lkPOSTFIX-9.6.1-7412.noarch.rpm	Postfix software provided with the supported Linux distributions installed and configured on each server. The same version of Postfix should be installed on each server.
<a href="#">LifeKeeper Samba Recovery Kit</a>	steeleye- lkSMB-9.6.1-7412.noarch.rpm	Standard Samba file services provided with the supported Linux distributions
<a href="#">LifeKeeper NFS Server Recovery Kit</a>	steeleye- lkNFS-9.6.1-7412.noarch.rpm	Linux kernel version 2.6 or later  The NFS Server and client packages must be installed on SLES systems.
<a href="#">LifeKeeper Network Attached Storage Recovery Kit</a>	steeleye- lkNAS-9.6.1-7412.noarch.rpm	NFS version of Mounted NFS file systems from an NFS server or Network Attached Storage (NAS) device v2, v3 and v4
<a href="#">LifeKeeper WebSphere MQ Recovery Kit</a>	steeleye- lkMQS-9.6.1-7412.noarch.rpm	IBM MQ v8.0, v9.0, v9.1 and v9.2  See <a href="#">Known Issues and Restrictions</a> > Installation.

<a href="#">Recovery Kit for Oracle Cloud Infrastructure</a>	steeleye- lkOCVIP-9.6.1-7412.noarch.rpm	Allows monitoring of the secondary IP address allocation status in an Oracle Cloud environment.
--------------------------------------------------------------	--------------------------------------------	-------------------------------------------------------------------------------------------------

## New Features of LifeKeeper Single Server Protection for Linux Version 9

Product	Feature
<b>New in Version 9.6.1</b>	
LifeKeeper Core	Red Hat Enterprise Linux 8.5 is supported.
	Oracle Linux 8.5 is supported.
	<a href="#">Bug Fixes</a>
Recovery Kit for Oracle Cloud Infrastructure	Communication via virtual IP address using the LifeKeeper IP Recovery Kit is allowed in the Oracle Cloud environment.
Network Attached Storage Recovery Kit	It is no longer necessary to set "NFS_RPC_PROTOCOL=tcp" when protecting an NFS shared file system with UDP disabled
LifeKeeper GUI	<a href="#">Bug Fixes</a>
<b>New in Version 9.6.0</b>	
LifeKeeper Core	Support SLES 15 SP3
	Supports Oracle Linux 8.4
Oracle	Supports Oracle 21c (21.3) running on-premises
PostgreSQL Recovery Kit	Supports PostgreSQL 14
	Supports FUJITSU Software Enterprise Postgres 13 (Advanced, Standard, Community)
	Supports EDB Postgres Advanced Server 14.0 (Certified in January 2022)
Bug Fixes	
<b>New in Version 9.5.2</b>	
LifeKeeper Core	Supports <b>Red Hat Enterprise Linux 8.4</b>
	Supports <b>CentOS 8.3</b>
	Supports <b>Oracle Linux 8.3</b>
	Supports <b>Oracle Linux 7 UEK 6</b> <b>Note:</b> The kernel should be updated to 5.4.17-2102.202.5 or higher.

	Supports <b>Oracle Linux 8 UEK 6</b> <b>Note:</b> The kernel should be updated to 5.4.17-2102.202.5 or higher.
	LKCLI has been enhanced to control the following Recovery Kits. <ul style="list-style-type: none"> <li>• SAP MaxDB</li> <li>• MQ RK</li> </ul>
	You can now run lkcli from non-root users who belong to the lk group.
	You can now use the “-i” option with lkcli stop as well as with lkstop.
	Bug Fixes
MySQL	Supports MariaDB 10.5
	Bug Fixes
MQ	Supports WebSphere MQ 9.2 for RHEL 7.9
	Supports WebSphere MQ 9.2 for RHEL 8.3
Install, NFS, IP, Oracle, Quorum/Witness, Filesystem, Generic, lksupport	Bug Fixes
<b>New in Version 9.5.1</b>	
LifeKeeper Core	Supports Red Hat Enterprise Linux 7.9 (Certified in December 2020)
	Supports CentOS 7.9 (Certified in December 2020)
	Supports Oracle Linux 7.9 (Certified in December 2020)
	<a href="#">lkstop -i</a> command stops LifeKeeper core but does not stop the protected resources. The user is prompted to confirm (yes/no) that they want to continue.
	LKCLI has been enhanced to control the following Recovery Kits. <ul style="list-style-type: none"> <li>• DB2 RDBMS</li> <li>• RAW</li> <li>• Postfix</li> <li>• VMDK as Shared Storage</li> <li>• SAP ASE</li> <li>• Samba</li> <li>• HULFT / HULFT HUB</li> <li>• OraclePDB for Oracle RDBMS</li> </ul>
	Supports Red Hat Enterprise Linux 8.3 (Certified in January 2021) <b>Note:</b> If you are using DataKeeper, login <a href="#">here</a> and follow <a href="#">these steps</a> when installing LifeKeeper.
	Supports Red Hat Enterprise Linux 8.2
	Supports CentOS 8.2
	Supports Oracle Linux 8.2 (UEK6 is not supported)

	Supports OpenSSL package to 1.1.1g
	Supports SLES15 SP2
	Supports cURL package to 7.68.0
	Bug Fixes
MQ	LifeKeeper for Linux now supports IBM MQ 9.2 (January 2021)
QSP	Apache Tomcat can be protected with the Quick Service Protection (QSP) Recovery Kit
install	Implemented a <a href="#">setup script improvement</a>
	Bug Fixes
PostgreSQL	Supports PostgreSQL 13 (Certified in December 2020)
	Supports EDB Postgres Advanced Server 13.0 (Certified in December 2020)
	Supports FUJITSU Software Enterprise Postgres 12 (Certified in December 2020)
SAP MaxDB, DB2, Filesystem, Generic Application, IP	Bug Fixes
<b>New in Version 9.5.0</b>	
LifeKeeper Core	Supports Red Hat Enterprise Linux 7.8 (Certified in July 2020)
	Supports CentOS 7.8 (Certified in July 2020)
	Supports Oracle Linux 7.8 (Certified in July 2020)
	Supports SUSE Linux Enterprise Server 12 SP5 (Certified in July 2020)
	Support VMware vSphere 7.0 (Certified in July 2020)
	Supports CentOS 8.0
	Supports Oracle Linux 8.0
	Supports Red Hat Enterprise Linux 8.1
	Supports CentOS 8.1
	Supports Oracle Linux 8.1
	The CLI has been enhanced to allow you to control LifeKeeper through the Command Line Interface. See <a href="#">LKCLI</a> for details.
Bug Fixes	
PostgreSQL	Support PostgreSQL 12
	EDB Postgres Advanced Server v12.0 is supported. (Certified in July 2020)
LifeKeeper Core,	Bug Fixes

Filesystem, NFS, DB2, MaxDB, Sybase ASE	
<b>New in Version 9.4.1</b>	
LifeKeeper Core	OpenJDK included with OS is installed. See <a href="#">Configuring the LifeKeeper GUI</a> for details.
	Supports SUSE Linux Enterprise Server 15 SP1
	Supports Oracle Linux 7.7
	Supports CentOS 7.7
	Supports AWS Nitro system
	Supports AWS Transit Gateway
	Bug Fixes
Install, IP, MaxDB	Bug Fixes
<b>New in Version 9.4.0</b>	
LifeKeeper Core	Oracle Linux 7 Unbreakable Enterprise Kernel Release 5 (UEK R5) is supported.
	Red Hat Enterprise Linux 8 is supported.
	<b>Note:</b> Upgrading from one kernel version to another major version such as from RHEL7 to RHEL8 is NOT supported. (i.e DataKeeper resource does NOT work when upgrading from RHEL7 to RHEL8.)
	Red Hat Enterprise Linux 7.7 is now supported. (Authorized in November 2019)
MySQL	MariaDB10.3 is supported.
DB2	DB2 11.5 is supported.
General maintenance	Bug Fixes
<b>New in Version 9.3.2</b>	
LifeKeeper Core	Red Hat Enterprise Linux 7.6 is supported.
	CentOS 7.6 is supported.
	Oracle Linux Version 7.6 is supported.
	SUSE Linux Enterprise Server 12 SP4 is supported.

	SUSE Linux Enterprise Server 15 is supported.
Install	The -s option for saving the current setup configuration has been added to the setup command.
PostgreSQL	PostgreSQL 11 is supported.
	EDB Postgres Advanced Server v11 is supported.
MQ	LifeKeeper for Linux now supports IBM MQ 9.1
Oracle	Support Oracle 19c (Certified in August 2019).
General Maintenance	Bug fixes
<b>New in Version 9.3.1</b>	
LifeKeeper Core	<p>Updated the OpenSSL package to 1.0.2p</p> <p>Support Red Hat Enterprise Linux 6.10</p> <p>Support CentOS 6.10</p> <p>Support Oracle Linux 6 Update 10</p>
MySQL	Support MySQL 8.0
Oracle	Support Oracle 18c (Certified in March 2019)
Install	Bug fixes
<b>New in Version 9.3</b>	
LifeKeeper Core	<p>Red Hat Enterprise Linux Version 7.5 is supported.</p> <p>CentOS7.5 is supported.</p> <p>Oracle Linux Version 7.5 is supported.</p> <p>Support VMware vSphere 6.7. (Certified in October 2018)</p> <p>Bug fixes</p>
Install	The installation script has been renewed. For details, please click <a href="#">here</a> .
Oracle, Samba, MQ, Sybase, Filesystem, Generic Application, QSP	Bug fixes

<b>New in Version 9.2.2</b>	
PostgreSQL	Support PostgreSQL 10 EDB Postgres Advanced Server v10.0 is now supported. (Certified in April 2018)
NAS	Bug fixes
<b>New in Version 9.2.1</b>	
LifeKeeper Core	Support Oracle Linux 7.4 Support CentOS 7.4 Support SUSE Linux Enterprise Server 12 SP3 <ul style="list-style-type: none"> <li>The kernel should be updated to 4.4.82-6.9.1 for SUSE Linux Enterprise Server 12 SP3</li> </ul> Bug fixes
PostgreSQL	Support EDB Postgres Advanced Server 9.6
MQ	Support IBM MQ 9.0
<b>New in Version 9.2</b>	
LifeKeeper Core	Support Red Hat Enterprise Linux 7.4 SNMP trap can be sent to multiple targets Virtualization environment Nutanix Acropolis Hypervisor is supported. (LifeKeeper is not supported) Bug fixes
IP	IP resources using real IP (primary IP address configured for NIC) can be created
PostgreSQL	Support PostgreSQL 9.6
MQ	Support IBM MQ 9.0 (Certified in December 2017)
SAP MaxDB, Install	Bug fixes
<b>New in Version 9.1.1</b>	

<p>LifeKeeper Core</p>	<p>SUSE Linux Enterprise Server 12 SP1 support.</p> <p>* SLES12.0 is not supported.</p> <p>* Btrfs is not supported.</p> <p>Red Hat Enterprise Linux Version 7.3 support.</p> <p>Oracle Linux Version 7.3 support.</p> <p>* UEK is not supported.</p> <p>vSphere 6.5 support (SMC feature is no longer supported with vSphere 6.5).</p> <p>Bug fixes</p>
<p>PostgreSQL</p>	<p>PostgreSQL 9.5 support</p> <p>EDB Postgres Advanced Server v9.5 support</p> <p>For the details, refer to the LifeKeeper Optional Recovery Software Requirements, <a href="#">PostgreSQL Recovery Kit Administration Guide &gt; Administration</a>.</p>
<p>Sybase ASE</p>	<p>Sybase ASE 16.0 support.</p>
<p>MySQL</p>	<p>MySQL 5.7 support on RHEL 7.x/CentOS 7.x/OEL 7.x.</p> <p>* MySQL 5.7 on other OS is already supported.</p>
<p><b>New in Version 9.1.0</b></p>	
<p>LifeKeeper Core</p>	<p>Red Hat Enterprise Linux 6.8 support (Certified in September 2016).</p> <p>CentOS 6.8, Oracle Linux 6.8 support (Certified in September 2016).</p> <p>*MD RecoveryKit is not supported on these OS.</p> <p>LifeKeeper API for Monitoring</p> <p>Added API to supply LifeKeeper status and log information.</p> <p>Quick Service Protection support</p> <p>Added functionality to easily protect OS services.</p> <p>Bug Fixes.</p>

<b>New in Version 9.0.2</b>	
LifeKeeper Core	<p>Support of Red Hat Enterprise Linux Version 7.2.</p> <p>※MySQL RK is not supporting RHEL 7.x/CentOS 7.x/OEL 7.x.</p> <p>※Support of RHEL 7.2 by each application must be confirmed by user.</p> <p>Update OpenSSL package to 1.0.1q</p> <p>Bug Fixes.</p>
MQ	<p>WebSphere MQ – Added support for Multi-version WebSphere MQ. With this support queue managers for 7.1, 7.5, and 8.x can all be protected on the same cluster node.</p> <p>Added the function that mqm group user can execute MQ command alternatively</p> <p>Bug Fixes.</p>
IP, Filesystem, PostgreSQL, SAP MaxDB, Oracle	Bug Fixes.
Licensing	Update the package of FlexNet
<b>New in Version 9.0.1</b>	
LifeKeeper Core	Bug Fixes
DataKeeper	Bug Fixes
<b>New in Version 9.0</b>	
LifeKeeper Core	<p>Combined documents of <a href="#">Parameters List</a>, and added the <a href="#">lkchkconf command</a>.</p> <p>vSphere 6 support (SMC feature is no longer supported with vSphere 6.)</p> <p>reiserfs filesystem is no longer supported.</p> <p>Arks supported with Red Hat Enterprise Linux Version 7.0/7.1, Community ENTerprise Operating System (CentOS) Version 7.0/7.1, and Oracle Linux Version 7.0/7.1 are the same as LifeKeeper for Linux v8.4.1. (Arks to be applied: PostgreSQL, MySQL, Oracle, DB2, Apache, Postfix, NFS, NAS, Samba)</p> <p>Bug Fixes.</p>

GUI	JRE 8u51 support. (JRE 7 is no longer supported.)  Chrome Browser is no longer supported  Bug Fixes.
FileSystem, PostgreSQL	Bug Fixes.

## Bug Fixes

The following is a list of the latest bug fixes and enhancements.

Bug	Description
PL-3402	An error is now displayed when autofs is detected when creating a filesystem resource.
PL-5679	Fixed the refresh button on the main frame that was grayed out and could not be selected when logging into the GUI as a user belonging to lkoper or lkguest.
PL-12204	Fixed a problem in which Sybase resources were displayed as ISP even though ASE was not running.

## Discontinued Features

Feature	Description
<b>Discontinued in Version 9.6.1</b>	
	None
<b>Discontinued in Version 9.6.0</b>	
	None
<b>Discontinued in Version 9.5.2</b>	
Core	SteelEye Management Console is no longer supported.
	Red Hat Enterprise Linux 6 is no longer supported.
	CentOS 6 is no longer supported.
	Oracle Linux 6 is no longer supported.
Oracle	Oracle virtual machine (OVM) latest 3.4.6 is no longer supported.
PostgreSQL	PostgreSQL 9.5 is no longer supported.
	EnterpriseDB Postgres Plus Advanced Server 9.5 is no longer supported.
VMware	VMware 5.5 and 6.0 are no longer supported.

<b>Discontinued in Version 9.5.1</b>	
	<b>None</b>
<b>Discontinued in Version 9.5.0</b>	
LifeKeeper Core	System log management using syslog-ng is no longer supported. Please use rsyslog.
	SUSE Linux Enterprise Server (SLES) 11.0 to SP4 is no longer supported.
Oracle	Oracle Database Enterprise Edition 11g R2 is no longer supported.
	Oracle Database Standard Edition 11g R2 is no longer supported.
	Oracle Database Standard Edition One 11g R2 is no longer supported.
MySQL	MariaDB 5.5, 10.0 is no longer supported.
PostgreSQL	PostgreSQL 9.4 is no longer supported.
	EnterpriseDB Postgres Plus Advanced Server 9.4 is no longer supported.
<b>Discontinued in Version 9.4.1</b>	
	<b>None</b>

## System Requirements

### LifeKeeper Single Server Protection Product Requirements

LifeKeeper Single Server Protection is supported on any Linux platform that satisfies the minimum requirements included in the table below.

Description	Requirement
Linux Operating System	See the <a href="#">Linux Configuration Table</a> for specific operating system information.
Virtual Environments	<p>The guest operating system running on the virtual machine must be one of the supported versions listed in the <a href="#">Linux Configuration Table</a>. The following virtual environment is an example where LifeKeeper for Linux is deployed. Please refer to the <a href="#">Support Matrix</a> for detailed versions of supported virtualization environments.</p> <ul style="list-style-type: none"> <li>• KVM</li> <li>• Oracle VM Server for x86</li> <li>• VMware vSphere v6.5, v6.7 and v7.0</li> <li>• Amazon EC2</li> <li>• Microsoft Azure</li> </ul>

	<ul style="list-style-type: none"> <li>• Nutanix Acropolis Hypervisor</li> <li>• Google Cloud</li> <li>• Oracle Cloud Infrastructure (<b>See Note3</b>)</li> </ul> <p>SAN configuration is supported for vSphere 6.5 or later except RDM which is not supported by VMWare.</p> <p>Fibre channel SAN and shared SCSI cluster configurations are not supported with LifeKeeper for Linux running in a KVM and Oracle VM Server for x86 virtual machine.</p> <p><b>Note1:</b> Some Amazon EC2 configurations have issues when the Shutdown Strategy is set to “Do not Switchover Resources”. For detailed information, see Troubleshooting &gt; <a href="#">Known Issues and Restrictions</a>.</p> <p><b>Note2:</b> On SLES v12 or later running on AWS or Azure, the dynamic change of the virtual IP address by the cloud network plug-in may affect the operation of the LifeKeeper cluster. For detailed information, see <a href="#">LifeKeeper Core – Known Issues / Restrictions</a>.</p> <p><b>Note3:</b> Refer to <a href="#">Support Configuration</a> for the supported configuration and the restrictions.</p>
Memory	<p>The LifeKeeper for Linux minimum memory requirement is the same as the OS minimum requirement. System memory should be sized for the applications that will be running on the LifeKeeper protected system as well. Refer to <a href="#">Application Configuration</a> for further information.</p>
Memory	<p>The minimum memory requirement for a system supporting LifeKeeper is 512 MB. This is the minimum amount required by LifeKeeper supported Linux distributions. System memory should be sized for the applications that will be running on the LifeKeeper protected system as well. Refer to <a href="#">Application Configuration</a> for further information.</p>
Disk Space	<p>The LifeKeeper Package Cluster requires the following disk space:</p> <p>/opt – approx 100MB (depending on kits installed)</p> <p>/ – approx 110MB</p>
Java Runtime Environment	<ul style="list-style-type: none"> <li>• OpenJDK 1.8, 10 or later</li> </ul>

## Upgrades

SSP for Linux can be upgraded to Version 9.6.1 from either SSP for Linux Version 9.4.x or Version 9.5.x. If upgrading from a version other than 9.4.x or 9.5.x, the older version will need to be uninstalled and SSP for Linux will have to be reinstalled. An alternative to uninstalling the older version would be to upgrade from the older version to 9.4.x or 9.5.x, then perform the upgrade to 9.6.1.

## LifeKeeper Single Server Protection Support Software Requirements

The following table of Supporting Software is only required in VMware VMs configured for VM and Application Monitoring.

Product	Requirement(s)	Disk Space Required
VMware	VMware vSphere Client (for LifeKeeper Single Server Protection vSphere Client Plug-in functionality)  VMware Tools installed and running on all protected virtual machines  VMware Application HA monitoring must be enabled and set to VM and Application Monitoring for all protected virtual machines	Approximately 175 KB in <i>/opt</i> (for VMware Tools)

## Open Source Packages

The following open source packages are included in the LifeKeeper installation image.

Name	Version	License Type and Version
curl-7.68.0-1	7.68.0	MIT
libcurl-7.68.0-1	7.68.0	MIT
gnutls-2.8.6-3.1	2.8.6	GPLv3+ and GPLv2+
gnutls-utils-2.8.6-3.1	2.8.6	GPLv3+
libgcrypt-1.5.0-2.1	1.5.0	LGPLv2+
libgpg-error-1.10-2.1	1.1	LGPLv2+
libxml2-2.7.8-7.1	2.7.8	MIT
libxml2-static-2.7.8-7.1	2.7.8	MIT
lighttpd-1.4.41-2	1.4.41	BSD

lighttpd-fastcgi-1.4.41-2	1.4.41	BSD
openssl-1.1.1g-1	1.1.1g	BSDish
openssl-perl-1.1.1g-1	1.1.1g	BSDish
pcre-4.5-2.1	4.5	distributable
pdksh-5.2.14-780.7.1	5.2.14	GPL, distributable
perl-5.8.8-8.2	5.8.8	Artistic License or GPL
perl-addons-5.8.8-26	5.8.8	Various(GPL, artistic v2, BSD, Open Market, GPLv2+, MIT)
readline-4.3-14.1	4.3	GPL
runit-2.0.0-4.11	2.0.0	BSD
util-linux-2.31.1-2	2.31.1	GPLv2 and GPLv2+ and LGPLv2+ and BSD with advertising and Public Domain
Perl Config::IniFiles (CPAN module)	2.27	GPL/Artistic (Same as Perl)
openjdk-12.0.2	12.0.2	GPLv2+
kconfig-frontends	4.11.0	GPLv2
balance	3.54	GPL
mdadm	3.2.6	GPL v2
nbd-client	1.0	GPL
nbd-server	1.3	GPL
HADR-CentOS-2.6.32	2.6.32	GPLv2
HADR-CentOS-3.10.0	3.10.0	GPLv2
HADR-CentOS-4.18.0	4.18.0	GPLv2
HADR-RHAS-2.6.32	2.6.32	GPLv2
HADR-RHAS-3.10.0	3.10.0	GPLv2
HADR-RHAS-4.18.0	4.18.0	GPLv2
HADR-OEL-2.6.32	2.6.32	GPLv2
HADR-OEL-3.10.0	3.10.0	GPLv2
HADR-OEL-4.18.0	4.18.0	GPLv2
HADR-OEL.UEK-4.14.35	4.14.35	GPLv2
HADR-SuSE-4.12.14	4.12.14	GPLv2
HADR-SuSE-5.3.18	5.3.18	GPLv2

## Known Issues

See the [Known Issues and Workarounds](#) section in [LifeKeeper Single Server Protection for Linux Technical Documentation](#) for known issues, workarounds and other troubleshooting information.

**Trademarks:**

- “Amazon Web Services, the “Powered by Amazon Web Services” logo, “AWS”, “Amazon EC2”, “EC2”, “Amazon Elastic Compute Cloud”, “Amazon Virtual Private Cloud”, and “Amazon VPC” are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries.
- Google Cloud, BigQuery and Google Compute Engine are trademarks of Google LLC.
- “Oracle Cloud” is a registered trademark of Oracle Corporation and its affiliates.  
Trademark symbols such as ® and ™ may be omitted from system names and product names in this document.

# 13.2. LifeKeeper Single Server Protection for Linux Installation Guide

## About LifeKeeper Single Server Protection for Linux

LifeKeeper Single Server Protection (SSP) allows for application monitoring in a single node configuration (i.e., no cluster requirements or restraints). Single node environments may be physical or virtual (vSphere, KVM, Amazon EC2). LifeKeeper SSP is built on the proven and stable architecture of SIOS LifeKeeper. LifeKeeper SSP provides superior application monitoring and can perform recovery of failed applications and system infrastructure items (e.g., NFS share, IP address, File System). If an application cannot be recovered for some reason, LifeKeeper SSP will initiate a restart of the node via a system reboot or via a VMware HA restart for VMware virtual machines configured for VM and Application Monitoring.

 **Note:** Because LifeKeeper SSP is built using the SIOS LifeKeeper technology, you will see references to LifeKeeper throughout the documentation as well as references to information found in the LifeKeeper for Linux documentation for topics common to both products. When referencing these common topics the following subject items do not apply to LifeKeeper SSP:

- Clustering
- Communication Paths
- Shared Storage (requirements, configuration, ...)
- Extending / Unextending resource hierarchies
- Quorum/.Witness
- Application Recovery Kits
  - Recovery Kit for EC2
  - LVM Recovery Kit
  - MD Recovery Kit
  - Route53 Recovery Kit
  - SAP Recovery Kit
  - SAP HANA Recovery Kit
  - VMDK as Shared Storage Recovery Kit
- Storage Kits (DR, DMMP, HDLM, LVM, MD, PPATH and NEC SPS)

 **Note:** Without the underlying storage kits in LifeKeeper SSP, steps must be taken to ensure the devices required to mount protected file systems are activated during system boot (e.g. if the file system is mounted on a logical volume the volume must be in the

active state before LifeKeeper SSP starts)

- Resource/Machine failovers (by default with LifeKeeper SSP these result in a node restart)
- Resource switchovers
- Switchable IP Addresses (with LifeKeeper SSP Switchable IP addresses are required for some protected applications but since there is only a single node no switching actually takes place)

 **Note:** When operating on Amazon EC2, assign a secondary private IP address to the NIC using the Amazon EC2 Management Console prior to creating the IP resource. Next, create the IP resource as the private IP address on the NIC that is using the LifeKeeper GUI client. An Elastic IP can now be associated with the IP resource

For more information on the SIOS LifeKeeper product, on which LifeKeeper SSP is built, please see the [LifeKeeper for Linux documentation](#) for the common release number. This documentation will provide detailed information on resource hierarchy creation, resource types, states and relationships, Graphical User Interface (GUI), as well as common and advanced tasks.

## 13.2.1. LifeKeeper Single Server Protection for Linux Introduction

---

LifeKeeper Single Server Protection (SSP) allows for application monitoring in a single node configuration (i.e., no cluster requirements or restraints). Single node environments may be physical or virtual (vSphere, KVM, Amazon EC2). LifeKeeper SSP is built on the proven and stable architecture of SIOS LifeKeeper. LifeKeeper SSP provides superior application monitoring and can perform recovery of failed applications and system infrastructure items (e.g., NFS share, IP address, File System). If an application cannot be recovered for some reason, LifeKeeper SSP will initiate a restart of the node via a system reboot or via a VMware HA restart for VMware virtual machines configured for VM and Application Monitoring.

 **Note:** Because LifeKeeper SSP is built using the SIOS LifeKeeper technology, you will see references to LifeKeeper throughout the documentation as well as references to information found in the LifeKeeper for Linux documentation for topics common to both products. When referencing these common topics the following subject items do not apply to LifeKeeper SSP:

- Clustering
- Communication Paths
- Shared Storage (requirements, configuration, ...)
- Extending / Unextending resource hierarchies
- Storage Kits (DR, DMMP, HDLM, LVM, MD, PPATH and NEC SPS)

 **Note:** Without the underlying storage kits in LifeKeeper SSP, steps must be taken to ensure the devices required to mount protected file systems are activated during system boot (e.g. if the file system is mounted on a logical volume the volume must be in the active state before LifeKeeper SSP starts)

- Resource/Machine failovers (by default with LifeKeeper SSP these result in a node restart)
- Resource switchovers
- Switchable IP Addresses (with LifeKeeper SSP Switchable IP addresses are required for some protected applications but since there is only a single node no switching actually takes place)

 **Note:** When operating on Amazon EC2, assign a secondary private IP address to the NIC using the Amazon EC2 Management Console prior to creating the IP resource. Next, create the IP resource as the private IP address on the NIC that is using the LifeKeeper GUI client. An Elastic IP can now be associated with the IP resource

For more information on the SIOS LifeKeeper product, on which LifeKeeper SSP is built, please see the [LifeKeeper for Linux documentation](#) for the common release number. This documentation will provide detailed information on resource hierarchy creation, resource types, states and relationships, Graphical User Interface (GUI), as well as common and advanced tasks.

## 13.2.2. Installing the LifeKeeper Single Server Protection Software

---

This document will guide you through the installation of the LifeKeeper Single Server Protection Software (SSP) and assumes the user has basic knowledge of the Linux operating system. Please refer to the [LifeKeeper Single Server Protection Software for Linux product documentation](#) for more information.

### Pre-Installation Requirements

Before installing SSP for Linux, please check the following:

- [SSP for Linux Release Notes](#) -The Release Notes include supported platforms, operating systems, applications, and storage. They also include the latest features and Bug Fixes.
- **TCP/IP Connection and Name Resolution** – In order to use the GUI function, both cluster nodes need to be able to resolve the name. Use the DNS service or `/etc/hosts` for name resolution. Also, localhost needs to be resolved to 127.0.0.1.
- **Firewall** – The following ports are used:
  - Communication Path (TCP): 7365/tcp
  - Communication of a GUI Server: 81/tcp, 82/tcp
  - RMI Communication between the GUI Server and Client: all the ports after 1024/tcp

#### More Firewall Information

- The port used for communication with the GUI server and a client needs to be open on the cluster node where SSP is installed and on all systems where the GUI client runs.
- For communication between the GUI server and a client, Java RMI (Remote Method Invocation) randomly uses ports 1024 and above. Please refer to the [Technical Documentation](#) for the setting details.
- Add the following to the port numbers you are using: *WebGUI server process and policy setting with the `lkipolicy` command : 778(SSL) /tcp*
- **Check the SELinux Setting** – When the SELinux setting is enabled, SSP for Linux may not be able to be installed depending on the mode..
  - enforcing mode – SSP for Linux cannot be installed
  - permissive mode – SSP for Linux can be installed (not recommended except in some ARK environments)
    - It is not recommended to use SELinux permissive mode unless it is required in an SAP environment. Please make sure that the application to be run on the cluster supports permissive mode. SELinux permissive mode has been tested for following ARKs: SAP MaxDB / Sybase / Oracle / DB2 / NFS / NAS / IP / FileSystem / MQ. Refer to [Linux Dependencies](#) for required packages.
  - disabled mode – SSP for Linux can be installed
    - Please refer to the OS distribution documentation on how to disable SELinux.
  - Install the appropriate package provided by your distribution.

- Check [Known Issues](#) – Please make sure that there are no known issues for your environment.

## Installing SSP for Linux

Install the LifeKeeper Single Server Protection software on each server in the LifeKeeper Single Server Protection configuration.

Packages that LifeKeeper is dependent on are installed automatically because the LifeKeeper installation setup script uses package manager tools (yum or zypper) to ensure installation of all dependent packages.

**!** **IMPORTANT:** A functional yum or zypper configuration is required for the successful installation of LifeKeeper. A non-functional configuration can result in an installation failure (see [Installation Known Issues](#) for more information). Additionally, the package manager repo or rpm database must not be locked as that could cause the install to hang. If the dependent packages cannot be installed automatically via the package manager, refer to Linux Dependencies and install the necessary packages in advance.

The LifeKeeper Single Server Protection image file (lkssp.img) provides a set of installation scripts designed to perform the user interactive system setup tasks that are necessary when installing SSP on your system (see [Interactive Mode](#) for more information). A non-user interactive install can be performed as well (see [Non-interactive Mode](#) for more information).

A licensing utilities package is also installed providing utilities for obtaining and displaying the Host ID or Entitlement ID of your server. Host IDs and/or Entitlement IDs are used to obtain valid licenses for running LKSSP.

**!** **IMPORTANT:** LifeKeeper Single Server Protection does not provide shared storage support or I/O fencing. Each server must use local disk storage for application data. All LifeKeeper Single Server Protection packages are installed in the directory `/opt/LifeKeeper`.

Please refer to [How to Use Setup Scripts](#) for the installation activities.

For upgrading, please refer to [Upgrading SSP](#).

## 13.2.3. How to Use Setup Scripts

---

To install or upgrade SSP, follow the steps below.

### How the Setup Scripts Works

1. Interactive installation

Configure and install LifeKeeper from the menu.

If you save the configuration information at this time, it can be used for the non-interactive installation described below.

2. Non-interactive installation

Install LifeKeeper using the saved configuration information.

Since no inquiry to the user occurs, you can perform this using a building tool (e.g. Ansible).

### How in Install / Upgrade SSP Using the Setup Script

#### Interactive Installation

1. After logging in as the root user, use the following command to mount the lkssp.img file:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

Where PATH is the path to the image

IMAGE\_NAME is the name of the image

MOUNT\_POINT is the path to mount location

2. Change to the directory where lkssp.img is mounted and enter the following:

```
./setup [-s <response_file>]
```

When the `-s` option is specified, you can save your configuration information in the `response_file`, which is used for a non-interactive installation.

3. The script collects information about the system environment and determines what you need to do to install SSP.

If the system requirements for installation or upgrade are not satisfied, then an error message is displayed and the installation / upgrade is cancelled.

Also, if some restrictions arise or a configuration change is required, a warning message will be displayed requiring the user to decide whether to continue or abort the installation.

4. Select the SSP features and Application Recovery Kits (ARKs) to install via the main dialog screen.

Please refer to [How to Use the Dialog Screen](#).

5. Once all the required SSP features and ARKs have been selected, select <Done> to begin the installation.

If any notifications are output when the installation completes, please take the necessary actions to correct them.

## Non-interactive Installation

1. After logging in as root user, use the following command to mount the lkssp.img file:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

PATH is the path to the image

IMAGE\_NAME is the name of the image

MOUNT\_POINT is the path to mount location

2. After copying the configuration file to the system where you want to install LifeKeeper, run the following command:

```
./setup -f <response_file> -q y
```

The “-q y” option gives the answer that the warning has been noted.

 ‘sh setup’ (bourne shell) cannot be used. Use **bash (** instead.

 If you use a configuration file saved with the `-s` option for a non-interactive installation, the system on which the file is used must be configured the same way as the system on which the file was generated. If the systems have too many differences the non-interactive installation may fail. The configuration file created with the `create_response_file` script has no such restrictions.

## Creating the Configuration Information

The configuration information file used for non-interactive installation can be created during setup with `setup -s <response_file>` or created in advance with the `create_response_file` script.

1. After logging in as root user, run the following command to mount the lkssp.img file:

```
mount <PATH/IMAGE_NAME> <MOUNT_POINT> -t iso9660 -o loop
```

PATH is the path to the image

IMAGE\_NAME is the name of the image

MOUNT\_POINT is the path to mount location

2. Change to the directory where lkssp.img is mounted and enter the following:

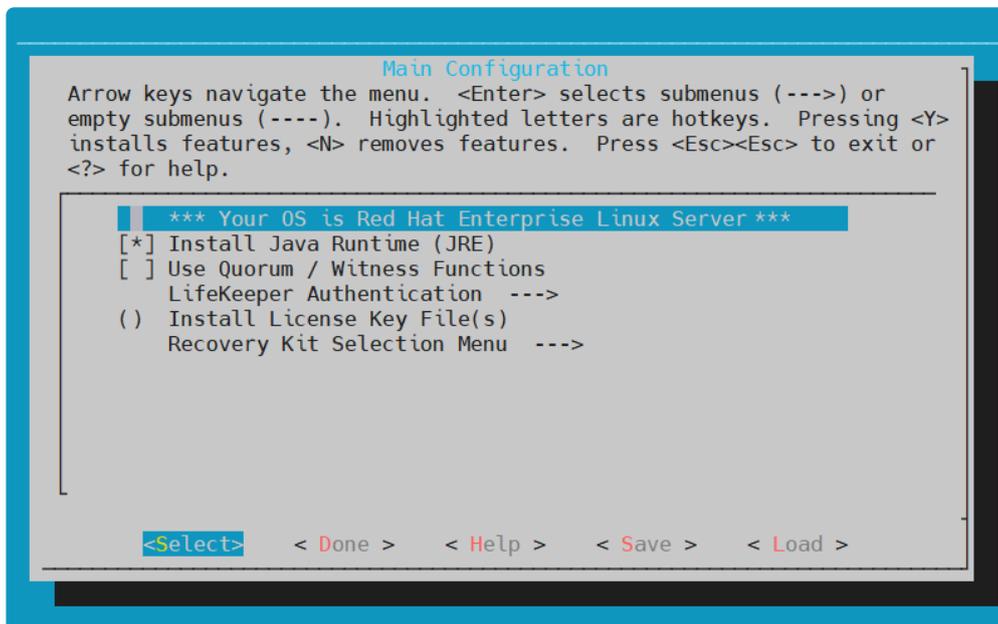
```
./create_response_file <response_file>
```

3. Select the LifeKeeper SPS features and Application Recovery Kits (ARKs) to install via the main dialog screen. Please refer to [How to Use the Dialog Screen](#).
4. Select the SPS features and Application Recovery Kits (ARKs) to install, select **Done** to save the configuration to response\_file and exit the script. The response\_file is copied to the destination system.

Repeat steps 2 through 4 to change the saved configuration information.

## How to Use the Dialog Screen

The dialog screen is displayed below.



Use the following keys to navigate the menu.

↑ ↓	Navigate between menu items
← →	Navigate between the menu buttons at the bottom of the screen
ENTER	Open the selected sub menu
Y / N / ENTER / SPACE	Turn on, turn off or invert the selected item

The menu buttons at the bottom of the screen are used for the following operations.

Select	Opens a sub menu dialog screen
Done	Closes the current screen and returns to the previous screen. Selecting this button on the main screen completes the configuration.
Help	Displays help for the highlighted item

Save	Saves the current settings in a configuration file. The saved configuration file can be used for non-interactive installations. The full path to the file where the configuration information is to be saved should be specified. (Disabled for create_response_file)
Load	Loads a saved configuration file (Disabled for create_response_file)

The “Save” and “Load” menu buttons display a dialog screen asking for a configuration file name for use in saving the current configuration or for loading a saved configuration. If you want to change the default file name provided, move to the file name field using the [TAB] key, and enter a new name. **Note:** The “Save” operation will prompt for confirmation before overwriting a file with the same name.

The items listed below are configurable during installation. During an upgrade only items that can be configured are listed. Using the hotkey <Z> will show those items that will remain unchanged during the upgrade.

- **Install Java Runtime (JRE)**

Install the Java runtime environment used by the LifeKeeper GUI.

- **Restart NFS Service**

When configuring High Availability NFS, restarting the NFS services is required. When this is selected, the services are restarted automatically after the configuration is completed.

**Note:** If you do not want to restart the NFS services automatically, a restart will need to be done to pick up the configuration changes before using the NFS Recovery Kit.

- **Use Quorum / Witness Functions**

Use Quorum / Witness for I/O fencing. For details, please refer to [Quorum/Witness](#) in the technical documentation.

- **LifeKeeper Authentication**

Specify the users allowed to log in to the SSP for Linux GUI along with their privilege levels. Multiple user accounts can be specified by separating them with blanks. For details, refer to [GUI User Settings](#).

- **Install License Key File(s)**

Install the licenses required to start SSP for Linux by entering the path name of the license file to install. Multiple files can be specified by separating them with spaces.

Please refer to [Licensing](#) for details.

- **Recovery Kit Selection**

Select the Application Recovery Kits to install.

Application Recovery Kits are broken into several categories based on common functionality.

Please refer to [Categories for Application Recovery Kits](#) for details.

- **LifeKeeper Startup After Install**

When selected, SSP for Linux will be started when the installation is completed.



**Note:** Because the SSP for Linux Data Replication package may install kernel modules for some of the supported OS distributions, a reinstall of SSP for Linux may be required when the kernel is upgraded. This applies to OS distributions for RedHat, CentOS and Oracle Enterprise Linux (including UEK kernels) running kernel versions 3.10.0-514 or later, and for SUSE Linux Enterprise Server running kernel version 4.12.14-95 or later.

## Adding / Removing Application Recovery Kits

To add Application Recovery Kits after completing an installation, simply execute setup, select the Recovery Kit in the Recovery Kit Selection, followed by the Application Recovery Kit Category and then select the desired kit. If you deselect an Application Recovery Kit which is no longer necessary, that kit will be removed. However, since the kit cannot be removed for the resources in use, delete the resources in advance.

## Repair Installation

To repair an SSP for Linux installation run setup with the “—force” option. A repair installation will update the installation replacing any lost or corrupted files.

## setup Script Options

The setup script can be executed with the following options:

- `-f <response_file>`

Install non-interactively. `<response_file>` contains the configuration information to use during the installation.

- `-s <response_file>`

Save a configuration file containing your menu selections. This file can then be used with the “-f” option to install the same LifeKeeper configuration to another system. For example, run:

```
setup -s <response_file>
```

Select the necessary packages and options and complete setup.

Then run:

```
setup -f <response_file> -q y
```

to run a silent installation of LifeKeeper (on another system) with the same options that were selected the first time setup was run.

- `-force`

Forcibly reinstall SSP for Linux.

- `-q <y/n>`

Specifies the response to any confirmation questions that may arise during non-interactive installation.

## Categories for Application Recovery Kits

Category	Description
Application Suite	A group of recovery kits that protect applications such as SAP and IBM MQ.
Networking	A group of recovery kits that protect network services in the cloud such as EC2 and Route53.
Database	A group of recovery kits that protect database applications, including, but not limited to, Oracle, PostgreSQL, and MaxDB.
File Sharing	A group of recovery kits that protect file sharing services such as NFS and Samba.
Mail Server	A group of recovery kits that protect email services such as Postfix.
Storage	A group of recovery kits that protect data storage methods, including, but not limited to, DataKeeper (replication), Device Mapper (DM) Multipath (DMMP), and Network Attached Storage (NAS).
Web Server	A group of recovery kits that protect web services such as Apache.

## 13.2.4. Upgrading LKSSP

---

LifeKeeper Single Server Protection (SSP) can be upgraded to future releases while maintaining existing hierarchies.

 **Note:** Only the previous two generations of LifeKeeper Single Server Protection can be upgraded to the latest version. If you are upgrading from older versions, you will need to uninstall the old version and reinstall LifeKeeper Single Server Protection. Instead of uninstalling the old version, you can also upgrade to the latest version after upgrading the older version to either of the previous two generations.

 **Note:** If using [lkbackup](#) during your upgrade, refer to [the known issues of lkbackup](#) for further information.

1. Upgrade your Linux operating system before upgrading SSP if necessary.
2. Upgrade LifeKeeper referring to [How to Use Setup Scripts](#).

## 13.2.5. Obtaining and Installing the License for LKSSP

LifeKeeper Single Server Protection requires a unique license for each server. The license is a run-time license, which means that you can install LifeKeeper Single Server Protection without it, but the license must be installed before you can successfully start and run the product.

The Installation script installs the Licensing Utilities package which obtains and displays all of the available Host IDs for your server during the initial install of your LifeKeeper Single Server Protection Software. Once your licenses have been installed the utility will return the Entitlement ID if it is available or the Host IDs if it is not.

**Note:** Host IDs, if displayed will always be based on the MAC address of the NICs.

Any LifeKeeper Single Server Protection licenses obtained from the SIOS Technology Corp. Licensing Operations Portal will contain your Entitlement ID and will not be locked to a specific node in the cluster. The Entitlement ID (Authorization Code) which was provided with your LifeKeeper Single Server Protection Software, is used to obtain the permanent license required to run the LifeKeeper Single Server Protection Software. The process is illustrated below.



**Note:** Each software package requires a license for each server.

Perform the following steps to obtain and install your licenses for each server:



lkkeyins and specify the filename (including full path) to the file.

6. Repeat on all additional servers. You must install a license on the other SSP server(s) using the unique Host ID for each server.
7. Restart LifeKeeper for Linux.

The Host ID used by the License Key Installer utility is obtained from the LifeKeeper Single Server Protection for Linux server's primary network interface card (NIC). SSP for Linux will check for a valid license each time it starts. If your SSP for Linux server should require a NIC replacement in the future that would cause the Host ID to change, then the next time SSP for Linux is stopped, a License Rehost must be performed before starting it again. Log in to the [SIOS Technology Corp. Licensing Operations Portal](#) and select **License Support, List Licenses, Action, Rehost**.

 **Note:** A rehost can be performed four times per six-month period (per Activation ID) without contacting support.

## 13.2.6. Resource Policy Management

---

Resource Policy Management in LifeKeeper Single Server Protection (SSP) provides behavior management of resource local recovery and failover. Resource policies are managed with the **lkpolicy** command line tool (CLI).

### LifeKeeper SSP Recovery Behavior

LifeKeeper SSP is designed to monitor individual applications and groups of related applications, periodically performing local recoveries or notifications when protected applications fail. Related applications, by example, are hierarchies where the primary application depends on lower-level storage or network resources. When an application or resource failure occurs, the default behavior is:

1. **Local Recovery:** First, attempt local recovery of the resource or application. An attempt will be made to restore the resource or application on the local server without external intervention. If local recovery is successful, then LifeKeeper SSP will not perform any additional action.
2. **Failover:** Second, if a local recovery attempt *fails* to restore the resource or application (or the *recovery kit* monitoring the resource has no support for local recovery), then a **failover** will be initiated (see Failover in the Standard Policies section below).

Please see [LifeKeeper Single Server Protection Fault Detection and Recovery Scenario](#) for more detailed information about our recovery behavior.

### Custom and Maintenance-Mode Behavior via Policies

LifeKeeper SSP supports the ability to set additional policies that modify the default recovery behavior. There are four policies that can be set for individual resources (see the section below about precautions regarding individual resource policies) or for an entire server. **The recommended approach is to alter policies at the server level.**

The available policies are:

#### Standard Policies

- **Failover** – For LifeKeeper SSP this policy setting can be used to turn on/off resource failover (which results in a reboot).
- **LocalRecovery** – LifeKeeper SSP by default, will attempt to recover protected resources by restarting the individual resource or the entire protected application prior to performing a failover (which would be a reboot). This policy setting can be used to turn on/off local recovery.
- **TemporalRecovery** – Normally, LifeKeeper SSP will perform local recovery of a failed resource. If local recovery fails, LifeKeeper SSP will perform a reboot. If the local recovery succeeds, failover (which would be a reboot) will not be performed.

There may be cases where the local recovery succeeds, but due to some irregularity in the server, the local recovery is re-attempted within a short time; resulting in multiple, consecutive local recovery attempts. This may degrade availability for the affected application.

To prevent this repetitive local recovery/failure cycle, you may set a temporal recovery policy. The temporal recovery policy allows an administrator to limit the number of local recovery attempts (successful or not) within a defined time period.

*Example:* If a user sets the policy definition to limit the resource to three local recovery attempts in a 30-minute time period, LifeKeeper SSP will failover(reboot) when a third local recovery attempt occurs within the 30-minute period.

Defined temporal recovery policies may be turned on or off. When a temporal recovery policy is off, temporal recovery processing will continue to be done and notifications will appear in the log when the policy would have fired; however, no actions will be taken.

 **Note:** It is possible to disable failover and/or local recovery with a temporal recovery policy also in place. This state is illogical as the temporal recovery policy will never be acted upon if failover or local recovery are disabled.

## Meta Policies

The “meta” policies are the ones that can affect more than one other policy at the same time. These policies are usually used as shortcuts for getting certain system behaviors that would otherwise require setting multiple standard policies.

- **NotificationOnly** – This mode allows administrators to put LifeKeeper SSP in a “monitoring only” state. **Both local recovery and failover(reboot) of a resource (or all resources in the case of a server-wide policy) are affected.** The user interface will indicate a **Failure** state if a failure is detected; but no recovery or failover(reboot) action will be taken. **Note:** The administrator will need to correct the problem that caused the failure manually and then bring the affected resource(s) back in service to continue normal LifeKeeper SSP operations.

## Important Considerations for Resource-Level Policies

Resource level policies are policies that apply to a specific resource only, as opposed to an entire resource hierarchy or server.

*Example:*

app

- IP

- file system

In the above resource hierarchy, app depends on both IP and file system. A policy can be set to disable local recovery or failover of a specific resource. This means that, for example, if the IP resource's local recovery fails and a policy was set to disable failover of the IP resource, then the IP resource will not fail over or cause a failover of the other resources. However, if the file system resource's local recovery fails and the file system resource policy does not have failover disabled, then the entire hierarchy will failover causing a reboot.



**Note:** It is important to remember that resource level policies apply only to the specific resource for which they are set.

This is a simple example. Complex hierarchies can be configured, so care must be taken when setting resource-level policies.



**See known issue.** (*“Resources removed in the wrong order during failover”*)

## The lkpolicy Tool

The `lcpolicy` tool is the command-line tool that allows management (querying, setting, removing) of policies on servers running LifeKeeper SSP. `lcpolicy` supports setting/modifying policies, removing policies and viewing all available policies and their current settings. In addition, defined policies can be set on or off, preserving resource/server settings while affecting recovery behavior.

The general usage is :

```
lcpolicy [-list-policies | -get-policies | -set-policy | -remove-policy] <name
value pair data...>
```

The `<name value pair data...>` differ depending on the operation and the policy being manipulated, particularly when setting policies. For example: Most on/off type policies only require `-on` or `-off` switch, but the temporal policy requires additional values to describe the threshold values.

## Example lcpolicy Usage

### Authenticating With Local and Remote Servers

The `lcpolicy` tool communicates with LifeKeeper SSP servers via an API that the servers expose. This API requires authentication from clients like the `lcpolicy` tool. The first time the `lcpolicy` tool is asked to access a LifeKeeper SSP server, if the credentials for that server are not known, it will ask the user for credentials for that server. These credentials are in the form of a username and password and:

1. Clients must have LifeKeeper SSP admin rights. This means the username must be in the `lkadmin` group according to the operating system's authentication configuration (via `pam`). It is **not** necessary to run as **root**, but the root user can be used since it is in the appropriate group by

default.

2. The credentials will be stored in the credential store so they do not have to be entered manually each time the tool is used to access this server.

The lkpolicy tool communicates with LifeKeeper SSP servers via an API that the servers expose. This API requires authentication from clients like the lkpolicy tool. The first time the lkpolicy tool is asked to access a LifeKeeper SSP server, if the credentials for that server are not known, it will ask the user for credentials for that server. These credentials are in the form of a username and password and:

See [Configuring Credentials for LfieKeeper for Linux](#) for more information on the credential store and its management with the `credstore` utility.

An example session with lkpolicy might look like this:

```
[root@thor49 ~]# lkpolicy -l -d v6test4
Please enter your credentials for the system 'v6test4'.
Username: root
Password:
Confirm password:
Failover
LocalRecovery
TemporalRecovery
NotificationOnly
[root@thor49 ~]# lkpolicy -l -d v6test4
Failover
LocalRecovery
TemporalRecovery
NotificationOnly
[root@thor49 ~]#
```

## Listing Policies

```
lkpolicy -list-policy-types
```

## Showing Current Policies

```
lkpolicy -get-policies
```

```
lkpolicy -get-policies tag=\*
```

```
lkpolicy -get-policies -verbose tag=mysql\* # all resources starting with
mysql
```

```
lkpolicy -get-policies tag=mytagonly
```

## Setting Policies

```
lkpolicy -set-policy Failover -off
```

```
lkpolicy -set-policy Failover -on tag=myresource
```

```
lkpolicy -set-policy Failover -on tag=\*
```

```
lkpolicy -set-policy LocalRecovery -off tag=myresource
```

```
lkpolicy -set-policy NotificationOnly -on
```

```
lkpolicy -set-policy TemporalRecovery -on recoverylimit=5 period=15
```

```
lkpolicy --set-policy TemporalRecovery --on --force recoverylimit=5 period=10
```

## Removing Policies

```
lkpolicy -remove-policy Failover tag=steve
```



**Note:** *NotificationOnly is a policy alias. Enabling NotificationOnly is the equivalent of disabling the corresponding LocalRecovery and Failover policies.*

## 13.2.7. Verifying LifeKeeper Single Server Protection Installation

---

You can verify that the LifeKeeper Single Server Protection packages were installed correctly by entering the following at the command line:

```
rpm -V <package name>
```

**Note:** If the package is installed correctly, no output will be displayed by this command.

To perform a query from the command line, type

```
rpm -qi <package name>
```

**Note:** The expected output for this command is the package information.

# 13.3. LifeKeeper Single Server Protection for Linux Technical Documentation

---

## About LifeKeeper Single Server Protection for Linux

LifeKeeper Single Server Protection (SSP) allows for application monitoring in a single node configuration (i.e., no cluster requirements or restraints). Single node environments may be physical or virtual (vSphere, KVM, Amazon EC2). LifeKeeper SSP is built on the proven and stable architecture of SIOS LifeKeeper. LifeKeeper SSP provides superior application monitoring and can perform recovery of failed applications and system infrastructure items (e.g., NFS share, IP address, File System). If an application cannot be recovered for some reason, LifeKeeper SSP will initiate a restart of the node via a system reboot or via a VMware HA restart for VMware virtual machines configured for VM and Application Monitoring.

**Note:** Because LifeKeeper SSP is built using the SIOS LifeKeeper technology, you will see references to LifeKeeper throughout the documentation as well as references to information found in the LifeKeeper for Linux documentation for topics common to both products. When referencing these common topics the following subject items neither apply to nor support LifeKeeper SSP:

- Clustering
- Communication Paths
- Shared Storage (requirements, configuration, ...)
- Extending / Unextending resource hierarchies
- Quorum/Witness
- Application Recovery Kits
  - Recovery Kit for EC2
  - LVM Recovery Kit
  - MD Recovery Kit
  - Route53 Recovery Kit
  - SAP Recovery Kit
  - SAP HANA Recovery Kit
  - VMDK as Shared Storage Recovery Kit
- Storage Kits (DR, DMMP, HDLM, LVM, MD, PPATH and NEC SPS)

**Note:** Without the underlying storage kits in LifeKeeper SSP, steps must be taken to ensure the devices required to mount protected file systems are activated during system boot (e.g. if the file system is mounted on a logical volume the volume must be in the active state before LifeKeeper SSP starts).

- Resource/Machine failovers (by default with LifeKeeper SSP these result in a node restart)

- Resource switchovers
- Switchable IP Addresses (with LifeKeeper SSP Switchable IP addresses are required for some protected applications but since there is only a single node no switching actually takes place)

**Note:** When operating on Amazon EC2, assign a secondary private IP address to the NIC using the Amazon EC2 Management Console prior to creating the IP resource. Next, create the IP resource as the private IP address on the NIC that is using the LifeKeeper GUI client. An Elastic IP can now be associated with the IP resource.

For more information on the SIOS LifeKeeper product, on which LifeKeeper SSP is built, please see the [LifeKeeper for Linux documentation](#) for the common release number. This documentation will provide detailed information on resource hierarchy creation, resource types, states and relationships, Graphical User Interface (GUI), as well as common and advanced tasks.

## 13.3.1. Documentation and Training

A complete reference providing instructions for installing, configuring, administering and troubleshooting SIOS LifeKeeper Single Server Protection for Linux is available in the [LifeKeeper Single Server Protection Documentation](#). The following sections cover every aspect of SIOS LifeKeeper Single Server Protection for Linux:

Section	Description
<a href="#">Introduction and Installation</a>	Provides useful information for planning and setting up your LifeKeeper Single Server Protection environment, installing and licensing LifeKeeper Single Server Protection and configuring the LifeKeeper graphical user interface (GUI).
<a href="#">Administration</a>	Discusses server-level tasks such as editing server properties and creating resources and resource-level tasks such as editing, extending or deleting resources.
<a href="#">User's Guide</a>	Contains detailed information on the LifeKeeper GUI, including the many tasks that can be performed within the LifeKeeper GUI.
<a href="#">Troubleshooting</a>	Describes known issues and restrictions and suggests solutions to problems that may be encountered during installation, configuration and/or use of SIOS LifeKeeper Single Server Protection for Linux.
<a href="#">Recovery Kits</a>	Contains planning and installation instructions as well as administration, configuration and user information for the Optional Recovery Kits that allow LifeKeeper Single Server Protection to manage and control specific applications.

### Training

LifeKeeper Single Server Protection training is available through SIOS Technology Corp. or through your LifeKeeper Single Server Protection provider. Contact your sales representative for more information.

### Technical Support

As a SIOS Technology Corp. customer with a valid Support contract, you are entitled to access the new [SIOS Technology Corp. Support Self-Service Portal](#).

[The SIOS Technology Corp. Support Self-Service Portal](#) offers you the following capabilities:

- Search our Solution Knowledge Base to find solutions to problems and answers to questions
- Always on 24/7 service with the SIOS Technology Corp. Support team to:
- **Log a Case** to report new incidents
- **View Cases** to see all of your open and closed incidents
- **Review Top Solutions** provides information on the most popular problem resolutions being

viewed by our customers.

Contact SIOS Technology Corp. Support at [support@us.sios.com](mailto:support@us.sios.com) to set up and activate your Self-Service Portal account.

You can also contact SIOS Technology Corp. Support at:

1-877-457-5113 (Toll Free)

1-803-808-4270 (International)

Email: [support@us.sios.com](mailto:support@us.sios.com)

## 13.3.2. Intergration with VMware HA

---

As noted in the Introduction Section, LifeKeeper Single Server Protection is designed for use in both physical and virtual environments. When LifeKeeper SSP is installed in a VMware VM the HA features of VMware can be used in conjunction with LifeKeeper SSP to monitor and recover from any protected resource or node failure. To enable these features see [Enabling VMware HA Integration with LifeKeeper Single Server Protection](#).

## 13.3.3. Administration

### LifeKeeper Single Server Protection Administration Overview

LifeKeeper Single Server Protection does not require administration during operation. LifeKeeper Single Server Protection works automatically to monitor protected resources and to perform the specified recovery actions if a fault should occur. You use the LifeKeeper Single Server Protection GUI in these cases:

- **Resource and hierarchy definition.** LifeKeeper Single Server Protection provides these interface options:
  - LifeKeeper Single Server Protection GUI.
  - LifeKeeper Single Server Protection command line interface.
- **Resource monitoring.** The LifeKeeper Single Server Protection GUI provides access to resource status information and to the LifeKeeper Single Server Protection logs.
- **Manual intervention.** You may need to stop servers or specific resources for maintenance or other administrative actions. The LifeKeeper Single Server Protection GUI provides menu functions that allow you to bring specific resources in and out of service. Once applications have been placed under LifeKeeper Single Server Protection, they should be started and stopped only through these LifeKeeper Single Server Protection interfaces. Starting and stopping LifeKeeper Single Server Protection is done through the command line only.

See [Administration Tasks](#), [GUI Tasks](#), and [Maintenance Tasks](#) in the LifeKeeper for Linux documentation for detailed instructions on performing administration, configuration, and maintenance operations including the creation of resource hierarchies.

 **Note:** All LifeKeeper executable scripts and programs run via the command line require super user authority. A super user (granted permissions by running the “su” or “sudo” command) is able to execute LifeKeeper commands. However, SIOS Technology Corp has tested executing LifeKeeper commands via the root user only.

## 13.3.3.1. Enabling VMware HA Integration with LifeKeeper Single Server Protection

---

By default LifeKeeper Single Server Protection integration with VMware HA is disabled when installed on a VMware VM. To enable integration requires the following steps:

1. Installation of VMware tools in the LifeKeeper Single Server Protection VM.
2. Edit */etc/default/LifeKeeper* and change the VMware HA integration tunable `HA_DISABLE` value from 1 to 0.
3. Set the VMware HA parameter *das.iostatsinterval* to 0. This setting disables the I/O stats interval for VM Monitoring sensitivity. The default is 120 seconds.
4. Restart LifeKeeper Single Server Protection. If LifeKeeper Single Server Protection is currently running, it must be stopped and restarted to pick up the above change in */etc/default/LifeKeeper*.

## 13.3.3.2. Enabled VMware HA Fault Detection and Recovery Scenario

---

The ability to provide detection and alarming for problems within an application is critical to building the best total fault resilient solution. Since every specific application varies on the mechanism and format of failures, no one set of generic mechanisms can be supplied. In general, however, many application configurations can rely on the Core system error detection provided within LifeKeeper Single Server Protection. This topic demonstrates the power of LifeKeeper Single Server Protection's core facilities.

Below is a recovery scenario to demonstrate how LifeKeeper Single Server Protection provides fault detection and recovery when an application fails.

1. LifeKeeper Single Server Protection will first attempt recovery by trying to restart the application.
2. If the recovery succeeds, the application should continue to run normally.
3. If the recovery attempt fails:
  - a. If the LifeKeeper Single Server Protection recovery attempt fails, and LifeKeeper Single Server Protection is installed in a VMware guest OS with HA enabled (HA\_DISABLE=0 in /etc/default/LifeKeeper), then LifeKeeper Single Server Protection will trigger VMware HA by withholding the heartbeat that LifeKeeper Single Server Protection sends down to the Application Monitoring Interface. VMware HA will then respond by restarting the server.
  - b. If the LifeKeeper Single Server Protection recovery attempt fails, and LifeKeeper Single Server Protection is not installed in a VMware guest OS or is installed in a VMware guest OS but has HA disabled (HA\_DISABLE=1 in /etc/default/LifeKeeper), then a system reboot will be forced.

## 13.3.3.2.1. Maintaining a LifeKeeper Single Server Protection Protected System

---

When performing system or application maintenance on a LifeKeeper Single Server Protection-protected server, you should either stop LifeKeeper Single Server Protection monitoring or place the protected resources into maintenance mode. This will stop LifeKeeper Single Server Protection from interfering with the system and application maintenance tasks by disabling both application recovery and triggering of VMware HA failure events.

To stop and restart LifeKeeper Single Server Protection:

1. **Stop LifeKeeper Single Server Protection.** Stop LifeKeeper Single Server Protection using the `/opt/LifeKeeper/bin/lkstop -f` command. The resources will remain running but will no longer be monitored by LifeKeeper Single Server Protection. Any failure will have to be handled manually.
2. **Perform maintenance.** Perform the necessary maintenance.
3. **Start LifeKeeper Single Server Protection.** Use the command `/opt/LifeKeeper/bin/lkstart` to start LifeKeeper Single Server Protection. Your resources are now protected.

**Alternatively**, place the resources in maintenance (a.k.a., notification only) mode:

1. **Place all resources in maintenance mode.** Use the command `/opt/LifeKeeper/bin/lkpolicy -s NotificationOnly —On`. Resources will not be recovered and VMware HA failure events will not be triggered.
2. **Perform maintenance.** Perform the necessary maintenance.
3. **Turn off maintenance mode.** Use the command `/opt/LifeKeeper/bin/lkpolicy -s NotificationOnly —Off`. Resources are now protected.

## 13.3.3.3. LifeKeeper Single Server Protection Heartbeat with VMware HA

---

The LifeKeeper Single Server Protection heartbeat is the signal sent to VMware HA (every 10 seconds if running in a VMware guest OS and if HA is enabled) indicating that the protected applications are OK. If an application fails, LifeKeeper Single Server Protection will first attempt to recover the application. If recovery fails, LifeKeeper Single Server Protection will withhold the heartbeat, which instructs VMware HA to reboot the VM.

## 13.3.3.4. Quick Service Protection (QSP) Recovery Kit

---

The QSP Recovery Kit provides a simplified method to protect the OS service. With the QSP Recovery Kit, users can easily create a LifeKeeper resource instance to protect an OS service provided that service can be started and stopped by the OS service command. The service can also be protected via the Generic Application Recovery Kit but the use of that kit requires code development, whereas the QSP Recovery Kit does not. Also, by creating a dependency relationship, protected services can be started and stopped in conjunction with the application that requires the service.

The QSP Recovery Kit quickCheck can only perform simple health checks (using the “status” action of the service command). QSP doesn’t guarantee that the service is provided or the process is functioning. If complicated starting and/or stopping is necessary, or more robust health checking operations are necessary, using a Generic Application is recommended.

For details, please refer to the following URL:

- [Quick Service Protection \(QSP\) Recovery Kit](#)

## 13.3.3.5. LifeKeeper API for Monitoring

---

### Introduction

The LifeKeeper API for Monitoring can obtain the operational status of LifeKeeper nodes and their protected resources by making status inquiries to the available nodes in the LifeKeeper cluster.

This API operates as CGI on lighttpd and the information that can be acquired by the lcdstatus command can be acquired via the API.

Information that can be obtained with the API:

- LifeKeeper node status is the node alive and processing or down
- Communication path status between nodes in the cluster, are communication path(s) up or down
- Status of protected resources

Please refer to [LifeKeeper API for Monitoring](#) for details.

## 13.3.3.6. Watchdog

Watchdog is a method of monitoring a server to ensure that if the server is not working properly, corrective action (reboot) will be taken so that it does not cause problems. Watchdog can be implemented using special watchdog hardware or using a software-only option.

\* **Note:** This configuration has only been tested with Red Hat Enterprise Linux Version 7. No other operating systems have been tested; therefore, no others are supported at this time.

### Components

- Watchdog timer – software driver or an external hardware component
- Watchdog daemon – rpm available through the Linux distribution
- LifeKeeper core daemon – installed with the LifeKeeper installation
- \* Health check script – Script to check the status of LifeKeeper SSP core



**LifeKeeper Interoperability with Watchdog**

Read the next section carefully. The daemon is designed to recover from errors and will reset the system if not configured carefully. Planning and care should be given to how this is installed and configured. This section is not intended to explain and configure watchdog, but only to explain and configure how LifeKeeper SSP interoperates in such a configuration.

### Configuration

The following steps should be carried out by an administrator with root user privileges. The administrator should already be familiar with some of the risks and issues with watchdog.

The health check script (LifeKeeper monitoring script) is the component that ties the LifeKeeper configuration with the watchdog configuration (`/opt/LifeKeeper/samples/watchdog/LifeKeeper-watchdog`). This script can monitor the basic parts of LifeKeeper core components.

1. If watchdog has been previously configured, enter the following command to stop it. If not, go to Step 2.

```
systemctl stop watchdog
```

2. Edit the watchdog configuration file (`/etc/watchdog.conf`) supplied during the installation of watchdog software.

- Modify test-binary:

```
test-binary = /opt/LifeKeeper/samples/watchdog/LifeKeeper-watchdog
```

- Modify test-timeout:

```
test-timeout = 5
```

- Modify interval:

```
interval = 7
```

The interval must be greater than or equal to the test-timeout value. The recommended value is between 5 and 10 because if the interval is too long, failure detection will be delayed.

3. Make sure LifeKeeper SSP has been started. If not, please refer to the [Starting LifeKeeper](#) topic.
4. Start watchdog by entering the following command:

```
systemctl start watchdog
```

5. To start watchdog automatically on future restarts, enter the following command:

```
systemctl enable watchdog
```

**Note:** Configuring watchdog may cause some unexpected reboots from time to time. This is the general nature of how watchdog works. If processes are not responding correctly, the watchdog feature will assume that LifeKeeper (or the operating system) is hung, and it will reboot the system (without warning).

## Uninstall

Care should be taken when uninstalling LifeKeeper. The above steps should be done in reverse order as listed below.

**!** **WARNING:** IF UNINSTALLING LIFEKEEPER BY REMOVING THE RPM PACKAGES THAT MAKE UP LIFEKEEPER, **TURN OFF WATCHDOG FIRST!** In Step 2 above, the watchdog config file was modified to call on the LifeKeeper-watchdog script; therefore, if watchdog is not turned off first, it will call on that script that is no longer there. An error will occur when this script is not found which will trigger a reboot. This will continue until watchdog is turned off.

1. Stop watchdog by entering the following command:

```
systemctl stop watchdog
```

2. Edit the watchdog configuration file (/etc/watchdog.conf) supplied during the installation of watchdog software.

- Modify test-binary and interval by commenting out those entries (add # at the beginning of each line):

```
#test-binary =  
#interval =
```

 **Note:** If interval was used previously for other functions, it can be left as-is

3. Uninstall LifeKeeper. See the [Removing LifeKeeper](#) topic.
4. Watchdog can now be started again. If only used by LifeKeeper, watchdog can be permanently disabled by entering the following command:

```
systemctl disable watchdog
```

## 13.3.3.7. LKCLI (LifeKeeper Command Line Interface)

---

LKCLI provides functions that can be performed with the LifeKeeper GUI through the command line interface. LifeKeeper provides export/import of communication paths and resource information which the GUI does not provide. Export/import functionality enables duplicating a created system and easy deployment from a testing environment to a production environment.

### Supported Environments

LKCLI is supported in the following environments:

- OS – All operating systems supported by LifeKeeper are supported.
- Application Recovery Kits – Except for those subject to the restrictions described below, only environments configured with the Application Recovery Kits in the “Supported ARK List” are supported. If unsupported ARK resources have been created on a node, the environment is not supported.

If you want to perform command line operations in an unsupported environment, please consider using the [Command Line Interface](#).

### Restrictions

For Single Server Protection for Linux, LKCLI is not available or supported with the following functions:

- Common option
  - [ --remote <str> ]
- Communication path operations
  - lkcli commpath ...
  - lkcli import (not available only for the creation of communication paths)
- Extending/unextending or switching resources
  - lkcli resource extend
  - lkcli resource unextend
  - lkcli resource eqv
  - lkcli resource reorder-priority
  - lkcli resource switchback
  - lkcli server ...
- Application Recovery Kits
  - SIOS DataKeeper for Linux
  - Recovery Kit for EC2
  - LVM Recovery Kit
  - MD Recovery Kit
  - Route53 Recovery Kit
  - SAP Recovery Kit

- SAP HANA Recovery Kit
  - VMDK as Shared Storage Recovery Kit
  - Multipath Recovery Kit (DMMP / HDLM / PPATH / NECSPS)
  - Ikcli export or Ikcli resource info command of the IP Recovery Kit
- If processing stops with the Ikcli export command or Ikcli resource, ensure that the hostname is configured properly.

You can use LKCLI with Single Server Protection for Linux in the same way as the LifeKeeper for Linux. See [LKCLI \(LifeKeeper Command Line Interface\)](#) for details.

## 13.3.4. Troubleshooting

---

This section contains restrictions and/or known issues open against LifeKeeper Single Server Protection.

## 13.3.4.1. Known Issues and Workarounds

Included below are the restrictions and/or known issues open against LifeKeeper Single Server Protection.

### Core

#### Bug 2257

**Access to LifeKeeper Single Server Protection and LifeKeeper for Linux nodes via credstore requires proper credstore key**

**Solution:** When storing credentials for a LifeKeeper Single Server Protection or LifeKeeper for Linux node using `credstore`, you must use the proper form of the hostname for the credstore credentials key (i.e. `credstore -k`):

For the LifeKeeper Single Server Protection plugin, `credstore` should be run using the hostname of the system as reported in the **Hostname:** field of the LifeKeeper Single Server Protection plugin display.

For LifeKeeper Single Server Protection, the hostname used to store credentials must be the same as the one you plan to use in the command line tool's (e.g., `lkipolicy`) `-d` argument. For example, if you want to run `lkipolicy -d mynode1`, then you must store credentials using `credstore -k mynode1`. You cannot store credentials using the FQDN in this case. If you do, you must run `lkipolicy -d FQDN`.

**Workaround:** If you've stored a default credential set (i.e., `credstore -k default`) that works for all your LifeKeeper Single Server Protection and/or LifeKeeper for Linux nodes, then you will not be affected by this issue.

#### Bug 2408

**HA heartbeat incorrectly enabled**

lkvmhad incorrectly enables the HA heartbeat after second resource failure

**Workaround:** Set `LKCHECKINTERVAL` in `/etc/default/LifeKeeper` greater than the VMware HA, VM Monitoring Failure Interval. **Note:** The `LKCHECKINTERVAL` default is 120 seconds. This is also the default 'low' monitoring sensitivity for VMware HA, VM Monitoring.

## GUI

### Refresh problem with LifeKeeper Single Server Protection GUI

The GUI may occasionally scramble the resource tree (i.e., resource dependencies may not be shown correctly).

**Workaround:** Perform a refresh of the GUI.

## Apache

### Apache resource creation fails

#### Example of Error message:

Error: valid\_http\_root: Since "/usr/sbin/httpd" is shareable on "/usr", "/etc/httpd" must be also

#### Cause:

Due to a defect, files in mount point "/"(root) cannot be detected appropriately.

For example, if "/etc/httpd" is in a same filesystem as the mount point "/", a resource creation will fail.

#### Workaround:

Mount one of the below workarounds to avoid this issue.

- (a) Transfer such as "/etc/httpd" under the other mount point.
- (b) Mount " /etc" to such as " /dev/sdb1".

## Oracle

### Bug 2387

**Cannot create an Oracle hierarchy on root file system in LifeKeeper Single Server Protection environment**

**Workaround:** Using the following procedure, copy Oracle to a new file system.

Create a new disk large enough for Oracle data (e.g. /dev/sdb). (Note: You can size up /oracle directory to get an idea how big this should be; multiply by at least 50% to allow for logs)

Using gdisk, create a new partition on that disk.

```
gdisk /dev/sdb
```

Make a file system.

```
mkfs -t ext3 /dev/sdb1
```

Mount this file system (example using /mnt/oracle).

```
mkdir /mnt/oracle
```

```
mount /dev/sdb1 /mnt/oracle
```

Stop Oracle, Listener.

Copy Oracle to new file system.

```
cd /oracle
```

```
cp -a * /mnt/oracle
```

**(Note:** This step may take some time based on the amount of data)

Unmount the new file system.

```
umount /mnt/oracle
```

Mount the new file system over /oracle.

```
mount /dev/sdb1 /oracle
```

Start Listener and then Oracle.

## SAP

### Bug 2388

**For SAP, hierarchies cannot be created using the GUI**

**Workaround:** Use the command line option to create hierarchies. However, at the end of the command line, specify the number 76 as follows:

- \$LKROOT/lkadm/subsys/appsuite/sap/bin/create <primary sys> <tag> <SAP SID> <SAP Instance> <switchback type> <IP Tag> <Protection Level> <Recovery Level> <Additional SAP Dependents> 76

See [Setting Up SAP from the Command Line](#) for further command line information.

Also, refer to [Known Issues and Restrictions](#).

## 13.4. Application Recovery Kits

---

LifeKeeper Single Server Protection for Linux Application Recovery Kits (ARKs) include tools and utilities that allow SSP to manage and control a specific application. The following optional recovery kits are available with this release of SSP.

[Apache Recovery Kit Administration Guide](#)

[DB2 Recovery Kit Administration Guide](#)

[IP Recovery Kit Administration Guide](#)

[MySQL Recovery Kit Administration Guide](#)

[WebSphere MQ Recovery Kit Administration Guide](#)

[NAS Recovery Kit Administration Guide](#)

[NFS Recovery Kit Administration Guide](#)

[Oracle Recovery Kit Administration Guide](#)

[PostgreSQL Recovery Kit Administration Guide](#)

[Postfix Recovery Kit Administration Guide](#)

[Samba Recovery Kit Administration Guide](#)

[SAP MaxDB Recovery Kit Administration Guide](#)

[Sybase Recovery Kit Administration Guide](#)

# 14. Product Support Schedule

For customers under an annual support agreement, SIOS Technology provides full support for its products for three years from their General Availability date. This support period is extended in situations where simple upgrade paths do not exist to later versions of SIOS products.

The table below shows products whose End of Support dates have been set. If these products are deployed within your IT infrastructure, we strongly recommend that you begin planning to upgrade to later versions. You can see these latest versions and their documentation on our [website](#). If you are using an earlier version than what is listed below, it is no longer supported.

Product	Product Release Date	End of Support
See the <a href="#">LifeKeeper for Linux Support Matrix</a> for supported configuration details.		
LifeKeeper for Linux v9.4.0	October 7, 2019	October 31, 2022
LifeKeeper for Linux v9.4.1	January 14, 2020	January 31, 2023
LifeKeeper for Linux v9.5.0	May 11,2020	May 31, 2023
LifeKeeper for Linux v9.5.1	October 6, 2020	October 31, 2023
LifeKeeper for Linux v9.5.2	August 11, 2021	August 31, 2024
LifeKeeper for Linux v9.6.0	November 3, 2021	November 30, 2024
LifeKeeper for Linux v9.6.1	April 20, 2022	April 30, 2025
LifeKeeper Single Server Protection for Linux v9.4.0	October 7, 2019	October 31, 2022
LifeKeeper Single Server Protection for Linux v9.4.1	January 14, 2020	January 31, 2023
LifeKeeper Single Server Protection for Linux v9.5.0	May 11,2020	May 31, 2023
LifeKeeper Single Server Protection for Linux v9.5.1	October 6, 2020	October 31, 2023
LifeKeeper Single Server Protection for Linux v9.5.2	August 11, 2021	August 31, 2024
LifeKeeper Single Server Protection for Linux v9.6.0	November 3, 2021	November 30, 2024
LifeKeeper Single Server Protection for Linux v9.6.1	April 20, 2022	April 30, 2025
See the <a href="#">LifeKeeper for Windows Support Matrix</a> for supported configuration details.		
LifeKeeper for Windows v8.7.0	November 22, 2019	November 30, 2022
LifeKeeper for Windows v8.7.1	March 26, 2020	March 31, 2023
LifeKeeper for Windows v8.7.2	November 20, 2020	November 30, 2023
LifeKeeper for Windows v8.8.0	July 6, 2021	July 31, 2024
LifeKeeper for Windows v8.8.1	November 16, 2021	November 30, 2024

LifeKeeper for Windows v8.8.2	January 13, 2022	January 31, 2025
LifeKeeper Single Server Protection for Windows v8.7.0	November 22, 2019	November 30, 2022
LifeKeeper Single Server Protection for Windows v8.7.1	March 26, 2020	March 31, 2023
LifeKeeper Single Server Protection for Windows v8.7.2	November 20, 2020	November 30, 2023
LifeKeeper Single Server Protection for Windows v8.8.0	July 6, 2021	July 31, 2024
LifeKeeper Single Server Protection for Windows v8.8.1	November 16, 2021	November 30, 2024
LifeKeeper Single Server Protection for Windows v8.8.2	January 13, 2022	January 31, 2025
See the <a href="#">DKCE Support Matrix</a> for supported configuration details.		
DataKeeper for Windows v8.7.0	November 22, 2019	November 30, 2022
DataKeeper for Windows v8.7.1	March 26, 2020	March 31, 2023
DataKeeper for Windows v8.7.2	November 20, 2020	November 30, 2023
DataKeeper for Windows v8.8.0	July 6, 2021	July 31, 2024
DataKeeper for Windows v8.8.1	November 16, 2021	November 30, 2024
DataKeeper for Windows v8.8.2	January 13, 2022	January 31, 2025
DataKeeper Cluster Edition for Windows v8.7.0	November 22, 2019	November 30, 2022
DataKeeper Cluster Edition for Windows v8.7.1	March 26, 2020	March 31, 2023
DataKeeper Cluster Edition for Windows v8.7.2	November 20, 2020	November 30, 2023
DataKeeper Cluster Edition for Windows v8.8.0	July 6, 2021	July 31, 2024
DataKeeper Cluster Edition for Windows v8.8.1	November 16, 2021	November 30, 2024
DataKeeper Cluster Edition for Windows v8.8.2	January 13, 2022	January 31, 2025