



SIOS Protection Suite for Windows

8.6.4 — Last update: 2019/11/15

SIOS TECHNOLOGY CORP.

Table of Contents

SIOS Protection Suite for Windows	16
SIOS Protection Suite for Windows Quick Start Guide	17
SIOS Protection Suite for Windows Release Notes	21
SIOS Protection Suite Installation Guide	41
Planning Your SIOS Protection Suite Environment.....	42
Planning Server Communication	43
Recovery Kit Requirements	46
Storage and Adapter Requirements	47
Verifying Server Specifications	48
Setting Up Your SIOS Protection Suite Environment	50
Configuring Your Storage	51
DNS Resource Requirements	53
Installing and Setting Up Database Applications	54
Safe Creation of Shared Disk Volume Instances	55
Verifying Network Configuration.....	57
Installing SIOS Protection Suite	62
Core Software.....	63
Installing Core	64
Licensing	68
Installing Localized Language Supplement	72
Silent Installation of SIOS Protection Suite	73
Third Party Product Files	75
Application Directory Anomaly	77
Uninstalling SIOS Protection Suite for Windows.....	78
Upgrading SIOS Protection Suite	81
Repair.....	84
Starting LifeKeeper	85
SIOS Protection Suite for Windows Technical Documentation	87
Introduction	88
Core	91
MSSQL.....	93
Communication Paths.....	94
Communication Path Types.....	95
SIOS Protection Suite Heartbeat	97

Heartbeat Interval	98
Safety Check	99
Resource Hierarchies	100
Hierarchy Relationships	101
Resource Hierarchy Information	102
Resource States	103
Shared Equivalencies	104
Configuration	105
SIOS Protection Suite Configuration Steps	106
Active-Active Grouping	107
Active-Standby Grouping	108
Intelligent Versus Automatic Switchback	110
SIOS Protection Suite Configurations	111
Common Hardware Components	112
System Grouping Arrangements	114
Configuring SPS for Multibyte Language Encodings	115
Administration	116
Administrator GUI Tasks	118
Editing Server Properties	119
Set Server Shutdown Strategy	120
Server Properties	121
Disabling Automatic Failover	127
Creating a Communication Path	129
Deleting a Communication Path	132
Working With Resource Hierarchies	133
Creating Resource Hierarchies	134
Creating a DNS Resource Hierarchy	136
Creating a File Share List Resource Hierarchy	140
Creating a Generic Application Resource Hierarchy	142
Creating a LAN Manager Resource Hierarchy	145
Creating a Volume Resource Hierarchy	147
Creating an IP Address Resource Hierarchy	150
IP Local Recovery Scenario	153
Editing Resource Priorities	154
Incomplete Resource Priority Modification	156
Editing Resource Properties	158
Extending Resource Hierarchies	159
Extending a DNS Resource Hierarchy	160
Extending a File Share Resource Hierarchy	161

Extending a Generic Application Resource Hierarchy	162
Extending a LAN Manager Resource Hierarchy	163
Extending a Volume Resource Hierarchy	164
Extending an IP Address Resource Hierarchy	165
Unextending a Hierarchy	166
Adding a Resource Dependency	167
Removing a Resource Dependency	168
Deleting a Hierarchy from All Servers	169
Man Pages	170
LCD – Miscellaneous LCD Programs	171
lcdrcp	173
lcdrecover	174
lcdremexec	176
lcdsync	177
lkstart	178
lkstop	179
LCDI Applications	180
app_create	181
app_list	182
app_remove	183
LCDI Instances	184
ins_list	186
Initialization Strategy	188
Initial State	189
ins_create	190
ins_gettag	191
ins_remove	192
ins_setas	194
ins_setchkint	195
ins_setin	196
ins_setit	197
ins_setlocalrecover	198
ins_setst	199
LCDI-relationship	200
dep_create	202
dep_list	203
dep_remove	204
eqv_create	205
eqv_list	206

eqv_remove	207
LCDI-resource_type	208
typ_create	210
typ_list	211
typ_remove	212
LCDI-systems	213
sys_create	215
sys_getds	216
sys_getst	217
sys_list	218
sys_remove	219
LifeKeeper Flags	220
flg_create	221
flg_list	222
flg_remove	223
flg_test	224
LCDI Flags	225
lk_chg_value	228
lk_err	231
perform_action	233
sendevent	238
volume	242
LKSUPPORT	244
Setting Browser Security Parameters	245
IP Local Recovery	246
Overview of SIOS Protection Suite Event Forwarding via SNMP	247
Java Upgrade	251
User Guide	253
LifeKeeper GUI	255
GUI Overview	256
Status Table	258
Properties Panel	259
Output Panel	260
Message Bar	261
Toolbars	262
Global Toolbar	263
Resource Context Toolbar	265
Server Context Toolbar	266
Menus	267

Resource Context Menu	268
Server Context Menu	270
File Menu	271
Edit Menu – Resource.....	272
Edit Menu – Server	274
View Menu	275
Help Menu	277
LifeKeeper GUI Server and Client Components.....	278
Running the SIOS Protection Suite Web Client.....	280
Running the GUI Application on a SIOS Protection Suite Server	283
LifeKeeper GUI User Accounts	284
Configuring GUI Users	285
Common Tasks.....	287
Connecting To A Cluster.....	288
Disconnecting From a Cluster.....	290
Viewing Connected Servers.....	291
Viewing The Status Of A Server	292
Viewing Server Log Files	293
Viewing Server Properties	296
Viewing Resource Tags and IDs.....	297
Viewing the Status of Resources	298
Viewing Resource Properties.....	300
Viewing Message History	301
Expanding and Collapsing A Resource Hierarchy Tree.....	302
Operator Tasks	304
Bringing A Resource In Service	305
Taking a Resource Out Of Service	306
Taking Volume Resources In and Out Of Service	307
Volume Shadow Copy (VSS).....	311
Volume Locking for Shared SCSI Volumes	312
Advanced Topics	313
LifeKeeper Configuration Database (LCD).....	314
LCD Directory Structure	315
Diagram of LCD Directory.....	316
LCD Configuration Data.....	317
LCD Resource Types	318
Resources Subdirectories.....	319
LCDI Commands	321
LifeKeeper Communications Manager (LCM)	322

Communication Status Information	323
Maintenance Tasks.....	325
Starting and Stopping LifeKeeper	326
Managing IP Resources	327
Managing DNS Resources	329
Displaying List of Protected File Shares	331
EditFileShareResource Utility	332
Transferring Resource Hierarchies	333
Performing Offline Maintenance On A Shared Disk.....	334
Maintaining a SIOS Protection Suite Protected System	335
Configuring Generic Application Scripts.....	336
Maintaining a Resource Hierarchy	339
Recovering After a Failover	340
Uninstalling SPS for Windows	341
Performing CHKDSK on a protected volume.....	344
Data Replication	346
Monitoring Replicated Volume Resources	347
Replication Settings.....	349
Performing Actions on Replicated Volumes	351
What is Split-Brain.....	353
Split-Brain Recovery.....	355
DataKeeper.....	359
Introduction.....	360
User Interface	362
DataKeeper Components	364
DataKeeper Service Log On ID and Password Selection.....	366
Understanding Replication.....	372
SIOS DataKeeper Intent Log.....	373
Relocation of Intent Log	375
Resynchronization.....	377
Synchronous and Asynchronous Mirroring	379
Read and Write Operations	386
Volume Considerations	388
Specifying Network Cards for Mirroring	389
Performance Monitor Counters.....	390
DataKeeper Configuration.....	397
Sector Size.....	398
Network Bandwidth.....	399
Network Adapter Settings	401

DataKeeper Service Log On ID and Password Selection	403
Firewall Configurations	410
High-Speed Storage Best Practices	416
Configuration of Data Replication From a Cluster Node to External DR Site	418
Performance Tuning	419
Disable “Automatically manage paging file size for all drives”	420
WAN Considerations	424
Initial Synchronization of Data Across the LAN or WAN.....	425
Compression	429
Bandwidth Throttle	430
DataKeeper Administration	431
DataKeeper Event Log Notification	432
Primary Server Shutdown	434
Secondary Server Failures	435
Extensive Write Considerations	436
CHKDSK Considerations	437
DKSUPPORT	438
DKHEALTHCHECK	439
Event Log Considerations.....	440
Using Disk Management.....	441
Registry Entries	442
Using EMCMD with SIOS DataKeeper.....	454
Mirror State Definitions.....	458
BREAKMIRROR.....	459
CHANGEMIRRORENDPOINTS	460
CHANGEMIRRORTYPE.....	466
CLEARBLOCKTARGET	469
CLEARSNAPSHOTLOCATION	470
CLEARSWITCHOVER	471
CONTINUEMIRROR	472
CREATEJOB.....	473
CREATEMIRROR	474
DELETEJOB	476
DELETELOCALMIRRORONLY	477
DELETEMIRROR.....	478
DROPSNAPSHOT	479
GETBLOCKTARGET.....	480
GETCOMPLETEVOLUMELIST	481
GETCONFIGURATION	482

GETEXTENDEDVOLUMEINFO	483
GETJOBINFO	484
GETJOBINFOFORVOL	485
GETMIRRORVOLUMEINFO	486
GETMIRRORVOLINFO	487
GETREMOTEBITMAP	489
GETRESYNCSTATUS	490
GETSERVICEINFO	492
GETSNAPSHOTLOCATION	493
GETSOURCEMIRROREDVOLUMES	494
GETTARGETMIRROREDVOLUMES	495
GETVOLUMEDRVSTATE	496
GETVOLUMEINFO	497
ISBREAKUSERREQUESTED	499
ISPOTENTIALMIRRORVOL	500
LOCKVOLUME	502
MERGETARGETBITMAP	503
PAUSEMIRROR	504
PREPARETOBECOMETARGET	505
READREGISTRY	506
REGISTERCLUSTERVOLUME	508
RESTARTVOLUMEPIPE	509
RESYNCMIRROR	510
SETBLOCKTARGET	511
SETCONFIGURATION	512
SETSNAPSHOTLOCATION	514
STOPSERVICE	515
SWITCHOVERVOLUME	516
TAKESNAPSHOT	517
UNLOCKVOLUME	518
UPDATEJOB	519
UPDATEVOLUMEINFO	520
Using DKPwrShell with SIOS DataKeeper	522
New-DataKeeperMirror	523
New-DataKeeperJob	525
Remove-DataKeeperMirror	527
Remove-DataKeeperJob	528
Add-DataKeeperJobPair	529
Get-DataKeeperVolumeInfo	531

DataKeeper User Guide	532
Getting Started	533
Disk-to-Disk	534
One-to-One	536
One-to-Many (Multiple Targets).....	538
Many-to-One	540
N-Shared-Disk Replicated to One	542
N-Shared-Disk Replicated to N-Shared-Disk	544
N-Shared-Disk Replicated to Multiple N-Shared-Disk Targets	546
Setting Up SIOS DataKeeper	548
Setting Up SIOS DataKeeper	549
Connecting to a Server.....	550
Disconnecting from a Server	551
Creating a Job.....	552
Configuring Mirrors	553
Creating a Mirror	554
Creating Mirrors With Shared Volumes.....	556
Safe Creation of a Shared-Storage Volume Resource	560
Creating Mirrors With Multiple Targets	562
Switchover and Failover with Multiple Targets	564
Working With Jobs.....	568
Jobs	569
Renaming a Job	571
Deleting a Job	572
Reassigning a Job.....	573
Switching Over a Mirror.....	574
Working With Mirrors	576
Managing Mirrors	577
Pause and Unlock	578
Continue and Lock	579
Partial Resync.....	580
Break	581
Resync.....	582
Deleting a Mirror	583
Replacing a Target.....	584
DataKeeper Volume Resize	585
Mirror Properties	588
Changing the Compression Level of an Existing Mirror.....	590
Working With Shared Volumes	592

Managing Shared Volumes	593
Adding a Shared System.....	595
Removing a Shared System	596
Using Microsoft iSCSI Target With DataKeeper on Windows 2012	597
Installation of the iSCSI Target.....	599
Mirror Creation and Cluster Configuration	601
Creation of iSCSI Virtual Disks	605
Setup of iSCSI Initiator on Windows 2012	608
DataKeeper Notification Icon	610
DataKeeper Target Snapshot	613
Using SIOS DataKeeper Standard Edition To Provide Disaster Recovery For Hyper-V Virtual Machines.....	626
FAQs	639
Awareness of Windows Filenames and Directory Names.....	640
AWS Issues and Workarounds	641
Change Mirror Endpoints.....	642
Change Mirror Type.....	643
Create a Mirror and Rename Job and Delete Job Actions Grayed Out.....	644
Data Transfer Network Protocols	645
Delete and Switchover Actions Grayed Out	646
Deleting a Mirror	647
Error Messages Log	648
Inability to Create a Mirror	649
Network Disconnect.....	650
Reclaim Full Capacity of Target Drive.....	652
Resize or Grow Mirrored Volumes	653
Split-Brain FAQs.....	654
Stop Replication Between Source and Target.....	657
Using Volume Shadow Copy	658
Volumes Unavailable for Mirroring	659
DataKeeper Troubleshooting	660
Known Issues and Workarounds	663
Access to Designated Volume Denied	665
DataKeeper Volume Not Available as Cluster Resource Type	666
Failed to Create Mirror	667
Hyper-V Host Cluster Error	668
Live Migration Failure.....	671
MaxResyncPasses Value	672
Mirroring with Dynamic Disks	673

New Resources Offline But Unlocked	674
Server Login Accounts and Passwords Must Be Same on Each Server in the Cluster	675
System Event Log – Create Mirror Failed in the GUI	677
Unable to Determine Previous Install Path	678
User Interface – Failed to Create Mirror	679
User Interface – Shows Only One Side of the Mirror	680
Windows Server 2012 Specific Issues	681
Windows Server 2012 MMC Snap-in Crash	682
Windows Server 2012 iSCSI Target Role Does Not Support Dynamic Disks	684
Windows Server 2012 NIC Teaming Issue	685
Windows Server 2016 Specific Issues	705
Occasional Job Creation Failure	706
WSFC 2016 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager	707
Restrictions	709
Bitlocker Does Not Support DataKeeper.....	710
CHANGEMIRRORENDPOINTS	711
CHKDSK	712
DataKeeper Volume Resize Restriction	713
Directory for Bitmap Must Be Created Prior to Relocation	714
Duplicate IP Addresses Disallowed Within a Job	715
Intensive I-O with Synchronous Replication.....	716
Resource Tag Name Restrictions	717
Troubleshooting	718
Applet Troubleshooting	720
Error When Attempting to Run SIOS Protection Suite Command From Command Prompt	721
Firewall	722
GUI Error Messages	723
GUI Network Related – Initial Connection to Server Failed (Error 117)	726
GUI Network Related – Long Connection Delays on Windows Platforms	727
GUI Network Related – NoRouteToHostException Message Generated During Connection Attempt	729
GUI Network Related – Unknown Host Exception Message Generated During Connection Attempt	730
GUI Server Troubleshooting	734
Health Check Timeouts.....	735
Incomplete Resource Creation	737
Installation – Access is Denied	738
IP Resource Create Issue	739
Java Mixed Signed and Unsigned Code Warning	740

LANMAN Name May Be Seen Twice in Browse List.....	742
Licensing – Licensed Recovery Kit Resource Fails to Come In Service	743
Licensing – License Key Not Found	744
LifeKeeper GUI Issues with Java 1.8.x.....	745
SIOS Protection Suite Web Client May Lock Up.....	746
SIOS Protection Suite Web Client May Lock Up – Multiple Changes Made to Existing Hierarchy ...	747
New Evaluation License Key Error	748
Recovering From a SIOS Protection Suite Server Failure in a 1 × 1 Configuration	749
Recovering Out-of-Service Hierarchies	757
Remove Hangs During Recovery Kit Uninstall	758
Replicated Volume Switchover Failure.....	759
Resource Tag Name Restrictions.....	760
Restore and Health Check Account Failures	761
SQL 2008	762
SQL Server Reporting Services (MSSQLSERVER).....	763
Two-Server Cluster Issue	764
Unknown User Name or Bad Password	765
Web Client Troubleshooting	766
IIS Resource Hierarchy Creation Error.....	767
Application Recovery Kits	768
SIOS Protection Suite Microsoft SQL Server Recovery Kit Introduction.....	769
SQL Server Services	770
Recovery Kit Requirements	771
SQL Server Installation.....	772
Recovery Kit Installation	773
Removal	774
Installing and Configuring SQL Server with SIOS Protection Suite.....	775
Installation – Shared Storage Systems	776
Using a Non-Admin Local System Account on Target System	779
Installation – Replicated Storage Systems	780
Additional Setup Tasks for Extended Configurations	782
Creating the SQL Hierarchy	783
Extending a SQL Hierarchy	786
Unextending a SQL Hierarchy	787
Deleting a SQL Hierarchy	788
SQL Server Configuration Considerations.....	789
Managing a SQL Server Configuration	793
Process Config Menu	794

Database Config Menu	795
Admin Actions Menu.....	796
Manage User Menu	797
Testing Your SQL Resource Hierarchy	798
SQL Server Hierarchy Administration.....	799
Access Via Protected Communication Paths	800
User Resource Name For Remote Access.....	801
Reserve Volumes For Exclusive SQL Use	802
Understand Manual Switchover Limitations in SQL Server Environment	803
Running Microsoft SQL Management Tools.....	804
Start and Stop SQL Server Only Through SIOS Protection Suite	805
Adding Microsoft SQL Server Volumes	806
Recovering From Databases in Suspect State After a Switchover	807
Pausing Microsoft SQL Server (MSSQLServer).....	808
Configuring SQL Server to Connect Using the Switchable IP Address	809
Maintaining SQL Server Login and Passwords	810
Monitoring Your SQL Hierarchy	811
Troubleshooting	812
Create and Extend Failure	813
Extend of a SQL Resource Fails.....	814
SQLAgent Service Fails to Start	815
Connecting ODBC Clients to Named Instances of SQL Server	816
SIOS Protection Suite PostgreSQL Server Recovery Kit Introduction	817
PostgreSQL Server Installation	818
Installation and Configuration Details – Adding LifeKeeper to an Existing PostgreSQL	
Configuration.....	819
Installation and Configuration – Adding PostgreSQL to an Existing LifeKeeper Configuration...	823
Additional Setup Tasks for Extended Configurations	826
Server Configuration Considerations.....	828
Creating the PostgreSQL Hierarchy.....	833
Extending a PostgreSQL Hierarchy	837
Unextending a PostgreSQL Hierarchy	839
Deleting a PostgreSQL Hierarchy	840
Testing Your PostgreSQL Resource Hierarchy	841
PostgreSQL Server Hierarchy Administration.....	842
Access Via Protected Communication Paths	843
Reserve Volumes For Exclusive PostgreSQL Use	844
Start and Stop PostgreSQL Server Only Through SIOS Protection Suite	845
Creating and Protecting Additional PostgreSQL Database Clusters	846

PostgreSQL Administrative Login	848
Configuration for Unattended Connections	849
Configuring the Postmaster Port Argument	851
Monitoring Your PostgreSQL Hierarchy	854
Troubleshooting	855
Create Fails	856
Restore Fails	858
Tunable Settings for the PostgreSQL Recovery Kit	860
SIOS Protection Suite Oracle Recovery Kit Introduction	862
Oracle Overview	863
Oracle Services	865
Resource Hierarchy for Oracle	866
Recovery Kit Requirements	867
Installation	868
Recovery Kit Installation	869
Installing and Configuring LifeKeeper with Oracle	870
Before Installing Oracle	871
Installing Oracle	873
Oracle Resource Configuration Tasks	876
Creating an Oracle Hierarchy	878
Extending an Oracle Hierarchy	880
Unextending an Oracle Hierarchy	881
Deleting an Oracle Hierarchy	882
Manage Oracle Database Configuration	884
Testing Your Oracle Resource Hierarchy	886
Oracle Hierarchy Administration	887
Oracle Hierarchy Administration Guidelines	888
LifeKeeper Oracle Recovery Kit Recovery Variables	890
Updating Oracle Username and Password for LifeKeeper	891
Manually Configure Oracle 11g DB Console	892
Oracle Troubleshooting Tips	893
Create Hierarchy Failed	894
Bring In Service Failed	895
Oracle TNSListener Service is Not Started or Stopped as it Should Be	896
Server Not Responding	897
Remote Users Cannot Log In ORA 12504 or ORA 12514 or ORA 12541	898
SIOS Protection Suite Microsoft Internet Information Services Recovery Kit Introduction	899
IIS Overview	900
Installation	902

Hardware and Software Requirements	903
Kit Installation.....	905
Kit Removal.....	906
Configuration Definitions and Restrictions	907
IIS Configuration Considerations	909
Installing and Configuring IIS with LifeKeeper.....	913
Install and Configure SIOS Protection Suite and Recovery Kits	914
Install and Configure Microsoft IIS on All Servers	916
Install and Configure SMTP Virtual Server.....	918
IIS Configuration Definitions and Restrictions	920
IIS Required Roles and Role Services and Features	921
IIS Active Active	922
Primary and Backup Designations	924
Naming Restrictions	925
Identical Primary Backup Web Sites	926
Configuring Secure Servers.....	927
IIS Configuration	928
Document Content Location	929
Use Different Volume for Multiple IIS Sites	930
IIS Resource Configuration Tasks.....	931
Create an IIS Resource Hierarchy	932
Extend an IIS Resource Hierarchy.....	933
Delete an IIS Resource Hierarchy.....	934
Unextend Your IIS Hierarchy	935
Testing Your IIS Resource Hierarchy.....	936
Performing a Manual Switchover from the GUI	937
Recovery Operations.....	938
IIS Hierarchy Administration	939
Modifying Quick Check Interval, Deep Check Interval and Local Recovery	940
Manual Switchover	941
IIS Failover	942
Protecting FTP Sites with Non-Anonymous Login	943
Using an FTP Login Script.....	944
Disabling the FTP Deep Check Process	945
Changing LifeKeeper Microsoft IIS Recovery Kit Configuration.....	946
Removing Microsoft IIS	947
Changing the Network Interface Card (NIC).....	948
IIS Troubleshooting.....	949
Symptoms and Solutions	950

SIOS Protection Suite for Windows Support Matrix 952

Product Support Schedule..... 958

SIOS Protection Suite for Windows

Your information resource for SIOS Protection Suites and DataKeeper Cluster Edition

SIOS Technology Corp. maintains documentation for all supported versions of SIOS Protection Suites. We welcome your suggestions and feedback. To help us continue to improve our documentation, please complete our brief [Documentation Feedback Survey](#).

SIOS Protection Suite for Windows Quick Start Guide

This topic provides step-by-step instructions for installing and configuring the SIOS Protection Suite (SPS) for Windows. SIOS Protection Suite is a software bundle that includes [DataKeeper \(DK\)](#), [LifeKeeper \(LK\)](#) and the optional [SIOS Protection Suite Microsoft SQL Server Recovery Kit](#).

The series of steps includes links into the documentation that describe each step in detail.

Prerequisites and Installation

1. Read the [SIOS Protection Suite for Windows Release Notes](#) for late breaking information.
2. [Firewall Configurations](#) – Make sure you understand what ports must be opened on any firewalls.
3. [Network Bandwidth](#) – If replicating across a WAN, it is critical that a rate of change analysis be done to ensure there is adequate bandwidth.
4. DataKeeper is a block-level volume replication solution and requires that each server have additional volume(s) (other than the system drive) that are the same size. Please review [Volume Considerations](#) for additional information regarding storage requirements.
5. Review the section [Understanding Replication](#) to help understand how DataKeeper works and the difference between synchronous and asynchronous replication.
6. Ensure that you have the correct version and/or capacity of all components for each SIOS Protection Suite server by reviewing [Verifying Server Specifications](#). **Note:** All servers within a cluster should be running the same version of Windows.
7. Plan your communication paths for your cluster. For optimal performance, multiple communication paths must be defined. Review [Planning Server Communication](#) for additional information.
8. LK supports either shared storage or replicated storage. Review [Storage and Adapter Requirements](#) and [Configuring Your Storage](#) for additional information.
9. In order for SIOS Protection Suite to function properly, it is important that your networking is configured as described in [Verifying Network Configuration](#). In particular, you must **make sure that**

your Public NIC is at the top of the binding order on all of the cluster nodes.

10. Before installing any applications (Oracle, SQL Server, Exchange), please review [Installing and Setting Up Database Applications](#).
11. Installing SIOS Protection Suite is as straight forward as running the setup executable. The setup received from SIOS will go through the process of installing LK, DK and any optional ARKs. At the end of the setup, you will be prompted to enter your license key. Evaluation customers should use the time limited key that was supplied by SIOS. If you purchased the software, follow the instructions in [Obtaining and Installing the License](#) to complete your installation. For complete details, see the [SIOS Protection Suite for Windows Installation Guide](#).
12. When using the SIOS Protection Suite (LifeKeeper and DataKeeper) in a Workgroup environment, the LifeKeeper Service must use the same account (ID and Password) as the DataKeeper Service on each system [DataKeeper Service Log On ID and Password Selection](#).

Configuration

Once you have installed SIOS Protection Suite and ensured that your networking and storage are configured as described in the [Prerequisites](#) section, it is time to configure SIOS Protection Suite to protect your business critical applications. The detailed steps to configure protection are found in [SIOS Protection Suite Configuration Steps](#).

A summary of the steps minimally required are described below with links for more information.

1. Create your [communication paths](#). Communication paths carry cluster information between nodes and also heartbeats that are used to detect when entire systems fail.
2. Create your [resource heirarchy](#). Depending upon the application being protected, create one or more of the following resources. Click on the resource type for detailed instructions on resource creation.
 - [Volume](#) – A volume resource can be a replicated volume or a shared physical disk. A volume resource(s) is used in just about every cluster configuration.
 - [IP](#) – An IP resource is used for client redirection between cluster nodes in the same subnet.
 - [DNS](#) – A DNS resource is used for client redirection between cluster nodes that are in different subnets.
 - [File Share](#) – A file share resource is used to protect and recover Windows file shares. The file

share resource assumes that a [Volume](#) resource has already been created.

- [LAN Manager](#) – A LAN Manager resource is used in environments that require NetBIOS name resolution. Alternatively, if DNS resolution is functioning properly, you can create an “A” record in DNS and have it resolve to the [IP](#) resource for client redirection.
 - [Generic Application](#) – The GenApp resource allows users to protect any application by writing a simple start, stop and recover script. The GenApp resource comes with an example script that easily allows the user to provide protection for any Windows Service.
 - [SIOS Protection Suite Microsoft SQL Server](#) – The SQL Recovery Kit allows users to build high availability and disaster recovery solutions for their SQL implementation. The entire SQL Server cluster creation process is also detailed in this [online video screen capture demonstration](#).
3. After creating your resources, the dependencies may need to be adjusted as described in [Adding a Resource Dependency](#). Creating dependencies ensures that resources always fail over together and also ensures that dependent resources start and stop in the proper order.
 4. At this point, you have a functioning cluster. Don't forget to take the time to read the rest of the [documentation](#) to discover other advanced features that you may wish to take advantage of in your environment.

Management

While there are many things you may wish to do with your cluster, one of the first things you will want to do is test a manual switchover to ensure that your application can be moved from the active cluster to the standby server. To do this, right-click on the top level resource on the standby system and choose [In Service](#). This will initiate a manual switchover.

Once you verify that a manual switchover completes as expected, you may wish to test [Local Recovery](#). If SIOS Protection Suite detects a failure of a resource, it will attempt to recover that resource locally if [Local Recovery](#) is enabled. One easy way to test this is to stop the service of a protected application. Once the [Quick Check](#) runs, it will detect the failure and automatically restart the service.

The final test you will want to perform is a system recovery. If the active server fails completely, the standby server will detect the failure as it stops receiving [heartbeats](#) over the [communication path\(s\)](#). Once the timeout period expires (heartbeat interval X max missed heartbeats), the standby server will start recovering the protected resources. The best way to test system recovery is to pull the power cord on the active server.

Simply shutting down the active server does not necessarily cause failover; it all depends on how the [Shutdown Strategy](#) is set. Once the standby server completes the recovery, turn the failed server back on, and it will come online as the new standby server.

For more information on SIOS Protection Suite Management, please refer to the [SIOS Protection Suite Technical Documentation](#).

Troubleshooting

Use the following resources to help troubleshoot issues:

- [DataKeeper Troubleshooting](#) and [SIOS Protection Suite Troubleshooting](#) sections
- For customers with a support contract – <http://us.sios.com/support/overview/>
- For evaluation customers only – [Pre-sales support](#)

SIOS Protection Suite for Windows Release Notes

SIOS Protection Suite for Windows

Release Notes

Version 8.6.4

(Version 8 Update 6 Maintenance 4)

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.

To maintain the quality of our publications, we welcome your comments on their accuracy, clarity, organization and value.

Introduction

This information is provided for the person who installs, configures and/or administers the SIOS Protection Suite (SPS) for Windows product and contains important information such as version requirements, last-minute changes to instructions and procedures, product restrictions and known issues. It is important that you review this document before installing and configuring the SIOS Protection Suite software.

SIOS Protection Suite Product Descriptions

SIOS Protection Suite for Windows is a software bundle that integrates high availability clustering and data replication functionality to protect mission-critical data and applications and includes DataKeeper (DK), LifeKeeper (LK) and optional Recovery Kits.

LifeKeeper for Windows

LifeKeeper for Windows continues SIOS Technology Corp.'s tradition of providing world-class reliability for mission critical applications. LifeKeeper for Windows leverages over a decade of experience with high availability platforms by providing customers the ability to cluster multiple servers in order to monitor and restore their applications. In the event of a failure, LifeKeeper recovers all network interfaces, data and applications. Recovery occurs automatically and is transparent to clients, thus minimizing downtime and loss of business.

LifeKeeper for Windows enables continuous operations during planned downtime as well as in the event of a system or application failure. With LifeKeeper, the amount of downtime required for common maintenance tasks and upgrades is significantly reduced or eliminated.

DataKeeper for Windows

SIOS DataKeeper is a highly optimized host-based replication solution which ensures your data is replicated as quickly and as efficiently as possible from your source server across the network to one or more target servers.

New Features of SIOS Protection Suite for Windows v8

Feature	Description
New in This Release	
Windows 2019 Support	SIOS Protection Suite for Windows now supports Windows 2019 (Restrictions)
Delay Writes when Write Queue limit is reached	Introduced the BlockWritesOnLimitReached tunable, which will delay writes when a mirror's queue reaches the HighWater or ByteLimit, rather than pausing the mirror and going into a resync state.
PostgreSQL 11	SIOS Protection Suite for Windows now supports PostgreSQL 11
Oracle 18c	SIOS Protection Suite for Windows now supports Oracle 18c
General maintenance	Bug Fixes .
New in Version 8.6.3	
Queue Current Age	This Performance Monitor Counter value is the age of the oldest write request in the write queue.
General maintenance	Bug Fixes.
New in Version 8.6.2	
Integration with SIOS iQ	DataKeeper now includes the DataKeeper Signal package to deliver events to SIOS

	iQ.
General maintenance	Bug Fixes
New in Version 8.6.1	
Microsoft SQL Server 2017 Support	The SIOS Protection Suite SQL Server Recovery Kit supports Microsoft SQLServer 2017.
General Maintenance	See Bug Fixes.
New in Version 8.6	
VSS Provider	SIOS VSS Provider has been disabled by default.
PostgreSQL Support	Added a new SIOS Protection Suite for Windows core Application Recovery Kit that provides protection for PostgreSQL database clusters.
Windows 2016 Support	SIOS Protection Suite for Windows now supports Windows 2016.
Tunable Write Queue Byte Limit	Users can specify the maximum number of bytes that can be allocated for the write queue of a mirror by changing the WriteQueueByteLimitMB registry value.
General maintenance	Bug Fixes.
New in Version 8.5	
CHANGEMIRRORTYPE	This EMCMD command is used to change the mirror type of a mirror that is part of a DataKeeper job.
Microsoft SQL Server 2016 Support	The SIOS Protection Suite SQL Server Recovery Kit supports Microsoft SQLServer 2016.
Tunable bitmap block size	Users can modify the effective size of an entry in the DataKeeper intent log (bitmap) by changing the BitmapBytesPerBlock registry value.
General maintenance	Bug Fixes.
New in Version 8.4	
Target Bitmap File	Target writes are now tracked in a persistent target bitmap file.
General maintenance	Bug Fixes.
New in Version 8.3	
DataKeeper Notification Icon	The DataKeeper Notification Icon shows a summary of your DataKeeper mirrors in the Windows Notification Tray. In addition to the display functions, the DataKeeper Notification Icon also serves as a shortcut to managing your DataKeeper mirrors.
Oracle 12c and Oracle 12c Standard Edition 2	SIOS Protection Suite for Windows Version 8.3 and later supports Oracle 12c and Oracle 12c Standard Edition 2 (excluding pluggable database).
Powershell cmdlet support	Powershell cmdlets that can be used to create job(s), create mirror(s), remove job(s), remove mirror(s) and fetch information about a volume used in DataKeeper (New-DataKeeperMirror, New-DataKeeperJob, Remove-DataKeeperMirror, Remove-DataKeeperJob, Add-DataKeeperJobPair, Get-DataKeeperVolumeInfo).

mirrorcleanup.cmd	This command will remove all remnants of a mirror for a selected volume on the local system only and should only be run when recommended by SIOS Support.
DKHEALTHCHECK	Support status and issue identification tool. Provides command line interface for basic mirror status and problem detection.
General maintenance	Bug Fixes.
New in Version 8.2.1	
General maintenance	Bug Fixes.
New in Version 8.2	
General maintenance	Bug Fixes.
New in Version 8.1	
Microsoft Windows Server 2012 R2 Support	LifeKeeper Version 8.1 and later supports Windows Server 2012 R2.
General maintenance	Bug Fixes.
New in Version 8.0.1	
General maintenance	Bug Fixes.
New in Version 8.0	
General maintenance	Bug Fixes.

Discontinued Features of SIOS Protection Suite for Windows v8

Feature	Description
Discontinued in This Release	
Data Rewind	Removed Rewind feature from SPS for Windows

Bug Fixes

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

	Description
4189	BSOD on shutdown if volume goes offline just after boot time on target node
4211	emcmd SETSNAPSHOTLOCATION command does not validate the snapshot path

4367	Ensure the PostgreSQL executable path lists all installed versions as choices and that help clarifies the need to select the correct PostgreSQL version
4381	Briefly delay writes to allow resync to complete
4392	Ensure a warning message is displayed when PostgreSQL is running without the port option
4394	Correct typo in message warning that the PostgreSQL datadir is not located on a protected volume
4408	Do not allow extension of a PostgreSQL resource to a target node that does not have PostgreSQL installed
4419	GenApp – remove the recover script should work and should disable local recovery
4420	Memory leak – shared disk taken offline / online
4421	Language Pack welcome screen has text cut off
4427	Correct typo in message warning that the PostgreSQL port is already being used
4431	SQL restore reports “/Binn not found”
4432	Fixed typo in PostgreSQL canextend informational message
4433	Implement BlockWritesOnLimitReached tunable
4437	License Manager not working properly in some locales (date calculations)
4440	License Manager incorrectly reports invalid host ID
4441	Removed rewind from DataKeeper
4442	Minor typos in the comment sections of a number of PostgreSQL scripts
4444	Removed ips_arpd.exe from LifeKeeper for Windows
4445	Warning message during install about NIC IPV4 metrics
4455	Prevent product installation on Windows 2019 systems with Secure Boot enabled
4458	Possible data corruption if unlocked target is rebooted

Product Requirements

Operating System

Important: SIOS Technology Corp. recommends that users use Domain accounts that have local administrator privileges on all servers running SIOS Protection Suite. If local accounts are being used, the user names and passwords must match on all servers running SIOS Protection Suite. This recommendation is for all editions and all platforms.

Note: All servers within a cluster should be running the same version of Windows.

Product	Operating Systems	Additional Software
SIOS Protection Suite (Server Components)	See the SPS Support Matrix	n/a
SIOS Protection Suite (User Interface)	See the SPS Support Matrix	MMC 3.0 – download from: https://download.cnet.com/Microsoft-Management-Console-3-0-for-Windows-XP-KB907265/3000-2206_4-10741230.html
Virtual Environments	<p>The operating system versions listed above are supported for guests running on the following virtual platforms:</p> <ul style="list-style-type: none"> • Amazon EC2 (AWS) • VMware vSphere 4.0 or later • Microsoft Hyper-V Server 2008 R2 or later • Citrix XenServer 5.5 or later • KVM with Kernel 2.6.32 or later 	
64-bit versions (x64, no Itanium) of all of the listed OS platforms are supported		

Requirements for Windows 2008 R2, 2012, 2012 R2 and 2016

While installing SIOS Protection Suite on Windows 2008 R2, a dialog box will prompt whether the installer should make the system configuration changes described below. If the installer is not allowed to make these changes, they will need to be made manually after installation is complete.

- [Windows Firewall](#)
- The **Distributed Link Tracking Client** must be **disabled**

For systems running SIOS Protection Suite for Windows and Microsoft FTP Service 7.5 for IIS 7.0, Windows 2008 R2 or later is required.

In addition, if your Windows 2008 R2, 2012, 2012 R2 and 2016 servers are not in a domain, the Local Security policy setting “**Network Access: Let Everyone permissions apply to anonymous users**” must be enabled. If the servers are in a domain, then this setting is not required.

SIOS Protection Suite Requirements

The following table shows requirements applicable to the SIOS Protection Suite core and recovery kits.

Core	Requirement(s)
SIOS Protection Suite License	One license is required for every server on which SIOS Protection Suite runs. This applies to both physical and virtual servers.
LAN Manager Recovery Kit	Requires the “ File and Print Sharing for Microsoft Networks ” component (lanmanserver) to be installed on the Windows server. NetBIOS must also be enabled. Otherwise, the LAN Manager resource will not come in service.
Memory Requirements	The minimum memory requirement for a system supporting SIOS Protection Suite for Windows is based on the memory requirements for the operating system being used. Additional memory is required to run user applications in addition to that required for SIOS Protection Suite.
GUI	<p>Ports:</p> <p>SIOS Protection Suite uses Port 82 for Remote Method Invocation (RMI) communication between the GUI server and client.</p> <p>The LifeKeeper GUI uses Port 81 for its administration web server which should be different from any public web server. This is used by the GUI when run as a Java applet on a remote client.</p> <p>In the event of conflict with an existing application, these ports can be changed by editing the RMI_PORT or WEB_PORT entries in the SIOS\LIFEKEEPER\JAVAGUI\SERVER registry key.</p>

Optional Recovery Kits

All optional SIOS Protection Suite Recovery Kits require a software license key in order to function with SIOS Protection Suite.

Kit Name	Versions/Requirement(s)
Microsoft SQL Server Recovery Kit	See the SPS Support Matrix
Oracle Recovery Kit	See the SPS Support Matrix
PostgreSQL Recovery Kit	See the SPS Support Matrix

GUI Requirements, Platforms and Browsers

LifeKeeper requires that the Java Runtime Environment (JRE) be installed on each server. The 32-bit Windows JRE 1.8.0_101 is installed with the SIOS Protection Suite Core software. JRE 1.8.0_101 has been fully tested with the LifeKeeper GUI Server and GUI Application components.

SIOS Protection Suite can be administered from a system outside the SIOS Protection Suite cluster by running the SIOS Protection Suite web client. Included in the following table is a list of the supported platforms and browsers for the SIOS Protection Suite web client. As in the case of the server, we have tested with JRE 1.8.0_101, but we expect that the client will work equally well with future JRE updates. Updating the client JRE only affects that machine, so it is not as critical to test for safety as when you are updating the server JRE. We do recommend that you test updates before committing to them, and that you prepare to roll them back if a problem occurs.

Operating System	Internet Explorer 5.5+, 6.0	Internet Explorer 7.0, 8.0	Internet Explorer 9.0	Internet Explorer 10.0	Internet Explorer 11.0	Mozilla Firefox 1.5, 2	Mozilla Firefox 3
Windows 2016				X	X		
Windows 2012 R2				X	X		
Windows 2012				X	X		
Windows 2008 R2		X					
Windows 7		X					
Linux	N/A	N/A	N/A	N/A	N/A	X	X

Note: Other recent platforms and browsers will likely work with the SIOS Protection Suite web client, but they have not been tested by SIOS Technology Corp.

Installing and Removing SIOS Protection Suite for Windows

SIOS Protection Suite for Windows uses InstallShield to provide a standard installation interface with choices for **Typical**, **Compact** or **Custom installation**. See the [SIOS Protection Suite Installation Guide](#) for details about installing, removing or upgrading your SIOS Protection Suite software.

IMPORTANT

- Customizations made to SIOS Protection Suite scripts must be reapplied after upgrading to all releases of SIOS Protection Suite for Windows v8.
- Make sure you obtain the correct licenses; the old licenses will remain on the system and can be deleted with the license installer tool.
- SIOS does not support upgrading SIOS Protection Suite from more than one major version back. If upgrading from a version prior to LifeKeeper for Windows v7.x to SIOS Protection Suite for Windows v8, uninstall the old version of LifeKeeper and reinstall SIOS Protection Suite for Windows v8.

Technical Notes

lkstart

This program starts LifeKeeper on the current system if it is not currently running. lkstart modifies entries in the `%LKROOT%\etc\LKinit.config` file pertaining to the LifeKeeper daemons so that they will be respawned if they die.

The `-w` option, with `waitperiod` in seconds, can be used to change the timeout interval. Use the `-w` argument to specify a wait period before the startup.

The LifeKeeper service can also be started using the Microsoft Services MMC under Administrative Tools or from a command prompt using either `"sc start LifeKeeper"` or `"net start LifeKeeper"`.

Note: This program must be run from the console.

Running CHKDSK.EXE on SIOS Protection Suite Protected Volume

Microsoft recommends running the utility `chkdsk.exe` to check and correct file system or disk errors on volumes that have not been cleanly shut down. However, depending on the extent of errors, the utility may

take a very long time to complete. It may take several hours or even days for `chkdsk` to completely check the volume, or it may hang while checking the volume. Due to these reasons, SIOS Protection Suite does not run the `chkdsk` utility on protected volumes. SIOS Protection Suite does run the Microsoft utility `chkntfs.exe` to check whether a volume is dirty or not before bringing the volume in service. If a protected volume is found dirty, SIOS Protection Suite will log an error to the event log.

It is recommended that administrators periodically run `chkdsk` on SIOS Protection Suite protected volumes on the server where the volume resource(s) are in service. Administrators should take all the applications using the volume resource(s) out-of-service prior to running `chkdsk`.

Running CHKDSK.EXE During System Boot

Protection Suite protected volumes are typically not eligible for the `chkdsk` utility to run on them at system boot time because LifeKeeper and DataKeeper need to be able to lock the volumes. If a SIOS Protection Suite protected volume needs to be checked at boot time, the steps below can be performed on the active node.

For Mirrored Volumes or SDRS Volumes (shared at one site, replicated to a remote site)

1. `"%ExtMirrBase%\emcmd" . getconfiguration <drv>` (save the number reported on the first line of output for later use after reboot)
2. `"%ExtMirrBase%\emcmd" . setconfiguration <drv> 32`
3. `"%LKBIN%\lkstop" -f`
4. `sc stop ExtMirrSvc`
5. `sc config lifekeeper start= demand`
6. `sc config ExtMirrSvc start= demand`
7. `chkntfs /D`
8. `chkntfs /c <drv>`
9. `reboot`

Perform the following steps after reboot.

10. `sc config lifekeeper start= auto`
11. `sc config ExtMirrSvc start= auto`
12. `sc start ExtMirrSvc`
13. `"%ExtMirrBase%\emcmd" . setconfiguration <drv>` (number reported by `emcmd` `getconfiguration` in step 1).
14. `reboot`

For Shared Volumes

1. `"%LKBIN%\volume" -U <drv>`
2. `"%LKBIN%\lkstop" -f`
3. `chkntfs /c <drv>`
4. `reboot`

Perform following steps after reboot.

5. `"%LKBIN%\volume" -p <drv>`
6. `"%LKBIN%\lkstop" -f`
7. `"%LKBIN%\lkstart"`

For Replicated Volumes

1. `"%LKBIN%\lkstop" -f`
2. `chkntfs /D`
3. `chkntfs /c <drv>`
4. `reboot`

Communication Paths Over Fibre Channel

When building a SIOS Protection Suite cluster using shared storage, it is important to maintain working communication paths between the nodes in the cluster. Communication paths should be created using TCP communication protocols. Normally, TCP communication paths are built on Ethernet network devices. SIOS Protection Suite, however, can use any type of connection on which the TCP protocol can run. If a shared storage cluster is being created using a Fibre Channel SAN, it is possible (and desirable) to use the Fibre Channel SAN as a SIOS Protection Suite communication path.

QLogic provides a miniport driver and an IP driver for Windows that will allow a QLogic Fibre Channel storage adapter to also run the TCP/IP protocol. This, in effect, allows the QLogic Fibre Channel adapter to function both as a storage adapter and as a network adapter. Once this driver is in place, the QLogic card can be configured, as any network card would, using standard network configuration techniques.

QLogic's driver can be downloaded from the following web site:

http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/DefaultNewSearch.aspx

Using iSCSI Storage with SIOS Protection Suite

iSCSI storage can be used as shared storage and protected by SIOS Protection Suite. For shared storage environments, the iSCSI target device must be configured so that all server initiators have access to the disk. The vendor of the iSCSI storage device provides the interface and commands needed to configure the iSCSI device. A dependency on the Microsoft iSCSI Initiator service (MSiSCSI) should be added to the LifeKeeper service. This will ensure that the shared volume is available before LifeKeeper attempts to access the volume.

To create a dependency on MSiSCSI for the LifeKeeper service, use the registry editor "*regedt32.exe*" and select the subkey representing the LifeKeeper service under *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LifeKeeper*. The service key has a value name "DependOnService" with one value "EISM". Double-click the value name "DependOnService" to open for editing. When the dialog box appears, add the service name "MSiSCSI" for Microsoft iSCSI Initiator service on a new line and click **OK**.

To verify that the dependency was created, open *Administrative Tools->Services* MMC snap-in. Go to LifeKeeper service and double-click to bring up the "**Properties**" dialog. When the dialog box appears, go to "**Dependencies**" page and verify that "**Microsoft iSCSI Initiator**" service is listed along with "**LifeKeeper External Interface**" in the "**depends on**" field.

System Load Considerations for Quickcheck and Deepcheck

SIOS Protection Suite launches a separate thread to monitor each protected resource in the system. These threads operate independently of one another. Typically, system load from `Quickcheck` and `Deepcheck` script execution will be randomly distributed. SIOS Protection Suite also works to distribute resource monitoring load by skipping a `Quickcheck` execution whenever a `Deepcheck` for the same resource is scheduled to run at the same time. However, because the check load is randomly distributed, there will occasionally be peaks in system load from resource monitoring. The more protected resources in the system, the larger these peaks will be and the more often they may occur. The largest peak will occur when LifeKeeper is started and `Deepcheck` scripts for each active resource are first launched. If the server can handle this first load peak in a satisfactory way, then there should not be a performance problem later.

[VSS Shadow Copy](#)

Protection Suite support for VSS Shadow Copy requires that shadow copies must NOT be stored on the SIOS Protection Suite protected volumes. However, shadow copies may be saved on another non-protected volume.

Restrictions and Known Issues

Restrictions

Windows 2019

Microsoft has introduced new driver signing restrictions with Windows Server 2019. Windows 2019 systems that are configured with the Secure Boot UEFI BIOS feature enabled will not boot if any SIOS product is installed unless Windows Driver Signing checks are disabled. Disabling driver signing checks is an interactive task that has to be done each time the system is booted. To use any SIOS product with Windows 2019 either disable Secure Boot in the UEFI BIOS settings or use a legacy BIOS bootloader.

Cloud environments such as AWS and Azure do not currently enable Secure Boot and it is not anticipated that this will change. However, given the uncertainty of the current and future status of Secure Boot use in the various cloud environments, we do not recommend or support running SIOS products on any production Windows Server 2019 workloads hosted in a 3rd-party Cloud environment.

The restriction regarding Secure Boot will be addressed in a future SIOS product release.

Antivirus Software

Antivirus software may mistakenly detect LifeKeeper binaries as malware. You should configure an exclusion for the folder LKROOT (C:\LK by default) to prevent LifeKeeper utilities from being mistakenly put into quarantine.

SCVMM 2012

If using DataKeeper with SCVMM 2012, you must use SCVMM 2012 SP1.

Server with Microsoft Failover Cluster Installed

SIOS Protection Suite is not supported on Enterprise or DataCenter class servers with Microsoft Cluster Server or Microsoft Failover Cluster features installed. It should never be the case that two “Clustering” solutions are deployed on the same group of servers. As part of this restriction, SIOS Protection Suite communication paths will not function using IP addresses (169.254.xxx.xxx) that are hosted by the Microsoft Failover Cluster Virtual Adapters (Virtual NICs).

FAT File System Support

SIOS Protection Suite does not support protection for volumes using the FAT or FAT32 file systems.

Fault Tolerant Disk Sets

While SIOS Protection Suite replicated volumes are supported using Windows fault tolerant disk sets (Software RAID), SIOS Protection Suite shared volumes are not compatible with Windows fault tolerant disk sets. Fault tolerant disk sets must be set up with dynamic disks and dynamic disks cannot be shared between two systems.

File Share Recovery Kit

- The File Share Recovery Kit is supported only in an Active Domain environment, not in a Workgroup environment. File share permissions granted to local machine accounts, either in a workgroup environment or a domain environment, will not be preserved during failover because local User IDs are valid only on the local system where they originated; other systems will not recognize them. Even if two local User IDs are spelled the same way on two different machines, they will be treated as two different accounts and valid only on the system where they originated. Domain accounts, on the other hand, are identifiable and usable on any system in the domain.
- The File Share Recovery Kit will not work if more than 9999 file shares are defined on the system.

Any attempt to protect eligible file shares under SIOS Protection Suite will fail if the total number of user-defined shares exceeds 9999. This restriction also applies to editing file share resources. You will not be able to alter the list of protected shares if more than 9999 shares are defined on the system.

LAN Manager Recovery Kit

Microsoft supports LAN Manager functions only over the first IP address per network interface card (Microsoft bug SRX#9704116-48). This prohibits using LAN Manager functions over SIOS Protection Suite protected IP addresses. Therefore, the only way to switch over an alias computer name using the TCP/IP protocol is to allow dynamic IP#-to-LAN Manager name mapping for your clients. The recommended solution is to use a WINS server. You will need to make the SIOS Protection Suite servers (and all computers accessing the protected LAN Manager name) WINS clients of the same WINS server.

Low Virtual Memory Degrades System State

SIOS Protection Suite depends on memory being available when it is needed. If your system is reporting that it is low on virtual memory, that need must be resolved immediately.

A virtual memory shortage serious enough to degrade or delay communications and other internal system functions will very likely cause SIOS Protection Suite to malfunction. For instance, `deepcheck` of TCP/IP communication resources may be impacted enough to cause a false failure, and thus a failover of the resource to the backup server.

If SIOS Protection Suite communication with other servers in the cluster is degraded, it could cause a manually initiated switchover to fail. However, this will not affect SIOS Protection Suite's ability to fail over protected resources when a server completely fails.

GUI interoperability

The LifeKeeper GUI may only be used to administer SIOS Protection Suite on Windows servers. Note that you can *connect* to and *monitor* a SIOS Protection Suite for Linux cluster. However, performing administrative tasks such as creating resources, editing properties, bringing servers in and out of service, is **not** supported at this time.

Discontinuing Serial Port Communication Paths

SIOS Protection Suite discontinued support for TTY communication paths in Version 7.2. Though SIOS does not recommend it, if currently using TTY communication paths, this option can be re-enabled by removing the (#) symbol on the TTYCA.EXE line in the `/etc/lkinit.config` file as shown below:

```
# ... /bin/TTYCA.EXE|-t 1 X X X X X X <=  
(TTY Comm Paths Disabled)  
... /bin/TTYCA.EXE|-t 1 X X X X X X <=  
(TTY Comm Paths Enabled)
```

To enable or disable the TTY communication path feature, the LifeKeeper service must be stopped and restarted after editing *lkinit.config*. To stop LifeKeeper, run command {c:\lk}\bin\lkstop.exe -f (c:\lk being the LifeKeeper installation path). Make sure the GUI is closed and all processes associated have stopped. Restart LifeKeeper by entering {c:\lk}\bin\lkstart.exe.

The TTY technology is obsolete. TTY communication paths are not supported and should be replaced with TCP/IP communication paths.

Console Application Management

Launching console applications from SIOS Protection Suite is not supported on Windows Server 2008 and later. Server architecture and security improvements in Server 2008 including UAC and memory management, prevent background processes such as SIOS Protection Suite from starting console applications.

Bitlocker Does Not Support DataKeeper

According to Microsoft, Bitlocker is not supported to work with Software RAID configurations. Since DataKeeper is essentially a software RAID 1, Microsoft does not support Bitlocker working with DataKeeper.

The specific article and section can be found here:

http://technet.microsoft.com/en-us/library/ee449438#BKMK_R2disks

Known Issues

Microsoft SQL Server 2017 Optional Services

- SQL Server Launchpad
- SQL Server PolyBase Data Movement
- SQL Server PolyBase Engine

- SQL Server CEIP service



Note: SIOS Protection Suite for Windows v8.6.2 does not protect these Optional Services.

For additional known issues, see the Troubleshooting section of [SIOS Protection Suite for Windows Technical Documentation](#).

Frequently Asked Questions

Can I change my SIOS Protection Suite configuration database setting including resource values without reinstalling SIOS Protection Suite or rebuilding my resources?

Yes. Use the `lk_chg_value.ksh` command.

Can I upgrade my existing SIOS Protection Suite hierarchies from a previous version of SIOS Protection Suite for Windows to v8?

You may upgrade your existing SIOS Protection Suite for Windows software while preserving your resource hierarchies. Please refer to the [Upgrading SIOS Protection Suite](#) topic for the correct upgrade procedure.

Note: SIOS does not support upgrading SIOS Protection Suite from more than one major version back. If upgrading from a version prior to LifeKeeper for Windows v7.x to SIOS Protection Suite for Windows v8, uninstall the old version of LifeKeeper and reinstall SIOS Protection Suite for Windows v8.

Does SIOS Protection Suite operate in a cluster with Microsoft Cluster Services (Windows 2003) or Windows Server Failover Cluster (Windows 2008 and later)?

No. SIOS Protection Suite v8.0.1 is an alternative clustering product and does not support either Microsoft Cluster Service or Windows Server Failover Clustering.

Does SIOS Protection Suite require that all servers in the cluster be identically configured?

No. As long as all servers are powerful enough to run any application that may run on them as the result of a failover operation and meet all other SIOS Protection Suite requirements, a cluster can be built. SIOS Protection Suite does not require identical hardware, but the software should be the same and configured with the same service pack levels.

Does SIOS Protection Suite for Windows support 64-bit environments?

Yes. SIOS Protection Suite for Windows supports only 64-bit platforms.

How do I change permissions on SIOS Protection Suite protected File Share resources?

The `EditFileShareResource` utility can be used to update a file share resource with all current file shares and permissions on the associated volume(s). This can be useful in environments where there are a large number of file shares, and file shares have been added or deleted since the resource was created or permissions have been modified. Using the utility can prevent the need to delete and re-create the file share resource. The `EditFileShareResource` utility is located under `%LKROOT%\bin` directory.

To invoke the utility, on the command line enter:

```
EditFileShareResource <Tag name>
```

where <Tag name> is the tag name of a file share resource that is currently in service.

The utility protects **all** eligible file shares defined on the volumes that are associated with the file share hierarchy. It deletes any previously protected shares that have been deleted from the system and adds newly defined shares (meeting the eligibility criteria) to the list. It will also update the file share permissions defined on the file share.

Documentation

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

Section	Description
Introduction	Provides an introduction to the SIOS Protection Suite for Windows product, including an overview of its components.
Installation	Provides useful information for planning and setting up your SIOS Protection Suite environment, installing and licensing SIOS Protection Suite and configuring the LifeKeeper GUI to run on a remote system.
Configuration	Contains detailed information and instructions for configuring the SIOS Protection Suite software on each server in your cluster.
Administration	Discusses server-level tasks such as editing server properties, creating resources and creating or deleting comm paths and resource-level tasks such as editing, extending or deleting resources.
Man Pages	Provides reference manual pages for the SIOS Protection Suite product.

User's Guide	Contains detailed information on the LifeKeeper GUI, including the many tasks that can be performed within the LifeKeeper GUI. Also includes information on Data Replication along with many more Advanced Topics.
DataKeeper	Provides an overview of how DataKeeper replication works and contains complete information on configuring and administering DataKeeper. Topics include network considerations, common configuration issues and requirements necessary to successfully install and configure DataKeeper.
Troubleshooting	Describes known issues and suggests solutions to problems that may be encountered during installation, configuration or use of SIOS Protection Suite for Windows.
Recovery Kits	Contains planning and installation instructions as well as administration, configuration and user information for the Optional Recovery Kits (SQL Server and Oracle) that allow LifeKeeper to manage and control specific applications.

Quick Start Guides

To get started using SIOS Protection Suite for Windows, refer to the [SIOS Protection Suite for Windows Quick Start Guide](#) and the [DataKeeper Quick Start Guide](#).

Training

SIOS Protection Suite training is available through SIOS Technology Corp. or through your SIOS Protection Suite 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 [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 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

SIOS Protection Suite Installation Guide

The topics in this Installation Guide will assist in defining your SIOS Protection Suite cluster environment. Once your requirements have been determined and your SIOS Protection Suite configuration has been mapped, these topics will assist you in setting up, licensing and installing SIOS Protection Suite.

[Planning Your SIOS Protection Suite Environment](#)

[Setting Up Your SIOS Protection Suite Environment](#)

[Installing SIOS Protection Suite](#)

Planning Your SIOS Protection Suite Environment

This section will assist you in defining your SIOS Protection Suite cluster environment enabling you to successfully achieve your high availability goals quickly and effectively.

The major subjects covered are listed below.

[Planning Server Communication](#)

[Recovery Kit Requirements](#)

[Storage and Adapter Requirements](#)

[Verifying Server Specifications](#)

Planning Server Communication

Determine and document server communication in a configuration map similar to the one below, using the following guidelines:

Cluster requirements – To avoid a single point of failure, SIOS Protection Suite requires at least two communication paths (also called “comm paths” or “heartbeats”) between servers in a cluster. See [Communication Path Considerations](#) below for more details.




Figure 1 – Sample Configuration Map for SIOS Protection Suite Pair

This is a very simple configuration map depicting a pair of SIOS Protection Suite servers sharing a disk array subsystem. Under normal circumstances, Server 1 runs the application(s) and is considered the primary or active server. Server2 is the secondary or standby server. In this case, there is no contention for

disk resources because only one server at a time reserves an entire volume of the disk array.

This sample cluster also shows TCP/IP communication paths configured on the public network and on the private network. On your configuration map, label the IP addresses associated with each TCP/IP comm path.

A pair of servers is the simplest SIOS Protection Suite configuration. When planning a cluster consisting of more than two servers, a configuration map is even more critical to ensure that the appropriate connections exist between and among servers. Each server must have a physical communication path to every other server in the cluster in order to provide cascading failover capability.

Note: If using replicated storage rather than shared storage, refer to [SIOS DataKeeper](#) for additional information on configuring hardware and software for replication.

Communication Path Considerations

SIOS Protection Suite comm paths are used to communicate the state of protected resources in a cluster and to manage failovers. Each comm path is assigned a priority number with the lowest number designated as the “highest” priority.

The recommended configuration is two separate LAN-based (TCP/IP) comm paths configured on independent subnets. The primary comm path should be configured on the private network. **A switchable IP address should not be configured on the Network Interface Card (NIC) carrying the primary comm path.**

Redundant Comm Paths

SIOS Protection Suite strongly recommends redundant comm paths whenever possible. If a single comm path is used and that comm path fails, then SIOS Protection Suite hierarchies may 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 can access and write data to the shared storage device.

Primary Comm Path (Private Network)

A private TCP/IP comm path provides reliable communication between systems that is not affected by any communication occurring on the public network. For this reason, it is recommended that the primary comm path be configured on a private network and the secondary comm path on the public network. Private network addresses must not be registered with DNS. The **“Register this connection’s address with DNS”** checkbox must not be checked for private network addresses.

TCP/IP comm paths are configured in SIOS Protection Suite using static IP addresses and subnet masks. The cabling may consist of either a crossover cable for a two-node cluster or a small hub for clusters of three or more nodes.

Note: It is very important that private network connections are not registered with DNS. DNS should normally publish only the public network connection for each server. This is essential when connecting a local LifeKeeper GUI admin client to a remote SIOS Protection Suite system. Refer to [Verifying Network Configuration](#) for network configuration details.

Recovery Kit Requirements

Each of the SIOS Protection Suite Recovery Kits has requirements that you should consider in planning and connecting all the components of your SIOS Protection Suite cluster. While the [SIOS Protection Suite for Windows Release Notes](#) provides technical requirements for each kit such as program versions and disk space requirements, you will find detailed configuration information in the [Recovery Kit](#) section.

The Core recovery kits (Volume, IP, LAN Manager, File Share, DNS, Microsoft Internet Information Services (IIS) and Generic Application) are documented throughout the [SIOS Protection Suite for Windows Technical Documentation](#).

Note: All separately packaged (optional) SIOS Protection Suite Recovery Kits require a software license key in order to function with LifeKeeper v4.3 and higher. You can install the license key by running the License Key utility from Start->All Programs->SIOS->LifeKeeper->License Key Installer.

Storage and Adapter Requirements

SIOS Protection Suite configurations may use the facilities of shared 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.

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). SIOS Protection Suite supports a number of hardware RAID peripherals for use in SIOS Protection Suite configurations. The primary requirement is that the device is supported by Microsoft. See the [Windows Server Catalog](#).

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

- SIOS Protection Suite manages resources at the volume level, making the resources on each volume 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 SIOS Protection Suite.

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 Microsoft so that there is a driver available. See the "Cluster Solutions" categories in the [Windows Server Catalog](#) for Microsoft-supported adapters and peripherals.

For reference purposes, you should add the host adapter specifications to your configuration map.

Verifying Server Specifications

Ensure that you have the correct version and/or capacity of the following components for each SIOS Protection Suite server:

- Windows 2008 R2, 2012, 2012 R2 and 2016 Operating System (64-bit only).

Note: All servers within a cluster should be running the same version of Windows.

Note: File and Printer Sharing is enabled for Lan Manager AND for use with DataKeeper replicated volumes. During the installation procedure, SIOS Protection Suite can automatically configure the Windows 2008 firewall so that ports it needs are opened, and so that ICMP is enabled.

Note: The Local Security Policy “**Network Access: Let Everyone permissions apply to anonymous users**” must be **Enabled** if you plan to use DataKeeper replicated volumes with SIOS Protection Suite. This policy will be enabled by LifeKeeper installation.

Note: By default, firewall is enabled. During installation, if a firewall is detected, the appropriate rules will be added to windows firewall. However, if the firewall is disabled during installation and re-enabled at a future time, the setup firewall script needs to run to add the rules. This script is installed as `LKROOT%\support\firewallSetup.bat`. To run the command from the command line, type `firewallSetup.bat %LKROOT`

- Ethernet TCP/IP-supported network interface card(s) for LAN-based cluster heartbeat(s).
- Disk arrays and storage adapters (SCSI or Fibre Channel) if you are using shared storage.
- Memory. See the [SIOS Protection Suite for Windows Release Notes](#) for minimum memory requirements for SIOS Protection Suite.

Note: Additional memory (beyond that required for SIOS Protection Suite) is required to run user applications.

- Disk space. See the [SIOS Protection Suite for Windows Release Notes](#) for minimum disk space requirements for SIOS Protection Suite and recovery kits.
- LifeKeeper Graphical User Interface (GUI) platforms and browsers.

- **Power Requirements.** To maximize the availability of your SIOS Protection Suite servers, it is strongly recommended that you use Uninterruptible Power Supplies (UPSs), or at a minimum, separate the electrical sources to your servers.
- Application software to be protected by SIOS Protection Suite.

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.

Setting Up Your SIOS Protection Suite Environment

Now that you have determined your requirements and mapped your SIOS Protection Suite configuration, you can start setting up the components of your SIOS Protection Suite environment.

The major tasks of this topic are:

[Configuring Your Storage](#)

[DNS Resource Requirements](#)

[Installing and Setting Up Database Applications](#)

[Safe Creation of Shared Disk Volume Instances](#)

[Verifying Network Configuration](#)

Note: Although it is possible to perform some setup tasks in a different sequence, this list provides the recommended sequence.

Configuring Your Storage

SIOS Protection Suite may be used with either shared storage or with replicated storage. Follow the instructions that apply to your configuration below.

Shared Storage Configuration

If you are using shared storage, then after your Windows environment is installed, you should set the host adapter and shared peripheral addressing.

Note: Dynamic disks are not supported with Shared Storage because the dynamic disk configuration is stored somewhere (undocumented) on each system, not on the disks themselves. There is currently no way to replicate that configuration between the two systems.

Refer to the documentation accompanying your adapter and storage device for specific details. Perform the following tasks to configure your shared storage for access by all servers in the SIOS Protection Suite cluster:

1. Because all disks placed under SIOS Protection Suite protection must be partitioned, your shared disk arrays must now be configured into partitions (volumes) using the Windows Disk Management utility. You should also format the partitions with the NTFS file system.

Note: To safely configure your shared storage, it is recommended that you follow the procedure in [Safe Creation of Shared Disk Volume Instances](#).

You should refer to your disk array software documentation for detailed instructions.

2. If you plan to use a Shared Disk comm path, designate a small raw (unformatted) partition to use for the comm path. One MB should be a sufficient size.
3. Power on the other server(s) in the cluster and verify that all servers recognize the shared disks. From the backup server(s), make drive assignments for the shared volumes exactly the same as the first server. It is recommended that you have the Disk Management utility open on only one server at a time.
4. If you have created file shares on the shared volumes, you will need to turn on the file sharing attribute of these folders on each server in the cluster.

Replicated Volume Configuration

If you are using SIOS DataKeeper for Windows, create your disk partitions (volumes) to be replicated using the Windows Disk Management utility. You should also format the partitions with the NTFS file system.

Be sure to assign the same drive letter to the source volume (on the primary server) and target volume (on the backup server).

DNS Resource Requirements

The DNS Recovery Kit included with the SIOS Protection Suite for Windows Core product provides a mechanism to update DNS A and PTR records of the primary server or a LifeKeeper alias name on the DNS servers in your configuration. The DNS resource allows the user to select the A record of the primary server or a LifeKeeper alias name in DNS which will be modified along with the PTR record (if exists) with the IP address of a backup server when failover or switchover occurs. Using a DNS resource allows clients to connect to the servers in a WAN environment when a failover or switchover occurs. When SIOS Protection Suite servers are in different network subnets, it is not possible to use a switchable IP address. In this type configuration, a DNS resource should be used to provide client connectivity. For details on creating DNS resources, refer to [Creating a DNS Resource Hierarchy](#).

Restriction: SIOS Protection Suite servers should not be configured as Domain Controllers or DNS Servers. Creating a DNS resource that points to a DNS server on the same system will fail with the following error message: "User credentials cannot be used for local connections."

TTL of DNS Records

When the SIOS Protection Suite for Windows DNS Recovery Kit updates the A record of the primary server or LifeKeeper alias name in DNS, the A record on the caching DNS servers' cache is not updated. These caching DNS servers are those who do not hold the zone that the SIOS Protection Suite protected A record belongs to. The A record in the cache stays until the TTL is expired or the cache is cleared manually. Therefore, the clients of those caching DNS servers will not get the updated value of the A record in timely fashion. For SIOS Protection Suite protected DNS resources, it is recommended that the TTL value of the A record of the primary server or LifeKeeper alias name should be set to a lower value.

If SIOS Protection Suite creates the A and PTR records for a DNS resource, the TTL of those records is set to 5 minutes. This value can be changed using the Microsoft DNS management console (dnsmgmt.msc). However, changing the value to a higher value will make the A record live in the cache longer on caching DNS servers.

For DNS A and PTR records created prior to creating the SIOS Protection Suite DNS resource hierarchy, it is recommended that the TTL value be set to a lower value like 5 minutes.

Installing and Setting Up Database Applications

If your environment includes a protected database application such as [SQL Server](#), you should install the application using the documentation provided with the database. Ensure that the database is on a shared or replicated file system and that the configuration files are on a shared or replicated file system. The executables may either be on each local or a shared file system. Refer to [SIOS Protection Suite Microsoft SQL Server Recovery Kit Technical Documentation](#) for additional installation and setup considerations regarding SQL Server.

Although it is possible to install your application after SIOS Protection Suite is installed, you should test the application to ensure it is configured and operating properly before placing it under SIOS Protection Suite protection.

Safe Creation of Shared Disk Volume Instances

In order to safely create a shared-storage volume resource, the user must ensure that only one system at a time has write access to the volume at any time. This includes the time prior to the creation of the SIOS Protection Suite instance.

Since SIOS Protection Suite cannot recognize that the volume is shared before an instance is created, manual steps must be taken to ensure that the volume is never writable on two or more systems at the same time.

To protect the volume from simultaneous write access, use the following procedure. In this example, two systems – SYSA and SYSB – are connected to shared storage. This storage is configured with two volumes which should be assigned drive letters E and F on both systems, then protected with SIOS Protection Suite volume instances.

1. Power on SYSA, while leaving SYSB powered off.
2. Install LifeKeeper if it has not been installed.
3. Assign drive letters E and F to the volumes; format with NTFS if not formatted yet.
4. Power off SYSA.
5. Power on SYSB.
6. Install LifeKeeper if it has not been installed.
7. Assign drive letters E: and F: to the shared volumes.
8. In a command prompt, run the following commands:

```
%LKBIN%\volume -p E
```

```
%LKBIN%\volume -p F
```

9. Reboot SYSB. It will come up with the E: and F: drives locked.

10. Power on SYSA. It will come up with the E: and F: drives writable.

11. Create volume resources for E: and F: on SYSA and extend to SYSB.

An alternative to powering the systems off is to use Disk Management to take the shared physical disk offline.

Verifying Network Configuration

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

1. You must ensure that every network interface card (NIC) has one permanent IP address in order to create a TCP/IP comm path or protect an IP address.
2. If your server has more than one NIC (recommended), you should configure them to be on different subnets. If the adapters are on the same subnet, TCP/IP cannot effectively utilize the second adapter.
3. Your IP addresses should be configured as follows, assuming at least two NICs in each server (one on a private network and one on the public network):
 - a. In the **Control Panel**, click on **Network Connections**. Right-click **Open**.
 - b. From the **Advanced** menu, select **Advanced Settings**.
 - c. Ensure that the NIC connected to the public network is in the topmost position of the **Connections** list.
 - d. Do not register private network connections with DNS. Uncheck the “**Register this connection’s address with DNS**” checkbox for the private network adapter as follows:

Internet Protocol (TCP/IP) Properties-> Advanced -> DNS Tab

Since no DNS servers are needed for the private network connection, none should be listed.

This prevents the browser from occasionally getting confused when switching over LAN Manager computer names.

4. 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.
5. To ensure that the LifeKeeper GUI server and client components can effectively communicate ensure that localhost is resolvable by each server in the cluster.

- If DNS is not implemented, edit the **%windir%\system32\etc\drivers\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.

- If DNS is implemented, verify the configuration to ensure the servers in your SIOS Protection Suite cluster can be resolved using DNS.

6. Ensure each server's hostname and networking addressing information is correct and will not change after SIOS Protection Suite is installed. If changing the hostname after SIOS Protection Suite is in operation, you must run the `lk_chg_value` utility to modify the computer name in the SIOS Protection Suite configuration files. If changing the networking configuration after SIOS Protection Suite is in operation, you must run the `lk_chg_value` utility to modify existing SIOS Protection Suite comm paths and resource hierarchies after re-configuring your network information.

Note: If you are using SIOS DataKeeper for Windows, refer to the [SIOS DataKeeper](#) section of the documentation for additional information on specifying the network cards to be used for replication and comm path considerations.

Switchable IP Address

Most SIOS Protection Suite configurations use the IP Recovery Kit, which defines 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 SIOS Protection Suite protection are associated with the switchable IP address. Then, if there is a failure on the primary server, the switchable 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 on the network.

Note: If using teaming software or if network cards are changed after creating a switchable IP resource, the switchable IP resource should be deleted and recreated as the associated index number for the card can change.

Switchable IP Addresses, DNS and LifeKeeper GUI

Considerations

Special network considerations must be made when a “virtual” IP address is used on the server’s main NIC and DNS registration is enabled on the NIC. When a “virtual” IP address is created by SIOS Protection Suite on a registered NIC, DNS will add this additional IP address for the server and start using it for host name resolution on the network. However, SIOS Protection Suite protected “virtual” IP addresses are switchable across cluster nodes. Therefore, precautions must be taken to prevent the LifeKeeper GUI from also using DNS registered “virtual” IP addresses to get updates from local and remote cluster nodes.

To keep LifeKeeper GUI connections to local and remote systems stable when using “virtual” IP addresses, there are two options:

1. Use a network hosts file on each SIOS Protection Suite node.
 - In the *hosts file*, identify the permanent IP address for every other remote cluster node.
 - Do this on every SIOS Protection Suite system in the cluster.

As explained above, these addresses must be on the highest priority network used for LifeKeeper GUI binding.

2. Use an alternate network and associated alternate NIC for LifeKeeper GUI connections to all other nodes in the cluster. This option differs from the simpler recommendations explained above.
 - Enable DNS registration on the alternate network and NIC.
 - Make the alternate network the highest priority in the *Network Connections -> Advanced -> Advanced Settings* selection in the **Adapters and Bindings** tab. The LifeKeeper GUI will use this highest binding network.
 - The highest priority SIOS Protection Suite comm path should also use this network.
 - Do this on every SIOS Protection Suite system in the cluster.

The LifeKeeper GUI will use this alternate network for connections to all cluster nodes. With no virtual IPs assigned to this alternate network, the address registration will be stable. DNS registration may also be used for the main/public NIC on the server as needed.

Note: After making network configuration changes, the “`ipconfig /flushdns`” command may be used to remove obsolete cached DNS information.

IP Local Recovery Configuration

SIOS Protection Suite provides the ability to monitor local switchable IP addresses and move them to another network adapter in the same system when a failure is detected. This functionality, called IP Local Recovery, imposes additional requirements and limitations on the system configuration.

The backup adapter, also known as the Local Recovery Adapter where the switchable address will become active after a failure of the primary adapter, must be configured as follows:

- Both adapters must be connected to the same physical subnet.
- For routing purposes, all addresses on the Local Recovery Adapter must be on a different logical subnet than any permanent addresses on the Primary adapter. They must also be on a different logical subnet than any SIOS Protection Suite-protected switchable addresses that are configured on the Primary adapter.
- IP Local Recovery can only be enabled at the time the IP resource is created. Local Recovery cannot be added to an IP resource by modifying its resource attributes after the resource has been created.
- IP Local Recovery may be disabled for an IP resource by using the `ins_setlocalrecovery` command line utility. This utility is located in the SIOS Protection Suite **bin** directory (`C:\LK\bin` by default). From a command prompt, type `ins_setlocalrecovery` for the usage and switch options.

How IP Local Recovery Works

When IP Local Recovery is enabled and the IP resource fails its quick check or deep check tests, then SIOS Protection Suite will do the following:

- First, SIOS Protection Suite will attempt to bring the IP address back in service on the current network interface.
- If that fails, SIOS Protection Suite will check the resource instance to determine if there is a backup (Local Recovery Adapter) available. If so, it will then attempt to move the IP address to the backup interface.
- If all local recovery attempts fail, SIOS Protection Suite will perform a failover of the IP address and all dependent resources to a backup server.

Even if you do not have a backup adapter, you can enable Local Recovery so that SIOS Protection Suite

will retry the primary network interface before initiating failover to a backup server.

Installing SIOS Protection Suite

If you have completed planning and setting up your SIOS Protection Suite environment, you should be ready to install the SIOS Protection Suite software on each server in your cluster.

[Core Software](#)

[Installing Core](#)

[Licensing](#)

[Installing Localized Language Supplement](#)

[Silent Installation of SIOS Protection Suite](#)

[Third Party Product Files](#)

[Application Directory Anomaly](#)

[Uninstalling SIOS Protection Suite for Windows](#)

[Upgrading SIOS Protection Suite](#)

[Repair](#)

[Starting LifeKeeper](#)

Core Software

SIOS Protection Suite Core Software

The SIOS Protection Suite Core software is available via ftp download. The SIOS Protection Suite Core is comprised of:

- The basic LifeKeeper software, including:
 - Perl (CPAN v5.8.8)
 - Cygwin
 - International version of Java Runtime Environment (JRE) v1.8.0 Update 101
 - LifeKeeper GUI (both server and client)
 - Microsoft Visual C++ 2008 Redistributable package (v 8.0.56336)
- Core recovery kits:
 - Volume
 - IP
 - DNS
 - LAN Manager
 - File Share
 - Generic Application
 - Internet Information Services (IIS)
 - PostgreSQL
- DataKeeper
 - DataKeeper Driver (ExtMirr.sys)
 - DataKeeper Service (ExtMirrSvc.exe)
 - Command Line Interface (EMCMD.exe)
 - DataKeeper GUI (Datakeeper.msc)
 - Packaging files, SIOS Protection Suite scripts, help files, etc.

Installing Core

Installing the SIOS Protection Suite Core Software

SIOS Protection Suite uses the Flexera InstallShield product to provide a standard installation interface. A license must be obtained and installed for each server in the cluster.

We recommend that you read the [SIOS Protection Suite for Windows Release Notes](#) before installing and configuring SIOS Protection Suite.

To install SIOS Protection Suite software, run the setup program delivered with the SIOS Protection Suite for Windows product. The InstallShield Wizard will first install LifeKeeper for Windows. Once the LifeKeeper installation is complete, SIOS DataKeeper for Windows will be installed. Follow the setup instructions on each screen. Some explanatory notes are included below.

LifeKeeper Installation Notes

- You must have administrative privileges to install the LifeKeeper software. While non-administrative users will not be prohibited from running the setup program, the installation will exit immediately due to lack of special permissions required during setup.
- Installing LifeKeeper on your shared storage is **not** supported. Each server should have its own copy installed on its local disk.
- The default LifeKeeper installation path is `C:\LK`. You may change this path, but due to some scripting issues, **be sure to choose a path with NO EMBEDDED SPACES and containing eight characters or less**. For instance, `C:\Program Files\LK` and `C:\LifeKeeper` are invalid choices that will result in application errors.
- Two Windows registry changes are made during the installation of LifeKeeper:
`DisableStrictNameChecking` and `DisableLoopbackCheck`. Both of these changes are required to allow access to servers using an alias name.

Setup Type

Choose one of the following:

- **Typical** installs the LifeKeeper Core and all Core recovery kits (recommended). **Note:** DHCP Media

Sense for TCP/IP will be disabled by default.

- **Compact** installs the LifeKeeper Core only (which includes the Volume Recovery Kit).
- **Custom** allows you to select from the list of LifeKeeper components to install: Core files (always required), IP Recovery Kit, DNS Recovery Kit, LAN Manager Recovery Kit, File Share Recovery Kit, Generic Application Recovery Kit and IIS Recovery Kit. The Custom option will ask the following questions:
 - “Disable DHCP Media Sense for TCP/IP?”
 - “Do you wish to start the LifeKeeper Services?” See [Starting LifeKeeper Services](#) below for details.

Firewall Change Prompt (Windows 2008 and later)

LifeKeeper cannot function properly if the firewall settings for the source and target machines are not configured correctly. During installation of LifeKeeper, you will be prompted to allow the installer to configure your firewall rules needed by LifeKeeper, as well as to configure other system settings that are required by LifeKeeper. If you choose to allow the installer to make these changes, you will not need to configure your firewall manually. Please refer to [Troubleshooting](#) for more information.

LifeKeeper requires the following ports / protocols / processes to be open or enabled:

TCP Ports: 81, 82, 1500, 3278, 3279

Process: %LKROOT%\jre1.8.0_101\bin\java.exe

Process: %LKROOT%\jre1.8.0_101\bin\jp2launcher.exe

Protocols: ICMP Echo

Starting LifeKeeper Services

If you choose the **Custom** installation option, you will be asked, “Do you wish to start the LifeKeeper Services?” In most cases you should answer **Yes** so that LifeKeeper will be started automatically when the system is booted. Answering **No** will cause LifeKeeper not to be started after installation, and it will set the **Startup Type** for the LifeKeeper services to **Manual**.

If you select **No** and you later wish to start the LifeKeeper services, you should do so using the **Services** tool in the **Windows Control Panel**. (You should start both LifeKeeper and LifeKeeper External Interfaces.) In addition, you can set the **Startup Type** to **Automatic** by right-clicking on each service and selecting

Properties, then changing the **Startup Type** option to **Automatic**. This will tell LifeKeeper to always start at system boot time.

Question: In what situation would it make sense to answer **No** to starting the LifeKeeper services?

Answer: Choosing not to start the LifeKeeper services may be useful in a staging environment where you are not ready to configure your network addresses but you wish to install LifeKeeper and replicate it across a number of systems prior to final installation of the cluster.

Explanation: When LifeKeeper is started the FIRST time, the system's network configuration information is written into the **LifeKeeper Configuration Database** (LCD). Changing your network configuration AFTER LifeKeeper is started requires deleting and re-creating your comm paths and resource hierarchies. Therefore, by choosing NOT to start the LifeKeeper services at install time, you can install LifeKeeper and associated recovery kits, then configure your network later.

DataKeeper Installation Notes

Once LifeKeeper installation is complete, the InstallShield Wizard will begin installing SIOS DataKeeper for Windows. You will be prompted to select the DataKeeper features to install. A typical installation includes both features.

- [DataKeeper Server Components](#)
- [DataKeeper User Interface](#)

During installation of DataKeeper Server Components:

1. [Configure firewall settings](#).
2. Select a [DataKeeper Service log on](#) account type.
 - If **Domain or Server account** is selected, provide DataKeeper Service log on ID and Password .
3. [Install licensing](#) via the **License Manager**.

Reboot your server and begin using DataKeeper. See the [DataKeeper Technical Documentation](#) for further information on using DataKeeper.

The **SIOS DataKeeper User Interface and Server Components Feature** can be installed independently,

and the installation can be modified later to include any feature that has not previously been installed.

Important: The SIOS DataKeeper User Interface feature and the target snapshot feature require Microsoft MMC 3.0 and Microsoft .NET Framework 3.5 SP1 to be installed. For Windows 2008 R2 and 2012, use “Server Manager” to enable the .NET Framework 3.5.1 features. If the SIOS Protection Suite install is attempted prior to installing these proper versions, an error will be received and the installer will be stopped. SIOS Protection Suite will need to be uninstalled and the SIOS Protection Suite install process will need to be restarted.

Exclusion list for Antivirus Software for LifeKeeper and DataKeeper for Windows

The following things should be excluded in your antivirus software for LifeKeeper and DataKeeper:

- For SPS C:\LK* directory (or the folder LifeKeeper is installed in).
- For DataKeeper C:\Program Files (x86)\SIOS\DataKeeper\ directory (or the folder DataKeeper is installed in).
- The bitmap file location (by default on the c: drive but it may be relocated – C:\Program Files (x86)\SIOS\DataKeeper\Bitmaps).

These locations have all of the executables and sometimes the Antivirus Software can quarantine them, thus rendering LifeKeeper or DataKeeper inoperable.

The list of registry keys that LifeKeeper and DataKeeper use is located [here](#).

AND

UpperFilters registry key located at:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\
Class\{71A27CDD-812A-11D0-BEC7-08002BE2092F}
```

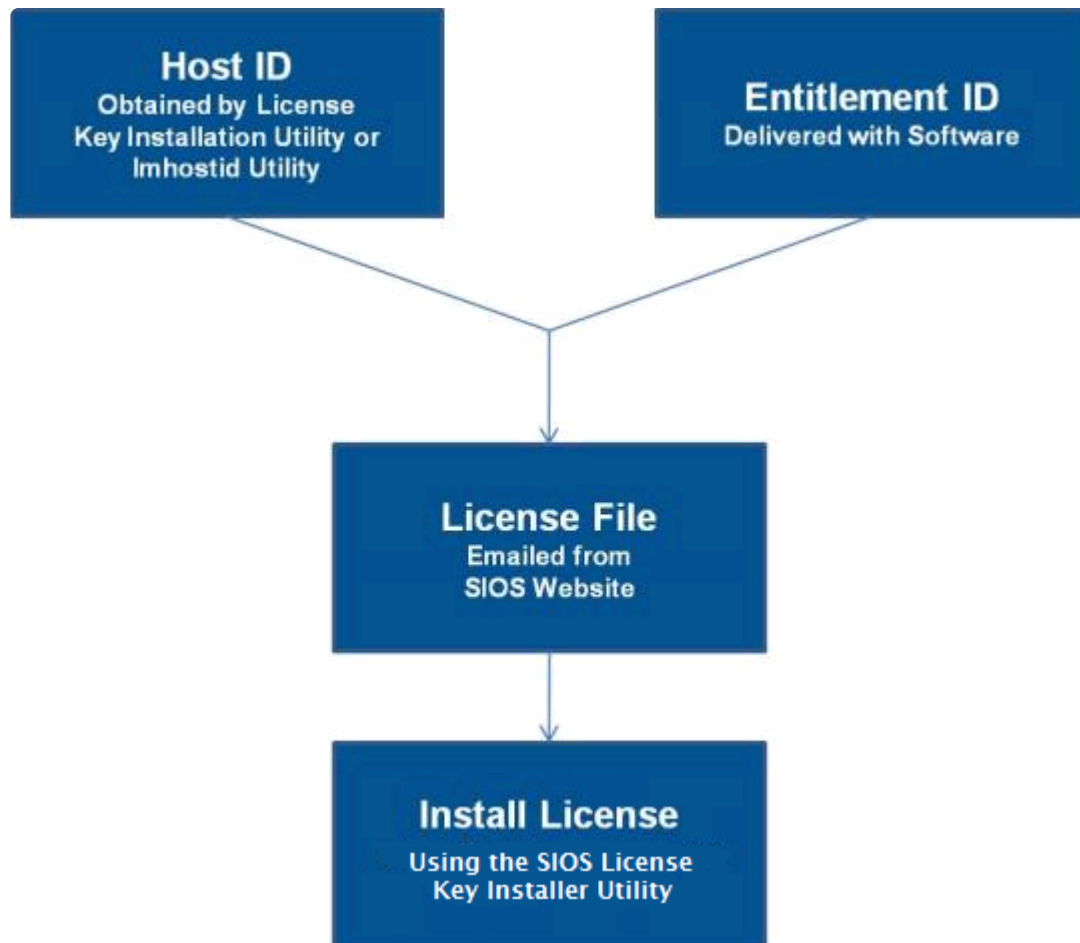
The contents of the UpperFilters key should be “NCR_LKF ExtMirr” when using both LifeKeeper and DataKeeper.

Licensing

Obtaining and Installing the License

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

The final screen of the **InstallShield installation utility** displays the Host ID of your server. The **Host ID**, along with the **Entitlement ID** (Authorization Code) that was provided with your SIOS Protection Suite for Windows software is used to obtain the license. The process is illustrated below.



License Key Manager

In addition to installing SPS for Windows product licenses, the **License Key Manager** allows you to perform

the following functions:

- View all licenses currently installed on your system.
- View all expiration notifications (days remaining) for each time-expiring license.
- Identify invalid licenses that are currently installed.
- Delete any installed licenses (right-click on the license and select **Delete**).
- Delete all expired licenses as a group (press the **Delete Expired License** button).
- **Refresh** the Installed License list when installing software or upgrades.

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

1. Get your **Host ID**. At the end of the SPS for Windows installation, make note of the **Host ID** displayed by the **License Key Installer** utility as shown below. The Host ID may also be obtained by running `%ExtMirrBase%\bin\lmhostid` (where `%ExtMirrBase%` is the SPS for Windows installation path, by default `C:\Program Files (x86)\SIOS\LifeKeeper`) on the system(s) that you are obtaining licenses for. (If you need to obtain your Host ID again at a later time, you may do so by running the **License Key Installer** utility from the **Start-Programs** menu **Start-All Programs-SIOS-LifeKeeper-License Key Installer**.)
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 SPS for Windows **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**.

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 key(s) to the C:\Windows\SysWOW64\LKLicense folder.

OR

- Run the **License Key Installer** from the **Start-Programs** menu (**Start-All Programs-SIOS-LifeKeeper-License Key Installer**).
- Press the **Install License File ...** button on the main screen of the **License Key Installer**.
- Browse to the location of the license file that you saved in **Step 4** above.
- Click on the license file name. It will become highlighted.
- Press the **Install License File ...** button that appears in that dialog box below the file names. A license detection confirmation popup will be displayed.

6. Repeat on all additional servers. You must install a license on the other SIOS Protection Suite for Windows server(s) using the unique Host ID for each server.

7. Restart SIOS Protection Suite for Windows.

Primary Network Interface Change May Require a License Rehost

The Host ID used by the License Key Installer utility is obtained from the SPS for Windows server's primary network interface card (NIC). SPS for Windows will check for a valid license each time it starts. If your SPS for Windows server should require a NIC replacement in the future that would cause the Host ID to change, then the next time SPS for Windows is stopped, a License Rehost must be performed before starting either again. Log in to the [SIOS Technology Corp. Licensing Operations Portal](#) and select **Support Actions/Rehost** from the **Manage Licenses** screen to perform this rehost. (**Note:** A rehost can be performed one time per six-month period without contacting support.)

Troubleshooting

If errors are encountered, please try the following before contacting Support:

- Review the error messages in the **Windows Event Viewer**.
- Verify credentials by logging in to the [SIOS Technology Corp. Licensing Operations Portal](#). Enter **User ID** and **Password**. Run %ExtMirrBase%\lmSubscribe.exe again using the correct **User ID** and **Password**.
- To force a manual check for a license renewal, stop and restart the service. (**Note:** To find the service, bring up the view for all of the Windows services and search for “**SIOS Subscription Licensing**”.)
- If ownership of the license certificate has changed, please [contact SIOS Technology Corp. Support](#) personnel to have the certificate moved to the new owner. Once ownership has been moved, the automatic license renewal service will need to be updated with these new credentials by running the above command again using the new **User ID** and **Password**.

Installing Localized Language Supplement

LifeKeeper for Windows Localized Language Supplements are available to support LifeKeeper running in a localized environment. The Localized Language Supplement contains translated LifeKeeper GUI text strings and context-sensitive help in the localized language. For LifeKeeper v7.2.1 and beyond, language supplements are available for Japanese language. The SIOS Protection Suite for Windows Core installation program installs the required version of JRE.

The Japanese Localized Language Supplement includes language content for the SIOS Protection Suite including LifeKeeper and DataKeeper products. The administrator can select which product to update. The DataKeeper mmc-based GUI requires the Windows language pack be installed unless the complete localized OS is already installed.

The LifeKeeper for Windows Localized Language Supplement, like the SIOS Protection Suite Core, is installed via InstallShield. The installation requires no selection for Typical/Compact/Custom options. To install the LifeKeeper for Windows Localized Language Supplement, run the setup program shipped with the Localized Language Supplement product.

To repair an existing installation of LifeKeeper for Windows Localized Language Supplement, run the setup program and choose **Repair** from the list of InstallShield options.

To remove LifeKeeper for Windows Localized Language Supplement, run **Add/Remove Programs** from the **Windows Control Panel**. The Localized Language Supplement must be removed before removing the LifeKeeper core product.

Silent Installation of SIOS Protection Suite

* Note: To perform the silent installation for SIOS Protection Suite for Windows, you must contact [Support](#) to get separate LifeKeeper and DataKeeper installation files.

You can install SIOS Protection Suite for Windows silently through the use of the `-silent` command line option. This option suppresses both the wizard and launcher user interfaces (UIs) resulting in what is considered a “silent installation.” This is how an installation is performed without any information displaying to or requiring any interaction with the end user. **Response files**, also known as “*options*” files, are used to pass command line options at installation. This is done as you would normally specify them on the command line to represent the responses to dialogs and/or to set the value of a property or variable. The options specified in the **response/options** file are executed after the execution of the options that were entered directly on the command line.

LifeKeeper Response File

To create a response file for LifeKeeper, open a command window and run the LifeKeeper setup program using the command `LK-{version}-Setup.exe /r /f1C:\setup.iss`. The responses entered to the dialogs will be recorded into the file `setup.iss`.

To perform a silent install using the created response file, open a command window and run the **LifeKeeper setup program** using the command:

```
LK-{version}-Setup.exe /s /f1C:\setup.iss /f2C:\setup.log
```

DataKeeper Response File

To create a response file for DataKeeper, open a command window and run the **SIOS DataKeeper setup program** using the command:

```
DK-{version}-Setup.exe /r /f1C:\setup.iss
```

The responses entered to the dialogs will be recorded into the file `setup.iss`.

Note: When creating the initial `setup.iss` file, if a local user server account is used for the DataKeeper service, you must edit the `setup.iss` file for use on other servers. This change can be made by opening the `setup.iss` file in Notepad and changing the name of the server found within the `szName`

field. (i.e.- szName=<serverName>\Administrator). When using the **Local Service account** or a **Domain account** that is the same across all installations, changing the *setup.iss* file is not required.

To perform a silent install using the created response file, open a command window and run the **SIOS DataKeeper setup program** using the command:

```
DK-{version}-Setup.exe /s /f1C:\setup.iss /f2C:\setup.log
```

Results from the silent install are stored in the file *setup.log*. "ResultCode=0" indicates a successful install. A negative result code indicates failure. Please check the operating system requirements for further information regarding the cause of failure.

When the SIOS Protection Suite install is finished, copy the license key(s) to the C:\Windows\SysWOW64\LKLicense folder or run the **License Key Installer** utility from the **Start-Programs** menu.

```
Start->All Programs->SIOS->DataKeeper->License Key Installer
```

Reboot the server.

Third Party Product Files

The following third party files were not developed by SIOS Technology Corp. but are installed during the SIOS Protection Suite/DataKeeper installation process.

Path and File Name	Provider	Purpose
<DK InstallPath>/lmdiag.exe <DK InstallPath>/lmhostid.exe <DK InstallPath>/lminstall.exe <DK InstallPath>/motdk_libFNP.dll	Flexera	License Management
<DK InstallPath>/SnapIn/IronPython.dll (.Net python language implementation) <DK InstallPath>/SnapIn/IronPython.Modules.dll (.Net python modules)	github.com/ IronLanguages/ ironpython2 (Microsoft open source)	Testing/Debugging
<DK InstallPath>/SnapIn/J832.Common.dll <DK InstallPath>/SnapIn/ J832.Wpf.BagOTricksLib.dll	Kevin Moore, http://j832.com/bagotricks/	Utilities and controls for WPF development
<DK InstallPath>/SnapIn/log4net.dll (.Net logging library)	Apache Software Foundation	Application logging
<DK InstallPath>/SnapIn/ Microsoft.Scripting.Core.dll <DK InstallPath>/SnapIn/Microsoft.Scripting.dll	github.com/ IronLanguages/ ironpython2 (part of IronPython)	

<DK InstallPath>/SnapIn/MMCFxCommon.dll <DK InstallPath>/SnapIn/ microsoft.managementconsole.dll	Microsoft	MMC managed snap-in library
<DK InstallPath>/VSSHelper/AlphaVSS-license.txt <DK InstallPath>/VSSHelper/ AlphaVSS.Common.dll <DK InstallPath>/VSSHelper/ AlphaVSS.Common.xml <DK InstallPath>/VSSHelper/AlphaVSS.x64.dll <DK InstallPath>/VSSHelper/log4net.dll <DK InstallPath>/VSSHelper/log4net.xml <DK InstallPath>/VSSHelper/cfg/ log4net.Config.xml	Pete Palotas, github.com/alphaleonis/AlphaVSS	Alpha VSS provider
<LK InstallPath>/Admin/kit/lpapp/bin/wpcap.dll <LK InstallPath>/Admin/kit/lpapp/bin/packet.dll	CACE Technologies	Gratuitous ARP Update
Note: By default is C:\Program Files (x86)\SIOS\DataKeeper		

Application Directory Anomaly

The following file is installed in a directory other than the default directory that you selected during the DataKeeper installation procedure. This exception occurs when the operating system installs [performance monitor counters](#).

Path and File Name	Purpose
<code><windows dir>/inf/ExtMirr/ ExtMirrCounters.h:</code>	Performance monitoring. This file contains counter names and definitions.

Uninstalling SIOS Protection Suite for Windows

Before Removing LifeKeeper

Included below are the requirements for removing LifeKeeper software.

1. **Move or stop applications.** Before removing the software, verify that applications requiring SIOS Protection Suite protection are not on the server. Never remove LifeKeeper from a server where an application resource hierarchy is in service. Removing LifeKeeper removes all configuration data, such as equivalencies, resource hierarchy definitions and log files. See [Transferring Resource Hierarchies](#) for additional information.
2. **Ensure LifeKeeper is running.** Recovery Kits may require LifeKeeper to be running when you remove the recovery kit software. Use the **Services MMC** snap-in to ensure that LifeKeeper services are running. If it is not running, the removal process cannot remove the resource instances from other SIOS Protection Suite servers in the cluster which would leave the servers in an inconsistent state.
3. **Remove resource hierarchies.** Unextend or delete any resource hierarchies from the server where LifeKeeper will be removed. Never remove a Recovery Kit from a server where the resource hierarchy is in service. This will corrupt current hierarchies and they will need to be recreated when reinstalling the Recovery Kit.
4. **Remove all packages.** If removing the LifeKeeper core, first remove other packages that depend upon LifeKeeper; for example, SIOS Protection Suite Recovery Kits. It is recommended that before removing a SIOS Protection Suite Recovery Kit, first remove the associated application resource hierarchy.

Before Removing DataKeeper

If planning to uninstall DataKeeper and reinstall a previous version, all jobs/mirrors must be deleted on each node prior to uninstalling. These will need to be recreated once software is reinstalled.

Uninstall SIOS Protection Suite

- In **Windows Control Panel**, find your list of installed programs and select **SIOS DataKeeper** or

LifeKeeper.

- Select **Uninstall**.

Once the uninstall process is complete, rebooting the system is required.

Note: Uninstalling automatically stops the SIOS DataKeeper and/or LifeKeeper services and clears the registry entries.

Once removed, the following files will not be removed by the uninstall procedure.

Path and File Name	Definition and Special Considerations
<code><windows dir>/SysWOW64/ LKLicense</code>	<p>Common license file directory for SIOS Technology Corp. products. This is where license files are installed and licenses for multiple SIOS Technology Corp. products may be installed here at any given time. We don't remove this at uninstall so as to not disturb the installed licenses.</p> <p>Safe to remove manually, but the license will need to be reinstalled if the software is reinstalled at a later time.</p>
<code><windows dir>/SysWOW64/ PerfStringBackup.ini</code>	<p>A backup file created by Windows when new performance monitor counters are installed. This is created when we install the perfmon counters.</p> <p>This should probably be left alone since it is a file created by Windows itself.</p>
<code><windows dir>/inf/ ExtMirr/0011/ ExtMirrCounters.ini</code>	<p>This file describes the DataKeeper performance monitor counters. This file can be removed or left alone. It is not an executable.</p>

Notes

- **Important:** Uninstallation of SIOS Protection Suite software requires that the Microsoft Visual C++ 2008 Redistributable package be installed. Do not remove this package until SIOS Protection Suite has been uninstalled.
- **Modify** or **Repair** must be run from the SIOS Protection Suite setup program.

- Removal of SIOS Protection Suite may NOT delete the SIOS Protection Suite directory. This directory can be deleted manually after the **Add/Remove** operation is complete.
- A reboot of the system is required to completely remove SIOS Protection Suite remnants.

Upgrading SIOS Protection Suite

You may upgrade from previous versions of SIOS Protection Suite for Windows while preserving your resource hierarchies and mirrors by using the procedure below.

Upgrade Procedure

The following scenario illustrates the upgrade process when upgrading both LifeKeeper and SIOS DataKeeper. The upgrade should be performed on LifeKeeper prior to upgrading SIOS DataKeeper. The LifeKeeper Services and SIOS DataKeeper Service will be stopped during the upgrade process. A system reboot is required after upgrading both LifeKeeper and SIOS DataKeeper.

Given two systems (Sys1 and Sys2), with Sys1 being the primary (active) server, perform the following steps to upgrade LifeKeeper and SIOS DataKeeper:

Upgrading the Backup Server

1. Exit the LifeKeeper GUI and SIOS DataKeeper GUI on backup server *Sys2*.
2. Open a command window and enter `$LKROOT\bin\lkstop` (where `$LKROOT` is the SIOS Protection Suite installation path, by default `C:\LK`) to stop all the LifeKeeper services. Wait until you see "LIFEKEEPER NOW STOPPED" before continuing.
3. Upgrade LifeKeeper for Windows on the backup server *Sys2* by running the setup program. Click **Yes** to continue upgrading LifeKeeper.
4. The existing LifeKeeper files will be overwritten by the LifeKeeper installation. You should install your new LifeKeeper license (if necessary) using the **License Manager** utility – pre-7.0 LifeKeeper licenses will not work with LifeKeeper 7.0. Do not reboot the backup server until SIOS Data Replication is upgraded to SIOS DataKeeper.
5. Upgrade SIOS DataKeeper for Windows on the backup server *Sys2* by running the setup program. Click **Yes** to continue upgrading SIOS DataKeeper. You should install your new DataKeeper license (if necessary) using the **License Manager** utility – SIOS Data Replication licenses will not work with SIOS DataKeeper.
6. [Upgrade the Language Supplement Package](#) (if required) and any optional recovery kits at this time

by running the appropriate installation program.

7. Reboot the backup server `Sys2`.

For additional backup servers in your cluster, follow these steps on each server.

Note: Newer versions of SIOS Protection Suite contain links to the SIOS Technical Documentation in lieu of being included in the install package. When performing an upgrade from previous versions which contained the Online Product Manual within the product, the upgrade will not uninstall the old Online Product Manual files. If you would like for these files to be removed, you will need to manually uninstall the OLPM package.

Upgrading the Primary Server

8. Once backup server has been rebooted, allow mirror(s) to resync and return to the **Mirroring** state.
9. Perform a switchover. This will bring the active resource hierarchies In Service on `Sys2` and will reverse the role of the mirror(s) allowing the primary server `Sys1` to be upgraded.
10. The above procedure will be repeated on the primary server `Sys1`. Exit the LifeKeeper GUI and SIOS DataKeeper GUI.
11. Open a command window and enter `$LKROOT\bin\lkstop` (where `$LKROOT` is the LifeKeeper installation path, by default `C:\LK`) to stop all the LifeKeeper services. Wait until you see "LIFEKEEPER NOW STOPPED" before continuing.
12. Upgrade LifeKeeper for Windows on the primary server `Sys1` by running the Setup program. Click **Yes** to continue upgrading LifeKeeper.
13. The existing LifeKeeper files will be overwritten by the LifeKeeper installation. You should install your new LifeKeeper license (if necessary) using the **License Manager** utility – pre-7.0 LifeKeeper licenses will not work with LifeKeeper 7.0. Do not reboot the server until SIOS DataKeeper is upgraded.
14. Upgrade SIOS DataKeeper for Windows on the primary server `Sys1` by running the Setup program. Click **Yes** to continue upgrading SIOS DataKeeper. You should install your new DataKeeper license (if necessary) using the **License Manager** utility – SIOS Data Replication licenses will not work with SIOS DataKeeper.
15. [Upgrade the Language Supplement Package](#) (if required) and any optional recovery kits at this time

by running the appropriate installation program.

16. Reboot the primary server `Sys1`.
17. Start the LifeKeeper GUI on `Sys1` by clicking **Start**, and then point to **Programs**, then **LifeKeeper**, then **LifeKeeper GUI** and log in to `Sys1`.

Upgrading from SIOS Data Replication v6.2x to DataKeeper

Because DataKeeper incorporated a new structure called a “job”, upgrading from SIOS Data Replication to DataKeeper requires that you delete your existing mirrors before upgrading to DataKeeper and then recreate them after the upgrade is complete. This insures that the job and mirror information gets set up properly for DataKeeper.

DataKeeper also requires updated licensing, so you will have to install your new DataKeeper licenses when the License Manager screen is presented. We also recommend removing your old SDR v6.2x licenses at this time.

The procedure is exactly the same as the upgrade procedure above, with two exceptions underlined below.

Upgrade Procedure

1. In the SIOS Data Replication UI, **delete all existing mirrors.**
2. Close the SIOS DataKeeper Replication UI if it is currently running.
3. Perform the upgrade procedure listed above and **apply new licensing on each server when prompted.**
4. Bring up the DataKeeper UI and recreate your mirrors.


Reinstalling SIOS Protection Suite

To reinstall SIOS Protection Suite, perform the same procedures as above, the only exception being that when Setup presents a list of InstallShield options, select **Repair**.

Repair

The Install process also allows repairing the SIOS Protection Suite software. Use this option if the software that was previously installed was accidentally deleted or if the user is performing a point release upgrade. This option copies all the files from the setup folder and prompts the user to reboot the system.

✿ **Note:** It is possible to get errors during repair, we noticed these errors specifically with Windows 2012. If you encounter such errors, retry the operation.



Repair SIOS Protection Suite in the Windows Control Panel

- In **Windows Control Panel**, find your list of installed programs and select **SIOS DataKeeper** or **LifeKeeper**.
- Select **Repair**.

Once the repair process is complete, rebooting the system is required.

Starting LifeKeeper

With a typical installation, LifeKeeper is started automatically when the server is booted. Your applications are brought up in a protected state.

When LifeKeeper starts, it also starts the LifeKeeper GUI Server. The LifeKeeper GUI client is launched from a web browser or from the **Start->All Programs->SIOS->LifeKeeper->LifeKeeper (Admin Only)**, and is described in detail in the [LifeKeeper GUI](#) section of [SIOS Protection Suite for Windows Technical Documentation](#).

Starting and Stopping LifeKeeper Processes

Because LifeKeeper is started automatically when the server is booted, you should not normally need to start/stop LifeKeeper. In the rare event that you need to start or stop LifeKeeper manually, you can do so in one of two ways:

Services MMC Snap-In

You can stop and start LifeKeeper services using the **Services MMC** snap-in under **Administrative Tasks**.

LifeKeeper consists of two services:

- LifeKeeper
- LifeKeeper External Interfaces

Generally, these two services should be stopped and started together. However, since LifeKeeper External Interfaces is a dependency of the LifeKeeper service, stopping it will also stop the LifeKeeper service. Likewise, it must be started before the LifeKeeper service can be started.

Command Line

When stopping LifeKeeper, there are a number of related services that must be stopped. This process can take several seconds, although the Services tool does not reflect exactly when all the services are stopped. Using the command line to enter `$LKROOT\bin\lkstop` will show the services as they are being stopped, and when completed, the message "LIFEKEEPER NOW STOPPED" will display as confirmation.



Stopping LifeKeeper takes all protected hierarchies out of service. This means that any protected applications will not be accessible.

SIOS Protection Suite for Windows Technical Documentation

Introduction

About SIOS Protection Suite for Windows

SIOS Protection Suite (SPS) for Windows integrates high availability clustering and data replication functionality to protect mission-critical data and applications. SIOS Protection Suite also integrates with SIOS' Application Recovery Kits, which provide application-aware agents for SQL Server and Oracle and more enabling you to make more informed recovery decisions.

SIOS Protection Suite for Windows Integrated Components

LifeKeeper for Windows provides a fault resilient software solution to provide high availability for data, applications, and communications resources. LifeKeeper does not require any customized, fault-tolerant hardware. You simply group two or more systems and enter site-specific configuration data. Then, LifeKeeper automatically provides fault detection and recovery.

In case of a failure, LifeKeeper migrates protected resources from the failed system to the backup system. Users experience a brief interruption during the actual switchover, but LifeKeeper automatically restores operations on the backup system when it completes the failover recovery.

SIOS DataKeeper is a highly optimized host-based replication solution which ensures your data is replicated as quickly and as efficiently as possible from your source server across the network to one or more target servers.

Optional Microsoft SQL Server Recovery Kit provides granular visibility into SQL Server so you can ensure your systems are properly responding – not just running.

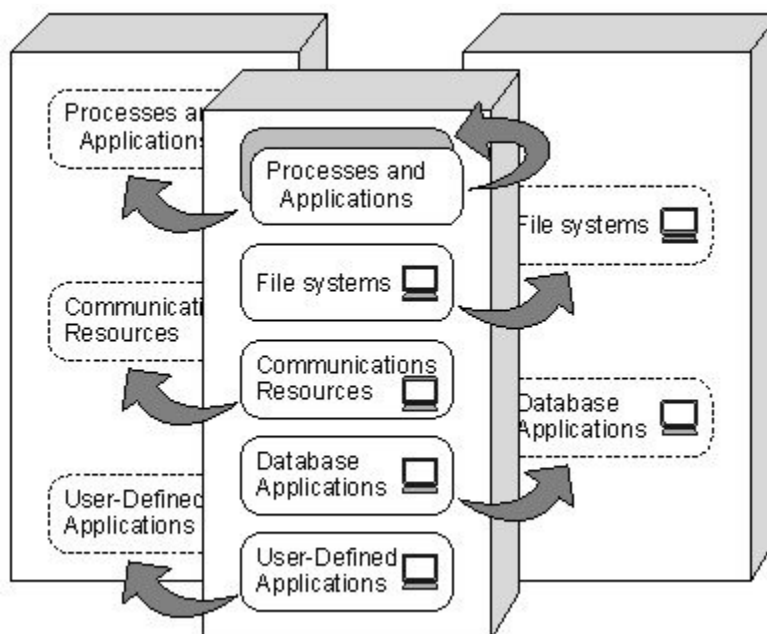
Protected Resources

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

- **Volume.** With LifeKeeper's volume resource type, you can protect data and applications on shared SCSI peripherals or replicated volumes by creating a resource in the SIOS Protection Suite hierarchy for the disk volume containing those resources.

- **File share.** SIOS Protection Suite's File Share resource lets you protect a specific folder or directory on a shared drive.
- **Computer alias name.** SIOS Protection Suite's LAN Manager Recovery Kit enables automatic failover of the computer alias name for applications that communicate to the server via NetBEUI.
- **Communications resources.**
 - SIOS Protection Suite's IP Recovery Kit allows you to create resources that enable switchover of IP addresses.
 - SIOS Protection Suite's DNS Recovery Kit provides a mechanism to update DNS A and PTR records.
- **Database application.** SIOS Protection Suite provides an optional Recovery Kit for Microsoft SQL Server.
- **Generic applications.** SIOS Protection Suite's Generic Application Recovery Kit allows the creation of resources for an application that has no predefined recovery kit.

SIOS Protection Suite supports N-Way recovery for a range of resource types. N-way recovery allows different resources to fail over to different backup servers in a cluster.



See [SIOS Protection Suite Core Software](#) for the core components available with SIOS Protection Suite.

[Core](#)[MSSQL](#)[Communication Paths](#)[Communication Path Types](#)[Heartbeat](#)[Heartbeat Interval](#)[Safety Check](#)[Resource Hierarchies](#)[Hierarchy Relationships](#)[Resource Hierarchy Information](#)[Resource States](#)[Shared Equivalencies](#)

Core

SIOS Protection Suite Core Software

SIOS Protection Suite for Windows Core includes LifeKeeper and the basic LifeKeeper software packages, DataKeeper and the Core Recovery Kits.

- **LifeKeeper** – The LifeKeeper for Windows core includes the basic LifeKeeper software packages plus the Core Recovery Kits.
 - **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.
 - **LifeKeeper GUI** – The LifeKeeper GUI is a client/server application that provides a graphical administration interface to LifeKeeper and its configuration data. The LifeKeeper GUI client is implemented as both a stand-alone application and as a Java applet invoked from a web browser.

- **DataKeeper** – The DataKeeper package provides real-time data replication that allows you to keep Windows Server data in sync across multiple servers and data centers.

SIOS Protection Suite Core Recovery Kits

The Core Recovery Kits provide all fault detection and recovery mechanisms and the definitions necessary to provide SIOS Protection Suite protection for volumes (drive letters), file shares, communications resources and Microsoft Internet Information Server.

- **Volume Recovery Kit** – Allows you to create a resource to protect an entire shared or mirrored drive (for instance, the K: drive).
- **LAN Manager Recovery Kit** – Enables automatic recovery of the computer alias name and Windows file share lists. The LAN Manager resource lets you create a “switchable” computer name for applications that communicate with the server via NetBEUI, or system name. It includes the File Share recovery component.
- **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 SIOS Protection Suite 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 SIOS Protection Suite 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. This kit also provides local recovery of the IP resource.
- **DNS Recovery Kit** – Provides a mechanism to update DNS A and PTR records of the primary server or an alias name. After a failover or switchover to the backup server, the A record and the associated PTR record (if exists) of the primary server or alias name will be updated with the IP address of the backup server.
- **Microsoft IIS Recovery Kit** – Protects Web, FTP and SMTP services of the Microsoft Internet Information Services (IIS). It continuously monitors the health of your Internet servers, and if a problem arises, provides automatic failover of the affected web server to a backup server.
- **Generic Application Recovery Kit** – Allows the creation of resources for an application that has no predefined recovery kit.

MSSQL

SIOS Protection Suite Microsoft SQL Server Recovery Kit

SIOS Protection Suite Microsoft SQL Server Application Recovery Kit includes tools and utilities that allow SIOS Protection Suite to manage and control the Microsoft SQL-based database application. When installed, SIOS Protection Suite is able to monitor the health of the application and automatically recover the application if it fails. The SIOS Protection Suite Recovery Kit is non-intrusive and requires no changes within the application in order for SIOS Protection Suite to protect it.

This recovery kit provides menu-driven definition of resources for automatic switchover of a Microsoft SQL Server instance. The kit provides options that allow you to easily create a complete resource hierarchy so that the recovery operation can include all disk resources used by the SQL Server as well as the Named Pipe or IP socket resources used to access the database.

For further information, refer to the [Microsoft SQL Server Recovery Kit Technical Documentation](#).

Note: Product Requirements and FAQ information is delivered in the OS specific Release Notes.

Communication Paths

Overview

SIOS Protection Suite monitors resource operations and provides failover using shared communication paths (comm paths) between servers. It is critical to SIOS Protection Suite fault detection and resource recovery that communication between the servers remains active. As a result, you need to define multiple comm paths using different protocols to guard against a system failover simply because a communication line fails.

Before you can define resources or resource hierarchies in SIOS Protection Suite, you must define your comm paths on each of the servers. SIOS Protection Suite uses the comm paths to coordinate resource definitions and states between the nodes and for fault detection and recovery.

The [Communication Path Types](#) section describes these comm path topics:

- **Comm path types.** SIOS Protection Suite supports two types of comm paths for two-server configurations: TCP/IP and Shared Disk. Configurations with greater than two servers support only TCP/IP comm paths.
- **SIOS Protection Suite heartbeat.** A key function of the comm path between the servers is the SIOS Protection Suite heartbeat. This periodic signal between the servers assures each server that the other is still alive and processing.
- **Safety check.** If all comm paths die, SIOS Protection Suite performs a safety check to verify system failure before switching over resources.

Communication Path Types

SIOS Protection Suite provides three different types of comm paths so that you can define redundant comm paths using different protocols. Although there is less value in defining multiple paths of the same type over the same media, redundant paths using different protocols or different media provide good protection against failover due to a comm path failure.

When you define your comm paths, you establish a priority for each path. SIOS Protection Suite uses the paths in priority order for inter-node communication. However, SIOS Protection Suite sends heartbeat signals over all active comm paths. These are the comm paths supported by SIOS Protection Suite and the default priority range assigned to each:

- **TCP/IP (socket).** The LAN comm path is given the highest priority. The default priority range for the socket path is 1 to 30. You can define multiple LAN comm paths between a pair of servers by defining each one over a different adapter to prevent false failovers.

Note: The LifeKeeper GUI uses TCP/IP for communicating status information about protected resources; if there are two TCP/IP comm paths configured, SIOS Protection Suite uses the comm path with the highest priority for communicating resource status.

- **Shared disk.** SIOS Protection Suite allows you to define a raw disk partition on a shared disk as a communication location for a pair of servers in the cluster. The shared disk path must be identified with the same drive letter on both servers and the drive letter must identify the same disk partition. The disk partition is usually small, typically one megabyte. The default priority range for the shared disk comm path is 61 to 99 (not supported in greater than two-server configurations).

Note: Shared Disk comm paths are used by SIOS Protection Suite only for detecting whether other servers in the cluster are alive. Therefore if the TCP/IP comm path used by the LifeKeeper GUI is down, the GUI will show hierarchies on other servers in an UNKNOWN state, even if the shared disk or secondary TCP/IP comm path is operational.

More About the Shared Disk Comm Path

The shared disk comm path can be used as a channel of last resort in the case where all other communication has been severed. If the shared disk comm path were to be lost as well, it is very likely that at least one of the servers would not be able to access the storage subsystem, thereby preventing a “split-brain” situation where both servers may access the same disk resource simultaneously.

CAUTIONS:

- A SIOS Protection Suite configuration should include no more than one shared disk comm path between any two servers in the cluster.
- Before using shared disk comm paths on JBOD or Host-based RAID, be sure to test the comm path for reliability when a member of the cluster is shut down or out of service. Sometimes, in configurations using JBOD or Host-based RAID, the comm path will fail to go down when a cluster member goes down, and therefore a failover is not initiated.

SIOS Protection Suite Heartbeat

The heartbeat is a key SIOS Protection Suite fault detection mechanism. The heartbeat is a periodic signal sent over the comm path between a pair of servers. The regular signals tell each server that the other is still active. When you define your comm path, the definition sets the [heartbeat signal interval](#) in seconds and specifies the number of consecutive heartbeats a server can miss before marking the comm path as dead.

When the SIOS Protection Suite servers mark a comm path as dead, inter-node communications immediately commence over the comm path with the next highest priority. Only when the server fails to receive the heartbeat signal on all comm paths does SIOS Protection Suite initiate the [safety check](#) to determine the need for failover recovery.

Heartbeat Interval

The SIOS Protection Suite heartbeat interval is the number of seconds between heartbeat signals that verify the servers are alive. The default (and recommended) interval is six seconds.

- If you wish to set the interval at the minimum allowed value of four seconds, then you should ensure that the communication path is configured on a private network and tested thoroughly since values lower than five seconds create a risk of false failovers due to network interruptions.
- The heartbeat interval works in conjunction with the maximum heartbeat misses, which has a default value of five (recommended). Setting the maximum heartbeat misses to a lower number (3 or 4) can create a risk of false failovers. Be sure to test thoroughly in your environment.
- Setting these values too high can effectively disable SIOS Protection Suite's ability to detect a failure.

Safety Check

When all the communications paths on a server are DEAD, SIOS Protection Suite assumes that the paired system is DEAD (or down) and attempts to fail over. However, SIOS Protection Suite performs a safety check to ensure that the failure occurred in the server rather than just the comm paths.

The safety check queries on the network (through LAN Manager) to see if the machine still exists. One of two events can occur:

- **System is alive.** If the check receives a response that the system does exist on the network, it aborts the failover and reports the following message to the LifeKeeper event log:

```
SAFETY CHECK FAILED: COMM_DOWN ABORTED
```

- **System is dead.** If the check does not receive a response within a specified time-out period (default 8 seconds), the machine is assumed to be down and the failover proceeds.

SIOS Protection Suite performs this check only once, after all comm paths go down. If the safety check detects that the system is alive, failover is aborted. SIOS Protection Suite does not re-initiate failover until all of the following events happen in sequence:

1. At least one of the comm paths comes back **ALIVE**.
2. All comm paths again go **DEAD**.
3. The safety check activates and does not detect that the paired system is alive.

Resource Hierarchies


The LifeKeeper GUI enables you to create a resource hierarchy on one server and extend that hierarchy to one or more backup servers. SIOS Protection Suite then automatically builds the designated hierarchies on all servers specified. SIOS Protection Suite 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, SIOS Protection Suite manages the stopping and starting of the resources within the hierarchy. The following topics provide background for hierarchy definition tasks:

- [Resource States](#)
- [Hierarchy Relationships](#)
- [Shared Equivalencies](#)
- [Resource Hierarchy Information](#)

Hierarchy Relationships

SIOS Protection Suite 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.



In the example above, `MSEch.0` is an Exchange resource, which has three dependencies – a DNS resource (`DNS.0`) and two volume resources (`Vol.L` and `Vol.X`).

The dependency relationships specified by the resource hierarchy tell SIOS Protection Suite the appropriate order for bringing resource instances in service and out-of-service. In the example resource hierarchy, SIOS Protection Suite cannot bring the `MSEch.0` resource into service until it successfully brings into service the DNS and volume instances.

Resource Hierarchy Information

A snapshot of information about all the resources defined for a server can be displayed in the **Server Properties** dialog box.





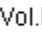













General CommPaths Resources			
Name	Application	Resource Type	State
DNS.0	comm	dns	ISP
Vol.L	filesys	volume	ISP
Vol.X	filesys	volume	ISP
Vol.M	filesys	volume	ISP
MSExch.0	mail	msexch	ISP

Other resource information can be viewed in the [Status Table](#) (main GUI window) or in the [Viewing Resource Properties](#) topics.

Resource States

The LifeKeeper GUI status display shows the resources that are defined across all servers to which it is connected. The left pane of the status window displays the **Resource Hierarchy Tree** which reflects the global resource status (that is, the status of the resource across all servers).

The right pane of the status window contains columns showing the status of each individual resource on each server.

Hierarchies			
 Unprotected			
  MSEXch.0 <ul style="list-style-type: none"> DNS.0 Vol.L Vol.X Vol.M	 CARDINAL		 BLUEJAY
	 Active	1	 StandBy 10
	 Active	1	 StandBy 10
	 Active	1	 StandBy 10
	 Active	1	 StandBy 10
	 Active	1	

The sample above shows a hierarchy MSEXch.0 with a status of **In Service, Protected (ISP)**. The resource *Vol.M* exists only on CARDINAL. Thus it is **In Service**, but it has no failover protection, which is indicated by the yellow triangle.

For more details on resource states, see [Viewing the Status of Resources](#).

Shared Equivalencies

When you create a SIOS Protection Suite resource hierarchy, you create the hierarchy initially on the primary server and extend the hierarchy to a backup server. Most resource instances can be active on only one server at a time. For such resources, SIOS Protection Suite 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, a shared equivalency exists between each hierarchy level on a pair of servers. For example, the `MSExch.0` resource exists on both servers, and there is a shared equivalency between the two instances of `MSExch.0` (just as there is between the one DNS instance and the two volume instances).

Hierarchies					
Active Protected		CARDINAL		BLUEJAY	
MSExch.0	Active	1	↓	StandBy	10
DNS.0	Active	1	↓	StandBy	10
Vol.L	Active	1	↓	StandBy	10
Vol.X	Active	1	↓	StandBy	10

Shared Equivalencies

Configuration

If the SIOS Protection Suite environment has been installed, the SIOS Protection Suite software can be configured on each server in the cluster. The topics in this section will help with this configuration.

[SIOS Protection Suite Configuration Steps](#)

[Active-Active Grouping](#)

[Active-Standby Grouping](#)

[Intelligent Versus Automatic Switchback](#)

[SIOS Protection Suite Configurations](#)

[Common Hardware Components](#)

[System Grouping Arrangements](#)

SIOS Protection Suite Configuration Steps

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

1. Ensure that the LifeKeeper services are running by checking the **Services** in the **Administrative Tools** on the **Control Panel**. You should see both *LifeKeeper* and *LifeKeeper External Interfaces* services. If they are not both running, start them now.

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

2. Users with administrator privileges on a SIOS Protection Suite server can run the application client from that server. Click **Start**, then point to **All Programs**, then **SIOS->LifeKeeper->LifeKeeper (Admin Only)**.

After the application is loaded, the **LifeKeeper GUI** appears and the **Cluster Connect** dialog is displayed. Enter the **Server Name** you wish to connect to, followed by the **login** and **password**.

3. [Create Communication Paths](#). Before you can activate SIOS Protection Suite protection, you must create the communication path (heartbeat) definitions within LifeKeeper.
4. Set your [Server Shutdown Strategy](#). This tells LifeKeeper whether to switch over resources when you initiate an orderly shutdown.
5. SIOS Protection Suite is now ready to protect your applications. The next step depends on which SIOS Protection Suite Recovery Kit(s) you will be using:
 - If you are using a Core Recovery Kit, then refer to the topics for creating [Volume](#), [DNS](#), [IP](#), [File Share](#), [LAN Manager](#), or [Generic Application](#) hierarchies.
 - If you are using the optional Microsoft SQL Server Recovery Kit, refer to the [Microsoft SQL Server Recovery Kit Administration Guide](#) for instructions on creating and extending your resource hierarchies.


Active-Active Grouping

In an active/active group, 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 *Server 1* and/or *Server 2* as a third or even fourth backup system.



Note: Because SIOS Protection Suite applies locks at the volume 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.

Active-Standby Grouping

In an active/standby group, 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 SIOS Protection Suite resource definitions specify the following active/standby paired relationships:

- *AppA* on *Server 1* fails over to *Server 2*.
- *AppB* on *Server 3* fails over to *Server 2*.
- *AppC* on *Server 4* fails over to *Server 2*.

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 slices on the same volume. SIOS Protection Suite applies locks at the volume level. When the SCSI locks are applied, only one system on the shared SCSI bus can access volumes on the disk device. 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.

- **SIOS Protection Suite administration.** In the example, *Server 2* provides backup for three other servers. In general, it is not desirable to administer the SIOS Protection Suite 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.

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. SIOS Protection Suite offers an **automatic switchback** option as an alternative to the normal *intelligent switchback* behavior described above. This option can be selected for individual resource hierarchies on individual servers. If *automatic switchback* is selected for a resource hierarchy in the In-Service-Protected (ISP) state running 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:

- If using data replication (DataKeeper), you must choose **intelligent switchback**. *Automatic switchback* is not supported.
- Checks for *switchback* are only made either when LifeKeeper starts or when a new server is added to the cluster; they are not performed during normal cluster operation.
- SIOS Protection Suite never performs an *automatic switchback* from a higher priority server to a lower priority server.

If there is a dependency between two resources with different switchback strategies, the switchback strategy of the parent resource takes precedence.

SIOS Protection Suite Configurations

SIOS Protection Suite works on the basis of resource hierarchies you define for groups of two or more servers. The following three topics introduce the SIOS Protection Suite failover configuration concepts.

[Common Hardware Components](#)


[System Grouping Arrangements](#)

[Resource Hierarchies](#)

Common Hardware Components

All SIOS Protection Suite configurations share these common components as illustrated in the diagram below:

1. **Server Groups.** The basis for the fault resilience provided by SIOS Protection Suite is clustered Windows servers. The servers, also referred to as SIOS Protection Suite nodes, do not have to be the same hardware platform.
2. **Communication paths for heartbeat.** It is strongly recommended that each pair of servers in the group share at least two communication paths (comm paths), although only one is required. To avoid unnecessary failover due to communication failure, you should configure your redundant comm paths using different protocols and communication media, for example TCP/IP (or socket). SIOS Protection Suite uses the comm paths to coordinate resource availability for the fault-detection heartbeat, a periodic message between nodes and for switchover of resources. (See [Overview of Communication Paths.](#))
3. **Shared data resources.** SIOS Protection Suite can recover and restore shared or mirrored data, applications and communication resources. SIOS Protection Suite controls access at the volume (drive letter) level. In case of a server failure, SIOS Protection Suite automatically switches availability of protected resources to an active server. Peripheral devices that are to be shared between systems must be packaged in external peripheral cabinets. See the [Configuring Your Storage](#) topic for information to help you configure your shared storage.
4. **Shared communication for user connections.** SIOS Protection Suite can also automatically manage the switchover of user communication resources, such as IP addresses, computer alias names and file share lists. Switchover of communication resources allows users to connect using their normal paths.




System Grouping Arrangements

A resource hierarchy is defined on a cluster of SIOS Protection Suite 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 SIOS Protection Suite 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.



Configuring SPS for Multibyte Language Encodings

LifeKeeper can operate in all system locales that are supported by Windows. However, locales that use multibyte character encodings (such as Japanese, Chinese, Korean, and other Eastern languages) may require that LifeKeeper be configured to use the correct encoding based on the system locale.

In order to do this correctly, a configuration file has been included in the core software which allows the user to specify the locale (including encoding) based on the system codepage, which is a numeric value that corresponds to the Windows locale for the system. The file is %LKROOT%\cygwin\usr\share\locale\locale.from.codepage. This text file can be customized if necessary – any lines that are not commented should be blank or should contain two values separated by a tab – the codepage numeric value, and the locale string to be used.

The locale.from.codepage file has already been populated with these values:

```
932 ja_JP.SJIS
```

```
936 zh_CN.GB2312
```

To find a system's codepage, you can run the powershell command:

```
[System.Text.Encoding]::Default.CodePage
```

Note that LifeKeeper will operate correctly in a multibyte locale without this file, as long as LifeKeeper does not need to access system resources whose name or value contains a multibyte character. For example, during creation of an IP Address, the name of the NIC is used. If the NIC is not named with a multibyte character, the IP Address creation will work regardless of whether this file has been configured for the system's codepage.

Administration

SIOS Protection Suite Administration Overview

SIOS Protection Suite provides two administration interface options:

- LifeKeeper GUI
- LifeKeeper command line interface

The LifeKeeper GUI is used for the following tasks which are listed in the typical sequence for configuring SIOS Protection Suite.

- **Communication path definition.** You must define the communication paths you want to use before you define any other resource instances or hierarchies in LifeKeeper. This can be done using the **Edit** menu or the **Create Comm Path** icon on the GUI toolbar.
- **Resource definition.** As you install recovery kits, the resource types supported by those kits appear in the **Create Resource Hierarchy** dialog box. For most recovery kits, the necessary dependencies will be created automatically.
- **Monitoring.** The LifeKeeper GUI's status display provides a visual status of resources protected by SIOS Protection Suite on the connected servers. In addition, SIOS Protection Suite maintains log files which you can view through the GUI.
- **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 SIOS Protection Suite protection, they should be started and stopped only through SIOS Protection Suite.

For initial step-by-step configuration instructions, see [SIOS Protection Suite Configuration Steps](#).

See the [GUI Tasks](#) and [Maintenance Tasks](#) topics for detailed instructions on performing SIOS Protection Suite administration, configuration and maintenance operations using the GUI.



Note: SIOS Protection Suite is set up so that the SIOS Protection Suite services are run by the local system account on each server. SIOS Protection Suite should not be changed to

run as any other user account.

[Administrator GUI Tasks](#)

[Working With Resource Hierarchies](#)

[Man Pages](#)

[LKSUPPORT](#)

[Setting Browser Security Parameters](#)

[IP Local Recovery](#)

[Overview of SIOS Protection Suite Event Forwarding via SNMP](#)

[Java Upgrade](#)

Administrator GUI Tasks

[Editing Server Properties](#)

[Set Server Shutdown Strategy](#)

[Server Properties](#)

[Disabling Automatic Failover](#)

[Creating a Communication Path](#)

[Deleting a Communication Path](#)

Editing Server Properties

1. To edit the properties of a server, begin just as you would for [viewing server properties](#).
2. If you are logged in to that server with the appropriate permissions, the following items will be editable.
 - Shutdown Strategy
 - Automatic Failover Configuration
 - Server Configuration (only for servers with specialized configuration settings)
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.

Set Server Shutdown Strategy

The Shutdown Strategy is a 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)	SIOS Protection Suite will not switch over resource hierarchies during an orderly shutdown.
Switch Over Resources	SIOS Protection Suite will switch over all resource hierarchies during an orderly shutdown.

Restriction: The **Switch Over on Shutdown** setting is not supported with SIOS DataKeeper resources.

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, and if you wish, change the Shutdown Strategy to ***"Switch Over Resources"***.

For each server in the cluster:

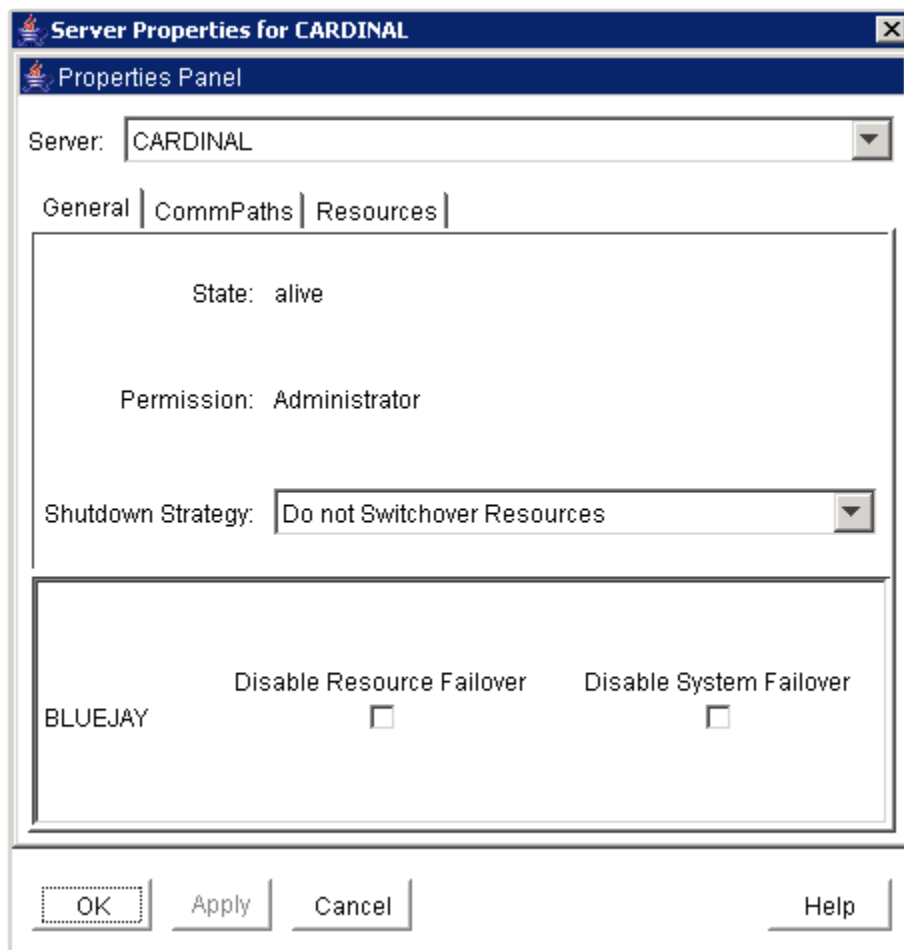
1. Bring up the **Server Properties** dialog just as you would for [viewing server properties](#).
2. On the [General Tab](#) of the **Server Properties** dialog, select the **Shutdown Strategy**.

Note: The LifeKeeper process must be running during an orderly shutdown for the Shutdown Strategy to have an effect. If LifeKeeper is not running or the resources are not currently in service, the resources will not switch over.

Server Properties

The **Server Properties** dialog is available from the [Edit Menu](#) or from a server popup menu. This dialog displays the properties for a particular server. When accessed from the **Edit** menu, you can select the server. The **Server Properties** dialog updates itself when the selected server changes.

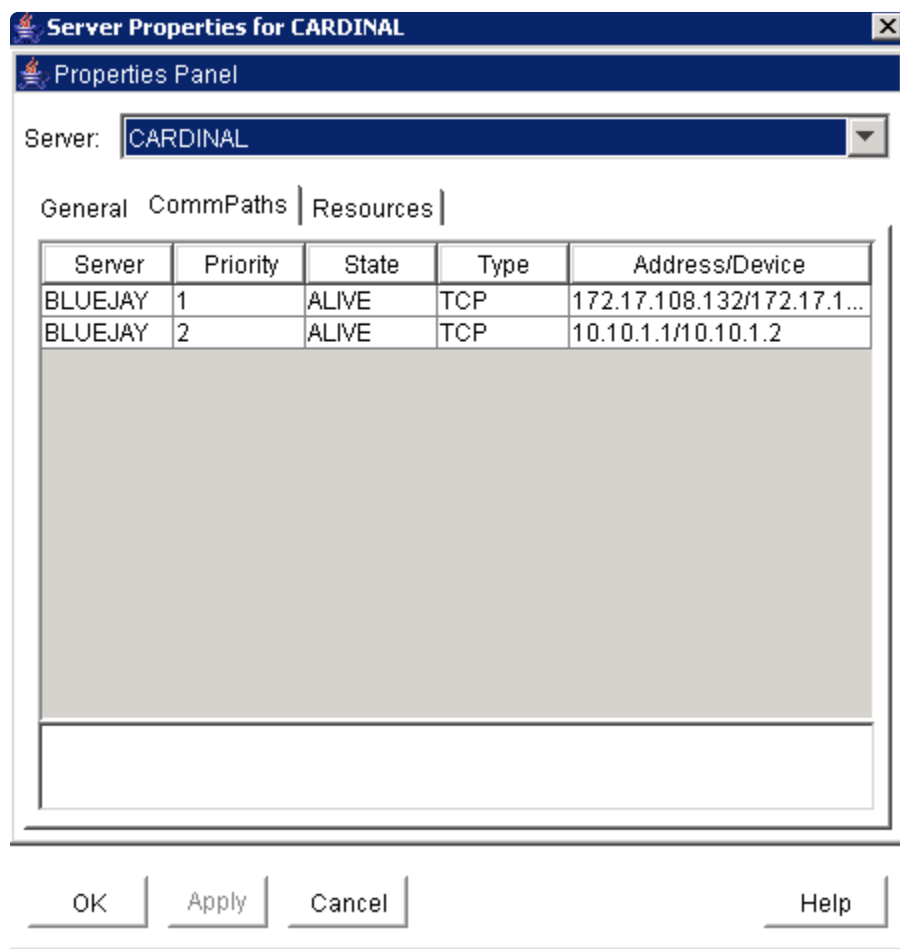
The **OK** button applies any changes that have been made and then closes the window. The **Apply** button applies any changes that have been made. The **Cancel** button closes the window without saving any changes made since **Apply** was last clicked.



- **Name.** Name of the selected server.
- **Permission.** The permission level of the user currently logged into that server. These are the possible permission values:

- **Administrator** – the user can perform any SIOS Protection Suite task.
- **Operator** – the user can monitor SIOS Protection Suite resource and server status, and can bring resources in service and take them out of service.
- **Guest** – the user can monitor SIOS Protection Suite resource and server status.
- **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.
- **ShutdownStrategy** (editable). The setting that governs whether or not resources which are in service are switched over to a backup server in the cluster when a server is shut down. The setting “**Switch Over Resources**” indicates that resources will be brought in service on a backup server in the cluster. The setting “**Do not Switch Over Resources**” indicates that resources will not be brought in service on another server in the cluster.
- **Server Name.** Automatic failover capabilities from the local server to other servers in the cluster may be configured here. All servers in the cluster should be operational (i.e. at least one SIOS Protection Suite comm path must be active) as inactive servers are not listed. The name of each active server in the cluster is listed, excluding the local server. For each server, two types of failover capabilities are configurable. By default, all failover capabilities are enabled.
 - **Disable Resource Failover** – Select the remote server(s) to be disqualified as a backup server for any failed resource hierarchy on the local server. When disabled, the designated server is disqualified as a failover site if a local resource fails. Unselect to re-enable automatic failover capabilities.
 - **Disable System Failover** – Select the remote server(s) to be disqualified as a backup server for a complete failure of the local server. When disabled, the designated server is disqualified as a failover site if the local server completely fails. Unselect to re-enable automatic failover capabilities.

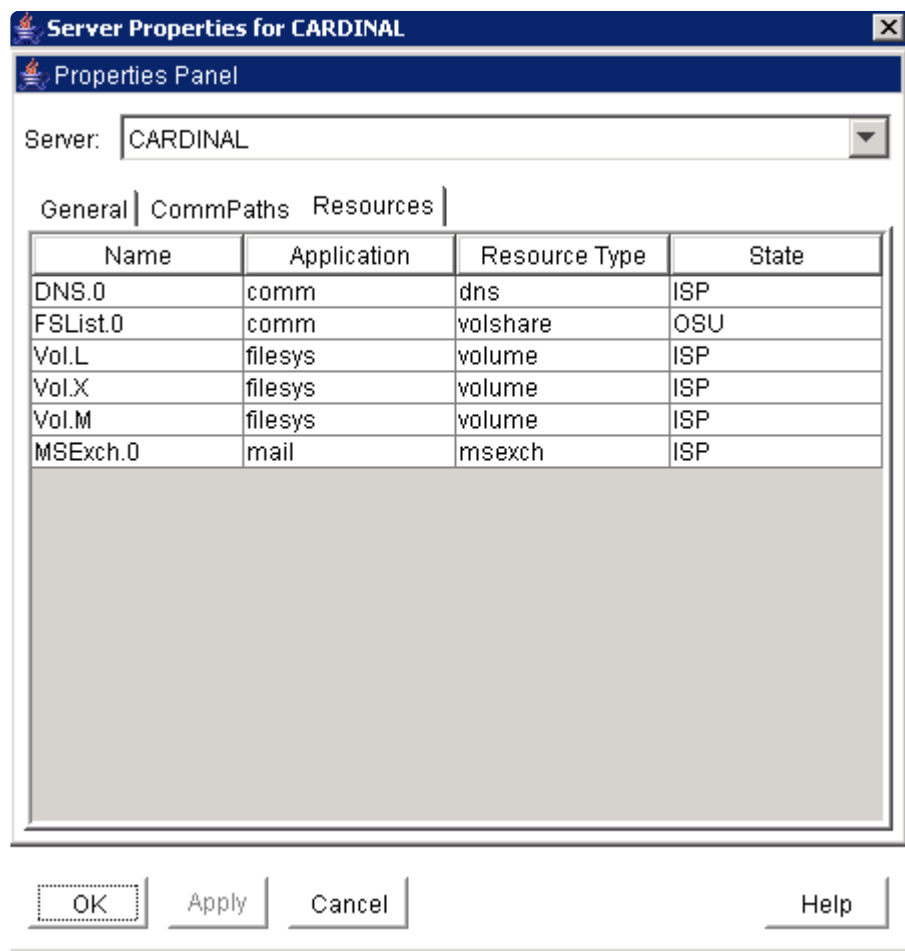
Note: If all remote servers are disabled for resource failovers, then the failed resource will be marked as “**Failed**” and no additional quick check or deep check monitoring will be performed for the failed resource. However, the failed resource as well as other dependent resources in the hierarchy will not be removed from service and no failover will be attempted.



- **Server.** The server name of the other server to which the communication path is connected in the SIOS Protection Suite cluster.
- **Type.** The type of comm path between the server in the list and the server specified in the **Server** field (TCP/IP or Shared Disk).
- **State.** State of the comm path in the LifeKeeper Configuration Database (LCD). These are the possible comm path state values:
 - **ALIVE** – functioning normally
 - **DEAD** – no longer functioning normally
 - **UNKNOWN** – state could not be determined. The GUI server may not be available.
- **Address/Device.** The IP address or device name that this comm path uses.
- **Comm Path Status.** Summary comm path status determined by the GUI based on the state of the comm paths in the LifeKeeper Configuration Database (LCD). These are the possible comm 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, or only one comm path exists.
- **DEGRADED** – one or more redundant comm paths to a given server are dead
- **NONE DEFINED** – no comm paths defined



- **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, volume, TCP/IP, SQL...)
- **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 SIOS Protection Suite startup sequence. Resources in this state are not under SIOS Protection Suite protection.
- **UNKNOWN** – Resource state could not be determined. The GUI server may not be available.

Server Properties – General

The screenshot shows a Windows-style dialog box titled "Server Properties for CARDINAL". It has a "Properties Panel" at the top with a tree view showing "General", "CommPaths", and "Resources", with "General" selected. Below the tree, there is a "Server:" dropdown menu showing "CARDINAL". The main area contains the following information: "State: alive", "Permission: Administrator", and "Shutdown Strategy:" with a dropdown menu showing "Do not Switchover Resources". At the bottom of the main area, there are two checkboxes: "Disable Resource Failover" and "Disable System Failover", both of which are unchecked. The label "BLUEJAY" is positioned to the left of the "Disable Resource Failover" checkbox. At the very bottom of the dialog, there are four buttons: "OK", "Apply", "Cancel", and "Help".

Server Properties for CARDINAL

Properties Panel

Server: CARDINAL

General | CommPaths | Resources

State: alive

Permission: Administrator

Shutdown Strategy: Do not Switchover Resources

BLUEJAY Disable Resource Failover ☐ Disable System Failover ☐

OK Apply Cancel Help

Disabling Automatic Failover

In the event that the primary server has attempted and failed local recovery or failed completely, most server administrators will want SIOS Protection Suite to automatically restore the protected resource(s) to a backup server. This is the default SIOS Protection Suite behavior. However, some administrators may not want the protected resource(s) to automatically go in service at a recovery site; for example, if SIOS Protection Suite 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 SIOS Protection Suite are enabled.

In the **Disable System Failover** column, select the server to be disqualified as a backup server for a complete failure of the local server.

In the **Disable Resource Failover** 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.

 Properties Panel

Server: CARDINAL

General | CommPaths | Resources

State: alive

Permission: Administrator

Shutdown Strategy: Do not Switchover Resources

BLUEJAY Disable Resource Failover Disable System Failover

☐ ☐

OK Apply Cancel Help

Creating a Communication Path

Before configuring a SIOS Protection Suite communication path between servers, verify the hardware and software setup. See the [Configuration](#) section for requirements.

Configuration Notes

- You should configure **no more than one shared disk comm path** between servers.
- Shared Disk comm paths are supported for two-server clusters only.
- For greater than two-server clusters, use multiple TCP/IP comm paths for heartbeat redundancy. A priority value is used to tell SIOS Protection Suite the order in which TCP/IP paths to a given remote server should be used.
- **IMPORTANT:** Supported configurations require that you define redundant comm paths so that the failure of a single communication line will not cause an unnecessary failover. If a single comm path is used and the comm path fails, SIOS Protection Suite hierarchies may come in service on multiple servers simultaneously. This is known as “split-brain”. Additionally, heavy network traffic on a TCP/IP comm path can result in unexpected behavior, including false failovers and SIOS Protection Suite initialization problems.

Creating a Comm Path

1. Select one of the servers, and then select **Create Comm Path** from the [server context menu](#) or [server context toolbar](#).
2. 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 Server**. 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**.
3. Select either *TCP* or *DISK* for **Device Type** and click **Next**.
4. Provide all required information for the **Device Type** that you selected and click **Next** after each step. Refer to the table below for additional information on each configuration field.


Field	Tips
-------	------


For TCP/IP Comm Path...	
Heartbeat Interval	Enter a value between 4 and 15 for the heartbeat interval , which is the number of seconds between heartbeat signals that verifies the servers are alive). The default = 6.
Maximum Heartbeat Misses	Enter a value between 3 and 99. This is the number of consecutive heartbeat signals that can be missed before the comm path is marked as dead. The default = 5.
Local IP Address	Enter the IP address to be used by the local server for this comm path.
Priority	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 the two servers will be used. Priority 1 is the highest, 99 is the lowest.
Remote IP Address	Enter the IP address to be used by the remote server for this comm path.
Port Number	Enter a unique port number to be used by the TCP/IP service. This number must be between 1500 and 10000. SIOS Protection Suite offers a default which you can change.
For Shared Disk Comm Path...	
Heartbeat Interval	Enter a value between 4 and 15 for the heartbeat interval , which is the number of seconds between heartbeat signals that verifies the servers are alive). The default = 6.
Maximum Heartbeat Misses	Enter a value between 3 and 99. This is the number of consecutive heartbeat signals that can be missed before the comm path is marked as dead. The default= 5.
Priority	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 the two servers will be used. Priority 1 is the highest, 99 is the lowest.
Drive Letter	The drive letter associated with the shared volume to be used for the shared disk comm path. This must be the same letter on both servers.

- Click **Create**. The dialog should display a message indicating the network connection is successfully created. If the output panel is enabled, the message will be displayed there as well. Click **Next**.
- If you selected multiple Local IP Addresses or multiple Remote Servers and the **Device Type** was set to **TCP**, then you will be taken back to Step 4 to continue with the next Comm Path. If you selected multiple Remote Servers and the **Device Type** was set to **DISK**, then you will be taken back to Step 3 to continue with the next Comm Path.
- Click **Done** when presented with the concluding message.

Verifying the Comm Path

You can verify the comm path by viewing the [Server Properties](#) dialog. 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. 

If the comm path does not activate after a few minutes, verify that the paired server's computer name is correct.

Deleting a Communication Path

1. Select one of the servers, and then select **Delete Comm Path** from the [server context menu](#) or [server context toolbar](#).
2. Select the communications path(s) that you want to delete and click **Delete Comm Path(s)**.
3. 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.

Working With Resource Hierarchies

The topics in this section describe the tasks that are common across any type of resource hierarchy. These tasks function very much the same regardless of whether you are working with a core Recovery Kit or an optional Recovery Kit.

The documentation for the optional [SIOS Protection Suite Recovery Kits](#) is available in the SIOS Protection Suite for Windows Technical Documentation.

[Creating Resource Hierarchies](#)

[Extending Resource Hierarchies](#)

[Unextending a Hierarchy](#)

[Adding a Resource Dependency](#)

[Removing a Resource Dependency](#)

[Deleting a Hierarchy from All Servers](#)

Creating Resource Hierarchies

1. Select the server, and then select **Create Resource Hierarchy** from the [server context menu](#) or [server context toolbar](#).
2. A dialog entitled **Create Protected Application** 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. Continue through the succeeding dialogs, entering whatever data is needed for the type of resource hierarchy that you are creating.
4. After receiving the message that the resource hierarchy has been created successfully, click **Next** to continue. If SIOS Protection Suite has detected a problem, an ERROR will appear in the information box, the partially created resource hierarchy will be removed and the **Next** button will be disabled. In that case, click **Cancel** to exit the Wizard.

SIOS Protection Suite Application Resource Hierarchies

If you install SIOS Protection Suite without any optional recovery kits, the **Application to Protect** list includes options **DNS**, **File Share List**, **Generic Application**, **IIS**, **IP Address**, **LAN Manager** and **Volume** by default. The **Generic Application** option may be used for applications that have no associated recovery kits.

See the following topics describing these available options:

- [Creating a DNS Resource Hierarchy](#)
- [Creating a File Share Resource Hierarchy](#)
- [Creating a Generic Application Resource Hierarchy](#)
- [Creating an IP Address Resource Hierarchy](#)
- [Creating a LAN Manager Resource Hierarchy](#)
- [Creating a Volume Resource Hierarchy](#)

Microsoft SQL Server Recovery Kit

Installing Microsoft SQL Server Recovery Kit adds entries to the **Application to Protect** list. Refer to the [Microsoft SQL Server Recovery Kit Administration Guide](#) for instructions on creating the required resource hierarchies.

Creating a DNS Resource Hierarchy

The DNS Recovery Kit provides a mechanism to create, update, and manage a DNS A record and associated PTR record for a virtual server name. The DNS resource allows the administrator to configure the managed virtual server name and also the IP address it will be associated with when placed In-Service on each specific LifeKeeper cluster node.

The example charts below show the changes that occur in DNS for a managed server name (SQLSERVER) when In-Service on primary and backup cluster nodes. In this example, the Primary node's public IP address is 172.17.10.24, and the Backup node's public IP address is 172.16.10.25. The managed virtual server name SQLSERVER will have its IP address changed based on which LifeKeeper node it is in service on. Clients use the FQDN "SQLSERVER.mydomain.com" to connect to the application that is associated with this virtual server name.

DNS Server Zone: **mydomain.com**

In-Service on Primary node (before Switchover or Failover):

A Record	SQLSERVER	172.17.10.24
PTR Record	24.10.17.172.in-addr.arpa	SQLSERVER.mydomain.com

In-Service on Backup node (after Switchover or Failover):

A Record	SQLSERVER	172.16.10.25
PTR Record	25.10.16.172.in-addr.arpa	SQLSERVER.mydomain.com

DNS resource configuration requires that all cluster nodes are members of a single domain and that domain includes at least one DNS server that is accessible to each LifeKeeper cluster node. During DNS resource creation and extension, one or more targeted DNS server names to be used by each cluster node are requested by the **SIOS Protection Suite Resource Configuration Wizard** as shown below.

To create a DNS resource hierarchy on the primary LifeKeeper server, you should complete the following steps:

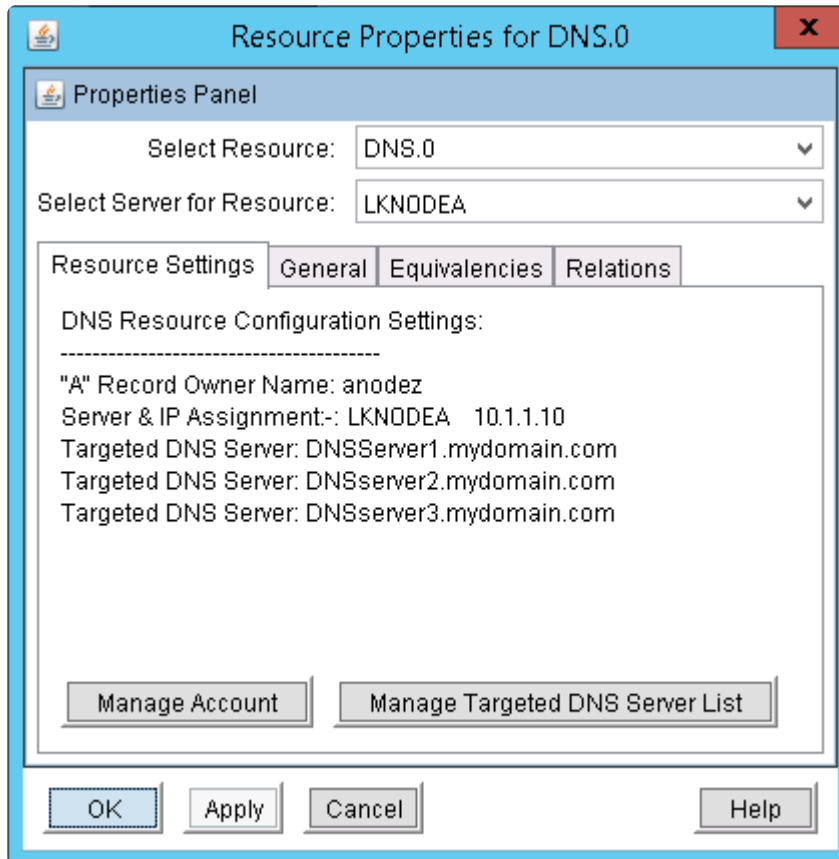
1. Select the server, and then select **Create Resource Hierarchy** from the [server context menu](#) or [server context toolbar](#).
2. The **Create Protected Application** window appears displaying the **Primary** and **Backup** Servers in your cluster. Select the correct systems for this configuration.

3. A dialog appears with a list of all recognized recovery kits installed within the cluster. Select **DNS** and click **Next**.
4. The **Create Protected Application** will prompt you to enter the information in the table below.
5. After receiving the message that the resource hierarchy has been created successfully on the primary node, click **Next** to continue and extend the DNS resource to the backup node.

Field	Tips
Resource DNS Tag	<p>Select or enter a unique tag for the DNS resource instance you are creating. This field provides a default tag name which you can change if desired.</p> <p>Note: The tag name must consist of printable ASCII characters.</p>
A Record Owner Name	Enter the name of the virtual server whose A and PTR records will be updated in DNS. This server name entry is a NetBIOS computer name, NOT a fully qualified server name.
IP Address	Enter the IP address to assign to the virtual server name when the DNS resource is in service on this node. This is usually the LifeKeeper node's public IP address. The A record mapping of this IP address to the virtual server name will be updated upon failover or switchover to this local node.
DNS Server Name (Fully Qualified)	<p>Enter the fully qualified name of one or more targeted DNS servers in the form of <i><DNS Server Name>. <mydomain>.com</i>, where the DNS Resource Records can be modified. If multiple DNS servers are being configured, their names must be space-separated. At least one DNS server must be accessible at all times from the primary or backup LifeKeeper servers when they are In-Service, preferably co-located at the site of each LifeKeeper node. The targeted DNS server lists may be the same or different on each LifeKeeper node. Upon failover or switchover, records on the NS (Name Servers) in the DNS environment will also be updated.</p> <p>You do not need to provide a complete list of all of your DNS servers during creation of the DNS resource. LifeKeeper only needs to connect to one of the servers in the list in order to complete its DNS operations – it will discover the other DNS servers in your configuration when it successfully connects to any DNS server in the list.</p>
DNS Administrative User Name	Enter the user name of the Windows DNS/Domain administrator. This user account should have privileges to make changes in the DNS configuration and should be a member of the "Domain Admins" group in the same domain as the DNS server. Enter the user ID in <i><DomainName>\<UserID></i> format where <i><DomainName></i> is the NetBIOS

	name of the domain.
DNS Administrator Password	Enter the password associated with the Windows DNS/Domain administrator account.

To modify a DNS resource configuration on each server, right click on the DNS resource and select “Properties”. A summary of the current DNS Resource configuration will be displayed as shown below.



Configuration of the DNS resource on each cluster node can be quickly inspected and/or modified by selecting the specific LifeKeeper cluster node in the “Select Server for Resource” drop-down box as shown above.

Resource configuration options include:

- Management of DNS/Domain Account ID and/or Password used to update DNS
- Targeted DNS Server Name List Management (Add and Delete)

The DNS deep check script, which monitors the managed DNS resource, will check for the existence of the

A record of the managed server name on the targeted DNS server (first successful connection) and then on discovered (NS) DNS servers. If the *A record* mapping to the correct IP address is not found on at least one of the DNS servers, the deep check script will fail, which will trigger local recovery (if enabled) and the *A* and *PTR* records will be recreated on the targeted and discovered (NS) DNS servers. If local recovery is not enabled on the DNS resource, or if enabled and not successful, then a failover will occur.

Creating a File Share List Resource Hierarchy

The Windows File Manager function allows you to define file shares. The File Share List Resource type allows you to create a resource that includes one or more of those file shares.

Criteria for File Share List Resources

Not all file shares are available to be shared. The following statements will help you to determine which files shares are available.

- The share name must reside on a volume that is shared between the machines.
- The shared volume can already be protected between the two machines where the file share list resource is being created; however, it should not exist on a third machine until you extend the file share hierarchy to that machine.
- If the share name already exists on the second machine then both share names must point to the exact same directory.
- If the share name is already protected on either machine, it is not eligible.
- It is the responsibility of the administrator to ensure that any share names created actually point to directories. It is possible to create a share name for a directory and then delete the directory. If this is the case then the administrator should ensure that the share name is deleted as well.



Note: After a file share has been brought in-service on a backup server it becomes a share on that machine. The share remains even after the hierarchy is deleted.

File Share List Resource Creation

To create a file share list resource hierarchy, follow the steps below.

1. Select the server, and then select **Create Resource Hierarchy** from the [server context menu](#) or [server context toolbar](#).
2. The **Create Protected Application** window appears displaying the **Primary** and **Backup** Servers in

your cluster. Select the correct systems for this configuration.

3. A dialog appears with a list of all recognized recovery kits installed within the cluster. Select **File Share List** and click **Next**.
4. The **Configuration Wizard** will prompt you 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, SIOS Protection Suite will cancel the entire creation process.

Field	Tips
File Share and Path Name	Select one or more file shares to be protected. If “none found” is displayed, verify that the volume where the file share exists is under SIOS Protection Suite protection.
File Share List Resource Tag	<p>Select or enter a unique tag for the File Share List Resource instance you are creating. This field provides a default tag name FSList.x (where x is a number assigned by SIOS Protection Suite, starting with) which you can change if desired.</p> <p>Note: The tag name must consist of printable ASCII characters.</p>

5. After all of the data is entered, the **Next** button will appear. When you click **Next**, SIOS Protection Suite will create and validate your resource hierarchy.
6. After receiving the message that the resource hierarchy has been created successfully, click **Next** to continue. If SIOS Protection Suite has detected a problem, an ERROR will appear in the information box, the partially created resource hierarchy will be removed, and the **Next** button will be disabled. In that case, click **Cancel** to exit the **Wizard**.

Creating a Generic Application Resource Hierarchy

Use this option to protect an application that has no associated recovery kit.

Before Creating a Resource Hierarchy

1. The first task is to create scripts for the five basic SIOS Protection Suite action functions:

- Restore
- Remove
- Quick Check
- Deep Check
- Local Recovery

Perl and VB Script templates are provided for these scripts in *\$LKROOT\admin\kit\app\templates*. Be sure to copy these templates to another directory on the same volume as *\$LKROOT* before customizing and testing them for the application that you wish to protect.

Note: If you want to use optional **Create**, **Extend** and **Delete** scripts, also include them in the folder with your other scripts. The script selection wizard will search for them by these names (and extension) and automatically enter them for you.

Note: A Restore script and a Remove script are required for a Generic Application solution. There is no time limit for the duration of the Restore script or a Remove script to complete. However, there are implications to duration. A resource In-Service operation can not complete until its restore script has completed successfully. A resource switchover operation can not proceed to the standby server until the resource remove script has completed successfully on the active server.

Note: If provided, Quick Check and Deep Check scripts are run at customer designated intervals. To avoid overlaps, the duration of Quick Check and Deep Check Scripts should not exceed those customer designated interval times. A Local Recovery script is also optional. In the case of a Quick Check or Deep Check failure, a Local Recovery script, if provided, should not extend a Quick Check

or Deep Check duration beyond the customer designated interval.

- For applications depending upon other resources such as a volume or IP address, create each of these resources separately before creating your Generic Application resource hierarchy. You can create the appropriate dependencies later using [Add Dependency](#).

Creating Your Resource Hierarchy

Now you are ready to create the Generic Application resource hierarchy using the modified scripts.

- Select the server, and then select **Create Resource Hierarchy** from the [server context menu](#) or [server context toolbar](#).
- The **Create Protected Application** window appears displaying the **Primary** and **Backup** Servers in your cluster. If not already selected, choose the appropriate systems to configure. Click **Next**.
- A dialog appears with a list of all recognized recovery kits installed within the cluster. Select **Generic Application** and click **Next**.
- The **Configuration Wizard** will prompt you to enter the following information. **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, SIOS Protection Suite will cancel the entire creation process.

Field	Tips
Restore Script	Enter the path and filename for the Restore Script for the application. This is the command that starts the application. A template restore script is provided in the templates directory. The restore script must not impact applications that are already started.
Remove Script	Enter the path and filename for the Remove Script for the application. This is the command that stops the application. A template remove script is provided in the templates directory.
Quick Check Script [optional]	Enter the path to the Quick Check Script for the application. This is the command that monitors the application. A template quickchk script is provided in the templates directory.
Deep Check Script [optional]	Enter the path to the Deep Check Script for the Application. This command monitors the protected application in more detail than the Quick Check Script. A template deepchk script is provided in the templates directory.

Local Recovery Script [optional]	Enter the path to the Local Recovery Script for the application. This is the command that attempts to restore a failed application on the local server. A template recover script is provided in the templates directory.
Application Information [optional]	Enter any Application Information next. This is optional information about the application that may be needed by the restore, remove, recover, and quickCheck scripts.
Resource Tag Name	<p>This field provides a default tag name <i>App.x</i> (where x is a number assigned by SIOS Protection Suite, starting with 0) which you can change if desired.</p> <p>Note: The tag name must consist of printable ASCII characters.</p>

- After all of the data is entered, the **Create Instance** button will appear. When you click **Create Instance**, SIOS Protection Suite will create and validate your resource hierarchy.
- After receiving the message that the resource hierarchy has been created successfully, click **Next** to continue. If SIOS Protection Suite has detected a problem, an ERROR will appear in the information box, the partially created resource hierarchy will be removed, and the **Next** button will be disabled. In that case, click **Cancel** to exit the **Wizard**.

Creating a LAN Manager Resource Hierarchy

The LAN Manager Recovery Kit provides a way to create a computer alias name with associated file shares. The computer alias name acts as a “switchable” computer name, and its associated file shares become available on the system that has the SIOS Protection Suite LAN Manager hierarchy in service. In addition, an IP address can be associated with the computer alias name as part of the hierarchy.

1. Select the server, and then select **Create Resource Hierarchy** from the [server context menu](#) or [server context toolbar](#).
2. The **Create Protected Application** window appears displaying the **Primary** and **Backup** Servers in your cluster. If not already selected, choose the appropriate systems to configure. Click **Next**.
3. A dialog appears with a list of all recognized recovery kits installed within the cluster. Select **LAN Manager** and click **Next**.
4. The **Configuration Wizard** will prompt you 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, SIOS Protection Suite will cancel the entire creation process.

Field	Tips
Computer Alias Name	Enter a name to be used for the computer alias, or you can accept the default name offered by SIOS Protection Suite.
LAN Manager Resource Tag	Select or enter a unique tag for the LAN Manager resource instance you are creating. This field provides a default tag name (the same as the computer alias name entered in the previous step) which you can change if desired. Note: The tag name must consist of printable ASCII characters.

5. After all of the data is entered, the **Next** button will appear. When you click **Next**, SIOS Protection Suite will create and validate your resource hierarchy.
6. After receiving the message that the resource hierarchy has been created successfully, click **Next** to

continue. If SIOS Protection Suite has detected a problem, an ERROR will appear in the information box, the partially created resource hierarchy will be removed, and the **Next** button will be disabled. In that case, click **Cancel** to exit the **Wizard**.

Creating a Volume Resource Hierarchy

When you want to protect resources on shared SCSI disks, you partition the shared disk into logical volumes using the Windows Disk Management tool. SIOS Protection Suite can protect shared volumes by defining a volume resource instance. Each instance is assigned a drive letter (for example, G:).

SIOS Protection Suite brings the volume resource instance into service on the primary server and provides software locks so that a backup server cannot access the volume while it is active on the primary server. In case of a failure of the primary server, SIOS Protection Suite automatically brings the volume resource into service on the backup server and locks the primary server from accessing the volume resource when it is repaired.

SIOS Protection Suite also automatically changes the primary and designations so that the failed server is now locked from access to the volume resource. In this way, the resource is protected from inappropriate access while you repair the failed server.

This dynamic redefinition of primary and backup servers is SIOS Protection Suite's intelligent switchback feature that allows you to select the appropriate time to bring the resource back into service on the repaired system.

To create a volume resource, follow the steps below. Since SIOS Protection Suite maintains the volume locks, do not stop SIOS Protection Suite after creating the resource, as this would disable the locks.



Before creating and extending a mirrored volume resource, be sure to exit from any DataKeeper GUI processes that are connected to any of the SIOS Protection Suite cluster systems.

1. Select the server, and then select **Create Resource Hierarchy** from the [server context menu](#) or [server context toolbar](#).
2. The **Create Protected Application** window appears displaying the **Primary** and **Backup** Servers in your cluster. If not already selected, choose the appropriate systems to configure.
3. A dialog appears with a list of all recognized recovery kits installed within the cluster. 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, SIOS

Protection Suite will cancel the entire creation process.

Field	Tips
Volume	Select the volume to be protected. If “none found” is displayed, verify that the volume is under SIOS Protection Suite protection.
Volume Tag	<p>The Volume tag is a resource identifier. SIOS Protection Suite provides a default volume tag name in the form: Volume.X, where X is the drive letter. You can change the tag name, but it must be unique.</p> <p>Note: The tag name must consist of printable ASCII characters.</p>

- After the data is entered, the **Next** button will appear. When you click **Next**, SIOS Protection Suite will create and validate your resource hierarchy.
- After receiving the message that the resource hierarchy has been created successfully, click **Next** to continue. The **Extend Volume Resource** window displays. Refer to the help topic, [Extending a Volume Resource Hierarchy](#) for additional information while completing this procedure.
- If SIOS Protection Suite has detected a problem, an ERROR will appear in the information box, the partially created resource hierarchy will be removed and the **Next** button will be disabled. In that case, click **Cancel** to exit the **Wizard**.

After the Volume Resource is Created

After a SIOS Protection Suite volume is created or deleted, the following command is executed for the SIOS Protection Suite protected volume:

```
chkntfs /x <vol_1> <vol_2> ... <vol_n>
```

This Windows command excludes the volumes listed from being checked by chkdsk at system startup. This is required for SIOS Protection Suite protected volumes so that they will not be accessed – particularly on backup systems – before LifeKeeper has a chance to start. If no SIOS Protection Suite volumes remain, `chkntfs /d` is executed to restore the Windows default settings.



The `chkntfs /x` command does not remember previous volumes it was applied to, so if a user executes this command, it could disable the SIOS Protection Suite settings, (and

likewise, SIOS Protection Suite could subsequently override the user's settings). If you wish to exclude a non-SIOS Protection Suite volume from checking at startup, you should also include all the SIOS Protection Suite volumes in the `chkntfs /x` command.

Creating an IP Address Resource Hierarchy

SIOS Protection Suite provides the ability to monitor local switchable IP addresses and moves them to another network adapter in the same system when a failure is detected. This can avoid an entire resource hierarchy failing over to a backup server.

IP Local Recovery imposes requirements and limitations on the system configuration.

Requirements for IP Local Recovery

IP local recovery allows you to specify a single backup network interface for each SIOS Protection Suite-protected IP address on a server. In order for the backup interface to work properly, it must be attached to the same physical network as the primary interface. The system administrator is expected to ensure that a valid interface is being chosen. Note that it is reasonable and valid to specify a backup interface on one server but not on another within the cluster (i.e., the chosen backup interface on one server has no impact on the choice of a backup on any other server).

The backup adapter, also known as the Local Recovery Adapter where the switchable address will become active after a failure of the primary adapter, must be configured in the following way:

- Both adapters must be connected to the same physical subnet.
- For routing purposes, all addresses on the Local Recovery Adapter must be on a different logical subnet than any permanent addresses on the Primary adapter. They must also be on a different logical subnet than any SIOS Protection Suite-protected switchable addresses that are configured on the Primary adapter.
- The IP Local Recovery feature requires that a network gateway exist on the network. Specifically, the default gateway field in the TCP/IP configuration for the system must contain the address of a network gateway. In addition, the local recovery adapter must also be configured with the same network gateway.
- IP Local Recovery can only be enabled at the time the IP resource is created. Local Recovery can not be added to an IP resource by modifying its resource attributes after the resource has been created.
- IP Local Recovery may be disabled for an IP resource by using the “ins_setlocalrecovery” command line utility. This utility is located in the LifeKeeper \bin directory (C:\LK\bin by default). From a

command prompt, type `ins_setlocalrecovery` for the usage and switch options.

Before you create and use IP Address resources in SIOS Protection Suite hierarchies, your network should be configured and tested as described in the [Verifying Network Configuration](#) topic.

Also verify that the switchable IP address you plan to use is unique using the ping command. The switchable IP address does not need to be created as a prerequisite; it is created when you create the IP address hierarchy.


To create an IP address resource hierarchy from the primary server, you should complete the following steps:

1. Select the server, and then select **Create Resource Hierarchy** from the [server context menu](#) or [server context toolbar](#).
2. A dialog entitled **Application to Protect** will appear with a list of all recognized recovery kits installed within the cluster. Select **IP Address** and click **Next**.
3. The **Wizard** will prompt you 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, SIOS Protection Suite will cancel the entire creation process.

Field	Tips
IP Address	This is the switchable IP address that SIOS Protection Suite will use for this resource. This is used by client applications to login to the parent application over a specific network interface.
Subnet Mask	The IP subnet mask which 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. Note: 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.
IP Resource Tag	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 that matches the resource name or IP address. You can change this tag if you want to. Note: The tag name must consist of printable ASCII characters.
Network Connection	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 TCP/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 in previous dialogs.
Local Recovery Network Connection	If you answered “ Yes ” to Local Recovery, you must select a network connection to use as the backup interface. Specify the backup NIC if one exists; otherwise, specify the primary NIC.

- After all of the data is entered, click **Next** and SIOS Protection Suite will create and validate your resource hierarchy.
- After receiving the message that the resource hierarchy has been created successfully, click **Next** to continue. If SIOS Protection Suite has detected a problem, an ERROR will appear in the information box, the partially created resource hierarchy will be removed, and the **Next** button will be disabled. In that case, click **Cancel** to exit the **Wizard**.

 **Note:** If using teaming software or if network cards are changed after creating a switchable IP resource, the switchable IP resource should be deleted and recreated as the associated index number for the card can change.

IP Local Recovery Scenario


When IP Local Recovery is enabled and the IP resource fails its quick check or deep check tests, then SIOS Protection Suite will do the following:

- First, SIOS Protection Suite will attempt to bring the IP address back in service on the current network interface.
- If that fails, SIOS Protection Suite will check the resource instance to determine if there is a backup interface available. If so, it will then attempt to move the IP address to the backup interface.
- If all local recovery attempts fail, SIOS Protection Suite will perform a failover of the IP address and all dependent resources to a backup server.

Even if you do not have a backup adapter, you can enable Local Recovery so that SIOS Protection Suite will retry the primary network interface again before initiating failover to a backup server.

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.

Incomplete Resource Priority Modification

A hierarchy in SIOS Protection Suite 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, SIOS Protection Suite 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 SIOS Protection Suite 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 SIOS Protection Suite. 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 SIOS Protection Suite 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.


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 Properties
3. Once you have made changes, the **Apply** button will be enabled. Clicking this button will apply your changes.
4. When you are finished, click **OK** to save any changes and close the window or **Cancel** to close the window without applying changes.


Extending Resource Hierarchies

The SIOS Protection Suite Extend Resource Hierarchy option copies an existing hierarchy from one server and creates a similar hierarchy on another SIOS Protection Suite 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 SIOS Protection Suite communication 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.

 **Note:** When you create a new resource hierarchy, you will be prompted to extend that hierarchy immediately afterwards.

1. To extend an existing resource hierarchy, select that server hierarchy you want to extend and then select **Extend Resource Hierarchy** from the [resource context menu](#) or [resource context toolbar](#).
2. Select the **Backup Server** and click **Next**.
3. A dialog will then display the results of SIOS Protection Suite's pre-extend checks. If these tests succeed, SIOS Protection Suite will display a message stating that the pre-extend scripts were successful. Click **Next** to enter any remaining data needed for the specific type of resource hierarchy that you are extending.

 **Note:** ALL roots in a multi-root hierarchy must be extended together, that is, they may not be extended as single root hierarchies.

Extending a DNS Resource Hierarchy

This operation can be started automatically after you have finished [creating a DNS resource hierarchy](#), or from an existing DNS resource hierarchy, as described in the section on [extending resource hierarchies](#). The following additional data is required to extend a DNS resource hierarchy.

Field	Tips
IP Address	Enter the IP address of the A record associated with the protected primary server or alias name. The record will be updated with this IP address when the DNS resource is brought in-service on this server.
DNS Server Name (Fully Qualified)	Enter fully qualified name of a DNS server, in format <i><DNS ServerName>.<mydomain>.com</i> , where the Resource Records can be modified. The DNS server should be accessible from the backup server, preferably in the same site.

Extending a File Share Resource Hierarchy

This operation can be started automatically after you have finished [creating a file share resource hierarchy](#), or from an existing file share resource hierarchy, as described in the section on [extending resource hierarchies](#). No additional data is required to extend a file share resource hierarchy.

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 hierarchy, as described in the section on [extending resource hierarchies](#). No additional data is required to extend a generic application resource hierarchy.

Extending a LAN Manager Resource Hierarchy

This operation can be started automatically after you have finished [creating a LAN manager resource hierarchy](#), or from an existing LAN manager resource hierarchy, as described in the section on [extending resource hierarchies](#). No additional data is required to extend a LAN manager resource hierarchy.

Extending a Volume Resource Hierarchy

This operation can be started automatically after you have finished [creating a volume resource hierarchy](#), or from an existing volume resource hierarchy, as described in the section on [extending resource hierarchies](#). The following additional data is required to extend a volume resource hierarchy.

Field	Tips
Volume Type	Select Shared Disk if using shared storage, Create Mirror if using SIOS DataKeeper and the mirror does not exist, or Existing Mirror if using SIOS DataKeeper and the mirror has already been created.
Network end points (Target/Source)	If Volume Type Create Mirror or Existing Mirror , select the network end points for the mirror. End points must be IP addresses.
Mode	<p>If Volume Type Create Mirror, then select the mode of the mirror.</p> <p>Asynchronous Mirror: Source writes are queued for transmission to the target, and return immediately. Less reliable than synchronous, but source writes are quicker.</p> <p>Synchronous Mirror: All writes to the source volume will be committed to the target volume immediately. Higher reliability, lower performance.</p>
When extending the volume resource to a third system in the cluster, you must specify the volume type for each of the equivalent systems in the cluster.	
Volume Type (Shared or SIOS DataKeeper)	Select Shared Disk or the network end points for the mirror between the equivalent systems.

Note: Mirrors created from the LifeKeeper GUI will be deleted when the volume resource hierarchy is deleted. To prevent the mirror deletion, set the [LifeKeeper Delete Mirror Flag](#) to **False**.

Extending an IP Address Resource Hierarchy

This operation can be started automatically after you have finished [creating an IP address resource hierarchy](#) or from an existing IP address resource hierarchy as described in the section on [extending resource hierarchies](#). The following additional data is required to extend an IP address resource hierarchy.

Field	Tips
Subnet Mask	Enter the subnet mask to use for the IP resource on the target server. SIOS Protection Suite will, by default, offer the subnet mask used on the template server.
Network Connection	Select the network connection to use on the target server.
Target Restore Mode	<p>This feature applies to three-node SIOS Protection Suite clusters where two nodes are on a LAN (same subnet) and the third node is on a WAN (different subnet). The restore mode of the IP resource would be enabled on the LAN nodes and disabled on the WAN node.</p> <p>Select the appropriate <i>Restore Mode</i> for this IP resource on the target system. In some situations a protected IP address should not be used on a remote target system. For example, the remote target system may be connected to a different subnet than other systems in the cluster. In this situation the IP resource may be extended using the “Disable” Restore Mode. When using the “Disable” Restore Mode option, SIOS Protection Suite will not configure the IP address on the target system when the resource is placed in-service there and monitoring for the IP resource will be disabled. In these situations, network redirection may be implemented some other way or by using a SIOS Protection Suite DNS resource. You may use the IP resource properties page on the target system to change your selection at a later time. See Managing IP Resources.</p>
Target Local Recovery	Click Yes if you wish to enable IP Local Recovery on the target server; otherwise choose No .
Target Local Recovery Network Connection	If you answered Yes to Local Recovery , you must select a network connection to use as the backup interface. Specify the backup NIC if one exists; otherwise, specify the primary NIC.

Unextending a Hierarchy

The **Unextend Resource Hierarchy** option removes a complete hierarchy, including all of its resources, from a single server. This is different than the [Delete Resource Hierarchy](#) selection which removes a hierarchy from all servers.

When using **Unextend Resource Hierarchy**, the server from which the existing hierarchy is to be removed is referred to as the target server.

The **Unextend Resource Hierarchy** selection can be used from any SIOS Protection Suite server that has active SIOS Protection Suite communication paths to the target server.

1. Select a server-specific resource instance from the hierarchy that you want to unextend, and then select **Unextend Resource Hierarchy** from the [resource context menu](#) or [resource context toolbar](#).
2. The dialog will display a message verifying the server and resource hierarchy that you have specified to be unextended. Click **Unextend** to perform the action.
3. If the [output panel](#) is enabled, the dialog closes, and the results of the commands to unextend the resource 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.

Adding 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. Select a server-specific resource instance as the parent to which you want to add a child dependency, and then select **Add Dependency...** from the [resource context menu](#) or [resource context toolbar](#).
2. Select a **Parent Resource IP Address** from the drop down box. Click **Next**.
3. 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.

4. The dialog will then confirm that you have selected the appropriate parent and child resource tags for your dependency creation. Click **Add Dependency** to create the dependency on all servers in the cluster to which the parent has been extended.
5. 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.

Removing a Resource Dependency

1. Select a server-specific resource instance as the parent from which you want to remove a child dependency, and then select **Remove Dependency** from the [resource context menu](#) or [resource context toolbar](#).
2. Select the **Child Resource** from the drop down box. This should be the 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 **Remove 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.

Deleting a Hierarchy from All Servers

1. Select a server-specific resource instance in the hierarchy that you want to delete, and then select **Delete Resource Hierarchy** from the [resource context menu](#) or [resource context toolbar](#).
2. The dialog will display a message verifying the hierarchy you have specified for deletion. Click **Delete** to perform the action. Deletion will begin on the sever that you initially selected.
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.

Man Pages

[LCD – Miscellaneous LCD Programs](#)

[LCDI Applications](#)

[LCDI Instances](#)

[LCDI-relationship](#)

[LCDI-resource_type](#)

[LCDI-systems](#)

[LifeKeeper Flags](#)

[lk_chg_value](#)

[lk_err](#)

[perform_action](#)

[sendevent](#)

[volume](#)

LCD – Miscellaneous LCD Programs

Synopsis

[lcdremexec](#) [-e] -d destname — cmd [arg1 arg2 ... argn]

[lcdsync](#) [-d destname]

[lcdrecover](#) -g {remote|restore|delete} — [arg1 arg2 ... argn] | -G {remote|restore|delete} — [arg1 arg2 ... argn] | -p primarytest /| [-o resource]

[lcdrcp](#) file1 file2 file3 ... {dest:ofile | dest:odir}

[lkstart](#) [-w waitperiod]

[lkstop](#) [-f or -r|-n]]

Description

These programs have various uses by application developers. They are all found in the directory %LKROOT%\bin.

Exit Codes

The following exit codes could be returned by these commands:

0	The operation has succeeded.
1	A system call or library call has internally returned failure.
2	A user-specified syntax error occurred.
3	LifeKeeper internal error.
4	A request to perform an operation on an object that already exists.
5	An argument specified is illegal.
6	Index out-of-range.
7	A request has been made on an object that does not exist.

8	A request was made to delete a resource instance on which another non-deleted resource instance depends.
9	An attempt to communicate with another system failed.

lcdrcp

```
lcdrcp file1 file2 file3 ... {dest:ofile | dest:odir}
```

lcdrcp is a general purpose program used to transfer the ASCII files *file1 file2 file3 ...* to another system using the LifeKeeper communications path. Binary files cannot be copied using **lcdrcp**.

LifeKeeper transfers the files to dest in the directory *odir*. If only one file is sent, the alternate form including the destination file name at location *ofile* on system dest is provided. Take extra caution while using Windows drive names (like D:), as destination arguments as they could be misinterpreted as destination names if a destination name is missing. However, if a destination system name is specified, drive names are interpreted properly.

lcdrecover

```
lcdrecover -g {remove|restore|delete} - [arg1 arg2 ... argn] | -G
{remote|restore|delete} - [arg1 arg2 ... argn] | -p primarytest | [-o resource]
```

The **-g** option takes one of three arguments, *remove*, *restore*, or *delete*. This option will run the preglobal scripts for the specified argument. The preglobal scripts are registered by applications to run before certain events. For example, with the *restore* argument, this option runs the prerestore script registered by LifeKeeper, then any prerestore scripts registered by all of the applications. Normally, **perform_action** [see [LRACI-perform_action](#)] automatically performs the prerestore scripts, except when the **-G** option is specified to **perform_action**.

The **-G** option of **perform_action** allows multiple **perform_action** commands to be run, with the preglobal scripts running only once before the first **perform_action** execution using **lcdrecover -g restore**. An application may register a preglobal script by installing the script at the path:

```
%LKROOT%\ subsys\ <appname>\actions\prerestore.ksh
```

arg1, *arg2*, ... *argn* are arguments that will be passed to the preglobal scripts when they are executed.

Similar scripts (preremove) exist for the *remove* argument. They can be run before a **perform_action -G -a remove** is run. They are run when **lcdrecover -g remove** is executed.

The predelete scripts are similar, but they are run before the **ins_remove -G ...** [see [LCDI-instance](#)] program is run, unless **-G** for **ins_remove** is left out.

The **-G** option for **lcdrecover** is analogous to **-g**, except that it specifies that the postglobal scripts should be run. The **-G** option should not be used without running an earlier **lcdrecover -g arg**, and it should be run after all of the **perform_action** or **ins_remove** programs are run. If you are executing the **-G** option within a **getlocks** protected region (after **getlocks** and before **rlslocks**), set *arg1* to *-m* to avoid executing a second instance of **getlocks**, which would cause the operation to hang.

The following example runs multiple **perform_action** commands where the preglobal and postglobal scripts run only once:

```
lcdrecover -g restore

# run "preglobal" restore scripts
```

```
perform_action -G -a restore -t tagname

# neither scripts are run

perform_action -G -a restore -t tagname2

# neither scripts are run

lcdrecover -G restore - -m

# run "postglobal" restore scripts

# use -m arg when in getlocks protected region of code
```

This example runs multiple prerestore and postrestore scripts:

```
perform_action -a restore -t tagname

# all scripts once

perform_action -a restore -t tagname2

# all scripts again
```

The **-p** option for **lcdrecover** is used to determine if a particular resource is on a resource hierarchy that is on the primary system or the secondary system. Specify the resource tag name with primary test, and it will print out to standard output the string primary if the resource is on the primary hierarchy, or secondary if it is not.

The **-o** option can be used to retrieve the remote system associated with the resource tag specified.

lcdremexec

```
lcdremexec [-e] -d destname - cmd [arg1 arg2 arg3 ... argn]
```

This program sends a remote request over the LifeKeeper communication paths to the system *destname*, to execute the command **cmd** remotely with arguments *arg1 arg2 arg3 ... of the* and returns the standard output and standard error of the remote command to standard output of the **lcdremexec** command. The exit code of the remote command is returned by **lcdremexec**.

Note: If *destname* is the current system, no messages are sent; **lcdremexec** will execute it locally.

The **-e** option will split standard output and standard error of the remote command and first print standard output of the remote command to standard output of **lcdremexec**, then print standard error of the remote command to standard error of the **lcdremexec** command. This option has no effect for local commands, which have their standard output and standard error unchanged.

cmd can be either a Korn shell script or a Win32 executable. It will be executed with %LKROOT% on *destname* as the current working directory, thus being able to accept path names relative to %LKROOT%.

Before executing, the directory %LKROOT%\BIN is always added to the head of the PATH variable on *destname*. If *destname* is DEAD or goes DEAD in the middle of the execution, **lcdremexec** returns a non-zero exit code.

lcdsync

```
lcdsync [-d destname]
```

This program checks to see if the LifeKeeper resource hierarchy configuration and communication path status data stored in shared memory has been modified. If it is different, the data is “synchronously” written to disk. Therefore, when this program returns, the data is guaranteed to be on disk properly. If *destname* is not specified, the current system is assumed.

Note: The commands used to modify resource hierarchy configurations or communication paths (such as **ins_create**, **dep_create**, **ins_setit**, **eqv_remove**,...) only modify the shared memory segment and are not reflected in the permanent file storage of LifeKeeper, until the **lcdsync** program is run.

lkstart

```
lkstart [-w waitperiod]
```

This program starts up LifeKeeper on the current system if it is currently not running. **lkstart** modifies entries in the `%LKROOT%\etc\LKinit.config` file pertaining to the LifeKeeper daemons so that they will be respawned if they die.

The `-w` option, with *waitperiod* in seconds, can be used to change the timeout interval. Use the `-w` argument to specify a wait period before the startup.

The LifeKeeper service can be started using the Services mmc under Administrative Tools, or from a command prompt using either `"sc start LifeKeeper"` or `"net start LifeKeeper"`.

Note: This program must be run from the console.

lkstop

```
lkstop [-n] [-f] [-r]
```

This script shuts down LifeKeeper on the system, if it is currently running. LifeKeeper will automatically restart at system boot.

The table below describes the actions taken by LifeKeeper when each **lkstop** option is entered:

Command Line	Action
lkstop	Resources in service are removed from service and are NOT switched over to a backup server.
lkstop -n	Same as lkstop with no options specified.
lkstop -f	Resources in service do not get removed from service.
lkstop -r	Same as -f.

The LifeKeeper services can also be stopped using the Services tool under Administrative Tasks in the Windows Control Panel.

LCDI Applications

Synopsis

[`app_create`](#) [-d destsys] -a appname

[`app_remove`](#) [-d destsys] -a appname

[`app_list`](#) [-d destsys]

Description

A LifeKeeper application is a group of related resource types. When an application is removed, all resource types installed under it are also removed.

These programs provide an interface for generating new applications in the configuration database and removing existing ones. All commands exit to 0 if they are successful. Commands exit with a nonzero code and print to standard error if they fail.

Exit Codes

The following exit codes could be returned by these commands:

0	The operation has succeeded.
1	A system call or library call has internally returned failure.
2	A user-specified syntax error occurred.
3	LifeKeeper internal error.
4	A request to perform an operation on an object that already exists.
5	An argument specified is illegal.
6	Index out-of-range.
7	A request has been made on an object that does not exist.
8	A request was made to delete a resource instance on which another non-deleted resource instance depends.
9	An attempt to communicate with another system failed.

app_create

```
app_create [-d destsys] -a appname
```

Creates a new application. The application is tied to system *destsys*, using the name called *appname*. If *destsys* is not supplied, the application is created locally.

app_list

```
app_list [-d destsys]
```

This command prints to standard output a list of applications that have installed software to work with LifeKeeper on the system *destsys*. If *destsys* is not specified, the current system is assumed.

app_remove

```
app_remove [-d destsys] -a appname
```

Removes the given application from the configuration database set of known applications. All resource types, resource instances, and equivalencies relating to this application are also removed. Failure can occur because the application is not known to the configuration database.

LCDI Instances

Synopsis

- [ins_gettag](#) [-d destsys] -i id
- [ins_create](#) [-d destsys] -a appname -r restyp [-l{AUTORES_ISP|INIT_ISP|INIT_OSU}] [-v info] -t tag -i id [-Q quickChkInt] [-DdeepChkInt] [-l localRecover{Y/N}] [-s AUTOMATIC/INTELLIGENT]
- [ins_remove](#) [-d destsys] [-R roottag] [-a appname] [-r restyp] [-ttag] [-i id] [-v] [-l] [-N] [-G]
- [ins_setin](#) [-d destsys] -t tag [-v info]
- [ins_setit](#) [-d destsys] -t tag -l {AUTORES_ISP|INIT_ISP|INIT_OSU}
- [ins_setst](#) [-d destsys] -t tag -S {ISP|ISU|OSU} [-R reason] [-A]
- [ins_list](#) [-d destsys] [-fC] [-R top] [-a appname] [-r typ] [-t tag] [-i id]
- [ins_setchkint](#) [-d destsys] -t tag -c {q=quick|d=deep} -vinterval
- [ins_setlocalrecover](#) [-d destsys] -t tag -l {Y=enable|N=disable}
- [ins_setas](#) [-d destsys] -t tag -s {INTELLIGENT|AUTOMATIC}

Description

Resources are used by LifeKeeper to represent volumes, applications, or system objects known by the system. Resource types are classifications of resources; resource instances are actual instances of a resource type. For example, resource types would include file system volumes, file shares, IP addresses, LAN Manager names and various servers like SQLServer. Generic, user-definable types permit users to build custom fault resilient setups. Multiple instances may exist for a resource type.

Resource instances may exist in a number of states. These states may take on the following values and meanings:

ISP	Resource is in service, protected. ISP is the normal state of resources on the primary node.
-----	--

OSU	Out of service, unimpaired. The resource is not available on this system because it was brought out of service by executing its remove script. The OSU state also is used for objects that have dependencies on children in the OSF or OSU state or when the equivalent object on the backup machine is in the ISP or ISU state. OSU is the normal state of resources on the secondary node.
OSF	Out of service due to a failure. The resource is not available on this system because a failure has occurred trying to restore the object.

Exit Codes

All commands exit to 0 if they are successful. Commands exit with a nonzero code and print to standard error if they fail. The following exit codes could be returned by these commands:

0	The operation has succeeded.
1	A system call or library call has internally returned failure.
2	A user-specified syntax error occurred.
3	LifeKeeper internal error.
4	A request to perform an operation on an object that already exists.
5	An argument specified is illegal.
6	Index out-of-range.
7	A request has been made on an object that does not exist.
8	A request was made to delete a resource instance on which another non-deleted resource instance depends.
9	An attempt to communicate with another system failed.

ins_list

ins_list [-d destsys] [-fC] [-R top] [-a appname] [-r typ][-t tag] [-i id]

This command prints lines relating to a set of resource instances to standard output. Each line contains all of the current information known about a particular resource instance. Examples of the lines are:

```
LKSYS1-filesys-volume-ISSUTILS-K:--ISP-restore action
hassucceeded-AUTORES_ISP--0-0-
```

Each line contains fields separated by a delimiter character. The default delimiter character is ^A (\001). If the -fC option is specified, the delimiter is changed to the specified character. The above example shows a dash (-) as a delimiter. The fields in the example command above are as follows:

LKSYS1	Name of the system the resource the instance resides on.
filesys	Application name of resource type.
volume	Resource type name.
ISSUTILS	User-defined resource instance tag identifier.
K:	LifeKeeper internal identifier for resource instance.
- -	If this field is not empty, as in the example, it provides additional instance information (type dependent).
ISP	Current state of resource instance ISP, ISU, OSU, or OSF.
restore action has succeeded	Reason for last state change.
AUTORES_ISP	Available resource initialization options are: AUTORES_ISP, INIT_ISP, and INIT_OSU.
- -	<div>If this field is not empty, as in the example, it indicates that the resource is currently being reserved for:</div> <div>RESTORE: restoring the resource to service</div> <div>REMOVE: removing the resource from service</div> <div>RECOVER: performing local recovery on resource</div>

0	Process ID of process that has reserved resource.
0	Reserved.
180	Quick check interval, in seconds.
300	Deep Check interval, in seconds.
0	Local recovery option. 0 = disabled; 1 = enable

The other arguments limit the number of resource instances included in the list. If none of the arguments are used, then all resources on *destsys* are listed. These are the limiting arguments:

destsys. If *destsys* is not specified, the current system is assumed; otherwise, data from the remote system is printed.

top. If *top* is the space string “ ”, only the root resources will be printed. If *top* is specified (but not the space string), the report lists the top resource and all children resources below it, recursively.

appname. If *appname* is specified, all resource instances associated with all resource types defined by this application are printed. If *appname* is not specified, all resource instances for all applications defined on the system are printed.

typ. If *typ* is specified, all resource instances of type, *typ*, in application *appname* are printed.

tag or *id*. If *tag* or *id* is specified, the resource instance associated with that *tag* or *id* is printed.

Initialization Strategy

It is recommended that you accept the default Auto ISP Initialization Strategy. These are the actions taken when LifeKeeper starts (initializes):

Autores ISP	Resource is automatically brought into service if it is not in service on the paired node.
Init ISP	Resource is always initialized into the ISP state.
Init OSU	Resource is always initialized into the OSU state.

Initial State

The state is the current processing status for the resource. For example, the normal state for a resource on the primary system is **ISP** – in service, protected. The normal state for a resource on the secondary system is **OSU** – out of service, unimpaired.

It is recommended that you accept the default initial state. If you set the Initial State to **OSU**, you must manually bring the resource into service.

ins_create

```
ins_create [-d destsys] -a appname -r restyp [-I
{AUTORES_ISP|INIT_ISP|INIT_OSU}][-v info] -t tag -i id [-Q quickChkInt][-D
deepChkInt][-l localRecover{Y|N}] [-s AUTOMATIC|INTELLIGENT]
```

Defines a new resource instance on system *destsys* in the configuration database. The resource instance is described by the arguments given. If *destsys* is not specified, the current system is assumed. The command offers the following string tag options:

- The *-a* and *-r* options indicate the preexisting application and resource type associated with this new instance.
- Initialization type field specified by the *-I* option indicates how the resource instance should be initialized if LifeKeeper restarts (for example, at boot time).
- Optional string info specified by the *-v* option is a field that can contain additional resource type specific information and does not necessarily have to be unique per resource type.
- String tag specified by the *-t* option is a string that names the resource instance and is unique on a system. It is a string that is meaningful externally to LifeKeeper. String tag specified by the *-t* option is a string that names the resource instance and is unique on a system. It is a string that is meaningful externally to LifeKeeper. **Note:** The tag name must consist of printable ASCII characters.
- String id specified by the *-i* option is also unique per system, but may be meaningful only internally to LifeKeeper.
- Quick check interval provided with *-Q* option should be in seconds. The value should be zero if **quickchk.ksh** script doesn't exist for the resource. LifeKeeper waits this interval time between two consecutive execution of **quickchk.ksh** script. Valid range of value: 0 – 604800.
- Deep check interval provided with *-D* option should be in seconds. The value should be zero if **deepchk.ksh** script doesn't exist for the resource. LifeKeeper wait this interval time between two consecutive execution of **deepchk.ksh** script. Valid range of value: 0 – 604800.
- Local recover option indicates whether resource should be recovered by executing **recover.ksh** script. This option should be “N” if **recover.ksh** script doesn't exist for the resource.

ins_gettag

```
ins_gettag [-d destsys] -i id
```

Prints to standard output the tag name that corresponds to the internal identifier provided in *id* on the system with name *destsys*. If *destsys* is not specified, the current system is assumed.

Note: The tag name and *id* name for a resource are unique on a system, but may be reused to indicate different resource instances on different systems.

The resource tag provides an understandable handle (human readable name) for a resource instance, for example, *user-partition*, whereas the *id* is an internal descriptor. The resource name *id* is used by the application software associated with the resource to uniquely describe the resource.

ins_remove

```
ins_remove [-d destsys] [-R roottag] [-a appname] [-r restyp] [-t tag] [-i id] [-v] [-I] [-N] [-G]
```

Removes resource instance(s) on system *destsys* from the configuration database. Associated dependencies and equivalencies will also be removed. If *destsys* is not specified, the current system is assumed.

Note: All resources that depend upon any of the selected resources directly or indirectly will also be removed before the selected resource is removed.

When an [ins_setchkint](#) resource instance is removed, and if a delete action was defined for the resource type of the instance being removed, the delete action is run before the instance is removed.

The command has the following options:

-R	The <i>-R</i> option is for removing entire sub-hierarchies and the resources that depend on them. The <i>roottag</i> string defines a list of instance tag names (separated by the ^A character) for which these resources and the resources below on the hierarchy will be recursively removed, until a resource is encountered for which a resource not being removed depends.
-a	If the <i>-a</i> option is specified, only resources from that application will be removed.
-r	If the <i>-r</i> option is specified, all resources of the specified resource type will be removed.
-t or -i	If the <i>-t</i> option or <i>-i</i> option is specified, the instance with the matching <i>tag</i> or <i>id</i> will be removed along with the resources that depend on them.
-v	If the <i>-v</i> option (verbose) is specified, the function prints a message to standard output including the tag names of all the removed resource instances.
-I	The <i>-I</i> option initializes the resource hierarchy so that ins_remove can work properly. This option should be used by the “highest-level” recursive call to ins_remove , but not necessarily in a lower-level recursive call such as inside a delete script. The <i>-I</i> option should NOT be used by a recursive call to ins_remove from inside a delete script.
-N	The <i>-N</i> option tells ins_remove NOT to reinitialize the resource hierarchy. The assumption when using this option is that a higher level call of ins_remove will perform the <i>-I</i> option. The <i>-N</i> option MUST be used inside a delete script, since a delete script is being called from a parent invocation of ins_remove and the <i>-N</i> option prevents hierarchy cycles from occurring.

-G	<p>The -G option indicates that the predelete and postdelete scripts [see LCD] should not be performed as part of this call to ins_remove. This option is useful if you wish to perform multiple top-level calls to ins_remove and have the predelete and postdelete run manually [using lcdrecover-g delete of LCD] before and after (respectively) the calls to ins_remove. It would also be wise to use the -G option in the delete scripts since the highest-level ins_remove should perform the predelete and postdelete scripts.</p>
----	---

ins_setas

```
ins_setas [-d destsys] -t tag -s {INTELLIGENT|AUTOMATIC}
```

Sets the switchback type of a root resource on system *destsys* with *tag* name *tag* to the strategy specified in the -s option. Use only on root resource to change the switchback type of the root resource and all its dependent resources.

ins_setchkint

```
ins_setchkint [-d destsys] -t tag -c {q=quick|d=deep} -v interval
```

This command modifies the quick check or deep check interval for the resource specified by the *tag* name “-t”. The interval must be entered in seconds.

Examples:

To change the quick check interval for the file share resource FSList.0 to two minutes, run the following command from \$LKROOT\bin:

```
ins_setchkint -t FSList.0 -c quick -v 120
```

To disable the deep check for the file share resource FSList.0, run the following command from \$LKROOT\bin:

```
ins_setchkint -t FSList.0 -c deep -v 0
```

ins_setin

```
ins_setin [-d destsys] -t tag [-v info]
```

The string *info* specified by the *-v* option is a field that can contain additional resource type specific information and does not necessarily have to be unique per resource type. For example, an instance of a resource of file share type will have the names of all shares managed by this instance in its *info* value.

ins_setit

```
ins_setit [-d destsys] -t tag -I {AUTORES_ISP | INIT_ISP | INIT_OSU}
```

Indicates to LifeKeeper how it should initialize the state of a resource when LifeKeeper itself initializes (for example, at system boot time). If you do not set this option, LifeKeeper sets the initialization state to default options.

These are the restore options you can specify:

AUTORES_ISP. If resource initialization is set to *AUTORES_ISP*, the resource is first set to the OSU state, then the restore action is performed and, if successful, the resource is put into the ISP state. If restore fails, the resource is placed into the OSF state.

INIT_ISP. If *INIT_ISP* is set, LifeKeeper assumes resource initialization by other means and places the resource into the ISP state.

INIT_OSU. If *INIT_OSU* is set, LifeKeeper assumes the resource is not started up during initialization and that the system administrator will manually start up the resource using LifeKeeper Graphical User Interface (GUI) application.

ins_setlocalrecover

```
ins_setlocalrecover [-d destsys] -t tag -l {Y=enable|N=disable}
```

This command modifies the local recovery setting for the resource specified by the tag name “-t”.

Example:

To disable local recovery for the file share resource FSList.0, run the following command from
\$LKROOT\bin:

```
ins_setlocalrecover -t FSList.0 -l N
```

ins_setst

```
ins_setst [-d destsys] -t tag -S {ISP|ISU|OSU} [-R reason] [-A]
```

Sets the resource state on system *destsys* with tag name *tag* to the state specified in the *-S* option. If *destsys* is not specified, the current system is assumed. Use this command cautiously because the resource state will be changed by the resource's action script (e.g. remove or restore script). The caller is responsible for making sure the new state reflects the actual state of the application.

Additional text explaining the reason for the change of state may be provided by the *-R* option. The *-A* option sets the state of the specified resource and all the resources that depend on it, recursively up the hierarchy.

LCDI-relationship

Synopsis

[`dep_create`](#) [-d destsys] -p partag -c chdtag

[`dep_remove`](#) [-d destsys] [-p partag] [-c chdtag]

[`dep_list`](#) [-d destsys] [-fC] [-C allchild | -P allparent | -c ofparenttag | -p ofchildtag] [-r typ] [-a app]

[`eqv_create`](#) [-d destsys] [-s sys] -t tag [-p sysPriority] [-Sothersys] -o othertag [-r othersysPriority] -e SHARED

[`eqv_remove`](#) [-d destsys] [-s sys] -t tag [-S othersys] -o othertag -e SHARED

[`eqv_list`](#) [-d destsys] [-s sys] [-t tag] [-fC]

Description

LifeKeeper resources exist in relationship to one another. Two resources may be unrelated or they may be in a dependency relationship. In a hierarchy, a resource may have several resources depending upon it, and it may depend upon several resources. Each resource also relates to a like resource on the paired system with a shared equivalency. This shared equivalency ensures that a resource is active on only one system at a time. Equivalency object also indicates priority of the system for the resource. This priority value determines the order of cascading failover. A higher priority system has precedence over a lower priority system in recovering the resource. A priority value of 1 is the highest. The higher the numerical value the lower the priority. Two systems can't be assigned same priority for a resource. Valid range is 1 to 1024.

Exit Codes

All commands exit to 0 if they are successful. Commands exit with a nonzero code and print to standard error if they fail. The following exit codes could be returned by these commands:

0	The operation has succeeded.
1	A system call or library call has internally returned failure.
2	A user-specified syntax error occurred.

3	LifeKeeper internal error.
4	A request to perform an operation on an object that already exists.
5	An argument specified is illegal.
6	Index out-of-range.
7	A request has been made on an object that does not exist.
8	A request was made to delete a resource instance on which another non-deleted resource instance depends.
9	An attempt to communicate with another system failed.

dep_create

```
dep_create [-d destsys] -p parent -c child
```

This function creates a dependency relationship between the resource instances with tags *parent* and *child*. Both resources must be on the same system *destsys*. If *destsys* 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 and must be in the same state (ISP or OSU) for proper operation.

dep_list

```
dep_list [-d destsys] [-fC] [-C allchild | -P allparent | -c ofparenttag | -p ofchildtag] [-r typ] [-a appname]
```

This function prints strings to standard output describing dependency relationships between resource instances. If *destsys* is not specified, the current system is assumed. Each string is in the following form:

```
LK0-LKSYS:135.66.249.201

LK0-LKSYSA:FSLIST.0

FSLIST0:fi.vo.0
```

There are two fields in each string that are separated by a delimiter character. The default delimiter character is ^A (\001). If the *-fC* option is specified, the delimiter is changed to the character specified. 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.

You can use options to limit the contents of the list. If you use **no** options, all dependencies are printed. The command has the following options:

-C	If the -C option is specified, this command will print out all direct and indirect child dependencies of the resource specified in <i>allchild</i> .
-P	If the -P option is specified, this command will print out all direct and indirect parent dependencies of the resource specified in <i>allparent</i> .
-c	If the -c option is specified, this command will print out only the direct child dependencies of the resource specified in <i>ofparenttag</i> .
-p	If the -p option is specified, this command will print out only the direct parent dependents of the resource specified in <i>ofchildtag</i> .
-r	Specifying the -r option lists all the dependencies of child <i>typ</i> .
-a	Specifying the -a option lists all the dependencies of application <i>appname</i> .

dep_remove

```
dep_remove [-d destsys] [-p parent] [-c child]
```

Removes the dependency relationship(s) from the database on system *destsys*. If *destsys* is not specified, the current system is assumed. If *child* is not specified, all dependencies of *parent* are removed. If *parent* is not specified, all dependents with tag *child* are removed. If both are not specified, all dependencies are removed.

eqv_create

```
eqv_create [-d destsys] [-s sys] -t tag [-p sysPriority] [-S othersys] -o  
othertag [-r othersysPriority] -e SHARED
```

Creates an equivalency in the configuration database on system *destsys* (local, if not specified).

LifeKeeper will automatically add a *SHARED* equivalency on a remote system, if either *sys* or *othersys* is specified. The resource specified as *tag* on system *sys* will be assumed by LifeKeeper to be the "primary" resource that runs under normal conditions; the resource specified as *othertag* on system *othersys* will be the "secondary" resource on the paired system. When LifeKeeper initializes, the primary resource is set depending upon resource initialization set up (see [LCDI_instances](#)). If the spare system boots, LifeKeeper on the spare checks the primary system to see if the primary system is functioning and if the primary resource is in the ISP state. If both cases are true, LifeKeeper puts the secondary resource into the OSU state (resource initialization ignored). If either case is false, the secondary resource will be initialized according to "resource initialization". The priority value specified with *-p* option is the priority of system *sys* for the resource *tag*. The priority value specified with *-r* option is the priority of system *othersys* for the resource *othertag*.

eqv_list

```
eqv_list [-d destsys] [-s sys] [-t tag] [-e SHARED] [-fC]
```

This function prints strings to standard output describing equivalency relationships between resource instances. If *destsys* is not specified, the current system is assumed. Each line contains fields separated by a delimiter character. The default delimiter character is ^A (\001). If the *-fC* option is specified, the delimiter is changed to C.

The example listings below show a colon (:) as a delimiter.

```
LKSYSA:135.66.249.201:LKSYSB:135.66.249.201:SHARED
```

```
LKSYSA:FSLIST.0:LKSYSB;FSLIST.0:SHARED
```

```
LKSYSA:LK0-LKSYSA:LKSYSB:LK0-LKSYSA:SHARED
```

Using `LKSYSA:fi.vo.0:LKSYSB;fi.vo.0:SHARED`, these are the fields:

LKSYSA	Primary system name where the resource resides.
fi.vo.0	Volume resource tag on the primary system.
LKSYSB	System name for the secondary system, where the resource equivalency resides.
fi.vo.0	Volume resource tag for the equivalent resource on the secondary system.
SHARED	Equivalency type.

The remaining arguments limit the information output as specified below:

- s *sys*. This option limits the output to include only the equivalencies relating to the system specified by the *sys* argument.
- t *tag*. This option limits the output to include only the equivalencies relating to the tag specified by the *tag* argument.
- e *SHARED*. This option prints all SHARED equivalency information.

eqv_remove

```
eqv_remove [-d destsys] [-s sys] -t tag [-S othersys] -o othertag [-e SHARED]
```

Removes equivalency from the configuration database on system *destsys* (local if not specified) of equivalency type, specified by the *-e* option, between the resources *tag* and *othertag* existing on systems *sys* and *othersys*, respectively. If *sys* or *othersys* is not specified, the current system is assumed.

LCDI-resource_type

Synopsis

[typ_create](#) [-d destsys] -a appname -r restyp

[typ_list](#) [-d destsys] -a appname -r restyp

[typ_remove](#) [-d destsys] [-fC] [-a appname]

Description

Resources are used by LifeKeeper to represent volumes, applications or other objects known by the system. Resource types are classifications of resources and are distinguished by a common set of recovery procedures that can be applied to all instances. Resource type examples would include:

- File system volumes, for example K:
- File shares, for example UTIL_SHARE
- IP addresses, for example 153.66.232.21.

The `typ_create` and `typ_remove` commands provide an interface for generating new types in the configuration database. The command `typ_list` provides an interface to the configuration database for listing all resource types existing on a specific system.

Exit Codes

All commands exit to 0 if they are successful. On failure, they return a nonzero code (see EXIT CODES section) and print to standard error. The following exit codes could be returned by these commands:

0	The operation has succeeded.
1	A system call or library call has internally returned failure.
2	A user-specified syntax error occurred.
3	LifeKeeper internal error.
4	A request to perform an operation on an object that already exists.

5	An argument specified is illegal.
6	Index out-of-range.
7	A request has been made on an object that does not exist.
8	A request was made to delete a resource instance on which another non-deleted resource instance depends.
9	An attempt to communicate with another system failed.

typ_create

```
typ_create [-d destsys] -a appname -r restyp
```

Creates a new resource type in the configuration database on system *destsys* (local, if not specified). The resource type is named *restyp* and is installed under the already-existing application, *appname*. Failure occurs if the system or application is not known or if the resource type already exists.

typ_list

```
typ_list [-d destsys] [-fC] [-a appname]
```

This command prints to standard output a list of resource types that have been defined on the application, *appname*, installed on system *destsys* (local, if not specified). If *appname* is not specified, all resource types for all applications are printed in the following format:

```
filesys:volume
```

```
comm:ipM
```

```
database:informix
```

The application name is to the left of the delimiter and the resource typename is to the right. Each line contains fields separated by a delimiter character. The default delimiter character is ^A (\001). If the *-fC* option is specified, the delimiter is changed to the specified character. The above example shows a colon (:) as the delimiter.

typ_remove

```
typ_remove [-d destsys] -a appname -r restyp
```

Removes the given resource type from the configuration database set of known resource types of system *destsys* (local, if not specified). All resource instances, dependencies, and equivalencies associated with this type are also removed. Failure occurs if the resource type is not known to the configuration database.

LCDI-systems

Synopsis

`sys create` [-d destsys] -s sys

`sys getds` [-d destsys] -s sys

`sys getst` [-d destsys] -s sys

`sys list` [-d destsys] -s sys

`sys remove` [-d destsys]

Description

The LifeKeeper configuration database knows about related systems. Because resources and resource types are specific to the systems on which they exist, it is necessary for the configuration database interface to contain the concept of a system.

The LCDI-systems commands return (or create) information into or remove information out of the database.

Exit Codes

All commands exit to 0 if they are successful. On failure, they return a nonzero code and print to standard error. The following exit codes could be returned by these commands:

0	The operation has succeeded.
1	A system call or library call has internally returned failure.
2	A user-specified syntax error occurred.
3	LifeKeeper internal error.
4	A request to perform an operation on an object that already exists.

5	An argument specified is illegal.
6	Index out-of-range.
7	A request has been made on an object that does not exist.
8	A request was made to delete a resource instance on which another non-deleted resource instance depends.
9	An attempt to communicate with another system failed.

sys_create

```
sys_create [-d destsys] -s sys
```

Creates a new system definition in the configuration database on system *destsys* (local, if not specified). The *-s sys* option is required to identify the system to which the system name is assigned.

sys_getds

```
sys_getds [-d destsys] -s sys
```

Prints to standard output the optional text description of why the system has gone to the current state from the database on system *destsys* (local, if not specified).

sys_getst

```
sys_getst -s sys
```

Prints the system's state to standard output as one of the strings:

DEAD	The system is believed to be unavailable.
ALIVE	The system is believed to be available.
UNKNOWN	System state is unavailable.

sys_list

```
sys_list [-d destsys]
```

This command prints to standard output a list of systems that LifeKeeper knows about from the database on system *destsys* (local, if not specified).

sys_remove

```
sys_remove [-d destsys] -s sys
```

Removes a system definition from the configuration database on system *destsys* (local, if not specified). The *-s sys* option is required to identify the system to which the system name is assigned.

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>:Restore_hierarchy.** The creation of a file system hierarchy produces a flag in this format in the status display. The filesys designation can be appdp for applications with disk partition hierarchies or appfs for applications with file system hierarchies.
- 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 Flags](#) for more detailed information on the possible flags.

flg_create

```
flg_create [-d destsys] -f flag
```

The flag *flag* is created on system *destsys*.



Note: This only modifies the "shared memory" segment of the LifeKeeper configuration database.

The LifeKeeper **lcdsync** command should be run after this command to ensure that the shared memory changes are reflected into the permanent storage onto a disk file.

flg_list

```
flg_list [-d destsys]
```

flg_list prints to standard output a short listing, one flag per line, of all of the flags currently defined on this system (unless *destsys* is specified). The listing is not in any particular order.

flg_remove

```
flg_remove [-d destsys] -f flag
```

The flag *flag* is removed on system *destsys*.



Note: This only modifies the shared memory segment of the LifeKeeper configuration database.

The LifeKeeper **lcdsync** command should be run after this command to ensure that the shared memory changes are reflected into the permanent storage onto a disk file.

flg_test

```
flg_test [-d destsys] -f flag
```

A check is made to see if the flag *flag* exists on system *destsys*. Returns 0 or 7.

LCDI Flags

Synopsis

[`flg create`](#) [-d destsys] -f flag

[`flg remove`](#) [-d destsys] -f flag

[`flg test`](#) [-d destsys] -f flag

[`flg list`](#) [-d destsys]

Description

LifeKeeper provides a facility to dynamically set flags to perform various tasks. The following special purpose flags can exist.

`!nofailover!uname`

If this flag exists, failover is inhibited for resources on the system with name, `uname`, that have defined the system with the flag on it as their backup system. **Note:** This is a temporary flag that will be removed automatically when LifeKeeper detects that system `uname` is ALIVE.

`!action!procid!timestamp!uname:identifier`

This is an example of an "admin lock flag" [see `getlocks`]. These flags are used for actions that require that no other action be performed at the same time on any of the systems in a LifeKeeper configuration. For example, you may not create a hierarchy on one system while creating a hierarchy on another. The "admin lock flags" are used to ensure that one of these "global" operations is not performed until the one currently running completes.

The identifier field of the "admin lock flag" identifies the kind of action being performed. The system that the process that requested the "admin lock flag" was running on is specified by `uname`. The flag was created at timestamp number of seconds after Jan 1, 1970, by a process with a process ID of `procid` that called `getlocks` [see `getlocks`].

An example of such a flag is as follows:

!action!01525!701120147!cindy:Create_Hierarchy

This flag indicates that the action Create_Hierarchy is in progress, indicating a hierarchy is being created. The process with process ID1525 requested the "admin lock flag," at the time 701120147 on system cindy.

!restore

This flag is set by LifeKeeper when the prerestore scripts (see [LCD](#)) are run. It indicates that the postrestore scripts should be run. Normally, this is a transitory condition that LifeKeeper automatically fixes when the postrestore scripts (see [LCD](#)) run. The only exception is if the postrestore scripts are being explicitly run using the following command:

```
%LKROOT%\bin\lcdrecover -G restore
```

!restore!uname

When this flag is set it indicates that the postrestore scripts (see [LCD](#)) should be run remotely on system uname. When the postrestore scripts are run on this system, LifeKeeper sends a remote request to system uname to run its postrestore scripts. Normally, this is a transitory condition that LifeKeeper automatically fixes. The only exception is if the postrestore scripts are being explicitly run using the %LKROOT%\bin\lcdrecover -G restore command.

!remove

This flag is set by LifeKeeper when the preremove scripts (see [LCD](#)) are run. It indicates that the postremove scripts should be run at a later time. Normally, this is a transitory condition that LifeKeeper automatically fixes when the postremove scripts (see [LCD](#)) are run at a later time. The only exception is if the postremove scripts are being explicitly run using the following command:

```
%LKROOT%\bin\lcdrecover -G remove
```

!remove!uname

When this flag is set it indicates that the postremove scripts (see [LCD](#)) should be run remotely on system uname. When the postremove scripts are run on this system, a remote request is sent to system uname to run its postremove scripts.

Normally, this is a transitory condition that LifeKeeper automatically fixes. The only exception is if the postremove scripts are being explicitly run using the %LKROOT%\bin\lcdrecover -G remove command.

!delete

This flag is set by LifeKeeper when the predelete scripts (see [LCD](#)) are run. It indicates that the postdelete scripts should be run at a later time. Normally, this is a transitory condition that LifeKeeper automatically fixes when the postdelete scripts (see [LCD](#) (1M)) are run. The only exception to this is if the postdelete scripts are being explicitly run by using the following command:

```
%LKROOT%\bin\lcdrecover -G delete
```

!delete!uname

When this flag is set it indicates that the postdelete scripts (see [LCD](#)) should be run remotely on system uname. When the postdelete scripts are run on this system, a remote request is sent to system uname to run its postdelete scripts. Normally, this is a transitory condition that LifeKeeper automatically fixes. The only exception is if the postdelete scripts are being explicitly run using the following command:

```
%LKROOT%\bin\lcdrecover -G delete
```

lk_chg_value

NAME

`lk_chg_value.ksh` -- changes specified values in local LifeKeeper configuration database files

SYNOPSIS

```
lk_chg_value.ksh {-o old_value -n new_value | -f filename} [-vFIMT]
```

DESCRIPTION

This command is to be used to modify arbitrary values in local LifeKeeper configuration database files (e.g. LifeKeeper uname, communication path addresses, resource tag names, etc.). **lk_chg_value.ksh** needs to be run locally with the Administrator login on each machine within a LifeKeeper configuration while LifeKeeper is not running. Also, you must use the LifeKeeper provided shell (`sh.exe`) to invoke the script as shown above. This command does not modify the system's uname or network interfaces. If the LifeKeeper uname or communication path addresses are to be modified, the system uname and network interfaces must be modified prior to the execution of this command using system utilities. In order for LifeKeeper to be properly updated, this command must be run on every system in the cluster.

The values to be modified may be specified on the command line using the `-o` and `-n` options or in a file using the `-f` option. The syntax of a file containing substitutions is `old_value=_new_value`, with one substitution per line (lines not in this form are ignored).

To see the changes **lk_chg_value.ksh** will make without modifying any LifeKeeper files, use the `-M` option. To see the files **lk_chg_value.ksh** is examining, use `-v`. To not modify tag names, use the `-T` option. To not modify resource ids, use the `-I` option.

Because a resource id may contain structured information, **lk_chg_value.ksh** does not allow substitutions that completely replace the id field. To override this

behavior, use the *-F* option.

EXAMPLES

Systems A, B, and C have been configured in a LifeKeeper configuration. Systems A and B manage a database resource and system A also manages a communication resource with System C. To modify the uname and comm path address of System A with comm path old address, the following must be performed:

1. Stop LifeKeeper by executing the `lkstop` command on each affected system. However, if the resources being managed are to stay available to the user community, execute `lkstop -f`.
2. Change System A's uname to X and change the network address to `new_address`. Create a substitutions file, `/tmp/lksubs`, containing the substitution pairs:

A=X old_address=new_address

As Administrator, login to System A and execute the following:

```
set LKROOT=<LKROOT> (i.e. set LKROOT=C:\LK)
```

```
<LKROOT>\bin\sh.exe lk_chg_value.ksh -vf /tmp/lksubs
```

This changes all local occurrences of A and `old_address` found within the LifeKeeper core and recovery kits on System A, which is now identified by System X, to refer to X and `new_address`, respectively.

3. Copy the substitutions file from System A to Systems B and C. As Administrator, login to Systems B and C and execute the following:

```
set LKROOT=<LKROOT> (i.e. set LKROOT=C:\LK)
```

```
<LKROOT>\bin\sh.exe lk_chg_value.ksh -vf /tmp/lksubs
```

This changes all occurrences of A and `old_address` found within the LifeKeeper configuration database on Systems B and C to refer to X and `new_address`, respectively.

EXIT CODES

0	Execution of command completed successfully.
1	Interrupt occurred... files restored.
2	Invalid arguments passed to command.
3	LifeKeeper processes are still running.
4	Command needs to be executed by an Administrator login.
5	ID field change attempted. Resource ID cannot be changed without using -I option.
6	LKROOT environment variable not set.
7	No matches were found.

NOTES

The **lk_chg_value.ksh** utility is located in the <LKROOT>\bin folder.

The **lk_chg_value.ksh** utility is case sensitive.

As shown above you must use the LifeKeeper provided shell (sh.exe) to invoke the **lk_chg_value.ksh** script.

<LKROOT> refers to the LifeKeeper home directory. The default home directory is C:\LK, but this can be modified during LifeKeeper installation.

FILES

<LKROOT>\bin\lk_chg_value.ksh

lk_err

Synopsis

```
lk_err -c Category -n Error number -p Process Name [-d {TO_LOG |TO_STDERR}]  
"Message"
```

Description

This utility is used within recovery scripts to log errors to the Microsoft Event Log. It also prints messages to stderr.

The arguments are:

Category. The following is a list of LifeKeeper message categories and their Event Log classifications:

LK Category	Event Category	Event Type
FRS_MES	General	Information
FRS_WARN	General	Warning
FRS_ERR	General	Error

Error Number. Must be a positive integer.

Process Name. Name of the script calling **lk_err**.

Destination. The destination parameter is optional. By default, events generated by **lk_err** will be directed to both the Windows Event Log (TO_LOG) and to the system console stderr message stream (TO_STDERR). However, with the **-d** option you may direct events specifically to one destination or the other.

Please note, however, that messages directed to the stderr (TO_STDERR) by programs or scripts that are executed by the LifeKeeper core will not display on the system console because the LifeKeeper core runs them as background tasks without interactive properties. Therefore, directing messages to stderr is useful only as a manual script testing and debugging aid.

Message. Message string must be enclosed in " ".

perform_action

Synopsis

```
perform_action [-G] [-s] [-b] [-n] -t tag-name -a action-name [- - arg1 arg2 . .  
. argn]
```

Description

The LRACI program **perform_action** performs processes in the following order:

- Finds the resource specified by the *tag-name* argument of the *-t* option.
- Finds the action script specified by the *action-name* argument of the *-a* option.
- Executes the action script on the *tag-name* resource instance.

The arguments after the *--* argument are passed unchanged to the action script(s). These are arguments that the developer of the action may optionally require to use the action.

The **perform_action** program finds the action script by the following algorithm: it first searches in the actions directory for the resource type of the resource instance specified by *tag-name*:

```
%LKROOT%\subsys\appname\resources\restypname\actions\action-name.ksh
```

and if not there, it checks the actions directory of the application the resource instance belongs to:

```
%LKROOT%\subsys\appname\actions\action-name.ksh
```

The restore and remove actions are special cases for LRACI. The restore action moves an application hierarchy that may be in-service on the remote system to the local system. For restore, LRACI first checks to make certain that the resource instance *tag-name* is not already in the ISP state. If it is not, it

recursively checks all of the resource instances that this resource depends upon. The check continues until a resource is found that either depends on no resources, or all of the resources it depends on are already in the ISP state. If the resource was in the ISU state, it is placed in the ISP state.

If the resource was in the OSU or OSF state, LRACI executes the remove script for any affected resources on the remote system. When this completes, LRACI finds the restore action using the above algorithm and runs it. If the script fails, the resource is placed in the OSF state and LRACI stops. If it succeeds, LRACI recursively "restores" the resources up the tree, until the resource specified by *tag-name* is restored. Then the LRACI recursively checks and "restores" the parent resource instances in a similar fashion until all related root resource instances are restored. In each case, LRACI uses the above algorithm to find the correct restore script using the resource application and resource type of the resource currently being scanned by LRACI, not the resource application and resource type of the *tag-name* resource.

For the remove action, the resources are moved recursively in the opposite direction. LRACI calls the remove script of all resources starting at the root resources that depend directly or indirectly on the *tag-name* resource down to, and including, the *tag-name* resource if any of those resources are in the ISP or ISU state. Resources not in the ISP or ISU state are ignored. If one of the remove scripts fails, LRACI places the failing resource into the OSF state and stops. In each case, LRACI uses the algorithm to find the correct remove script using the resource application and resource type of the resource currently being scanned by LRACI, not the resource application and resource type of the *tag-name* resource.

The remove and restore actions automatically have the *-t tag-name* and *-i ident-field* arguments added to the argument list that corresponds to the resource instance being acted upon.

The following sections discuss the arguments accepted by **perform_action**.

-G	This option is only used if action-name is remove, restore, or delete. If this option is not specified, LRACI performs the preglobal and postglobal scripts before and after the actions are performed [see
----	--

	<p>lcdrecover in LCD]. If the option is specified, LRACI does not run the preglobal and postglobal scripts.</p> <p>This option is useful if you need to run perform_action more than once, but you only want to run the preglobal and postglobal scripts once. It is also useful if you need to run perform_action while creating a resource hierarchy. The preglobal and postglobal scripts should not be run by perform_action during hierarchy create because the hierarchy creation scripts should be set up to obtain the "admin lock flags" [see LCDI_flag] and postrestore also requires the "admin lock flags" which would lead to contention problems.</p>
-s	<p>The default behavior for the restore action is to bring all objects above and below the specified tag into service, and the default behavior for the remove action is to bring all objects above the specified tag out of service. The -s option limits the scope of the operation to just the specified flag.</p>
-b	<p>The default behavior for the restore action is to bring all objects above and below the specified tag into service. The -b option changes this behavior to just objects below the specified tag. This option has no effect on the remove action.</p>
-n	<p>This option is only used if the <i>action-name</i> is remove or restore. If this option is specified, the resource reserves are not to be checked and the actions are performed whether the resources are reserved or not.</p> <p>WARNING: EXTREME CAUTION SHOULD BE TAKEN WHEN USING THIS OPTION!</p> <p>If this option is not specified, before any remove or restore scripts are executed, LRACI checks to see if any of the resources on which any of the actions will be run are currently reserved by another process. A resource can be reserved while the following operations are being performed on them: a resource "remove from service" is in progress, a "resource restore to service" is in progress, or a resource "recovery" is in progress.</p> <p>If any resource is so reserved, LRACI waits a specified period of time</p>

	for the process to remove the reserve on the resource. If that period expires, LRACI removes the reserve. In either case, LRACI reserves all of the resources, then follows the specified algorithm to perform the action(s).
<code>-t tag-name</code>	This is the last resource instance the action will be performed on.
<code>-a action-name</code>	This is the resource action that will be performed.
<code>-- arg1 arg2 ... argn</code>	Argument(s) the resource action developer can optionally define to be passed to the action script. When executing perform_action within a getlocks protected region and the <code>-G</code> option is not used, set <code>arg1</code> to <code>-m</code> to avoid executing a second instance of getlocks, which would cause the operation to hang.

Example

The following is an example of calling an action:

```
perform_action -t SCSI-USR-DISK -a reset-heads - -h 7
```

The LRACI program **perform_action** would find the action corresponding to `reset-heads` and execute it with the arguments:

```
reset-heads -t SCSI-USR-DISK -h 7
```

Exit Codes

The following exit codes could be returned by LRACI:

0	The operation has succeeded.
1	A system call or library call has internally returned failure.
2	A user-specified syntax error occurred.
3	LifeKeeper internal error.
4	A request to perform an operation on an object that already exists.
5	An argument specified is illegal.

6	Index out-of-range.
7	A request has been made on an object that does not exist.
8	A request was made to delete a resource instance on which another non-deleted resource instance depends.
9	An attempt to communicate with another system failed.

sendevent

Synopsis

```
%LKROOT%\bin\sendevent -C class-of-event -E event -m monitor-name -nname-of-obj-  
inst [-s severity]
```

Description

The event notification facility consists of two parts: an event notification mechanism (`%LKROOT%\bin\sendevent`) and an application registration environment. Applications wishing to use the event facility should "register" to get notification of specific events or alarms (or all occurrences of event/alarms).

The **sendevent** command is a program invoked by a daemon monitor process when the monitor has detected an event (failure or recovery) in the objects that it is monitoring. This command is not intended to be run directly at the shell level by a regular user or by a system administrator (only by a daemon process or another command).

The **sendevent** command is used to notify a "registered" application of the occurrence of an event. For example, an application may want to be notified of an impending system shutdown so it can appropriately save files and data; or, in a client-server environment, the application may need to reconfigure itself to an alternate service provider. The application is responsible for providing the appropriate command support to handle the event.

The **sendevent** command passes all of its options to the event-response commands of the application.

An application registers to receive notification of events or alarms by installing its event-response commands in a specific registration directory, `%LKROOT%\events`. This should be done at application installation time. The events under `%LKROOT%\events` are further categorized in classes of events. Create separate subdirectories, `genclass` (for general events) and `allclass` to be used by applications to register to be notified upon occurrence of any event.

* **Note:** If an event occurs which causes an application to place application-response commands in both the specific event location and the all location, both scripts run.

Each class directory contains a subdirectory for each event within that class. Each add-on package that monitors events and uses this event notification mechanism documents events it monitors and supports.

It is the responsibility of the application object monitor package to maintain a file called ACTIVE in the events subdirectories. If the ACTIVE file exists, it is a signal to the applications that a monitor is currently actively running and monitoring its objects for that specific event. If the package containing the monitor program is removed, the files named ACTIVE for the affected monitored events are removed too (by the package remove script) to indicate to applications that the event is no longer being monitored. The removal of the package should not remove event-response commands or event directories even if they are empty.

For those applications that may depend upon standard commands from another application, the application registration environment provides other application-specific directories, `%LKROOT%\subsys\application-name\actions`, for applications to place "sharable" action commands. For example, application X may depend upon application Y being up and running after an event recovery. If this is not the case, application X may invoke the start command for application Y from the `LKROOT%\subsys\Y\actions\start` directory. Interdependencies between applications must be resolved and specified by the application developers.

The `-C` (class of event), `-E` (event), `-m` (monitor name), and `-n` (name of object instance) options are required. If the `-s` (severity) option is not specified, `sendevent` will default to a severity of MAJOR alarm.

Upon invocation of the **sendevent** command by a monitoring process, **sendevent** determines which event class and event has occurred based upon the arguments to the `-C` and `-E` options. The **sendevent** command executes in the background until it finishes processing all the event-response commands (if any) placed in the registration directory corresponding to that class/event pair and all of the commands registered in the all directory.

The following options are supported:

-C class-of-event

Events are grouped together into classes. This required option indicates which class the event belongs to.

-E event

This required option indicates which event in a class is being triggered.

-m monitor-name

Each application object monitor that can send alarms/events is identified by a name in the form:

OM-product-name:OM-component-name

OM-product-name is an ASCII string, of up to eight characters. It is an abbreviated identifier specifying the product that monitors the objects that cause alarms or events. OM-component-name is an ASCII string, of up to 16 characters. It is defined by the object monitor to identify the component of the object monitor that detected the alarm or event.

The monitor names are used to distinguish between different products that may be used to monitor the same object.

-n name-of-obj-inst

This option is used to name a specific instance of an application object. It is an ASCII string with a maximum length of 64 characters. For example, D: may be the name of a volume application object, whereas 1234 could be used to identify a specific process object.

-s severity

Each alarm or event must specify the severity of the problem it is reporting. If this option is not specified, **sendevent** internally adds the default severity for MAJOR alarm. Severity is an ASCII represented integer interpreted as follows:

0	CLEARED alarm specified by "id-of-alarm/event" has been recovered
1	INFORMATIONAL alarm (INFO message or cmn_err() NOTICE message)
2	WARNING alarm (WARNING message)
3	MINOR alarm (MINOR message)
4	MAJOR alarm (MAJOR or ERROR message) (default)
5	CRITICAL alarm (CRITICAL message or cmn_err() PANIC or HALT message)

Output

The output this command generates occurs in one of two conditions:

- Error messages are printed to standard error and a nonzero exit code is returned.
- The identifier for the alarm\event called id-of-alarm/event is printed to standard output at each call to sendevent.

Exit Codes

The following exit codes are returned by **sendevent**:

0	The sendevent command has completed successfully without errors.
1	Syntax error in the argument list.
2	No class corresponding to the string passed with the -C option exists in the %LKROOT%\events directory.
3	No event corresponding to the string passed with the -E option exists in the %LKROOT%\events\<class= directory.
4	The -A option is internally generated and may not be specified directly.
5	The -i option must be specified if the -s 0 (severity CLEARED) option is used.

volume

Synopsis

```
volume [-d | -D] [-l | -u | -p | -U volume_letter]
```

Description

This command is used to lock and unlock volumes on the Windows server. It may also be used to register with the LifeKeeper Service. When used in this fashion, it determines which volumes should be protected (locked) by LifeKeeper at startup. The lock provides LifeKeeper with exclusive access to the volume and will not allow any other process to access the volume.

LifeKeeper must be running in order for this command to succeed. The command interfaces with the LifeKeeper Service to provide the locking mechanism.

The following options are available where *volume_letter* is the drive letter to be locked/unlocked or protected/unprotected (i.e. C to Z).

-d	Display the currently locked volumes.
-D	Display the volumes that are registered with LifeKeeper. This would display volumes that have been added with the -p option. Generally, -D displays a different list than the one shown by the -d option.
-l	Lock the volume for exclusive access. The lock will fail if a remote user has opened the volume or a local application has opened the volume for a write operation.
-u	Unlock the volume from exclusive access.
-p	Register the volume with LifeKeeper, so that on subsequent reboots or restarts of LifeKeeper, the volume is automatically locked.
-U	Unregister the volume with LifeKeeper so that it is not automatically locked on LifeKeeper startup.

Example

The following illustrates how the volume command should be used:


```
#

# Register drive volume e: to be locked by LifeKeeper

#

ret=`volume -p E`

if [ $ret -gt 0 ]

then

    # Report error that it wasn't protected

fi

#

# Lock volume e: for exclusive access

#

ret=`volume -l E`

if [ $ret -gt 0 ]

then

    # Report error that it wasn't locked

fi
```

Exit Codes

The following exit codes could be returned by this command:

0	The operation has succeeded.
greater than 0	The operation has failed. An error message is printed to standard error.

LKSUPPORT

LKSUPPORT, found in the *LK/SUPPORT* directory, is used to collect important configuration information and event log files and put them in a zip file. SIOS Support Engineers will commonly request this zip file as part of the Support process. To run this utility, simply double-click LKSUPPORT and the zip file will be created in the same Support directory.

Setting Browser Security Parameters

In order to run the GUI web client, you must set your browser security settings to low. For Internet Explorer, follow the procedures below.



Be careful of other sites you visit with low security settings.

Internet Explorer

The most secure method for using Internet Explorer is to add the SIOS Protection Suite server to the Trusted Sites zone as follows:

1. From the **Tools** menu, click **Internet Options**.
2. Click the **Security** tab.
3. Select **Trusted Sites** zone and click **Custom Level**.
4. Under **Reset custom settings**, select **Medium/Low**, then click **Reset**.
5. Click **Sites**.
6. Enter the **server name** and **port number** for the SIOS Protection Suite server(s) to which you wish to connect (for instance: http://server1:81).

An alternative, but possibly less secure method is to do the following:

1. From the **Tools** menu, click **Internet Options**.
2. Select either **Internet** or **Local Intranet** (depending upon whether your remote system and the SIOS Protection Suite cluster are on the same intranet).
3. Adjust the **Security Level** bar to **Medium** (for Internet) or **Medium-low** (for Local Intranet). These are the default settings for each zone.
4. Click **OK**.

IP Local Recovery

When IP Local Recovery is enabled and the IP resource fails its deepcheck (a periodic extensive check of the IP resource), then SIOS Protection Suite will do the following:

- First, SIOS Protection Suite will attempt to bring the IP address back in service on the current network interface.
- If that fails, SIOS Protection Suite will check the resource instance to determine if there is a backup interface available. If so, it will then attempt to move the IP address to the backup interface.
- If all local recovery attempts fail, SIOS Protection Suite will perform a failover of the IP address and all dependent resources to a backup server.

Even if you do not have a backup adapter, you can enable Local Recovery so that SIOS Protection Suite will retry the primary network interface before initiating failover to a backup server.

Overview of SIOS Protection Suite 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, thereby 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).

SIOS Protection Suite has an event notification mechanism for registering applications that wish to be notified of specific events or alarms (see the [sendevent](#) man page). SIOS Protection Suite can be easily enabled to send SNMP trap notification of key SIOS Protection Suite events to a third party network management console wishing to monitor SIOS Protection Suite activity. SIOS Protection Suite installs an MIB file under `%LKROOT%\include\LifeKeeper-MIB.txt` which describes SIOS Protection Suite trap definitions.

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

SIOS Protection Suite Events Table

The following table contains the list of SIOS Protection Suite 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 SIOS's 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**

SIOS Protection Suite Event/Description	Trap #	Object ID
---	--------	-----------

LifeKeeper Startup Complete Sent from a node when LifeKeeper is started on that node	100	.1.3.6.1.4.1.7359.1.0.100
LifeKeeper Shutdown Initiated Sent from a node beginning LifeKeeper shutdown	101	.1.3.6.1.4.1.7359.1.0.101
LifeKeeper Startup Complete Sent from a node completing LifeKeeper shutdown	102	.1.3.6.1.4.1.7359.1.0.102
LifeKeeper Manual Switchover Initiated on Server Sent from the node from which a manual switchover was requested	110	.1.3.6.1.4.1.7359.1.0.110
LifeKeeper Manual Switchover Complete – recovered list Sent from the node where the manual switchover was completed	111	.1.3.6.1.4.1.7359.1.0.111
LifeKeeper Manual Switchover Complete – failed list Sent from the node where the manual switchover was completed	112	.1.3.6.1.4.1.7359.1.0.112
LifeKeeper Node Failure Detected Sent from each node within the cluster when a node in that cluster fails	120	.1.3.6.1.4.1.7359.1.0.120
LifeKeeper Node Recovery Complete – recovered list	121	.1.3.6.1.4.1.7359.1.0.121

Sent from each node within the cluster that has recovered resources from the failed node		
LifeKeeper Node Recovery Complete – failed list Sent from each node within the cluster that has failed to recover resources from the failed node	122	.1.3.6.1.4.1.7359.1.0.122
LifeKeeper Resource Recovery Initiated Sent from a node recovering a resource; a 131 or 132 trap always follows to indicate whether the recovery was completed or failed.	130	.1.3.6.1.4.1.7359.1.0.130
LifeKeeper Resource Recovery Failed 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
LifeKeeper Resource Recovery Complete 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
Mirror State Change Sent from the node, who is the source of the mirror, when the mirror state changes. Displays the volume letter, mirror state and IP address of the target node. Valid Mirror States: -1: Invalid State 0: No Mirror	150	.1.3.6.1.4.1.7359.1.0.150

1: Mirroring 2: Mirror is resyncing 3: Mirror is broken 4: Mirror is paused 5: Resync is pending		
LifeKeeper replicated volume Split-Brain detected Sent from the node where LifeKeeper has detected mirror is Source on both sides. Displays volume letter and IP address of the target node.	160	.1.3.6.1.4.1.7359.1.0.160
The following variables are used to “carry” additional information in the trap PDU:		
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.

Java Upgrade

Steps to Upgrade Java Runtime Environment (JRE) Version for SIOS Protection Suite for Windows

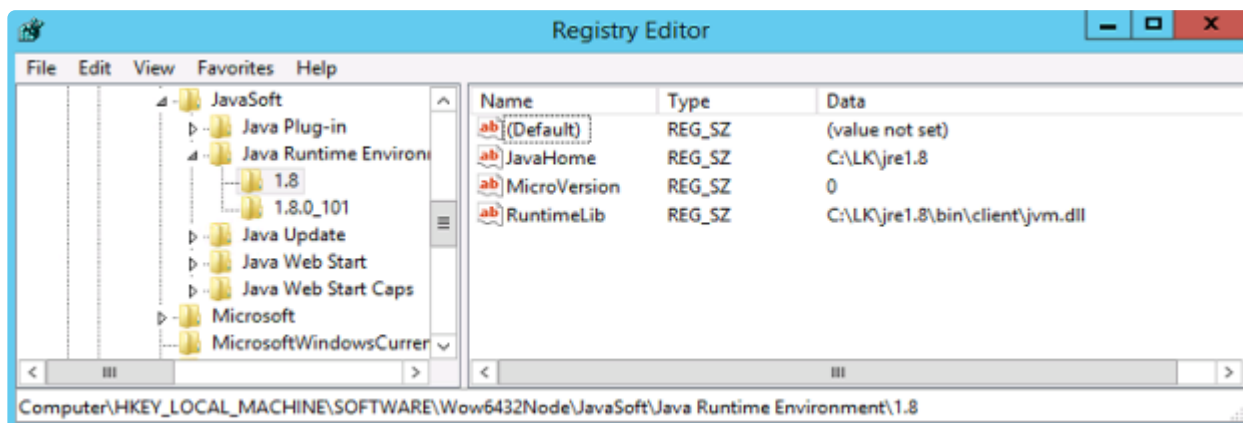
1. Download and install the latest 32-bit JRE from java.com. (**Note:** Make sure you are downloading the 32-bit Java.)
2. Run the installer. The first screen for the installer will have a check box to change the location of the destination folder. This box SHOULD be checked.
3. As the Java installer proceeds select a new subfolder in `c:\lk` (e.g. `c:\lk\jre1.8`) as the destination folder.
4. Repeat the above procedure (Steps 1-3) on the other node(s).
5. Stop LifeKeeper on the backup node(s).

```
net stop lifekeeper
```

6. Edit the registry keys.

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\SteelEye\LifeKeeper\JavaGUI :  
JavaVersion from "1.8.0_101" to "1.8"
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\SteelEye\LifeKeeper\JavaGUI\Server  
: JavaVersion from "1.8.0_101" to "1.8"
```



✿ **NOTE:** The 1.8 points to the key that JavaSoft installed in the previous step as shown in the picture, make the appropriate adjustment if the key is different.

7. If firewall is enabled, open the “[Windows firewall inbound rule](#)” **LifeKeeper Java**” and edit the property “**Programs and Services**” tab and select the Java executable (e.g. `c:\lk\jre1.8\bin\java.exe`).

Change the description to state that it is for Java JRE <version number>.

Also update or create a “**LifeKeeper Java Launcher**” rule and edit the property “**Programs and Services**” tab and select the Java executable (e.g. `c:\lk\jre1.8\bin\jp2launcher.exe`).

Verify that all firewall rule changes are correct.

8. Start LifeKeeper.

```
net start lifekeeper
```

9. Repeat Steps 5-8 on other backup nodes (if any).
10. Log in to SIOS Protection Suite and switch the hierarchy to the backup node.
11. Repeat the procedure (Steps 5-8) on the primary node.
12. Bring the hierarchy in service on the primary node.

At this point, SIOS Protection Suite has been upgraded using the latest version of Java.

Notes

1. If a repair is done on your SIOS Protection Suite installation using the installer/repair option, repeat the above [Step 6](#) as a repair will reset the registry settings. The directory created in [Step 3](#) should be intact.
2. When uninstalling SIOS Protection Suite, the directory created in [Step 3](#) will need to be manually removed.
3. The above steps assume that SIOS Protection Suite was installed in the default folder `c:\lk`. If not, change accordingly and replace `c:\lk` with the appropriate subdirectory.

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.

The User Guide is broken up into the following areas:

[LifeKeeper GUI](#) – These topics give an overview and description of the LifeKeeper Graphical User Interface.

[Common Tasks](#) – These topics cover 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](#) – This section covers more advanced tasks that require Operator permission such as bringing resources in and out of service.

[Advanced Topics](#) – This section details information on the LifeKeeper Configuration Database and LifeKeeper Communications.

[Maintenance Tasks](#) – The topics in this section cover the tasks necessary for maintaining SIOS Protection Suite.

[Data Replication](#) – The topics in this section provide details on using data replication with SIOS Protection Suite.

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

LifeKeeper GUI

The GUI components should have already been installed as part of the SIOS Protection Suite Core installation.

The LifeKeeper GUI uses Java technology to provide a graphical user interface to SIOS Protection Suite and its configuration data. Since the LifeKeeper GUI is a client/server application, a user will run the client component to monitor or administer SIOS Protection Suite. The client and the server components may or may not be run on the same system.

[GUI Overview](#)

[Toolbars](#)

[Menus](#)

[LifeKeeper GUI Server and Client Components](#)

[Running the SIOS Protection Suite Web Client](#)

[Running the GUI Application on a SIOS Protection Suite Server](#)

[LifeKeeper GUI User Accounts](#)

GUI Overview

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 is initialized on each SIOS Protection Suite server at system startup. It communicates with GUI clients using Hypertext Transfer Protocol (HTTP) and Remote Method Invocation (RMI).

GUI Client

The GUI client can be run either as a web client on any Java-enabled system or as an application on a SIOS Protection Suite server.

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.

Starting GUI Clients

Starting the Web Client

To run the web client on a SIOS Protection Suite server, click **Start** then point to **All Programs**, then point

to **SIOS->LifeKeeper->LifeKeeper**. This will invoke a web browser and connect to the local GUI server using `http://localhost:81`.

On systems outside the SIOS Protection Suite cluster, open a web browser and go to the URL `http://<server name>:81` where <server name> is the name of a SIOS Protection Suite server. This will load the web client from the GUI server on that machine.

After the web client has finished loading, you should see the [Cluster Connect Dialog](#) which allows you to connect the web client to any GUI server.

Note: When you run the web client, if your system does not have the required Java Plug-in, you will be automatically taken to the web site for downloading the plug-in. See the [Java Upgrade](#) topic for steps to upgrade. You must also [set your browser security parameters to enable Java](#).

If you have done this and the client still is not loading, see [Web Client Troubleshooting](#).

Starting the Application Client

Users with administrator privileges on a SIOS Protection Suite server can run the application client from that server. Click **Start**, then point to **All Programs**, then **SIOS->LifeKeeper->LifeKeeper (Admin Only)**.

If you have done this and the client still is not loading, see [Network-Related Troubleshooting](#).

Exiting GUI Clients

Select **Exit** from the [File Menu](#) to disconnect from all servers and close the client.

Status Table

The status table provides a visual representation of the status of connected servers and their resources. It shows the following.

- 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.

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.

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 **Refresh** button queries the server for the current values of all properties clearing any changes that you may have made. 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](#).

Output Panel

The output panel collects output from commands issued by the 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 GUI will return to its default behavior.

Message Bar

The message bar appears beneath the status window. It is used for displaying messages in a single text line. Messages 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](#).

Toolbars

SIOS Protection Suite for Windows Toolbars

[Global Toolbar](#)













[Resource Context Toolbar](#)



[Server Context Toolbar](#)

Global Toolbar

This toolbar is a combination of the default [server context toolbar](#) and [resource context toolbar](#) 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.



	Connect . Connect to a cluster.
	Disconnect . Disconnect from a cluster.
	Refresh . Refresh GUI.
	View Logs . View log messages.
	Create Resource Hierarchy . Create a resource hierarchy.
	Delete Resource Hierarchy . Remove a resource hierarchy from all servers.
	Create Comm Path . Create a communication path between servers.
	Delete Comm Path . Remove communication paths from a 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.








	Add Dependency . Create a parent/child relationship between two resources.
	Remove Dependency . Remove a parent/child relationship between two resources.

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 default toolbar is described here but this toolbar might be customized for specific resource types in which case the custom toolbar will be described in the appropriate resource kit documentation.

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.



	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.
	Add Dependency . Create a parent/child relationship between two resources.
	Remove Dependency . Remove a parent/child relationship between two resources.
	Delete Resource Hierarchy . Remove a resource hierarchy from all servers.

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.



	Disconnect . Disconnect from a cluster.
	Refresh. Refresh GUI.
	View Logs . View log messages.
	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.

Menus

SIOS Protection Suite for Windows Menus

[Resource Context Menu](#)

[Server Context Menu](#)

[File Menu](#)

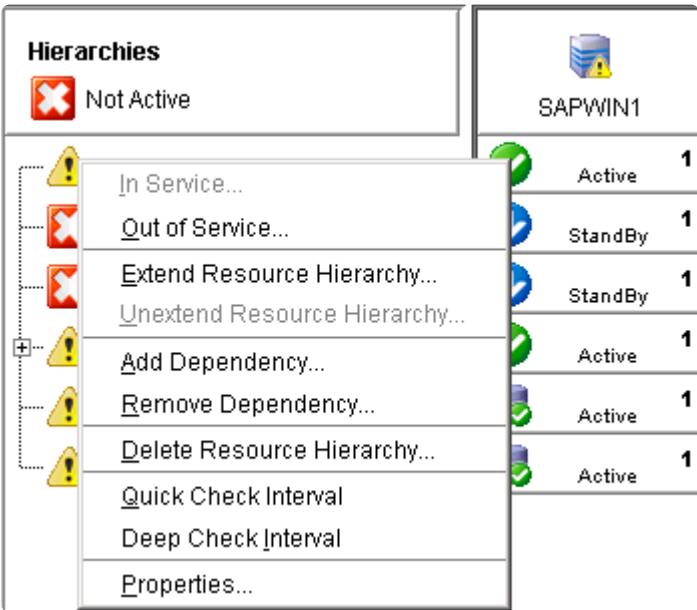
[Edit Menu – Resource](#)

[Edit Menu – Server](#)

[View Menu](#)

[Help Menu](#)

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.

[Add Dependency](#). Create a parent/child relationship between two resources.

[Remove Dependency](#). Remove a parent/child relationship.

[Delete Resource Hierarchy](#). Remove a resource hierarchy from all servers.

Local Recovery – Select **Yes** to enable Local Recovery for this Resource. Local recovery for a file share means that if the folder becomes inaccessible, SIOS Protection Suite will attempt to re-share the folder.

Quick Check Interval – Enter the interval (in minutes) between basic checks of the resource's availability. Different values can be specified for each system. The default value is 3 minutes. The value range is between 0 and 10080. Setting the interval value to 0 will disable the quick check feature.

Deep Check Interval – Enter the interval (in minutes) between extensive checks of the resource's availability. This program utilizes Quickcheck for its Deepcheck implementation. Different values can be specified for each system. The default value is 5 minutes. The valid entry range is between 0 to 10080. Setting the interval value to 0 will disable the Deepcheck feature.

[Properties](#). Display the resource properties dialog.

Server Context Menu

The server context menu appears when you right-click on a server in the [status table](#). The actions are always invoked on the server that you select.

View Logs...
Create Resource Hierarchy...
Create Comm Path...
Delete Comm Path...
Disconnect...
Refresh...
Properties...

[View Logs](#). View SIOS Protection Suite log messages.

[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.

[Disconnect](#). Disconnect from a cluster.

Refresh. Refresh GUI.

[Properties](#). Display the [server properties dialog](#).

File Menu




[Connect](#). Connect to a SIOS Protection Suite cluster (requires login authentication on each server).

Exit. Disconnect from all servers and close the GUI window.

Edit Menu – Resource

This submenu of the main menu bar is the same as the default [resource context menu](#) except that you must select a resource and server when you invoke actions from this menu. The **Edit > Resource** menu cannot be customized.



[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.

[Add Dependency](#). Create a parent/child relationship between two resources.

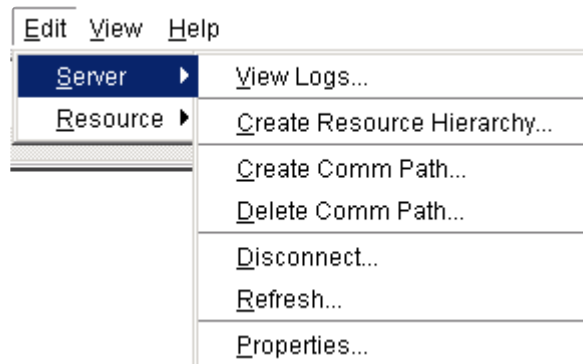
[Remove Dependency](#). Remove a parent/child relationship between two resources.

[Delete Resource Hierarchy](#). Remove a resource hierarchy from all servers.

[Properties](#). Display the [server properties dialog](#).

Edit Menu – Server

This submenu of the main menu bar is the same as the default server context menu except that you must select a server when you invoke actions from this menu. The **Edit > Server** menu cannot be customized.



[View Logs](#). View SIOS Protection Suite log messages.

[Create Resource Hierarchy](#). Create a resource hierarchy.

[Create Comm Path](#). Create a communication path between servers.

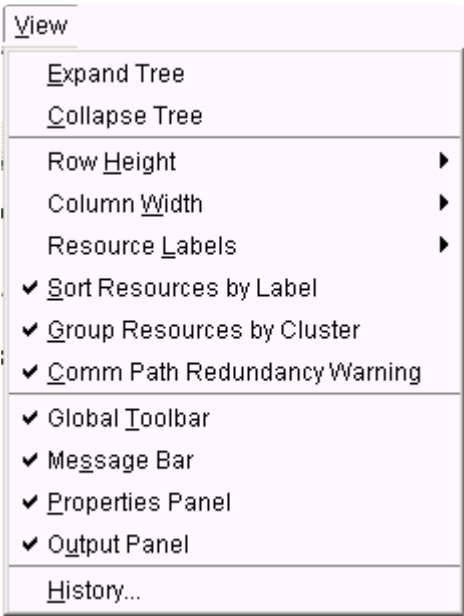
[Delete Comm Path](#). Remove communication paths.

[Disconnect](#). Disconnect from a cluster.

Refresh. Refresh GUI.

[Properties](#). Display the [server properties dialog](#).

View Menu



[Expand Tree](#). Expand the status table to show all resources in all hierarchies.

[Collapse Tree](#). Collapse the status table to show only the top resource in each hierarchy.



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.

By tag name:

**By 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 message bar messages in the message history dialog.

Help Menu



The **Help Menu** provides links to the **Release Notes** and **Technical Documentation** as well as descriptions of this documentation as follows:

Release Notes: Each product version provides Release Notes containing not only new features but also important information such as package versions, last-minute changes to instructions and procedures, 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 your software as it contains last minute information that must be considered before, during and after installation.

Technical Documentation: This online documentation resource is designed to provide the most up-to-date, detailed information about your SIOS product in an easy-to-use format. SIOS Technology Corp. maintains documentation for all supported versions of SIOS products on this site. For older versions of products, please request documentation from support@us.sios.com.

Select **About** to display the GUI version number.


LifeKeeper GUI Server and Client Components

The LifeKeeper GUI server is initialized on each SIOS Protection Suite server at system startup. It communicates with LifeKeeper GUI clients using Hypertext Transfer Protocol (HTTP) and Remote Method Invocation (RMI).

You can connect to the LifeKeeper GUI server with a web client that can be run from any system that can connect to Ports 81 and 82 of all servers in the cluster or with an application client that ships with SIOS Protection Suite and is designed to run on a server in the cluster.

Both SIOS Protection Suite clients include the same graphical components:

- Pop-up server and resource context menus provide access to server- and resource-related actions.
- The menu bar provides access to all LifeKeeper GUI actions.
- The toolbar provides quick access to many SIOS Protection Suite actions.
- The status window displays a graphical representation of the servers connected in the cluster, resource hierarchies and the status of resources and servers.
- The message bar at the bottom of the window displays processing information to the user.



Running the SIOS Protection Suite Web Client

If you wish to administer SIOS Protection Suite from a system outside your cluster, you must use the web client. This is possible for remote systems running any operating system. SIOS Protection Suite for Windows cannot manage both Linux and Windows servers in a single session, but it can manage either Windows or Linux systems no matter what OS it is running on. Whichever type of server OS you first connect to will determine the type of OS you can manage in that session. If you need to simultaneously manage both Linux and Windows servers, you will need to open up two browser windows, one for each.

The remote system's browser must provide JRE 1.8.0_101 support. Refer to the [SIOS Protection Suite for Windows Release Notes](#) for information on the supported platforms and browsers for the SIOS Protection Suite web client. The following sections explain steps for configuring the web browser on a remote system.

Follow the procedure below to run the SIOS Protection Suite web client.

1. Open the URL `http://<server name>:81` for the SIOS Protection Suite web page (where `<server name>` is the name of the SIOS Protection Suite server). The web page contains the SIOS Protection Suite splash screen and applet.

When you run the web client for the first time, if you are using Internet Explorer and your system does not have the required Java plug-in, you will be automatically taken to the appropriate web site for downloading the plug-in. See the [Java Upgrade](#) topic for steps to upgrade.

Notes:

- You should close and restart your browser after installing the plug-in and whenever plug-in properties are changed. Thus you will need to enter the SIOS Protection Suite server's URL again as stated above.

When the web page is opened, the following actions take place:

- the splash screen is displayed
- the applet is loaded
- the Java Virtual Machine is started
- some server files are downloaded

- the applet is initialized

Depending upon your network and system configuration, these actions may take up to 20 seconds. Typically, browsers provide some minimal status as the applet is loading and initializing.

Note: You may receive a Java Plug-In Security Warning stating “Unable to verify the certificate – code will be treated as unsigned.” Click **OK**.

Next, a Start button should appear in the applet area at the bottom of the splash screen. If the splash screen does not display a Start button or you suspect that the applet failed to load and initialize, refer to the [GUI Network-Related Troubleshooting](#) section in this guide.

2. Click **Start**. The SIOS Protection Suite web client appears and the [Cluster Connect dialog](#) is automatically displayed. Enter the **Server Name** you wish to connect to followed by the **login** and **password**. Once a Server Name has been entered and connection to the cluster established, the GUI window appears.

Note: Some browsers add “Warning: Applet Window” to windows and dialogs created by an applet. This is normal and should be ignored.

Configuring the Browser Security Level

In order to run the SIOS Protection Suite web client, you may need to modify your browser security settings. Follow the procedures below.

Internet Explorer

Internet Explorer will likely automatically put your SIOS Protection Suite servers in the Local intranet zone. If not, you should manually add all SIOS Protection Suite servers to the Local intranet zone as follows:

1. From the Tools menu, click **Internet Options**.
2. Click the **Security** tab.
3. Select **Local intranet zone**.
4. Click **Sites**.
5. Click **Advanced**.

6. Enter the server name(s) and port number(s) for all SIOS Protection Suite server(s) to which you wish to connect (for instance: *http://server1:81*), clicking **Add** after each.
7. Click **OK** until you're done.

Mozilla Firefox

1. From the **Tools** menu, select **Options**.
2. In the Options dialog box, click the **Content Category**.
3. Select the "**Enable Java**" and "**Enable Java Script**" options.
4. Click **OK**.
 - You should close and restart your browser after installing the plug-in and whenever plug-in properties are changed. Thus you will need to enter the SIOS Protection Suite server's URL again as stated above.

Running the GUI Application on a SIOS Protection Suite Server

You can also run the LifeKeeper GUI as an application on a SIOS Protection Suite server. By doing so, you are, in effect, running the GUI client and server on the same system. Only users with administrator privileges on the SIOS Protection Suite server are allowed to run SIOS Protection Suite applications.

1. Start the LifeKeeper GUI by clicking **Start->AllPrograms->SIOS->LifeKeeper->LifeKeeper (Admin only)**.
2. After the application is loaded, the LifeKeeper GUI appears and the [Cluster Connect dialog](#) is displayed. Enter the **Server Name** you wish to connect to followed by the login and password. See [LifeKeeper GUI User Accounts](#) for additional information on logging in.
3. Once a connection to the cluster is established, the GUI window appears.

To run the LifeKeeper GUI on a SIOS Protection Suite server using the web client, click **Start->All Programs-> SIOS->LifeKeeper->LifeKeeper**. This will invoke a web browser and connect to SIOS Protection Suite using *localhost:81*.

LifeKeeper GUI User Accounts

All LifeKeeper GUI users must belong to SIOS Protection Suite security groups. The SIOS Protection Suite administrator for a cluster can use local groups and user accounts on each server, or you can set up domain groups and users with local logon privileges.

Logging In

If the SIOS Protection Suite account is the same for each server in the cluster (same login and password), then logging in to one server in the cluster will allow you to access the other servers without additional logins. You may need to enter the domain name with the user name, for example: "Southdomain\john".

If the SIOS Protection Suite account is different for each server in the cluster (different login names and/or passwords), then upon logging in to the first server in the cluster, you will receive the following message when SIOS Protection Suite attempts to use the login for the next server in the cluster:

```
Access denied: invalid user name or bad password. Only users with local
privileges can use LifeKeeper. Would you like to re-enter the
authentication data?
```

Click **Yes** for a prompt to login to the next server.

Configuring GUI Users

There are three classes of GUI users with different permissions for each.

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

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.

User administration is performed by assigning users to local user groups on each server. Users assigned to the local Administrators group have Administrator permission, users in the local *LK_OPERATOR* group have Operator permission and users in the local *LK_GUEST* group have Guest permission. The local Administrators group is built in to all Windows machines, but the other two local groups are not, so you will need to create them.

The group names can be configured on English-language machines by editing the entries in the file *Server_RB_en.properties* which can be found in the folder *\$LKROOT/htdoc/com/SIOS/LifeKeeper/locale*. You can also localize the group names by creating a file *Server_RB_xx.properties* in the same folder, where "xx" is your locale code, and editing the entries in that file.

If you are working in a Domain Controller environment with no local groups or users on your servers, you can create the *LK_OPERATOR* and *LK_GUEST* groups as trusted global security groups. You will then need to set the group security policy to allow local logon to those groups.

To enable a user or a group to login locally on a Windows server, follow the instructions described below.

1. Log in to the machine using an account with local Administrator privileges.
2. Open the **Local Security Policy MMC** in the Administrative Tools program group.
3. Scroll down to **Local Policies -> User Rights Assignment**.

4. In the details pane, double-click **Allow Logon Locally** policy.
5. Use the **Add User or Group...** button to add domain groups *LK_OPERATOR* and *LK_GUEST* previously created for local login right.



IMPORTANT: Please ensure that the domain GPO does not overwrite these local policy changes.

Finally, you need to propagate these changes by executing the command `SECEDIT /REFRESHPOLICY USER_POLICY gpupdate` (for more details, see <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/gpupdate>). Once you have done this, SIOS Protection Suite will be able to recognize members of those groups and assign them the appropriate permissions.

Note: If you create these groups and users locally on your server, the assignments affect GUI permissions only for that server. In that case, you should repeat the assignment on all servers in the cluster. This takes more work but does make the cluster more robust as it is then not dependent on access to the domain controller.

Note: GUI group names, *Administrators*, *LK_OPERATOR*, and *LK_GUEST* referenced above cannot be changed.

Common Tasks

This section covers basic tasks that can be performed by any user.

[Connecting To A Cluster](#)

[Disconnecting From a Cluster](#)

[Viewing Connected Servers](#)

[Viewing The Status Of A Server](#)

[Viewing Server Log Files](#)

[Viewing Server Properties](#)

[Viewing Resource Tags and IDs](#)

[Viewing the Status of Resources](#)

[Viewing Resource Properties](#)


[Viewing Message History](#)

[Expanding and Collapsing A Resource Hierarchy Tree](#)

Connecting To A Cluster

1. From the [File Menu](#) or the [Global Toolbar](#), select Connect.
1. 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.



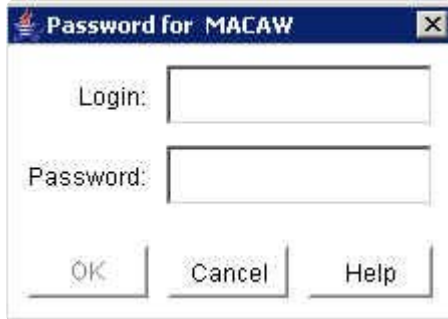
3. In the **Login** and **Password** fields, enter the login name and password of a user with SIOS Protection Suite authorization on the specified server.
4. Click **OK**.

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.

Password Dialog

* **Note:** 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.

Disconnecting From a Cluster

This task disconnects your client from all servers in a cluster.

1. Select a server from which you want to disconnect and then select **Disconnect** from the [Server Context Menu](#) or [Server Context Toolbar](#).
2. A **Confirmation dialog** listing all servers in the cluster is displayed. Click **OK** in the **Confirmation dialog** to confirm that you want to disconnect from all servers in the cluster.

After disconnecting from a cluster, all servers in that cluster are removed from the [Status Table](#).

Viewing Connected Servers

The state of a server can be determined by looking at the graphic representation of the server in the GUI as shown below. See [Viewing the Status of a Server](#) for an explanation of the server states indicated visually by the server icon.



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.



CARDINAL







BLUEJAY



MACAW



SCARLET

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 SIOS Protection Suite state is ALIVE.

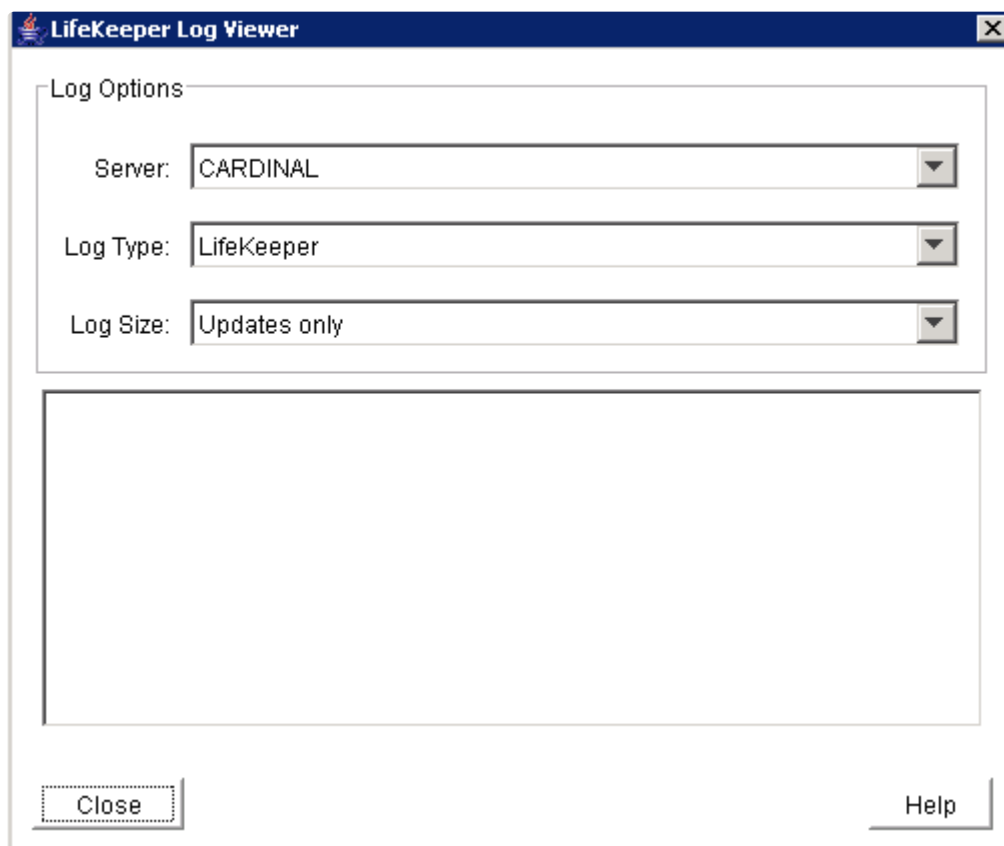
Viewing Server Log Files

To view server log files:

1. Select a server and then select **View Logs** from the [Server Context Menu](#) or [Server Context Toolbar](#). This will bring up the [Log Viewer Dialog](#).
2. When you are finished, click **OK** to close the dialog.

Log Viewer Dialog

The Log Viewer dialog may be accessed from the server context menu or the toolbar. This dialog displays a limited view of the log files maintained by LifeKeeper. When accessed from the toolbar, you can look at log files for any server by changing the selected server in the Server list.



Server. A drop-down list of servers connected to the cluster. Select the server whose log files you want to view. This list is not available if the dialog is invoked from the [server context menu](#).

Log Type. A drop-down list of log files available on the selected server. Select one of the options to indicate which log files you want to see:

- LifeKeeper

Log Size. A drop-down list containing four options. Select one of the options to indicate how much of the log file you want to see:

- Updates Only
- Last 100 lines
- Last 500 lines
- Last 1000 lines

Viewing Server Properties

To view server properties:

- If the [Properties Panel](#) is enabled, simply select the server in the [Status Table](#) and the properties will be displayed in the Properties Panel.
- If the Properties Panel is disabled, select the server and then select **Properties** in the [Server Context Menu](#).

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:





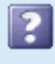
```
Resource Tag = F-Drive, Resource ID = F:
```

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.





Viewing the Status of Resources

The status or state of a resource is displayed in two formats: **global resource status** (across all servers) and **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

Server State	Visual State	What it Means
Active		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		Backup resource that can take over operation from the active resource (OSU)
Failed		Problem with resource detected on this server. For example, an attempt to bring the resource in-service failed (OSF)
UNKNOWN		Resource has not been initialized (ILLSTATE) or SIOS Protection Suite is not running on this server
	Empty Panel	Server does not have the resource defined

Global Resource Status

Description	Visual State	What it Means / Causes
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 Cannot determine state from		More than one server is claiming to be active Lost connection to server

available information		All server resource instances are in an unknown state
-----------------------	--	---

Viewing Resource Properties

To view resource properties:

- If the [Properties Panel](#) is enabled, simply select the server-specific resource instance in the [Status Table](#) and the properties will be displayed in the **Properties Panel**.
- If the Properties Panel is disabled, select the server-specific resource instance and then select **Properties** in the [Resource Context Menu](#).

Viewing Message History

1. On the [View Menu](#), click **History**. The **Message History** dialog is displayed (see below).
2. If you want to clear all messages from the history, click **Clear**.
3. Click **OK** to close the dialog box.

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"
```

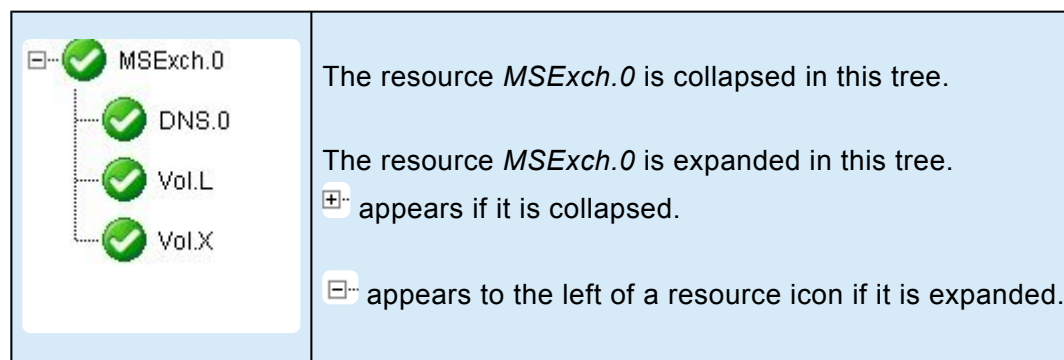
```
—>"server name":"app res": "action"
```

```
—>"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.

Expanding and Collapsing A Resource Hierarchy Tree



To expand a resource hierarchy,

- click the or
- double-click the resource icon to the right of a .

To expand all resource hierarchies,

- On the [View Menu](#), click **Expand Tree** or
- Double-click the **Hierarchies** button in the top left corner of the [Status Table](#).

Note: The resource tag/ID shown in the resource hierarchy 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,

- click the or
- double-click the resource icon to the right of a .

To collapse all resource hierarchies,

- On the [View Menu](#), click **Collapse Tree**, or
- Double-click the **Hierarchies** button in the top left corner of the [Status Table](#).

Operator Tasks

This section covers more advanced tasks that require Operator permission such as bringing resources in and out of service.

[Bringing A Resource In Service](#)

[Taking a Resource Out Of Service](#)

[Taking Volume Resources In and Out Of Service](#)

[Volume Shadow Copy](#)

[Volume Locking for Shared SCSI Volumes](#)

Bringing A Resource In Service

To bring a resource in service:

1. Select a server-specific resource instance that you want to bring in service and then select **In Service** from the [Resource Context Menu](#) or [Resource Context Toolbar](#).
2. A dialog appears confirming the server and resource(s) that you have selected to bring into service. This dialog will include a warning if you are bringing a dependent child resource into service without bringing its parent resource into service as well. 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.
4. Errors that occur while bringing a resource in service are logged in both the LifeKeeper log and the GUI log of the server on which you want to bring the resource into service.

Taking a Resource Out Of Service

To take a resource out of service:

1. Select a server-specific resource instance that you want to take out of service and then select **Out of Service** from the [Resource Context Menu](#) or [Resource Context Toolbar](#).
2. A dialog appears confirming the selected resource(s) to be taken out of service. This dialog will include a warning if you are taking a dependent child resource out of service without taking its parent resource out of service as well. Click **Out of Service** to take the resources out of service.
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 both the LifeKeeper log and the GUI log of the server on which you want to take the resource out of service.

Taking Volume Resources In and Out Of Service

Some background processes such as virus scanners and Windows Services will try to access a shared volume for write access. These processes may be required to run applications and cannot be arbitrarily turned off for long periods of time.

In most cases, these applications will have no effect on the manual switchover of SIOS Protection Suite volumes. However, if you see error messages similar to the following during a manual switchover (actual failovers are not affected), you should use the Volume Remove Stop/Restart feature described below.

```
*ERROR* [No. 12035] Unable to lock volume <volume ID> on <system name>
machine at this time as it may be in use by some application. Please free
this volume and try again.
```

If the volume has users doing non-write access only (remote links, local opens), removing the volume from service succeeds as does restoring the volume on the other system (for example, manual switchover). The existing user “opens” are, of course, no longer valid. However, they prevent the volume from being restored to service on the original system whether it is manually switched back over or automatically failed back over to the original system. Attempts to do so result in the following error message:

```
*ERROR* [No. 12046] LifeKeeper RESTORE VOLUME <volume ID> FAILED(err=<error
number>) .
```

The end result is that removing or restoring a volume, switching a volume over or back or switching any hierarchy that includes a volume resource over or back fails if the volume has any users, local or remote.

In addition, the system's PATH variable must not contain any file shares that exist on the volume to be protected by SIOS Protection Suite. File shares in the PATH variable may also cause volume operations to fail. Remove any such shares prior to creating the volume resource in SIOS Protection Suite. The PATH variable may be modified by selecting **System** then **Environment** from the **Windows Control Panel**.

Volume Remove Stop/Restart Programs and Services

SIOS Protection Suite provides configurable registry keys that permit the user to specify programs and services to be stopped and restarted on a system during a volume remove operation. This feature starts and stops user specified programs and services when SIOS Protection Suite detects open handles on any

volume that is being removed and taken out of service on that system. However, if SIOS Protection Suite detects no open handles on the volume being removed from service, it will not attempt to stop anything on that system. Instructions for using this feature are as follows:

1. Determine which programs and services are accessing the volume in question and preventing successful volume failovers.
2. Specify programs to stop and start by adding a subkey for each program under:

- **64-Bit:** `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\SIOS\LifeKeeper\VolumeRemoveStopPrograms*`

For example, to stop a program called “myapp.exe”, add the following subkey:

`HKEY_LOCAL_MACHINE\SOFTWARE\SIOS\LifeKeeper\VolumeRemoveStopPrograms\myapp.exe\`

3. Under the subkey for each program, add the following values (all are REG_SZ values):

ProgramName	the program name only (Note: Must match subkey name from Step 2.)
ProgramPath	the program name, including full path
Restart	a value of 0 indicates that the program should not be restarted; a value of 1 indicates that the program should be restarted
StartCmdLine	optional command line arguments to be used when starting the program
WasRunning	used by SIOS Protection Suite to save the number of instances of the program that were running before stopping them; should be initialized to 0

For example, values could be entered to stop and start “myapp.exe /a /t /p” with its associated arguments as follows:

ProgramName	myapp.exe
ProgramPath	C:\mydir\myapp.exe
Restart	1
StartCmdLine	/a /t /p
WasRunning	0

Note: By default, SIOS Protection Suite includes a subkey for stopping perfmon.exe and the restart option is disabled.

1. Specify services to stop and restart by adding a subkey for each service under:

- **64-Bit:** `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\SIOS\LifeKeeper\VolumeRemoveStopServices\`

For example, to stop a service called “mysvc”, you would add the following subkey:

`HKEY_LOCAL_MACHINE\SOFTWARE\SIOS\LifeKeeper\VolumeRemoveStopServices\mysvc\`

2. Under the subkey for each service, add the following values (all are REG_SZ values):

ServiceName	The service display name (Note: Must match subkey name from Step 4.)
Restart	A value of 0 indicates that the program should not be restarted; a value of 1 indicates that the program should be restarted
WasRunning	Used by SIOS Protection Suite to save the number of instances of the program that were running before stopping them; should be initialized to 0
StopWait	The number of seconds to wait for the service to reach the STOPPED state. If StopWait is negative, then it (and the failover) will wait indefinitely for the service to reach the STOPPED state
StopWait	The number of seconds to wait for the service to reach the RUNNING state. If the service does not reach the RUNNING state in the configured period, an error is logged in the Event Log. If StartWait is negative, then the failover will wait indefinitely for the service to start. If it is set to 0, then the service will be started but will not wait for it to reach the RUNNING state (and thus doesn't generate an Event Log message if the service can't be restarted)

For example, values could be entered to stop and start “mysvc” as follows:

ServiceName	mysvc
Restart	1
WasRunning	0
StopWait	120
StopWait	120

Volume Restore Stop/Restart Programs and Services

SIOS Protection Suite provides configurable registry keys that permit the user to specify programs and services to be stopped and restarted on a system during a volume restore operation. This feature is similar to the Volume Remove Stop/Restart Programs and Services feature described above with the following

differences:

- Programs to stop and restart are specified as subkeys under:
 - **64 Bit:** `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\SIOS\LifeKeeper\VolumeStopPrograms\`
- Services to stop and restart are specified as subkeys under:
 - **64 Bit:** `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\SIOS\LifeKeeper\VolumeStopServices\`

Volume Shadow Copy (VSS)

Using Volume Shadow Copy (VSS) with DataKeeper/SIOS Protection Suite Volumes

On Windows 2008 R2, VSS Shadow Copy can be enabled for SIOS Protection Suite-protected (shared or replicated) volumes. However, the following guidelines apply:

- VSS snapshot images must not be stored on a SIOS Protection Suite-protected volume. Storing VSS snapshots on a SIOS Protection Suite-protected volume will prevent SIOS Protection Suite from being able to lock the volume and switch it over to another node.
- When a SIOS Protection Suite-protected volume is switched or failed over, any previous snapshots that were taken of the SIOS Protection Suite protected volume are discarded and cannot be reused.
- VSS snapshot scheduling is not copied between the SIOS Protection Suite servers. If snapshots are scheduled to be taken twice a day on the primary server and a switchover occurs, this schedule will not be present on the backup server and will need to be redefined on the backup server.
- There is a slight difference in behavior when switching back to a server where snapshots were previously enabled:
 - If the volume is a shared volume, VSS snapshots must be re-enabled.
 - If the volume is a replicated volume, VSS snapshots are automatically re-enabled.

Volume Locking for Shared SCSI Volumes

When you want to protect resources on shared SCSI disks, you partition the shared disk into logical volumes using the **Windows Disk Management** tool. SIOS Protection Suite can protect shared volumes by defining a volume resource instance. Each instance is assigned a drive letter (for example, G:).

SIOS Protection Suite brings the volume resource instance into service on the primary server and provides software locks so that a backup server cannot access the volume while it is active on the primary server. In case of a failure of the primary server, SIOS Protection Suite automatically brings the volume resource into service on the backup server and locks the primary server from accessing the volume resource when it is repaired.

SIOS Protection Suite also automatically changes the primary and designations so that the failed server is now locked from access to the volume resource. In this way, the resource is protected from inappropriate access while you repair the failed server.

This dynamic redefinition of primary and backup servers is SIOS Protection Suite's intelligent switchback feature that allows you to select the appropriate time to bring the resource back into service on the repaired system.

Since SIOS Protection Suite maintains the volume locks, do not stop SIOS Protection Suite, as this would disable the locks.

Advanced Topics

This section details information on the LifeKeeper Configuration Database and LifeKeeper Communications.

[LifeKeeper Configuration Database](#)

[LCD Directory Structure](#)

[Diagram of LCD Directory](#)

[LCD Configuration Data](#)

[LCD Resource Types](#)

[Resources Subdirectories](#)

[LCDI Commands](#)

[LifeKeeper Communications Manager](#)

[Communication Status Information](#)

LifeKeeper Configuration Database (LCD)

The LifeKeeper configuration database (LCD) maintains the object-oriented resource hierarchy information and stores recovery direction information for all resource types known to SIOS Protection Suite. 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 topics for information on the LCD directory structure, types of data stored, resource types available and use of application scripts:

- [LCD Directory Structure](#)
- [Diagram of LCD Directory in \\$lkroot/LifeKeeper](#)
- [LCD Configuration Data](#)

LCD Directory Structure

Major subdirectories under `$LKROOT` (by default, `c:\LK`):

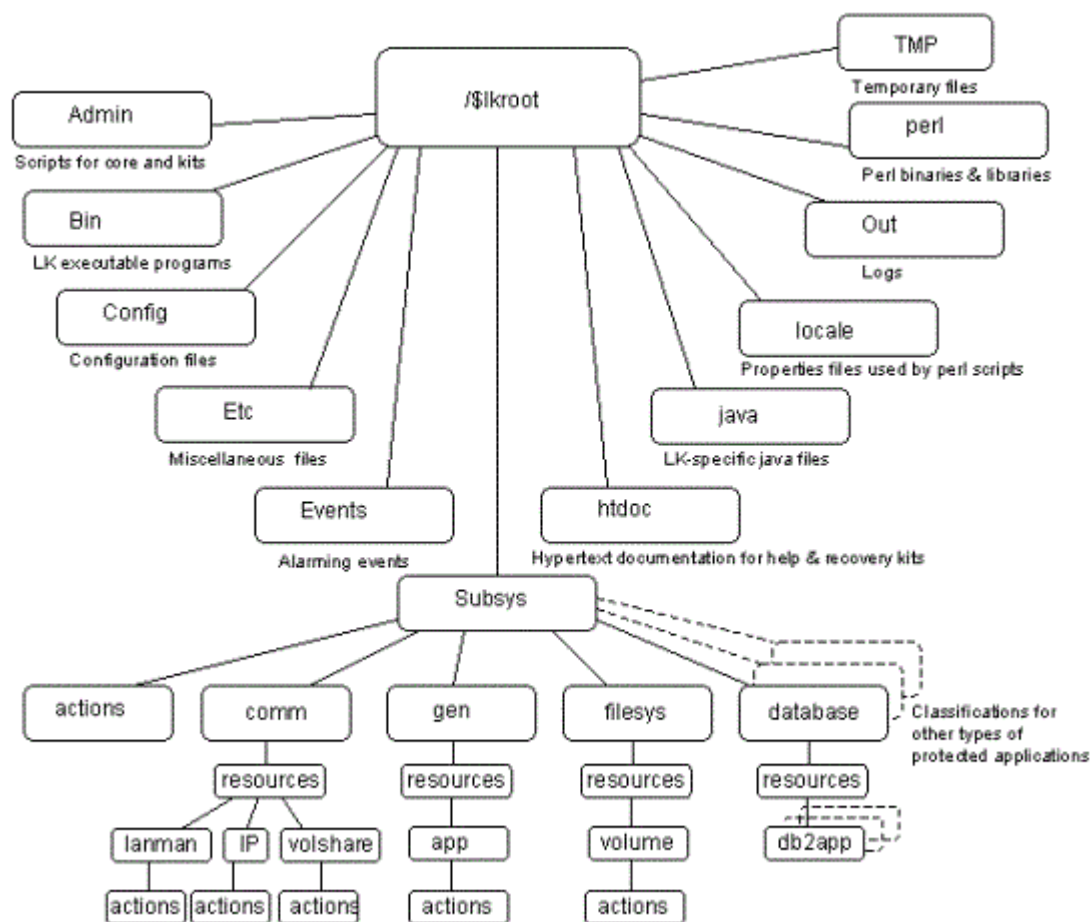
- **Admin.** Scripts for SIOS Protection Suite core and Recovery Kits.
- **Config.** SIOS Protection Suite configuration files, including shared equivalencies.
- **Bin.** SIOS Protection Suite executable programs.
- **Subsys.** Resources and types. SIOS Protection Suite provides resource and type definitions in subdirectories of Subsys. For instance, communications resources are stored in `comm`, volume resources are stored in `filesys`, and generic application resources are stored in `gen`. Optional Recovery Kits may create different resource types stored in different directories. For example, database application resources are stored in `database`.
- **Events.** Event alarms.
- **Out.** LifeKeeper logs. SIOS Protection Suite sends a variety of error and status messages to several different logs in this directory.
- **perl.** Perl binary executables and libraries.

The structure of the LCD directory in `$LKROOT` is shown in the topic [Diagram of LCD Directory](#).

Note: The location of these subdirectories can be changed by modifying the value of `LKROOT` in the environment.

Diagram of LCD Directory

The following diagram shows the directory structure of \\$.lkroot.



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, SIOS Protection Suite maintains a list of dependencies and a list of dependents (resources depending on a resource). For more information, see the [LCDI_relationship](#) and [LCDI_instances](#) manual pages.

Resource Status Information

LCD maintains status information in memory for each resource instance. The [resource states](#) recognized by LCD are **ISP**, **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 represent the same physical entity. When two servers have a resource with a shared equivalency relationship, SIOS Protection Suite 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.

Furthermore, the dependency relationships within a hierarchy guarantee that all resources that depend upon the volume, such as a file share, are in service on only one server at a time.

LCD Resource Types

The LCD is maintained in both shared memory and in the \$LKROOT directory. As highlighted on the [directory structure diagram](#), subsys contains application resource sets you can use to define your application interface:

- filesys – file system related resources like volume
- comm – communications related resources like IP, volshare (fileshare) and lanman
- database – database resources such as Oracle

These subdirectories are discussed in [Resources Subdirectories](#).

Resources Subdirectories

The *filesys*, *comm*, *WebServer*, *database*, *mail* and *appsuite* directories each contain a *resources* subdirectory. The content of those directories provides a list of the resource types that are currently defined and managed by SIOS Protection Suite:

- **filesys resource types.** You find these resource types in the `$LKROOT\LifeKeeper\subsys\filesys\resources` directory:
 - **volume** — disk partitions or virtual disk devices
- **comm resource types.** You find these resource types in the `/$LKROOT/LifeKeeper/subsys/comm/resources` directory:
 - **IP** — created by the IP Recovery Kit
 - **DNS** — created by the DNS Recovery Kit
 - **volshare** — fileshare resources created by the LAN Manager Recovery Kit
 - **lanman** — computer alias created by the LAN Manager Recovery Kit
- **WebServer resource types.** You find these resource types in the `$LKROOT\LifeKeeper\subsys\WebServer\resources` directory:
 - **IIS** — created by the IIS Recovery Kit
- **database resource types.** You find these resource types in the `$LKROOT\LifeKeeper\subsys\database\resources` directory:
 - **Microsoft SQL Server**
 - **PostgreSQL**

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.
- **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](#) man page.

LCDI Commands

SIOS Protection Suite 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 SIOS Protection Suite 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.

LifeKeeper Communications Manager (LCM)

The LifeKeeper Communication Manager (LCM) provides reliable communication between processes on one or more SIOS Protection Suite servers. This process can use redundant communication paths between systems so that failure of a single communication path does not cause failure of SIOS Protection Suite or its protected resources. The LCM supports a variety of communication alternatives including TCP/IP and shared disk connections.

The LCM provides the following:

- **SIOS Protection Suite Heartbeat.** Periodic communication with other connected SIOS Protection Suite systems to determine if the other systems are still functioning. SIOS Protection Suite 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 SIOS Protection Suite 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 facilities allow the LCD to maintain consistent resource information between the primary and secondary systems.
- **Failover Recovery.** If a resource fails on a system, the LCM notifies SIOS Protection Suite to recover the resource on a backup system.

Communication Status Information

The **Communication** tab of the [Server Properties dialog](#) lists the servers known to SIOS Protection Suite and their current state followed by information about each communication path.

Server Properties – CommPath

Server Properties for CARDINAL

Properties Panel

Server: CARDINAL

General CommPaths Resources

Server	Priority	State	Type	Address/Device
BLUEJAY	1	ALIVE	TCP	172.17.108.132/172.17.1...
BLUEJAY	2	ALIVE	TCP	10.10.1.1/10.10.1.2

OK Apply Cancel Help

Maintenance Tasks

The topics in this section cover the tasks necessary for maintaining SIOS Protection Suite.

[Starting and Stopping LifeKeeper](#)

[Managing IP Resources](#)

[Managing DNS Resources](#)

[Displaying List of Protected File Shares](#)

[EditFileShareResource Utility](#)

[Transferring Resource Hierarchies](#)

[Performing Offline Maintenance On A Shared Disk](#)

[Maintaining a SIOS Protection Suite Protected System](#)

[Configuring Generic Application Scripts](#)

[Maintaining a Resource Hierarchy](#)

[Recovering After a Failover](#)

[Uninstalling SPS for Windows](#)

[Performing CHKDSK on a protected volume](#)

Starting and Stopping LifeKeeper

Because LifeKeeper is typically started automatically after installation and each time the server is booted, you should not normally need to start/stop LifeKeeper. (The only exception is if you chose to do a Custom installation and opted not to start LifeKeeper at that time.)

In the event that you need to start or stop LifeKeeper manually, you should do so using the **Services** tool under **Administrative Tasks** in the Windows Control Panel.

Starting LifeKeeper

LifeKeeper consists of two services:

- LifeKeeper
- LifeKeeper External Interfaces

Generally, these two services should be stopped and started together. However, since LifeKeeper External Interfaces is a dependency of the LifeKeeper service, stopping it will also stop the LifeKeeper service. Likewise, it must be started before the LifeKeeper service can be started.

Select **LifeKeeper** and click **Start**. This will automatically start the **LifeKeeper External Interfaces** service.

Stopping LifeKeeper

In the **Services** tool, select **LifeKeeper External Interfaces** and click **Stop**. This will stop both services. Note that the length of time that it takes to stop LifeKeeper will vary depending upon the hierarchies currently configured although the Services tool shows the services as stopped immediately.


Using the command line to enter `$LKROOT\bin\lkstop` will more accurately show the services being stopped, and it will confirm with the message **"LIFEKEEPER NOW STOPPED"**.



Note: Stopping LifeKeeper takes all protected hierarchies out of service. This means that any protected applications will not be accessible.

Managing IP Resources

To view configuration information associated with a protected IP resource from the LifeKeeper GUI, right-click on the IP resource (on the right-hand side of the LifeKeeper GUI) and select **Properties**, then select the **IP Configuration** tab. The example below shows the configuration details for the SIOS Protection Suite protected IP resource 172.17.100.118.



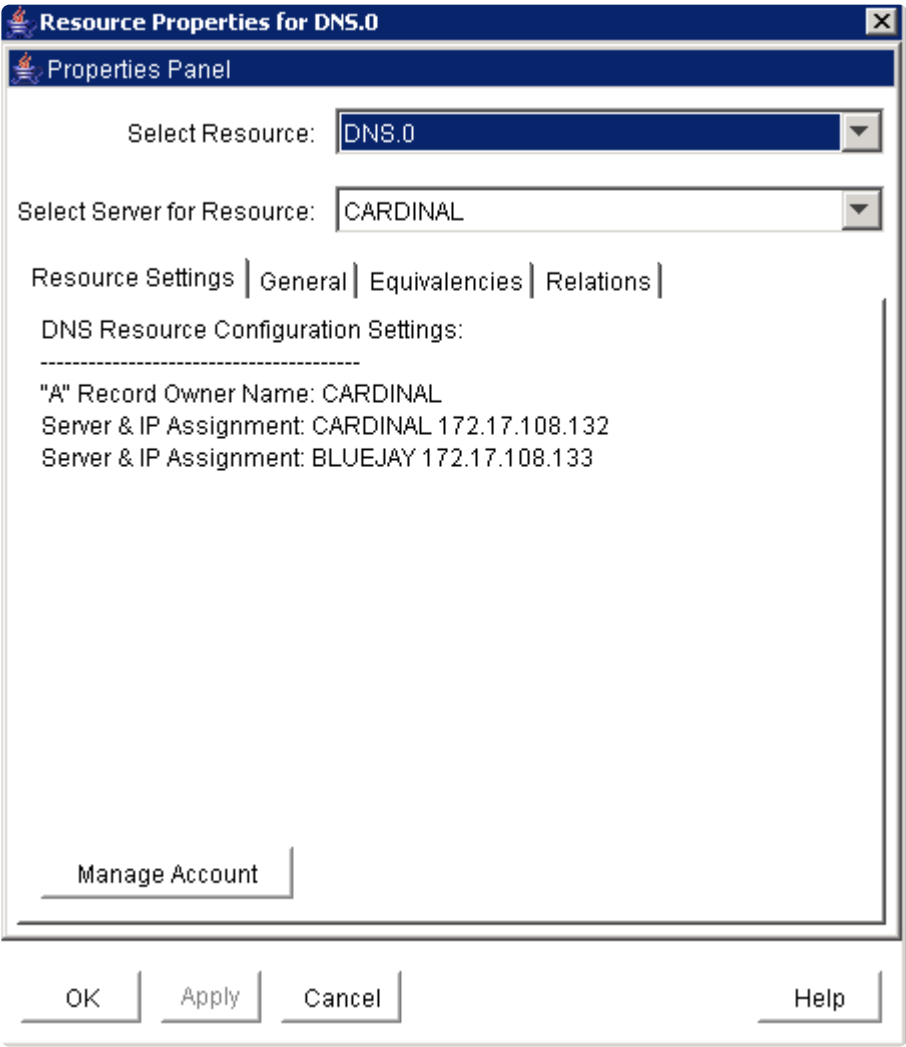
To enable or disable the IP address restore capabilities on the selected server while still allowing the SIOS Protection Suite IP resource to report a successful in-service operation, click the **Modify** button, then select **Enable** or **Disable** for the restore mode. This feature applies to three-node SIOS Protection Suite clusters where two nodes are on a LAN (same subnet) and the third node is on a WAN (different subnet). The restore mode of the IP resource would be enabled on the LAN nodes and disabled on the WAN node.

Pinglist functionality can be enabled by creating a file containing valid IP address. This function will attempt to ping the addresses in the file. It will succeed if any of the IP addresses respond to the ping, and will fail if none of them respond. The pinglist check will be executed at the end of an IP Restore operation, as well as during an IP Quick Check. The pinglist is optional – if no file is found, Restore and Quick Check skip that test.

The file required to enable this feature needs to be placed in %LKROOT%/Subsys/comm/resources/ip/, and should be named “pinglist.<tag>” where <tag> is the name given to the protected IP resource. IP addresses should be entered one per line in the file. The IP addresses used should also be reachable from the protected IP address’s subnet, and should respond to ICMP PING requests.

Managing DNS Resources

To change the Domain administrative user and password associated with a protected DNS resource from the LifeKeeper GUI, right-click on the DNS resource (on the right-hand side of the LifeKeeper GUI) and select **Properties**, then select the **Resource Settings** tab. Select **Manage Account** on the Resource Settings page to change the Domain administrative user and password for your DNS resource.



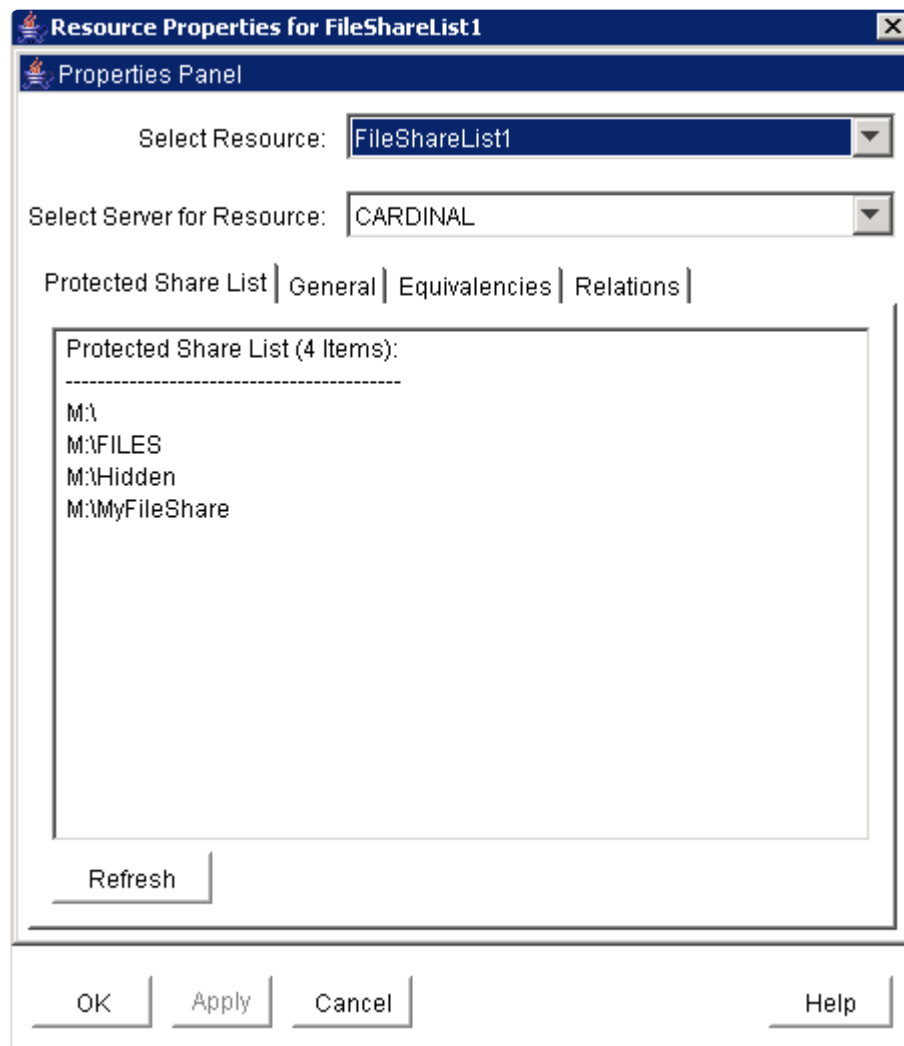
Manage Account :

Field	Tips
Enter User ID (Domain\)	Enter the user name of the Windows DNS/Domain administrator. This user account should have privileges to make changes in the DNS configuration and should be a member of the "Domain Admins" group in the same domain as the DNS server. Enter the user ID in

UserID)	<DomainName>\<UserID> format where <DomainName> is the NetBIOS name of the domain.
Enter Password	Enter the password for the account previously entered.

Displaying List of Protected File Shares

To display the list of file shares associated with a protected file share resource from the LifeKeeper GUI, right-click on the File Share resource (on the right-hand side of the LifeKeeper GUI) and select **Properties**, then select the **Protected Share List** tab.



EditFileShareResource Utility

The EditFileShareResource utility can be used to update a file share resource with all current file shares on the associated volume(s). This can be useful in environments where there are a large number of file shares and file shares have been added or deleted since the resource was created. Using the utility can prevent the need to delete and re-create the file share resource.

To invoke the utility, on the command line enter:

```
EditFileShareResource <Tag name>
```

where <Tag name> is the tag name of a file share resource that is currently in service.

The utility protects **all eligible file shares** defined on the protected volumes that are associated with the file share hierarchy. It deletes any previously protected shares that have been deleted from the system and adds newly defined shares (meeting the eligibility criteria) to the list. It will also update the file share permissions defined on the file share.

Transferring Resource Hierarchies

When you need to perform routine maintenance or other tasks on a SIOS Protection Suite 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 SIOS Protection Suite protection depending on the number of servers in the cluster.

Performing Offline Maintenance On A Shared Disk

When performing offline maintenance on a shared SCSI host adapter or a disk on a shared bus, you must stop LifeKeeper and power down all servers and shared disks. Perform these actions in the following order:

1. **Stop LifeKeeper.** Use the **Services** tool to stop the LifeKeeper and LifeKeeper External Interfaces services on each SIOS Protection Suite server. Your resources are now unprotected.
2. **Shut down Windows.** Shut down the Windows operating system on all servers in the cluster.
3. **Power down all servers.**
4. **Power OFF all shared disks.**
5. **Perform maintenance.** Perform the necessary maintenance on the shared SCSI host adapter or shared disk.
6. **Power ON all shared disks.**
7. **Power ON all servers, one at a time.** Let each server boot the Windows operating system completely before powering on the next server.
8. **Start LifeKeeper.** Log on as administrator, then use the **Services** tool to start the LifeKeeper and LifeKeeper External Interfaces services on each SIOS Protection Suite server. SIOS Protection Suite automatically mounts all shared file systems and restarts and brings into service all databases on shared disks.

Maintaining a SIOS Protection Suite Protected System

When performing shutdown and maintenance on a SIOS Protection Suite-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. For off-line maintenance of a shared disk, see [Off-Line Maintenance of a Shared Disk](#).

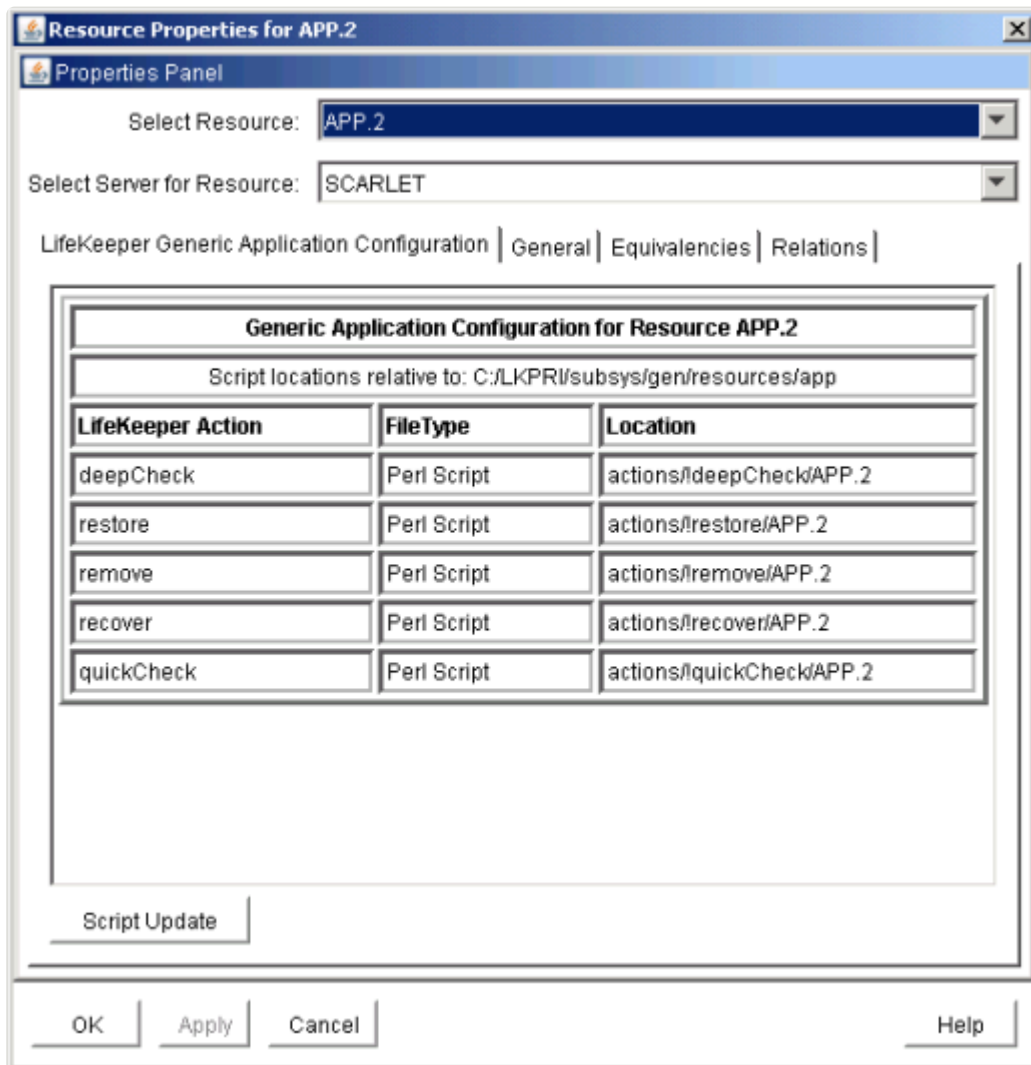
Perform these actions in the order specified where Server A is the primary system in need of maintenance and Server B is the backup server:

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 SIOS Protection Suite protection. See [Bringing a Resource In Service](#) for instructions.
2. **Stop LifeKeeper on Server A.** In the Services tool, select **LifeKeeper External Interfaces** and click **Stop**. This will stop both services. Your resources are now unprotected.
3. **Shut down Server A.** Shut down the Windows 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.** Power on Server A and bring up the Windows operating system.
6. **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.

Configuring Generic Application Scripts

Use this feature to update a script that has been created to protect an application that has no associated SIOS Protection Suite Recovery Kit.

1. Right-click on the generic application resource and select **Properties**. Select the **SIOS Protection Suite Generic Application Configuration** tab.



2. Select the **Script Update** button. Use the following table to complete the fields in the Generic Application Configuration procedure.

Field	Tips
-------	------

<p>Select the Action Script to Update</p>	<p>Select the SIOS Protection Suite action name for the resource that will be updated. Select:</p> <ul style="list-style-type: none"> • restore to update the script responsible for in-service operations. • remove to update the script responsible for out-of-service operations. • quickCheck to update the script responsible for monitoring the application. • deepCheck to update the script that performs in-depth monitoring of the application. • recover to update the script responsible for resource recover operations. • delete to update the script that performs any additional actions required to remove the application from SIOS Protection Suite protection. • extend to update the script responsible for additional actions required to prepare the application for protection with SIOS Protection Suite on the target server(s)
<p>Full Path to New Script</p>	<p>Enter the pathname for the shell script or object program for the application.</p> <ul style="list-style-type: none"> • The restore script is responsible for bringing a protected application resource in-service. (Required) • The remove script is responsible for bringing a protected application resource out-of-service. (Required) • The quickCheck script is responsible for monitoring a protected application resource after a failure event. • A copy of this script or program will be saved by SIOS Protection Suite in the resource hierarchy on the server. • There may be a short wait while SIOS Protection Suite validates the pathname to remove monitoring or recovery.

	<ul style="list-style-type: none">• Do not specify a shell script or object program.• Valid characters allowed in the script pathname are letters, digits and the following special characters: – _ ! . /
--	--

3. The **Basic File Statistics** dialog displays old and new configuration information about the current script. Click **Continue**.
4. The **Update All Systems** dialog displays. Select Yes to update all systems in this cluster. Select **No** to only update the current system. If you choose **No**, you must separately update the corresponding script for the configuration on the backup servers. Click **Next**.
5. Click **Done** to complete.

Maintaining a Resource Hierarchy

You can perform maintenance on a resource hierarchy while maintaining SIOS Protection Suite 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.

Recovering After a Failover

After SIOS Protection Suite performs a failover recovery from a primary server (ServerA) to a backup server (Server B), perform the following steps:

1. **Monitor failover.** When SIOS Protection Suite on Server B performs a failover recovery from Server A, status messages are displayed during the failover ending with the following message:

FAILOVER RECOVERY OF MACHINE *Server A*

FINISHED AT: *date time year*

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, go to the **Windows Services tool**, select **LifeKeeper** and click **Start**. This will automatically start the **LifeKeeper External Interfaces** service.
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.

Uninstalling SPS for Windows

Before Removing LifeKeeper

Included below are the requirements for removing LifeKeeper software.

1. **Move or stop applications.** Before removing the software, verify that applications requiring SIOS Protection Suite protection are not on the server. Never remove LifeKeeper from a server where an application resource hierarchy is in service. Removing LifeKeeper removes all configuration data, such as equivalencies, resource hierarchy definitions and log files. See [Transferring Resource Hierarchies](#) for additional information.
2. **Ensure LifeKeeper is running.** Recovery Kits may require LifeKeeper to be running when you remove the recovery kit software. Use the **Services MMC** snap-in to ensure that LifeKeeper services are running. If it is not running, the removal process cannot remove the resource instances from other SIOS Protection Suite servers in the cluster which would leave the servers in an inconsistent state.
3. **Remove resource hierarchies.** Unextend or delete any resource hierarchies from the server where LifeKeeper will be removed. Never remove a Recovery Kit from a server where the resource hierarchy is in service. This will corrupt current hierarchies and they will need to be recreated when reinstalling the Recovery Kit.
4. **Remove all packages.** If removing the LifeKeeper core, first remove other packages that depend upon LifeKeeper; for example, SIOS Protection Suite Recovery Kits. It is recommended that before removing a SIOS Protection Suite Recovery Kit, first remove the associated application resource hierarchy.

Before Removing DataKeeper

If planning to uninstall DataKeeper and reinstall a previous version, all jobs/mirrors must be deleted on each node prior to uninstalling. These will need to be recreated once software is reinstalled.

Uninstall SIOS Protection Suite

- In **Windows Control Panel**, find your list of installed programs and select **SIOS DataKeeper or LifeKeeper**.

- Select **Uninstall**.

Once the uninstall process is complete, rebooting the system is required.

Note: Uninstalling automatically stops the SIOS DataKeeper and/or LifeKeeper services and clears the registry entries.

Once removed, the following files will not be removed by the uninstall procedure.

Path and File Name	Definition and Special Considerations
<code><windows dir>/SysWOW64/LKLicense</code>	<p>Common license file directory for SIOS Technology Corp. products. This is where license files are installed and licenses for multiple SIOS Technology Corp. products may be installed here at any given time. We don't remove this at uninstall so as to not disturb the installed licenses.</p> <p>Safe to remove manually, but the license will need to be reinstalled if the software is reinstalled at a later time.</p>
<code><windows dir>/SysWOW64/PerfStringBackup.ini</code>	<p>A backup file created by Windows when new performance monitor counters are installed. This is created when we install the perfmon counters.</p> <p>This should probably be left alone since it is a file created by Windows itself.</p>
<code><windows dir>/inf/ExtMirr/0011/ExtMirrCounters.ini</code>	<p>This file describes the DataKeeper performance monitor counters. This file can be removed or left alone. It is not an executable.</p>

Notes

- **Important:** Uninstallation of SIOS Protection Suite software requires that the Microsoft Visual C++ 2008 Redistributable package be installed. Do not remove this package until SIOS Protection Suite has been uninstalled.
- **Modify** or **Repair** must be run from the SIOS Protection Suite setup program.
- Removal of SIOS Protection Suite may NOT delete the SIOS Protection Suite directory. This directory can be deleted manually after the **Add/Remove** operation is complete.

- A reboot of the system is required to completely remove SIOS Protection Suite remnants.

Performing CHKDSK on a protected volume

Disk volumes that are under LifeKeeper protection do not have automatic boot-time CHKDSK performed on them. Administrators should occasionally perform CHKDSK on their volumes manually. To do this, follow these instructions:

1. Take all resources that depend on the volume out of service. During CHKDSK, the volume will be locked and inaccessible, so any applications that are using the volume should be taken out of service.
2. Leave the volume resource in service – CHKDSK needs to be able to open the volume, so it should be left in-service and unlocked.
3. If the volume is a DataKeeper mirrored volume, we recommend that the mirror be put into a Paused state.
4. Put the volume resource in maintenance mode by running these commands from the command line on the node where the volume is being checked:
 - `cd %LKBIN%`
 - `flg_create -f Maintenance_<drvletter>`
 - `<drvletter>` is the volume drive letter – to put volume F: in maintenance mode, run the command: `flg_create -f Maintenance_F`

While the volume is in maintenance mode, volume check scripts will not perform health checks.

5. Check the volume and fix errors (for example, run “`chkdsk.exe /f F:`”).

When CHKDSK completes:

1. Take the volume resource out of maintenance mode by running these commands from the command line on the node where the chkdsk completed.
 - `cd %LKBIN%`
 - `flg_remove -f Maintenance_<drvletter>`
 - `<drvletter>` is the volume drive letter – to take volume F: out of maintenance mode, run the command: `flg_remove -f Maintenance_F`
2. Continue any mirrors that were paused to facilitate performing chkdsk. This allows filesystem changes to be resync'ed with target nodes.

3. Bring hierarchies in service. These can be brought online while the mirror is still resyncing.

Data Replication

The topics in this section provide details on using data replication with SIOS Protection Suite.

[Monitoring Replicated Volume Resources](#)

[Replication Settings](#)

[Performing Actions on Replicated Volumes](#)

[What is Split-Brain](#)

[Split-Brain Recovery](#)

Monitoring Replicated Volume Resources




The state of all SIOS Protection Suite protected replicated volume resources is displayed in the LifeKeeper GUI. Refer to the SIOS DataKeeper topic, [Mirror State Definitions](#), for details on mirror states.

The example below shows that the mirror state of the replicated volume resource Vol.L is **Resync** and that the mirror state of the replicated volume resource Vol.Y is **Mirroring**.


Hierarchies					
Active Protected		CARDINAL		BLUEJAY	
MSEXCH.0		Active	1	StandBy	10
172.17.108.130		Active	1	StandBy	10
Vol.L		In Service	1	Resync	10
Vol.Y		In Service	1	Mirroring	10

The table below describes the different states for replicated volume resources and their meaning.

Resource State	Visual State	What it Means
Active		Resource is operational on the primary server and protected (ISP).
Degraded		Resource is operational on the primary server but not protected by a backup resource (ISU).
Unknown		Resource has not been initialized (ILLSTATE) or SIOS Protection Suite is not running on this server.
Failed		Problem with resource detected on this server. For example, an attempt to bring the resource in-service failed (OSF).
Offline		Resource is out of service on this server. Volume is not accessible for read/write operations.
Resync Pending		Resource state on the backup server is Resync Pending .
Mirroring		Resource state on the backup server is Mirroring .

Paused		Resource state on the backup server is Paused .
Resync		Resource state on the backup server is Resync .
Broken		Resource state on the backup server is Broken .

To view the configuration information for a replicated volume resource from the LifeKeeper GUI, right-click on the volume resource and select **Properties**, then select the **Mirror Status** tab. The example below shows Vol.G is the source on CAE-QA-V100 and has one target – 10.200.8.213, and that it is resyncing to the target CAE-QA-V213.



Replication Settings





From the **Volume Resource Properties** page, select the **Replication Settings** button to set the compression level, the network throttling or the SIOS Protection Suite Delete Mirror flag for a replicated volume.


Field	Tips
Select Targets	Select the target server to which the action should be applied.
Set Compression Level	<p>Specify the compression level for the selected replicated volume.</p> <p>Valid values are 0 to 9. Level 0 is “no compression”. Values from 1 to 9 specify increasingly CPU-intensive levels of compression. Compression level 1 is a “fast” compression – it does not require as much CPU time to compress the data but results in larger (less compressed) network packets. Level 9 is the maximum amount of compression – it results in the smallest network packets but requires the most CPU time. The level can be set to somewhere in between to balance CPU usage and network efficiency based on your system, network and workload.</p> <p>Default is 0.</p>
Set Network Throttling	<p>The Bandwidth Throttle setting (specified in kilobits per second) limits the amount of network bandwidth that the replicated volume can use for resync and normal volume writes.</p> <p>Default is 0.</p>
Set SIOS Protection Suite Delete Mirror Flag	<p>The SIOS Protection Suite Delete Mirror Flag controls the behavior during delete of the SIOS Protection Suite resource for the replicated volume. When deleting the SIOS Protection Suite volume resource, if the flag is set to True, then SIOS Protection Suite will delete the mirror; otherwise, the mirror will remain.</p> <p>Select True if you want the mirror deleted when the volume resource is unextended or removed from SIOS Protection Suite.</p>

	<p>Select False if you want the mirror to remain intact.</p> <p>Default is True if mirror is created using LifeKeeper GUI. Default is False if mirror is created outside of LifeKeeper GUI.</p>
--	--

Performing Actions on Replicated Volumes

To perform actions on a replicated volume resource using the LifeKeeper GUI, right-click on the replicated volume resource and select the action you wish to perform from the context menu. If you have the **Properties Panel** enabled (View->Properties Panel), the resource toolbar will be displayed for the selected volume.

Action	Icon	Meaning
Select Target		Select the target system to which the action should be applied.
Pause Mirror		Select Pause Mirror to temporarily stop the data from being mirrored. A partial resync will be performed when you click Continue to un-pause the mirror. After pausing a mirror, it is possible to unlock the target volume using the Unlock Target action.
Continue Mirror/ Lock Target		<p>To continue a mirror after the mirror has been paused, select the Continue/Lock Target action. This un-pauses the mirror, re-locks the target volume (if unlocked) and resumes the mirroring process.</p> <p>While pause temporarily stops the writes from being mirrored, the writes are recorded during the pause interval. When the mirror is resumed, the recorded writes are sent to the target volume and the mirror is automatically re-synchronized (partial resync).</p>
Unlock Target		<p>To unlock the target volume of a mirror, select the Unlock Target action. This pauses the mirror (if not already paused) and unlocks the mirrored volume on the target system. This allows read/write access to the data on the volume.</p> <p>Continue Mirror will relock the target volume, perform a partial resync and resume the mirroring process.</p> <p>Warning: Any writes to the target volume while unlocked are lost when the mirror is resynchronized.</p>
Break Mirror/ Unlock Target		Breaking a mirror discontinues the mirror for the selected volumes and unlocks the target volume but does not remove the mirror from the volume list. A full resync must be

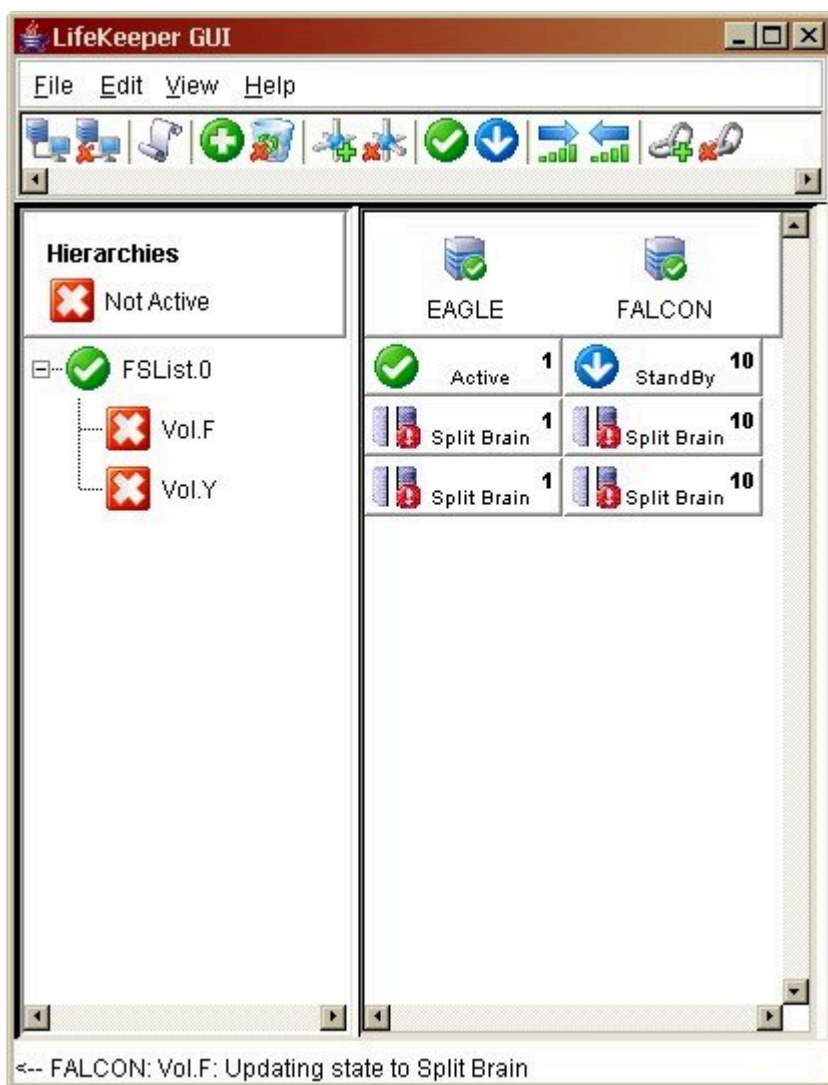
		<p>performed in order to re-establish the mirror after a Break Mirror/Unlock Target action.</p> <p>Warning: Any writes to the target volume while unlocked are lost when the mirror is resynchronized.</p>
Resync Mirror/ Lock Target		<p>To re-establish a broken mirror, select Resync Mirror/Lock Target action. A full resync will be performed.</p>

What is Split-Brain

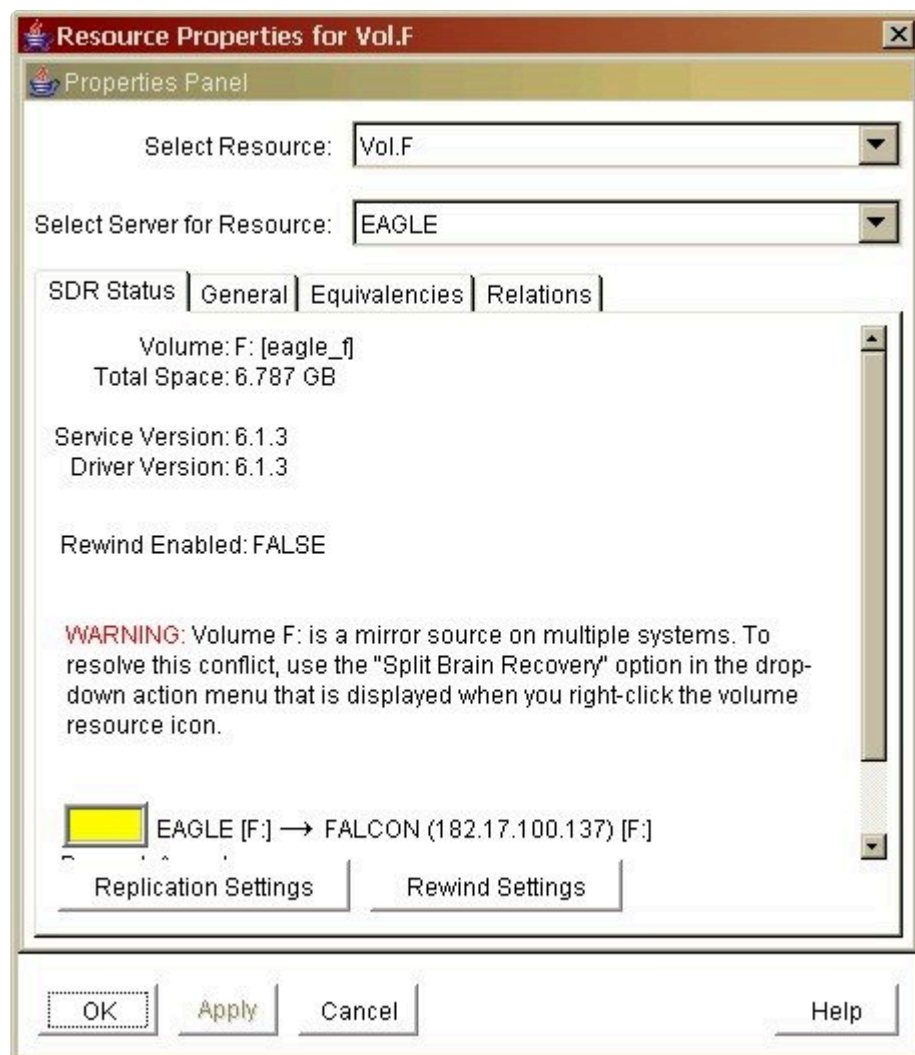
When all of SIOS Protection Suite's comm paths are disconnected and if **Automatic Node Failover** is enabled, each side of SIOS Protection Suite assumes that the other side is dead and attempts to bring all the resources in service. In the case of a SIOS DataKeeper resource, both sides become mirror sources and allow data to be written to the volume. This condition is defined as "split-brain" and will be indicated in the LifeKeeper GUI with the following icon:



Refer to the topic [Split-Brain Recovery](#) for the steps required to resolve this situation.



The **Properties Panel** for the selected volume displays additional information about the split-brain condition and instructions for resolving this problem.

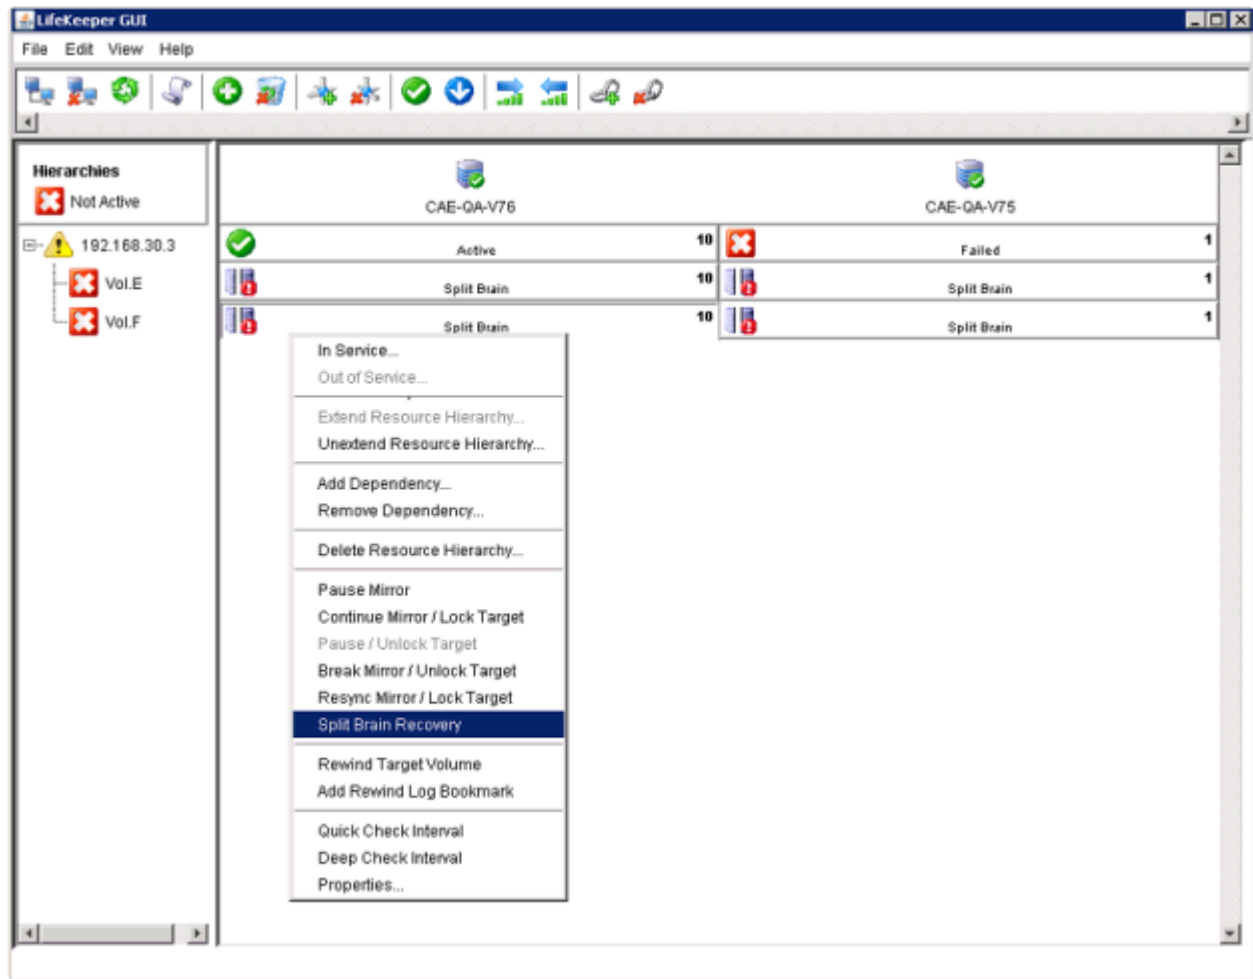


Split-Brain Recovery

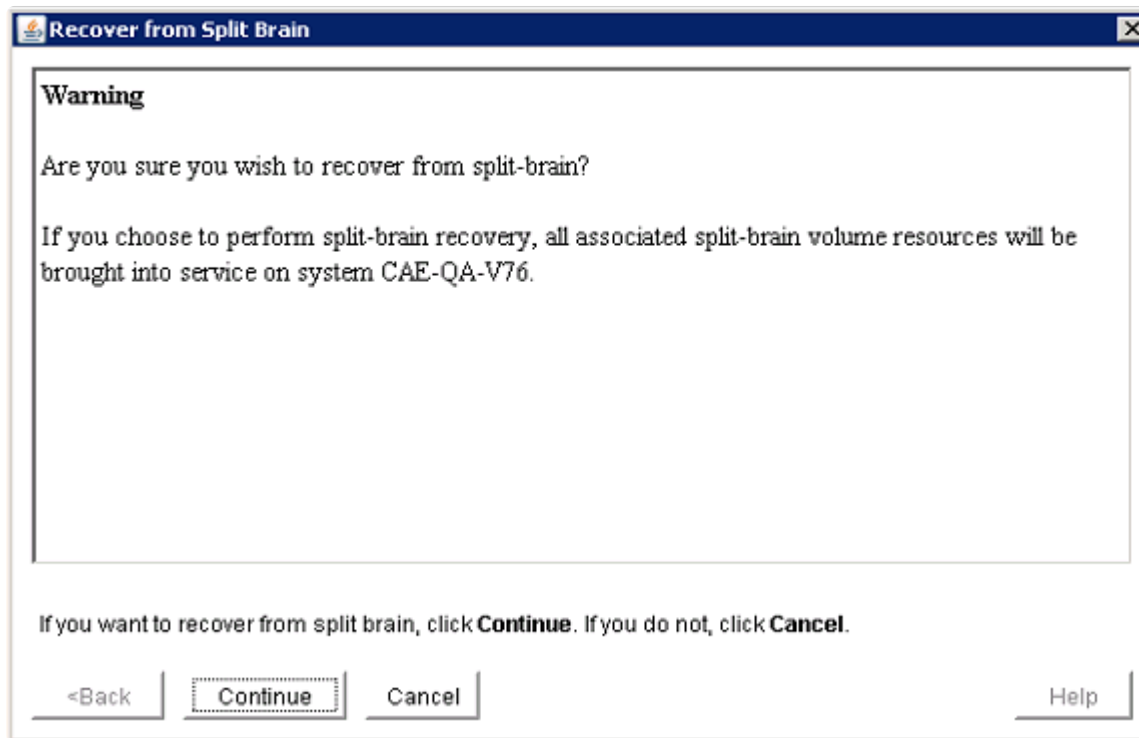
After the system's comm paths have been restored and the servers detect that the volumes are in the Split-Brain state, you will have to perform the **Split-Brain Recovery** procedure below.

Note: If multiple volumes are detected in different resource hierarchies, you will have to perform the Split-Brain Recovery procedure on each volume. Split-Brain volumes that are in the same hierarchy will be recovered together.


1. Right-click on the volume instance icon under the system that will be the source. The **Resource Context** menu displays. You can also right-click on the volume instance icon in the **Hierarchies** list in the far left panel. Choose **Split-Brain Recovery** from the menu, and you will be prompted to select which server should be the mirror source.
2. Select **Split-Brain Recovery** from the menu.



- The following warning message will display. Select the **Continue** button to complete the Split-Brain Recovery process.



4. Select **Finish** to complete.



5. Once the recovery is complete, the recovered resources will appear in the GUI as follows. The Split-Brain Recovery process will be completed when the mirror is re-synced.

LifeKeeper GUI

File Edit View Help

192.168.30.3

Active Protected

192.168.30.3

Vol.E

Vol.F

CAE-QA-V76			CAE-QA-V75		
✓	Active	10	↓	StandBy	1
✓	Active	10	↔	Mirroring	1
✓	Active	10	↔	Mirroring	1

Resource Tag= Vol.F, Resource ID= F:

DataKeeper


Your information resource for SIOS DataKeeper Cluster Edition

SIOS Technology Corp. maintains documentation for all supported versions of SIOS DataKeeper Cluster Edition. We welcome your suggestions and feedback. To help us continue to improve our documentation, please complete our brief [Documentation Feedback Survey](#).

Introduction

SIOS DataKeeper Overview

SIOS DataKeeper is a highly optimized host-based replication solution which ensures your data is replicated as quickly and as efficiently as possible from your source server across the network to one or more target servers.



Features

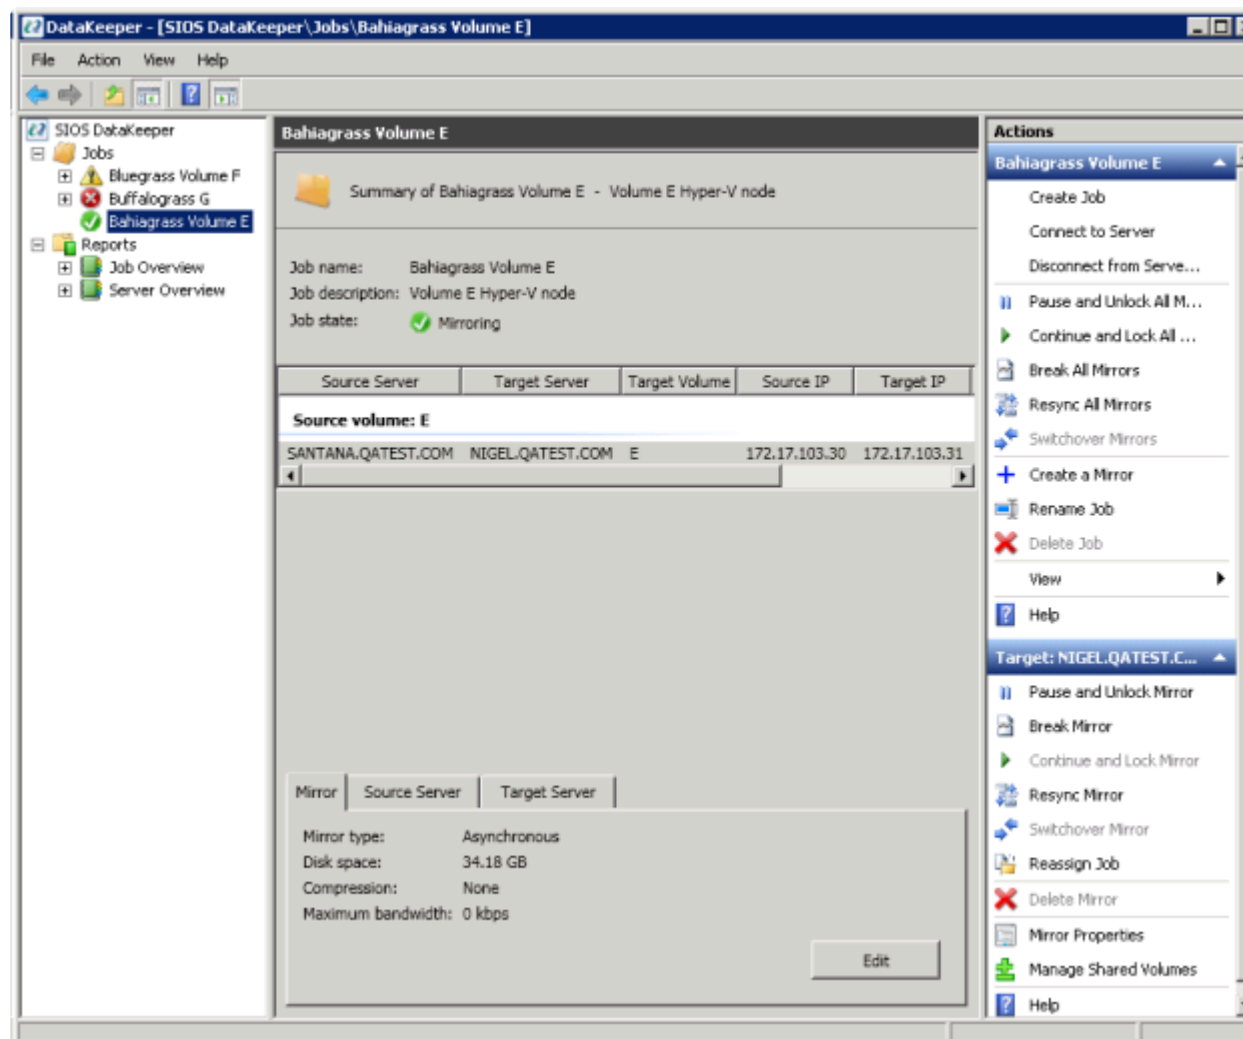
Some of the features include the following:

- Synchronous or Asynchronous block level volume replication.
- Built-in WAN optimization enabling SIOS DataKeeper to fully utilize a high speed/high latency network connection without the need for WAN accelerators.
- Efficient compression algorithms optimizing use of available bandwidth.
- Intuitive MMC 3.0 GUI.

User Interface

SIOS DataKeeper User Interface

The SIOS DataKeeper User Interface uses a standard MMC snap-in interface.



- The left pane displays the Console Tree view. This includes the **Jobs** and **Reports**. Currently, there are two reports available – **Job Overview** and **Server Overview**. The **Job Overview** report provides a summary of all the jobs on the connected servers. The **Server Overview** report provides a summary of all the mirrors on the connected servers.
- The middle pane is the **Summary** view. This includes information about the selected item.

- The right column is the **Actions** view. This pane appears when activated through the * View menu. The options available from this pane are the same options available from the **Action** menu. This column is divided into two sections. The **Actions** in the top section apply to the job and every mirror within the job. The **Actions** in the bottom section apply only to the selected mirror.
- At the bottom of the main window, three tabs appear: **Mirror**, **Source Server** and **Target Server**. These tabs provide information on the mirror that has been selected.
- The icon shows the state of the mirror, which provides more information than the icons and states provided in the Failover cluster UI.

DataKeeper Components

SIOS DataKeeper for Windows is comprised of the following components:

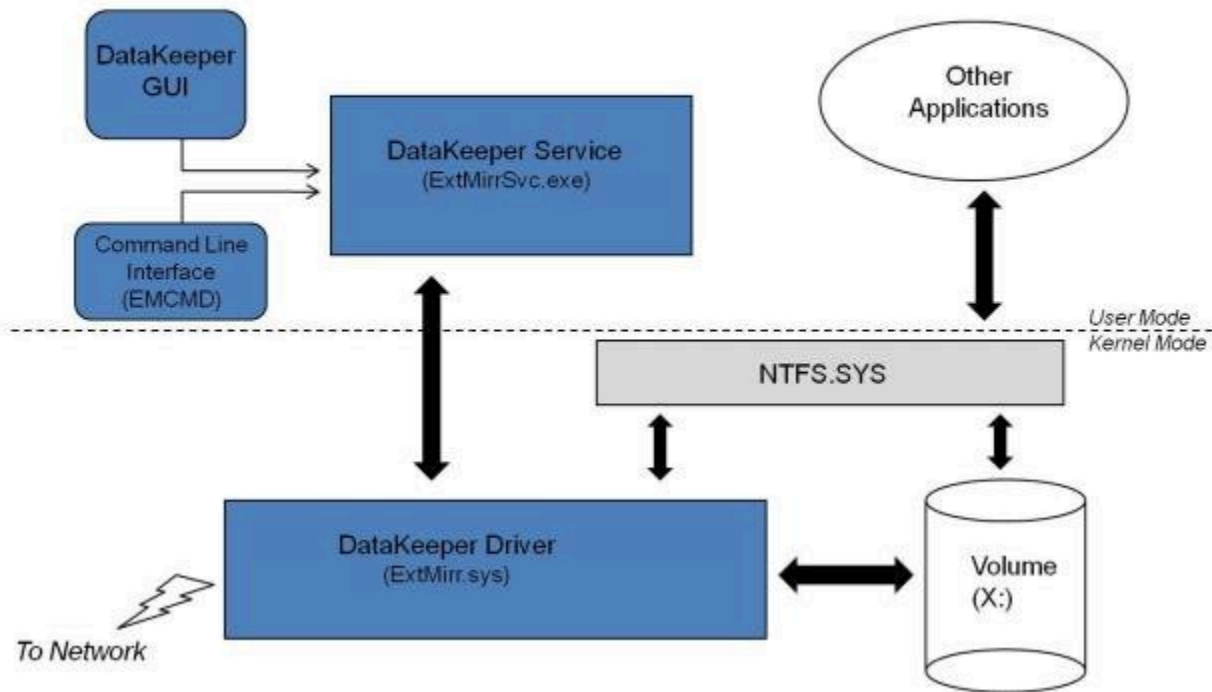
- **DataKeeper Driver (ExtMirr.sys)** – The DataKeeper Driver is a kernel mode driver and is responsible for all mirroring activity between the mirror endpoints.
- **DataKeeper Service (ExtMirrSvc.exe)** – The DataKeeper Service links the DataKeeper GUI and Command Line Interface to the DataKeeper Driver. All commands to manipulate the mirror are relayed through the DataKeeper Service to the DataKeeper Driver.

✳ **Important:** Stopping the DataKeeper Service does not stop mirroring. Sending the driver a PAUSE mirror, BREAK mirror or DELETE mirror command is the only way to interrupt mirroring.

- **DataKeeper Service Log On ID and Password Selection** – The DataKeeper Service Log On ID and Password Selection allows you to select the type of account to be used to start the service. Domain and Server account IDs with administrator privileges allow improved disaster recovery when network disruptions occur.
- **Command Line Interface (EMCMD.exe)** – There is an entire suite of EMCMD command options that can be used to operate DataKeeper.
- **DataKeeper GUI (Datakeeper.msc)** – The DataKeeper GUI is an MMC 3.0 (Microsoft Management Console) based user interface which allows you to control mirroring activity and obtain mirror status.
- **Packaging files, SIOS Protection Suite scripts, help files, etc.**

The following diagram displays how the DataKeeper components interface with the NTFS file system and each other to perform data replication.

DataKeeper Architecture



DataKeeper Service Log On ID and Password Selection

During a new DataKeeper installation setup, the user will be prompted for a DataKeeper Service Log On ID and Password.

The DataKeeper Service uses authenticated connections to perform volume switchovers and make mirror role changes across multiple servers. The Log On ID account chosen to run the DataKeeper Service will determine how much authority and permission is available to establish connections between servers and perform volume switchovers, especially when server or network disruptions occur.

Several types of Service Log On ID accounts are available as follows:

- A **Domain Account** with administrator privileges, valid on all connected servers in the domain (*recommended*)
- A **Server Account** with administrator privileges, valid on all connected servers
- The **Local System Account** (*not recommended*)

Note: For Workgroups, use the **Server Account** option and use the server name \ administrator on each system as the Service Account for DataKeeper. **You should also log on to all servers using this same Log On ID and Password** (see related [Known Issue](#)).

Note: The domain or server account used must be added to the Local System Administrators Group. The account must have administrator privileges on each server that DataKeeper is installed on.

Please note that the Local System account cannot be authenticated properly in a domain when network connectivity with Active Directory is lost. In that situation, connections between servers cannot be established with the Local System account causing DataKeeper volume switchover commands, via the network, to be rejected. IT organizations requiring fault tolerance during a disaster recovery, including network disruptions, should not use the Local System account.

DataKeeper Installation – Service Logon ID Type Selection:

SIOS DataKeeper for Windows

Service Setup

Service Logon Account Setup

The DataKeeper Service requires a logon account with Administrator privileges. The service logon account and password must be the same on all servers where DataKeeper is running. A Domain account is recommended.

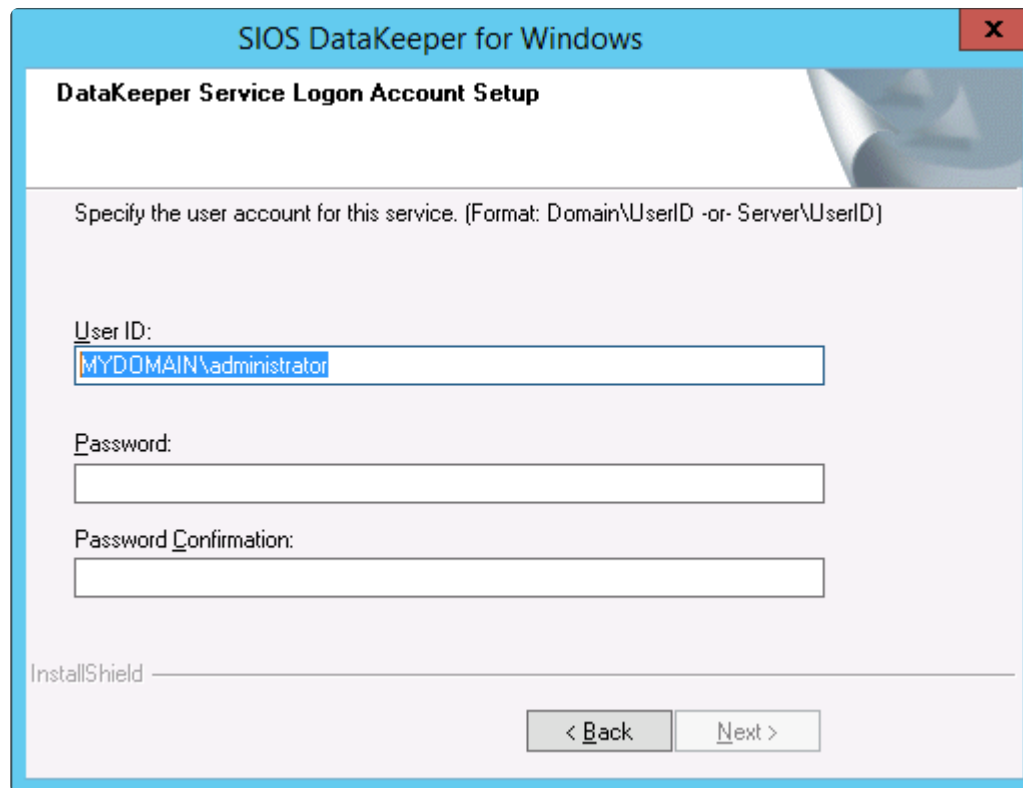
☒ Domain or Server account (recommended)

☐ LocalSystem account

InstallShield

< Back Next >

If a Domain or Server account is selected above, the DataKeeper Service Log On ID and Password Entry Form is displayed to enter that information.



The screenshot shows a Windows-style dialog box titled "SIOS DataKeeper for Windows" with a red close button in the top right corner. The main heading inside the window is "DataKeeper Service Logon Account Setup". Below this, a text label reads: "Specify the user account for this service. (Format: Domain\UserID -or- Server\UserID)". There are three input fields: the first is labeled "User ID:" and contains the text "MYDOMAIN\administrator"; the second is labeled "Password:" and is empty; the third is labeled "Password Confirmation:" and is also empty. At the bottom left, the text "InstallShield" is visible. At the bottom right, there are two buttons: "< Back" and "Next >".

It is recommended that the LifeKeeper and DataKeeper service accounts are synchronized on each system to ensure more reliable switchovers and failovers.



SIOS DataKeeper for Windows

Service Setup

Service Logon Account Setup

For optimum network connectivity DataKeeper and LifeKeeper services should use the same service logon accounts. Currently, the LifeKeeper service logon account does not match the DataKeeper service logon account. Make your selection below.

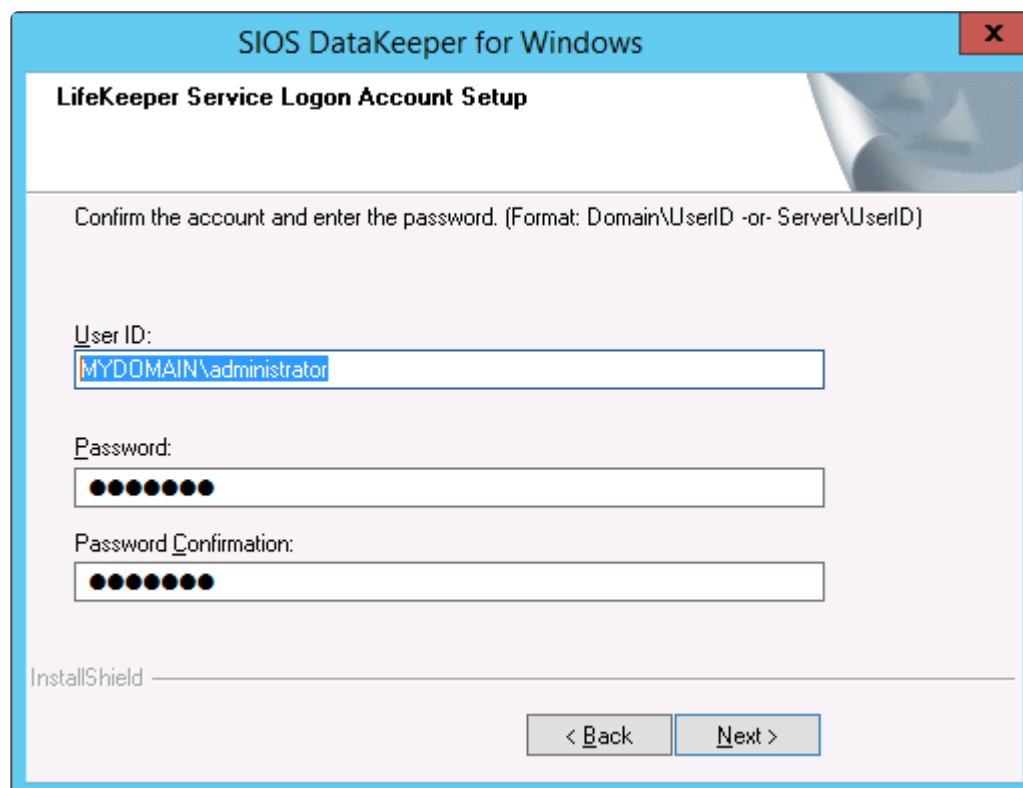
☒ Synchronize LifeKeeper Account (recommended)

☐ Do Not Synchronize Account

InstallShield

< Back Next >

LifeKeeper Service Logon:



SIOS DataKeeper for Windows

LifeKeeper Service Logon Account Setup

Confirm the account and enter the password. (Format: Domain\UserID -or- Server\UserID)

User ID:
MYDOMAIN\administrator


Password:
●●●●●●●●

Password Confirmation:
●●●●●●●●

InstallShield

< Back Next >

If the DataKeeper Service has previously been configured with a Service Log On ID and Password, the setup program will omit the Service ID and Password selection dialogs. However, at any time, an administrator can modify the DataKeeper Service Log On ID and Password using the Windows Service Applet. Be sure to restart the DataKeeper Service after changing the Log On ID and/or Password.



The following table outlines these requirements:

Environment	DataKeeper Service Requirements	DataKeeper UI Requirements
Same Domain or Trusted Domain	<ul style="list-style-type: none">Run the DK Service on all systems as the same account with the same credentials	<ul style="list-style-type: none">Log in as a domain admin and run the DK GUIOr use “run as” Administrator option to run DK GUI

Environment	<ul style="list-style-type: none">• Okay to use the default = Local System Account	
Mixed Environment Servers in a Mixture of Domain and WorkGroup or Servers in Separate Domains	<ul style="list-style-type: none">• Create a local account on each system with same account name and password• Add this local account to the Administrator Group• Run the DK Service on all systems with the local account	<ul style="list-style-type: none">• Log in using the local account you created to run the DK Service• Run the DK GUI <p>You should also log on to all servers using this same Log On ID and Password (see related Known Issue).</p>

Understanding Replication

How SIOS DataKeeper Works

At the highest level, DataKeeper provides the ability to mirror a volume on one system (source) to a different volume on another system (target) across any network. When the mirror is created, all data on the source volume is initially replicated to the target volume, overwriting it. When this initial synchronization (also referred to as a full resync of the data) of the volumes is complete, the target volume is an exact replica of the source volume in terms of size and data content. Once the mirror is established, DataKeeper intercepts all writes to the source volume and replicates that data across the network to the target volume.

Replication is performed at the block level in one of two ways:

- [Synchronous replication](#)
- [Asynchronous replication](#)

In most cases, asynchronous mirroring is recommended on WANs and synchronous mirroring is recommended on LANs.

SIOS DataKeeper Intent Log

SIOS DataKeeper uses an intent log (also referred to as a bitmap file) to track changes made to the source, or to target volume during times that the target is unlocked. This log is a persistent record of write requests which have not yet been committed to both servers.

The intent log gives SIOS DataKeeper the ability to survive a source or target system failure or reboot without requiring a full mirror resync after the recovery of the system.

There is a performance overhead associated with the intent log, since each write to the volume must also be reflected in the intent log file. To minimize this impact, it is recommended that the intent logs be stored on a physical disk that is not involved in heavy read or write activity. See [Relocation of Intent Log](#) for more information.

DataKeeper Bitmap File Size Calculation

- One bit per block (each block is 64KB)
- One bit represents 64KB of volume space/

Example: For a 640GB volume the bitmap is 1,310,720 bytes

$$640 \times 1073741824 / 65536 / 8$$

That is 1,310,720 bytes for the bitmap.

Non-Shared Volumes

By default, this intent log feature is enabled, and the intent log files are stored in a subdirectory called "Bitmaps" under the directory where SIOS DataKeeper was installed.

To create the intent log file in a directory other than the default location, set the [BitmapBaseDir](#) registry entry to a directory where SIOS DataKeeper will create the file. See [Relocation of Intent Log](#) for more information.

To disable the intent log feature, clear the [BitmapBaseDir](#) registry entry (set it to an empty string) on all current and potential mirror endpoint servers. **Disabling the intent log requires a reboot on each of these systems in order for this setting to take effect.** Keep in mind that if this feature is disabled, a full resync will be performed in the event of a source system failure.

Shared Volumes

When replicating shared volumes, the intent log files are stored in a subdirectory called “ReplicationBitmaps” on the replicated volume itself. This is necessary to allow switchover to the other shared source servers without resulting in a full resync of the data.

SIOS does not recommend relocating intent logs from their default locations.

Configuration Issue

When configuring a [BitmapBaseDir](#) registry entry, make sure that the folder and drive letter specified exist. If configured with a drive letter that does not exist, the following message will be received upon system boot up:

```
Global bitmap volume {drive letter}: has not been detected yet. Mirror source  
threads may hang if this volume does not exist. Check to make sure that the  
BitmapBaseDir registry entry specifies a valid volume for storage of bitmaps.
```

Relocation of Intent Log

To relocate the Intent Log (bitmap file), please perform the following on all servers involved:

 LEAVE THE MIRROR IN THE MIRRORING STATE! Do not pause it and then move the bitmap file.

1. If you have more than one DataKeeper mirror, move all mirrors to a single system so that it is source for all mirrors.
2. On all systems, create the directory for the new location of the bitmap files (*i.e.* `R:\Bitmaps`). **Important:** If you choose to relocate the bitmap file from the default location (`%EXTMIRRBASE%\Bitmaps`), you must first create the new directory before changing the location in the registry and rebooting the system.
3. Modify the [BitmapBaseDir](#) registry value on all systems other than the mirror source system to reflect the new location. This includes mirror targets and any systems that share the volume with the mirror source or share with any of the targets.

Edit Registry via regedit:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters
```

Modify the "BitmapBaseDir" parameter, change to the new location (*i.e.* `R:\Bitmaps`)

4. Reboot each of the non-source systems. If this volume is part of a Windows cluster, be sure that you do not shut down too many nodes simultaneously or you may lose the cluster quorum and cause the cluster to shut down on the remaining nodes.
5. Switch any volumes on the source system over to another system (target or shared source). Repeat Steps 2 and 3 on the system that was previously source.

6. After rebooting the original source system, all volume resources can be switched back to that system.

Resynchronization

SIOS DataKeeper Resynchronization

SIOS DataKeeper performs resynchronization through the use of a bitmap file ([intent log](#)). It allocates memory that is used to keep track of “dirty” or “clean” blocks. When a full resync begins, SIOS DataKeeper initializes the bit for each block that is in use by the file system to 1 (“dirty”), indicating that it needs to be sent to the target system. A full resync occurs at the initial creation of a mirror and during the resync operation after a mirror is broken. It then starts at the beginning of the bitmap, finds the first block whose bit is set to 1 or dirty, reads the corresponding block from the local hard disk, and sends it to the remote system. After this has completed successfully, it sets the block to 0 (“clean”). SIOS DataKeeper then finds the next dirty bit and repeats this process.

As new writes come in during a resync, the corresponding blocks are set to 1 or dirty.

Once resync gets to the end of the bitmap, it looks to see if there are still any dirty blocks. It does this through a counter that is incremented when one is made dirty and decremented when cleaned. If any blocks are dirty, it resets its pointer to the beginning of the bitmap and starts again, only sending the dirty blocks to the remote system.

This process continues for multiple passes until all blocks are clean. When this happens, the mirror will go from the **Resynchronizing** state to the **Mirroring** state, and at that point, every write is mirrored (the bitmap is no longer necessary at that point).

You can follow the resynchronization process by viewing the resynchronization control counters in Performance Monitor.

This same resynchronization mechanism is used when you CONTINUE a PAUSED mirror.



If the target system is rebooted/shut down via the DK GUI when

mirrors are paused and unlocked, a full resync will occur. To prevent the full resync in this case, be sure to perform a "Continue and Lock" prior to rebooting or shutting down the target system.

Initial Creation of a Mirror

When the mirror is created, DataKeeper must perform an [initial synchronization](#) of the data from the source volume to the target volume. This is referred to as a full resync. However, prior to this initial full resync of the data, DataKeeper first performs a process called "**whitespace elimination**" where all blocks of currently unused space on the source volume are eliminated from the initial synchronization and those blocks do not have to be replicated to the target volume.

Example: Whitespace Elimination

Source Volume Capacity	80 GB
Source Volume Free Space	35 GB
Amount of data to be resynced from source volume to target volume during initial creation of the mirror.	55 GB

Synchronous and Asynchronous Mirroring

SIOS DataKeeper employs both asynchronous and synchronous mirroring schemes. Understanding the advantages and disadvantages between synchronous and asynchronous mirroring is essential to the correct operation of SIOS DataKeeper.

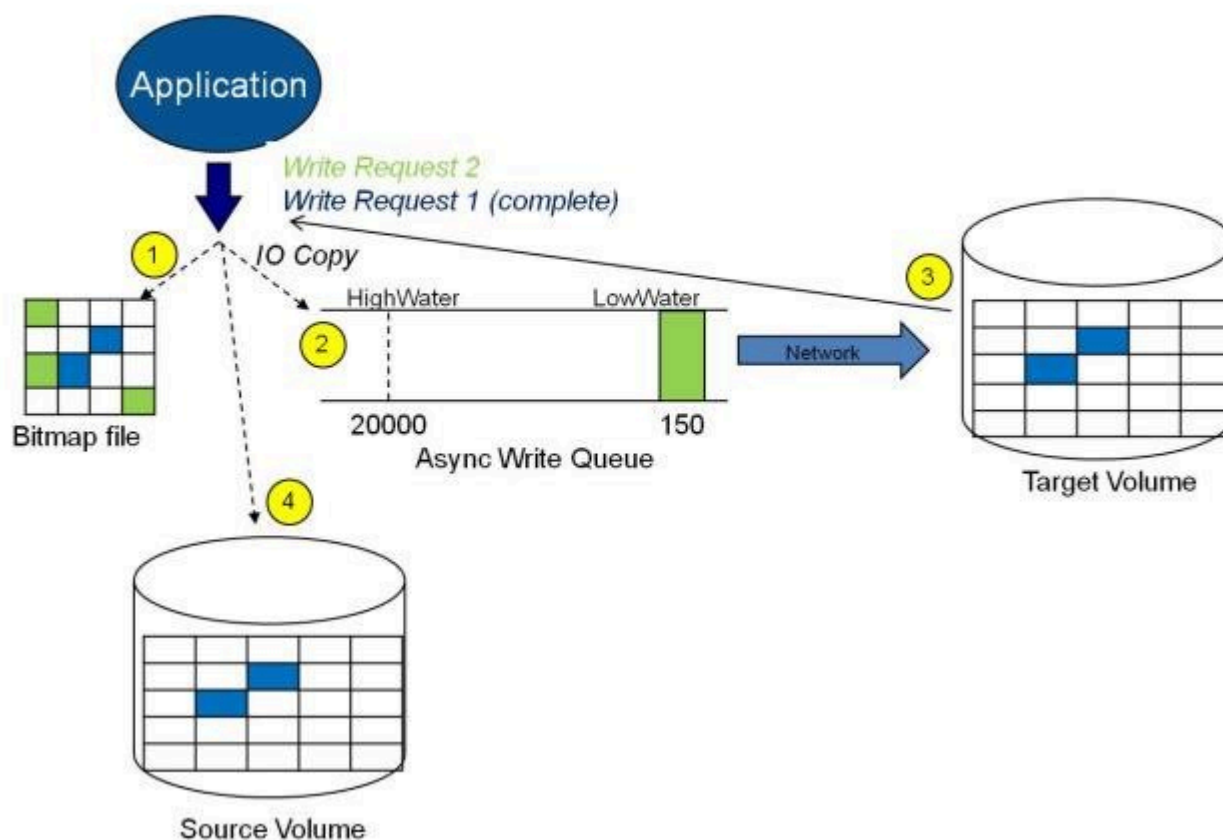
Synchronous Mirroring

With synchronous mirroring, each write is intercepted and transmitted to the target system to be written on the target volume at the same time that the write is committed to the underlying storage device on the source system. Once both the local and target writes are complete, the write request is acknowledged as complete and control is returned to the application that initiated the write. Persistent bitmap file on the source system is updated.

The following sequence of events describes what happens when a write request is made to the source volume of a synchronous mirror.

1. The following occur in parallel.
 - a. A copy of the write is put on the mirror Write Queue.
 - b. The write is sent to the local volume for completion.
2. The write returns a completion status to the caller after both operations above complete.
 - a. If any condition prevents the write from completing on the Target (network transmission error or write error on the target system), the mirror state is changed to Paused. However, the status of the volume write which is returned to the caller is not affected.
 - b. The status of the local volume write is returned to the caller.

Synchronous Replication



In this diagram, Write Request 1 has already completed. Both the target and the source volumes have been updated.

Write Request 2 has been sent from the application and the write is about to be written to the target volume. Once written to the target volume, DataKeeper will send an acknowledgment that the write was successful on the target volume, and in parallel, the write is committed to the source volume.

At this point, the write request is complete and control is returned to the application that initiated the write.

While synchronous mirroring insures that there will be no data loss in the event of a source system failure, synchronous mirroring can have a significant impact on the application's performance, especially in WAN or slow network configurations, because the application must wait for the write to occur on the source and across the network on the target.

Asynchronous Mirroring

With asynchronous mirroring, each write is intercepted and a copy of the data is made. That copy is queued to be transmitted to the target system as soon as the network will allow it. Meanwhile, the original write request is committed to the underlying storage device and control is immediately returned to the application that initiated the write.

To maintain data consistency across multiple volumes (such as database Log and Data files), some applications send Flush requests to the volume. DataKeeper honors Flush requests on a volume with a mirror in the Mirroring state by waiting for all writes in the queue to be sent to the target system and acknowledged. To prevent performance from being impacted in such cases, the registry entry "[DontFlushAsyncQueue](#)" may be set, or you may consider locating all files on the same volume.

At any given time, there may be write transactions waiting in the queue to be sent to the target machine. But it is important to understand that these writes reach the target volume in time order, so the integrity of the data on the target volume is always a valid snapshot of the source volume at some point in time. Should the source system fail, it is possible that the target system did not receive all of the writes that were queued up, but the data that has made it to the target volume is valid and usable.

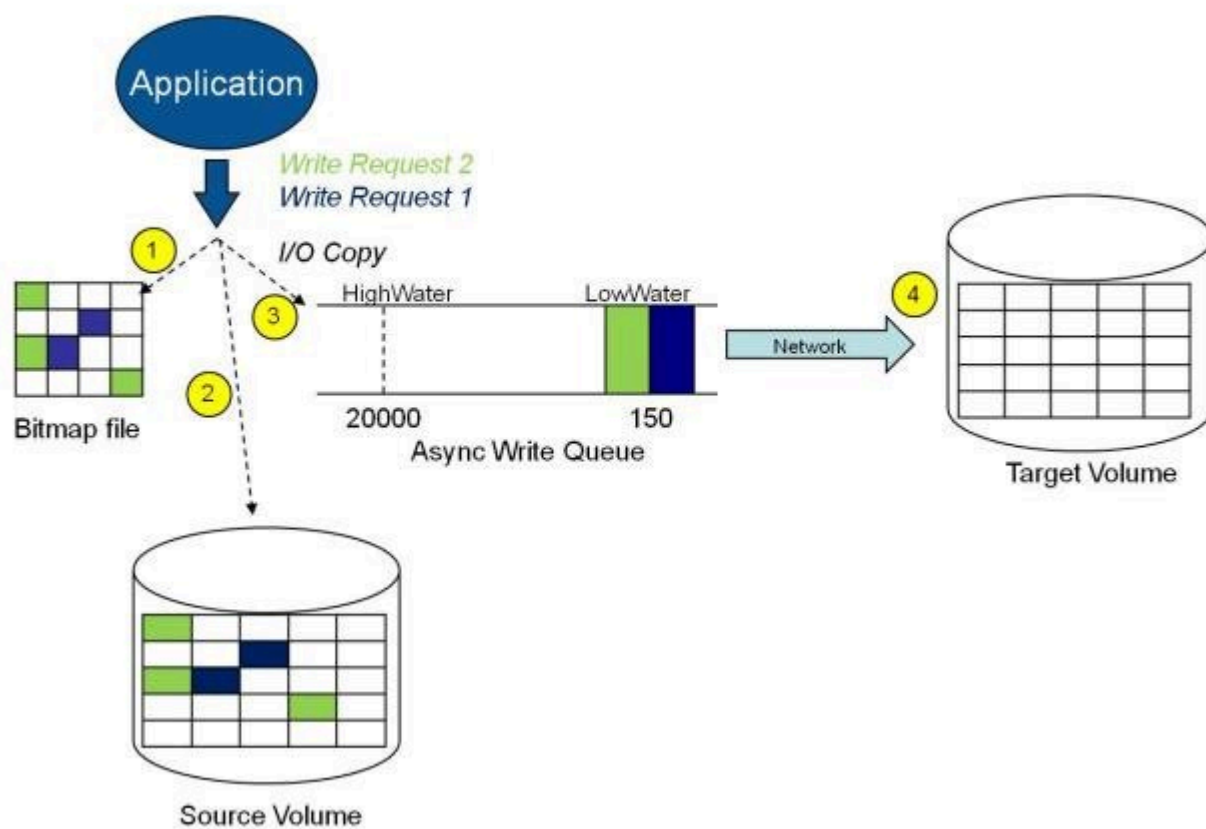
The following sequence of events describes what happens when a write request is made to the source volume of an asynchronous mirror.

1. Persistent bitmap file on the source system is updated.
2. Source system adds a copy of the write to the mirror Write Queue.
3. Source system executes the write request to its source volume and returns to the caller.
4. Writes that are in the queue are sent to the target system. The target system executes the write request on its target volume and then sends the status of the write back to the primary.
5. If the mirror's write queue reaches a configured limit (WriteQueueHighWater or WriteQueueByteLimit reached), the mirror's "BlockWritesOnLimitReached"

setting will be used to determine what the behavior should be. If `BlockWritesOnLimitReached` is "0", the mirror is paused and a partial resync starts a short time later. If `BlockWritesOnLimitReached` is "1", the incoming write is delayed until there is space on the write queue for it. The mirror remains in the **Mirroring** state, but the application throughput slows down to match the speed of the network and the remote node's volume.

6. Should an error occur during network transmission or while the target system executes its target volume write, the write process on the secondary is terminated. The state of the mirror then changes from **Mirroring** to **Paused**.

Asynchronous Replication: Mirroring

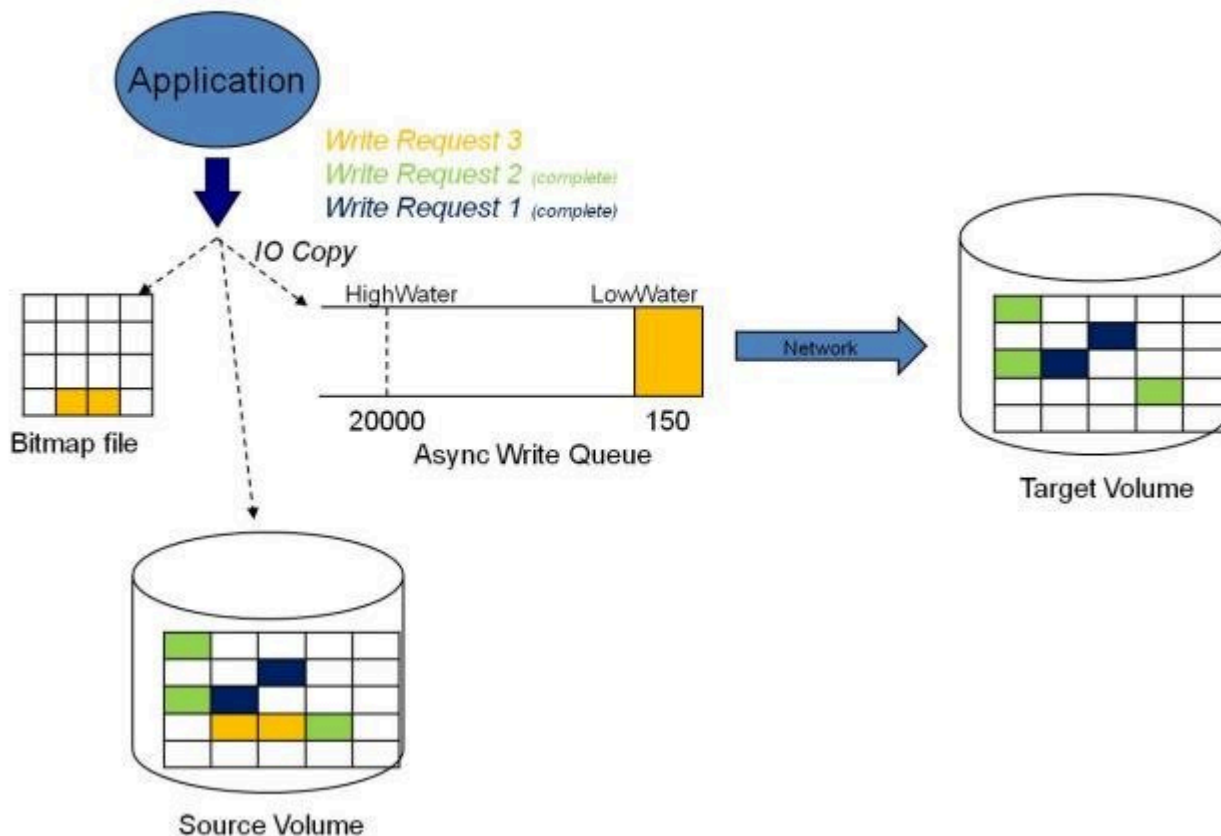


In the diagram above, the two write requests have been written to the source volume and are in the queue to be sent to the target system. However, control has already returned back to the application who initiated the writes.

In the diagram below, the third write request has been initiated while the first two writes have successfully been written to both the source and target volumes.

While in the mirroring state, write requests are sent to the target volume in time order. Thus, the target volume is always an exact replica of the source volume at some point in time.

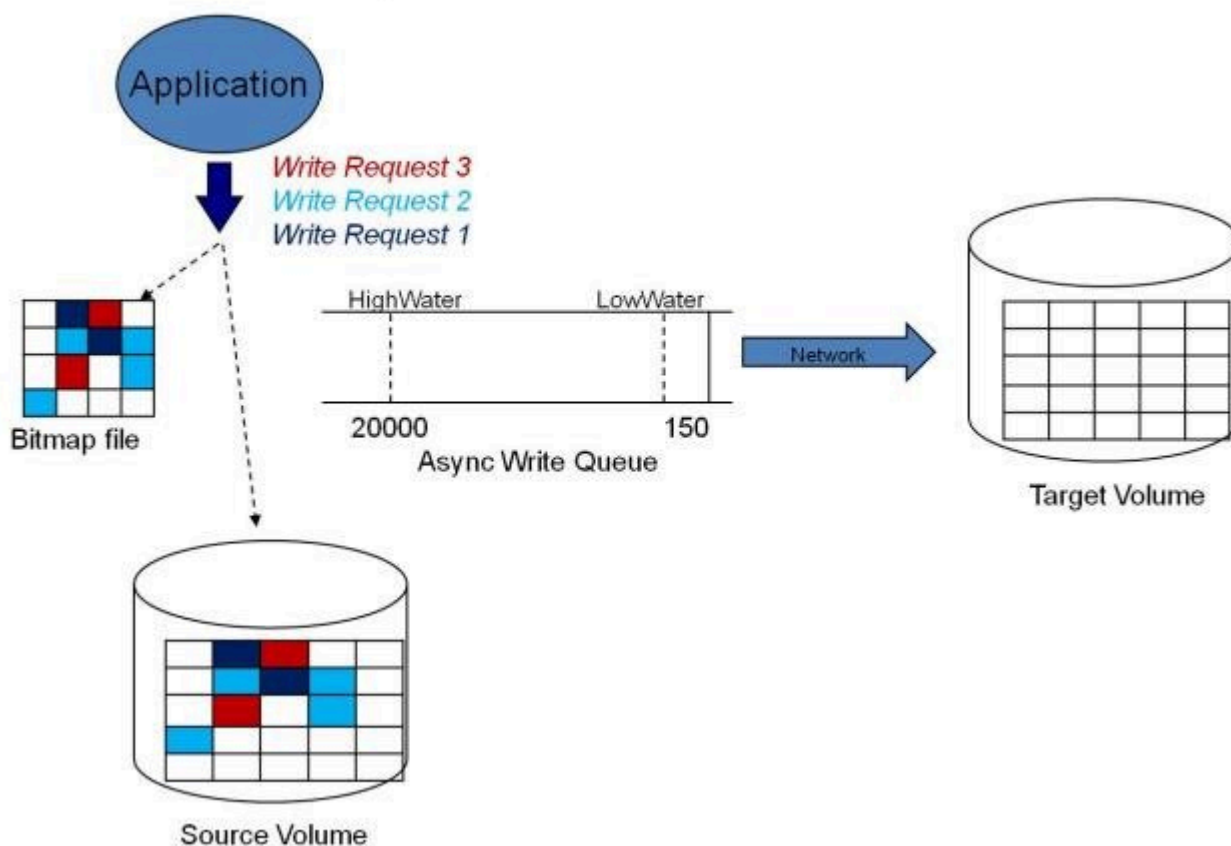
Asynchronous Replication: Mirroring



Mirror PAUSED

In the event of an interruption to the normal mirroring process as described above, the mirror changes from the **MIRRORING** state to a **PAUSED** state. All changes to the source volume are tracked in the persistent bitmap file only and nothing is sent to the target system.

Replication: Mirror Paused



Mirror RESYNCING

When the interruption of either an Asynchronous or Synchronous mirror is resolved, it is necessary to resynchronize the source and target volumes and the mirror enters into a **RESYNC** state.


DataKeeper reads sequentially through the persistent bitmap file to determine what blocks have changed on the source volume while the mirror was **PAUSED** and then resynchronizes only those blocks to the target volume. This procedure is known as a partial resync of the data.

The user may notice a **Resync Pending** state in the GUI, which is a transitory state and will change to the **Resync** state.

During resynchronization, all writes are treated as Asynchronous, even if the mirror is a Synchronous mirror. The appropriate bits in the bitmap are marked dirty and are later sent to the target during the process of partial resync as

described above.

Replication: Resynchronization



Read and Write Operations

After the volume mirror is created and the two drives on the primary and secondary servers are synchronized, the following events occur:

- The system locks out all user access to the target volume; reads and writes are not allowed to the target volume. The source volume is accessible for both reads and writes.
- Both mirrored and non-mirrored volume read operations arriving at the driver on the primary server are passed on and allowed to complete normally without intervention. Reads of a mirrored volume on the secondary system are not allowed, i.e., the secondary has not assumed the role of a failed primary.
- Whenever the primary server receives a write request, the system first determines whether the request is for a mirrored volume. If not, the write is allowed to complete normally without any further intervention. If the write request is for a mirrored volume, the request is handled depending on the mirroring type:
 - If the type is [synchronous](#), then the write request is put on the mirror Write Queue for transmission to the target system, and simultaneously sent to the local source volume. The write operation is not acknowledged as complete to the process that issued the write until the source disk write completes **and** notification from the target is received (success or failure). Should an error occur during network transmission or while the target system executes its write, the write process on the target is terminated and the state of the mirror is changed to **Paused**. The source volume completes the write regardless of the target write status.

If the type is [asynchronous](#), then the primary executes the write request to its source volume, puts a copy of the write on the asynchronous write queue and returns to the caller. Writes that are in the queue are sent to the target volume. The secondary system executes the write request on the target volume and then sends the status of the write back to the primary. Should an error occur during network transmission or while the secondary executes its mirrored volume write, the write process on the secondary is terminated. The state of the mirror then changes from **Mirroring** to **Paused**.

To ensure uninterrupted system operation, SIOS DataKeeper momentarily pauses the mirror and automatically continues it (i.e., performs a partial resync) in the following cases:

- When the mirror's BlockWritesOnLimitReached settings is 0, and the mirror write queue length reaches the WriteQueueHighWater limit, or the number of bytes in the queue reaches the WriteQueueByteLimitMB limit, due to a large number of writes to the volume in a short period of time (e.g., database creation). The user can monitor the mirroring behavior using the SIOS DataKeeper Performance Monitor counters and adjust the WriteQueueHighWater and/or the WriteQueueByteLimitMB value if necessary. See [Registry Entries](#) for more details.
- When transmission of a write to the target system times out or fails due to resource shortage (e.g., source system resource starvation due to a flood of writes/network transmissions in a short period of time).

Volume Considerations

SIOS DataKeeper primary and secondary systems have three types of volumes: system, non-mirrored and mirrored. During mirroring operations, system and non-mirrored volumes are not affected and the user has full access to all applications and data on these volumes.

What Volumes Cannot be Mirrored

The SIOS DataKeeper service filters out the following types of disk partitions:

- Windows system volume
- Volume(s) that contain the Windows pagefile
- Non-NTFS formatted volumes (e.g. FAT, FAT32, Raw FS, ReFS)
- Non-fixed drive types (e.g. CD-ROMs, diskettes)
- Target volumes that are smaller than the source volume

Volume Size Considerations

The source and target systems are not required to have drives of the same physical size. When the mirror is established, the target volume must be the same size, or larger than the source volume.

There is no limit on the size of volumes that can participate in a SIOS DataKeeper mirror. However, you should be aware that on initial mirror creation, all data that is in use by the file system on the source volume must be sent to the target. For instance, on a 20 GB volume with 2 GB used and 18 GB free, 2 GB of data must be synchronized to the target. The speed of the network connection between the two systems, along with the amount of data to be synchronized, dictates how long the initial mirror creation will take.

Specifying Network Cards for Mirroring

SIOS DataKeeper allows the administrator to specify which IP addresses should be used as mirror end-points. This allows the replicated data to be transmitted across a specific network which permits the user to segment mirrored traffic away from the client network if desired.

Dedicated LAN for Replication

While it is not required, a dedicated (private) network between the two servers will provide performance benefits and not adversely affect the client network.

Performance Monitor Counters

SIOS DataKeeper provides counters that extend Performance Monitor with statistics about the status of mirroring on volumes. The counters are installed during the full installation of SIOS DataKeeper software.

To access the counters, do the following:

1. On a **Microsoft Windows 2008 R2** system, start the **Windows Performance Monitor** through the **Start** menu in the **Reliability and Performance** group.

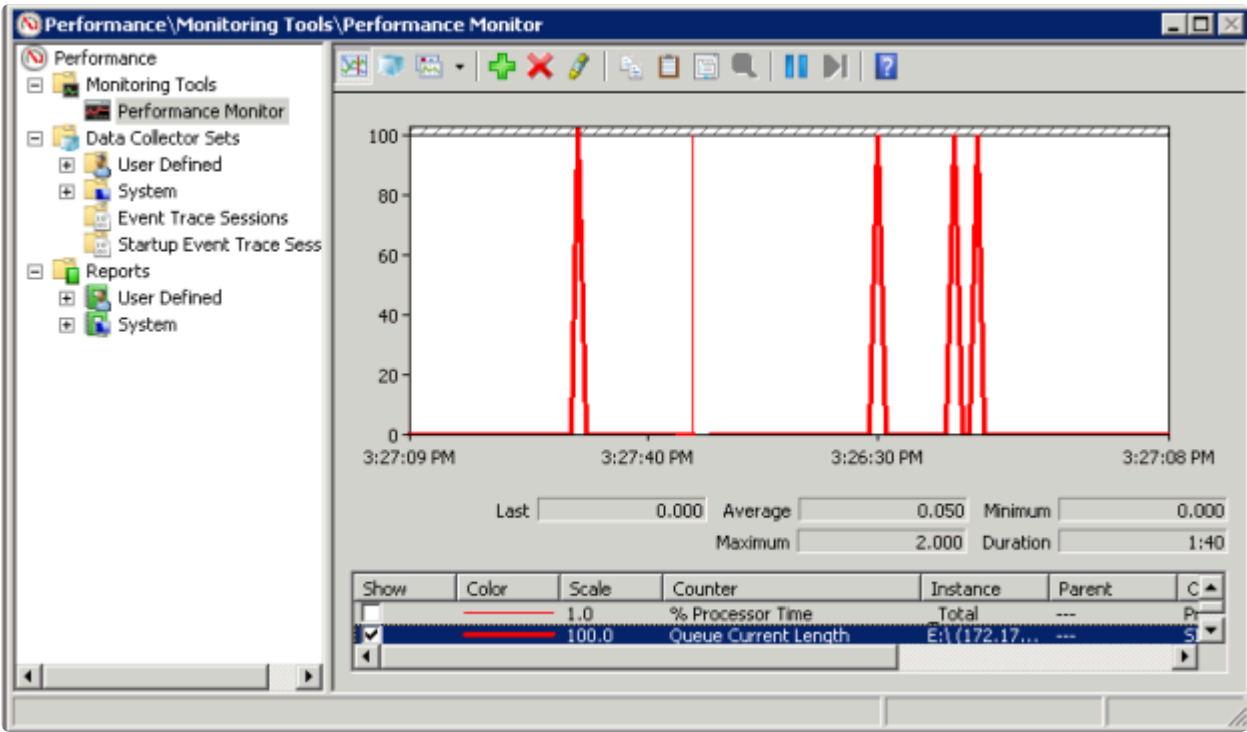
On a **Microsoft Windows 2012** system, start the **Windows Performance Monitor** through the **Performance Monitor** option in the **Administrative tools**.

On all versions of Windows, you can start performance monitor through entering **perfmon.msc** using the command line.

2. Select **Monitoring Tools, Performance Monitor**.
3. Click the **+** button in the chart pane to open the **Add Counters** dialog box.
4. Select the **SIOS Data Replication** object.

* The Performance Monitor Counters for DataKeeper (DK) are only accessible for a system that is in the **source** role for the mirror. On a system with a mirror in the source role, there will be one instance of the specified counter available for each target of that mirror. SIOS DataKeeper performance counters are not available on a system where the mirror is in the target role.

SIOS DataKeeper provides 17 counters that allow the monitoring of various operations related to the product. These counters allow the monitoring of such things as status, queuing statistics and general mirror status.



Mirror State Counters

Mirror Elapsed Time

Default Value: 0

Range: 0 - MAX_ULONG

This value represents the amount of time, in seconds, that the volume has been in Mirror state. This value will be 0 for volumes that are not currently involved in a mirror, volumes that are currently undergoing mirror creation (and synchronization), and volumes for which a mirror has been broken or deleted.

Mirror State

Default: 0

Range: 0 - 5

This value represents the current mirroring state of a volume. The following values are defined:

0 None - The volume is not currently involved in a mirror.

1 Mirroring - The volume is currently mirroring to a target.

2 Resynchronizing - The volume is currently being synchronized with its target.

3 Broken - The mirror exists but the source and target volumes are not in sync. New writes to the volume are not tracked.

4 Paused - The mirror exists but the source and target volumes are not in sync. The source server keeps track of any new writes.

5 Resync Pending - The source volume is waiting to be resynchronized.

Mirror Type

Default: 0

Range: 0-2

This value represents the type of mirroring this volume is engaged in. The following values are defined for this release:

0 None - The volume is not currently involved in a mirror.

1 Synchronous - Data is put on the Write Queue to be sent to the target, and written to the local volume, simultaneously. The write is not acknowledged as complete until both operations complete.

2 Asynchronous - Data is put on the Write Queue to be sent to the target, and written to the local volume, simultaneously. The write is acknowledged when the local volume write completes.

Network Number of Reconnects

Default: 0

Range: 0 - MAX_ULONG

This value is the number of network reconnections that have been made while the

volume has been mirrored. A network reconnection occurs when communication is lost with the target.

Write Queue Counters

Queue Byte Limit

Default Value: 0

This value displays the write queue byte limit as set in the WriteQueueByteLimitMB registry value. This value is displayed in bytes, and is therefore 1048576 times the value set in the registry.

Queue Current Age

Default Value: 0

Range: 0 -

This value is the age (in milliseconds) of the oldest write request in the write queue.

Queue Current Bytes

Range: 0 -

This value displays the number of bytes allocated for the given mirror's Write Queue.

Queue Current Length

Default Value: 0

Range: 0 -

This value represents the current length, in terms of number of writes, of the write queue for the selected mirror.

Queue High Water

Default: 20000

This counter displays the write queue high water mark as set in the mirror WriteQueueHighWater registry value.

Resynchronization Control Counters

Resync Reads

Default: 20

This value represents the maximum number of disk blocks that can be in the process of being read and sent to the target system during mirror resynchronization.

Resync Current Block

Default: 0

Range: 0 -

During the synchronization process, this value represents the current block that is being sent to the target. At other times (i.e. when mirror state is not EmMirrorStateResync), this value will be 0.

During synchronization, a given block may be sent to the target multiple times if writes are ongoing to the volume. This is based on the number of resync passes that are required.

Resync Dirty Blocks

Default Value: 0

Range: 0 -

This value is the number of total blocks that are dirty during mirror resynchronization. "Dirty" blocks are those that must be sent to the target machine before synchronization is complete. This value will be 0 for all states

other than EmMirrorStateResync.

When a mirror synchronization is begun, this value will be initially equal to the value of Resync Total Blocks. Please note that during a mirror synchronization, Resync Dirty Blocks may actually increase if a large number of incoming writes are made to the volume.

NOTE: This counter can be used to view the amount of data remaining to be resync'ed.

The amount of data remaining is "Resync Dirty Blocks" X 65536 (where 65536 is the number of bytes in a block).

Resync Elapsed Time

Default Value: 0

Range: 0 - MAX_ULONG

While the mirror is being synchronized, this value represents the elapsed time in seconds that the synchronization has been occurring. After a mirror is successfully resynchronized, the value represents the total amount of time the previous synchronization operation took since the last system boot. The value will be 0 for volumes that either never have been synchronized or volumes that were not synchronized during the last boot.

Resync New Writes

Default: 0

Range: 0 - MAX_ULONG

This value represents the number of writes that have occurred on the volume since a synchronization operation has begun. This value will directly affect the number of dirty blocks, the number of passes required to synchronize the mirror and the amount of time the synchronization takes to complete.

Resync Pass

Default Value: 10

Range: 0 - MaxResyncPasses (Registry)

This value is the number of passes that have currently been made through the volume during the resynchronization process to update the target. The number of passes required to complete the synchronization process will increase based on the amount of writing that is being performed during synchronization. While writing to the source volume is allowed during synchronization, heavy writes will cause the synchronization to take longer, thus resulting in a much longer time until it is finished.

Resync Total Blocks

Default Value: 0

Range: 0 - MAX_ULONG

This value represents the number of 64k blocks used for resynchronization of the mirrored volume. The value is approximately equal to the file system size of the volume divided by 64K. Please note that the file system size is less than the partition size of the volume that is shown in the Windows Disk Management program. To see the file system size, type CHKDSK X: (where X is the drive letter).

Resync Phase

Default Value: 0

Range: 0 - 3

This value has been deprecated and is no longer used.

DataKeeper Configuration

Requirements/Considerations

The topics in this section identify several prerequisites to be aware of before implementing your DataKeeper configuration.

[Sector Size](#)

[Network Bandwidth](#)

[Network Adapter Settings](#)

[DataKeeper Service Log On ID and Password Selection](#)

[Firewall Configurations](#)

[High-Speed Storage Best Practices](#)

[Configuration of Data Replication From a Cluster Node to External DR Site](#)

[Performance Tuning](#)

[Disable Automatically manage paging file size for all drives](#)

[WAN Considerations](#)

[Initial Synchronization of Data Across the LAN or WAN](#)

[Compression](#)

[Bandwidth Throttle](#)

Sector Size

Beginning with DataKeeper Version 7.2.1, disks with sector size not equal to 512 bytes are supported. However, DataKeeper requires that the mirror source volume be configured on disk(s) whose sector size is the same as the disk(s) where the mirror target is configured. NTFS Metadata includes the disk sector size. DataKeeper replicates the entire NTFS file system from source to target, so the sector sizes must match.



Note: For DataKeeper Version 7.2 and prior, only disk devices whose sector size is the standard 512 bytes are supported.

Network Bandwidth

Because DataKeeper can replicate data across any available network, special consideration must be given to the question, "Is there sufficient bandwidth to successfully replicate the volume and keep the mirror in the **mirroring** state as the source volume is updated throughout the day?"

Keeping the mirror in the **mirroring** state is critical, because a switchover of the volume is not allowed unless the mirror is in the **mirroring** state.

Determine Network Bandwidth Requirements

Prior to installing SIOS DataKeeper, you should determine the network bandwidth requirements for replicating your data. Use the method below to measure the rate of change for the data that you plan to replicate. This value indicates the amount of network bandwidth that will be required to replicate that data.

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 must consider one or more of the following options:

- Enable compression in DataKeeper, or in the network hardware, if possible
- Create a local, non-replicated storage repository for temporary data and swap files if you are replicating Hyper-V virtual machines
- Reduce the amount of data being replicated
- Increase your network capacity

If the network capacity is not sufficient to keep up with the rate of change that occurs on your disks, DataKeeper mirrors will remain in a resynchronizing state for considerable periods of time. During resynchronization, data on the target volume is not guaranteed to be consistent.

Measuring Rate of Change

Use [Performance Monitor](#) (perfmon) to measure the rate of change that occurs on your volumes that are to be replicated. The best way to do this is to create a log of 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,

- use perfmon to create a user-defined data collector set on Windows 2008 or Windows 2012.
- add the counter "Disk Write Bytes/sec" for each volume - the volume counters can be found in the logical disks group.
- start the log and let it run for the predetermined amount of time, then stop and open the log.

An alternative to creating a log of disk writes is to use perfmon to track disk write bytes/sec interactively, in the Performance Monitor tool, and to observe the maximum and average values there.

SIOS DataKeeper handles short bursts of write activity by adding that data to its 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 DataKeeper and your network can transmit.

SIOS DataKeeper can handle the following average rates of change, approximately:

Network Bandwidth	Rate of Change
1.5 Mbps (T1)	182,000 Bytes/sec (1.45 Mbps)
10 Mbps	1,175,000 Bytes/sec (9.4 Mbps)
45 Mbps (T3)	5,250,000 Bytes/sec (41.75 Mbps)
100 Mbps	12,000,000 Bytes/sec (96 Mbps)
1000 Mbps (Gigabit)	65,000,000 Bytes/sec (520 Mbps)

Network Adapter Settings

DataKeeper requires that "**File and Printer Sharing for Microsoft Networks**" be enabled on the network interfaces to make a NAMED PIPE connection and be able to run DataKeeper's command line tool (EMCMD).

To test if you can make a Named Pipe connection, try to map a network drive on the TARGET system. If that fails, you have a Named Pipe issue.

DataKeeper also requires that **NetBIOS over TCP/IP** and **SMB** protocols be enabled. If the GUI does not operate correctly, make sure the following network configurations are enabled:

- Enable **NetBIOS over TCP/IP** and **SMB** protocols as in the following example:

```
My Computer->Manage->System Tools->Device Manager->View->Show Hidden  
Devices->Non-Plug and Play Drivers->NetBIOS over Tcpip (Enable)
```

- Enable **NetBIOS over TCP/IP** on each network adapter carrying mirror traffic as in the following example:

```
Start->Settings->Network and Dial-up Connections->->Properties->  
>Internet Protocol(TCP/IP)->Properties->Advanced...button->WINS tab->  
>Enable NetBIOS over TCP/IP radio button (Checked)
```

- Enable the Microsoft "**Client for Microsoft Networks**" component on each system where the DataKeeper Administrator GUI will be used. This must be on the same adapter with **NetBIOS over TCP/IP** enabled (above). For example:

```
Start->Settings->Network and Dial-up Connections->->Properties->Client  
for Microsoft Networks(checked)
```

- Enable the Microsoft "**File and Printer Sharing for Microsoft Networks**" component on each system which the DataKeeper Administrator GUI will connect to locally and remotely. This must be on the same adapter with **NetBIOS over TCP/IP** enabled (above). For example:

```
Start->Settings->Network and Dial-up Connections->->Properties->File
```

and Printer Sharing for Microsoft

DataKeeper Service Log On ID and Password Selection

During a new DataKeeper installation setup, the user will be prompted for a DataKeeper Service Log On ID and Password.

The DataKeeper Service uses authenticated connections to perform volume switchovers and make mirror role changes across multiple servers. The Log On ID account chosen to run the DataKeeper Service will determine how much authority and permission is available to establish connections between servers and perform volume switchovers, especially when server or network disruptions occur.

Several types of Service Log On ID accounts are available as follows:

- A **Domain Account** with administrator privileges, valid on all connected servers in the domain (*recommended*)
- A **Server Account** with administrator privileges, valid on all connected servers
- The **Local System Account** (*not recommended*)

Note: For Workgroups, use the **Server Account** option and use the server name \ administrator on each system as the Service Account for DataKeeper. **You should also log on to all servers using this same Log On ID and Password** (see related [Known Issue](#)).

Note: The domain or server account used must be added to the Local System Administrators Group. The account must have administrator privileges on each server that DataKeeper is installed on.

Please note that the Local System account cannot be authenticated properly in a domain when network connectivity with Active Directory is lost. In that situation, connections between servers cannot be established with the Local System account causing DataKeeper volume switchover commands, via the network, to be rejected. IT organizations requiring fault tolerance during a disaster recovery, including network disruptions, should not use the Local System

account.

DataKeeper Installation - Service Logon ID Type Selection:

SIOS DataKeeper for Windows

Service Setup

Service Logon Account Setup

The DataKeeper Service requires a logon account with Administrator privileges. The service logon account and password must be the same on all servers where DataKeeper is running. A Domain account is recommended.

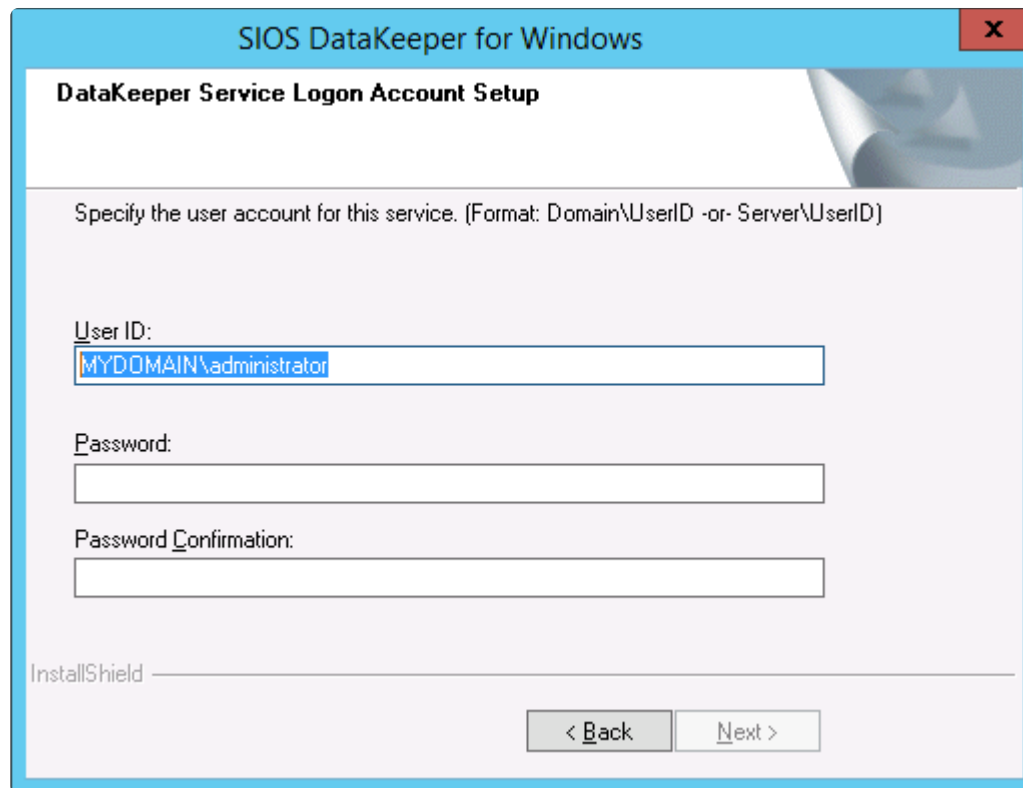
☒ Domain or Server account (recommended)

☐ LocalSystem account

InstallShield

< Back Next >

If a Domain or Server account is selected above, the DataKeeper Service Log On ID and Password Entry Form is displayed to enter that information.



The screenshot shows a Windows-style dialog box titled "SIOS DataKeeper for Windows" with a standard close button (X) in the top right corner. The main heading inside the dialog is "DataKeeper Service Logon Account Setup". Below this, a text label reads: "Specify the user account for this service. (Format: Domain\UserID -or- Server\UserID)". There are three input fields: the first is labeled "User ID:" and contains the text "MYDOMAIN\administrator"; the second is labeled "Password:" and is empty; the third is labeled "Password Confirmation:" and is also empty. At the bottom left, the text "InstallShield" is visible. At the bottom right, there are two buttons: "< Back" and "Next >".

It is recommended that the LifeKeeper and DataKeeper service accounts are synchronized on each system to ensure more reliable switchovers and failovers.



SIOS DataKeeper for Windows

Service Setup

Service Logon Account Setup

For optimum network connectivity DataKeeper and LifeKeeper services should use the same service logon accounts. Currently, the LifeKeeper service logon account does not match the DataKeeper service logon account. Make your selection below.

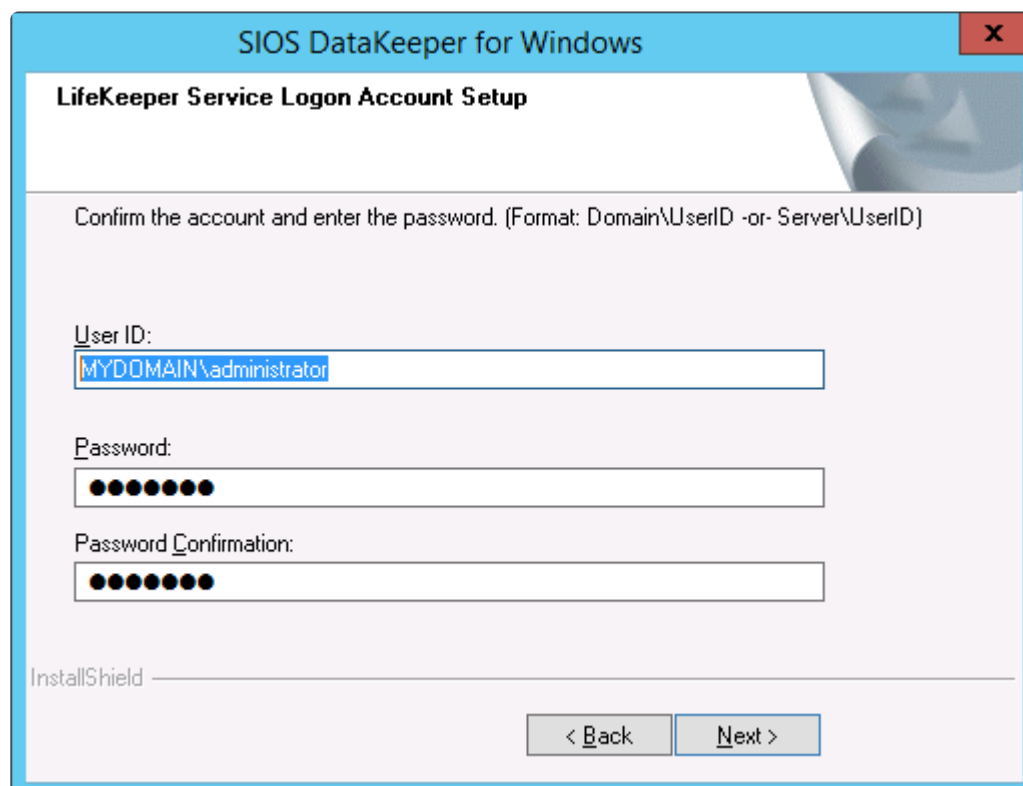
☒ Synchronize LifeKeeper Account (recommended)

☐ Do Not Synchronize Account

InstallShield

< Back Next >

LifeKeeper Service Logon:



SIOS DataKeeper for Windows

LifeKeeper Service Logon Account Setup

Confirm the account and enter the password. (Format: Domain\UserID -or- Server\UserID)

User ID:


Password:

Password Confirmation:

InstallShield

< Back Next >

If the DataKeeper Service has previously been configured with a Service Log On ID and Password, the setup program will omit the Service ID and Password selection dialogs. However, at any time, an administrator can modify the DataKeeper Service Log On ID and Password using the Windows Service Applet. Be sure to restart the DataKeeper Service after changing the Log On ID and/or Password.



The following table outlines these requirements:

Environment	DataKeeper Service Requirements	DataKeeper UI Requirements
Same Domain or	<ul style="list-style-type: none">Run the DK Service on all systems as the same account	<ul style="list-style-type: none">Log in as a domain admin and run the DK GUI

Trusted Domain Environment	<p>with the same credentials</p> <ul style="list-style-type: none"> • Okay to use the default = Local System Account 	<ul style="list-style-type: none"> • Or use "run as" Administrator option to run DK GUI
<p>Mixed Environment Servers in a Mixture of Domain and WorkGroup</p> <p>or</p> <p>Servers in Separate Domains</p>	<ul style="list-style-type: none"> • Create a local account on each system with same account name and password • Add this local account to the Administrator Group • Run the DK Service on all systems with the local account 	<ul style="list-style-type: none"> • Log in using the local account you created to run the DK Service • Run the DK GUI <p>You should also log on to all servers using this same Log On ID and Password (see related Known Issue).</p>
DataKeeper Cluster Edition Environment	<ul style="list-style-type: none"> • Create or use a domain account for use by the DataKeeper Service (preferred) <p>or</p> <ul style="list-style-type: none"> • Create a local account on each system with same account name and password 	<ul style="list-style-type: none"> • Log in using the local administrator account you created to run the DK Service • Run the DK GUI

	<ul style="list-style-type: none">• Add this local account to the Administrator Group• Run the DK Service on all systems with this local administrator account	
--	---	--

Firewall Configurations

SIOS DataKeeper cannot function properly if the firewall settings for the source and target machines are not configured correctly. This means you will need to configure a rule for inbound and outbound connections on each server running SIOS DataKeeper as well as any network firewalls that replication traffic must traverse.

During installation of SIOS DataKeeper, you will be prompted to allow the installer to configure your firewall rules needed by DataKeeper on Windows 2008 and Windows 2012. If you choose to allow the installer to make these changes, you will not need to configure your firewall manually. If you choose not to allow the installer to make these changes, you will need to configure your system manually as described in this section.

The ports that are required to be open for replication are as follows: 137, 138, 139, 445, 9999, plus ports in the 10000 to 10025 range, depending upon which volume letters you plan on replicating. The chart below shows the additional ports which must be open depending upon which drive letters you plan on replicating.

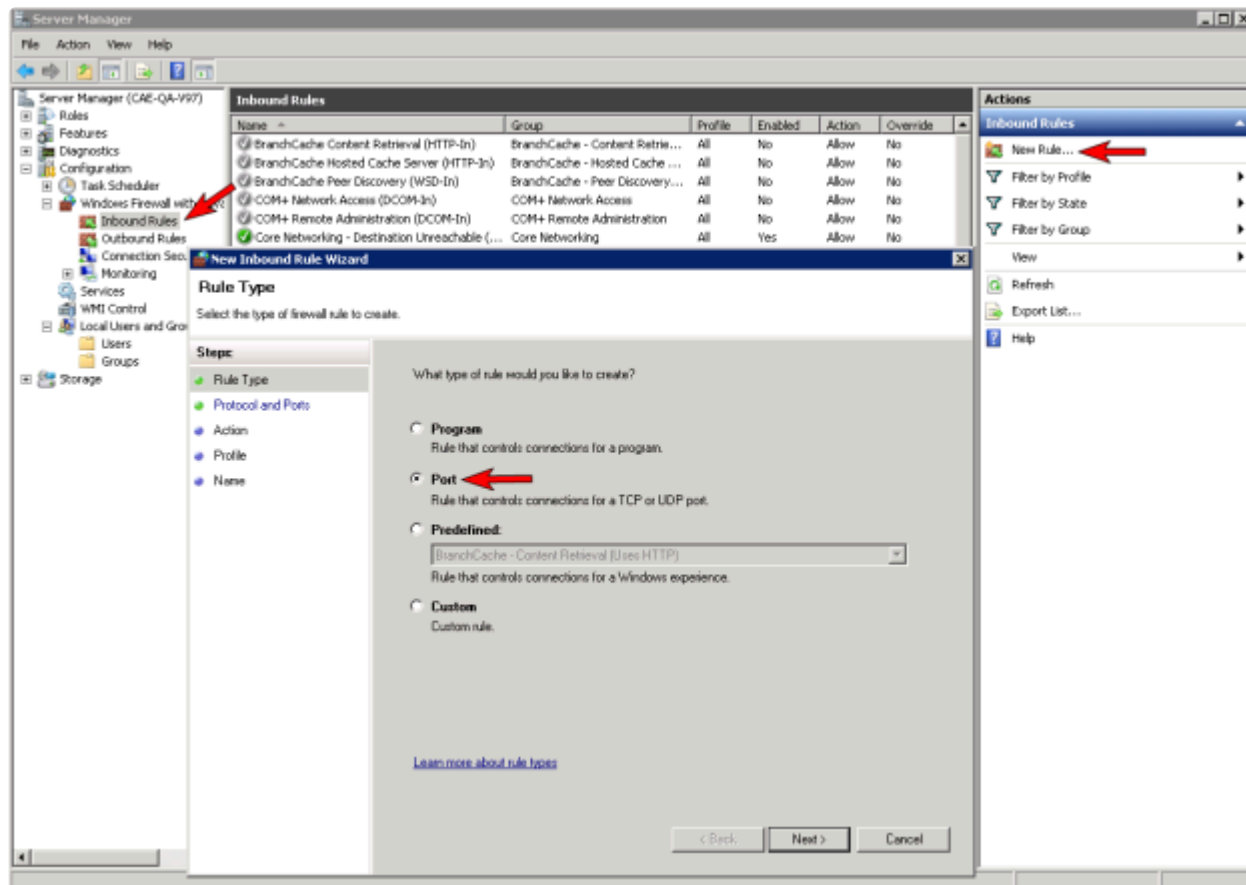
Port #	Volume Letter	Port #	Volume Letter
10000	A	10013	N
10001	B	10014	O
10002	C	10015	P
10003	D	10016	Q
10004	E	10017	R
10005	F	10018	S
10006	G	10019	T
10007	H	10020	U
10008	I	10021	V
10009	J	10022	W
10010	K	10023	X
10011	L	10024	Y

10012	M	10025	Z
-------	---	-------	---

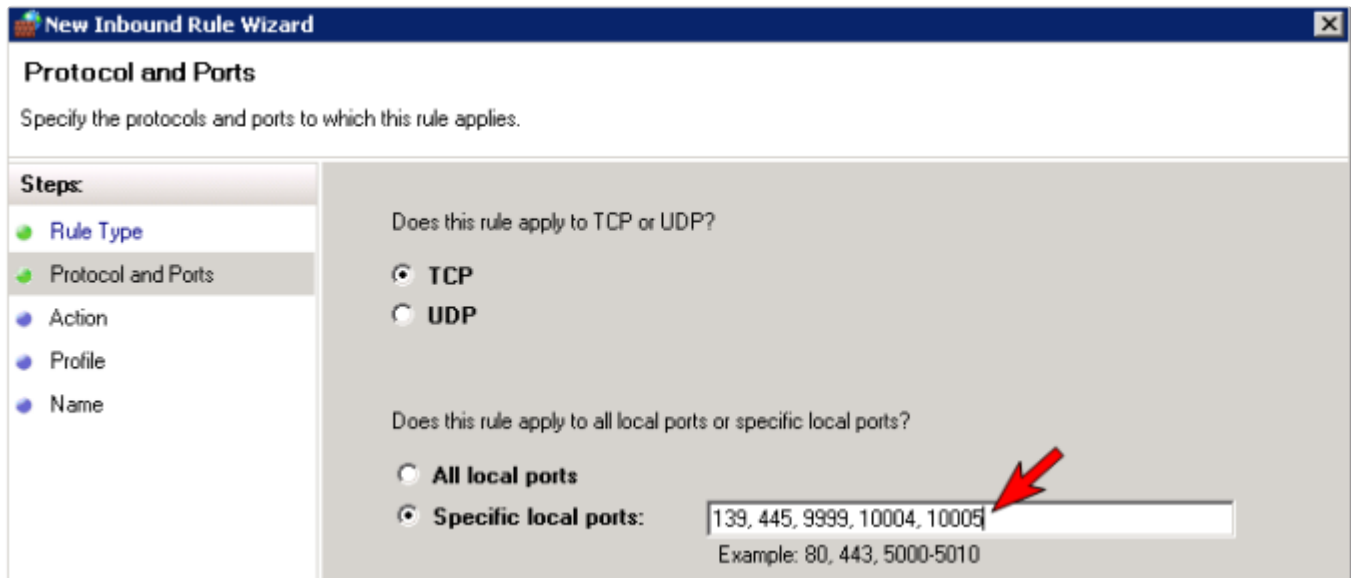
Configuring Microsoft's Windows Firewall with Advanced Security - Example

The exact steps required to configure the firewall for each cluster is as varied as each possible cluster configuration, but the following procedure and screen shots will give you one example to follow when using SIOS DataKeeper to replicate the E: and F: volumes. Note the Port # and Volume Letter table listings in the previous section.

1. Open Microsoft's **Defender Firewall with advanced security** and select **Inbound Rules** to create a rule for the TCP protocol as well as the UDP protocol.
2. Select **New Rule** from the **Actions** panel in the right column of the window. Select **Port** as the type of rule to be created. Select **Next**.




3. Select **TCP** for the type of protocol impacted by this rule. Select the **Specific local ports** button and enter the following ports: **139, 445, 9999, 10004** (for the E drive) and **10005** (for the F drive). Select **Next**.



The screenshot shows the 'New Inbound Rule Wizard' window, specifically the 'Protocol and Ports' step. The window title is 'New Inbound Rule Wizard'. The main heading is 'Protocol and Ports' with the instruction 'Specify the protocols and ports to which this rule applies.' On the left, a 'Steps' pane shows 'Rule Type', 'Protocol and Ports' (selected), 'Action', 'Profile', and 'Name'. The main area contains two questions. The first is 'Does this rule apply to TCP or UDP?' with radio buttons for 'TCP' (selected) and 'UDP'. The second is 'Does this rule apply to all local ports or specific local ports?' with radio buttons for 'All local ports' and 'Specific local ports' (selected). Below the second question is a text box containing '139, 445, 9999, 10004, 10005', with a red arrow pointing to it. Below the text box is the example text 'Example: 80, 443, 5000-5010'.

4. For the action, select **Allow the Connection**. Select **Next**.
5. For the profile, select **Domain, Private** and **Public** for the conditions when this rule applies. Select **Next**.



6. Enter a **Name** and **Description** for the new **Inbound Rule** and select **Finish**.

New Inbound Rule Wizard

Name

Specify the name and description of this rule.

Steps:

- Rule Type
- Protocol and Ports
- Action
- Profile
- Name**


Name: DataKeeperTCP

Description (optional): DataKeeper TCP Inbound Rule

< Back Finish Cancel

7. Select **New Rule** again to create the rule for **UDP protocol**. Select **Port** as the type of rule to be created. Select **Next**.
8. Select **UDP** for the type of protocol impacted by this rule. Select the **Specific local ports** button and enter the following ports in the Specific local ports field: **137, 138**. Select **Next**.
9. For the action, select **Allow the Connection**. Select **Next**.
10. For the profile, select **Domain, Private** and **Public** for the conditions when this rule applies. Select **Next**.
11. Enter a **Name** and **Description** for the new **Inbound Rule** and select **Finish**.

12. Your new DataKeeper rules will appear in the **Inbound Rules list** and the **Action** panel column. You can select the DataKeeper rule in the center panel and click the right mouse button to view the rule **Properties**.



High-Speed Storage Best Practices

Configure Bitmaps

If the DataKeeper default bitmap location (%ExtMirrBase%\Bitmaps) is not located on high-speed storage, you should move the bitmaps to a high-speed storage device in order to eliminate I/O bottlenecks with bitmap access. To do this, allocate a small disk partition, located on the high-speed storage drive, on which to place the bitmap files. Create the folder in which the bitmaps will be placed, and then [Relocate the bitmaps](#) (intent logs) to this location.

Disk Partition Size

The disk partition size must be big enough to contain all bitmap files for every mirror that will exist on your system. Each bit in the DataKeeper bitmap represents 64 KB of space on the volume, so to determine the bitmap size for a bitmap file, use the following formula:

$$\text{<volume size in bytes>} / 65536 / 8$$

Example:

For a 765 GB volume, convert the 765 GB to bytes

$$765 * 1,073,741,824 = 821,412,495,360 \text{ bytes}$$

Divide the result by 64K (65,536 bytes) to get the number of blocks/bits

$$821,412,495,360 / 65,536 = 12,533,760 \text{ blocks/bits}$$

Divide the resulting number of blocks/bits by 8 to get the bitmap file size in bytes

$$12,533,760 / 8 = 1,566,720$$

So a mirror of a 765 GB volume would require 1,566,720 bytes for its bitmap file, or approximately 1.5 MB.

A simple rule of thumb to use is that each GB of disk space requires 2 KB of bitmap file space.

Remember to reserve room for all mirror targets (if you have multiple target systems, each one needs a bitmap file). Also remember to reserve room for all mirrored volumes.


Handling Unmanaged Shutdown Issues

Unmanaged shutdowns due to power loss or other circumstances force a consistency check during the reboot. This may take several minutes or more to complete and can cause the drive not to reattach and can cause a dangling mirror. Use the ioAdministrator console to re-attach the drives or reboot the system again and make sure the check runs. For further information, refer to the ioXtreme User Guide for Windows.

Other Recommendations/Suggestions

- Check the Network Interface configuration settings. Increasing the Receive and Transmit buffers on the interfaces often improves replication performance. Other settings that may also affect your performance include: Flow Control, Jumbo Frames and TCP Offload. In certain cases, disabling Flow Control and TCP Offload can result in better replication performance. Enabling larger ethernet frames can also improve throughput.
- Check the location of the NICs on the bus (the slot that they're physically plugged into) as this can also affect the speed.
- Use Iometer, an I/O subsystem measurement and characterization tool available free on the internet, to test network throughput. Iometer can be set up in a client/server configuration and can test network throughput directly. Another alternative is to set up a file share using the replication IP address, and then copy large amounts of data over that share while monitoring the network throughput using Perfmon (Network Interface / Bytes Sent Per Second) or the Task Manager "Networking" tab.
- Make sure you have the latest drivers and firmware for the network adapters.

Configuration of Data Replication From a Cluster Node to External DR Site



Performance Tuning

Please refer to the following topics for ways to improve SIOS DataKeeper performance.

[SIOS DataKeeper Intent Log](#) - SIOS DataKeeper uses an intent log (also referred to as a bitmap file) to track changes made to the source, or to the target volume when the target is unlocked. This log is a persistent record of write requests which have not yet been committed to both servers. The intent log gives SIOS DataKeeper the ability to survive a source or target system failure or reboot without requiring a full mirror resync after the recovery of the system. There is a performance overhead associated with the intent log, since each write to the volume must also be reflected in the intent log file. To minimize this impact, it is recommended that the intent logs be stored on a physical disk that is not involved in heavy read or write activity. See [Relocation of Intent Log](#) for more information.

[High-Speed Storage Best Practices](#) - If the DataKeeper default bitmap location (%ExtMirrBase%\Bitmaps) is not located on high-speed storage, you should move the bitmaps to a high-speed storage device in order to eliminate I/O bottlenecks with bitmap access. To do this, allocate a small disk partition, located on a high-speed storage drive, on which to place the bitmap files. Create the folder in which the bitmaps will be placed, and then [Relocate the bitmaps](#) (intent logs) to this location.



Note: Ephemeral storage is recommended when the system is in Azure or AWS.

Disable "Automatically manage paging file size for all drives"

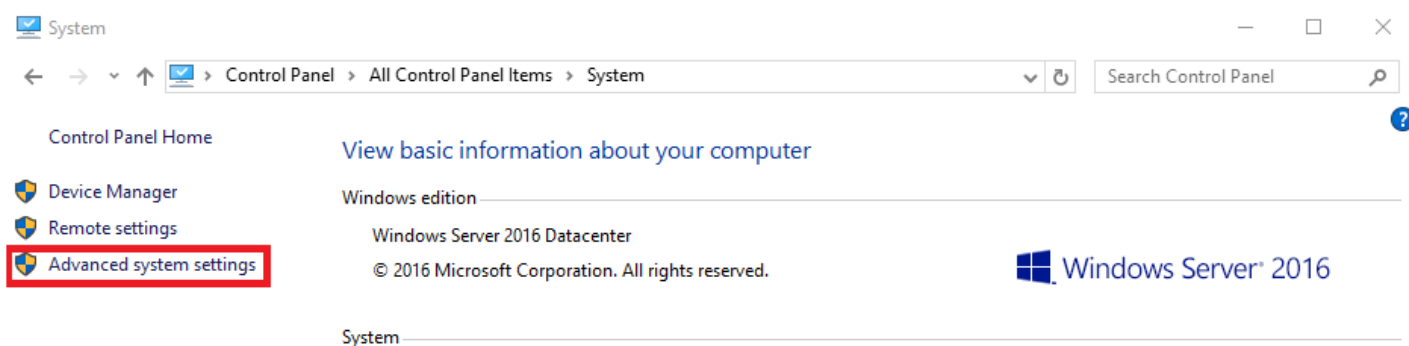
By default, Windows configures virtual memory so that page files are automatically created on volumes as the Operating System determines is best. This Virtual Memory setting is called "Automatically manage paging file size for all drives".

When this setting is enabled, page files sometimes are created by the Operating System on volumes that are part of DataKeeper mirrors. When this occurs, DataKeeper is not able to perform operations on the volume that are necessary for full protection. This setting needs to be disabled on all systems that have DataKeeper mirrors.


How to disable "Automatically manage paging file size for all drives"

This setting can be found in Control Panel "System" dialog.


First, click the **Advanced system settings** option:



From the **System Properties** dialog, choose the **Advanced** tab and click the **Settings** button in the **Performance** section.



Choose the **Advanced** tab in the Performance Options dialog, and click the **Change...** button in the Virtual Memory section.



In the Virtual Memory dialog, **uncheck** the "Automatically manage paging file size for all drives". Then configure pagefiles so that any DataKeeper-protected volumes have no page file configured.

Virtual Memory

☐ Automatically manage paging file size for all drives

Paging file size for each drive

Drive (Volume Label)	Paging File Size (MB)
C:	System managed
E: [VolE]	None
F: [VolF]	None
G: [New Volume]	None

Selected drive: C:
Space available: 16674 MB

☐ Custom size:
Initial size (MB):
Maximum size (MB):

☒ System managed size

☐ No paging file

Set

Total paging file size for all drives

Minimum allowed:	16 MB
Recommended:	1407 MB
Currently allocated:	704 MB

OK Cancel

* **Note:** Adjusting Windows Virtual Memory configuration can affect system performance. Be sure to consult Microsoft documentation on recommendations regarding these modifications.

WAN Considerations

Replicating data across the network to a remote server located miles away from the source server is the most common use of DataKeeper. Typically, this configuration relies on a WAN of some sort to provide the underlying network that DataKeeper uses to replicate the data. If the bandwidth of the WAN is limited, there are a number of additional factors to consider including:

[Initial Synchronization of Data Across the LAN/WAN](#)

[Compression](#)

[Bandwidth Throttle](#)

Initial Synchronization of Data Across the LAN or WAN

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. DataKeeper avoids almost all full resyncs by using its bitmap technology. However, the initial synchronization of the data, which occurs when the mirror is first created, cannot be avoided.

In WAN configurations, one way to avoid the initial full synchronization of data across the WAN is to configure both systems on a LAN, create the mirror and allow the initial full synchronization to occur across the LAN. Once the initial synchronization is complete, update the IP addresses for the source and target, which will place the mirror in the **Paused** state. Move the target system to its new location. Once the target system is in place, power it on and verify all network settings, including the IP address that was updated. On the source system, run the **CHANGEMIRRORENDPOINTS** command. The mirror will be **CONTINUED** and only a [partial resync](#) (the changes that have occurred on the source volume since the mirror was **PAUSED**) of the data is necessary to bring the TARGET volume in sync with the SOURCE.


✳ **Note:** This command supports changing the endpoints of a mirrored volume that is configured on 3 nodes or fewer. For configurations greater than three nodes, create mirrors with the final endpoint at the local site and use route adds to get the mirrors created and resynced before moving the server to the final location/address/DR site.

Example:

In the example below, a mirror is created locally in the primary site, and then the target will be moved to remote site. The source server is assigned the IP address 172.17.100.1, and the target server is assigned the IP address 172.17.100.2. The WAN network IP is 88.17.100.x,

- Using the DataKeeper UI, create a mirror on Volume X from 172.17.100.1 to 172.17.100.2. **Note:** Connecting to the target by name is recommended so DNS

name resolution later will automatically resolve to the new IP address.



Once the initial sync of the data is complete,

- Update the IP address for the network adapter for the source to 88.17.100.1 and update the IP address for the network adapter on the target to 88.17.200.2. This will place the mirror on the source side into the PAUSED state.
- Ship the target machine to its new location.
- Power on the target machine and verify all network settings, including the IP address updated above.
- On the source system, open a DOS command window and change directory to the DataKeeper directory by executing the following command:


```
cd %EXTMIRRBASE%
```

- Run the following command to update the existing mirror endpoints to the new IP addresses:

```
EMCMD 172.17.100.1 CHANGEMIRRORENDPOINTS X 172.17.100.2 88.17.100.1
```


88.17.200.2

- DataKeeper will resync the changes that have occurred on the source server while the target server was unreachable.
- When this partial resync is complete, the mirror will change to the **MIRRORING** state.



Verifying Data on the Target Volume

By design, DataKeeper locks the target volume. This prevents the file system from writing to the target volume while the replication is occurring. However, DataKeeper does provide a mechanism to unlock the target volume and allow read/write access to it while the mirror is still in place. There are two methods to do this:

1. Pause the mirror and unlock the target volume via the [Pause and Unlock](#) mirror option in the DataKeeper UI.
2. Use the DataKeeper command line interface (EMCMD) to pause the mirror ([PAUSEMIRROR](#)) and unlock the target volume ([UNLOCKVOLUME](#)).

Once unlocked, the target volume is completely accessible. When finished inspecting the target volume, be sure to continue the mirror to re-lock the

target volume and allow DataKeeper to resync any changes that occurred on the source volume while the mirror was paused. Any writes made to the target volume while it was unlocked will be lost when the mirror is continued.



If a reboot is performed on the target system while the target volume is unlocked, a full resync will occur when the target system comes back up.

Compression

DataKeeper allows the user to choose the compression level associated with each mirror. Enabling compression can result in improved replication performance, especially across slower networks. A compression level setting of 3-5 represents a good balance between CPU usage and network efficiency based on the system, network and workload.



Note: The compression level of a mirror can be changed after the mirror is created. See [Changing the Compression Level of an Existing Mirror](#).

Bandwidth Throttle

DataKeeper attempts to utilize all of the available network bandwidth. If DataKeeper is sharing the available bandwidth with other applications, you may wish to limit the amount of bandwidth DataKeeper is allowed to use. DataKeeper includes a feature called **Bandwidth Throttle** that will do this. The feature is enabled via a registry setting.

✿ **Note:** For additional information on both **Compression** and **Bandwidth Throttle**, see the topics below.

- [Registry Entries](#)
- [Changing the Compression Level of an Existing Mirror](#)

DataKeeper Administration

The topics in this section provide detailed instructions for performing DataKeeper administration tasks.

[DataKeeper Event Log Notification](#)

[Primary Server Shutdown](#)

[Secondary Server Failures](#)

[Extensive Write Considerations](#)

[CHKDSK Considerations](#)

[DKSUPPORT](#)

[DKHEALTHCHECK](#)

[Event Log Considerations](#)

[Using Disk Management](#)

[Registry Entries](#)

[Using EMCMD with SIOS DataKeeper](#)

[Using DKPwrShell with SIOS DataKeeper](#)


DataKeeper Event Log Notification

The **Event Log notification** is a mechanism by which one or more users may receive email notices when certain events occur. The **Windows Event Log** can be set up to provide notifications of certain DataKeeper events that get logged.

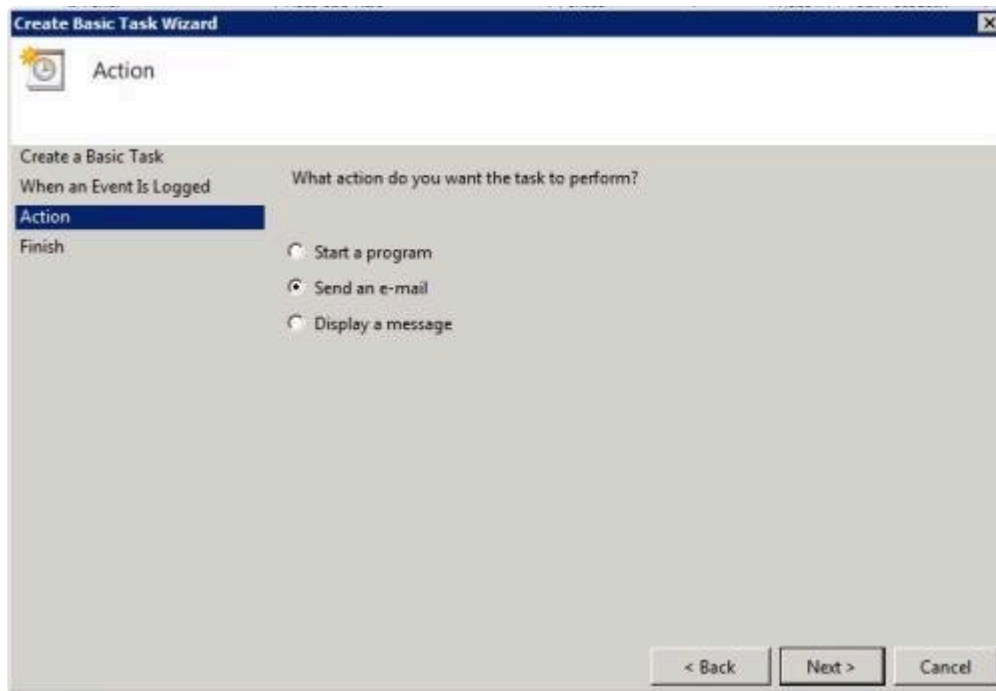
✿ **Note:** This option is only available for Windows Server 2008 R2.

To set up the **Windows Event Log email task** for DataKeeper events, perform the following steps:

1. Open Event Viewer, go to the System or Application log and highlight the event in which you want to be notified.
2. Right-click the event and select **Attach Task To This Event...**



3. Follow the **Task Wizard** directions, choosing the **Send an e-mail** option when prompted and filling in the appropriate information.



4. When you click **Finish** at the end of the **Task Wizard**, the new task will be created and added to your Windows schedule.

✿ **Note:** These email tasks will need to be set up on each node that will generate email notification.

Primary Server Shutdown

On a graceful shutdown of the source server, all pending writes to the target are completed. This ensures that all data is present on the target system.

On an unexpected source server failure, the [Intent Log](#) feature eliminates the need to do a full resync after the recovery of the source server. If the Intent Log feature is disabled or if SIOS DataKeeper detected a problem accessing the volume's Intent Log file, then a full resync will occur after the source server is restored to service.

Secondary Server Failures

In the event there is a failure affecting the secondary (target) system, the affected mirror is marked **Paused**. It is necessary to correct the condition that caused the secondary to fail and then resync the volumes. There are no write attempts made to the target after the secondary server fails.

When the secondary server comes back online after a failure, the source side of the mirror will automatically reconnect to the target side of the mirror. A partial resync follows.

Extensive Write Considerations

SIOS DataKeeper allows a volume that is being resynced to be accessed and written to. During Resync, DataKeeper makes multiple passes through the bitmap, synchronizing the data corresponding to each dirty bit found. While this is happening, writes that occur on the volume can dirty blocks that have already been synchronized. Under certain conditions, a resync can fail to complete due to constant write activity on the source volume.

In order to allow a resync to complete while writes are occurring on the source volume, DataKeeper may temporarily delay incoming writes in order to allow the last dirty blocks to be synchronized. The [ResyncLowWater](#) registry value specifies the maximum number of dirty blocks that can be present in the bitmap in order for DataKeeper to delay writes. The [ResyncBlockWritesTimeoutMs](#) registry value specifies the maximum amount of time that writes will be delayed (in milliseconds).

Note that DataKeeper will not delay writes until resync has attempted several resync passes - specifically, 10% of the value specified in the [MaxResyncPasses](#) registry value. For example, if MaxResyncPasses is set to the default of 200, DataKeeper will not delay writes until the 21st pass through the bitmap, and will only delay writes if the number of dirty blocks is less than or equal to ResyncLowWater.

If writes are delayed but resync does not complete within the amount of time specified by ResyncBlockWritesTimeoutMs, the resync fails, writes are permitted to complete, and the mirror goes into the Paused state.

To disable write delay completely, set ResyncLowWater to 0.

CHKDSK Considerations

If you must run CHKDSK on a volume that is being mirrored by SIOS DataKeeper, it is recommended that you first **pause** the mirror. After running CHKDSK, **continue** the mirror. A partial resync occurs (updating those writes generated by the CHKDSK) and mirroring will continue.

Failure to first **pause** the mirror may result in the mirror automatically entering the **Paused** state and performing a **Resync** while CHKDSK is in operation. While this will not cause any obvious problems, it will slow the CHKDSK down and result in unnecessary state changes in SIOS DataKeeper.

SIOS DataKeeper automatically ensures that volumes participating in a mirror, as either source or target, are not automatically checked at system startup. This ensures that the data on the mirrored volumes remains consistent.

Note: The bitmap file (for non-shared volumes) is located on the C drive which is defined by [BitmapBaseDir](#) as the default location. Running CHKDSK on the C drive of the Source system will cause an error due to the active bitmap file. Therefore, a switchover must be performed so that this Source becomes Target and the bitmap file becomes inactive. The CHKDSK can then be executed on this system as the new target (original source).

DKSUPPORT

DKSUPPORT .cmd, found in the <DataKeeper Installation Path>\SUPPORT directory, is used to collect important configuration information and event log files and put them in a zip file. SIOS Support Engineers will commonly request this zip file as part of the support process. To run this utility, double-click the file DKSUPPORT from the explorer window or right click the [DataKeeper Notification Icon](#) and then click on 'Gather Support Logs'.

This utility may also be executed from the command prompt using the following procedure.

- Launch an Administrator command prompt
- Type "cd %extmirrbase%"
- You will now be placed in the DataKeeper directory or c:\Program Files (x86) \SIOS\DataKeeper
- From the aforementioned directory type "cd support"
- From within the support directory, execute the following command "dksupport.cmd"
- Run this command on all systems that are participating in DataKeeper mirroring

The zip file will be created in the same Support directory, and can either be emailed to support@us.sios.com or File transferred (FTP) to support engineering

Note: This command may take some time to execute.

DKHEALTHCHECK

DKHealthCheck.exe, found in the \DKTools directory, is a tool that can provide basic mirror status and problem detection of mirror issues. SIOS Support may request that you run this tool as part of the Support process.

Note: DKHEALTHCHECK output is captured by DKSupport automatically and does not need to be run separately if you are already running DKSupport.

You can run this tool by right clicking the [DataKeeper Notification Icon](#) and then clicking on 'Launch Health Check' or by following the below procedure.

Open a command prompt

- Type `cd %extmirrbase%`
- You will now be placed in the DataKeeper directory or `c:\Program Files (x86)\SIOS\DataKeeper`
- From the aforementioned directory type `cd DKTools`
- From within the DKTools directory, execute the following command
`DKHealthCheck.exe`

The results of the tool can be copied and pasted from the command prompt and emailed to support@us.sios.com.

Alternatively, you may direct the output to a file, by running this command inside of the DKTools directory.

- `DKHealthCheck.exe > HealthCheck.txt`

This file can then be attached and sent as part of an email.

Note: This command may take some time to execute.

Event Log Considerations

It is important that SIOS DataKeeper be able to write to the Event Log. You should ensure that the Event Log does not become full. One way to accomplish this is to set the Event Log to overwrite events as needed:

1. Open the **Event Log**.
2. Right-click on **System Log** and select **Properties**.
3. Under **Log Size**, select **Overwrite Events as Needed**.

Using Disk Management

When using the Windows Disk Management utility to access SIOS DataKeeper volumes, please note the following:

- Using Disk Management to delete partitions that are being mirrored is not supported. Deleting a partition that is part of a SIOS DataKeeper mirror will yield unexpected results.
- Using Disk Management to change the drive letter assigned to a partition that is a part of a SIOS DataKeeper mirror is not supported and will yield unexpected results.
- The Windows Disk Management utility will take longer to start on the target node based on the number of drives. Because the Windows operating system has error condition retries built in when a volume is locked, the speed with which it starts on the "locked" target node is affected.

Registry Entries

The following registry entries are associated with the SIOS DataKeeper service or driver and can be viewed and edited using Regedt32.

* No values in any DataKeeper registry key should be modified unless listed they are listed here.

Registry Entries that MAY be Modified

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters

The SIOS DataKeeper driver uses the Parameters key and those below it. The values within the Parameters key (denoted with *) are global for all volumes on the system. The values under each of the Target IP registry keys (denoted with †) are specific to a mirror only. Values denoted with both * and † appear under both keys. (The target-specific value overrides the global value in this case.)

BandwidthThrottle †

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\BandwidthThrottle</i>		
Name	Type	Default Data
BandwidthThrottle	REG_DWORD	0
Specifies the maximum amount of network bandwidth (in kilobits per second) that a particular mirror is allowed to use. A value of 0 means unlimited.		

BitmapBaseDir *

*Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\BitmapBaseDir</i> *		
Name	Type	Default Data
BitmapBaseDir	REG_SZ	<i>C:\%EXTMIRRBASE%\Bitmaps</i> (usually <i>C:\Program Files\SIOS\DataKeeper\Bitmaps</i> but may be different when upgrading a system or if you install SIOS DataKeeper to a different path)

Specifies a directory where SIOS DataKeeper stores its Intent Log files. (**Note:** The drive letter must be in uppercase.) To disable the intent log feature, clear this registry entry (set it to an empty string) on all current and potential mirror endpoint servers. **Disabling the intent log requires a reboot on each of these systems in order for this setting to take effect.**

BitmapBytesPerBlock * †

Locations:

For New Mirrors: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\BitmapBytesPerBlock`

For Existing Mirrors: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\BitmapBytesPerBlock`

Note: If editing this entry under **Parameters**, all NEW mirrors created will inherit this value. If editing this entry under a **{Target IP}**, the value pertains to that one Target only. **{Target IP} values override Parameter values.**

Name	Type	Default Data
BitmapBytesPerBlock	REG_DWORD	65536 (0x10000)

Specifies the number of bytes that are represented as dirty in a [DataKeeper Intent](#) Log bitmap when a write request occurs. A single bit in the bitmap represents 65536 bytes, and the BitmapBytesPerBlock indicates the effective block size, which may be represented as multiple bits. Increasing this value can improve replication performance in some environments - in particular with workloads that perform sequential writes, on systems with relatively high-latency Bitmap storage. A larger block size means that fewer writes to the bitmap file will occur with sequential writes that are smaller than the adjusted block size. A larger block size will not noticeably help performance in environments where writes are primarily random, and may not help on systems with fast, low-latency bitmap storage. Also, a larger block size may result in larger amounts of data to resync in the event of a system failure.

Note: The minimum value of BitmapBytesPerBlock is 65536 - any value less than

this is treated as 65536. There is no maximum value enforced.

Note: BitmapBytesPerBlock does not affect the rate of mirror resync.

BlockWritesOnLimitReached * †

Locations:

For New Mirrors: *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\BlockWritesOnLimitReached*

For Existing Mirrors: *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\BlockWritesOnLimitReached*

Note: If editing this entry under **Parameters**, all NEW mirrors created will inherit this value. If editing this entry under a **{Target IP}**, the value pertains to that target only. **Any {Target IP} values override the global Parameter values.**

Name	Type	Default Data
BlockWritesOnLimitReached	REG_DWORD	0

This value determines what the behavior of a mirror is when the mirror's Write Queue reaches a defined limit (WriteQueueHighWater or WriteQueueByteLimit reached). If BlockWritesOnLimitReached is "0", the mirror is paused and a partial resync starts a short time later. If BlockWritesOnLimitReached is "1", the incoming write is delayed until there is space on the write queue for it. The mirror remains in the Mirroring state, but the application throughput slows down to match the speed of the network and the remote node's volume. After updating this registry value, execute the [READREGISTRY](#) command so that DataKeeper immediately starts using the new value.

Compression Level †

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\CompressionLevel</i>		
Name	Type	Default Data
CompressionLevel	REG_DWORD	0
<p>Specifies the compression level for the given mirror. Valid values are 0 to 9. Level 0 is "no compression". Values from 1 to 9 specify increasingly CPU-intensive levels of compression. Compression level 1 is a "fast" compression - it does not require as much CPU time to compress the data, but results in larger (less compressed) network packets. Level 9 is the maximum amount of compression - it results in the smallest network packets but requires the most CPU time. The level can be set to somewhere in between, to balance CPU usage and network efficiency based on your system, network and workload.</p>		

DontFlushAsyncQueue *

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\DontFlushAsyncQueue</i>		
Name	Type	Default Data
DontFlushAsyncQueue	REG_SZ	empty <drive letter>> [<drive letter>]
<p>Allows the user to specify a volume or volumes that should not flush their async queues when the driver receives a flush request. This value should contain the drive letter(s) of the volume(s) to which this applies. Drive letters may be adjacent to each other (i.e. XY), or space separated (i.e. X Y), with no colons. After updating this registry value, execute the READREGISTRY command so that DataKeeper immediately starts using the new value. (Note: When setting DontFlushAsyncQueue, data and database logs should be on the same partition.)</p>		

MaxResyncPasses *

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\MaxResyncPasses</i>		
Name	Type	Default Data
MaxResyncPasses	REG_DWORD	200 (0xc8)
<p>Specifies the maximum number of resync passes before SIOS DataKeeper temporarily interrupts the resync process while there is constant write activity on the source volume. Resync will be automatically resumed after this interruption, typically after 60 seconds. In every pass, SIOS DataKeeper marks</p>		

the volume blocks that were written to during the pass. In the next pass, it will send to the target only the marked blocks.

Note: In order for any changes to take effect a system reboot is required.

NotificationIconUpdateStatus *

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\NotificationIconUpdateStatus</i>		
Name	Type	Default Data
NotificationIconUpdateStatus	REG_SZ	true
Allows the user to turn off status update checks performed by all instances of the DataKeeper Notification Icon on a machine. This value should contain either true or false . Disabling the Notification Icon via its context menu will set this entry to false .		

PingInterval *

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\PingInterval</i>		
Name	Type	Default Data
PingInterval	REG_DWORD	3000 (0xBB8)
Specifies the interval in milliseconds between pings. Use a higher value for Wide Area Networks (WANs) or unreliable networks. Along with the MaxPingMisses , you may customize them to adjust mirroring to the network performance.		

ResyncBlockWritesTimeoutMs *†

Locations:		
For New Mirrors:		
<i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\ResyncBlockWritesTimeoutMs</i>		
AND		

For Existing Mirrors:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\ResyncBlockWritesTimeoutMs

Note: If editing this entry under **Parameters**, all NEW mirrors created will inherit this value. If editing this entry under **{Target IP}**, the value pertains to that target only. **Any {Target IP} values override the global Parameter values.**

Name	Type	Default Data
ResyncBlockWritesTimeoutMs	REG_DWORD	15000 (0x3a98)

Specifies the maximum amount of time that DataKeeper will delay writes during resync. See the [Extensive Write Considerations](#) topic for more information.

ResyncLowWater *†

Locations:

For New Mirrors:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\ResyncLowWater

AND

For Existing Mirrors:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\ResyncLowWater

Note: If editing this entry under **Parameters**, all NEW mirrors created will inherit this value. If editing this entry under **{Target IP}**, the value pertains to that target only. **Any {Target IP} values override the global Parameter**

values.

Name	Type	Default Data
ResyncLowWater	REG_DWORD	150 (0×96)
Specifies the maximum number of dirty blocks that can remain during resync before DataKeeper delays writes to allow the resync to complete. See the Extensive Write Considerations topic for more information.		

ResyncReads *

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\ResyncReads</i>		
Name	Type	Default Data
ResyncReads	REG_DWORD	20 (0×14)
This value represents the maximum number of disk blocks that can be in the process of being read and sent to the target system during mirror resynchronization. Changing this value may change the speed of mirror resynchronizations.		
Note: This tunable applies to synchronous and asynchronous mirrors.		

SnapshotLocation †

Location: <i>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\SnapshotLocation</i>		
Name	Type	Default Data
SnapshotLocation	REG_SZ	<drive letter>
Specifies the folder where the target snapshot file for this volume will be stored.		

TargetDispatchPort *

Locations:

On Target System:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\TargetDispatchPort

On Source System Creating Mirror to Above Target:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Targets{Target IP} (i.e. create a key whose name is the IP Address of the target system, or update the TargetDispatchPort value in that key if it already exists.)

Name	Type	Default Data
TargetDispatchPort	REG_DWORD	9999

There are two places where this should be set if you are changing the dispatch port from 9999. On the target system, place it in the *ExtMirr\Parameters* key. The new setting will apply to all existing and new targets on that server. **A target reboot is required when the target Parameters key has been changed for this setting to take effect.** On any source system that will be creating the mirror to this target, place it in the *ExtMirr\Parameters\Targets\{Target IP}* key if the mirror already exists. Create that key if it does not already exist. **Note:** Make sure the ports are the SAME on both the source and the target.

A firewall port must also be opened manually on all source and target servers for the new dispatch port to work.

TargetPortBase *

Location: ***HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\TargetPortBase***

Name	Type	Default Data
TargetPortBase	REG_DWORD	10000

Specifies the base TCP port number for target volume connections. This number may need to be adjusted if the default port is used by another service or is blocked by a firewall. The actual port that the target listens on is calculated as follows:

$$\text{Port} = \text{TargetPortBase} + (\text{Volume Letter} - \text{A:})$$

For example:

TargetPortBase = 10000

Volume Letter = H

$$\text{Port} = 10000 + (\text{H:} - \text{A:}) = 10007$$

TargetPortIncr *

Location: *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\TargetPortIncr*

Name	Type	Default Data
TargetPortIncr	REG_DWORD	256

Specifies the increment to the base TCP port number. This is used only when a TCP port is found to be in use. For example, if the target is attempting to listen on port 10005 and that port is in use, it will retry listening on port 10005 + **TargetPortIncr**.

TargetSnapshotBlocksize *

Location: *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\TargetSnapshotBlocksize*

Name	Type	Default Data
TargetSnapshotBlocksize	REG_DWORD	None

DataKeeper target snapshot uses a default block size of 64KB for all entries

that are written to the snapshot file. This block size can be modified by creating this TargetSnapshotBlocksize registry key.

The value should always be set to a multiple of the disk sector size, which is usually 512 bytes. Certain workloads and write patterns can benefit from changing the block size. For example, a volume that is written in a sequential stream of data (e.g. SQL Server log files) can benefit from a larger block size. A large block size results in fewer reads from the target volume when consecutive blocks are written. But a volume that is written in a random pattern may benefit from a smaller value or the default 64KB. A smaller block size will result in less snapshot file usage for random write requests.

VssQuiesceWaitTimeoutMs *

Location: *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\VssQuiesceWaitTimeoutMs*

Name	Type	Default Data
VssQuiesceWaitTimeoutMs	REG_DWORD	60000
Specifies the amount of time (in milliseconds) the DataKeeper service will wait for a VSS Snapshot Source Initiate request to complete. The VSS Snapshot Source Initiate request uses VSS to quiesce the data on snapshotted volumes.		

WriteQueueByteLimitMB †

Location: *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\WriteQueueByteLimitMB*

Name	Type	Default Data
WriteQueueByteLimitMB	REG_DWORD	0
Specifies the maximum number of bytes that can be allocated for the write queue of this mirror (expressed in megabytes - multiples of 1048576 bytes). The value "0" means "no limit". During periods of high disk write activity, if this mirror's write queue grows to a level which reaches the WriteQueueByteLimitMB, the SIOS DataKeeper driver momentarily pauses the mirror, drains the queue and automatically starts a partial resync. After updating this registry value, execute the READREGISTRY command so that DataKeeper immediately starts using		

the new value.

This value is used during transmission of volume data to the target, when the mirror is in the Mirroring state as well as when the mirror is in the Resync state. You should ensure that the ResyncReads value (see below), which specifies the number of 64KB (65536 byte) blocks that can be put on the Write Queue during resync, does not exceed the limit specified by WriteQueueByteLimitMB. Multiply ResyncReads by 65536, then divide by 1048576 - the resulting value must not exceed WriteQueueByteLimitMB if WriteQueueByteLimitMB is not set to 0.

This value can be used in conjunction with WriteQueueHighWater (see below). If both limits are set to nonzero values, then the mirror will be paused if either of them is reached. If one is set to 0 and one is not, then the nonzero limit is the only one that is enforced. If both are set to 0, then the mirror's write queue is not limited at all (this is not recommended - the WriteQueue uses Nonpaged memory).

Note: This tunable applies to synchronous and asynchronous mirrors. You can monitor the mirroring behavior using the SIOS DataKeeper Performance Monitor counters - specifically the Queue Current Bytes value - and set this limit accordingly.

WriteQueueHighWater * †

Locations:

For New Mirrors:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\WriteQueueHighWater

AND

For Existing Mirrors:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ExtMirr\Parameters\Volumes\{Volume GUID}\Targets\{Target IP}\WriteQueueHighWater

Note: If editing this entry under **Parameters**, all NEW mirrors created will inherit this value. If editing this entry under **Target**, the value pertains to that one Target only. **Any Target values override Parameter values.**

Name	Type	Default Data
WriteQueueHighWater	REG_DWORD	20000 (0x4e20)

Specifies the maximum number of write requests - not the number of bytes - that can be stored in this mirror's write queue. The value "0" means "no limit". During periods of high disk write activity, if this mirror's write queue length reaches this value, the SIOS DataKeeper driver momentarily pauses the mirror, drains the queue and automatically starts a partial resync. After updating this registry value, execute the [READREGISTRY](#) command so that DataKeeper immediately starts using the new value.

This value is used during transmission of volume data to the target, when the mirror is in the Mirroring state as well as when the mirror is in the Resync state. You should ensure that the ResyncReads value (see below), which specifies the number of blocks that can be put on the Write Queue during resync, does not exceed the limit specified by WriteQueueHighWater if WriteQueueHighWater is not set to 0.

This value can be used in conjunction with WriteQueueByteLimitMB. If both limits are set to nonzero values, then the mirror will be paused if either of them is reached. If one is set to 0 and one is not, then the nonzero limit is the only one that is enforced. If both are set to 0, then the mirror's write queue is not limited at all (this is not recommended - the WriteQueue uses Nonpaged memory).

Note: This tunable applies to synchronous and asynchronous mirrors. You can monitor the mirroring behavior using the SIOS DataKeeper Performance Monitor counters - specifically the Queue Current Length value - and set this limit accordingly.

Using EMCMD with SIOS DataKeeper

The EMCMD utility that ships with SIOS DataKeeper provides users with a command line method to manipulate the mirror. Because these scripts run in situations where the “normal” validation rules may not apply, EMCMD does not perform the same kinds of sanity checks that the user would experience using the SIOS DataKeeper User Interface. EMCMD simply passes commands to the SIOS DataKeeper Replication service allowing the service to make any decisions. It is this lack of checks that also makes this a useful diagnostic and support tool – though it is potentially dangerous for someone not very experienced with the inner workings of SIOS DataKeeper.

The following sections detail the operation of the EMCMD SIOS DataKeeper Command Line.

- Launch an Administrator command prompt
- Type “cd %extmirrbase%”
- You will now be placed in the DataKeeper directory or c:\Program Files (x86) \SIOS\DataKeeper

Note: The following style conventions will be utilized throughout.

<system>	Use the system’s NetBIOS name, IP address or fully qualified domain name to attach to a given system. You can also use a period (.) to attach to the local system where emcmd is being executed.
<drive>	Refers to the drive letter that is being referenced. EMCMD parses out everything after the first character, therefore, any “:” (colon) would be extraneous.

In some cases a series of EMCMD commands should be run to perform a function.

Example: To clean up a deleted mirror the following three commands should be run on each cluster node.

- emcmd . deletelocalmirroronly <volume letter of mirror to clean up>

- `emcmd . clearswitchover <volume letter of mirror to clean up>`
- `emcmd . updatevolumeinfo <volume letter of mirror to clean up>`

Then, you can recreate the mirror by using the `emcmd createmirror` command (example: `emcmd <address of source of mirror> createmirror <volume letter> <address of target of mirror> <Type of Mirror - either S for sync or A for async>`). This command will recreate the mirror and connect it to the existing DataKeeper Job.

Note: Run these commands with caution. If you have any questions please contact Support at support@us.sios.com.

Mirror State Definitions

BREAKMIRROR

CHANGEMIRRORENDPOINTS

CHANGEMIRRORTYPE

CLEARBLOCKTARGET

CLEARSNAPSHOTLOCATION

CLEARSWITCHOVER

CONTINUEMIRROR

CREATEJOB

CREATEMIRROR

DELETEJOB

DELETELOCALMIRRORONLY

[DELETEMIRROR](#)

[DROPSNAPSHOT](#)

[GETBLOCKTARGET](#)

[GETCOMPLETEVOLUMELIST](#)

[GETCONFIGURATION](#)

[GETEXTENDEDVOLUMEINFO](#)

[GETJOBINFO](#)

[GETJOBINFOFORVOL](#)

[GETMIRRORATYPE](#)

[GETMIRRORVOLINFO](#)

[GETREMOTEBITMAP](#)

[GETRESYNCSTATUS](#)

[GETSERVICEINFO](#)

[GETSNAPSHOTLOCATION](#)

[GETSOURCEMIRROREDVOLUMES](#)

[GETTARGETMIRROREDVOLUMES](#)

[GETVOLUMEDRVSTATE](#)

[GETVOLUMEINFO](#)

[ISBREAKUSERREQUESTED](#)

[ISPOTENTIALMIRRORVOL](#)

[LOCKVOLUME](#)

[MERGETARGETBITMAP](#)

[PAUSEMIRROR](#)

[PREPARETOBECOMETARGET](#)

[READREGISTRY](#)

[REGISTERCLUSTERVOLUME](#)

[RESTARTVOLUMEPIPE](#)

[RESYNCMIRROR](#)

[SETBLOCKTARGET](#)

[SETCONFIGURATION](#)

[SETSNAPSHOTLOCATION](#)

[STOPSERVICE](#)

[SWITCHOVERVOLUME](#)

[TAKESNAPSHOT](#)

[UNLOCKVOLUME](#)

[UPDATEJOB](#)

[UPDATEVOLUMEINFO](#)

Mirror State Definitions

The following numbers are used by the system to internally describe the various states. They are used by EMCMD, and they are also the state numbers found in event log entries.

-1: Invalid State

0: No Mirror

1: Mirroring

2: Mirror is resyncing

3: Mirror is broken

4: Mirror is paused

5: Resync is pending

BREAKMIRROR

EMCMD <system> **BREAKMIRROR** <volume letter> [<target system>]

This command forces the mirror into a **Broken** state. Breaking the mirror will cause a full resync to occur when the mirror is continued or resynced.

The parameters are:

<system>	This is the source system of the mirror to break. Running the BREAKMIRROR command on the target has no effect.
<volume letter>	The drive letter of the mirror that you want to break.
<target system>	This is the IP address of the target system of the mirror to break. This optional parameter may be used if multiple targets are associated with the mirror. If not specified, the mirror will be broken to all targets.

CHANGEMIRRORENDPOINTS

```
EMCMD <NEW source IP> CHANGEMIRRORENDPOINTS <volume letter>
<ORIGINAL target IP> <NEW source IP> <NEW target IP>
```

This command is used to change the replication IP addresses within systems that are already part of a DataKeeper job for the given volume.

Note: This command supports changing the endpoints of a mirrored volume that is configured on 3 nodes or fewer. If your configuration consists of more than three nodes, the mirrors must be deleted and recreated.

Refer to the examples below.

See "[WAN Considerations](#)" and "[Initial Synchronization of Data Across the LAN/WAN](#)" in the "[Configuration](#)" section.

<new source IP> OR <system name>	This is the system that has the new source IP address available for the mirror.
<volume letter>	The drive letter of the mirror to be changed.
<original target IP>	The previous IP address of the target system.
<new source IP>	The new IP address of the source system.
<new target IP>	The new IP address of the target system.

Notes:

- A job may contain multiple volumes and multiple mirrors. The CHANGEMIRRORENDPOINTS command will modify endpoints on one mirror each time it is used. For a 1×1 mirror (1 source, 1 target), only one command is required. For a 2×1 mirror (2 nodes with a shared volume with one target node) or a 1×1×1 (1 source, two target nodes), two commands are required to

change the necessary mirror endpoints.

- If an existing mirror whose endpoints are being changed is currently an active mirror, it must be put into the [Paused](#), [Broken](#) or **Resync Pending** state before the endpoints can be changed.

! **CAUTION:** Using the Break command will cause a **full resync**. It is recommended that the mirror be [Paused](#) instead.

- Before making changes, it will be helpful to display **Job Information** for the volume. For example, `emcmd . getJobInfoForVol D .`
- While making endpoint changes, the **Job** icon in the DataKeeper GUI may turn red. However, it will return to green after the `ContinueMirror` command is performed.

In the following examples, we move mirrors from the 172.17.103 subnet to the 192.168.1 subnet. The basic steps are as follows:

1. **Display job information** for the volume
2. **Pause the Mirror** using the EMCMD command line
3. **Change the IP address** on the system(s) (if necessary)

*** IMPORTANT:** If you haven't already done so, prior to performing the **CHANGEMIRRORENDPOINTS** command, **update the IP addresses** for the source and target. This will automatically place the mirror into the **Paused** state.

4. **Run EMCMD CHANGEMIRRORENDPOINTS** to change to the new IP address
5. **Run EMCMD CONTINUEMIRROR** to resume mirroring

! **CAUTION:** If the source system is rebooted before the mirrors are continued a full resync will occur on the mirrored volumes.

1×1 Mirror CHANGEMIRRORENDPOINTS Command Example

For a 1×1 mirror (source and target only), one command is required.

```
emcmd SYS1.MYDOM.LOCAL getJobInfoForVol D

ID = caa97f9f-ac6a-4b56-8f25-20db9e2808a8

Name = Mirr Vol D

Description = Mirror Volume D

MirrorEndpoints =
SYS3.MYDOM.LOCAL;D;172.17.103.223;SYS1.MYDOM.LOCAL;E;172.17.103.221;A

emcmd SYS1.MYDOM.LOCAL PauseMirror D

emcmd SYS1.MYDOM.LOCAL ChangeMirrorEndpoints D 172.17.103.223
192.168.1.221 192.168.1.223

emcmd SYS1.MYDOM.LOCAL getJobInfoForVol D

. . .

MirrorEndpoints =
SYS3.MYDOM.LOCAL;D;192.168.1.223;SYS1.MYDOM.LOCAL;D;192.168.1.221;A

emcmd SYS1.MYDOM.LOCAL ContinueMirror D
```

2×1 Mirror CHANGEMIRRORENDPOINTS Command Example

For a 2×1 mirror that includes a shared source volume and a target volume, two commands are required.

```
emcmd SYS1.MYDOM.LOCAL getJobInfoForVol E

ID = caa97f93e-ac6a-4b56-8f25-20db9e2808a8

Name = Mirr Vol E
```

```

Description = Mirror Volume E

MirrorEndPoints = SYS1.MYDOM.LOCAL;E;0.0.0.0;SYS2.MYDOM.LOCAL;E
;0.0.0.0;D

MirrorEndPoints =
SYS3.MYDOM.LOCAL;E;172.17.103.223;SYS2.MYDOM.LOCAL;E;172.17.103.222;A

MirrorEndPoints =
SYS3.MYDOM.LOCAL;E;172.17.103.223;SYS1.MYDOM.LOCAL;E;172.17.103.221;A

emcmd SYS1.MYDOM.LOCAL PauseMirror E

emcmd SYS1.MYDOM.LOCAL ChangeMirrorEndPoints E 172.17.103.223
192.168.1.221 192.168.1.223

emcmd SYS2.MYDOM.LOCAL ChangeMirrorEndPoints E 172.17.103.223
192.168.1.222 192.168.1.223

emcmd SYS1.MYDOM.LOCAL getJobInfoForVol E

. . .

MirrorEndPoints =
SYS1.MYDOM.LOCAL;E;0.0.0.0;SYS2.MYDOM.LOCAL;E;0.0.0.0;D

MirrorEndPoints =
SYS3.MYDOM.LOCAL;E;192.168.1.223;SYS2.MYDOM.LOCAL;E;192.168.1.222;A

MirrorEndPoints =
SYS3.MYDOM.LOCAL;E;192.168.1.223;SYS1.MYDOM.LOCAL;E;192.168.1.221;A

emcmd SYS1.MYDOM.LOCAL ContinueMirror E

```

1×1×1 Mirror CHANGEMIRRORENDPOINTS Command Example

For a 1×1×1 mirror that includes 2 Target volumes, 2 commands are required.

```
emcmd SYS1.MYDOM.LOCAL getJobInfoForVol J
```

```
ID = caa97f93j-ac6a-4b56-8f25-20db9j2808a8
```

```
Name = Mirr Vol J
```

```
Description = Mirror Volume J
```

```
MirrorEndpoints =
```

```
SYS1.MYDOM.LOCAL;J;172.17.103.221;SYS3.MYDOM.LOCAL;J;172.17.103.223;A
```

```
MirrorEndpoints =
```

```
SYS3.MYDOM.LOCAL;J;172.17.103.223;SYS2.MYDOM.LOCAL;J;172.17.103.222;A
```

```
MirrorEndpoints =
```

```
SYS1.MYDOM.LOCAL;J;172.17.103.221;SYS2.MYDOM.LOCAL;J;172.17.103.222;A
```

In this example the system "SYS3.MYDOM.LOCAL" will be moved to another site.

SYS1 and SYS2 will now use a new subnet (192.168.1.*) to communicate with SYS3.

However, SYS1 and SYS2 will continue using 172.17.103.* to communicate with each other.

```
emcmd SYS1.MYDOM.LOCAL PauseMirror J
```

```
emcmd SYS1.MYDOM.LOCAL ChangeMirrorEndpoints J 172.17.103.223
192.168.1.221 192.168.1.223
```

```
emcmd SYS2.MYDOM.LOCAL ChangeMirrorEndpoints J 172.17.103.223
192.168.1.222 192.168.1.223
```

```
emcmd SYS1.MYDOM.LOCAL getJobInfoForVol J
```

```
. . .
```

```
MirrorEndpoints =
```

```
SYS1.MYDOM.LOCAL;J;192.168.1.221;SYS3.MYDOM.LOCAL;J;192.168.1.223;A
```

```
MirrorEndpoints =
```

```
SYS3.MYDOM.LOCAL;J;192.168.1.223;SYS2.MYDOM.LOCAL;J;192.168.1.222;A
```

```
MirrorEndpoints =  
SYS1.MYDOM.LOCAL;J;172.17.103.221;SYS2.MYDOM.LOCAL;J;172.17.103.222;A  
  
emcmd SYS1.MYDOM.LOCAL ContinueMirror J
```

CHANGEMIRRORTYPE

EMCMD <system> CHANGEMIRRORTYPE <volume letter> <remote ip> <A/S>

This command is used to change the mirror type of a mirror that is part of a DataKeeper job.

Refer to the examples below.

See [Synchronous and Asynchronous Mirroring](#) for information about the supported DataKeeper mirror types.

<system>	The source or target system on which to initiate the changing of the mirror type.
<volume letter>	The drive letter of the mirror to be changed.
<remote IP>	The IP address of the remote system.
<A/S>	The new mirror type (Asynchronous or Synchronous).

Notes:

- A job may contain multiple volumes and multiple mirrors. The CHANGEMIRRORTYPE command will modify the type of one mirror each time it is used.
- The volume must be online on each system in the mirror to change the mirror type of an existing mirror.
- The mirror type of an existing mirror can be changed while the mirror is in the active Mirroring state. The type change takes effect immediately.
- The mirror type of non-existing mirrors can be changed. See the 1×1×1 example below.

- The mirror type of a mirror that is in the Split-Brain state cannot be changed - the Split Brain must be resolved first.
- If a job contains multiple mirrors, individual mirror types can be modified. Having mixed mirror types within a job, and within the mirrors for an individual volume in the job, is supported.

1×1 Mirror CHANGEMIRRORTYPE Command Example

```
emcmd SYS1.MYDOM.LOCAL getJobInfoForVol D
```

```
ID = caa97f9f-ac6a-4b56-8f25-20db9e2808a8
```

```
Name = Mirr Vol D
```

```
Description = Mirror Volume D
```

```
MirrorEndpoints =
```

```
SYS3.MYDOM.LOCAL;D;172.17.103.223;SYS1.MYDOM.LOCAL;E;172.17.103.221;A
```

```
emcmd SYS1.MYDOM.LOCAL ChangeMirrorType D 172.17.103.223 S
```

The above example changes the mirror of D: between SYS1 and SYS3 to Synchronous.

1×1×1 Mirror CHANGEMIRRORTYPE Command Example

```
emcmd SYS1.MYDOM.LOCAL getJobInfoForVol J
```

```
ID caa97f93j-ac6a-4b56-8f25-20db9j2808a8
```

```
Name = Mirr Vol J
```

```
Description = Mirror Volume J
```

```
MirrorEndpoints =
```

```
SYS1.MYDOM.LOCAL;J;172.17.103.221;SYS2.MYDOM.LOCAL;J;172.17.103.222;A
```

```
MirrorEndpoints =
```

```
SYS1.MYDOM.LOCAL;J;172.17.103.221;SYS3.MYDOM.LOCAL;J;172.17.103.223;A
```

```
MirrorEndpoints =  
SYS3.MYDOM.LOCAL;J;172.17.103.223;SYS2.MYDOM.LOCAL;J;172.17.103.222;A  
  
emcmd SYS1.MYDOM.LOCAL GetMirrorVolInfo J  
  
J: 1 SYS1.MYDOM.LOCAL 172.17.103.222 1  
  
J: 1 SYS1.MYDOM.LOCAL 172.17.103.223 1  
  
emcmd SYS1.MYDOM.LOCAL ChangeMirrorType J 172.17.103.222 S  
  
emcmd SYS1.MYDOM.LOCAL ChangeMirrorType J 172.17.103.223 S  
  
emcmd SYS2.MYDOM.LOCAL ChangeMirrorType J 172.17.103.223 S
```

In this example, all mirror types will be changed to Synchronous. The third command changes the mirror type of the non-existing mirror between SYS2 and SYS3.

CLEARBLOCKTARGET

EMCMD <system> CLEARBLOCKTARGET <volume letter>

This command sets the state of the block target flag to FALSE. The block target flag when set to FALSE will allow that system to become a target for the selected volume. This command is for internal use only. No output is produced when running this command.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume for which you want to set the state of the block target flag to FALSE.

CLEARSNAPSHOTLOCATION

EMCMD <system> CLEARSNAPSHOTLOCATION <volume letter>

This command clears the snapshot location (directory path) for the given volume on the given system. Once this command executes successfully, snapshots will be disabled for the given volume.

The parameters are:

<system>	This is the system name/IP address of snapshot location.
<volume letter>	This is the drive letter of the volume to be snapshotted.

Sample output:

```
Status = 0
```

When the command is successful, it will return a status of 0. Otherwise, it will report a non-zero status.

CLEARSWITCHOVER

EMCMD <system> CLEARSWITCHOVER <volume letter>

This command should be run on a target system where a mirror has been previously deleted with the [DELETELOCALMIRRORONLY](#) command and now needs to be re-established. This command clears the SIOS DataKeeper switchover flag that is set for a volume that has been deleted from the Target role using DELETELOCALMIRRORONLY. If you delete a target using DELETELOCALMIRRORONLY and do not run CLEARSWITCHOVER, you will not be able to re-establish a mirror target unless you reboot the system.

<system>	This is the target system where you just ran DELETELOCALMIRRORONLY.
<volume letter>	The drive letter of the mirror.

CONTINUEMIRROR

EMCMD <system> CONTINUEMIRROR <volume letter> [<target system>]

This command forces a paused or broken mirror to resume mirroring. On successful completion of the resync (full or partial), the mirror state is changed to **Mirroring**. This command will not automatically relock the target volume if it is unlocked.

Note: If target volume is unlocked, it must be [rellocked](#) prior to running this command.

The parameters are:

<system>	This is the source system of the mirror to resume mirroring.
<volume letter>	The drive letter of the mirror that you want to resume mirroring.
<target system>	This is the IP address of the target system of the mirror to resync. This optional parameter may be used if multiple targets are associated with the mirror. If not specified, a resync will be performed to all targets.

CREATEJOB

```
EMCMD . CREATEJOB <JobName> <Description> <FQDN Source>  
<DrvLetter1> <IP SourceReplication> <FQDN Target> <DrvLetter2>  
<IP Target for Replication><MirrorType> . . .
```

This command is for internal use only.

CREATEMIRROR

```
EMCMD <system> CREATEMIRROR <volume letter> <target system>
<type> [options]
```

This command creates a mirror between two machines, using the same drive letter on each.

The parameters are:

<system>	This is the IP address of the source system (see Note below).
<volume letter>	This is the drive letter that is being mirrored. This will be both the source and target drive letter.
<target system>	This is the IP address of the target system (see Note below).
<type>	<p>This is the type of mirror, where type is a single character:</p> <p>A - Create an Asynchronous Mirror</p> <p>S - Create a Synchronous Mirror</p>
[options]	<p>Optional arguments that specify behavior deviant from the norm. These can be OR'd together to create a set of options (add decimal values - for example, for option 1 + option 4, place a 5 in the command). They are:</p> <p>1: Create the mirror without doing a full resync operation.</p> <p>2: Do not wait for the target side of the mirror to be created before returning.</p> <p>4: Create with boot-time restrictions in place - essentially treat the create as you would a mirror re-establishment as part of the boot process. This option will check to see if the remote system is</p>

	already a source and fail the creation if it determines that it was a source.
--	---

✿ **NOTE:** Disk sector size must match on both source and target volumes. See [Sector Size](#) for more information.

✿ **NOTE:** Both source and target IP addresses must be of the same protocol. A mirror can only be created using two IPV4 or two IPV6 addresses. DataKeeper does not currently support mirror endpoints with different protocols.

IPv4 Example:

```
EMCMD 192.168.1.1 CREATEMIRROR E 192.168.1.2 A 5
```

IPv6 Example:

```
EMCMD 2001:5c0:110e:3304:a6ba:dbff:feb2:f7fd CREATEMIRROR F  
2001:5c0:110e:3304:a6ba:dbff:feb2:afd7 A 5
```

DELETEJOB

```
EMCMD <system> DELETEJOB <JobId>
```

This command is for internal use only.

DELETELOCALMIRRORONLY

EMCMD <system> DELETELOCALMIRRORONLY <volume letter> [<target system>]

This command deletes the mirror only on the <system> it is issued on. It handles the case when a mirror ends up with a target and no source or source and no target.

The parameters are:

<system>	This can be either the source or the target system.
<volume letter>	The drive letter of the mirror that you want to delete.
<target system>	This is the IP address of the target system of the mirror to delete. This optional parameter may be used if multiple targets are associated with the mirror.

DELETEMIRROR

EMCMD <system> DELETEMIRROR <volume letter> [<target system>]

This command deletes the mirror from both the source and the target if <system> is a source. If <system> is a target, it will delete the target side of the mirror only if the source system is down.

The parameters are:

<system>	This can be either the source or the target system.
<volume letter>	The drive letter of the mirror that you want to delete.
<target system>	This is the IP address of the target system of the mirror to delete. This optional parameter may be used if multiple targets are associated with the mirror. If not specified, the mirror will be deleted for all targets.

DROPSNAPSHOT

EMCMD <system> DROPSNAPSHOT <volume letter> [<volume letter> ...]

This command will notify DataKeeper to lock the volume and clean up the snapshot files that it created.

The parameters are:

<system>	This is the IP address of the system containing the snapshot.
<volume letter>	This is the drive letter of the snapshotted volume on the target server. If dropping multiple snapshots, the drive letters should be separated by spaces.

GETBLOCKTARGET

EMCMD <system> GETBLOCKTARGET <volume letter>

This command provides the current state of the block target flag, either TRUE or FALSE. The block target flag if set to TRUE will prevent that system from ever becoming a target for the selected volume. This command is for internal use only.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume for which you want to retrieve the state of the block target flag.

Sample output:

```
c:> EMCMD . GETBLOCKTARGET E
```

```
FALSE
```

GETCOMPLETEVOLUMELIST

EMCMD <system> GETCOMPLETEVOLUMELIST

This command displays information on all volumes eligible to be mirrored or already in a mirror. Sample output:

Volume 1 information:

Volume Root	= F:
Volume Label	= New Volume
Volume File System	= NTFS
Volume Total Space	= 2151608320
Mirror Role	= 01
Number of targets	= 2
Target 0 information:	
Volume State	= 0001
Target System	= 10.1.1.133
Target Drive Letter	= F
Target 1 information:	
Volume State	= 0002
Target System	= 10.1.1.134
Target Drive Letter	= F

GETCONFIGURATION

EMCMD <system> GETCONFIGURATION <volume letter>

This command retrieves and displays the net alert settings (also referred to as "volume attributes") for the volume.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume you want information on.

Sample output:

**** Calling GetConfiguration [Volume F] ****

All Net Alert bit	IS NOT enabled
Net Alert	IS NOT enabled
Broken State Alert	IS NOT enabled
Resync Done Alert	IS NOT enabled
Failover Alert	IS NOT enabled
Net Failure Alert	IS NOT enabled
LK Config	IS NOT enabled
Auto Resync	IS NOT enabled
MS Failover Cluster Config	IS NOT enabled
Shared Volume	IS NOT enabled

GETEXTENDEDVOLUMEINFO

EMCMD <system> GETEXTENDEDVOLUMEINFO <volume letter>

This command returns extended volume information about the selected volume such as disk signature, physical disk offset and internal disk id.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume you want information on.

Sample output:

-----EXTENDED INFO -----

Physical Disk Signature = {217abb5a-0000-0000-0000-000000000000}

Physical Disk Offset = 32256

Internal Disk ID = 0xf2fa

GETJOBINFO

EMCMD <system> GETJOBINFO [<JobId>]

This command displays job information for a specific JobId or all defined jobs.

GETJOBINFOFORVOL

EMCMD <system> GETJOBINFOFORVOL <DrvLetter>[<FullSysname>|<IP>]

This command displays job information related to a specific volume on a specific system.

GETMIRRORTYPE

EMCMD <system> GETMIRRORTYPE <volume letter>

This command provides a numeric output of the type of mirror.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The driver letter of the volume you want information on.

Output format:

```
c:> EMCMD . GETMIRRORTYPE F
```

```
Target system 10.1.1.133, Type 2
```

```
Target system 10.1.1.134, Type 2
```

Mirror Type:

```
-1:   Invalid Type (EMCMD cannot get the requested information.)
```

```
0:   No mirror
```

```
1:   Synchronous Mirror
```

```
2:   Asynchronous Mirror
```

GETMIRRORVOLINFO

EMCMD <system> GETMIRRORVOLINFO <volume letter>

This command provides a very terse output of the state of mirror. The command GETMIRRORVOLINFO can return multiple lines of output (one per target). It provides essentially the same information as the [GETVOLUMEINFO](#) command does.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The driver letter of the volume you want information on.

Sample output:

```
c:> EMCMD . GETMIRRORVOLINFO F
```

```
F: 1 CARDINAL10.1.1.133 1
```

```
F: 1 CARDINAL10.1.1.134 1
```

Output format:

```
[Volume Letter] {Mirror Role} [Source System] [Target System] [Mirror State]
```

Mirror Role: 1 = source; 2 = target

Mirror Type:

```
-1: Invalid State
```

```
0: No mirror
```

```
1: Mirroring
```

```
2: Mirror is resyncing
```

```
3: Mirror is broken
```

4: Mirror is paused

5: Resync is pending

GETREMOTEBITMAP

```
EMCMD <system> GETREMOTEBITMAP <volume letter> <targetsystem>  
<local file>
```

This command is for internal use only.

GETRESYNCSTATUS

EMCMD <system> GETRESYNCSTATUS <volume letter>

This command returns information indicating the overall status of a resync operation.

The parameters are:

<system>	This can only be performed on the Source system.
<volume letter>	The drive letter of the volume you want to set the configuration on.

Sample output:

Resync Status for Volume F:

Target 0 (Target System 10.1.1.133)

ResyncPhase : 3
BitmapPass : 1
NumberOfBlocks : 32831
DirtyBlocks : 0
CurrentBlock : 0
NewWrites : 1803
ResyncStartTime : Fri Nov 05 13.57.51 2008
LastResyncTime : Fri Nov 05 13.57.51 2008

Target 1 (Target System 10.1.1.134)

ResyncPhase : 2
BitmapPass : 0
NumberOfBlocks : 32831
DirtyBlocks : 2124
CurrentBlock : 29556
NewWrites : 0
ResyncStartTime : Fri Nov 05 15:09:47 2008
LastResyncTime : Fri Nov 05 15:09:47 2008

The **ResyncPhase** is used internally and has little meaning outside of the

development environment. The values are: 0 (unknown), 1 (initial), 2 (update), and 3 (done).

The **BitmapPass** is the number of times we have passed through the bitmap indicating the number of dirty blocks. We count from zero. If we do a resync in one pass, then this never increments.

The **NumberOfBlocks** is the number of 64K data blocks on the volume.

The **DirtyBlocks** parameter is the number of blocks that the bitmap indicates need to be updated (and have not already been).

The **CurrentBlock** parameter indicates the current location in the bitmap.

The **NewWrites** parameter indicates the number of writes that have occurred on the volume since we have been resyncing.

The **ResyncStartTime** and **LastResyncTime** parameters describe the time that the resync was begun and the last time a resync write operation was sent across the network.

GETSERVICEINFO

EMCMD <system> GETSERVICEINFO

This command retrieves version and other information about the SIOS DataKeeper service and driver that is running on the specified machine.

The parameters are:

<system>	This can be either the source or the target systems.
----------	--

Sample output:

```
Service Description: = SIOS DataKeeper Service
Service Build Type: = Release
Service Version = 7.0
Service Build = 1
Driver Version = 7.0
Driver Build = 1
Volume Bit Map = 1000070h
Service Start Time = Fri Oct 06 11:20:45 2008
Last Modified Time = Fri Oct 06 15:11:53 2008
```

GETSNAPSHOTLOCATION

EMCMD <system> GETSNAPSHOTLOCATION <volume letter>

This command retrieves the currently configured snapshot location (directory path) for the given volume on the given system. It will return an empty result if the snapshot location is not configured on the given volume.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	This is the drive letter of the volume to be snapshotted

Sample output:

C:\Temp

When the command is successful, it will report the snapshot directory path on stdout, which will be empty if snapshot location is not yet configured.

GETSOURCEMIRROREDVOLUMES

EMCMD <system> GETSOURCEMIRROREDVOLUMES

This command displays information about the volumes on the system that are currently the source in a mirror.

Sample output:

Status = 0

Source Volume = F:

Source Label = New Volume

Source #Targs = 2

Target 0

Target System = 10.1.1.133

Mirror State = 0001

Target 1

Target System = 10.1.1.134

Mirror State = 0001

GETTARGETMIRROREDVOLUMES

EMCMD <system> GETTARGETMIRROREDVOLUMES

This command displays information about the volumes on the system that are currently the target in a mirror.

Sample output:

```
** Calling GetTargetMirroredVolumes **
```

```
Returned 1 Target Volumes
```

```
Target Volume 1 information:
```

```
Volume Root = F:
```

```
Volume State = 1
```

```
Source = 10.1.1.132
```

```
Target = BLUEJAY
```

GETVOLUMEDRVSTATE

EMCMD <system> GETVOLUMEDRVSTATE <volume letter>

This command retrieves the current state of the SIOS DataKeeper device driver.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume you want to get the configuration on.

The output is a number indicating the state. The output is purposely terse as it is designed to be parsed in a DataKeeper recovery script. The 4

- 1:** Invalid State
- 0:** No mirror
- 1:** Mirroring
- 2:** Mirror is resyncing
- 3:** Mirror is broken
- 4:** Mirror is paused
- 5:** Resync is pending

The output also provides the address of the mirror end point (source or target).

GETVOLUMEINFO

EMCMD <system> GETVOLUMEINFO <volume letter> <level>

This command returns information about the selected volume.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume you want information on.
<level>	A number between 1-3 indicating the amount of detail you want.

Sample output:

————— LEVEL 1 INFO —————

Volume Root = F:

Last Modified = Fri Nov 05 15:24:14 2008

Mirror Role = SOURCE

Label = New Volume

FileSystem = NTFS

Total Space = 2151608320

Num Targets = 2

Attributes : 20h

————— LEVEL 2 INFO —————

>> Remote [0] = 10.1.1.133, F:

Mirror State = MIRROR

Mirror Type = ASYNCHRONOUSLY

>> Remote [1] = 10.1.1.133, F:

Mirror State = MIRROR

Mirror Type = ASYNCHRONOUSLY

----- LEVEL 3 INFO -----

>> Remote [0] = 10.1.1.133, F:

No Resync or CompVol Statistics to report

>> Remote [1] = 10.1.1.134, F:

No Resync or CompVol Statistics to report

ISBREAKUSERREQUESTED

EMCMD <system> ISBREAKUSERREQUESTED <volume letter>

This command checks whether a broken mirror is a result of a user request.

The parameters are:

<system>	This can be either the source or the target system.
<volume letter>	The drive letter of the volume that you want to check.

Output:

<TRUE>	The mirror was broken because of a user request.
<FALSE>	<p>The mirror was broken by SIOS DataKeeper (e.g., network failure, failure to write data on target side, etc).</p> <p>The volume is not in a BROKEN (3) state.</p>

ISPOTENTIALMIRRORVOL

EMCMD <system> ISPOTENTIALMIRRORVOL <volume letter>

This command checks to determine if a volume is a candidate for mirroring. The command may only be run on the local system.

The parameters are:

<system>	This should be the local system.
<volume letter>	The drive letter of the volume that you want to check.

Output:

TRUE - The volume is available for mirroring.

Otherwise, the output may be some combination of the following:

System Drive

RAW filesystem

FAT filesystem

ACTIVE partition

Contains PageFile

GetDriveType not DRIVE_FIXED

Contains DataKeeper bitmap files

If the drive letter points to a newly created volume (i.e. SIOS DataKeeper driver not attached yet), or a non-disk (network share, CD-ROM), the output will be:

Unable to open - SIOS DataKeeper driver might not be attached (you may need to reboot) or this might not be a valid hard disk volume.

If there is an internal error getting volume information, you may see the message:

Unable to retrieve the volume information for use in determining the potential use as a mirrored volume. The volume may be locked by another process or may not be formatted as NTFS.

LOCKVOLUME

EMCMD <system> LOCKVOLUME <volume letter>

This command forces an exclusive lock on the volume specified. This call will fail if a process owns open handles into the volume.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume you want to lock.

MERGETARGETBITMAP

EMCMD <system> MERGETARGETBITMAP <volume letter> <target system>

This command is for internal use only.

PAUSEMIRROR

EMCMD <system> **PAUSEMIRROR** <volume letter> [<target system>]

This command forces the mirror into a **Paused** state.

The parameters are:

<system>	This is the source system of the mirror to pause. Running the PAUSEMIRROR command on the target has no effect.
<volume letter>	The drive letter of the mirror that you want to pause.
<target system>	This is the IP address of the target system of the mirror to pause. This optional parameter may be used if multiple targets are associated with the mirror. If not specified, the mirror to all targets will be paused.

PREPARETOBECOMETARGET

EMCMD <system> PREPARETOBECOMETARGET <volume letter>

This command should only be used to recover from a [Split-Brain](#) condition. It should be run on the system where the mirror is to become a target and is only valid on a mirror source. This command causes the mirror to be deleted and the volume to be locked.

To complete split-brain recovery, run [CONTINUEMIRROR](#) on the system that remains as the mirror source.

Example Scenario

If volume F: is a mirror source on both SYSA and SYSB, you can use emcmd to resolve this split-brain situation. Choose one of the systems to remain a source - for example, SYSA. Make sure there are no files or modifications on SYSB that you want to save - if so, these need to be copied manually to SYSA. To re-establish the mirror, perform the following steps:

```
EMCMD SYSB PREPARETOBECOMETARGET F
```

The mirror of F: on SYSB will be deleted and the F: drive will be locked.

```
EMCMD SYSA CONTINUEMIRROR F
```

Mirroring of the F: drive from SYSA to SYSB will be established, a partial resync will occur (overwriting any changes that had been made on SYSB), and the mirror will reach the **Mirroring** state.

READREGISTRY

EMCMD <system> READREGISTRY <volume letter>

This command tells the SIOS DataKeeper driver to re-read its registry settings.

The parameters are:

<system>	This can be either the source system or the target system.
<volume letter>	The drive letter of the mirror for which you want to re-read settings.

This command causes the following registry settings to be re-read and any changes to take effect.

Source system (changes to these parameters take effect immediately):

BandwidthThrottle

BitmapBytesPerBlock

BlockWritesOnLimitReached

CompressionLevel

ResyncReads

WriteQueueByteLimitMB

WriteQueueHighWater

WriteQueueLowWater (This value is deprecated and no longer used.)

DontFlushAsyncQueue

Target system (changes take effect the next time the source and target systems reconnect):

TargetPortBase

TargetPortIncr

REGISTERCLUSTERVOLUME

EMCMD <system> REGISTERCLUSTERVOLUME <volume letter>

This command is used to register a DataKeeper protected volume in a WSFC cluster.

The parameters are:

<system>	This is the source system of the mirror.
<volume letter>	The drive letter of the volume you want registered.

RESTARTVOLUMEPIPE

EMCMD <system> RESTARTVOLUMEPIPE <volume letter>

This command is for internal use only.

RESYNCMIRROR

EMCMD <system> RESYNCMIRROR <volume letter> [<target system>]

This command forces the mirror to be fully resynced.

The parameters are:

<system>	This is the source system name.
<volume letter>	This is the drive letter of the mirror that should be resynced.
<target system>	This is the IP address of the target system of the mirror to resync. This optional parameter may be used if multiple targets are associated with the mirror. If not specified, a resync to all targets will be performed.

SETBLOCKTARGET

EMCMD <system> SETBLOCKTARGET <volume letter>

This command sets the state of the block target flag to TRUE. The block target flag when set to TRUE will prevent that system from ever becoming a target for the selected volume. This command is for internal use only. No output is produced when running this command.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume for which you want to set the state of the block target flag to TRUE.

SETCONFIGURATION

EMCMD <system> SETCONFIGURATION <volume letter> <configuration mask>

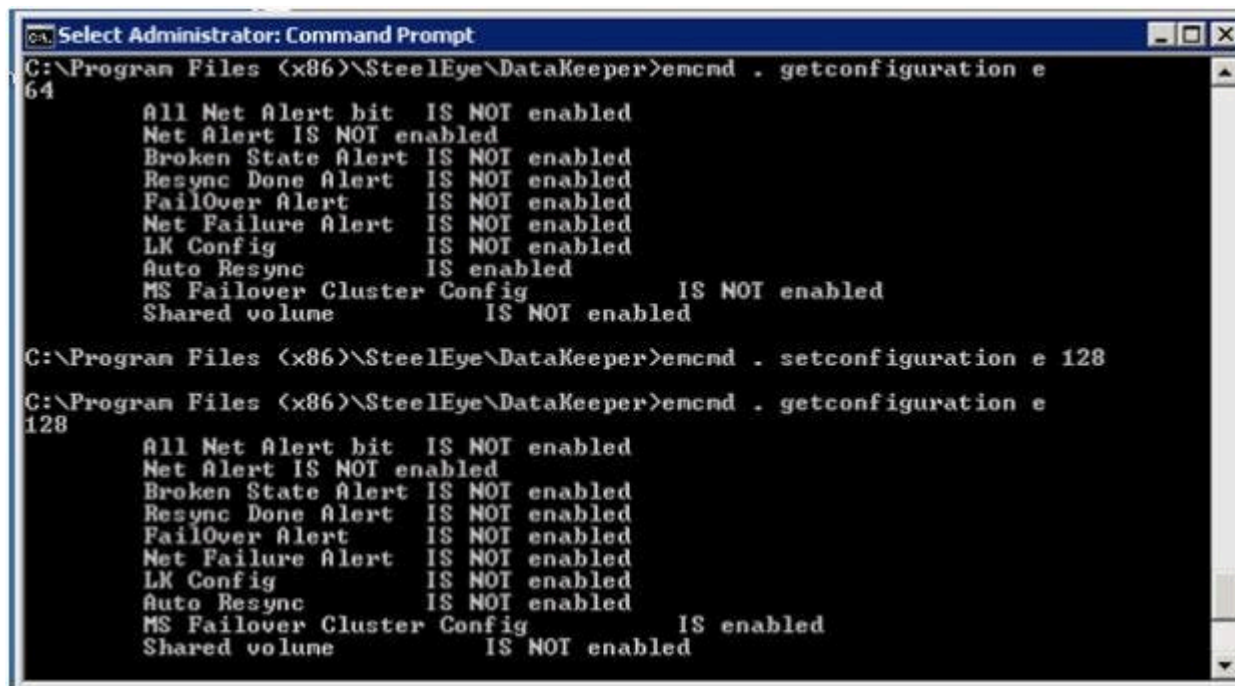
This command sets the net alert settings (also referred to as "volume attributes") for the volume.

The parameters are:

<system>	This can be either the source or the target system.
<volume letter>	The drive letter of the volume you want to set the configuration on.
<configuration mask>	<p>This is a bitmask indicating the net alert settings. These bits are defined:</p> <p>1 – 0x01: All Net Alerts is enabled</p> <p>2 – 0x02: Broken State Alert is enabled</p> <p>4 – 0x04: Resync Done Alert is enabled</p> <p>8 – 0x08: Failover Alert is enabled</p> <p>16 – 0x10: Net Failure Alert is enabled</p> <p>32 – 0x20: LifeKeeper Config is enabled</p> <p>64 – 0x40: Auto Resync is enabled</p> <p>128 – 0x80: MS Failover Cluster Config is enabled</p> <p>256 – 0x100: Shared Volume is enabled</p>

Example to enable MS Failover Cluster Config:

EMCMD . SETCONFIGURATION E 128



```

C:\Program Files (x86)\SteelEye\DataKeeper>emcmd . getconfiguration e
64
All Net Alert bit IS NOT enabled
Net Alert IS NOT enabled
Broken State Alert IS NOT enabled
Resync Done Alert IS NOT enabled
FailOver Alert IS NOT enabled
Net Failure Alert IS NOT enabled
LK Config IS NOT enabled
Auto Resync IS enabled
MS Failover Cluster Config IS NOT enabled
Shared volume IS NOT enabled

C:\Program Files (x86)\SteelEye\DataKeeper>emcmd . setconfiguration e 128

C:\Program Files (x86)\SteelEye\DataKeeper>emcmd . getconfiguration e
128
All Net Alert bit IS NOT enabled
Net Alert IS NOT enabled
Broken State Alert IS NOT enabled
Resync Done Alert IS NOT enabled
FailOver Alert IS NOT enabled
Net Failure Alert IS NOT enabled
LK Config IS NOT enabled
Auto Resync IS NOT enabled
MS Failover Cluster Config IS enabled
Shared volume IS NOT enabled

```

Example to clear all flags:

EMCMD . SETCONFIGURATION E 0

Multiple Configuration Example to enable Shared Volume and MS Failover Cluster Config (add decimal values 256 + 128):

EMCMD . SETCONFIGURATION E 384

SETSNAPSHOTLOCATION

EMCMD <system> SETSNAPSHOTLOCATION <volume letter> "<directory path>"

This command sets the snapshot location (directory path) for the given volume on the given system. The directory must be valid on the system in question, must be a local drive/path, must be an absolute path and cannot be left blank (see [CLEARSNAPSHOTLOCATION](#)). If no snapshot location is currently configured, executing this command will have the effect of enabling target snapshots on the given volume.

The parameters are:

<system>	This is the system name/IP address containing volume to be snapshotted.
<volume letter>	This is the drive letter of the volume to be snapshotted.
<directory path>	This is the absolute directory path, local to <system>, for the snapshot file location. Note that this value must be enclosed in quotes if the path contains a space character.

Sample output:

Status = 0

When the command is successful, it will return a status of 0. Otherwise, it will report a non-zero status.

STOPSERVICE

EMCMD <system> STOPSERVICE

This command stops the DataKeeper service.

SWITCHOVERVOLUME

EMCMD <system> SWITCHOVERVOLUME <volume letter> [-f]

This command attempts to make the given system become the source for the requested volume. **This command is for internal use only.**

The parameters are:

<system>	This is the IP address of the system to become source. Note: Use the system's NetBIOS name, IP address or fully qualified domain name to attach to a given system. You can also use a period (.) to attach to the local system where emcmd is being executed.
<volume letter>	This is the drive letter of the requested volume.
[-f]	This option may be used for a fast (<i>unsafe</i>) switchover. This option should only be used if the status of the current source is known. Incorrect usage of this can result in a split-brain condition.

TAKESNAPSHOT

EMCMD <target system> TAKESNAPSHOT <volume letter> [<volume letter>...]

This command, run on the target system, will notify DataKeeper to establish a snapshot of the given volume(s) on the given system. If no snapshot location has been configured, the command will fail.

The parameters are:

<target system>	This is the target system name/IP address containing the volume to be snapshot.
<volume letter>	This is the drive letter(s) of the volume(s) to be snapshot on the target server. If multiple volumes are to be snapshot, the drive letters should be separated by spaces.

Note: All target volumes must have the same source system.

UNLOCKVOLUME

EMCMD <system> UNLOCKVOLUME <volume letter>

This command forces the volume specified to unlock.

The parameters are:

<system>	This can be either the source or the target systems.
<volume letter>	The drive letter of the volume you want to unlock.

UPDATEJOB

```
EMCMD <system>UPDATEJOB <JobId> <Name> <Descr> [<SysName1>  
<DrvLetter1> <IP1> <SysName2> <DrvLetter2> <IP2> <MirrorType>]...
```

This command is for internal use only.

UPDATEVOLUMEINFO

EMCMD <system> UPDATEVOLUMEINFO <volume letter>

This command causes the SIOS DataKeeper service to query the driver for the correct mirror state. This command is useful if the DataKeeper GUI displays information that appears to be incorrect or not up-to-date.

Note: The SIOS DataKeeper service updates the volume information automatically based on new messages in the system [Event Log](#).

The parameters are:

<system>	This can be either the source or the target system.
<volume letter>	The drive letter of the volume that you want to update its info.

If there is an internal error updating volume information, you may see the message:

Unable to update the volume information. The volume may be locked by another process or may not be formatted as NTFS.

Using the `-proxy` option with EMCMD

All EMCMD requests can be routed through a “proxy” DataKeeper service. To do this, append the options

```
-proxy <proxy_system>-
```

to the end of the EMCMD command line. The `<proxy_system>` should be given using the same format as the `<system>` option. EMCMD will open a connection to the `<proxy_system>` first, and will request that it forward the EMCMD to `<system>`. The DataKeeper Service on `<proxy_system>` opens a connection to `<system>`, and sends the requested EMCMD to `<system>`, returning the response to the user.

The `-proxy <proxy_system>` option allows you to verify that DataKeeper communication between nodes is working.

Example

```
EMCMD DK_NODE_2 GETSERVICEINFO -proxy DK_NODE_1
```

Opens a connection to the DataKeeper service running on `DK_NODE_1`, which in turn opens a connection to `DK_NODE_2`, forwards the `GETSERVICEINFO` request, and returns the service information from `DK_NODE_2`. This can be used to validate that the DataKeeper service on `DK_NODE_1` is able to communicate with the DataKeeper service on `DK_NODE_2`.

Using DKPwrShell with SIOS DataKeeper

SIOS DataKeeper includes a powershell module (DKPwrShell) that allows a user to manipulate a DataKeeper mirror using Microsoft Powershell. Commands are passed to a SIOS DataKeeper service and will fail if the service is not running.

With Microsoft Powershell v3 or later the SIOS DataKeeper powershell module is loaded automatically when starting Powershell. For Microsoft Powershell versions prior to 3.0 the SIOS DataKeeper powershell module must be loaded via the `import-module` command by using the following syntax:

```
import-module "<DK InstallPath>\DKPwrShell"
```

Note: By default <DK InstallPath> is C:\Program Files (x86)\SIOS\DataKeeper

[New-DataKeeperMirror](#)

[New-DataKeeperJob](#)

[Remove-DataKeeperMirror](#)

[Remove-DataKeeperJob](#)

[Add-DataKeeperJobPair](#)

[Get-DataKeeperVolumeInfo](#)

New-DataKeeperMirror

This cmdlet is used to create a new DataKeeper mirror. Mirrors created with this cmdlet will be visible in the DataKeeper SnapIn (Reports > Server Overview). If a job exists that includes information that matches this mirror (systems, IP Addresses, Volumes, and Sync Type), the mirror will be displayed in the DataKeeper SnapIn as part of that job.

Parameters

Parameter	Type	Required	Position	Notes
SourceIP	String	Yes	0	IP address on the source to be used for DataKeeper mirror data.
SourceVolume	String	Yes	1	The source volume to mirror.
TargetIP	String	Yes	2	IP address on the target to be used for DataKeeper mirror data.
TargetVolume	String	Yes	3	The target volume to become the mirror target. If not specified it will be the same volume indicated by the SourceVolume parameter.
SyncType	String	Yes	4	Valid options are: Sync - A synchronous mirror Async - An asynchronous mirror
CreateFlags	uint	No	5	Optional arguments that specify behavior deviate from the norm. These can be OR'd together to create a set of options (add decimal values. Example: for option 1 + option 2, place a 3 in the command). 1. Create the mirror without doing a full resync operation. 2. Do not wait for the target side of the mirror to be created before returning.

Inputs

None

Outputs

An integer value representing the status of the command. 0 means that the

command succeeded, any other value is a Windows error code.

* NOTE: Both source and target IP addresses must be of the same protocol. A mirror can only be created using two IPV4 or two IPV6 addresses. DataKeeper does not currently support mirror endpoints with different protocols.

Example:

```
New-DataKeeperMirror -SourceIP 10.200.8.55 -SourceVolume E -TargetIP 10.200.8.56  
-TargetVolume E -SyncType Async
```

```
New-DataKeeperMirror 10.200.8.55 E 10.200.8.56 E Async
```

* NOTE: Disk sector size must match on both source and target volumes. See [Sector Size](#) for more information.

New-DataKeeperJob

This cmdlet is used to create a DataKeeper job consisting of two nodes. Jobs created using this cmdlet will be added to the DataKeeper SnapIn the next time it is loaded.

Parameters

Parameter	Type	Required	Position	Notes
JobName	String	Yes	0	The name of the job.
JobDescription	String	Yes	1	A brief description of the job.
Node1Name	String	Yes	2	The FQDN of the first node.
Node1IP	String	Yes	3	The IP address of the first node that is used for DataKeeper Replication.
Node1Volume	String	Yes	4	The volume of the first node that is involved in replication.
Node2Name	String	Yes	5	The FQDN of the second node.
Node2IP	String	Yes	6	The IP address of the second node that is used for DataKeeper Replication.
Node2Volume	String	Yes	7	The volume of the second node that is involved in replication.
SyncType	String	Yes	8	Valid options are: Sync - A synchronous mirror Async - An asynchronous mirror Disk - These two volumes are a single shared disk

Inputs

None

Outputs

On success, an object representing the created job. On failure, an exception containing a Windows error code.



NOTE: Both IP addresses must be of the same protocol (IPv4 or IPv6). DataKeeper does not currently support mirror endpoints with different protocols.

Example:

```
New-DataKeeperJob -JobName "name" -JobDescription "desc" -Node1Name  
example1.domain.com -Node1IP 10.200.8.55 -Node1Volume E -Node2Name  
example2.domain.com -Node2IP 10.200.8.56 -Node2Volume F -SyncType Async
```

```
New-DataKeeperJob "name" "desc" example1.domain.com 10.200.8.55 E  
example2.domain.com 10.200.8.56 F Async
```

Remove-DataKeeperMirror

This cmdlet will remove a DataKeeper mirror. It will attempt to remove the mirror from all nodes for this mirror. This command will not remove the mirror from any down or network inaccessible node.

Parameters

Parameter	Type	Required	Position	Notes
Source	String	Yes	0	The source node of the mirror.
Volume	String	Yes	1	The mirror volume letter (on the source node) that you want removed.
Target	String	No	2	The IP address of the target system of the mirror. If this parameter is left empty all targets of the source volume will be removed.

Inputs

None

Outputs

An integer value representing the status of the command. 0 means that the command succeeded, any other value is a Windows error code.

Example:

```
Remove-DataKeeperMirror -Source example1.domain.com -Volume E -Target  
10.200.8.56
```

```
Remove-DataKeeperMirror -Source 10.200.8.55 -Volume E -Target 10.200.8.56
```

```
Remove-DataKeeperMirror 10.200.8.55 E
```

Remove-DataKeeperJob

This cmdlet will remove a DataKeeper job of a given ID. It will remove this job from all systems contained within the job.

Parameters

Parameter	Type	Required	Position	Notes
JobID	String	Yes	0	The unique job GUID assigned to it when the job was created.
Node	String	Yes	1	The FQDN or IP address of a node containing the job specified by JobID.

Inputs

None

Outputs

An integer value representing the status of the command. 0 means that the command succeeded, any other value is a Windows error code.

Example:

```
Remove-DataKeeperJobPair -JobID a1flecc6-649e-476b-bbff-286b815fdd30 -Node  
example1.domain.com
```

```
Remove-DataKeeperJobPair a1flecc6-649e-476b-bbff-286b815fdd30 10.200.8.55
```

Add-DataKeeperJobPair

This cmdlet will add a node pair to an existing DataKeeper Job. It is used to expand the nodes and volumes contained within an existing job. For example, if a job exists for a volume between nodes A and B, and you want to add node C, run AddDataKeeperJobPair twice:

- for the new relationship definition between node A and node C
- for the new relationship definition between node B and node C

Parameters

Parameter	Type	Required	Position	Notes
JobID	String	Yes	0	The unique job GUID assigned to it when the job was created.
Node1Name	String	Yes	1	The FQDN of the first node
Node1IP	String	Yes	2	The IP address of the first node that is used for DataKeeper Replication.
Node1Volume	String	Yes	3	The volume of the first node that is involved in replication.
Node2Name	String	Yes	4	The FQDN of the second node.
Node2IP	String	Yes	5	The IP address of the second node that is used for DataKeeper Replication.
Node2Volume	String	Yes	6	The volume of the second node that is involved in replication.
SyncType	String	Yes	7	Valid options are: Sync - A synchronous mirror Async - An asynchronous mirror Disk - These two volumes are a single shared disk

Inputs

None

Outputs

An integer value representing the status of the command. 0 means that the

command succeeded, any other value is a Windows error code.

Example:

```
Add-DataKeeperJobPair -JobID a1f1ecc6-649e-476b-bbff-286b815fdd30 -Node1Name  
example1.domain.com -Node1IP 10.200.8.55 -Node1Volume E -Node2Name  
example2.domain.com -Node2IP 10.200.8.56 -Node2Volume F -SyncType Async
```

```
Add-DataKeeperJobPair a1f1ecc6-649e-476b-bbff-286b815fdd30 example1.domain.com  
10.200.8.55 E example2.domain.com 10.200.8.56 F Async
```


Get-DataKeeperVolumeInfo

This cmdlet is used to fetch information about a volume used in DataKeeper. It reports DataKeeper volume information.

Parameters

Parameter	Type	Required	Position	Notes
Node	String	Yes	0	Use the Node parameter to specify the system containing the volume being replicated. This parameter can take the form of an IPv4 address, FQDN, or simply `.` for the local system.
Volume	String	Yes	1	The mirror volume letter (on the system node).

Inputs

None

Outputs

VolumeInfo object

Example:

```
Get-DataKeeperVolumeInfo -Node example.domain.com -Volume E
```

```
Get-DataKeeperVolumeInfo 10.200.8.55 E
```

```
Get-DataKeeperVolumeInfo . E
```

DataKeeper User Guide

The topics in this section are designed to be a reference for you as you get started using SIOS DataKeeper, helping you identify the type of configuration you are interested in implementing and providing detailed instructions for effectively using your SIOS DataKeeper software.

[Getting Started](#)

[DataKeeper Setup](#)

[Configuring Mirrors](#)

[Working With Jobs](#)

[Working With Mirrors](#)

[Working With Shared Volumes](#)

[Using Microsoft iSCSI Target With DataKeeper on Windows 2012](#)

[DataKeeper Notification Icon](#)

[DataKeeper Target Snapshot](#)

[Using SIOS DataKeeper Standard Edition To Provide Disaster Recovery For Hyper-V Virtual Machines](#)

Getting Started

Choose Your Configuration

DataKeeper can be utilized in a number of different configurations to facilitate a number of different functions including:

- Provide a second physical copy of your data
- Extend an existing WSFC cluster to a remote DR site
- Eliminate the Single Point of Failure associated with traditional WSFC clusters

Review the following replication configurations and their example USE CASES to familiarize yourself with just some of DataKeeper's capabilities. Then use the topics associated with the configuration you are interested in to obtain detailed information about that configuration.

[Disk-to-Disk](#)

[One-to-One](#)

[One-to-Many](#)

[Many-to-One](#)

[N-Shared-Disk Replicated to One](#)


[N-Shared-Disk Replicated to N-Shared-Disk](#)

[N-Shared-Disk Replicated to Multiple N-Shared-Disk Targets](#)

Disk-to-Disk

This is a simple one server, two disks configuration, mirroring Volume X on Server A to Volume Y on Server A. The volumes that are used for Disk-to-Disk replication can't also be configured to replicate to another system.

* Disk to Disk does not support mirrors with multiple targets.



Example:	Replicate data from one volume on a server to another volume on the same
-----------------	--


USE CASE	server. These disks can be different storage arrays, protecting against data loss should the primary SAN fail.
-----------------	--

Additional topics of interest include:

- [Creating Mirrors](#)
- [Managing Mirrors](#)
- [Extensive Write Considerations](#)
- [Frequently Asked Questions](#)

One-to-One

This is a simple one source, one target configuration, mirroring Volume X across the network. In addition to providing a second physical copy of the data, DataKeeper also provides the ability to switch over the mirror which allows the data to become active on the backup server.




Example:	Replicate data on one or more volumes from a server in one city to another server in another city.
USE CASE	

Additional topics of interest include:

- [Primary Server Shutdown](#)
- [Secondary Server Failures](#)
- [Using DataKeeper Standard Edition To Provide Disaster Recovery For Hyper-V Virtual Machines](#)
- [Frequently Asked Questions](#)

One-to-Many (Multiple Targets)

This configuration involves one primary (source) system replicating one (or more) volume(s) to two different target systems across the network. This is referred to as a multiple target configuration.



Note that there are two mirrors that are completely independent of each other. The mirrors might be using different networks, they may have different compression or bandwidth throttle settings and they may be in completely different states (e.g. Mirror 1 – Mirroring, Mirror 2 – Resyncing).


Example: USE CASE	Replicate data to one target server that resides locally in the same site with the primary server and replicate another copy of the data to a remote site for disaster recovery purposes should something happen to the first site.
Example: USE CASE	To periodically replicate or “push” data to multiple target systems from a single source system.

Additional topics of interest include:

- [Primary Server Shutdown](#)
- [Secondary Server Failures](#)
- [Creating Mirrors with Multiple Targets](#)
- [Switchover and Failover with Multiple Targets](#)
- [Using DataKeeper Standard Edition To Provide Disaster Recovery For Hyper-V Virtual Machines](#)
- [Frequently Asked Questions](#)

Many-to-One

This configuration involves multiple source servers replicating one (or more) volumes to the same target system. In this configuration, each volume being replicated to the target server must have a unique drive letter.



Note: This is actually two One-to-One mirrors.

Example:	Users may wish to replicate multiple branches back to a single data center for backup consolidation and disaster recovery purposes.
USE CASE	

Additional topics of interest include:


- [Primary Server Shutdown](#)
- [Secondary Server Failures](#)
- [Using DataKeeper Standard To Provide Disaster Recovery For Hyper-V Virtual](#)

[Machines](#)

- [Frequently Asked Questions](#)

N-Shared-Disk Replicated to One

This configuration allows you to replicate the shared volume(s) of the primary site to a remote system across the network.



This configuration is ideal for providing local failover within the Primary Site and disaster recovery protection should the entire Primary Site go down.

Example: USE CASE	Extend your WSFC cluster to a DR site by replicating the shared volume to a remote target. In the event of a primary site outage, the remote server becomes the active server.
------------------------------	--

Additional topics of interest include:


DataKeeper Standalone

- [Creating Mirrors with Shared Volumes](#)
- [Managing Shared Volumes](#)
- [Adding a Shared System](#)

- [Removing a Shared System](#)
- [Frequently Asked Questions](#)

N-Shared-Disk Replicated to N-Shared-Disk

This configuration replicates data between sites where each site utilizes shared storage.



Note that the number of systems in the Primary Site does not have to equal the number of systems in the Remote Site.

Also note that only the Source Server has access to the Source Volume. Shared Source systems and all systems on the target side cannot access the volume and are locked from the file system’s perspective.

Example: USE CASE	Users who wish to provide the same level of availability in their DR site will deploy this configuration to ensure that regardless of what site is in service, the availability level stays the same.
Example: USE CASE	Where Hyper-V clusters are configured with virtual machines distributed across many cluster nodes, it is important to have a similar number of cluster nodes available in the disaster recovery sight to ensure that the resources are available to run all of the virtual machines in the event of a disaster.


Additional topics of interest include:

DataKeeper Standalone

- [Creating Mirrors with Shared Volumes](#)
- [Managing Shared Volumes](#)
- [Adding a Shared System](#)
- [Removing a Shared System](#)
- [Frequently Asked Questions](#)

N-Shared-Disk Replicated to Multiple N-Shared-Disk Targets

This is a complex configuration which combines the aspects of replicating a shared storage environment to multiple shared targets.



Note that the number of systems in the Primary Site does not have to equal the number of systems in the Remote Site.

Also note that only the Source Server has access to the Source Volume. Shared Source systems and all systems on the target side cannot access the volume and are locked from the file system's perspective.

Example: USE CASE	Users who wish to provide the same level of availability in their DR site will deploy this configuration to ensure that regardless of what site is in service, the availability level stays the same.
Example: USE CASE	Where Hyper-V clusters are configured with virtual machines distributed across many cluster nodes, it is important to have a similar number of

	cluster nodes available in the disaster recovery sight to ensure that the resources are available to run all of the virtual machines in the event of a disaster.
--	--

Additional topics of interest include:

DataKeeper Standalone

- [Creating Mirrors with Shared Volumes](#)
- [Managing Shared Volumes](#)
- [Adding a Shared System](#)
- [Removing a Shared System](#)
- [Frequently Asked Questions](#)

Setting Up SIOS DataKeeper

Follow these steps to start using SIOS DataKeeper:

1. [Connect to the servers](#) you wish to configure for replication. You can select **Connect to Server** from the **Action** pull down menu, right-click on the job folder in the left panel tree display and select **Connect to Server** or choose **Connect to Server** from the **Actions** pane.
2. [Create a Job](#). From the right **Actions** pane, select **Create Job** or you can right-click on the job folder in the left panel tree and select **Create Job**.
3. [Create a mirror](#) for the new job.

Setting Up SIOS DataKeeper


Setting Up SIOS DataKeeper

Follow these steps to start using SIOS DataKeeper:

1. [Connect to the servers](#) you wish to configure for replication. You can select **Connect to Server** from the **Action** pull down menu, right-click on the job folder in the left panel tree display and select **Connect to Server** or choose **Connect to Server** from the **Actions** pane.
2. [Create a Job](#). From the right **Actions** pane, select **Create Job** or you can right-click on the job folder in the left panel tree and select **Create Job**.
3. [Create a mirror](#) for the new job.

Connecting to a Server


Use this dialog to connect to the server of your choice. You may enter the IP address, system NetBIOS name or the full system domain name for the server. Click **Connect** to select it.



Disconnecting from a Server

Use this dialog to disconnect from a server. You may use this option if you no longer wish to view the server in the Administration Window.

From the list of servers, select the server(s) that you wish to disconnect from and click **Disconnect**.



Creating a Job

1. If not already connected, [connect to the server](#) where you want to create a [job](#).
2. From the right **Actions** pane, select **Create Job**. The **Job Wizard** will prompt you for a **Job Name** and **Description**.
3. Enter the appropriate information and select **Create Job** to finish.
4. You will immediately be prompted to [Create a Mirror](#) for this job.

Configuring Mirrors

[Creating a Mirror](#)

[Creating Mirrors With Shared Volumes](#)

[Safe Creation of a Shared-Storage Volume Resource](#)

[Creating Mirrors With Multiple Targets](#)

[Switchover and Failover with Multiple Targets](#)

Creating a Mirror

Before creating a mirror, ensure the following:

- You have [created a job](#) to hold the mirror.
- The volume on both the source and target systems must be of the **NTFS** file system type.
- The target volume must be greater than or equal to the size of the source volume.
- If the volume will be configured on a **Dynamic Disk**, create the dynamic volume first, then reboot the system before continuing with mirror creation (see the [Mirroring with Dynamic Disks](#) Known Issue for further information).
- See [Volume Considerations](#) for more information, including what volumes cannot be mirrored.
- You must be connected to both the source and target server before creating the mirror. Use the [Connect to Server](#) link in the **Actions** pane or in the **Mirror Create** dialog box.

Creating the Mirror

1. Select **Create a Mirror** from the right column **Actions** task pane. The **Choose a Source** dialog box appears.
2. Enter or choose the **Server Name** for the source volume. You can select the **Connect to Server** link below this field to connect to the server at this time.
3. Choose the **IP address** that is on the subnet you wish to use for the replication traffic. The IP address that you choose must not be used for replication by any other node that is part of this job (see [Duplicate IP Addresses Disallowed Within a Job](#) for further information).
4. Enter or choose the **Volume** to be used on the selected server. Select **Next**.

The **Choose a Target** dialog box appears.

5. Enter or choose the server with the **Target Volume**. If necessary, you can select the **Connect to Server** link at this time.
6. Choose the **IP address** that is on the subnet you wish to use for the replication traffic. The IP address that you choose must not be used for replication by any other node that is part of this job (see [Duplicate IP Addresses Disallowed Within a Job](#) for further information).
7. Enter or choose the **Volume** to be used on the selected server. Press **Next** to continue. The **Configure Details** dialog box will display.
8. Use the slide bar to set the **data compression level** for data sent from the source to the target system. **Note:** Compression is only recommended to be used when replicating across WAN connections.
9. Select how ([Asynchronously or Synchronously](#)) the source volume data should be sent to the target volume.
10. If you wish to limit the amount of bandwidth used by replication, enter the **maximum bandwidth** for transmission; otherwise, leave the default setting. Select **Done**. The job with the new mirror will appear in the left tree pane and the main window displays.

Note: After creating a mirror, its initial state may be displayed as **Resync Pending** in the **Summary** pane. When the initial mirror resynchronization completes, its state will automatically switch to the **Mirror** state.

Creating Mirrors With Shared Volumes

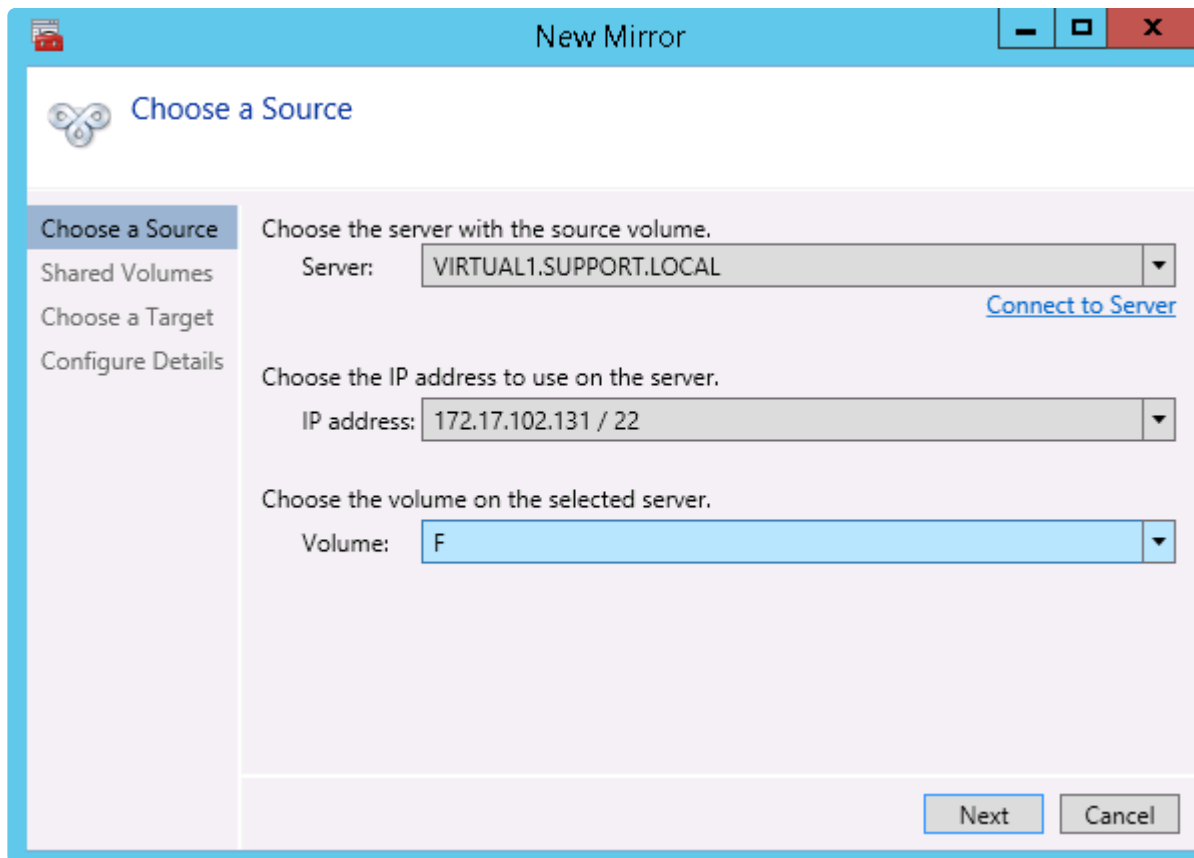
In order to properly configure DataKeeper in a shared volume configuration, use the **DataKeeper GUI** to connect to all systems where the shared volumes are configured. When connected, the DataKeeper GUI uses hardware signatures to automatically detect which volumes are shared and which are not.

Important: If the GUI is not connected to a system, the GUI cannot detect shared volumes on that system.

Note: Dynamic disks are not supported with Shared Storage because the dynamic disk configuration is stored somewhere (undocumented) on each system, not on the disks themselves. There is currently no way to replicate that configuration between the two systems.

Note: DataKeeper allows mirrors to be created on shared volumes where more than one system has access to the same physical storage. To prevent simultaneous access, see [Safe Creation of a Shared-Storage Volume Resource](#) prior to performing the following steps.

1. Connect to all systems via the **DataKeeper GUI**.
2. Choose [Create Job](#).
3. Define a job name and job description and select **Create Job**. The **Choose a Source** dialog box appears.



The screenshot shows a Windows-style dialog box titled "New Mirror". It has a blue header bar with standard window controls (minimize, maximize, close). Below the header, there's a sub-header "Choose a Source" with a small icon. On the left, a vertical sidebar contains four options: "Choose a Source" (highlighted), "Shared Volumes", "Choose a Target", and "Configure Details". The main area of the dialog is divided into three sections. The first section is titled "Choose the server with the source volume." and contains a "Server:" label followed by a dropdown menu showing "VIRTUAL1.SUPPORT.LOCAL" and a "Connect to Server" link. The second section is titled "Choose the IP address to use on the server." and contains an "IP address:" label followed by a dropdown menu showing "172.17.102.131 / 22". The third section is titled "Choose the volume on the selected server." and contains a "Volume:" label followed by a dropdown menu showing "F". At the bottom right, there are two buttons: "Next" and "Cancel".

New Mirror

Choose a Source

Choose a Source

Choose the server with the source volume.

Server: [Connect to Server](#)

Choose the IP address to use on the server.


IP address:

Choose the volume on the selected server.

Volume:

Next **Cancel**

4. Choose a **Source System**, **IP Address** and **Volume**.
5. Select **Next**. The **Shared Volumes** dialog box appears.



- Choose the systems that have volumes which are shared with the source system.

Note: All systems connected to the shared volumes must be configured with IP addresses on the same subnet. The **Next** button will not be enabled until all included systems have a valid IP address.

While it is possible to uncheck the **Include** box for a given system, the user should be very careful to make sure that the volume listed really is not a shared volume. It is possible (although unlikely) that the hardware signatures of two volumes will match even if they are not shared. In this case, it is valid for the user to uncheck the **Include** box.

- Select **Next**. The **Choose a Target** dialog box appears.
- Choose a **Target System**, **IP Address** and **Volume**.
- Select **Next**.

Note: If there are volumes on other systems that are shared with this target volume, the **Shared Volumes** dialog will appear next. Configure these shared target volumes as you would for shared source volumes, described above.

10. Select **Next** to continue. The **Configure Details** dialog box appears.

11. Use the slide bar to set the **data compression level** for data sent from the source to the target system.

Note: Compression is only recommended to be used when replicating across WAN connections.

12. Select how ([Asynchronously or Synchronously](#)) the source volume data should be sent to the target volume.

13. If you wish to limit the amount of bandwidth used by replication, enter the [maximum bandwidth](#) for transmission; otherwise leave the default setting.

14. Select **Done**. The job with the new mirror will appear in the left tree pane and the main window displays.

Safe Creation of a Shared-Storage Volume Resource

DataKeeper allows mirrors to be created on shared volumes where more than one system has access to the same physical storage. The shared volume can be on the source side of the mirror or on the target side.

Note: Dynamic disks are not supported with Shared Storage because the dynamic disk configuration is stored somewhere (undocumented) on each system, not on the disks themselves. There is currently no way to replicate that configuration between the two systems.

In order to safely create a shared-storage volume resource, the user must ensure that only one system has write access to the volume at any time. This includes the time prior to the creation of the DataKeeper mirror. Since DataKeeper doesn't know that the volume is shared before a mirror is created, manual steps must be taken to ensure that the volume is never writable on two or more systems at the same time.

To protect the volume from simultaneous write access, use the following procedure. In this example, two systems - *SYSA* and *SYSB* - are connected to shared storage, then replicated to a third system, *SYSC*, the target system. This storage is configured with two volumes which should be assigned drive letters *E:* and *F:* on all three systems.

1. Power on *SYSA*, while leaving *SYSB* powered off.
2. Install DataKeeper if it has not been installed.
3. Assign drive letters *E:* and *F:* to the volumes; format with NTFS if not formatted yet.
4. Power off *SYSA*.
5. Power on *SYSB*.
6. Install DataKeeper if it has not been installed and reboot the system after

the installation.

7. Assign drive letters *E:* and *F:* to the shared volumes.
8. In a command prompt, run the following commands to set the "shared" config flag:


```
"%ExtMirrBase%\emcmd" . setconfiguration E 256
```



```
"%ExtMirrBase%\emcmd" . setconfiguration F 256
```
9. Reboot *SYSB*. It will come up with the *E:* and *F:* drives locked.
10. Power on *SYSA*. It will come up with the *E:* and *F:* drives writable.
11. Use the DataKeeper GUI to [create a job and mirror](#) from *SYSA E:* (source) to *SYSC E:* (target) and from *SYSA F:* (source) to *SYSC F:* (target). DataKeeper will detect that *SYSB* is a shared source system.

Note: If using WSFC, see [Creating a DataKeeper Volume Resource in WSFC](#).

An alternative to powering the systems off is to use **Disk Management** to take the shared physical disk offline.

This procedure can also be used to safely create a mirror on a shared target volume. In the example above, the mirror could have been created from *SYSC* to *SYSA* - in that case, the volume **SYSB** would be a shared target.

If you have more than two shared systems at a site, this same procedure can be used to lock the volume on all systems that will not be part of the initial mirror.

Creating Mirrors With Multiple Targets

SIOS DataKeeper provides the ability to replicate data from a single source volume to one or more target volumes. In addition, DataKeeper also allows you to switch over control and make any of the target volumes become the source. Assuming you have already created a job with a mirror using the [Create a Mirror](#) procedure, use the following procedure to create a second mirror from the same source volume to a different target volume:

1. Right-click on an existing job.
2. Choose the **Create a Mirror** action.
3. Choose the **source** of the existing mirror (as this will also be the source of the new mirror).
4. Choose the **target** for the new mirror.
5. Select **Done**.

The next dialog displayed prompts you for additional information that DataKeeper requires to be able to properly switch over the source volume to one of the target volumes. You already specified the network endpoints between the source system and the first target system when you created the first mirror. You also specified the network endpoints between the source system and the second target system when you created the second mirror.

The final piece of information DataKeeper requires is the network endpoints of a (potential) mirror between the first target system and the second target system so that no matter which system becomes the source of the mirrors, mirrors can be properly established between all three systems.

6. On the **Additional Information Needed** dialog, choose the **network endpoints** that will be used to create a mirror between the first target system and the second target system.

This mirror will not be created now. DataKeeper is simply storing these mirror endpoints for future use.

7. Select **OK**.

Note: If you are replicating a single source volume to more than two target volumes, you will have to provide network endpoints for mirrors between all of the systems involved.

Examples:

3 Nodes (A,B,C) - Define Endpoints for Mirrors	
Created Mirrors	Additional Mirror Relationships
A → B	B → C
A → C	

4 Nodes (A,B,C,D) - Define Endpoints for Mirrors	
Created Mirrors	Additional Mirror Relationships
A → B	B → C B → D C → D
A → C	
A → D	

Switchover and Failover with Multiple Targets

In a multiple target configuration, it is important to understand how DataKeeper mirrors will work in the following scenarios:

- Manual switchover to a target server
- Source server failure followed by a manual switchover to a target server


Example:

In the following scenario, there are three servers:

- Server A (source)
- Server B (target 1)
- Server C (target 2)

Note that there are two separate mirrors and Server A is replicating to two different target volumes.

- Mirror 1: Server A → B
- Mirror 2: Server A → C




Manual Switchover to a Target Server

In the event the administrator wants to make Server B become the active (source) server, the following actions will occur:

1. Administrator initiates a switchover to Server B via the **Switchover Mirror** option in the DataKeeper UI.
2. Server A flushes its data to the source volume.
3. Mirror 1 is automatically deleted and recreated from Server B to Server A.
4. The mirror between Server A and Server C is also automatically deleted. (**Note:** There will be a few seconds delay noticed in the DataKeeper GUI; this delay can take some time based on [network bandwidth](#) and server performance.)
5. A new mirror is established between Server B and Server C. The [intent log](#) from Server A is copied to Server B. Only a partial resync of the data between Server B and Server C is required to bring them in sync. (A partial resync is the resynchronization of only the necessary data to establish the new end points and is usually much quicker than a full resync.)

RESULT

- Mirror 1: Server B → A (partial resync)
- Mirror 2: Server B → C (copy intent log from Server A, partial resync)



Source Server Failure – Manual Switchover to a Target Server


In the event the active (source) server fails, DataKeeper allows you to make Server B become the active (source) server. The following actions will occur:

1. Server A fails.
2. Administrator initiates a switchover to Server B via the "Switchover Mirror" option in the DataKeeper UI.
3. Server B deletes the local side of the mirror and creates a new mirror from Server B to Server A.
4. The mirror between Server A and Server C is deleted.
5. A new mirror is established between Server B and Server C.
6. When Server A comes back up, Server A detects that Server B became the source of the mirror while Server A was down and Server A

automatically becomes the target of the mirror.

RESULT

- Mirror 1: Server B → A (partial resync when Server A comes back up)
- Mirror 2: Server B → C (partial resync)



Working With Jobs

[Jobs](#)

[Renaming a Job](#)

[Deleting a Job](#)

[Reassigning a Job](#)


[Switching Over a Mirror](#)


Jobs


For ease of use and configuration, SIOS DataKeeper does much of its management of mirrors through an entity called a job. A job is a logical grouping of related mirrors and servers. This feature allows you to create a job for complex repetitive tasks and run them quickly from the SIOS DataKeeper user interface.

Mirrors that are related should be placed in a single job. For instance, multiple mirrors protecting an application like SQL Server should be placed in the same job. Mirrors that are unrelated should be placed in separate jobs.

Note: Mirrors created in previous versions of SIOS Data Replication will be imported as individual jobs. The administrator must take care to edit these jobs to ensure that mirrors are logically grouped together.

 Summary of Test 1 - Creating Mirrors
Test 1 has 1 mirrors

Job name: Test 1
Job description: Creating Mirrors
Servers: HERON, EGRET
Job state:  Mirroring

Source System	Target System	Target Volume	Source IP	Target IP	State	Resync Remaining
Source volume Y						
EGRET	HERON	Y	172.17.108.164	172.17.108.163	 Mirroring	0.00 KB

Mirror type: Asynchronous
File system: NTFS
Disk space: 146.68 GB
Compression: None
Maximum bandwidth: 0 kbps

Edit

Renaming a Job

1. Select the job in the left **Console Tree** pane of the main DataKeeper window.
2. You can select **Rename Job** from the **Actions** pane or right-click on the selected job and choose **Rename Job** from the menu that displays.
3. Enter the new **Job Name** and new **Job Description**.

Deleting a Job

1. Select the job in the left **Console Tree** pane of the main DataKeeper window.
2. You can select **Delete Job** from the **Actions** pane or right-click on the selected job and choose **Delete Job** from the menu that displays.
3. Select **Yes** to delete the selected job and associated mirror(s).

Reassigning a Job

Use the **Reassign Job** function to move an existing mirror from one job to another without deleting the mirror.

1. Select the job from the middle **Summary** panel.
2. Right-click and select **Reassign Job** or select **Reassign Job** from the **Actions** panel.
3. Select an existing job from the **Existing Jobs** dropdown list and press the **Assign Job** button. The new job assignment will display in the middle **Summary** panel.

Note: You can also choose to **Create a New Job** from this dialog if you do not want to use an existing job.

Switching Over a Mirror

The Switchover Mirror function enables you to switch over all the mirrors in a job or just one of the mirrors in a job. A "mirror" includes all variants for mirrors such as standard single-target replication and complex geometries such as multi-target replication and shared node sources and targets. These complex mirror configurations and geometries actually implement a related collection of individual mirrors working as a single unit.

Note: Before switching over a mirror to the current target system, the mirror must be in the **Mirroring** state. Please see the **Requirements for Switchover** table below to understand switchover requirements in multiple target and shared source/target configurations. Please use the DataKeeper GUI to view the state of the mirror; the WSFC GUI will not provide that level of detail and will state that the resources are on-line (Green) even when the mirrors are not in the mirroring state.

1. Select the job in the left column tree pane.
2. Right-click on the selection and select **Switchover Mirrors**.
3. A dialog displays allowing you to designate which node/host(s) in the selected job or mirror should become the new mirror source.

In the case of complex mirrors, it is valid to choose either a shared peer of the current mirror source or any one of the active targets that are currently in the mirroring state. Choosing a shared peer of an active target or one that is not currently mirroring will result in an error and leave the current mirror status and configuration unchanged.

4. An hour glass will appear over the mirror icon in the left tree panel.
5. You can confirm the switchover is complete by checking the mirror status in the **Summary** panel.

Note: If the **Switchover** option is grayed out (not available), this could mean the volume is under clustering protection (Microsoft clustering or SIOS Protection Suite clustering).

Requirements for Switchover

Configuration Type	Example Configuration	Switchover Action	Requirements for Switchover
Single Target	A → B	Switchover to B	Allowed if mirror is in MIRRORING STATE
Multiple Target	A → B (mirroring)	Switchover to B	Allowed because A→B mirror is in MIRRORING state
	A → C (paused)	Switchover to C	Not allowed
	*S1,S2,S3 → *T1,T2	Switchover to shared source (S2 or S3)	Always allowed
Shared Source/Targets	(S1 is current source)	Switchover to current target (T1)	Only allowed if mirror in MIRRORING state
	(T1 is current target)	Switchover to shared target (T2)	Not allowed – Switchover will fail

Working With Mirrors

[Managing Mirrors](#)

[Pause and Unlock](#)

[Continue and Lock](#)

[Partial Resync](#)

[Break](#)

[Resync](#)

[Deleting a Mirror](#)

[Replacing a Target](#)

[DataKeeper Volume Resize](#)

[Mirror Properties](#)

[Changing the Compression Level of an Existing Mirror](#)

Managing Mirrors

From the **Actions** pane, you can select a job and manage all the mirrors in a job, or you can perform an action on a single mirror in a job.

After selecting a job, you can:

- [Pause and Unlock](#) All Mirrors
- [Continue and Lock](#) All Mirrors
- [Break](#) All Mirrors
- [Resync](#) All Mirrors
- [Switchover](#) All Mirrors

The target-level actions (at the bottom of the **Actions** pane) are for individual mirrors. For example, if you have a job with two mirrors and you select one of the mirrors then choose the target **Pause and Unlock Mirror** action, only the selected mirror would be paused.

Pause and Unlock

This command pauses the mirror and unlocks the volume on the target system. You may wish to unlock the target volume in order to make a backup of the volume.

Warning: Any writes to the target volume while it is unlocked will be lost when the mirror is continued.

Note: If replacing the target volume, either [break the mirror](#) or [delete the mirror](#) in order to ensure a full resync of the data from the source volume to the new target volume when the new target volume is in place. See [Replacing a Target](#) for further information.

The [Continue and Lock](#) command will relock the target volume, perform a partial resync.

1. Select the job that contains the mirror you want to unlock.
2. Right-click on the job selection and choose **Pause and Unlock All Mirrors** or select **Pause and Unlock All Mirrors** from the **Actions** task pane.
3. Select **Yes** to pause and unlock all mirrors in the selected job.

Continue and Lock

This action locks the volume on the target system and then resumes the mirroring process.

While the mirror is paused, writes on the source system are recorded in the SIOS DataKeeper [Intent Log](#). When the **Continue and Lock** operation occurs, these changed blocks - along with any blocks that also changed on the target volume - are sent from the source to the target, and the mirror is resynchronized in what is called a [Partial Resync](#).

Warning: Any writes to the target volume while unlocked are lost when the mirror is continued.

Note: If *replacing* the target volume, either [Break](#) the mirror or [Delete the Mirror](#), which requires either a **Resync** or **Recreate** instead of **Continue and Lock**. See [Replacing a Target](#) for further information.

1. Select the job that contains the mirror you want to continue.
2. Right-click on the job selection and choose **Continue and Lock All Mirrors** or select **Continue and Lock All Mirrors** from the **Actions** task pane.
3. Select **Yes** to continue and lock all mirrors in the selected job.
4. The mirror state will change to **Mirroring** in the **Mirror Summary** window.

Partial Resync

A partial resync is the resynchronization of only the necessary data to establish the new end points and is usually much quicker than a full resync.

Break

Breaking a mirror is similar to the **Pause and Unlock** function. It suspends mirror operation and unlocks the target volume for read/write access. The difference is that the **Break** operation marks all bits in the DataKeeper [Intent Log](#) as dirty, which forces a full resync to occur when the mirror is resync'ed to resume mirroring.

Warning: Do not write to the target volume while the mirror is broken. Any writes to the target while the mirror is broken will be lost when the mirror is resynchronized.

1. Select the job that contains the mirror you want to break.
2. Right-click on the job selection and choose **Break All Mirrors** or select **Break All Mirrors** from the **Actions** task pane.
3. Select **Yes** to break all mirrors in the selected job.
4. The mirror state will change to **Broken** in the **Mirror Summary** window.

Note: The **Resync** command will relock the Target volume, perform a **full resync** and resume the mirroring process.

Resync

Use this command to re-establish a broken mirror. A full resync will be performed.

1. Select the job that contains the mirror you want to resync.
2. Right-click on the job selection and choose **Resync All Mirrors** or select **Resync All Mirrors** from the **Actions** task pane.
3. Select **Yes** to resync all mirrors in the selected job.
4. The mirror state will change to **Mirroring** in the **Mirror Summary** window.

Deleting a Mirror

This action discontinues replication and removes the mirror from the associated job. The target volume is unlocked and made fully accessible.

1. Select the job that contains the mirror you want to delete.
2. Right-click on the mirror and choose **Delete Mirror** or select **Delete Mirror** from the **Actions** task pane.
3. Select **Yes** to delete the mirror.
4. The mirror will be deleted and removed from the associated job.

Note: If the **Delete Mirror** option is grayed out (not available), this could mean the volume is under clustering protection (Microsoft clustering or SIOS Protection Suite clustering).

Replacing a Target

When replacing the target volume, you must either [break the mirror](#) or [delete the mirror](#) in order to ensure a full resync of the data from the source volume to the target volume when the target volume is back in place. Though similar to the [Pause and Unlock](#), breaking the mirror marks all bits in the DataKeeper Intent Log as dirty which forces a full resync to occur. Deleting the mirror discontinues replication altogether removing the mirror from the job so that when your mirror is recreated with the new target, a full resync will be performed.

Using the BREAK Command

1. Select the mirror that contains the target you want to replace.
2. Right-click on the mirror and choose **Break Mirror** or select **Break Mirror** from the **Actions** task pane.
3. Select **Yes** to break the mirror.
4. Once new target is in place, right-click on the job that contains the replaced volume and choose **Resync All Mirrors**.
5. The target volume will be locked, a full resync will be performed and the mirroring process is resumed.

Using the DELETE Command

1. Select the mirror that contains the target you want to replace.
2. Right-click on the mirror and choose **Delete Mirror** or select **Delete Mirror** from the **Actions** task pane.
3. Select **Yes** to delete the mirror.
4. Once new target is in place, [recreate the mirror](#).

DataKeeper Volume Resize

DataKeeper allows users to extend and shrink their DataKeeper volumes dynamically while retaining mirror settings. Once the resize is complete, a partial resync will be performed.

Note: This resize procedure should be performed on only one volume at a time.

!WARNING Do NOT attempt to perform the resize in releases prior to DataKeeper for Windows v7.4.

Restrictions

- DataKeeper does not support changing the disk type of the physical disk where a mirrored volume is located (for example, **Basic Disk** to **Dynamic Disk** – mirror must be deleted prior to creating your dynamic disk).
- DataKeeper does not support third-party partition resizing products.
- DataKeeper does not support volume resizing on shared volumes configured on **Dynamic Disks**. Windows cannot reliably use a shared Dynamic Disk.

Non-Shared Volume Procedure

Example configurations for using this procedure include the following:

[Disk-to-Disk](#)

[One-to-One](#)

[One-to-Many 'Multiple Targets'](#)

[Many-to-One](#)

To resize your DataKeeper volume in a non-shared volume configuration, perform the following steps.

1. Pause all mirrors and unlock all target volumes via the [Pause and Unlock](#) mirror option in the DataKeeper UI.

Note: A mirror must be in a "mirroring" state in order to Pause or Unlock it.

2. Using the **Windows Disk Management** utility, increase (or decrease if allowed by the Operating System) the volume size on the source system by selecting "**Extend Volume**" or "**Shrink Volume**" in the **Resizing Wizard**. Once that resize is complete and verified, **resize the target system(s)**. Make sure that the raw volume size of each target is greater than or equal to the size of the source volume.

Note: The Windows Disk Management utility will take longer to start on the target node based on the number of drives. Because the Windows operating system has error condition retries built in when a volume is locked, the speed with which it starts on the "locked" target node is affected.

Note: After the resizes on the source and target you will need to run a Rescan in Disk Management. Then, you will need to run the following on each system in the cluster so that DataKeeper sees the new volume size:

- Go to a command prompt (run as administrator)
- `cd %extmirrbase%`
- `emcmd . updatevolumeinfo <enter-volume-letter>`

3. [Continue and Lock](#) the mirrors after volumes have been resized. The mirroring process should resume and a partial resync should occur.

Shared Volume Procedure - Basic Disk

This resizing procedure will work on shared volumes if the shared volume is configured on a **Basic Disk**. Example configurations for using this procedure include the following:

*Shared Volume - More than one system has access to the same physical storage. This shared volume can be either on the source side of the mirror or on the target side.

[N-Shared-Disk Replicated to One](#)

[N-Shared-Disk Replicated to N-Shared-Disk](#)

[N-Shared-Disk Replicated to Multiple N-Shared-Disk Targets](#)

If there is free space on the disk, the volume can be extended to use the additional space.

1. Pause all mirrors and unlock all target volumes via the [Pause and](#)

[Unlock](#) mirror option in the DataKeeper UI.


2. Shut down (power off) all shared source and/or shared target systems. (**Note:** Current source and current target systems should not be shut down.)
3. Change the volume sizes as noted above in the Non-Shared Volume procedure.
4. [Continue and Lock](#) the mirrors after resizing has completed.
5. Power on all shared systems. The new volume configuration will automatically be recognized.

Error Handling:

1. After performing the **Continue and Lock**, if the GUI abnormally maintains the "**Paused**" mirror state, check the system logs on both source and target nodes.
2. DataKeeper will prevent a mirror resync from starting if the target volume is smaller than the source volume. If the system logs show such an error, the target volume must be unlocked manually via the [UNLOCKVOLUME](#) command, and the volume must again be resized making sure that the volume size of the target is greater than or equal to the size of the source volume. Then proceed with the Continue and Lock step above.
3. DataKeeper, upon continuing the mirror, will reallocate the bitmap file and in-memory bitmap buffer using the new volume size. In the event DataKeeper is unsuccessful in reallocating the bitmap buffer - due to insufficient memory resources on the source or target - the mirror will be placed into a '**Broken**' state which will require a FULL resync.
4. Once resizing a volume has begun, there is no way to back out of the resizing feature and the associated error handling as DataKeeper will have to reallocate the bitmap file and in-memory bitmap buffer. Any failure of this reallocation procedure will break the mirror and force a FULL resync.

Mirror Properties

Select a job in the **Job Summary** pane and right-click to choose **Mirror Properties**.



This dialog displays the following information about the job, source and target systems:

- **Job Name**
- **State** (current state of the job; for example, Active)
- **Source System**
 - Server - name of source server
 - Source IP - IP address of source server
 - Disk Space - capacity of the source volume

- Shared Hosts - other systems that have access to this volume via shared storage
- **Target System**
 - Server - name of target server
 - Target IP - IP address of target server

You can modify the following settings through the **Mirror Properties** dialog:

- [Compression Level](#) - specifies the compression level for the given mirror. The value can be set from lowest to highest. We recommend a level of "**Medium low**", but users should test several different settings to see what level works best in their specific environment. Compression is typically not required for LAN connections > 100 Mbps.

Note: Any changes made to the compression level setting are automatically propagated to all the systems listed in the **Mirror Properties** display.

- [Maximum Bandwidth](#) - Specifies the maximum amount of network bandwidth (in kilobits per second) that a particular mirror is allowed to use. A value of 0 means unlimited.

Note: In a multi-target configuration where A is mirroring to B and C, the properties of the mirror between B and C cannot be set until B or C becomes the source.

Changing the Compression Level of an Existing Mirror

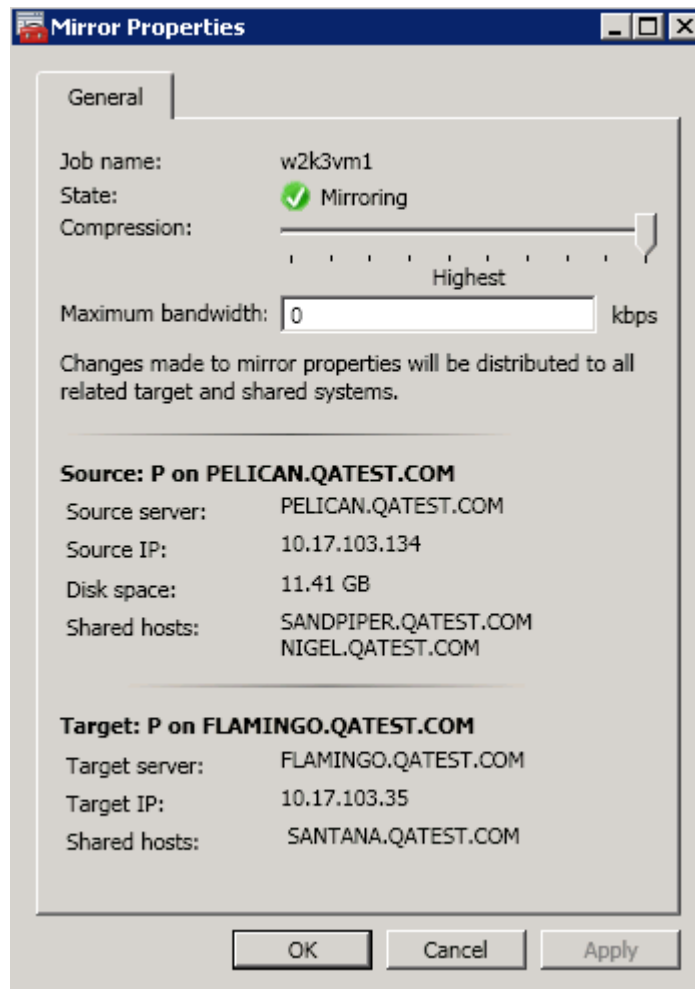
The compression level of the mirror is set during mirror creation and applies to that specific mirror only.

To change the compression level of an existing the mirror, edit the properties of the mirror from within the DataKeeper GUI.

1. Select the mirror and click on **Edit**.
2. Change the compression level by dragging on the slider button.

The values change from lowest to the highest. We recommend a level of "Medium low", but users should test several different settings to see what level works best in their specific environment.

Also note that by changing the parameter as the comment suggests in the dialog, the compression properties will be propagated to all the systems listed in the [Mirror Properties](#) display.



The image shows a Windows-style dialog box titled "Mirror Properties". It has a "General" tab selected. The dialog contains fields for "Job name" (w2k3vm1), "State" (Mirroring, indicated by a green checkmark), and "Compression" (a slider set to "Highest"). There is a "Maximum bandwidth" field set to "0 kbps". A note states: "Changes made to mirror properties will be distributed to all related target and shared systems." Below this, there are two sections: "Source: P on PELICAN.QATEST.COM" and "Target: P on FLAMINGO.QATEST.COM". Each section lists "Source server", "Source IP", "Disk space", and "Shared hosts". At the bottom are "OK", "Cancel", and "Apply" buttons.

General	
Job name:	w2k3vm1
State:	<input checked="" type="checkbox"/> Mirroring
Compression:	<input type="range"/> Highest
Maximum bandwidth:	0 kbps
Changes made to mirror properties will be distributed to all related target and shared systems.	
Source: P on PELICAN.QATEST.COM	
Source server:	PELICAN.QATEST.COM
Source IP:	10.17.103.134
Disk space:	11.41 GB
Shared hosts:	SANDPIPER.QATEST.COM NIGEL.QATEST.COM
Target: P on FLAMINGO.QATEST.COM	
Target server:	FLAMINGO.QATEST.COM
Target IP:	10.17.103.35
Shared hosts:	SANTANA.QATEST.COM

OK Cancel Apply

Working With Shared Volumes

[Managing Shared Volumes](#)

[Adding a Shared System](#)

[Removing a Shared System](#)


Managing Shared Volumes


Once your mirrors have been created, DataKeeper allows you to manage your shared volumes. By choosing **Manage Shared Volumes** from the DataKeeper GUI, you can [add another system](#), which is sharing a mirrored volume, to a job. It also allows you to [remove a shared system](#) from a job. These systems can exist on either the source side or the target side of the mirror.

To add or remove a system that is sharing a mirrored volume on either the source or target end of a mirror, select the job that you want to manage and highlight the mirror that contains the volume that is to be edited.

If a volume is mirrored to more than one target and you want to add or remove a shared system on the source side of the mirror, you can choose any of the mirrors, since they all refer to the same source volume. Choose the **Manage Shared Volumes** action for that mirror, and the **Shared Volumes** dialog will appear.

If you want to add or remove a shared system on the target side of the mirror, you must select that specific mirror.

 **Shared Volumes**

 **Source Shared Volumes**

Source Shared Volumes

Target Shared Volumes

Source server: first1.simulated.org

Source IP address: 110.1.0.1

Source volume: W

Choose the systems that have volumes which are shared with the system above. Uncheck the "Include" box if any system should not be included in the job. Shared volumes are required to be configured using the same subnet as their peers. If any systems do not have an IP address on the selected subnet, their IP address will be displayed as empty. These systems cannot be included in the job using this subnet.

Include	Server	Volume	Source IP / Mask
<input checked="" type="checkbox"/>	first2.simulated.org	W	110.1.0.2 / 8
<input checked="" type="checkbox"/>	first3.simulated.org	W	110.1.0.3 / 8
<input checked="" type="checkbox"/>	first4.simulated.org	W	110.1.0.4 / 8

[Connect to Server](#)

NextCancel

Adding a Shared System

To add a shared system to either the source or target side of a mirror, you must be connected to that system. You can connect to the system prior to starting the **Manage Shared Volumes** dialog, or you can click **Connect to Server** from within the dialog. In either case, if there are shared volumes that exist on that system that match either the source or target volume, the system and its matching IP address will be displayed in the correct page of the dialog. Leave the **Include** box checked to include the system in the job configuration and choose the correct IP address to be used for that system.

If a shared system does not have an IP address whose subnet matches the existing mirrored systems, the IP Address field will be blank and the **Include** box will be unchecked. You must reconfigure the system so that it has an IP address on that subnet. Then try adding the shared volume again.

When you click **Done** after adding a new shared system, it will be added to the job. If there are multiple mirrors in place, you will be asked to provide the network addresses to be used between the newly-added system and all other targets.

Removing a Shared System


To remove a shared system from either side of the mirror, bring up the **Manage Shared Volumes** dialog and uncheck the **Include** box for the system to be removed. When you click **Done**, the job will be updated so that the system is not part of the job.

Warning: If a shared system is removed from the source side of the mirror, the source volume is now accessible on multiple systems and simultaneous access of the source volume could result in data corruption.

Using Microsoft iSCSI Target With DataKeeper on Windows 2012

The following topics will guide you in setting up Microsoft iSCSI Target with DataKeeper via the user interface.

*NOTE: This configuration is not supported in a VMware ESX environment.



[Installation of the iSCSI Target](#)

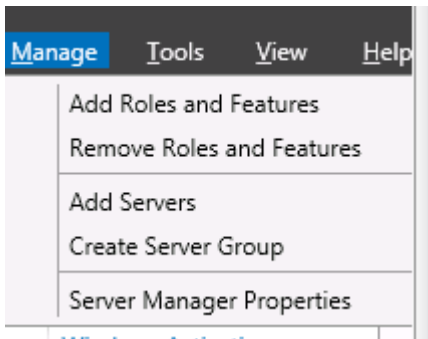
[Creation of Mirror and Configuration of Cluster](#)

[Creation of iSCSI Virtual Disks](#)

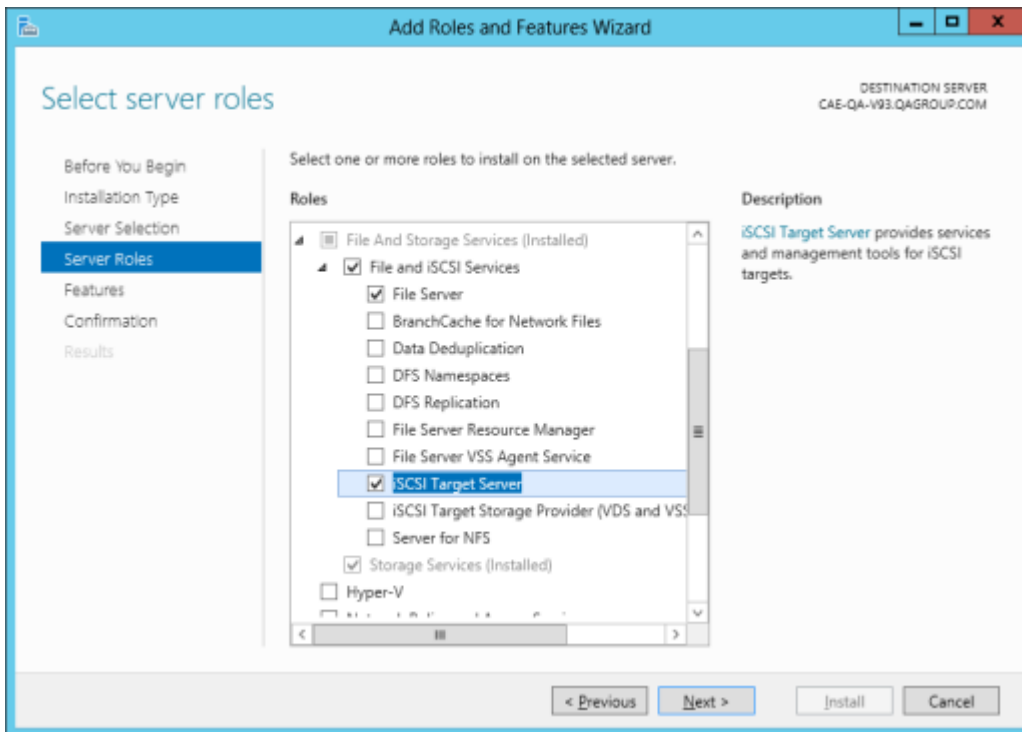
[Setup of iSCSI Initiator on Windows 2012](#)

Installation of the iSCSI Target

1. From the **Server Manager** menu, select **"Add Roles and Features"** from the **"Manage"** drop-down.



2. Select the **"Role-based or feature-based installation"** option.
3. From the list of servers presented, select the appropriate server.
4. On the **"Select Server Roles"** screen under **"Server Roles"**, navigate to and select **"File and iSCSI Services"** / **"iSCSI Target Server"**.
Note: **"File and iSCSI Services"** is in the tree hierarchy under **"File and Storage Services"** which is typically shaded and difficult to find.



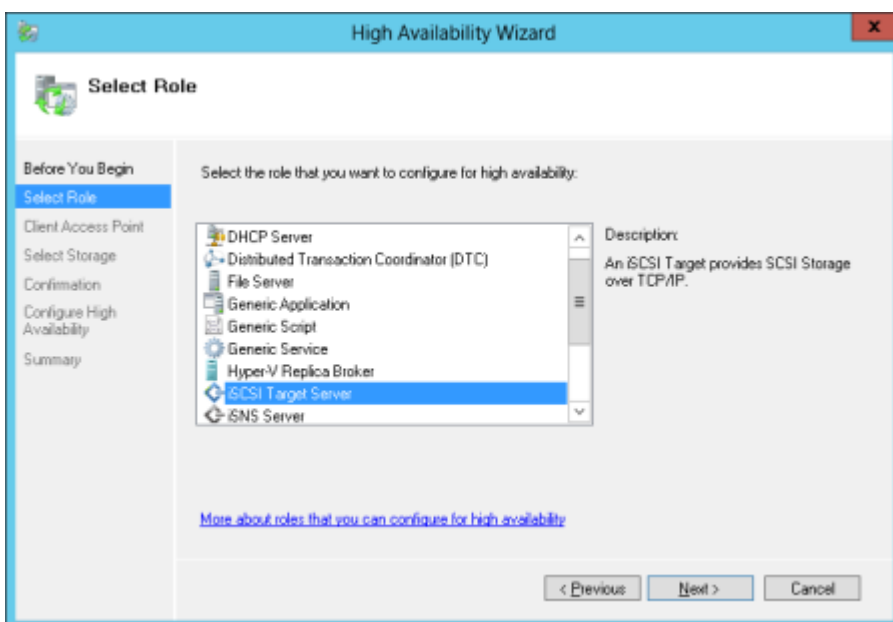
5. Click "**Next**" twice to get to the "**Install**" button to be able to install the role.
6. The feature will install and the progress will be shown.
7. Upon completion, the message "**Installation succeeded**" will be displayed.
8. Repeat these steps for all servers in the cluster.

Mirror Creation and Cluster Configuration

1. Create your **DataKeeper volumes** and your **cluster**. See [Creating a DataKeeper Volume Resource in WSFC](#) for reference.

***IMPORTANT:** The iSCSI Target Role only supports DataKeeper Volumes that are mirrors of **Simple Volumes** placed on **Basic Disks**. If any of your mirrors are using volumes such as Striped or Spanned volumes on a Dynamic Disk on either the source or target system, then you cannot create an iSCSI Target role that uses those DataKeeper Volume resources for storage.

2. From the **Windows Failover Cluster Manager UI** (cluadmin.msc), select **Configure Role** and navigate to the screen to select the **iSCSI target role**.

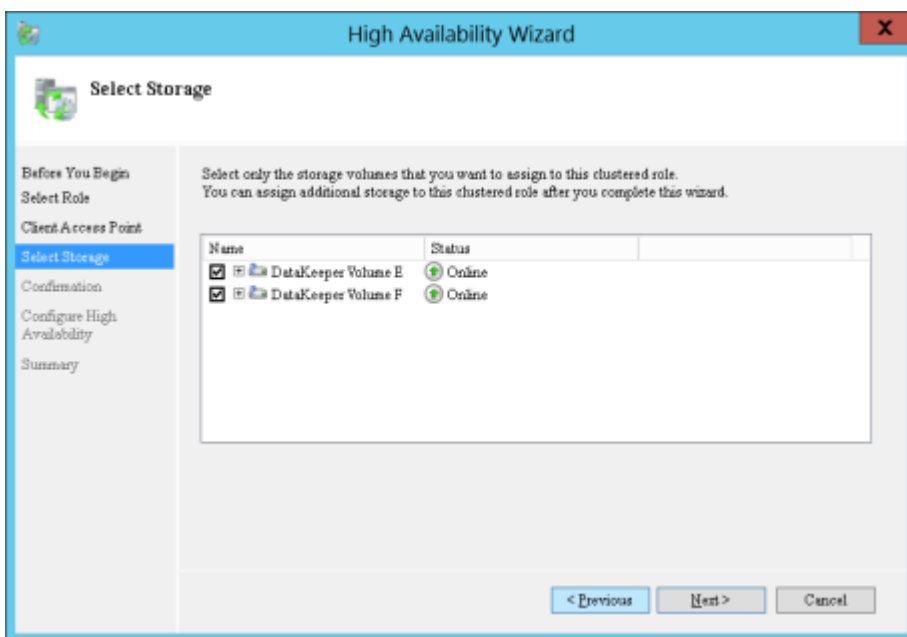


3. Select **iSCSI Target Server** role and select **Next**.
4. The **Client Access Point** page appears. Type the **Client Access Point name** and **IP address** for the iSCSI Target Server instance.


***IMPORTANT:** This name and IP address will be used later by clients to access the server address, so it should be recorded in DNS. This is very

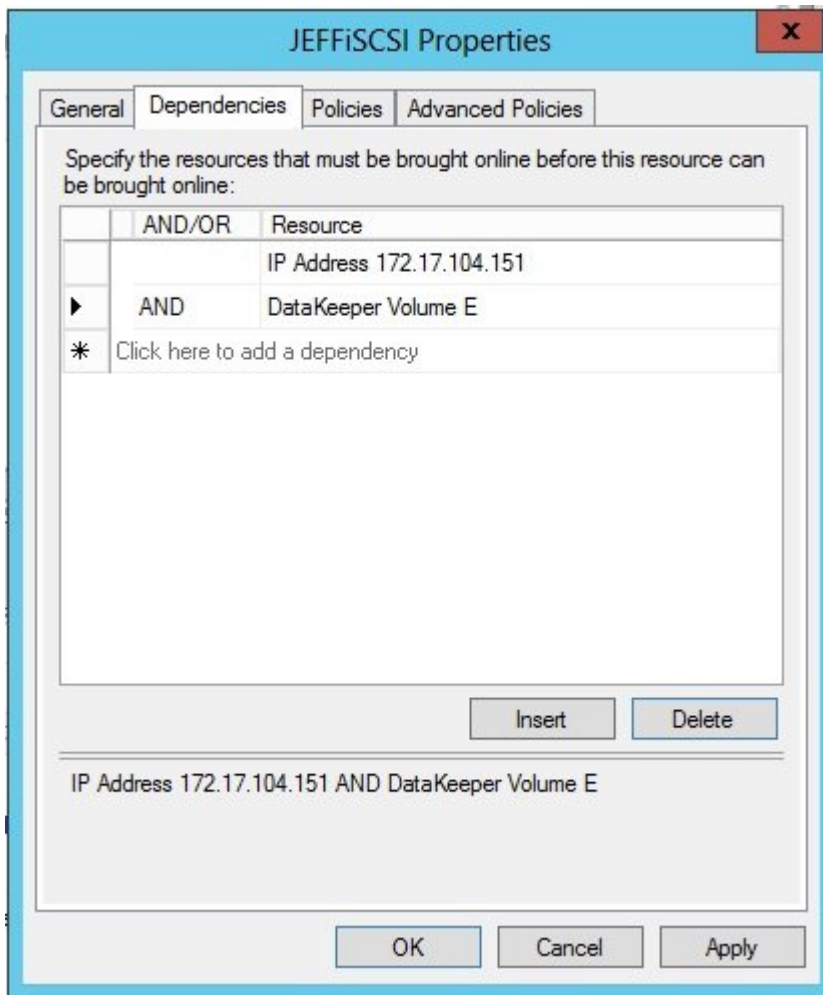
important for the servers to be able to resolve these names.

5. On the **Select Storage** dialog, select your **DataKeeper volume(s)**.



6. With the next set of screens, you should be able to complete the configuration.
7. Following setup, from the **Failover Cluster UI**, add dependencies for the DataKeeper volume(s).
 - a. Click on **Roles** in the left pane, then click on the **iSCSI Target Server** resource in the top center pane.
 - b. In the lower center pane, select the **Resources** tab, then right-click on the **Name: <client access point name>** under the **Server Name** heading and select **Properties**.
 - c. Select the **Dependencies** tab and add the appropriate DataKeeper volume(s) as dependencies.



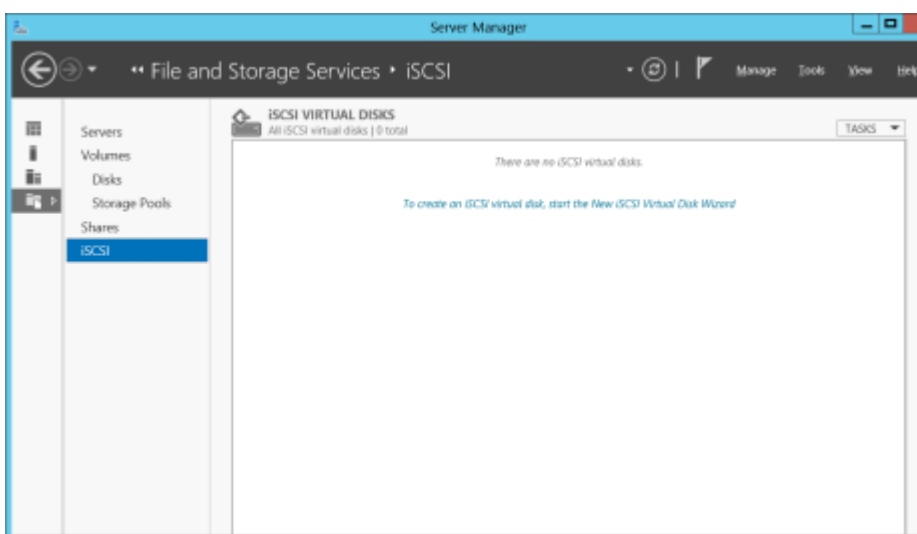


8. Setup is complete. Proceed to the [iSCSI Virtual Disks](#) configuration.


Creation of iSCSI Virtual Disks

Perform the following on the **primary server**, wherever the **iSCSI Target server is online at the moment**.

1. From **Server Manager**, navigate to **File and Storage Services** and select **iSCSI**. Click on the link "**To create an iSCSI virtual disk, start the New iSCSI Virtual Disk Wizard**". (Alternatively, select **New iSCSI Virtual Disk** from the **TASKS** drop-down menu on the upper right of the screen.) **Note:** Windows Server 2012 **Server Manager** inherently takes some time to display or update the information presented to the user.



2. On the **New iSCSI Virtual Disk Wizard**, you will see the server and the volume. Select the **DataKeeper volume** and click **Next**. (**Note:** The server name is the name created in the [previous step](#) and the volume is the DataKeeper volume exposed.)



3. Follow the next panel to configure the **iSCSI Virtual Disk**.

a. Specify **iSCSI Virtual Disk Name**.

b. Specify **iSCSI Virtual Disk Size**. (**Note:** Multiple files can be created. If file size spans the entire disk, the OS may warn that disk is low since the VHD file(s) created can consume the entire disk.)

c. Designate whether the iSCSI Virtual Disk will be assigned to an **Existing iSCSI Target** or a **New iSCSI Target** on the **Assign iSCSI Target screen**. (See [below](#) for an explanation on when to select **Existing iSCSI Target**.)

d. Specify **iSCSI Target Name**.

e. On the **Access Servers** screen, select **Add**. Add the **iSCSI Initiators** that will be accessing this **iSCSI Virtual Disk**. **Note:** The iSCSI Initiators should be added one at a time.

4. Once all the answers have been provided, the iSCSI virtual disk/target creation is complete. Proceed to configuration of the [iSCSI Initiator](#).

Setting Up Multiple Virtual Disks Within the Same Target Name

It is also possible to set up multiple iSCSI virtual disks within the same iSCSI target name. Whenever an iSCSI initiator connects to such a target, it will connect to all of the virtual disks that have been assigned to that name.


You need to have a plan ahead of time that describes which files you want to create and whether those files should all be accessed simultaneously or if the disks need to be accessed separate from one another.

To set up multiple virtual disks within the same target name, on Step 3c, instead of selecting **New iSCSI Target** on the **Assign iSCSI Target** screen, select **Existing iSCSI Target** and specify the iSCSI target name that was created previously. This target name will appear in the list of "targets" when an iSCSI Initiator connects to the iSCSI Target Server. If a target has more than one virtual disk associated with it, then the initiator will get a connection to each of those disks (they will appear as a new Disk in Disk Management).


Setup of iSCSI Initiator on Windows 2012

Once the virtual disks/targets are created, each of the cluster servers must initiate a connection to them via Microsoft's iSCSI Initiator.

1. From "**Administrator Tools**" in "**Server Manager**", start "**iSCSI Initiator**".
2. Select the "**Targets**" tab and enter the **Network Name** or **IP address** of the **clustered iSCSI Target** created in the [previous step](#). Select "**Quick Connect**".








3. New panel should indicate that "**Login has succeeded**". Click **OK** to hide the panel.
4. Start "**Disk Manager**". The new iSCSI virtual disk will be displayed and can be initialized.





5. Right-click on the disk(s) to bring it online.
6. Initialize the disk(s).
7. Create the new volume and assign the drive letter.
8. Configuration is now complete.


DataKeeper Notification Icon

The DataKeeper Notification Icon is an application that will show a summary of your DataKeeper mirrors in the Windows Notification Tray. The icon displayed indicates what conditions have been detected, with the following priority.

-  Error: An error condition, such as split brain (requiring manual intervention), has been detected.
-  Warning: This indicates that there is a condition that may require administrative intervention, such as a mirror being paused or broken, or a transient split-brain condition.
-  Resync: This indicates that a mirror is in the resync or resync pending state.
-  Mirroring: This indicates that all mirrors are in the mirroring state.
-  Disabled: This indicates that status updates are no longer occurring. During this state none of the other status conditions will be displayed.

More details about these conditions can be found by hovering over the DataKeeper Notification Icon, such as how many mirrors are in each state, or the nature of the detected error condition. Some examples are shown below.





Note: The DataKeeper Notification Icon uses DataKeeper jobs to determine which remote systems to poll for information. Only the status of mirrors in jobs that contain the node on which the DataKeeper Notification Icon is running will be reported.

In addition to the display functions, the DataKeeper Notification Icon also serves as a shortcut to managing your DataKeeper mirrors. Double clicking on the DataKeeper Notification Icon will launch the DataKeeper GUI.

Right clicking will bring up a menu with the following options:

- Launch DataKeeper GUI - Launches the DataKeeper GUI.
- Launch License Manager - Launches the SIOS License Manager.
- Launch Health Check - Opens a command prompt and runs [DKHealthCheck](#).
- Gather Support Logs - Runs [DKSupport](#) and opens an explorer window to the location containing the new archive.
- Set Refresh Rate - Will allow you to set how often the icon refreshes its state information.
- Disable/Enable Status Updates - Disables and Enables Status Updates. Requires EmTray to be run with administrator permissions.
- Exit - Stops and closes the DataKeeper Notification Icon.

Auto-Start at Login


The Notification Icon should automatically appear in the Windows Notification Tray upon logging in to the node.

To disable this functionality, delete the shortcut to EmTray.exe from the following location:

- C:\ProgramData\Microsoft\Windows\Start Menu\Programs\Startup

To re-enable this functionality, simply create a shortcut from the EmTray.exe (located at \DKTools) to the same location above.

Note: By default, Microsoft Windows will hide Notification Tray Icons. You can change this by going to the 'Notification Area Icons' option in Control Panel and changing the settings for the DataKeeper Notification Icon to 'Show icon and notifications'



DataKeeper Target Snapshot

Overview

DataKeeper's target snapshot feature, integrated with both DataKeeper and DataKeeper Cluster Edition, is the process of creating point in time copies of replicated volumes allowing access to data on a standby cluster node without interrupting data replication from the source system. Data protection is not lost for any period of time. Enabling target snapshot allows data to be used on an otherwise idle target node without negatively impacting the performance of the source.

Without target snapshot, DataKeeper and DataKeeper Cluster Edition are able to maintain a real-time replica of their source system's data on the target system. However, this replica cannot be accessed without pausing the mirror and unlocking the target system. Mirror failover and switchover cannot occur while in this paused and unlocked state, making the protected application less highly available. Application-consistent target snapshot allows access to data on the target system while maintaining high availability of the running application on the source system. The mirror remains in the **mirroring state** and continues to update the target volume with all writes from the source. Target snapshot integrates with Volume Shadow Copy Service (VSS) to ensure that the data which is exposed on the target system is in an application-consistent state.

When To Use Target Snapshot

DataKeeper Target Snapshot is an alternative to using the "Pause and Unlock" command to access data on your target system. Target Snapshot provides the following benefits that Pause and Unlock does not:

- The mirror remains in the Mirroring state, and data continues to be replicated from the source system with no interruption.
- Multiple volumes can be snapshotted simultaneously.
- VSS-aware applications (like MS SQL Server) that are running on the source system are briefly quiesced using VSS in order to ensure that the data exposed on the Target system is in an application-consistent state.

How To Use Target Snapshot

Define Snapshot Location on Target system

In order to use the Target Snapshot feature, a Snapshot Location must be defined on the target system for each volume you plan to access. The Snapshot Location can be defined in the DataKeeper GUI, in the Mirror Properties dialog. It can also be defined by running the [EMCMD SETSNAPSHOTLOCATION](#) on the target system.

```
EMCMD <system> SETSNAPSHOTLOCATION <volume letter> "<directory path>"
```

Enable SIOS VSS Provider on Source system

DataKeeper target snapshot uses VSS to quiesce data on the mirror source system. DataKeeper has a VSS Provider component which is used to accomplish this. However, due to reported interference of the SIOS VSS Provider with some backup products, the provider is shipped in a disabled state. In order to take a snapshot, the VSS Provider on the mirror source system must be activated.

To activate the SIOS VSS Provider, run the script "install-siosprovider.cmd" which is located in "%ExtMirrBase%\VSSProvider".

After you have taken a Target Snapshot, you may choose to de-activate it on the mirror source system by running the command "uninstall-siosprovider.cmd" in the same folder. If you are using a backup product that is incompatible with the SIOS VSS Provider, you should de-activate it using this command (see [Known Issues](#) below for incompatible products). However, if you are not using a product with such an incompatibility, you can leave the VSS Provider activated. **Note:** Any DataKeeper update will disable the provider, it must be re-enabled in order to take a target snapshot after this occurs.

The SIOS VSS Provider is only needed at the time that a snapshot is taken. The snapshot can be left in place on the target system after the provider has been deactivated, and the snapshot can be dropped while the provider is deactivated.

Execute the TAKESNAPSHOT command

After the Snapshot Location has been defined for each mirrored volume, and the SIOS VSS Provider is activated on the source system, the volumes can be made accessible on the target system by running the [EMCMD TAKESNAPSHOT](#) command:

```
EMCMD <target_system> TAKESNAPSHOT <volume letter> [<volume letter>...]
```

where <target_system> is the name or IP address of the target system, <volume letter> is the drive letter of one of the volumes to be snapshotted, [<volume letter>...] is the (optional) drive letter of another drive to be snapshotted at the same time, etc.

How Target Snapshot Works

DataKeeper target snapshot uses a copy-on-write strategy to maintain and expose a view of the volume at a particular point in time. A snapshot file is used to store the volume information. Configuring the location of this snapshot file is the first step toward enabling target snapshot.

When the EMCMD TAKESNAPSHOT command is run, DataKeeper will create and mount a snapshot file in the configured snapshot folder. A request is then sent to the source system telling it to use VSS to quiesce any VSS writers on the given volume and notify the target when all write operations to the disk are stopped and the volumes are in a well-defined state.

Quiescing the Database/Application

The application-consistent capabilities of this feature integrate with Volume Shadow Copy Service (VSS) to ensure that the data that is exposed on the target system is in an application-consistent state. Once a snapshot is requested, the VSS service pauses the systems and ensures that all applications modifying data on disk bring all their files into a consistent state prior to the creation of the snapshot. This is called quiescing the database/application. Rather than shutting down the database and reopening it in restricted mode, quiescing temporarily freezes application write I/O requests (read I/O requests are still possible) for the short time required to create the snapshot. Once in the quiesced state, the snapshot on each volume is initiated by adding the snapshot message to the driver mirror write queue(s). VSS will then unfreeze the applications and the volume is unlocked, thus minimizing the amount of time the apps are quiesced. The user can now perform actions on the target system while the mirror remains in the **Mirroring** state and the application on the source system remains highly available.

Read and Write I/O Requests

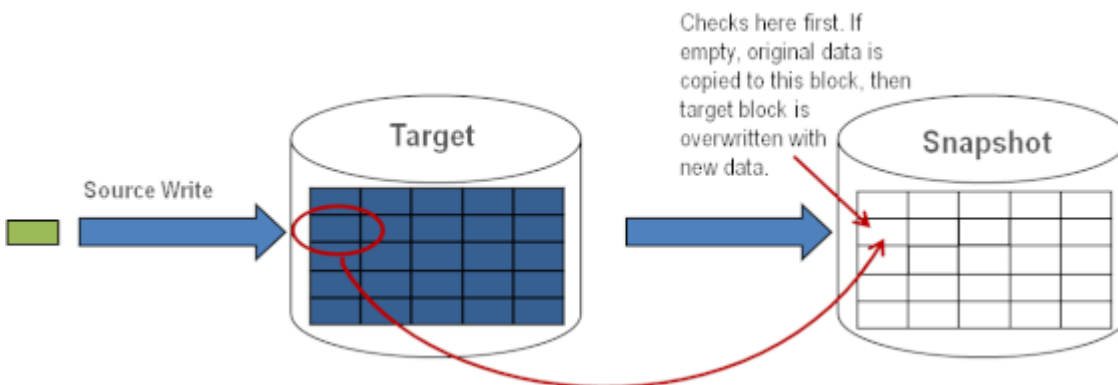
The snapshot exists in parallel with the live copy of the volume to be backed up, so except for the brief period of the snapshot's preparation and creation, an application can continue its work. Writes to the target, however, will now be processed differently while the target is in this

state.


Data mirroring from the source system will continue uninterrupted, but any new data from the source that is received after the snapshot is taken will not be visible on the target system until the snapshot is dropped. This allows an application on the target system to run, using (and updating) data that represents the source system's data at the point in time that the snapshot was taken.

Source Write

In order to accomplish source writes, when new data comes from the source, DataKeeper first determines if that particular block of data has already been written to the snapshot file.



If the block has not been written to, as shown above, that **original** block is written to the snapshot file in order to preserve the snapshot data, then the new data is written to the target. The result is shown below.




If DataKeeper determines that this block has already been written to the snapshot file, then this step is skipped and the block is

just written to the target. For blocks on the source volume that are overwritten frequently, the snapshot file only has to be updated once, the first time that block is written after the snapshot is taken.


Local Write

If local writes are performed on the target (from applications on the target system), these writes are stored in the snapshot file and do not overwrite any blocks on the replicated volume itself. (**Note:** Any local writes stored in the snapshot file will be lost when snapshot is dropped.)



Target Read Request

Read requests on the target volume will return snapshot data. This is accomplished by first reading data from the blocks written in the snapshot file. Any blocks that have not been saved to the snapshot file will be read from the target volume.



Using Target Snapshot

There are three tasks that must be performed when using target snapshot. The [snapshot location must be configured](#), the [snapshot must be initiated](#), then once target reporting actions are complete, the [snapshot must be dropped](#).

Configuring the Snapshot Location

When target snapshot is initiated, DataKeeper creates and mounts a file in the snapshot location to hold the snapshot data. This location must be configured prior to initiating a snapshot. See [Files / Disk Devices / Registry Entries](#) below for more information about the mounted snapshot disk(s).

When configuring the snapshot location, make sure it meets the following criteria:

- Is only used when a snapshot is requested.
- Cannot be stored on a DataKeeper mirrored volume.
- Can store multiple snapshot files for different volumes.
- Must have enough free space to create and accommodate a file that will grow depending on the source mirrored volume size and writes during snapshot use.

Note: Do not change the snapshot location during a snapshot.

Snapshot Location Size

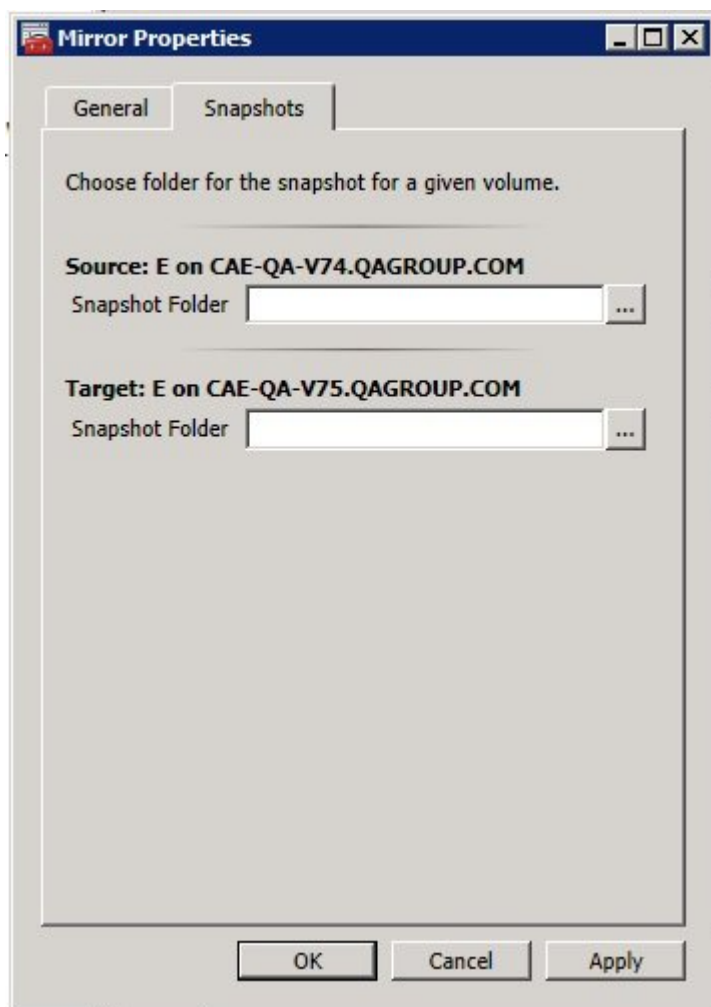
The size of the snapshot location should be determined on an individual basis based on several criteria. In practice, the size required for the snapshot file will be far less than the size of the volume being snapshot. The storage required needs to be big enough to contain any data that changes on the source system while the snapshot is being used. All snapshot files will be zeroed out each time a snapshot is initiated and will incrementally grow in size during use. The files will be deleted when the snapshot is dropped. Given that the copy on write process only writes "changed" blocks to the snapshot file, consideration should be given to the duration of the snapshot as well as the rate of

change in the volume being mirrored. Once an historical view can be established of snapshots from past activity, the size can be re-evaluated.

* **BEST PRACTICE:** Be conservative in your estimate, assuring that there is excess space available. If enough space is not allocated and the limit is reached, your snapshot will be dropped.


Snapshot Location Selection


1. Right-click on the appropriate mirror and select **Mirror Properties**.
2. From the **Mirror Properties** dialog, select the **Snapshots** tab.



***NOTE:** DataKeeper will use the snapshot location configured on the target

node; however, since either node in the mirror can become target, the snapshot location may be configured on both the source and the target.

3. Use the **browse**  button to choose the location for the snapshot or type the **path** into the text box.



When clicking the **browse** button that corresponds to the system where the GUI is running, the **Browse for Folder** dialog will appear. When clicking the **browse** button that corresponds to a system that is not the system where the GUI is running, the **Browse for Folder On Remote** dialog will appear.

4. Select your **snapshot location** for the source and the target. Make sure this volume has sufficient free space in order for the operation to complete successfully. Refer back to [Snapshot Location Size](#) for further details when estimating the volume size for your snapshot. Click **Apply**.

Note: Each volume on a given system can either use the same location or a different location can be selected.



In order to Bypass the GUI*, the location of the snapshot file can be set via command line using the [SETSNAPSHOTLOCATION](#) command. In order to view the current snapshot location of a given volume, use the [GETSNAPSHOTLOCATION](#) command

Taking a Snapshot

Once a **snapshot location** has been configured on the target system, a snapshot can be taken. From the target node, run the EMCMD command [TAKESNAPSHOT](#).

Dropping a Snapshot

When the snapshot is no longer needed, volume snapshots must be dropped in order to return to normal processing. Run the EMCMD command [DROPSNAPSHOT](#) which will lock the volume and clean up the snapshot files that were created. The volume will then return to a normal target where writes from the source will go directly to the volume with no copy-on-write storage.

Note: In Windows 2012R2, you will see the warning message "Disk # has been surprise removed."

Disabling Target Snapshot for a Given Volume

To disable target snapshot for a given volume, the snapshot location must be cleared. This can be accomplished via the GUI.

1. Right-click on the appropriate mirror and select **Mirror Properties**.
2. From the **Mirror Properties** dialog, select the **Snapshots** tab.
3. Remove the snapshot folder of the volume you would like target snapshot disabled on.
4. Click **Apply**.



The snapshot file location can also be cleared via command line by executing the [CLEARSNAPSHOTLOCATION](#) command.

Once successfully executed, a snapshot location will have to be reconfigured in order to initiate another snapshot of that volume.

Target Snapshot Notes

Supported Configurations

DataKeeper target snapshot is currently supported in non-shared (1×1 and 1×1×1) environments on all Windows OS versions supported by SPS.

Source Out of Service

DataKeeper target snapshot cannot be initiated when the source is out of service. However, if the source is taken out of service after snapshot is initiated, the snapshot will continue to work as expected. You can continue to use the snapshot, and drop it when you are done, while the source is out of service.

Switchovers and Failovers

If a snapshot is in progress, the volume being snapshotted cannot become the mirror source until the snapshot has been dropped. You must perform a [DROPSNAPSHOT](#) in order to allow a switchover or failover of the volume to the local node. Any processes that access data on the snapshotted volume will have their handles invalidated when the snapshot is dropped. However, if the volume is subsequently unlocked, you must make sure that those processes do not re-open their handles. At this point the data will be "**live**" application data and not the snapshotted data.

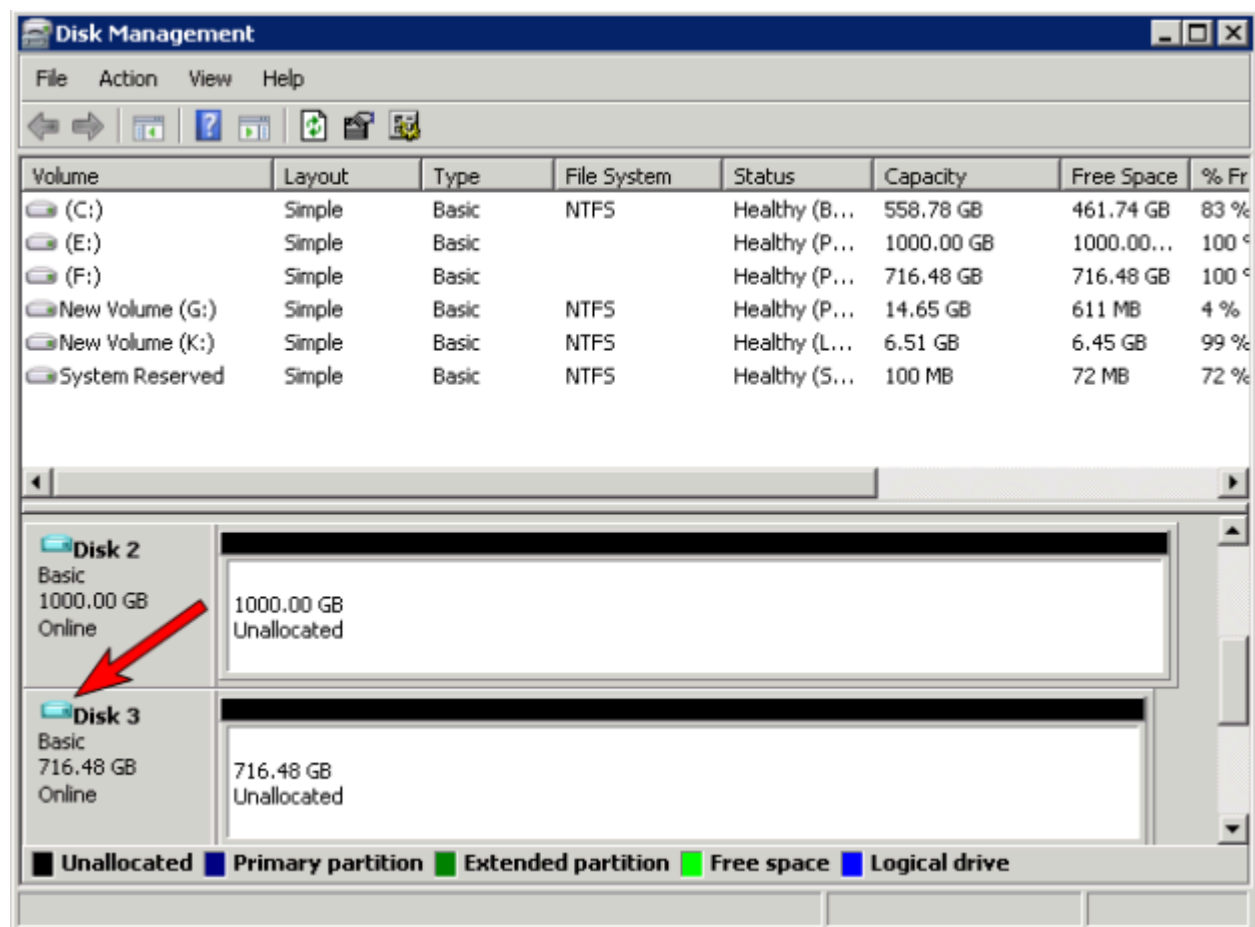
Note: To provide protection during SQL Server recovery we provide a generic script that needs to be added to stop the reporting SQL instance on the target node. Instructions are located in DKSnapshotCleanup.vbs script located in "\support". Please review the script code on how to add to your WSFC hierarchy.

Files / Disk Devices / Registry Entries

When a snapshot is taken, a snapshot file is created for each snapshotted volume in that volume's snapshot location. The name of the file that is created is datakeeper_snapshot_vol<X>.vhd, where <X> is the drive letter. This VHD file gets attached as a virtual disk device which can be seen in Windows Disk Management.


***NOTE:** The colored icon next to the disk number represents this disk as a

VHD.



! CAUTION: The virtual disk devices that are created will appear as unpartitioned Basic disks. They should be used for **snapshot data only** and should not be detached or partitioned while snapshots are in progress. Doing so may result in corruption of the snapshot data. **Make sure that they are not mistaken for available disks to be partitioned and formatted.**

Once these virtual disk devices are attached, a registry entry named SnapshotDevice is created in the volume's key. The value is set to \\.\PHYSICALDRIVE<x> where <x> is the disk number, as shown below:



TargetSnapshotBlocksize Registry Value

DataKeeper target snapshot uses a default block size of 64KB for all entries that are written to the snapshot file. This block size can be modified by creating a REG_DWORD value named [TargetSnapshotBlocksize](#) in the Volume registry key.

The value should always be set to a multiple of the disk sector size, which is usually 512 bytes. Certain workloads and write patterns can benefit from changing the block size. For example, a volume that is written in a sequential stream of data (e.g. *SQL Server log files*) can benefit from a larger block size. A large block size results in fewer reads from the target volume when consecutive blocks are written. But a volume that is written in a random pattern may benefit from a smaller value or the default 64KB. A smaller block size will result in less snapshot file usage for random write requests.

SQL Server Notes

If you are using DataKeeper target snapshot with SQL Server in a SIOS Protection Suite environment, it is recommended that you use a separate SQL Server instance to attach database(s) to the snapshot.

For a clustered SQL Server environment, you must use a separate SQL Server instance to attach database(s) to the snapshot.

Known Issues

SIOS VSS Provider Incompatible with some backup products

The SIOS VSS Provider component has been reported to cause backups to fail when using the following backup products:

- IBM Tivoli Storage Manager
- Microsoft Data Protection Manager

Microsoft .NET Framework 3.5 SP1 Requirement

The target snapshot feature requires Microsoft .NET Framework 3.5 SP1 to be installed - download from: <http://www.microsoft.com/net>.

NTFS File System Message

If an internal snapshot error occurs after target snapshot is initiated (such as the snapshot file running out of space or being detached by the user), snapshot will be disabled, the volume will be locked and snapshot files for any failed volumes will be deleted. While the snapshot error is being handled, you may receive NTFS file system errors. These messages are normal and can be ignored.

Application Data Using Snapshot

When using target snapshot data with your application, if the target snapshot is refreshed, you may need to close and reopen your application(s) to refresh the data.

Volume Shadow Copy Service (VSS) Free Disk Space Requirements

If your target snapshot volume has insufficient disk space, VSS operations involving that volume may fail with an "unexpected error". To avoid this, your snapshot volume should follow the guidelines from the Microsoft article [Troubleshoot VSS issues that occur with Windows Server Backup \(WBADMIN\) in Windows Server 2008 and Windows Server 2008 R2](#).

This article provides the following requirements for free disk space:

For volumes less than 500 megabytes, the minimum is 50 megabytes of free space. For volumes more than 500 megabytes, the minimum is 320 megabytes of free space. If the volume size is more than 1 gigabyte, a minimum of at least 1 gigabyte of free disk space on each volume is recommended.

Using SIOS DataKeeper Standard Edition To Provide Disaster Recovery For Hyper-V Virtual Machines

Considerations

When preparing a Hyper-V environment that spans subnets, subnetting may need to be taken into consideration for any applications that are running inside the virtual machine. Some applications "hard code" IP addresses into their configurations. When these types of applications are loaded in a virtual machine that is replicated (via a DataKeeper replicated volume) to a target server on a different subnet, they may not operate as expected due to the difference in the network settings.

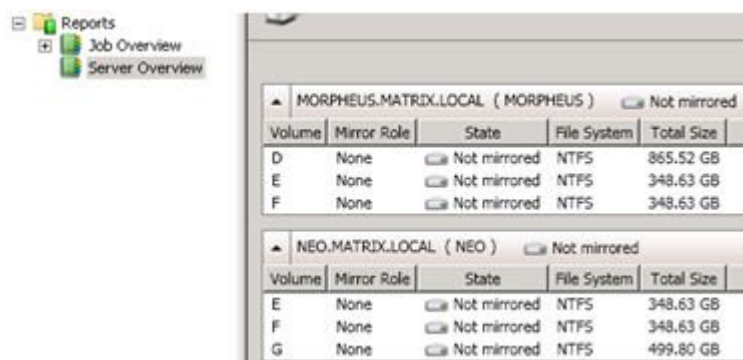
Preparing the Environment

1. Install Windows on two servers with at least two partitions, one for the OS and one for the Hyper-V virtual machine (VM) files. The partition for the files on the target server must be of equal or greater size to the source server's "data" partition. Install and configure the Hyper-V role on each server as described in Microsoft's "[Hyper-V Planning and Deployment Guide](#)" and the "[Hyper-V Getting Started Guide](#)", but wait to create your virtual machine until the DataKeeper replicated volume has been created.
2. Complete the installation requirements for the SIOS DataKeeper software.
3. [Connect to the Servers](#).

Once you connect, new options will appear in the center pane.

You can also optionally review the **Server Overview** report to see the status of your volumes.

When you connect to multiple servers that have DataKeeper installed and licensed, you will see multiple servers and volumes listed here.



MORPHEUS.MATRIX.LOCAL (MORPHEUS) <input type="checkbox"/> Not mirrored				
Volume	Mirror Role	State	File System	Total Size
D	None	<input type="checkbox"/> Not mirrored	NTFS	865.52 GB
E	None	<input type="checkbox"/> Not mirrored	NTFS	348.63 GB
F	None	<input type="checkbox"/> Not mirrored	NTFS	348.63 GB

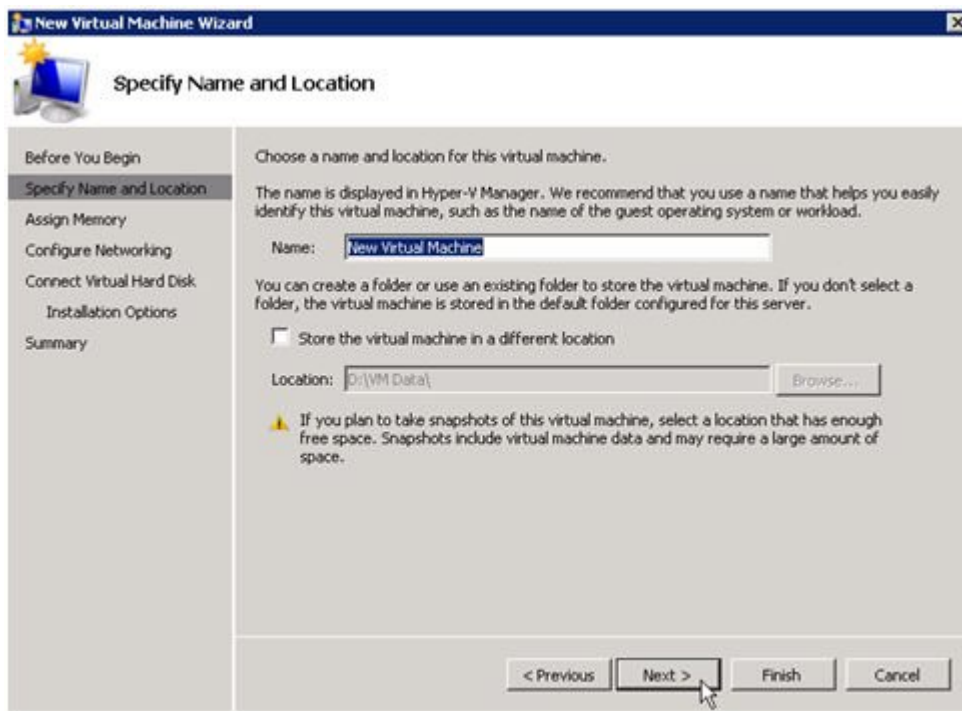
NEO.MATRIX.LOCAL (NEO) <input type="checkbox"/> Not mirrored				
Volume	Mirror Role	State	File System	Total Size
E	None	<input type="checkbox"/> Not mirrored	NTFS	348.63 GB
F	None	<input type="checkbox"/> Not mirrored	NTFS	348.63 GB
G	None	<input type="checkbox"/> Not mirrored	NTFS	499.80 GB

4. [Create a Job](#) / [Mirrored Volume](#).


Note: When you select your source server, ensure you select the server whose volume you want to replicate from. Reversing the source and target in these steps will completely overwrite your source volume with whatever is on the target server's volume, even if it is empty, causing you to lose any and all data stored on the source volume.

Create and Configure a Hyper-V Virtual Machine


1. Launch the **Hyper-V Console** from **Start - Administrative Tools - Hyper-V Manager**.
2. Start the **New Virtual Machine Wizard**.



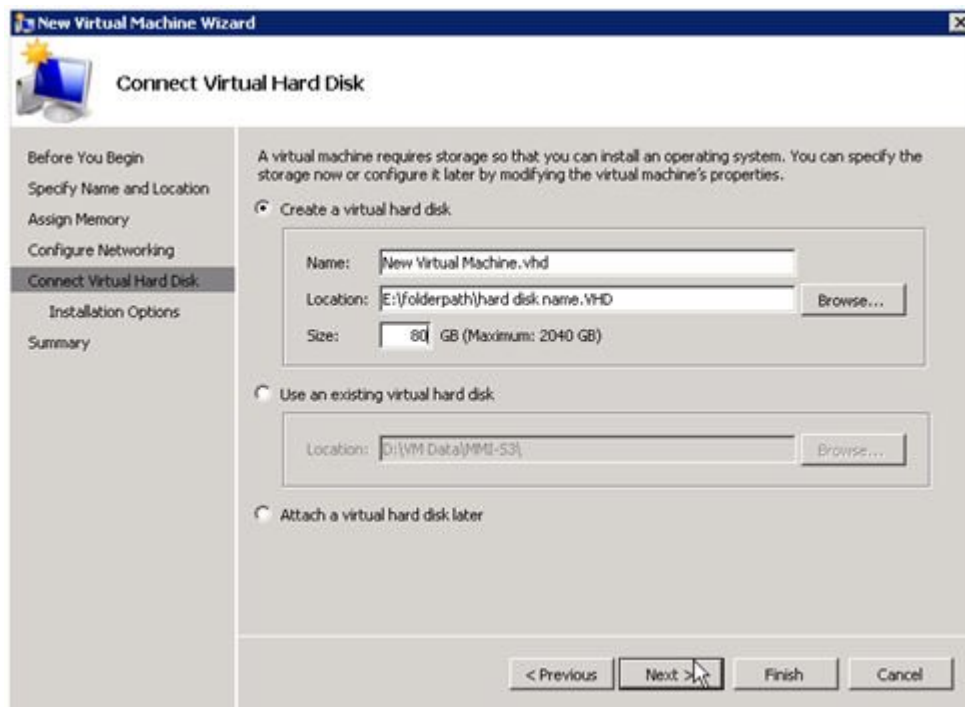
3. Specify the amount of **RAM** to use.




4. Select a **network adapter** to use.



5. Create a new **Virtual Hard Disk** on the replicated volume (or copy an existing VHD onto the replicated source volume and point the creation wizard at it to use as the virtual disk).



6. Specify the **operating system installation options**.



7. **Finish** the wizard and start the **virtual machine**.


Install an Operating System and Any Required Applications in the Virtual Machine

1. Load the operating system into the virtual machine as dictated by industry or vendor specified best practices.
2. Configure the networking within the virtual machine to use DHCP addresses. Use DHCP reservations and name resolution (DNS or WINS) records as well if necessary for address consistency for client connections.
3. Install any necessary applications in the virtual machine.

Configure the Target Server to Run the Virtual Machine

1. On the source Hyper-V host server, open **Hyper-V Manager**, connect to the virtual machine and do a full shutdown of the virtual machine. These actions will quiesce the data on the disk and will maintain data integrity on the target server.

2. Start the **DataKeeper console** as described previously.
3. Ensure the volume has been fully mirrored by checking the mirror status. The status must indicate **Mirroring** with the zero **KB Resync Remaining**.

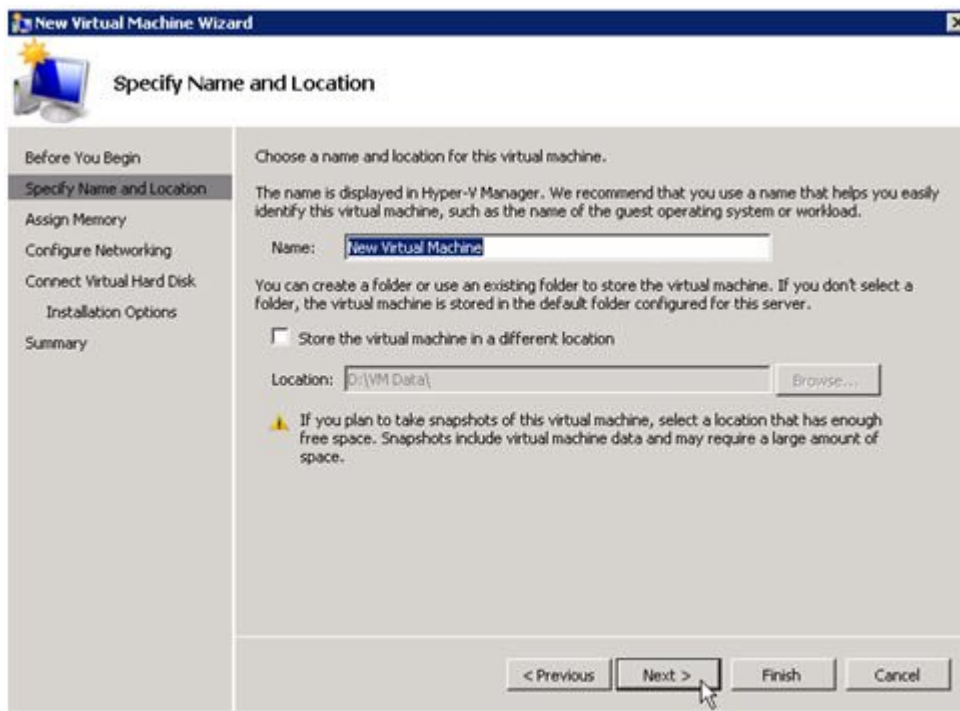


4. Select the mirror and click **Switchover** in the **Actions** pane.

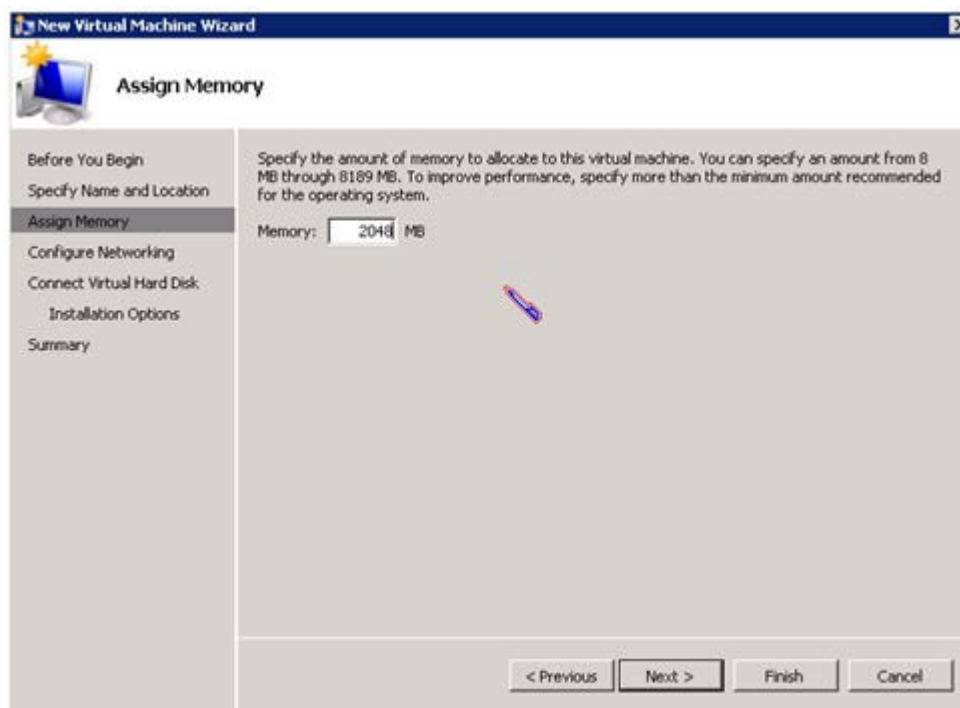


This will reverse the source and target and allow you to provision the virtual machine on the target server.


5. On the target server, start the **Hyper-V Manager**.
6. Start the **New Virtual Machine Wizard**.



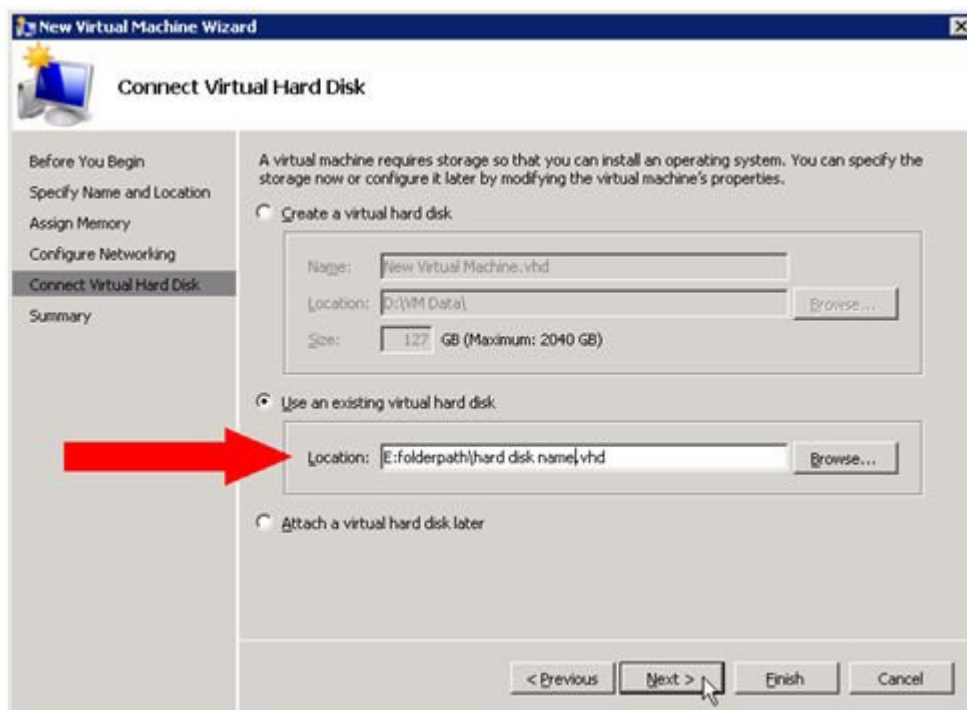
7. Specify the amount of **RAM** to use.




8. Select a **network adapter** to use.



***IMPORTANT:** Use the existing virtual hard disk on the replicated volume.



9. Click **Finish** to finalize the virtual machine creation process.



Start your virtual machine and test it to make sure it operates as expected.

Planned/Unplanned Switchover


Initiate a **Planned Switchover** to migrate the virtual machine back to your source server.

Initiating a switchover for testing or in the event of an actual outage on the primary server can be completed simply by doing a **Planned Switchover**. There are two types of switchovers, **planned** and **unplanned**.


Planned Switchover

A planned switchover is typically done in a maintenance window when the user community can be advised of planned downtime.

1. On the server on which the virtual machine is running, start **Hyper-V Manager**, as previously described, and connect to the **virtual machine**.



2. From inside the virtual machine, **Shut Down** the virtual machine.



3. On the same server, start the **DataKeeper console** as described previously.

Ensure the volume is in **mirroring** state by checking the **mirror status**. The status must indicate **Mirroring** with the **zero KB Resync Remaining** before switchover occurs.


State	Resync Remaining
Mirroring	0.00 KB

4. Select the mirror and click **Switchover** in the **Actions** pane.



Wait until the mirror has completely switched over and the DataKeeper user interface (UI) indicates the roles have been reversed properly.

5. Log into the **Hyper-V host server** that just became the source server in the DataKeeper interface.
6. Start **Hyper-V Manager** as described previously.
7. Start the virtual machine.



Unplanned Switchover

An unplanned switchover is necessary when a failure of some sort occurs and either the source system is unavailable or the connection between the systems is broken and requires that the virtual machine be brought online on the target server.

Since, in this scenario, the source server is unavailable for some reason, quiescing the data on the source server is not possible and as such, only the following steps are necessary on the target server to bring the virtual machine online.


1. On the target server, start the **DataKeeper console** as described previously.
2. Select the mirror and click **Switchover** in the **Actions** pane.



Wait until the mirror has completely come into service on the server and the DataKeeper user interface (UI) indicates the functional server is the source server.

3. On the same server, start **Hyper-V Manager** as described previously.

Start the virtual machine.



Switchback

Switchback is a planned event which transfers the virtual machine from the target server back to the source server and, in process, is exactly the same as the planned switchover process. Please refer to the steps

previously listed in the [Planned Switchover](#) section to affect a switchback.

FAQs

Refer to this section for answers to the most frequently asked questions about SIOS DataKeeper.

[Awareness of Windows Filenames and Directory Names](#)

[AWS Issues and Workarounds](#)

[Change Mirror Endpoints](#)

[Change Mirror Type](#)

[Create a Mirror and Rename Job and Delete Job Actions Grayed Out](#)

[Data Transfer Network Protocols](#)

[Delete and Switchover Actions Grayed Out](#)

[Deleting a Mirror](#)

[Error Messages Log](#)

[Inability to Create a Mirror](#)

[Network Disconnect](#)

[Reclaim Full Capacity of Target Drive](#)

[Resize or Grow Mirrored Volumes](#)

[Split-Brain FAQs](#)

[Stop Replication Between Source and Target](#)

[Using Volume Shadow Copy](#)

[Volumes Unavailable for Mirroring](#)

Awareness of Windows Filenames and Directory Names

Question

Is SIOS DataKeeper aware of Windows filenames and directory names?

Answer

SIOS DataKeeper is implemented with a Windows kernel mode filter driver that sits above the physical disk driver but below the file system. As a result, the SIOS DataKeeper driver knows nothing about individual files or the file system itself. It is only aware of raw writes to the disk.

AWS Issues and Workarounds

Question

What is the best practice for shutting down clustered VMs in AWS?

Answer

If shutting down the primary source node, all cluster roles depending on a SIOS DataKeeper Volume Resource should be placed in the Offline state. Also make sure all mirrors are in the mirroring state prior to shutting down any VM. The node shutdown order does not matter as long as the previous steps have been taken.

Change Mirror Endpoints

Question

Can I change the mirror endpoints (IP address) of a system currently associated with an existing mirror?

Answer

Yes. The EMCMD called [CHANGEMIRRORENDPOINTS](#) allows you to change the endpoints of a mirrored volume that is configured on 3 nodes or fewer. (If your configuration consists of more than three nodes, the mirrors must be deleted and recreated.)

Change Mirror Type

Question

Can you change the mirror type of an existing mirror from Synchronous to Asynchronous or vice-versa?

Answer

Yes, you can change the mirror type using the EMCMD [CHANGEMIRRORTYPE](#) command.

Create a Mirror and Rename Job and Delete Job Actions Grayed Out

Question

Why are the Create a Mirror, Rename Job and Delete Job actions grayed out?

Answer

If a node that is part of the job is down, these actions will not be enabled.

Data Transfer Network Protocols

Question

What are the network protocols used for SIOS DataKeeper Data Transfer?

Answer

SIOS DataKeeper uses named pipe communication and TCP Sockets.

Delete and Switchover Actions Grayed Out

Question

Why are the Delete and Switchover actions grayed out on the DataKeeper User Interface?

Answer

If the volume is under clustering protection (Microsoft clustering or SIOS LifeKeeper clustering), these actions are disabled.

Deleting a Mirror

Question

What actually happens when you delete a mirror?

Answer

The data remains on both sides, but the target and source data are no longer synchronized. The target volume is unlocked and made fully accessible.

Error Messages Log

Question

Where does DataKeeper log error messages?

Answer

DataKeeper events are logged in the **Windows Application Event Log** and the **Windows System Event Log**. Here is a breakdown of the messages you can look for.

Application Event Log:

- Source = ExtMirrSvc - events related to the DataKeeper service.
- Source = DataKeeperVolume - events related to DataKeeper Volume Resources defined in Windows Failover Clustering (WSFC).
- Source = SIOS.SDRSnapIn - events related to the DataKeeper GUI connecting to the DataKeeper systems.

System Event Log:

- Source = ExtMirr - events directly related to mirror creation, mirror manipulation and replication.

*Note: The **System Event Log** should always be set to "**Overwrite events as needed**". If the System Event Log fills up or becomes corrupted, it will prevent DataKeeper from properly recognizing mirror state changes.

Inability to Create a Mirror

Question

Why can't I create a mirror?

Answer

- The common cause of this problem is that the volume on either source or target is in use by another process. Stop the process that is accessing the volume and try again. The SIOS DataKeeper software requires exclusive access to the target volume during the creation of the mirror.
- The target volume must be as large, or larger, than the source volume. It is recommended that the user compare the target volume size with the source volume size using the Disk Management utility. If the sizes are not the same, recreate the target partition a little larger. See [Volume Considerations](#) for more information.
- An error experienced during [Create Mirror](#) could indicate that the target volume is corrupt. If this occurs, format the target volume and attempt to create the mirror again.

!WARNING: Drive letters on the target and source must match when using Windows Server Failover Clustering.

Network Disconnect

Scenario #1

In a 2-Node, non-clustering configuration (1×1) replicating a 100TB volume between Source server and Target server over a WAN connection, the network goes down for twenty minutes.

Question

In this scenario, what would happen to the **Mirror State** with DataKeeper Standard Edition?

Answer

After a couple of minutes, the Source server will detect that the network is down and the mirror will go from the **MIRRORING** state to the **PAUSED** state.

Question

Does DataKeeper continue to track changes on the Source server?

Answer

Yes. The Bitmap (# of Dirty Sectors) will continue to be updated on the Source server while the mirror is in the **PAUSED** state.

Question

Once the network is resumed, will a partial resync to the Target server occur?

Answer

Yes. The mirror will go to the **RESYNC** state and remain there until all dirty sectors are written to the Target server. It will be a partial resync.

Scenario #2

In a 2-Node, non-clustering configuration (1×1) replicating a 100TB volume between Source and Target over a WAN connection, the network goes down for twelve hours. The Source server is rebooted while the network is down.

Question

In this scenario, what would happen to the status of the Source server in DataKeeper Standard Edition?

Answer

The Bitmap on the Source server is persistent (on disk), so it will not be affected by a Source reboot. Only a partial resync is needed if the Source server is rebooted. The Target server will report that it is in the **MIRRORING** state until it is reconnected to the Source server. Then it will go to the **RESYNC** state while the resync is proceeding.

Reclaim Full Capacity of Target Drive

Question

How do I reclaim the full capacity of my target drive when I no longer need it for mirroring?

Answer

The file system on the target drive is overlaid by SIOS DataKeeper, thereby making it smaller than the actual partition size. Although Disk Management indicates the full partition size, SIOS DataKeeper and Windows Explorer indicate the smaller mirror size. To reclaim full capacity of the drive, reformat the partition or use a partition resizing utility such as GParted (<http://gparted.sourceforge.net/>).

Resize or Grow Mirrored Volumes

Question

Can you resize or grow mirrored volumes?

Answer

Yes, beginning with Version 7.4, users can extend and shrink their DataKeeper volumes dynamically while retaining mirror settings. See [DataKeeper Volume Resize](#) for more information.

Split-Brain FAQs

Scenario

I am using DataKeeper in a non-cluster environment. I am mirroring from Server1 at one site to Server2 at a second site. Communication is broken due to site-to-site VPN, and I need to fail over from Server1 to Server2. I cannot access Server1 from anywhere. Server1 is actually still on but not reachable internally or externally, and there may be some processes still running in the backend.

Question

How can I fail over from Server1 to Server2?

Answer

Using the [SWITCHOVERVOLUME](#) command or the **Switchover Mirror** option in the DataKeeper UI, switch the source of the mirror to Server2. There will be a delay while the Target tries to connect to the Source, but that should complete in 30-40 seconds or so.

Question

During the switchover period, both Server1 and Server2 are writing new data to the disk (Volume F on both Server1 and Server2). When the connection comes back online, will Server1 automatically become the Target?

Answer

No. This scenario will cause a [split-brain](#) condition. Perform one of the following to resolve this issue:

- Using the DataKeeper User Interface, perform the [Split-Brain Recovery Procedure](#).

or

- Run the EMCMD [PREPARETOBECOMETARGET](#) command on the system that is going to become the Target, and then run the [CONTINUEMIRROR](#) command on the system that is going to become the Source.

Question

Which of the two methods above do you recommend for resolving the split-brain issue?

Answer

Whichever you prefer - they both perform the same functions.

Question

Can the command for the Target server be run from the Source server?

Answer

Yes, the command for the Target server can be run from the Source server.

Question

How does DataKeeper sync the changed and unchanged blocks?

Answer

When resolving a split-brain condition, any changes on the system that is becoming the Target will be overwritten and lost. If there are changes on that system that you want to retain, manually copy those changes to the system that is going to become the Source.

Question

When running the [PREPARETOBECOMETARGET](#) command to resolve a split-brain condition, will a full resync or partial resync occur from the Source?

Answer

The **PREPARETOBECOMETARGET** command will delete the mirror(s) on that system but will leave the volume locked. The bitmap will remain intact so that a partial resync can be performed in the next step ([CONTINUEMIRROR](#)).

Question

How can I simulate a split-brain scenario?

Answer

To simulate a split-brain scenario, unplug the network between two systems so they cannot communicate. Run the [SWITCHOVERVOLUME](#) command (or select the **Switchover Mirror** option in the DataKeeper UI) on the Target so they both become Source, then reconnect the network. You are in a split-brain condition at that point.

Question

Should I wait for the **PREPARETOBECOMETARGET** command to complete before running **CONTINUEMIRROR** on the Source?

Answer

The **PREPARETOBECOMETARGET** command completes immediately.

Stop Replication Between Source and Target

Question

How do I stop the replication between the Source and Target volumes?

Answer

Replication occurs at the driver level and can only be stopped or interrupted by sending a command from the DataKeeper GUI or the DataKeeper command line (EMCMD) to the DataKeeper driver to do one of the following:

- [PAUSE the mirror](#) - Mirror endpoints still exist, but all replication is suspended. Writes are tracked on the source system so only a partial resync of the data is necessary to bring the target volume back into sync when the mirror is CONTINUED.
- [BREAK the mirror](#) - Mirror endpoints still exist, but all replication is suspended. Writes to the source system are not tracked. RESYNCING the mirror will initiate a full resync of the data which is required to bring the target volume back into sync with the source.
- [DELETE the mirror](#) - Mirror endpoints are deleted and replication stops.

***Note:** Stopping the DataKeeper service does not stop replication.

Using Volume Shadow Copy

Question

Can Volume Shadow Copy (VSS) be Used with DataKeeper Volumes?

Answer

VSS Shadow Copy can be enabled for DataKeeper volumes. However, the following guidelines apply:

- VSS snapshot images must not be stored on a DataKeeper volume. Storing VSS snapshots on a DataKeeper volume will prevent DataKeeper from being able to lock the volume and switch it over to another node.
- When a DataKeeper volume is switched or failed over, any previous snapshots that were taken of the DataKeeper volume are discarded and cannot be reused.
- VSS snapshot scheduling is not copied between the DataKeeper servers. If snapshots are scheduled to be taken twice a day on the primary server and a switchover occurs, this schedule will not be present on the backup server and will need to be redefined on the backup server.
- When switching back to a server where snapshots were previously enabled, VSS snapshots are automatically re-enabled; HOWEVER, any previous snapshots that were taken of the DataKeeper volume are discarded and cannot be reused.

Volumes Unavailable for Mirroring

Question

Why are some of my volumes not available for mirroring?

Answer

The SIOS DataKeeper service filters out the following types of disk partitions:

- Windows system volume
- Volume(s) that contain the Windows pagefile
- Non-NTFS formatted volumes (e.g. FAT, Raw FS)
- Non-fixed drive types (e.g. CD-ROMs, diskettes)
- Target volumes that are smaller than the source volume

DataKeeper Troubleshooting

The topics in this section contain important information about known issues and restrictions offering possible workarounds and/or solutions.

[Known Issues and Workarounds](#)

[Access to Designated Volume Denied](#)

[DataKeeper Volume cannot come Online after Network failure with clustered IP Address on Replication network](#)

[DataKeeper Volume Not Available as Cluster Resource Type](#)

[Failed to Create Mirror](#)

[Hyper-V Host Cluster Error](#)

[Live Migration Failure](#)

[MaxResyncPasses Value](#)

[Mirroring with Dynamic Disks](#)

[New Resources Offline But Unlocked](#)

[Server Login Accounts and Passwords Must Be Same on Each Server in the Cluster](#)

[System Event Log - Create Mirror Failed in the GUI](#)

[Unable to Determine Previous Install Path](#)

[User Interface - Failed to Create Mirror](#)

[User Interface - Shows Only One Side of the Mirror](#)

[WSFC - MS DTC Resource Failure](#)

[WSFC 2008 R2 SP1 Procedure Change](#)

[Windows Server 2012 Specific Issues](#)

[Windows Server 2012 MMC Snap-in Crash](#)

[Windows Server 2012 – Simultaneous Move of Multiple Clustered File Server Roles Can Result in DataKeeper Switchover Failures](#)

[Windows Server 2012 iSCSI Target Role Does Not Support Dynamic Disks](#)

[Windows Server 2012 Default Information Missing During Mirror Creation](#)

[Windows Server 2012 NIC Teaming Issue](#)

[WSFC 2012 Cluster Creation Default Setting Issue](#)

[WSFC 2012 Failover Cluster Manager UI Defect](#)

[WSFC 2012 File Server Resource Manager Event Log Errors](#)

[WSFC 2012 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager](#)

[WSFC 2012 New File Server Type Not Supported](#)

[WSFC 2012 Server Manager – Incorrect Volume Display](#)

[WSFC 2012 Server Manager – DataKeeper “Disk” Not Shown as Clustered](#)

[Windows 2012 File Share](#)

[Windows Server 2016 Specific Issues](#)

[Occasional Job Creation Failure](#)

[Restrictions](#)

[Bitlocker Does Not Support DataKeeper](#)

[CHANGEMIRRORENDPOINTS Restriction](#)

[CHKDSK](#)

[DataKeeper Volume Resize Restriction](#)

[Directory for Bitmap Must Be Created Prior to Relocation](#)

[Duplicate IP Addresses Disallowed Within a Job](#)

[Intensive I-O with Synchronous Replication](#)

[Resource Tag Name Restrictions](#)

Known Issues and Workarounds

Included below are known issues open against DataKeeper and DataKeeper Cluster Edition as well as possible workarounds and/or solutions.

[Access to Designated Volume Denied](#)

[DataKeeper Volume cannot come Online after Network failure with clustered IP Address on Replication network](#)

[DataKeeper Volume Not Available as Cluster Resource Type](#)

[Failed to Create Mirror](#)

[Hyper-V Host Cluster Error](#)

[Live Migration Failure](#)

[MaxResyncPasses Value](#)

[Mirroring with Dynamic Disks](#)

[New Resources Offline But Unlocked](#)

[Server Login Accounts and Passwords Must Be Same on Each Server in the Cluster](#)

[System Event Log - Create Mirror Failed in the GUI](#)

[Unable to Determine Previous Install Path](#)

[User Interface - Failed to Create Mirror](#)

[User Interface - Shows Only One Side of the Mirror](#)

[WSFC - MS DTC Resource Failure](#)

[WSFC 2008 R2 SP1 Procedure Change](#)

[Windows Server 2012 Specific Issues](#)

[Windows Server 2012 MMC Snap-in Crash](#)

[Windows Server 2012 – Simultaneous Move of Multiple Clustered File Server Roles Can Result in DataKeeper Switchover Failures](#)

[Windows Server 2012 iSCSI Target Role Does Not Support Dynamic Disks](#)

[Windows Server 2012 Default Information Missing During Mirror Creation](#)

[Windows Server 2012 NIC Teaming Issue](#)

[WSFC 2012 Cluster Creation Default Setting Issue](#)

[WSFC 2012 Failover Cluster Manager UI Defect](#)

[WSFC 2012 File Server Resource Manager Event Log Errors](#)

[WSFC 2012 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager](#)

[WSFC 2012 New File Server Type Not Supported](#)

[WSFC 2012 Server Manager – Incorrect Volume Display](#)

[WSFC 2012 Server Manager – DataKeeper “Disk” Not Shown as Clustered](#)

[Windows 2012 File Share](#)

[Windows Server 2016 Specific Issues](#)

[Occasional Job Creation Failure](#)

[WSFC 2016 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager](#)

Access to Designated Volume Denied

If access to the designated volume is denied, then check whether you are attempting to create the mirror while other applications are accessing the volume. During Mirror Creation, the volumes must be locked on the target system for exclusive access by the SIOS DataKeeper software.

In particular, the Distributed Tracking Client service, which is set to run by default in Windows, keeps two file handles open for each volume. If the volume houses a SIOS DataKeeper target, the SIOS DataKeeper driver cannot lock the volume. You must therefore stop the Distributed Tracking Client service and set its startup policy to Manual.

DataKeeper Volume Not Available as Cluster Resource Type

WSFC Server - The DataKeeper Volume is Not Available as a Cluster Resource Type After DataKeeper is Installed in a Microsoft WSFC Environment

Error/Message

The DataKeeper Volume is not available as a cluster resource type after DataKeeper is installed in a Microsoft WSFC environment.

The Event Log will include the following message: **"Failed to register the 'DataKeeper Volume' Resource DLL (DataKeeperVolume.dll). Error: 70"**

Description

Resource DLL registration requires that all cluster nodes are up and online. In the case where one node of an existing cluster is currently unavailable (offline, cluster service stopped, etc.), automatic DataKeeper Resource DLL registration may fail during installation/update.

Suggested Action

The problem is normally corrected automatically when the other cluster node goes online. As soon as the DataKeeper service is started there, Resource DLL registration will be attempted from that node and registration will occur cluster-wide. In the event that automatic Resource DLL registration does not occur, restart the DataKeeper service on any node after all cluster nodes are up and online. The registration process begins 60 seconds after the DataKeeper service starts.

Failed to Create Mirror

User Interface - Failed to Create Mirror - Application Event Log

Error/Message

Logged in the **Application Event Log**:

File: .\GuiThread.cpp Line: 3099 Attempt to connect to remote system REMOTESERVER failed with error 5. Please ensure that the local security policy for "**Network Access: Let Everyone permissions apply to anonymous users**" is enabled on all the servers running DataKeeper.

Check: Local security policy setting on the specified system.

Description

Failed to create the mirror. Mirror is created but not stored in the job.

Suggested Action

Make local security policy change, open command prompt and run "%EXTMIRRBASE %\emcmd. deletemirror <volume>", then perform the mirror creation action again.

Hyper-V Host Cluster Error

Failover Cluster Error After Changing Virtual Machine Configuration While VM Is Clustered

Description

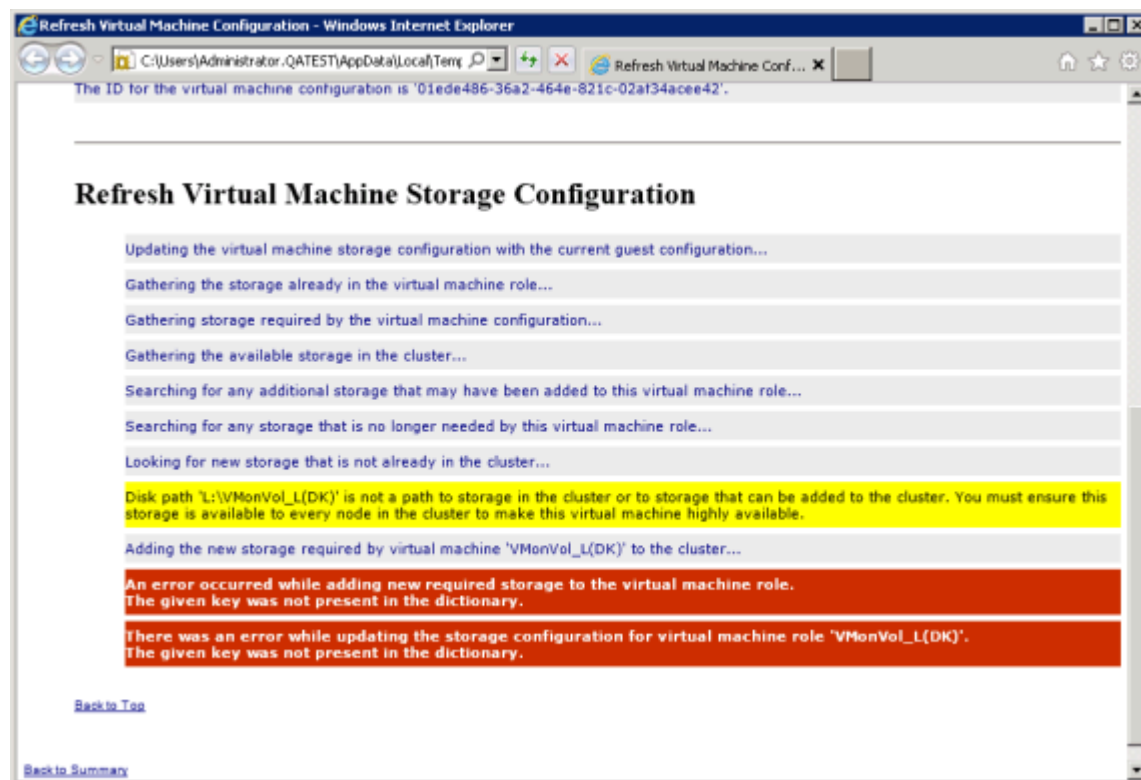
If Failover Cluster Manager is used to modify the VM configuration while the VM is clustered such as adding a Network Interface to the VM, **"Refresh Virtual Machine Storage Configuration"** errors may be generated and the VM will fail Quick Migration and/or Live Migration to another cluster node.

This problem occurs only when the following criteria are met:

1. The VM is in the cluster
2. Failover Cluster Manager is used to change the VM network configuration
3. Storage other than Cluster Shared Disk is used for VM storage, such as DataKeeper Volume replicated storage

All three criteria must be met for this error to occur. This error does not occur if Hyper-V Manager is used to change VM network configurations when the VM is out of the cluster.

Here is what to look for:




Suggested Action

Contact Microsoft for information on a solution.

To make Virtual Machine network adapter changes without installing the Microsoft Hotfix, perform the following:

1. Take the VM out of the cluster
2. Verify that Virtual Network Names for NIC connections are identical on all cluster nodes
3. Use Hyper-V Manager to make virtual machine network configuration changes
4. Return the VM to the cluster and bring the DataKeeper Volume resource into the VM resource hierarchy
5. Re-create the Virtual Machine Configuration resource to DataKeeper Volume resource dependency (shown in yellow below)



Live Migration Failure

Live Migration Fails if Virtual Network Names Differ

Description

When attempting to perform a live migration of a virtual machine(s) to another node in **Failover Cluster Manager**, if **Virtual Network Names** for NIC connections are not identical on all cluster nodes, the migration issues a "failed" status.

Suggested Action

Make sure that **Virtual Network Names** for NIC connections are identical on all cluster nodes.

MaxResyncPasses Value

If, during a volume resynchronization, the number of passes made through the intent log exceeds the **MaxResyncPasses** registry value (200 by default), SIOS DataKeeper logs a message to the **Event Log** indicating that the resync process is taking too many passes and requests that the administrator stop whatever process is writing to the drive being resynchronized. The mirror then goes to the **Paused** state. You can increase the **MaxResyncPasses** value from the registry to give the resync process more time.

Mirroring with Dynamic Disks

When changing from a **Basic Disk** to a **Dynamic Disk**, the underlying volume GUID may be changed by the OS upon reboot. This will cause a DataKeeper mirror to break.

Suggested Action

When mirroring with dynamic disks, your **dynamic** volumes should be created and a reboot should be performed PRIOR to creating your mirror. If the mirror has already been created, it must be deleted prior to creating your dynamic volumes.

New Resources Offline But Unlocked

WSFC Server – Newly Created Resources Appear Offline But Are Unlocked

Error/Message

Newly created resources appear offline but are unlocked.

Description

The new resource is always offline and unlocked before it is used.

Suggested Action

Switch the resource to online.

Server Login Accounts and Passwords Must Be Same on Each Server in the Cluster

The DataKeeper GUI cannot connect to the target server in a cluster if server **Login Accounts** and **Passwords** are different on each server.

Error Message

An Error Code 1326 will appear in the Application log (**Note:** The Error Code may also be a 2 with Event ID 0):

```
SteelEye.Dialogs.AddServerWindow: Failed to connect to server:
172.17.105.112 System.ApplicationException: Failed to open a
connection to 172.17.105.112 (error_code = 1326) at
SteelEye.DAO.Impl.DataReplication.ClientLibrarySDRService.ThrowIfNonZero(U
errorCode, String message) at
SteelEye.DAO.Impl.DataReplication.ClientLibrarySDRService.GetServiceInfo(S
serverName) at
SteelEye.DAO.Impl.DataReplication.CachingSDRService.<>c__DisplayClass2.b__
at SteelEye.DAO.Impl.DataReplication.Cacher`1.fetch(String typekey,
String datakey, Fetcher fetcher) at
SteelEye.DAO.Impl.DataReplication.CachingSDRService.GetServiceInfo(String
serverName) at
SteelEye.DataKeeper.SDR.SDRDataKeeperService.ConnectToServer(String
serverName) at
SteelEye.Dialogs.AddServerWindow.<>c__DisplayClass4.b__0(Object s,
DoWorkEventArgs e) at
System.ComponentModel.BackgroundWorker.WorkerThreadStart(Object
argument)
```

net helpmsg 1326 shows:

Logon failure: unknown user name or bad password

Description/Cause

The Service Account User Names and Passwords being used to start DataKeeper are the same on both servers and the firewalls are disabled on the servers; however, the Passwords used to log in to the servers

themselves are different.

Suggested Action

The DataKeeper GUI uses the server Login ID and Password; therefore, the User Name and Password used to log in to the servers themselves must be the same on each server and must have administrator privileges.

System Event Log – Create Mirror Failed in the GUI

Error/Message

Create Mirror Failed in the GUI.

Description

This can result if a vmms.exe program is holding on to volume and preventing SIOS DataKeeper from locking it.

Unable to Determine Previous Install Path

Installation – Fatal Error: Unable to Determine Previous Install Path

Error/Message

Fatal Error: Unable to determine previous install path. DataKeeper cannot be uninstalled or reinstalled.

Description

When performing a **"Repair"** or **"Uninstall"** of DataKeeper, the **"ExtMirrBase"** value is missing in the installation path of DataKeeper in the registry under ***HKLM\System\CurrentControlSet\Control\Session Manager\Environment***.

Suggested Action

Perform one of the following:

- Under the **Environment** key, create **"ExtMirrBase"** as a REG_SZ and set the value to the DataKeeper installation path (i.e. ***C:\Program Files(x86)\SIOS\DataKeeper***).
- To force InstallShield to perform a new install of DataKeeper, delete the following registry key:

```
HKLM\Software\Wow6432Node\Microsoft\Windows\
CurrentVersion\Uninstall\
{B00365F8-E4E0-11D5-8323-0050DA240D61}.
```

This should be the installation key created by InstallShield for the DataKeeper product.

User Interface - Failed to Create Mirror

User Interface - Failed to Create Mirror, Event ID 137

Error/Message

Failed to create the mirror.

Event Id: 137

System Event Log

Unable to initialize mirror on the target machine.

Volume Device:

Source Volume: E

Target Machine: 10.17.103.135

Target Volume: E

Failed operation: Target reports error

Error Code: 0xC0000055

Description

DataKeeper cannot lock the Target volume during mirror creation.

Suggested Action

1. Verify the Distributed Link Tracking Client service is not running on either system.
2. Stop any other processes that may prevent DataKeeper from locking the Target volume (e.g. anti-virus software).
3. Recreate the mirror.

User Interface – Shows Only One Side of the Mirror

If the SIOS DataKeeper UI shows a volume as a source and its corresponding target as available or a volume as a target with the corresponding source volume as available, you can use the command line utility to force an update to the SIOS DataKeeper GUI or delete the orphaned side of the mirror. From a command prompt, go to the SIOS DataKeeper directory on the server which is displaying unexpected mirror status and perform the following steps:

1. Make sure that the mirror is not in a **Paused** or **Broken** state on the source. If so, continue the mirror on the source. This should result in the mirror being re-established to the target.
2. Run EMCMD <system name> UpdateVolumeInfo <volume letter>

Where

<system name> is the name of the system;

<volume letter> is the letter of the volume.

3. If the problem is not resolved in Step 1, then stop and restart the SIOS DataKeeper service.

Windows Server 2012 Specific Issues

For issues related to **Windows Server 2012**, see the following topics:

[Windows Server 2012 MMC Snap-in Crash](#)

[Windows Server 2012 DataKeeper Switchover Failures](#)

[Windows Server 2012 iSCSI Target Role Does Not Support Dynamic Disks](#)

[Windows Server 2012 Default Information Missing During Mirror Creation](#)

[Windows Server 2012 NIC Teaming Issue](#)

[WSFC 2012 Cluster Creation Default Setting Issue](#)

[WSFC 2012 Failover Cluster Manager UI Defect \(Delete Action Missing\)](#)

[WSFC 2012 File Server Resource Manager Event Log Errors](#)

[WSFC 2012 File Shares Cannot be Created for File Server Resource](#)

[WSFC 2012 New File Server Type Not Supported](#)

[WSFC 2012 Server Manager – Incorrect Volume Display](#)

[WSFC 2012 Server Manager – DataKeeper “Disk” Not Shown as Clustered](#)

[WSFC 2012 File Share](#)

Windows Server 2012 MMC Snap-in Crash

Description

When using the DataKeeper user interface (MMC Snap-in) on Windows Server 2012, the mmc.exe process may crash unexpectedly due to an internal .Net or Windows Presentation Foundation (WPF) issue. The error may show up on the screen and/or the event viewer.

Suggested Action

This crash does not affect the server(s) to which the snap-in was connected or any DataKeeper mirrors established at the time of the crash. The MMC Snap-in may be safely relaunched. Simply close the UI and restart it.

The following are examples of **Application Event Log messages** that may be logged during this failure.

Log Name: Application
Source: Desktop Window Manager
Date: 11/28/2012 8:34:00 AM
Event ID: 9009
Task Category: None
Level: Information
Keywords: Classic
User: N/A
Computer: CAE-QA-V96.QAGROUP.COM
Description:
The Desktop Window Manager has exited with code (0xd00002fe)

Log Name: Application
Source: .NET Runtime
Date: 11/28/2012 8:34:00 AM
Event ID: 1026
Task Category: None
Level: Error


Keywords: Classic
User: N/A
Computer: CAE-QA-V96.QAGROUP.COM
Description:
Application: mmc.exe
Framework Version: v4.0.30319
Description: The process was terminated due to an unhandled exception.

Log Name: Application
Source: Application Error
Date: 11/28/2012 8:34:00 AM
Event ID: 1000
Task Category: (100)
Level: Error
Keywords: Classic
User: N/A
Computer: CAE-QA-V96.QAGROUP.COM
Description:
Faulting application name: mmc.exe, version: 6.2.9200.16384, time stamp:
0x50109efd
Faulting module name: KERNELBASE.dll, version: 6.2.9200.16384, time
stamp: 0x5010ab2d
Exception code: 0xe0434352
Fault offset: 0x000000000000189cc
Faulting process id: 0xdc4
Faulting application start time: 0x01cdccd27c68a1c6
Faulting application path: C:\Windows\system32\mmc.exe
Faulting module path: C:\Windows\system32\KERNELBASE.dll
Report Id: 443c3ed3-3960-11e2-9400-0050569b131b
Faulting package full name:
Faulting package-relative application ID:

Windows Server 2012 iSCSI Target Role Does Not Support Dynamic Disks

Description


The iSCSI Target Role only supports DataKeeper Volumes that are mirrors of **Simple Volumes** placed on **Basic Disks**. If any of your mirrors are using volumes such as Striped or Spanned volumes on a Dynamic Disk on either the source or target system, then you cannot create an iSCSI Target role that uses those DataKeeper Volume resources for storage.



Windows Server 2012 NIC Teaming Issue

If you use the **NIC Teaming** feature of Windows Server 2012, Windows 2012 will report only one adapter MAC address for the license. If you have many underlying adapters, the MAC address will arbitrarily change and Windows may pick one of the adapters that may no longer be licensed.

To resolve this issue, configure the **MAC address** property of the virtual team adapter. This property can be changed using the **Advanced** tab of the **Adapter Properties** as shown in the diagram:



Windows Server 2012 – Simultaneous Move of Multiple Clustered File Server Roles Can Result in DataKeeper Switchover Failures

Description

If more than one File Server role is created in Failover Clustering, each of which is using one or more DataKeeper Volume resources for storage, errors can occur if two or more roles are manually moved from one node to another simultaneously. In some cases, one or more DataKeeper Volume resources can fail to come online.

It is also possible that an error message is logged but the switchover works successfully; in that case, the message logged will be Event ID 196:

```
Attempt to connect to remote system failed with error 64. Please ensure that the local security policy for "Network Access: Let Everyone permissions apply to anonymous users" is enabled on all the servers running DataKeeper.
```

In this case, this event message can be ignored.


Suggested Action

If more than one File Server needs to be manually moved to another node, each one should be moved independently. Make sure that the File Server has completely come online before attempting to move any other File Servers.

Windows Server 2012 Default Information Missing During Mirror Creation

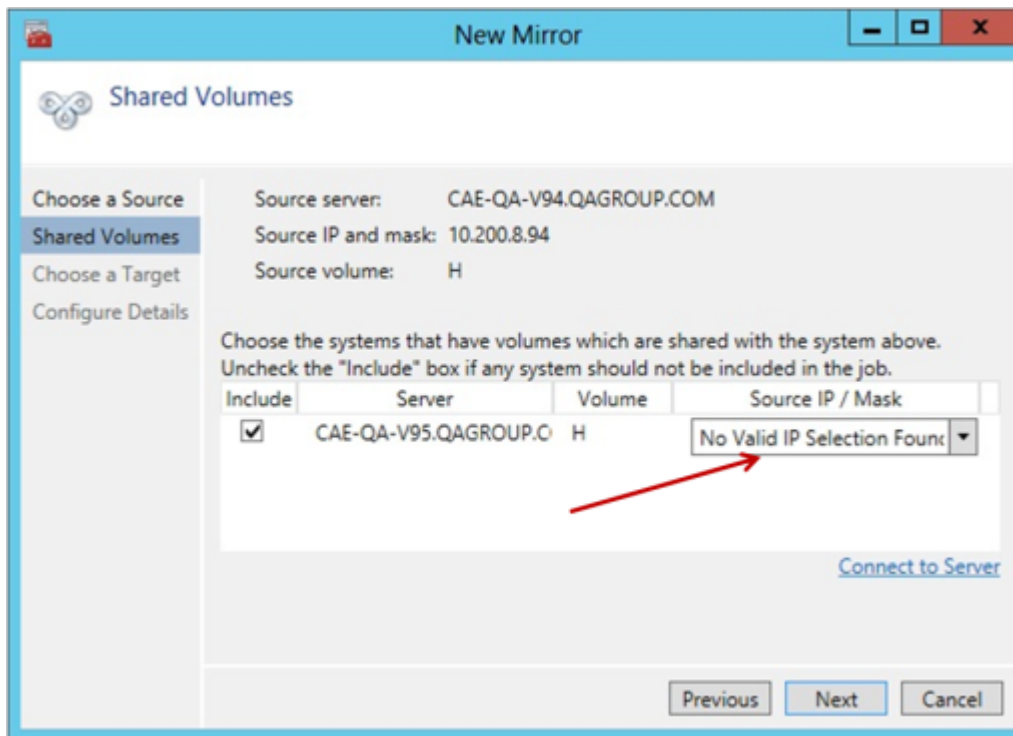
Creating Mirrors with Multiple Targets

The first issue is during mirror creation in a multi-target configuration. In the final step, the user is prompted for secondary relationship information. In previous OS versions, a default Source IP is provided on this **Additional Information Needed** dialog. In Windows Server 2012, however, this default IP is not provided, but the correct IP address must still be selected. If **OK** is clicked without selecting the IP address, the mirror will still create, but key relationship information will be missing.



Creating Mirrors with Shared Volumes

The other issue is with the **Shared Volumes** dialog box when creating mirrors with shared volumes. In previous OS versions, a default Source IP is provided on this screen. In Windows Server 2012, however, this dialog will display "**No Valid IP Selection Found.**" The correct Source IP will still need to be selected.




WSFC 2012 Cluster Creation Default Setting Issue

Description

During the cluster creation process in **Windows Server 2012**, Microsoft has added a new option to automatically consume all disks and manage them through Failover Clustering. As a result, any attempts to create a mirror in DataKeeper will fail and a message will be received in **Disk Manager** that the disk is being managed by **Failover Clustering**.

Suggested Action

To prevent this from happening, in the **Add Node Wizard**, uncheck the box marked "**Add all eligible storage to the cluster**" (shown below). Specific disks can then be added after the cluster is created.



To remedy if already being managed by **Failover Clustering**, remove the

disks from **Available Storage**, then online the disks in **Disk Manager** and use **DataKeeper** to manage the volumes.

WSFC 2012 Failover Cluster Manager UI Defect (Delete Action Missing)

Description

On Windows Server 2012, the Microsoft Failover Cluster Manager UI tool has a defect. When a "right-click" is performed on a DataKeeper Volume resource from the Available Storage group, the drop-down action list does not appear. That action list would normally include the **"Delete"** command (among other actions) to delete the resource from the cluster.

Unfortunately, when the Admin is finished using a DataKeeper Volume storage resource, the resource cannot be removed from the cluster with the **Failover Cluster Manager UI** tool. This appears to affect only non-Microsoft storage resources. Microsoft is working on a correction for this.

Suggested Action

Microsoft has released a Server 2012 Hotfix for this defect. Microsoft Article [2804526](#) provides a high level overview of several WSFC Server 2012 issues including this problem. This article will refer you to several hotfixes for Server 2012. Installing Microsoft Hotfix 2795997 will correct this particular problem. Windows Update KB2803748 must also be installed (this normally occurs automatically). If KB2803748 is not installed, the cluster will become unstable

When requesting this hotfix, click on **"Show hotfixes for all platforms and languages"** and check the **x64** selection. Be sure to also update your 2012 Server with all Windows Updates after installing this hotfix.

The workaround for this issue without installing the Microsoft hotfix is to delete the **"DataKeeper Volume"** resource using **Windows PowerShell**. To remove a DataKeeper Volume resource from a cluster with PowerShell, perform the following command:

```
remove-clusterResource "<DataKeeper Resource Name>"
```

For example:

```
PS C:\> remove-clusterResource "New DataKeeper Volume"

Remove-ClusterResource
Are you sure you want to remove cluster resource 'New DataKeeper Volume'?
[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): Y
PS C:\>
```

In the above example, replace **"New DataKeeper Volume"** with the actual name of your DataKeeper Volume resource.

WSFC 2012 File Server Resource Manager Event Log Errors

Description

In Windows 2012, if a File Server Role is created which uses one or more DataKeeper Volume resources and the **File Server Resource Manager** feature is enabled on the system, then a series of "SRMSVC" errors (ID 8228) will be received on the offline node:

```
File Server Resource Manager was unable to access the following file
or volume: 'E:'. This file or volume might be locked by another
application right now, or you might need to give Local System access
to it.
```

Note: If any DataKeeper Volume resource is offline, this message will be received every ten seconds.

Suggested Action

These messages can be ignored; however, to prevent these messages from being received, **File Server Resource Manager Service** may be disabled.

WSFC 2012 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager

Description

Once a cluster File Server role is created, neither **Server Manager** nor **Failover Cluster Manager** can be used to initially create the share.

Suggested Action

Microsoft Article [2804526](#) provides a high level overview of several WSFC Server 2012 issues including this problem. This article will refer you to several hotfixes for Server 2012.

When using Failover Cluster Manager on Server 2012, the File Share Wizard would not start when right-clicking the "Add File Share" short-cut or when using the right panel "Add File Share" button if third party storage was used. Installing Microsoft Hotfix 2795993 will correct this problem.

Alternatively, installing the following Windows Update modules for Server 2012 will also correct this problem:

KB2815769 KB2803676 KB2785094 KB2779768 KB2771744 KB2761094

KB2812829 KB2800088 KB2784160 KB2779562 KB2771431 KB2758246

KB2812822 KB2795944 KB2783251 KB2778171 KB2770917 KB2756872

KB2811660 KB2790920 KB2782419 KB2777166 KB2769165 KB2751352

KB2803748 KB2788350 KB2780342 KB2771821 KB2764870

The **Server 2012 Windows Update List** shown above was cumulative as of 4/2/2013. Our lab tests showed that Hotfix 2795993 may not install on every Server 2012 system. In that case, we recommend installing at least the Windows Update modules listed above.

On Server 2012, the Server Manager tool could not be used to create

shares on clustered volumes if third party storage was used. Installing Microsoft Hotfix 2796000 will correct this problem. Alternatively, installing the same set of Windows Update modules listed above will also correct this problem.

The workaround if not installing the above is to create the share using **Windows Explorer**. Once the share is created through Windows Explorer, adjusting permissions or other aspects of the file share can be performed normally through **Server Manager** or **Failover Cluster console**.


WSFC 2012 New File Server Type Not Supported

Description

Windows Server 2012 now offers two options for **File Server Resources**:

- File Server for General Use
- **Scale-Out File Server for Application Data (NEW)**

The new option, "**Scale-Out File Server for Application Data**", is not currently supported.



Suggested Action

When selecting a **File Server Type**, the first option, "**File Server for General Use**", must be selected. This File Server type existed in failover

clusters prior to Windows Server 2012. It can be used to increase the availability of files that are shared for use by users or by applications that open and close files frequently.

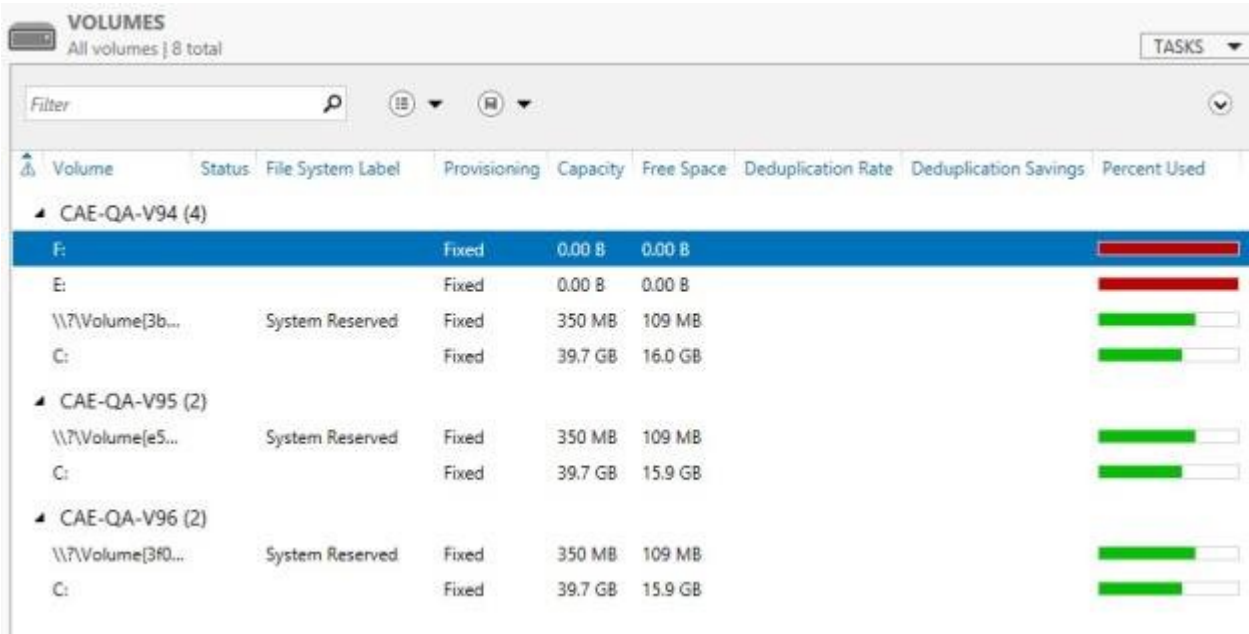
Note: Windows Server 2012 ReFS (Resilient File System) is also not currently supported.

WSFC 2012 Server Manager – Incorrect Volume Display

Description




In Windows Server 2012, volume status can be viewed and volume manipulation can be performed within **Server Manager > File and Storage Services > Volumes**. However, when using DataKeeper volumes with cluster resources, this interface will not accurately reflect volume status.

In the following example, the DataKeeper Volumes E and F are split. One is Cluster Owner\Source on CAE-QA-V95 and the other is Cluster Owner\Source on CAE-QA-V96; however, the **Server Manager "Volumes"** display shows the volumes (E & F) on CAE-QA-V94 with red "percent used" progress bars and does not show any volumes on CAE-QA-V95 or CAE-QA-V96.



If both resources share the same Cluster Owner\Source as in the following example (CAE-QA-V96), the **Server Manager** shows the correct information.

VOLUMES
All volumes | 8 total


Filter   

Volume	Status	File System Label	Provisioning	Capacity	Free Space	Deduplication Rate	Deduplication Savings	Percent Used
CAE-QA-V94 (2)								
\\?\Volume{3b...}	System Reserved		Fixed	350 MB	109 MB			<div><div></div></div>
C:			Fixed	39.7 GB	16.0 GB			<div><div></div></div>
CAE-QA-V95 (2)								
\\?\Volume{e5...}	System Reserved		Fixed	350 MB	109 MB			<div><div></div></div>
C:			Fixed	39.7 GB	15.9 GB			<div><div></div></div>
CAE-QA-V96 (4)								
\\?\Volume{3f0...}	System Reserved		Fixed	350 MB	109 MB			<div><div></div></div>
C:			Fixed	39.7 GB	15.9 GB			<div><div></div></div>
E:	Volume1		Fixed	4.97 GB	3.95 GB			<div><div></div></div>
F:	Volume2		Fixed	4.97 GB	3.95 GB			<div><div></div></div>

WSFC 2012 Server Manager – DataKeeper “Disk” Not Shown as Clustered


There are some inconsistencies in how cluster disks are shown in the “Disks” display under **Server Manager > File and Storage Services > Volumes > Disks**.

The following screen shot shows an iSCSI Shared Disk and two DataKeeper Volumes in a cluster. There is a check mark under the “**Clustered**” column heading for the iSCSI Disk, however, the DK Clustered Volumes do not have a check mark even though they are in a cluster. There will be nothing displayed in either the “**Read Only**” column or the “**Clustered**” column. This is due to the fact that DataKeeper operates with **Volumes**, not **Disks**.



Windows 2012 File Share

When creating a fileshare using "SMB - Basic" on Windows 2012, by default the "Enable Continuous availability" flag is checked which prohibits creating a fileshare resource. To solve this problem, uncheck the box as shown in the picture.



DataKeeper Volume cannot come Online after Network failure with clustered IP Address on Replication network

If you have multiple cluster networks, IP Addresses shouldn't be set up on the same network that DataKeeper Volume resources are using for replication. Network errors may cause the DataKeeper mirror to go into a Paused state. If the network error also causes the cluster IP Address resource to fail its health checks, any resource hierarchy that contains both a DataKeeper Volume resource and the cluster IP Address will not be brought Online on remote nodes due to the DataKeeper Volume mirror state being in a non-Mirroring state.

WSFC – MS DTC Resource Failure

Error/Message

The cluster resource host subsystem (RHS) stopped unexpectedly. An attempt will be made to restart it. This is usually due to a problem in a resource DLL. Please determine which resource DLL is causing the issue and report the problem to the resource vendor.

Description

In Windows Failover Clustering, an MS DTC Resource fails to come online if configured with a DataKeeper volume resource.

Log Name: System
Source: Microsoft-Windows-FailoverClustering
Date: <Date Time>
Event ID: 1146
Task Category: Resource Control Manager

Suggested Action

Install Service Pack 1 for Windows 2008 R2. This will allow the MS DTC resource to operate properly with a DataKeeper volume resource.

WSFC 2008 R2 SP1 Procedure Change

Description

When using WSFC 2008 R2 SP1, the procedure for Extending a Traditional 1×1 2-Node WSFC Cluster to a Shared-Replicated 3-Node Cluster has changed. The WSFC mmc GUI must not be used for adding a node that is hosting a DataKeeper shared volume.

Suggested Action

When using WSFC 2008 R2 SP1, additional nodes with shared DataKeeper volumes can be safely added to an existing cluster only with the WSFC Command line tool "cluster /add /node:<standby node name>. **VERY IMPORTANT.** Please refer to the topic [Extending a Traditional 2-Node Cluster to a Shared-Replicated Configuration](#) for more details.

Windows Server 2016 Specific Issues

For issues related to **Windows Server 2016**, see the following topic:

-
- [Occasional Job Creation Failure](#)
 - [WSFC 2016 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager](#)

Occasional Job Creation Failure

Occasionally, new job creation on Windows 2016 systems can fail. If this occurs, retry the create.

WSFC 2016 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager

Description

Once a cluster File Server role is created, neither **Server Manager** nor **Failover Cluster Manager** can be used to initially create the share.

Suggested Action

Microsoft Article [2804526](#) provides a high level overview of several WSFC Server 2016 issues including this problem. This article will refer you to several hotfixes for Server 2012.

When using Failover Cluster Manager on Server 2016, the File Share Wizard would not start when right-clicking the "Add File Share" short-cut or when using the right panel "Add File Share" button if third party storage was used. Installing Microsoft Hotfix 2795993 will correct this problem.

Alternatively, installing the following Windows Update modules for Server 2012 will also correct this problem:

KB2815769 KB2803676 KB2785094 KB2779768 KB2771744 KB2761094

KB2812829 KB2800088 KB2784160 KB2779562 KB2771431 KB2758246

KB2812822 KB2795944 KB2783251 KB2778171 KB2770917 KB2756872

KB2811660 KB2790920 KB2782419 KB2777166 KB2769165 KB2751352

KB2803748 KB2788350 KB2780342 KB2771821 KB2764870

The **Server 2012 Windows Update List** shown above was cumulative as of 4/2/2013. Our lab tests showed that Hotfix 2795993 may not install on every Server 2012 system. In that case, we recommend installing at least the Windows Update modules listed above.

On Server 2012, the Server Manager tool could not be used to create

shares on clustered volumes if third party storage was used. Installing Microsoft Hotfix 2796000 will correct this problem. Alternatively, installing the same set of Windows Update modules listed above will also correct this problem.

The workaround if not installing the above is to create the share using **Windows Explorer**. Once the share is created through Windows Explorer, adjusting permissions or other aspects of the file share can be performed normally through **Server Manager** or **Failover Cluster console**.

Restrictions

Included below are restrictions associated with DataKeeper and DataKeeper Cluster Edition as well as possible workarounds and/or solutions.

[Bitlocker Does Not Support DataKeeper](#)

[CHANGEMIRRORENDPOINTS Restriction](#)

[CHKDSK](#)

[DataKeeper Volume Resize Restriction](#)

[Directory for Bitmap Must Be Created Prior to Relocation](#)

[Duplicate IP Addresses Disallowed Within a Job](#)

[Intensive I-O with Synchronous Replication](#)

[Resource Tag Name Restrictions](#)

Bitlocker Does Not Support DataKeeper

According to Microsoft, Bitlocker is not supported to work with Software RAID configurations. Since DataKeeper is essentially a software RAID 1, Microsoft does not support Bitlocker working with DataKeeper.

Note: EFS (Encrypting File System) and TDE (Transparent Disk Encryption) are compatible with DataKeeper and can be used to encrypt data. In addition, both will also encrypt the data sent over the network by DataKeeper.

The specific article and section can be found here:

https://technet.microsoft.com/en-us/library/ee449438#BKMK_R2disks

CHANGEMIRRORENDPOINTS

Description:

This command, which is used to move a DataKeeper protected volume to another network location, only supports changing the endpoints of a mirrored volume that is configured on 3 nodes or fewer.

Workaround:

For configurations greater than three nodes, the mirrors must be deleted and recreated with the final endpoint at the local site and use route adds to get the mirrors created and resynced before moving the server to the final location/address/DR site.

CHKDSK

Description

If you must run CHKDSK on a volume that is being replicated by SIOS DataKeeper, it is recommended that you **PAUSE** the mirror before initiating the CHKDSK. After running CHKDSK, **CONTINUE** the mirror. A [partial resync](#) occurs (updating those writes generated by the CHKDSK) and replication will continue.

Note: The bitmap file (for non-shared volumes) is located on the C drive which is defined by [BitmapBaseDir](#) as the default location. Running CHKDSK on the C drive of the **Source** system will cause an error due to the active bitmap file. Therefore, a switchover must be performed so that this Source becomes Target and the bitmap file becomes inactive. The CHKDSK can then be executed on this system as the new target (original source).

DataKeeper Volume Resize Restriction

The DataKeeper volume resize procedure should be performed on only one volume at a time.

Directory for Bitmap Must Be Created Prior to Relocation

Description

If you choose to relocate the bitmap file from the default location (**%EXTMIRRBASE%\Bitmaps**), you must first create the new directory before changing the location in the registry and rebooting the system.

Duplicate IP Addresses Disallowed Within a Job

A DataKeeper job contains the endpoint information for all mirrors that are part of the job. This information includes the host name, IP address, and drive letter of each mirror endpoint.

Within a job, an IP address cannot be duplicated on more than one node. For example, in a 4-node job, nodes "A" and "B" may be configured with a private network connection, and nodes "C" and "D" may be configured with a separate private network connection. However, the IP addresses on those private networks must be unique for each node. If nodes A and B use 192.168.0.1 and 192.168.0.2 for replication, then nodes C and D cannot also use 192.168.0.1 or 192.168.0.2 for replication.

Intensive I-O with Synchronous Replication

Description

Due to the nature of synchronous replication (blocking volume writes while waiting for a response from the target system), you may experience sluggish behavior with any applications that are writing to the mirrored volume. The frequency of these events could be high depending on the ratio of "Volume I/O traffic" to "system resource". It is recommended that you use asynchronous replication when continuous and intensive I/O traffic is anticipated for the volume or when SIOS DataKeeper is used on a low bandwidth network.

Resource Tag Name Restrictions

Tag Name Length

All tags within DataKeeper may not exceed the 256 character limit.

Valid "Special" Characters

`-_./`

However, the first character in a tag should not contain "." or "/".

Invalid Characters

`+ ; : ! @ # $ * = "space"`

Troubleshooting

The topics in this section contain important information about known issues and restrictions offering possible workarounds and/or solutions.

[Applet Troubleshooting](#)

[Error When Attempting to Run SIOS Protection Suite Command From Command Prompt](#)

[Firewall](#)

[GUI Error Messages](#)

[GUI Network Related - Initial Connection to Server Failed](#)

[GUI Network Related - Long Connection Delays on Windows Platforms](#)

[GUI Network Related - NoRouteToHostException Message Generated During Connection Attempt](#)

[GUI Network Related - Unknown Host Exception Message Generated During Connection Attempt](#)

[GUI Server Troubleshooting](#)

[Health Check Timeouts](#)

[Incomplete Resource Creation](#)

[Installation - Access is Denied](#)

[IP Resource Create Issue](#)

[Java Mixed Signed and Unsigned Code Warning](#)

[LANMAN Name May Be Seen Twice in Browse List](#)

[Licensing - Licensed Recovery Kit Resource Fails to Come In Service](#)

[Licensing - License Key Not Found](#)

[LifeKeeper GUI Issues with Java 1.8.x](#)

[SIOS Protection Suite Web Client May Lock Up](#)

[SIOS Protection Suite Web Client May Lock Up - Multiple Changes Made to Existing Hierarchy](#)

[New Evaluation License Key Error](#)

[Recovering From a SIOS Protection Suite Server Failure in a 1 × 1 Configuration](#)

[Recovering Out-of-Service Hierarchies](#)

[Remove Hangs During Recovery Kit Uninstall](#)

[Replicated Volume Switchover Failure](#)

[Resource Tag Name Restrictions](#)

[Restore and Health Check Account Failures](#)

[SQL 2008](#)

[SQL Server Reporting Services](#)

[Two-Server Cluster Issue](#)

[Unknown User Name or Bad Password](#)

[Web Client Troubleshooting](#)

[Win2008 - IIS Resource Hierarchy Creation Error](#)

Applet Troubleshooting

Description

If the web client does not display the [Cluster Connect dialog](#), try the following:

1. Check whether the applet failed. Usually a message is printed somewhere in the browser window specifying the state of the applet. In Internet Explorer, an icon may appear instead of the applet, in addition to some text status. Clicking this icon may bring up a description of the failure.
2. Open the Java Console.
 - For Firefox and older versions of Internet Explorer, run the **Java Plug-In applet** from your machine's **Control Panel** and select the option to show the console, then restart your browser.
 - For recent versions of Internet Explorer, select **Tools > Sun Java Console**. If you do not see the Sun Java Console menu item, select **Tools > Manage Add-Ons** and enable the console, after which you may need to restart your browser before the console will appear.
 - For Mozilla, select **Tools > Web Development > Sun Java Console**.
3. If the web client is not open, reopen the URL *http://<server name>:81* to start it.
4. Check the console for any messages. The messages should help you resolve the problem. If the problem appears to be network related, refer to the Network-Related Troubleshooting section.

Error When Attempting to Run SIOS Protection Suite Command From Command Prompt

Symptom

When attempting to run a SIOS Protection Suite command from a command prompt, you receive the following error:

```
[File:lock.CLine:1610] Win32 Error: 2  
CRITICAL (No. 472) Can't run this application without LCD Daemon  
running.
```

Solution

SIOS Protection Suite commands require "console" rights to run. When the SIOS Protection Suite core is running on Server 2008 or later, invoke Remote Desktop connections with the "/admin" switch.

e.g. %SystemRoot%\system32\mstsc.exe/console

You may also run the SIOS Protection Suite command from the command prompt on the SIOS Protection Suite system itself.

Firewall

Symptom

Firewall was disabled during installation but is now enabled. How do I add the rules to Windows firewall?

Solution

There is a script in the LifeKeeper directory that will allow you to enable the firewall rules. The script is located in:

```
<LifeKeeper Root Directory>\support\firewallsetup.bat
```

By opening a command prompt and executing **firewallsetup.bat <LifeKeeper Root directory>**, you can add the rules. If the rules had been added, the script will not add duplicate rules.

If you open the Windows firewall (wf.msc), you will see the inbound rules prefixed with the LifeKeeper label.

Note: If you have created specific rules that have disabled the ports required by SIOS Protection Suite, the installation program will disable but will not delete those rules.

GUI Error Messages

Description

Error 101: Illegal argument was passed.

Error 102: This program requires a Java Virtual Machine Version 1.5 or greater to run properly. Please refer to the [LifeKeeper GUI](#) documentation to verify your setup.

Error 103: Could not set Look and Feel for LifeKeeper GUI.

Error 104: <filename> Image could not be loaded.

Error 106: Error trying to get data over RMI. Could not complete action.

Error 107: Failed to create Global Resource Instance.

Error 108: Failed to create Global Resource.

Error 109: Dialog requires a Server to be selected.

Error 112: Could not match Resource Instance to Global Equivalency.

Error 114: <server name> Security Exception caused connection failure to this server. Please note that this failure may result in degraded cluster resource rendering. Verify that your Java Policy file is installed properly. See [Running the SIOS Protection Suite Web Client](#).

Error 115: <server name> Name of this server could not be resolved resulting in a connection failure to this server. Please note that this failure may result in degraded cluster resource rendering. Please refer to the [LifeKeeper GUI](#) documentation to verify network naming conventions. See [Unknown Host Exception](#).

Error 116: <server name> This server could not resolve the name of this client host resulting in a connection failure to this server. Please note that this failure may result in degraded cluster resource rendering. Please refer to the [LifeKeeper GUI](#) documentation to verify network naming conventions. See [Unknown Host Exception](#).

Error 117: Initial connection to server failed. LifeKeeper GUI will continue trying to connect. Please refer to the [LifeKeeper GUI](#)

documentation to verify that SIOS Protection Suite and the LifeKeeper GUI server are active on this system. See [Java RMI Binding Problem](#).

Error 118: Incompatible client and server packages caused connection failure. Please verify that the versions are compatible between the target server and the server from which the client was started.

Error 119: Could not export remote object.

Error 120: Encountered exception when invoking remote method.

Error 121: Administrative java bean could not be initialized.

Error 122: Administrative java bean has no dialog classes to load. The properties file describing the administrative java bean is missing the "list" property.

Error 123: The properties file describing the administrative java bean has a missing property.

Error 124: Failed to find property bundle.

Error 125: Security Exception trying to create URLClassLoader. Please verify that the .java.policy file grants the proper permissions. You should typically create a .java.policy file in your home directory. The contents of the .java.policy file are case sensitive, so it is best to copy the sample file that is distributed with the LifeKeeper GUI package.

If you are using a browser plug-in for Java, then the user home directory that is being used for the java environment can be verified by enabling the Java console and examining the first few lines that are displayed. Refer to [Configuring the LifeKeeper GUI](#) for more information on configuring the GUI client.

Error 126: Could not find resource on server.

Error 127: Could not find extend properties file for this kit.

Error 128: Internal properties file error.

Error 129: Cannot establish an RMI connection to the server. Verify the LifeKeeper GUI Server is running on the server.

Error 130: The tag entered is being used by another resource. Please enter another tag.

Error 131: Exception calling invokeAndWait method to update the user interface.

Error 132: Encountered exception when invoking administrative java bean.

Error 133: Invalid value entered for equivalency priority. The priority value must be in the range of 1 through 999.

Error 134: The equivalency priority value conflicts with another priority in the table. Each equivalency priority value must be unique in the table.

GUI Network Related – Initial Connection to Server Failed (Error 117)

Symptom

Initial Connection to server failed (Error 117).

If you are attempting to connect to a server that has two or more network interface cards (NICs), it could indicate a Java RMI binding problem where the first NIC (the one that appears first in the output of ipconfig utility) has a non-reachable IP address.

Solution

You may need to reorder the protocol binding for use by the network services of the SIOS Protection Suite server. On each SIOS Protection Suite server, open "Network and Dial-up Connections", and on the **Advanced** menu, select **Advanced Settings**. The **List Box** at the top of the dialog shows the current order of the NIC cards. Click the **arrow button** to reorder them so that the reachable NIC is at the top of the list. This should enable Java RMI to allow client to connect to the server. A reboot of the server is required for this to take effect.

GUI Network Related – Long Connection Delays on Windows Platforms

There are two ways to address this issue. One is described here in "Solution 1 – IPV4 Interface Metric", the other in "Solution 2 – local hosts file".

Symptom

Long Connection Delays on Windows Platforms.

Solution 1 – IPV4 Interface Metric

This applies to LifeKeeper systems with more than one NIC connected to different subnets. Java's RMI host lookup uses each interface's IPV4 interface metric to determine the host IP address. In order to ensure that RMI communicates on the same network in both directions of a connection, the IPV4 interface metric should be specified so that a host's public NIC has the lowest interface metric. This will ensure that the host's public IP address will be reported.

By default, Windows network interfaces are configured with an "Automatic" interface metric. In some cases, this results in the wrong network being used for RMI communication between an LK GUI client and an LK host. The Automatic metric should be disabled, and an appropriate metric should be specified. For the public NIC, specify the metric "1", for each additional NIC specify increasing metric values (2, 3, 4...).

Solution 2 – local hosts file

From Sun FAQ:

"Most likely, your host's networking setup is incorrect. RMI uses the JavaAPI networking classes, in particular `java.net.InetAddress`, which will cause TCP/IP host name lookups for both host to address mapping and address to hostname. On Windows, the lookup functions are performed by the native Windows socket library, so the delays are not happening in RMI but in the Windows libraries. If your host is set up to use DNS, then

this could be a problem with the DNS server not knowing about the hosts involved in communication and what you are experiencing are DNS lookup timeouts. If this is the case, try specifying all the hostnames/addresses involved in the *local file\winnt\system32\drivers\etc\hosts* or *\windows\hosts*. The format of a typical host file is:

IPAddress Server Name

e.g.: 208.2.84.61 homer.somecompany.com

This should reduce the time it takes to make the first lookup.”

In addition, incorrect settings of the Subnet Mask and Gateway address may result in connection delays and failures. Verify with your Network Administrator that these settings are correct.

GUI Network Related - NoRouteToHostException Message Generated During Connection Attempt

Symptom

NoRouteToHostException Message Generated During Connection Attempt.

A socket could not be connected to a remote host because the host could not be contacted.

Solution

Typically, this indicates that some link in the network between the local and remote server is down or that the remote server is behind a firewall.

GUI Network Related – Unknown Host Exception Message Generated During Connection Attempt

Symptom

Unknown Host Exception Message Generated During Connection Attempt.

The LifeKeeper GUI Client and Server use Java RMI (Remote Method Invocation) technology to communicate. For RMI to work correctly, the client and server must use resolvable hostname or IP addresses. When unresolvable names, WINS names, or unqualified DHCP names are used, this causes Java to throw an UnknownHostException.

This error message may also occur under the following conditions:

- Server name does not exist. Check for misspelled server name.
- Misconfigured DHCP servers may set the fully qualified domain name of RMI servers to be the domain name of the resolver domain instead of the domain in which the RMI server actually resides. In this case, RMI clients outside the server's DHCP domain will be unable to contact the server because of the incorrect domain name.
- The server is on a network that is configured to use Windows Internet Naming Service (WINS). Hosts that are registered under WINS may not be reachable by hosts that rely solely upon DNS.
- The RMI client and server reside on opposite sides of a firewall. If your RMI client lies outside a firewall and the server resides inside of it, the client will not be able to make any remote calls to the server.

Solution

When using the LifeKeeper GUI, the hostname supplied by the client must be resolvable from the server and the hostname from the server must be resolvable by the client. The LifeKeeper GUI catches this exception and alerts the user. If the client cannot resolve the server hostname, this exception is caught and Message 115 is displayed. If the server cannot

resolve the Client hostname, this exception is caught and Message 116 is displayed. Both of these messages include the part of the Java exception which specifies the unqualified hostname that was attempted.

Included in the following sections are some procedures that may be used to test or verify that hostname resolution is working correctly.

From Windows

1. Verify communication with the SIOS Protection Suite server. From a prompt, ping the target using the hostname:

```
ping<TARGET_NAME>
```

For example;

```
ping homer
```

A reply listing the target's qualified hostname and IP address should be seen.

2. Verify proper configuration.
 - a. Check configuration of DNS or install a DNS server on your network.
 - b. Check the settings for **ControlPanel->Network->Protocols->TCP/IP**. Verify with your Network Administrator that these settings are correct. Note that the hostname in the DNS tab should match the name used on the local name server. This should also match the hostname specified in the GUI error message.
 - c. Try editing the hosts file to include entries for the local host and the SIOS Protection Suite servers that it will be connected to.

On Windows 2008 R2 and 2012 systems, the hosts file is:

```
%SystemRoot%\system32\drivers\etc\HOSTS (e.g. C:\windows\
system32\drivers\etc\HOSTS)
```

Note: On Windows 2008 R2 and 2012, if the last entry in the hosts file is not concluded with a carriage-return/line-feed, then the hosts file will not be read at all.

For example, if my system is called HOSTCLIENT.MYDOMAIN.COM and uses the IP address 153.66.140.1, add the following entry to the hostsfile:

```
153.66.140.1 HOSTCLIENT.MYDOMAIN.COM
```

3. Try setting the hostname property to be used by the GUI client. To do this from a browser with the Plug-in, open the **Java Plug-In Control Panel** and set the host name for the client by adding the following to "Java Run Time Parameters."

```
-Djava.rmi.server.hostname=<MY_HOST>
```

4. Check for Microsoft network-related patches at www.microsoft.com.

From Linux

1. Verify communication with the other server by pinging the target server from Linux using its hostname or IP address:

```
ping -s<TARGET_NAME>
```

For example:

```
ping -s homer
```

A reply listing the target's qualified hostname should be seen.

2. Verify that *localhost* is resolvable by each server in the cluster using ping with its hostname or IP address. 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).
3. Check that DNS is specified before NIS. DNS should be put before NIS in the host's line of */etc/nsswitch.conf*, and */etc/resolv.conf* should point to a properly configured DNS server(s).
4. If DNS is not to be implemented or no other method works, edit the */etc/hosts* file to add an entry for the hostname.
5. Try setting the hostname property to be used by the GUI client. This will need to be changed for each administrator.

To do this from a browser with the Plug-in, open the **Java Plug-In**

Control Panel and set the hostname for the client by adding the following to **Java RunTime Parameters**:

```
-Djava.rmi.server.hostname=
```

To do this from the HotJava browser, append the following to the hotjava command line:

```
-Djava.rmi.server.hostname=<MY_HOST>
```

For Example:

```
-Djava.rmi.server.hostname=153.66.140.1
```

```
-Djava.rmi.server.hostname= homer.somecompany.com
```

GUI Server Troubleshooting

Symptom

The LifeKeeper GUI uses Ports 81 and 82 on each server for its administration web server and Java remote object registry. If another application is using the same ports, the LifeKeeper GUI will not function properly.

Solution

These values may be changed by editing the following registry entries:

```
GUI_WEB_PORT=81
```

```
GUI_RMI_PORT=82
```

These entries are located in the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\SIOS\LifeKeeper\JavaGUI\Server
```

***Note:** The port values are initialized in the GUI server when it is started. If you alter them, you will need to stop and restart the GUI server. These values must be the same across all clusters to which you connect.

Health Check Timeouts

SYMPTOM: Occasionally, SIOS Protection Suite is installed on a system that is not performing well. There are tell-tale signs of abnormal system behavior that SIOS Protection Suite can detect such as the health check processes not starting or ending properly. The most common problem causing check process timeouts is incorrect system memory optimization for databases or mail servers. There must be enough memory to start health check processes and initiate failovers at all times.

SOLUTION: The SIOS Protection Suite Release Notes include guidelines that identify system memory requirements for SIOS Protection Suite. It also briefly explains how to use the Windows Performance Monitors to verify that enough memory is available for applications such as SIOS Protection Suite in your system. To help identify this situation, there are two Resource Monitoring options available to record abnormal behavior and three options to take corrective action when check process timeouts are occurring. All five options can be enabled or disabled as SIOS Protection Suite registry settings to meet specific customer requirements and preferences:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\STEELEYE\LifeKeeper\General\
ResMon_RecordTimeout
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\STEELEYE\LifeKeeper\General\
ResMon_RecordMemory
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\STEELEYE\LifeKeeper\General\
ResMon_ResFail
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\STEELEYE\LifeKeeper\General\
ResMon_RebootWaitInSec
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\STEELEYE\LifeKeeper\General\
ResMon_ResFailMaxWaitInMin
```

The first 4 options (**RecordTimeout**, **RecordMemory**, **ResFail**, and **RebootWaitInSec**) are triggered by any SIOS Protection Suite Quick Check or Deep Check process monitoring a protected resource that does not complete in the expected amount of time.

ResMon_RecordTimeout - This option records any Quick Check or Deep Check process that times out. This information is logged in *<LifeKeeper Root*

Folder>\Out\ResMonTimeout.log file. This option is enabled (=1) by default. To disable, set the value to 0.

ResMon_RecordMemory - This option records system memory usage and process memory usage for every active process whenever a Quick Check or Deep Check process times out. Memory usage is logged in **<LifeKeeper Root Folder>\Out\ResMonTimeout.log file.** This option is enabled (=1) by default. To disable, set the value to 0.

ResMon_ResFail - This option causes a resource hierarchy failover whenever a Quick Check or Deep Check process times out. All dependent resources in the affected hierarchy are failed over. This option is disabled (=0) by default. To enable, set the value to 1.

ResMon_RebootWaitInSec - This option causes a system reboot whenever a Quick Check or Deep Check process times out. This option is disabled (=0) by default. To enable, enter any non-zero number in the registry setting. This number will become the countdown displayed on the system console when the system is rebooted by this feature. The reboot sequence is totally automatic and designed for unattended system operation. Once the countdown is started, the reboot cannot be stopped.

ResMon_ResFailMaxWaitInMin - This option will monitor the SIOS Protection Suite failover process whenever any failover is occurring. The value of this registry setting is the number of minutes SIOS Protection Suite will wait for a resource hierarchy failover to complete. If the failover process cannot be started or if the failover does not complete in the specified number of minutes, SIOS Protection Suite will attempt to reboot the system. If the resources were not failed over, they will come in service again on the same server that was rebooted. This option is disabled (=0) by default.

The **Resource Monitoring Options** apply to all protected resources and they can be changed at any time. Changed settings take affect the next time a Quick Check or Deep Check process is started.

Incomplete Resource Creation

Description

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` and `dep_remove` to clean up the partial hierarchies.

Installation - Access is Denied

Symptom

During upgrade or re-installation, SIOS Protection Suite generates "Access is denied" error message.

Solution

The SIOS Protection Suite services have not all stopped. This can occur if Setup was unable to stop the SIOS Protection Suite services. Open a command window and enter `$LKROOT\bin\lkstop` to stop all the SIOS Protection Suite services, then wait until you see "LIFEKEEPER NOW STOPPED" before running Setup.

IP Resource Create Issue

Symptom

IP resource ***Pre-Extend Wizard*** can fail if new IP address partially matches an existing IP address.

Workaround

Currently no workaround is available. This will be addressed in a future release.

Java Mixed Signed and Unsigned Code Warning

Symptom

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.

Solution

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.

LANMAN Name May Be Seen Twice in Browse List

Symptom

After creating a LANMAN resource, the LANMAN name may be seen twice in the browse list.

Solution

One of these entries is an unusable Workstation record. Please disregard this entry and use the other LANMAN Server name.

Licensing – Licensed Recovery Kit Resource Fails to Come In Service

Symptom

After upgrade, licensed recovery kit resource fails to come in service and the following error is logged to the **Application Event Log** by SIOS Protection Suite:

```
"Process: lcdmachfail(3176) *ERROR* (No. 1001) resource <tag name>
requires a license (for Kit <recovery kit type>) but none is
installed."
```

Solution

Use the SIOS Protection Suite licensing utility to install your recovery kit license key. See [Obtaining and Installing the License](#) for information on installing a license.

Licensing – License Key Not Found

Symptom

After installing licensed recovery kit, the following error is logged to the **Application Event Log** by SIOS Protection Suite:

```
"Process: Lkinit:(1832) *ERROR* (No. 20042) SPS Recovery Kit  
<licensed recovery kit> license key NOT FOUND".
```

Solution

Use the SIOS Protection Suite licensing utility to install your recovery kit license key. See [Obtaining and Installing the License](#) for information on installing a license.

LifeKeeper GUI Issues with Java 1.8.x

Symptom

Unable to load the LifeKeeper GUI via a web browser on a client system using JRE v7 and JRE v8

Solution

In the Java Control Panel on the client system, edit security settings, define an Exception Site List, and enter in the cluster nodes.

SIOS Protection Suite Web Client May Lock Up

Symptom

The SIOS Protection Suite web client may lock up when used from a server machine if the "-" (reduce resource height) or "+" (increase resource height) accelerator keys are hit during the initial paint of SIOS Protection Suite resources for that server.

Solution

To recover, open Windows Task Manager and select **"End Task"** for the **"LifeKeeper - <web browser, e.g., Microsoft Internet Explorer>"** application. It may take up to one minute for the processes to end. Restart the SIOS Protection Suite web client from the **Start->All Programs->SIOS** shortcut and wait for initial screen paint to complete before using accelerator keys.

SIOS Protection Suite Web Client May Lock Up - Multiple Changes Made to Existing Hierarchy

Symptom

The SIOS Protection Suite web client may lock up when used from the server machine if multiple changes are made to an existing hierarchy (i.e. create/delete dependencies), closing and reopening the SIOS Protection Suite client between changes.

Solution

To recover, open **Windows Task Manager** and select **End Task** for the **"LifeKeeper - <web browser, e.g., Microsoft Internet Explorer>"** application. It may take up to one minute for the processes to end. For hierarchy administration on the server, use the LifeKeeper GUI (Admin Only) application. Shortcut is **Start->All Programs->SIOS->LifeKeeper->LifeKeeper** (Admin Only).

New Evaluation License Key Error

Symptom

When evaluating SIOS Protection Suite for Windows, an error may occur if a new evaluation license key is not used. The old evaluation licenses will not work on this release.

Solution

A new evaluation license key must be obtained. Restart the **License Key Manager** and enter the new, properly formatted license key.

Recovering From a SIOS Protection Suite Server Failure in a 1 × 1 Configuration

If a server in your SIOS Protection Suite cluster experiences a failure that causes re-installation of the operating system (and thus SIOS Protection Suite), you will have to re-extend the resource hierarchies from each server in the cluster. If a server in the cluster has a shared equivalency relationship with the re-installed server, however, SIOS Protection Suite will not allow you to extend the existing resource hierarchy to the re-installed server. SIOS Protection Suite 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:

After reinstalling your operating systems and all related patches as well as SIOS Protection Suite, begin the following steps for recovery (**Note:** The examples in the Suggested Action below are using test system names "BENHOGAN" and "GPLAYER"):

1. On each server where the resource hierarchies are configured, use the [eqv_list](#) command to obtain a list of all the shared equivalencies.

```
eqv_list [-d destsys] [-s sys] [-t tag] [-e SHARED] [-fC]
```

This function prints strings to standard output describing equivalency relationships between resource instances.

Example:

```
c:\LK\Bin>eqv_list
```

```
BENHOGAN2008 LK1-BENHOGAN2K8 GPLAYER2008 LK1-BENHOGAN2K8 SHARED  
1 10
```

```
BENHOGAN2008 Vol.E GPLAYER2008 Vol.E SHARED 1 10
```

See LifeKeeper GUI below of hierarchies, resources, etc.



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.

This function removes equivalency from the configuration database on system destsys (local if not specified) of equivalency type, specified by the **-e** option, between the resources tag and othertag existing on systems sys and othersys respectively.

```
eqv_remove [-d destsys] [-s sys] -t tag [-S othersys]-o
othertag [-e SHARED]
```

```
eqv_remove -s {this system} -t {TAGNAME} -S {othersys that has
gone away} [-e SHARED]
```

Example:

```
c:\LK\Bin>eqv_remove -s BENHOGAN2008 -t LK1-BENHOGAN2K8 -S
GPLAYER2008 -e SHARED
```

```
c:\LK\Bin>
```

3. Execute the [eqv_list](#) command again. There should be "no" list of shared equivalencies, etc.

Notice on the LifeKeeper GUI below how the target resources have been removed:



4. If there are any DataKeeper mirrored volumes configured in your cluster, clean up the **LKDRInfo** file for that volume.

SIOS Protection Suite will create an **LKDRInfo.<volume>** file in the folder **%LKROOT%\subsys\filesys\resources\volume** for each mirrored volume resource. This file should be deleted.

```
C:\LK\BIN> cd %LKROOT%\subsys\filesys\resources\volume
```

```
C:\LK\subsys\filesys\resources\volume> dir LKDRInfo.E
```

```
05/23/2012 01:02 PM 39 LKDRInfo.E
```

```
C:\LK\subsys\filesys\resources\volume> del LKDRInfo.E
```

5. If there are any DataKeeper mirrored volumes configured in your cluster, clean up the mirror and delete the DataKeeper job that contains that mirror.

DataKeeper mirrors and jobs must be recreated when the cluster is re-extended to the reinstalled server. Therefore, the local end of any mirrors that are configured must be deleted and any jobs that are configured with those mirrors must be deleted.

```
C:\LK\subsys\filesys\resources\volume> cd %ExtMirrBase%
```

```
C:\Program Files\SIOS\DataKeeper> emcmd . getjobinfoforvol E
```

```
ID = e829700c-27b0-447f-b852-1a3135da31a7
```

```
Name = E Vol
```

```
Description =
```

```
MirrorEndpoints = BENHOGAN2008;E;10.200.8.25;GPLAYER2008;E  
;10.200.8.26;A
```

```
C:\Program Files\SIOS\DataKeeper> reg /delete HKLM\System\  
CurrentControlSet\Services\ExtMirr\Parameters\Jobs\  
e829700c-27b0-447f-b852-1a3135da31a7
```

```
Permanently delete the registry key HKEY_LOCAL_MACHINE\System\  
CurrentControlSet\Services\ExtMirr\Parameters\Jobs\  
1b5e8715-a488-4030-8166-45c9232bc04e (Yes/No)? y
```

The operation completed successfully.

```
C:\Program Files\SIOS\DataKeeper> emcmd . getjobinfoforvol E
```

```
C:\Program Files\ SIOS\DataKeeper> emcmd .  
deletelocalmirroronly E
```

```
Status = 0
```

```
C:\Program Files\ SIOS\DataKeeper> emcmd . clearswitchover E
```

```
Status = 0
```

Repeat these steps for all mirrored volumes. **The job must be deleted from the registry directly; attempting to delete the job using "emcmd . deletejob" will fail because DataKeeper tries to delete the job on all nodes. This will not work since one of the nodes no longer exists and the job will be left intact.**

6. You are now ready to remove the old system.

a. Display all systems that were in the SIOS Protection Suite cluster:

```
sys\_list [-d destsys]
```

Example:

```
c:\LK\Bin>sys_list
```

```
BENHOGAN2008
```

```
GPLAYER2008
```

- b. Remove all old targets/systems.

[sys_remove](#) [-d destsys] -s sys

Example:

```
c:\LK\Bin>sys_remove -s GPLAYER2008
```

Do this for all systems participating in the SIOS Protection Suite cluster

Note: If attempting to execute this command for the local/source system, the following error will be received:

```
c:\LK\Bin>sys_remove -s BENHOGAN2008
```

```
(null)Process: sys_remove(2728)
```

```
[File:sys.C Line:81]
```

```
*ERROR* (No. 424) can't remove entry for local system  
"BENHOGAN2008"
```

7. Remove communication paths.

- a. Verify the comm paths that are present:

```
net_list
```

- b. Remove these comm paths:

```
net_remove
```

8. Write the changes.

[lcdsync](#) [-d destname]

Example:

```
c:\LK\Bin>lcdsync
```

This function checks to see if the SIOS Protection Suite resource hierarchy configuration and communication path status data stored in shared memory has been modified. If it is different, the data is "synchronously" written to disk. Therefore, when this program returns, the data is guaranteed to be on disk properly.


Optional Note: To completely remove all resources, first generate a list of all resource instances:

```
ins_list [-d destsys] [-fC] [-R top] [-a appname] [-r typ] [-t tag] [-i id]
```

then remove these instances:

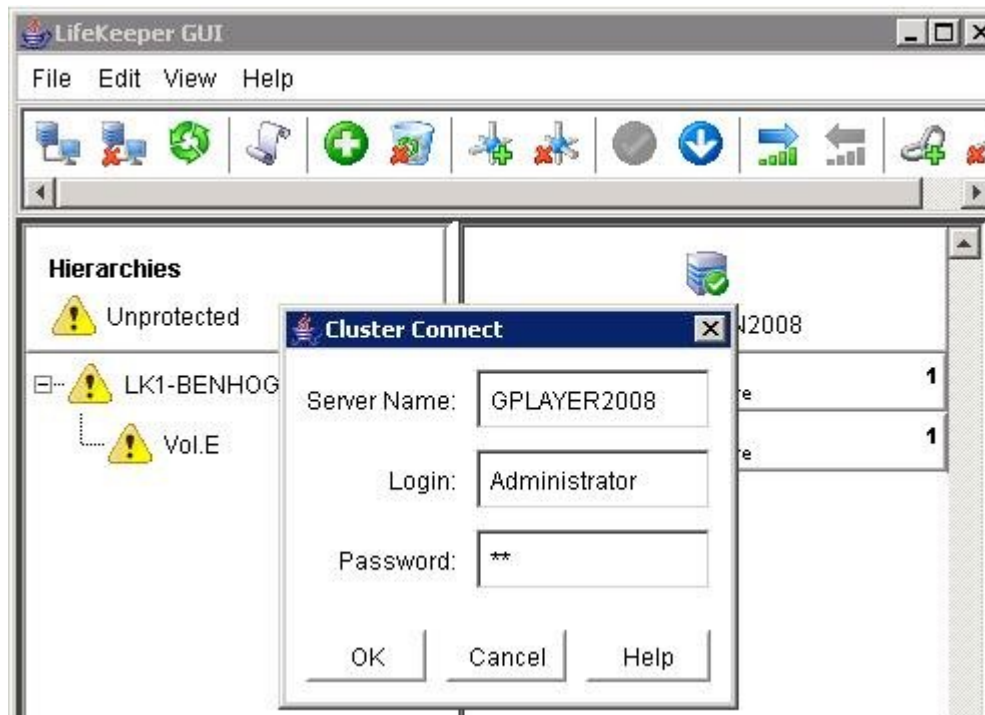
```
ins_remove [-d destsys] [-R roottag] [-a appname] [-r restyp] [-t tag] [-i id] [-v] [-I] [-N] [-G]
```

9. Close the LifeKeeper GUI and reopen.




10. Extend each resource hierarchy from the server where the resource hierarchy is in service to the re-installed server using the GUI.

a. Connect to your reinstalled server/target. **Note:** This may also require a re-creation of the comm paths.



b. Extend resource hierarchy.



c. Add/extend mirror resource.



Complete — All protected resources have been recovered and are protected once again.

Recovering Out-of-Service Hierarchies

Description

As a part of the recovery following the failure of a SIOS Protection Suite 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.

Remove Hangs During Recovery Kit Uninstall

Symptom

If a recovery kit is uninstalled while there are resource hierarchies of that kit in service, the **Remove** hangs.

To avoid this situation, it is recommended to always take a recovery kit's resource hierarchies Out of Service and delete them before uninstalling the recovery kit software.

Solution

If you encounter this situation, you will most likely need to re-boot your system since there are many related processes that hang, and clearing them all can be difficult.

Replicated Volume Switchover Failure

Description

If the **DELETEMIRROR** command fails during switchover of a replicated volume, the following error message will display in the Output panel:

***WARNING* (No. 12104) Delete mirror action for volume R: failed with status 51 (command=C:/SDR/EMCmd10.10.1.2 DELETEMIRROR R: 10.10.1.1)**

Refer to the table below for the Error Status number, description and recommended action. You can also use the following command to obtain more information about the error status:

Command: net helpmsg {status}

Example: net helpmsg 51

Error Status	Description	Recommended Action
5	Permission issues on the current SOURCE are not allowing the mirror to be deleted.	Check both systems for permission differences that may exclude the Local System Account from accessing the mirrored volume.
46	Mismatched user/password combination between systems. This will probably only occur if running DELETEMIRROR from the command prompt.	Use a domain account or make sure the local user accounts you are signing on with have the same password.
51	Windows cannot find the network path.	Make sure all network cards in both systems have File & Print Sharing for Microsoft Networks checked.
53	Cannot access the IP Address specified.	Verify your network configuration (including HOSTS files and DNS) are all resolving the IP address consistently.
207	The ring 2 stack is in use.	Make sure all the network cards in both systems have Client for Microsoft Networks checked.

Resource Tag Name Restrictions

Tag Name Length

All tags within SIOS Protection SuiteDataKeeper may not exceed the 256 character limit.

Valid "Special" Characters

`-_./`

However, the first character in a tag should not contain "." or "/".

Invalid Characters

`+ ; : ! @ # $ * = "space"`

Restore and Health Check Account Failures

Symptom

Some recovery kits monitor protected resources by performing query operations that simulate user and/or client activity. This provides SIOS Protection Suite with accurate status information about a protected application or service. It also requires that a valid user account ID and password with login privileges be provided during resource object creation. If the user account does not have login privileges on a particular system, the following error message will be recorded in the **Windows Application Event Log**:

Error Number 1385 - "Logon failure: the user has not been granted the requested logon type at this computer."

Solution

Have the domain administrator provide login privileges for the user account. Also, most recovery kits that require an ID and password have a resource properties or configuration tab available for administrators to change the account information for the resource. Right-click on the resource object and select the appropriate properties or configuration tab. If the resource does not have an account update feature, the resource object must be deleted and a new one created with updated account information.

SQL 2008

Symptom

When using SQL 2008, after a switchover or a failover, the variable `servername` still points to the primary system.

Solution

You can use `"select SERVERPROPERTY('ServerName')"` instead of using the variable `@servername`. This query will return the correct name of the machine after a switchover or failover.

or

1. Execute the following commands on the new backup server:

```
sp_dropserver @server='sys-A'
```

```
sp_addserver @server='sys-B', @local='LOCAL'
```

2. Restart the service.

SQL Server Reporting Services (MSSQLSERVER)

Symptom

When protecting SQL Server 2008 R2 services, the "SQL Server Reporting Services (MSSQLSERVER)" may be selected as an optional protected service. However, if the time required to start this service exceeds the default Windows service timeout, you may get Error 1053, the service may fail to start and the SIOS Protection Suite resource in-service operation will fail.

Solution

This problem may be related to system performance and configuration issues. The recommended action is to not protect this service. However, if it must be protected, the following registry setting will extend the time available for services to start. In the "HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control" Registry key, add a "ServicesPipeTimeout" value as a DWORD, and set the value to 60000 (decimal, = 60 secs.).

Two-Server Cluster Issue

Symptom

In a two-server cluster, when the primary server fails or is shut down which causes the hierarchies to fail over to the backup server, and the backup server also fails or is shut down before the hierarchies are entirely failed over to the backup server, the following behavior has been detected:

When both servers are rebooted, some of the resources in the hierarchies will be in service on one server and some will be in service on the other server. Some of the higher-level parent resources may not be in service on either server.

Solution

After both servers have been restarted and have completed SIOS Protection Suite initialization, select the parent resource in a hierarchy that did not come in-service from the Hierarchy Administration interface and bring it in-service manually. Repeat this task until all hierarchies are in-service.

Unknown User Name or Bad Password

Access Denied: Unknown User Name or Bad Password

Symptom

If a SIOS Protection Suite client tries to communicate with a server that is in the process of shutting down, the server may abort the validation process by refusing to allow the client to log on. In that case, the client will display a message stating "Access Denied: unknown user name or bad password. Only members of the SIOS Protection Suite-authorized security groups can use SIOS Protection Suite. Would you like to re-enter the authentication data?"

Solution

Click **Yes** to input new credentials, and then click either **Cancel** or re-enter the credentials and click **OK**.

***Note:** If you click **No** initially, the LifeKeeper GUI will disconnect from that server and will not reconnect automatically.

Web Client Troubleshooting

Background

The web client does not display the [Cluster Connect dialog](#).

Answer

If the web client does not display the [Cluster Connect dialog](#), try the following:

1. Check whether the applet failed. Usually a message is printed somewhere in the browser window specifying the state of the applet. In Internet Explorer, an icon may appear instead of the applet in addition to some text status. Clicking this icon may bring up a description of the failure.
2. Open the **Java Console**.
 - For **Firefox** and **older versions of Internet Explorer**, run the Java Plug-In applet from your machine's Control Panel and select the option to show the console. Restart your browser.
 - For **recent versions of Internet Explorer**, select **Tools > Sun Java Console**. If you do not see the Sun Java Console menu item, select **Tools > Manage Add-Ons** and enable the console, after which you may need to restart your browser before the console will appear.
 - For **Mozilla**, select **Tools > Web Development > Sun Java Console**.

If the web client is not open, reopen the URL, **http://<server name>:81**, to start it.

Check the console for any messages. The messages should help you resolve the problem. If the problem appears to be network related, refer to the [Network-Related Troubleshooting](#) section.

IIS Resource Hierarchy Creation Error

Symptom

When creating an IIS Resource Hierarchy, you receive the following message:

No qualified sites were found...

and the create fails.

BACKGROUND / TROUBLESHOOTING:

Run the following command:

```
C:\LK\Admin\kit\webapp\bin>enumiis query all
```

Look for the following error message:

ERROR: CoCreateInstance Failed! Error: -2147221164 (80040154)

ERROR: W3Service Com Object Failed to initialize

Solution

The IIS 6 Management Compatibility (Role Service) is required for the SIOS Protection Suite IIS Kit. You should install all of this option (Metabase Compatibility, WMI Compatibility, Scripting Tools, Management Console).

Application Recovery Kits

SIOS Protection Suite for Windows Application Recovery Kits (ARKs) include tools and utilities that allow SPS to manage and control a specific application. The following optional recovery kits are available with this release of SPS.

[SIOS Protection Suite Microsoft SQL Server Recovery Kit](#)

[SIOS Protection Suite PostgreSQL Server Recovery Kit](#)

[SIOS Protection Suite Oracle Recovery Kit](#)

[SIOS Protection Suite Microsoft Internet Information Services Recovery Kit](#)

SIOS Protection Suite Microsoft SQL Server Recovery Kit Introduction

The SIOS Protection Suite Microsoft SQL Server Recovery Kit software lets you tie the data integrity of Microsoft SQL-based databases to the increased availability provided by SIOS Protection Suite for Windows.

The LifeKeeper GUI allows you to easily create a SQL resource hierarchy. SIOS Protection Suite can then protect all of the disk resources used by the SQL Server instance, as well as the IP socket resources used to access the database.

***Important Note:** This kit is incompatible with the following SQL features: SQL Replication (Snapshot, Merge and Transactional), SQL Log Shipping, SQL Database Mirroring and SQL Server AlwaysOn Availability Groups.

SQL Server Services

The SIOS Protection Suite Microsoft SQL Server Recovery Kit will monitor and protect the following services:

SQL2008 (R2) , 2012, 2014, 2016 or 2017:

Core Services	Optional Services
SQL Server (MSSQLSERVER)	SQL Server Agent
	SQL Server Reporting Services
	Distributed Transaction Coordinator
	SQL Server Browser
	SQL Server VSS Writer
	SQL Server Integration Services
	SQL Full-text Filter Daemon Launcher (available starting in SQL Server 2014)
	SQL Server Launchpad (available starting in SQL Server 2016)
	SQL Server PolyBase Data Movement (available starting in SQL Server 2016)
	SQL Server PolyBase Engine (available starting in SQL Server 2016)
	SQL Server CEIP service (available starting in SQL Server 2016)

Note: "SQL Server Integration Services" is only displayed in the list of optional services to protect when this optional SQL feature is installed.

All data files are stored on shared or replicated volumes. Thus, upon detecting a failure, SIOS Protection Suite switches the database along with its associated volumes and IP socket resources to a backup server. Once SIOS Protection Suite switches all dependent resources to the backup server, it starts the Microsoft SQL service and any protected optional services.

Recovery Kit Requirements

Before installing and configuring the SIOS Protection Suite Microsoft SQL Server Recovery Kit, be sure that your configuration meets the following requirements:

- SIOS Protection Suite supports the versions of Windows operating systems listed in the [Operating System](#) section of the SIOS Protection Suite for Windows Release Notes.
- **SIOS Protection Suite software.** You must install the same version of SIOS Protection Suite for Windows on all servers in the cluster. If you plan to use Microsoft SQL Server with replicated volumes rather than shared storage, make sure you install the SIOS DataKeeper software on each server.
- **Microsoft SQL Server RDBMS software.** Refer to the [SIOS Protection Suite for Windows Support Matrix](#) for a list of supported Microsoft SQL versions. The same version of Microsoft SQL Server must be installed on all systems in the cluster.
- **Communication protocol.** TCP/IP is strongly recommended by Microsoft for use in a clustered environment. Although SIOS Protection Suite supports LAN Manager, this document will assume you are using TCP/IP and will refer to switchable IP resources (rather than LAN Manager resources) in its configuration instructions.

Consult your SIOS Protection Suite sales representative for release and ordering information.

SQL Server Installation

Proper operation of the SIOS Protection Suite Microsoft SQL Server Recovery Kit depends upon correct setup of the hardware and software.

Before continuing, please preview the [Hierarchy Administration](#) section of this guide. This section provides general guidelines, configuration details and troubleshooting hints to help you administer Microsoft SQL Server in a SIOS Protection Suite environment.

Recovery Kit Installation

The SIOS Protection Suite Microsoft SQL Server Recovery Kit is distributed via ftp download. Installation is simple and quick using InstallShield to provide a standard installation interface.

Before installing the SIOS Protection Suite Microsoft SQL Server Recovery Kit software, be sure you are familiar with the [product prerequisites](#). A SIOS Protection Suite Microsoft SQL Server Recovery Kit license key must be installed in order to protect a SQL resource using SIOS Protection Suite.

Removal

To remove the SIOS Protection Suite Microsoft SQL Server Recovery Kit software, choose **Microsoft SQL Server Recovery Kit** in the **Add/Remove Programs** or **Programs and Features** applet in the control panel.

!CAUTION: Be sure there are no SQL instances or resources in service when the kit is removed. Once the kit is removed, these resources will be unusable. All SQL hierarchies should be deleted before the kit is removed.

Installing and Configuring SQL Server with SIOS Protection Suite

Proper operation of the SIOS Protection Suite Microsoft SQL Server Recovery Kit depends upon correct setup of the hardware and software.

This section provides general guidelines, configuration details and troubleshooting hints to help you administer Microsoft SQL Server in a SIOS Protection Suite environment. Please remember to review the [Hierarchy Administration](#) tasks.

Before Installing SQL Server

Before you install the SQL Server software, the servers and storage must be configured and SIOS Protection Suite must be installed on each server in the cluster.

Installation – Shared Storage Systems

On the Primary Server

1. Power down the backup server so that there is no chance of simultaneous access to your shared storage.
2. Use the **Windows Disk Management** tool to configure your disk resources and define the shared volumes that you want to use. (Be sure the volume size is adequate.)
3. It is recommended that you use **Windows Explorer** to unshare from the network all volumes to be used by the SQL Server Instance.
4. Configure your networking to support the SIOS Protection Suite TCP/IP comm path(s) and, if applicable, the switchable IP address.
5. Install the SIOS Protection Suite Core software on a local disk, followed by the SIOS Protection Suite SQL Server Recovery Kit.

On the Backup Server

1. Bring up the backup server and use the **Disk Management** utility to assign the same drive letter to the shared volume as assigned on the primary server.
2. Install the SIOS Protection Suite Core software on a local disk, followed by the SIOS Protection Suite SQL Server Recovery Kit.

On the Primary Server

Now that you have SIOS Protection Suite installed on both servers, go back to the primary server and do the following:

1. In SIOS Protection Suite, create comm paths between the primary and backup servers.
2. In SIOS Protection Suite, create your volume resource and IP communication resource and extend them to the backup server. Later

when you create your SQL Server resource hierarchy, SIOS Protection Suite will automatically bring these resources into the hierarchy as dependencies.

Install the SQL Server Software

1. If using shared volumes, bring the volume resource hierarchy In Service on the backup server using the LifeKeeper GUI.
2. On the backup server, install Microsoft SQL Server using the following guidelines:
 - Install the database engine, along with any additional features. Configure the instance data and log files so that they are stored on volumes that are protected by the SIOS Protection Suite.
 - Select "**Mixed Mode**" database authentication, and enter a non-blank password for the SA account. The passwords **MUST** be the same on all servers in the cluster.

When the installation is complete, use* Microsoft SQL Server Configuration Manager* to verify that SQL Server can start properly on the backup server. Stop all Microsoft SQL Services on the backup server.

For shared volumes, do the following steps.

1. Bring the volume resource hierarchy In Service on the primary server.
2. On the primary server, open **Explorer** and access the drive associated with the shared volume.
3. Delete the directory where you previously installed the SQL data files. (You will re-install them in the next step).
4. Install Microsoft SQL Server on the primary server EXACTLY as you did on the backup server (program files on the local disk

and data files on the shared volume).

When the installation is complete, use **Microsoft SQL Server Configuration Manager** to verify that SQL Server can start properly on the primary server.

Using a Non-Admin Local System Account on Target System

There are certain cases where the SQL Server service account (sql_svc) cannot be added to the local admin or domain admin groups. In such cases, there will be permission problems on the files for that service account. Users will notice "Access Denied" errors like the following message:

"Open failed: Could not open file E:\SQL1\MSSQL10_50.SQL1\MSSQL\DATA\master.mdf for file number 1. OS error: 5 (Access is denied)."

Solution: Perform the following:

1. Copy the following script code to a file, e.g. c:\file.ksh

```
$LKROOT/bin/find . > $LKROOT/out/file
while read filename
do
icaccls "${filename}" /grant $1:F
done < $LKROOT/out/file
$LKROOT/bin/rm $LKROOT/out/file
```

2. Run the following command on each drive (E, F), starting at the volume. This command should be run as an admin (local or domain). This gives file permissions to that user only. SQL Server service user will then have access to the files and will not need to be added as an admin account.

```
E:\>c:\lk\bin\sh c:\file.ksh domain\sql_svc
```

***Note:** In this example c:\lk is where SIOS Protection Suite is installed, and domain\sql_svc is the name of the SQL Server service userid.

Installation – Replicated Storage Systems

On the Primary Server

1. Use the **Windows Disk Management** tool to configure your disk resources and define the replicated volumes that you want to use. (Be sure the volume size is adequate.)
2. It is recommended that you use **Windows Explorer** to unshare from the network all volumes to be used by SQL Server.
3. Configure your networking to support the SIOS Protection Suite TCP/IP comm path(s) and, if applicable, the switchable IP address.
4. Install the SIOS Protection Suite Core software on a local disk followed by the SIOS Protection Suite SQL Server Recovery Kit.
5. Install the SIOS DataKeeper software to the local disk now. Refer to the [SIOS Protection Suite for Windows Installation Guide](#) for details.
6. Using the LifeKeeper GUI, create comm paths between the primary and backup servers.
7. In SIOS Protection Suite, create your IP communication resource and extend them to the backup server. Later when you create your SQL Server resource hierarchy, SIOS Protection Suite will automatically bring the resource into the hierarchy as a dependency.

***Note:** When the SQL Server Hierarchy is created, the SIOS DataKeeper resource will automatically be created and brought into the SQL Server resource hierarchy as a dependency.

On the Backup Server

1. Bring the volume resource hierarchy In Service using the **LifeKeeper GUI**.
2. Install Microsoft SQL Server using the following guidelines:

- Install the database engine, along with any additional features. Configure the instance data and log files so that they are stored on volumes that are protected by the SIOS Protection Suite.
 - Select "Mixed Mode" database authentication, and enter a non-blank password for the SA account. The passwords MUST be the same on all servers in the cluster.
3. When installation is complete, use **Microsoft SQL Server Configuration Manager** to verify that SQL Server can start properly on the backup server. Stop all Microsoft SQL Services on the backup server. **Note:** For replicated volumes, you may wish to move the SQL *tempdb* database to a volume that is not protected by SIOS Protection Suite to improve performance.
 4. Install Microsoft SQL Server on the primary server EXACTLY as you did on the backup server (program files on the local disk and data files on the replicated volume).
 5. When installation is complete, use **Microsoft SQL Server Configuration Manager** to verify that SQL Server can start properly on the primary server.

***Note:** If the data files are installed to a replicated volume, you may wish to move the SQL *tempdb* database to a volume that is not protected by SIOS Protection Suite to improve performance.

On the Primary Server

1. Bring the **Communication resource** into service on the primary server.
2. Start the **SQL Server Services** on the primary server.
3. Create the **SQL Server hierarchy** on the primary server and extend it to the backup server. During the **Extend Volume Resource** process, choose **Create Replicated Mirror** then select **Next** to complete the wizard and finish the configuration. See [Creating the SQL Hierarchy](#) for details.

Test the new **SQL Server hierarchy** by performing a manual failover.

Additional Setup Tasks for Extended Configurations

If your configuration uses a shared storage device or you are using SIOS DataKeeper, you may choose a configuration that will be extended to a third (or more) server(s).

1. Configure two systems following the steps given in [Installing and Configuring SQL Server with SIOS Protection Suite](#).
2. Switch your protected volumes to the third server.
3. Install the Microsoft SQL Server software on the local drive and the master database on the same shared/replicated volume as used by the other servers. This will permit you to extend the hierarchy and utilize SIOS Protection Suite's cascading feature.

Creating the SQL Hierarchy

After you have completed the necessary setup tasks outlined in the [SIOS Protection Suite for Windows Installation Guide](#), use the following steps to define the SQL Server hierarchy to protect your database(s).

!Important: If you have an existing SQL database installed, you must close any client applications (local or remote) that are accessing the SQL database prior to completing this procedure. During the automated move operation, the database will need to be started and restarted numerous times and running applications may interfere with the commands following each service action. Restart the applications once this procedure has been completed.

1. From the LifeKeeper GUI menu, select **Edit** then **Server**. From the menu, select **Create Resource Hierarchy**.
2. The **Create Protected Application** dialog box will display. Select the **Primary** and **Backup** servers from the pull-down list. Select **Next** to continue.
3. The dialog box will appear with a drop down list box displaying all recognized recovery kits installed within the cluster. Select **MS SQL Server** and click **Next**.
4. 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
Select Microsoft SQL Server Instance	Select the instance of Microsoft SQL Server you wish to place under SIOS Protection Suite protection. SIOS Protection Suite will read the configuration data for this instance and pull the associated volumes into the hierarchy.
Enter Microsoft SQL Administrative	Enter the administrative user name that is used for Microsoft SQL on this server. This user account must include SA permissions to the master database.

User Name	
Enter Password	Enter the administrative password for the user account just entered.
Verify Current Locations	Displays the current location of database files. If any of the detected files for the specified instance to be protected by SIOS Protection Suite are located on the System Drive (c:), they will be highlighted in the table displayed on the screen.
The following fields only appear if an existing database needs to be relocated.	
Select Destination for Database Relocation	When database migration is required, specify the volume destination to relocate affected databases.
Verify the Move Operation	Verify the location and confirm the intent to move the specified databases.
Relocating the Databases	An action window displays showing the progress of the databases being relocated during the move operation.
Select Optional Services for Protection	Select optional SQL services to be protected in this hierarchy. The list includes only those services eligible for SPS protection.
Protected IP Address	Select an IP address to protect with this instance. IP Address is not required if only named pipes are used (though this is NOT recommended).
Named Pipe Alias	Named Pipe Alias
Microsoft SQL Server Resource Name	Enter a unique tag name, or you can accept the default tag name offered by SPS. Note: The tag name must consist of printable ASCII characters.

- After you click **Create**, the **Wizard** will create your SQL resource. SIOS Protection Suite will validate the data entered. If SIOS Protection Suite detects a problem, an error message will appear in the information box.
- Another information box will appear indicating that you have successfully created a SQL resource hierarchy, and you must **Extend** that hierarchy to another server in your cluster in order to achieve failover protection. Click **Next**.

7. After you click **Continue**, SIOS Protection Suite will launch the **Pre-Extend Wizard**.

Extending a SQL 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. 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.
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
Backup Priority	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. SIOS Protection Suite 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.

4. Click **Extend**.

!Important: After migrating the database using the automated tool, you should verify that you can access your SQL application and database files. All files are relocated using the system copy utility. After you have validated the success of this procedure, you can remove these data and log files.

Unextending a SQL Hierarchy

To remove a resource hierarchy from a single server in the SIOS Protection Suite 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 SQL resource. It cannot be the server where the SQL 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 SQL 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 SQL resource hierarchy you have chose to unextend. Click **Unextend**.
5. Another information box appears confirming that the SQL resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

Deleting a SQL Hierarchy

Before deleting a SQL hierarchy or instance, make sure that the hierarchy is active (green) on its primary server. You may also wish to remove the dependencies before deleting the hierarchy; otherwise, the dependencies will be deleted also.

To delete a resource hierarchy from all the servers in your SIOS Protection Suite 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 SQL 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 **Next**.
5. Another information box appears confirming that the SQL resource was deleted successfully.
6. Click **Done** to exit.


SQL Server Configuration Considerations

Before you install and configure your clusters, it is important to understand the concepts of Active/Standby configuration, and how multiple instances can be set up in a SQL configuration.

SQL Configuration


A configuration is Active/Standby when there is only one master database for each SQL Server, and it is located on a shared or replicated volume. The services run on only one system at a time. The servers are assigned priorities within SIOS Protection Suite which determine the order of failover for a particular hierarchy.

The figure below depicts a single SQL instance installed on a pair of servers. The instance contains two databases, databaseK and databaseX residing on separate volumes. Note that there is a single master database which resides on shared volume X.




When you create the SQL hierarchy within SIOS Protection Suite, you are asked to specify the SQL instance to be protected, and the IP resource that will be used to connect to the database. SIOS Protection Suite then reads the configuration data for that instance and pulls the associated volumes into the hierarchy.

Once the hierarchy is created, it will appear as follows in the LifeKeeper GUI.



Failover

In the event of failure, SIOS Protection Suite brings the SQL Server hierarchy In Service on the backup Server. SQL Server is started on the backup server and it takes over protection of all defined databases as depicted in the figure below.

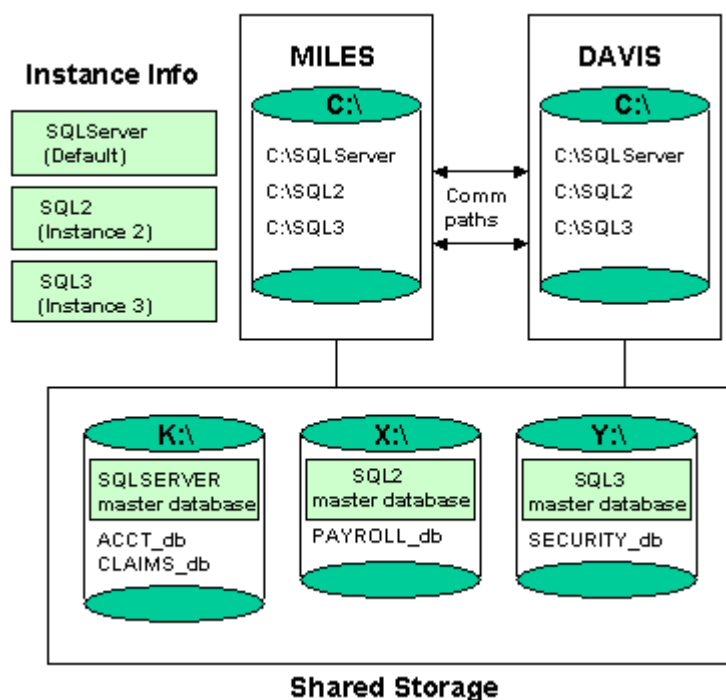


Multiple SQL Instances

SQL Server can be installed multiple times, which creates multiple SQL instances. SIOS Protection Suite can protect multiple instances of SQL Server. SIOS Protection Suite identifies each instance by the unique name given during SQL installation.

One SQL instance may contain multiple SQL databases. Each instance is protected in a single SIOS Protection Suite hierarchy. Thus, if the SQL instance contains two databases, the corresponding SIOS Protection Suite hierarchy will protect two databases (along with the associated IP and volume resources).

The figure below depicts three SQL instances: SQLServer (the default instance), SQL2, and SQL3. These are installed on a pair of servers, MILES and DAVIS.



Notes:

- The databases are located on three different shared storage volumes, K, X and Y. Note that the default instance contains two databases and the other two instances contain one database each.
- Each server can be the primary and backup server for multiple instances.
- It would be possible for MILES to be the primary server for the

default instance and DAVIS to be the primary server for the SQL2 and SQL3 instance.

Managing a SQL Server Configuration

To administer a protected SQL Server resource from the LifeKeeper GUI, right-click on the SQL Server resource (on the right-hand side of the LifeKeeper GUI) and select **properties**, then select the **SQL Server Configuration** tab. Use the **SQL Server Configuration** page to view or change information about your SQL resource.

Resource Properties for SQL.TS

Properties Panel

Select Resource: **SQL.TS**

Select Server for Resource: **SPARTA**

SQL Server Configuration | General | Equivalencies | Relations

LifeKeeper Resource Information

MS SQL Version: 9

Protected Instance: TS

Optional Protected Services

Service Name	Display Name
SQLAgent\$TS	SQL Server Agent (TS)
SQLBrowser	SQL Server Browser

Optional Protected Databases

Database Name	LifeKeeper Monitored
master	Yes
gatest	Yes
test1	Yes
model	No
msdb	No
test_2	No
test-3	No
1bigC	No

Process Config | Database Config | Admin Actions

OK | Apply | Cancel | Help

Process Config Menu

This menu allows users to modify the list of optional SQL Processes that are protected under the resource hierarchy. SIOS Protection Suite will monitor all protected optional services (see [Monitoring Your SQL Hierarchy](#)).

Select **Action**:

- *Add Process* - Add an additional process to the protected configuration. SIOS Protection Suite will start monitoring new SQL service
- *Delete Process* - Remove a process from the protected configuration.

Field	Tips
Service Name	Enter the service name for the process to Add or Delete from the protected configuration.
Update All Systems	Select Yes to update all systems in this cluster. Otherwise, select No to only update the current system. If you choose No , you must manually add the process to the backup servers.

Database Config Menu

This menu allows users to modify the list of optional SQL Databases that are protected under the resource hierarchy. SIOS Protection Suite will monitor all protected optional SQL Databases by performing a SQL query to test the connection to the database (see [Monitoring Your SQL Hierarchy](#)).

Select **Action**:

- *Add Database* - Add an additional database to the protected configuration. **Note:** The option is available only on the server where the SQL resource is active (In Service).
- *Delete Database* - Remove a database from the protected configuration.

Field	Tips
Enter Database Name	Enter the database name for the process to Add or Delete from the protected configuration.
Update All Systems	Select Yes to update all systems in this cluster. Otherwise, select No to only update the current system. If you choose No , you must manually add the process to the backup servers.

Admin Actions Menu

This menu allows users to manage the SQL administrator user used during SIOS Protection Suite operations or resolve ID conflicts between the primary and backup servers encountered during the extend operation.

Select **Administrative Action**:

- *Manage User* - Display or update the current user name used by the protected resource hierarchy.

Manage User Menu

Select **Management Action**:

- *Show Current User* - Display the current user name used by the protected resource hierarchy.
- *Change Password* - Update the user password for the current user associated with the protected resource hierarchy.
- *Change User and Password* - Update both the user and password to be used during SIOS Protection Suite operations to administer and monitor the SQL instance. The user must have sql admin privileges for all databases under protection.

Note: The SQL instance must be running to change the user and/or password.

Field	Tips
Enter User Name	Enter the administrative user name. This user account must include SA permissions to all databases under SIOS Protection Suite protection.
Enter Password	Enter the administrative password for the user account being updated.

Testing Your SQL Resource Hierarchy

You can test your SQL 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.

SQL Server Hierarchy Administration

Follow these guidelines when administering your SQL Server.

[Access Via Protected Communication Paths](#)

[User Resource Name For Remote Access](#)

[Reserve Volumes For Exclusive SQL Use](#)

[Understand Manual Switchover Limitations in SQL Server Environment](#)

[Running Microsoft SQL Management Tools](#)

[Start and Stop SQL Server Only Through SIOS Protection Suite](#)

[Adding Microsoft SQL Server Volumes](#)

[Recovering From Databases in Suspect State After a Switchover](#)

[Pausing Microsoft SQL Server](#)

[Configuring SQL Server to Connect Using the Switchable IP Address](#)

[Maintaining SQL Server Login and Passwords](#)

[Monitoring Your SQL Hierarchy](#)

Access Via Protected Communication Paths

All remote access of the service should be done through the hierarchy's protected IP addresses. This will ensure that users can access the SQL service regardless of which server it is currently running on.

User Resource Name For Remote Access

Unless the application is cluster-aware, when using **Microsoft SQL Enterprise Manager** to administer the service, you should register it by the switchable resource name (the name by which users access the server using TCP/IP). This gives you a continuous monitor of the viability of this path.

If you register the SQL Server by the system name, you can also monitor the system.

Reserve Volumes For Exclusive SQL Use

The volumes containing the protected SQL files should be reserved for use by SQL exclusively.

A SIOS Protection Suite protected volume may fail to switch over if it is accessed by an application, process or remote user.

Understand Manual Switchover Limitations in SQL Server Environment

Any manual action requires that all users be logged off of the SQL Server resources.

Local processes that have read-only access to volumes do not prevent removal of a resource from service but may cause a restore to fail when you try to switch back. Examples might be the Performance Monitor, which periodically polls each volume, or any running process which is installed on the shared or replicated volume. You can minimize your potential for this type of restore failure by installing the Microsoft SQL Server on local drives and putting only the database on shared or replicated volumes.

Running Microsoft SQL Management Tools

Open the **Microsoft SQL Server Configuration Manager** only when needed and do not run it constantly.

If the **Microsoft SQL Server Configuration Manager** is open and active at the database level, it may prevent the SQL hierarchy from coming into service properly and the failover will not complete successfully. If this occurs, close the Microsoft SQL Server Configuration Manager and manually bring the SQL resource into service.

Start and Stop SQL Server Only Through SIOS Protection Suite

Although much of your administration of the Microsoft SQL Server is done through the Microsoft SQL Enterprise Manager, you derive two distinct benefits from bringing the Microsoft SQL Server in service and out of service using the SIOS Protection Suite administration options:

1. **Consistent view.** When SIOS Protection Suite stops and starts Microsoft SQLServer, it maintains a consistent view of the server on all nodes in the configuration.
2. **Configuration details saved.** If you change your Microsoft SQL Server configuration, you can stop and start the server through SIOS Protection Suite or perform a manual switchover and SIOS Protection Suite automatically replicates the configuration changes on the paired node.

***Note:** If you do not use these options to replicate new configuration information on the paired server, the backup server will use old configuration information in a failover situation.

3. **Protected Microsoft SQL services** should be set to **Manual** startup mode through the **Control Panel "Services"** tool.

Adding Microsoft SQL Server Volumes

As your environment grows, you may need to add new Microsoft SQL Server databases on new shared or replicated volumes. You should perform the following tasks to add the new volumes to the SIOS Protection Suite hierarchy before administering the new databases within Microsoft SQL Server.

To add a new volume resource to an existing SQL Server hierarchy:

1. **Create the resource.** On the server where the SQL Server hierarchy is in service, create and extend the volume resource with the same priority order as that of the SQL hierarchy.
2. **Create dependency.** Right click on the SQL Server resource, and then select Create Dependency from the pop-up menu. For **Child Resource Tag**, select the new Volume resource.

When you have completed these SIOS Protection Suite tasks, you can perform the administration tasks to add the SQL Server database. Adding new databases to volumes that are already part of the resource hierarchy requires no SIOS Protection Suite specific administration.

Recovering From Databases in Suspect State After a Switchover

A database which gets marked as suspect may be caused by starting Microsoft SQL Server when the volume(s) on which it resides is unavailable to this system. When SIOS Protection Suite is used to protect SQL databases, the starting and stopping of SQL should be performed only as a function of bringing hierarchies in or out of service. In those situations where a database has been marked suspect and it is known that the database is fine, perform the following steps to correct the problem.

For those databases which are suspect on a primary/secondary server:

1. Make sure that the volume(s) where the database resides is actively (green) being protected by this server.
2. Use `sp_resetstatus` to change the state of the database. Execute the following commands from a query window to reset the status of the suspect database:

```
While in master database, execute sp_configure 'allow updates',  
1
```

```
Reconfigure with override
```

```
Sp_resetstatus 'dbname'
```

```
Sp_configure 'allow updates', 0
```

```
Reconfigure with override
```

3. **Stop** Microsoft SQL Server.
4. **Start** Microsoft SQL Server.

Pausing Microsoft SQL Server (MSSQLServer)

It is possible for the SQL Administrator to manually put Microsoft SQL Server into a PAUSED state whereby existing connections to the Microsoft SQL Server can continue processing, but no new connections are allowed. In this case, SIOS Protection Suite detects that the MSSQLServer service is not RUNNING, but will NOT attempt to restart the service locally or fail the SQL hierarchy to the backup server. Neither option is the appropriate action, so monitoring of the SQL resource is essentially forfeited when Microsoft SQL Server is in the PAUSED state.

Because manual intervention was required to put Microsoft SQL Server into this state, the SQL administrator must manually move SQL Server out of this state. Once out of the PAUSED state, SIOS Protection Suite can resume monitoring the SQL resource as outlined above.

Configuring SQL Server to Connect Using the Switchable IP Address

By default, TCP/IP sockets are set as the default network protocol when SQL is installed. If this setting is modified at any time, use the **SQL Server Configuration Manager** tool to re-enable the TCP/IP network setting.

Maintaining SQL Server Login and Passwords

During the creation of a SIOS Protection Suite SQL resource, the user must enter a SQL administrative username and password for that instance of Microsoft SQLServer. Should the password of this username change at some point in the future, the SQL resource must be updated on all systems in the cluster with this new password. Failure to do so will prevent SIOS Protection Suite from fully monitoring the status of SQL Server on these systems. A warning message will be logged in the Application Event Log, but the SQL resource will not fail as a result of this.

The SQL administrative username and password associated with the SIOS Protection Suite SQL resource can be changed using the **SQL Server Configuration** tab in the LifeKeeper GUI. To administer a protected SQL Server resource from the LifeKeeper GUI, right-click on the SQL Server resource (on the right hand side of the LifeKeeper GUI) and select **properties**, then select the **SQL Server Configuration** tab. Click the **Admin Actions** button and select **Manage User** on the **SQL Server Configuration** page to view or change information about your SQL resource.

See the section [Managing a SQL Server Configuration](#) for more details on the **SQL Server Configuration** page.

Monitoring Your SQL Hierarchy

For every SQL resource, SIOS Protection Suite monitors the Microsoft SQL Server service and any optional services selected during the creation of the SQL hierarchy or added at a later time. Should any of these services stop, the monitoring process associated with the SQL resource will detect this and attempt to restart the service on the local server if Local Recovery is enabled. If Local Recovery is disabled, the resource will fail over to the backup server.

SIOS Protection Suite will also perform a SQL query to test the connection to all protected SQL databases. If the SQL query fails to the master database, the monitoring process associated with the SQL resource will detect this and attempt to restart the services, if Local Recovery is enabled. If Local Recovery is disabled, the resource will fail over to the backup server. For other protected databases added after creation of the SQL hierarchy, an error will be logged to the **Application Event log**.

Troubleshooting

This section is intended to provide suggestions and insights into occurrences that are not specifically related to the SIOS Protection Suite software but have a relationship with the total environment.

[Create and Extend Failure](#)

[Extend of a SQL Resource Fails](#)

[SQLAgent Service Fails to Start](#)

[Connecting ODBC Clients to Named Instances of SQL Server](#)

Create and Extend Failure

Create and Extend Fail if Master or User Databases on System Drive

Symptom:

The **Create** and **Extend** of a Microsoft SQL Server resource will fail if the Master or any user databases are located on the system drive or a volume that cannot be protected by SIOS Protection Suite. SIOS Protection Suite does not require that the tempdb database be located on a SIOS Protection Suite protected volume as this database is recreated each time SQL Server starts.

Suggested Actions:

The following web sites provide information on moving the Master and user databases to a volume that can be protected by SIOS Protection Suite.

<http://support.microsoft.com/kb/224071/en-us>

<http://www.databasejournal.com/features/mssql/article.php/3379901>

Extend of a SQL Resource Fails

Symptom:

The Extend of a Microsoft SQL resource will fail if the primary and backup configurations are different (i.e. Microsoft SQL ID mismatch). Below is an example of the error message displayed in the LifeKeeper GUI during the can extend operation:

```
Process: canextend.ksh(1292)
```

```
*ERROR* (No. 14003) The target value for Database file location does not match the template value (target=- DQ:\SQLDEFAULT\MSSQL.1\MSSQL\DATA\MASTER.MDF,template=DQ:\SQLDEFAULT\MSSQL.4\MSSQL\DATA\MASTER.MDF)
```

```
Process: canextend.ksh(1572)
```

```
*ERROR* (No. 14003) The target value for Temp DB location does not match the template value (target=-EQ:\SQLDEFAULT\MSSQL.1\MSSQL\LOG\ERRORLOG,template=-EQ:\SQLDEFAULT\MSSQL.4\MSSQL\LOG\ERRORLOG)
```

```
Process: canextend.ksh(1156)
```

```
*ERROR* (No. 14003) The target value for Log file location does not match the template value (target=-LQ:\SQLDEFAULT\MSSQL.1\MSSQL\DATA\MASTLOG.LDF,template=-LQ:\SQLDEFAULT\MSSQL.4\MSSQL\DATA\MASTLOG.LDF)
```

```
Error - extmgr (HAWK, MSSQL.0, MSSQL.0, OSPREY) - canextend failed
```

Suggested Actions:

Use the ID Conflict Resolution option to resolve Microsoft SQL ID mismatches between the primary and backup servers. In the LifeKeeper GUI, right-click on the **SQL Server resource** (on the right-hand side of the LifeKeeper GUI) and select **properties**, then select the **SQL Server Configuration** tab. Click the **Admin Actions** button and select **ID Conflict Resolution** on the **SQL Server Configuration** page.

See the topic [Managing a SQL Server Configuration](#) for more details on the SQL Server Configuration page.

SQLAgent Service Fails to Start

SQLAgent Service Fails to Start Sometimes for Named Instances

Symptom:

On named instances of Microsoft SQL Server where SIOS Protection Suite is protecting the SQLAgent service, when the resource is originally brought in service, a SQL problem prevents this service from starting and forces a MAXWAIT situation (300 second delay) before the SQL gives up trying to start the SQLAgent service.

This message indicates that the INFO field of the SQL resource has become corrupted. You must delete and re-create the SQL resource. Note that you should remove any IP and volume dependencies prior to deleting the resource. Upon creating the new SQL resource, SIOS Protection Suite will re-create the dependencies.

If the Microsoft SQL Server service is already started on the system where the SQLAgent service is trying to start, you'll likely see this scenario.

If the SQLAgent service and the Microsoft SQL Server service are both started when the SQL hierarchy creation occurs, you will not see this issue.

Suggested Actions:

Stopping and starting the Microsoft SQL Server service usually clears up the problem and the SQL Agent service then starts. However, stopping the SQL Server service is not a good option.

Both the MSSQLServer and the SQLServerAgent service should start up properly using the Local System Account on the default instance or a named instance. They will both start up using a Domain Administrator account, provided you have added that Domain Admin account to the Local Administrator Group on each system.

Connecting ODBC Clients to Named Instances of SQL Server

Symptom:

After creating an ODBC connection to your SIOS Protection Suite cluster via the protected IP address, the connection fails after switching the SQL resource to the backup server.

Suggested Action:

1. Take each instance out of service and bring it back into service on the PRIMARY. Examine the application event log to determine which IP:PORT that particular SQL Server instance is listening on.
2. Bring each hierarchy In Service on the SECONDARY server and note the IP:PORT each SQL instance is listening on.
3. To insure your clients can connect via ODBC to either server, make sure the PORT each instance is listening on is the same on both the PRIMARY and SECONDARY servers.
4. To do this, use the Microsoft SQL Server Network Utility. Select the SQL instance (it must be running on that machine), highlight the TCP/IP protocol and look at the properties to determine the current default port it is listening on.
5. Change this value so the default PORT is the same on both systems for this instance.
6. Create your ODBC connections for each instance using the protected IPs:PORTs you just set up in Step 5.

SIOS Protection Suite PostgreSQL Server Recovery Kit Introduction

The SIOS Protection Suite PostgreSQL Server Recovery Kit software lets you tie the data integrity of PostgreSQL-based databases to the increased availability provided by SIOS Protection Suite for Windows.

The LifeKeeper GUI allows you to easily create a PostgreSQL resource hierarchy. SIOS Protection Suite can then protect all of the disk resources used by the PostgreSQL Server instance, as well as the LifeKeeper network resources used by clients to access the database.

PostgreSQL Server Installation

Proper operation of the SIOS Protection Suite PostgreSQL Server Recovery Kit depends upon correct setup of the hardware and software.

Before continuing, please preview the [Hierarchy Administration](#) section of this guide. This section provides general guidelines, configuration details and troubleshooting hints to help you administer PostgreSQL Server in a SIOS Protection Suite environment.

Installation and Configuration Details - Adding LifeKeeper to an Existing PostgreSQL Configuration

This section covers installation and configuration of the SIOS Protection Suite and PostgreSQL software when adding LifeKeeper to an existing PostgreSQL configuration. The Primary server is assumed to be the location of the active PostgreSQL database cluster that will be protected. These steps assume the database cluster to be protected is the version created by the default installation of PostgreSQL (e.g. postgresql-x64-9.6 with PostgreSQL v9.6) These steps below must be followed in order.

1. Install the SIOS Protection Suite software on the Primary and Backup Server.
2. Using the **LifeKeeper GUI** on the Primary Server, create comm paths between the primary and backup server.
3. On the Primary Server:
 1. Use the **Disk Management** utility to configure your disk resources and define the shared or replicated volumes that you want to use. Be sure the volume size is adequate. If you are configuring shared volumes, power down the backup server during configuration to avoid simultaneous access to your storage.
 2. It is recommended that you use **Windows Explorer** to unshare all volumes to be used by the PostgreSQL Server Instance from the network.
 3. Configure your networking to support the SIOS Protection Suite TCP/IP comm path(s) and the switchable IP address if applicable.
4. On the Backup Server:
 1. Start the backup server if it was stopped to configure shared volumes in Step 3.

2. Use the **Disk Management** utility to configure your disk resources and define the shared or replicated volumes that you want to use. If you are using shared storage, assign the same drive letter to the shared volume as assigned on the primary server. For replicated storage be sure the volume size is adequate.

5. On the Primary Server:

1. In SIOS Protection Suite, create your shared or replicated Volume resource (where the PostgreSQL database cluster will reside) and extend it to the backup server. Later when you create your PostgreSQL Server resource hierarchy. SIOS Protection Suite will automatically bring the resource into the hierarchy as a dependency.

6. On the Backup Server:

1. Bring the volume resource hierarchy In Service using the **LifeKeeper GUI**.
2. Install PostgreSQL to the same program folder as it is installed to on the primary server using the following guidelines:
 1. Using the `-extract-only` argument to the PostgreSQL installer is not recommended as it does not configure all of the information required by the PostgreSQL Recovery Kit.
 2. By default the installation of the PostgreSQL software creates a single database cluster (a database cluster is a collection of databases that is managed by a single instance of a running PostgreSQL database server). During installation the Data Directory prompt determines where the database cluster will be created. The default instance can be installed at any location since it will be deleted in step 6d.
3. The default database cluster instance created during installation is not required and can be removed.
 1. Stop the database cluster instance created during installation.

2. Open **Explorer** and access the drive associated with the replicated volume.
 3. Delete the PostgreSQL database cluster directory created during the installation.
 4. Delete the PostgreSQL service that was created during installation. You can use the Windows "sc delete " command to do this.
7. On the Primary Server:
1. Bring the volume resource hierarchy In Service using the **LifeKeeper GUI**.
 2. Stop the PostgreSQL database cluster instance that is to be protected by the PostgreSQL Recovery Kit.
 3. **Optional** - Perform a backup of the PostgreSQL database cluster data directory prior to moving it to the protected volume created above.
 4. Move the database cluster data directory to the protected volume created above.
 5. Set the access rights on the database cluster data directory. The user account setup to control the Window's Service for this instance must have full control file permissions on the data directory.
 6. Follow the steps outlined in [Configuring the Postmaster Port Argument](#) on the primary server. The configuration should also include modification of -D argument which specifies the location of the data directory. It should be changed to the path on the protected volume.
 7. Follow the steps outlined in [Configure for Unattended Connections](#) on the primary server.
 8. When the installation and configuration is complete, start the Windows PostgreSQL service to verify that the PostgreSQL Server can start properly on the primary server with the Postmaster port argument and the data directory now located on the protected volume.

9. Create the PostgreSQL Server hierarchy on the primary server and extend it to the backup server. See [Creating the PostgreSQL Hierarchy](#) for more information. Test the new PostgreSQL Server hierarchy by performing a manual switchover.

Installation and Configuration – Adding PostgreSQL to an Existing LifeKeeper Configuration

This section covers installation and configuration of the SIOS Protection Suite and PostgreSQL software when adding PostgreSQL to an existing LifeKeeper cluster. These steps must be followed in order.

1. On the Primary Server:

1. Use the **Window Disk Management** tool to configure your disk resources and define the shared or replicated volumes that you want to use. Be sure the volume size is adequate. If you are configuring shared volumes, power down the backup server during configuration to avoid simultaneous access to your storage.
2. It is recommended that you use **Windows Explorer** to unshare all volumes to be used by the PostgreSQL Server Instance from the network.
3. Configure your networking to support the SIOS Protection Suite TCP/IP comm path(s) and the switchable IP address if applicable.

2. On the Backup Server:

1. Start the backup server if it was stopped to configure shared volumes in Step 1.
2. Use the **Disk Management** utility to configure your disk resources and define the shared or replicated volumes that you want to use. If you are using shared storage, assign the same drive letter to the shared volume as assigned on the primary server. For replicated storage be sure the volume size is adequate.

3. On the Primary Server:

1. In SIOS Protection Suite, create your shared or replicated Volume resource (where the PostgreSQL database cluster will

reside) and extend it to the backup server. Later when you create your PostgreSQL Server resource hierarchy, SIOS Protection Suite will automatically bring the resource into the hierarchy as a dependency.

4. On the Backup Server:

1. Bring the volume resource hierarchy In Service using the **LifeKeeper GUI**.
2. Install the PostgreSQL Server software using the following guidelines:
 1. Using the `-extract-only` argument to the PostgreSQL installer is not recommended as it does not configure all of the information required by the PostgreSQL Recovery Kit.
 2. By default the installation of the PostgreSQL software creates a single database cluster (a database cluster is a collection of databases that is managed by a single instance of a running PostgreSQL database server). During installation the Data Directory prompt determines where the database cluster will be created. The location entered should be on the protected volume created above in Step 3a.
 3. The database service created during installation does not configure the postmaster process to start with the port argument (`-p port`) which is required by the SIOS Protection Suite PostgreSQL Server Recovery Kit to properly manage the instance. The Windows service (e.g. `postgresql-x64-9.6` with PostgreSQL v9.6) created for the default database cluster will need to be updated to include this option if it will be protected by the SIOS Protection Suite. See [Configuring the Postmaster Port Argument](#) for more information.
3. Follow the steps outlined in [Configure for Unattended Connections](#) if the `pgpass.conf` will be used for authentication (if setting up a trust relationship via the `pg_hba.conf` file this step can be skipped as that will be done as part of the configuration on the primary server).
4. Verify the Windows PostgreSQL database cluster will start once

the installation and configuration (for the Postmaster port argument and for unattended connections) is complete. This requires stopping and restarting the default Windows Service for PostgreSQL. Once the verification is complete, stop the PostgreSQL service.

5. On the Primary Server:

1. Bring the volume resource hierarchy In Service using the **LifeKeeper GUI**.
2. Open **Explorer** and access the drive associated with the replicated volume.
3. Delete the PostgreSQL database cluster directory created during the installation on the backup server. (You will recreate it in the next step).
4. Install the PostgreSQL Server software EXACTLY as you did on the backup server (program files in the same directory on the local disk and data files in the same location on the protected volume).
5. Follow the steps outlined in [Configuring the Portmaster Port Argument](#) on the primary server.
6. Follow the steps outlined in [Configure for Unattended Connections](#) on the primary server.
7. When the installation is complete, start the Windows PostgreSQL service to verify that the PostgreSQL Server can start properly on the primary server.
8. Create the PostgreSQL Server hierarchy on the primary server and extend it to the backup server. See [Creating the PostgreSQL Hierarchy](#) for more information. Test the new PostgreSQL Server hierarchy by performing a manual failover.

Additional Setup Tasks for Extended Configurations

If your configuration uses a shared storage device or you are using SIOS DataKeeper, you may choose a configuration that will be extended to a third (or more) server(s).

1. If it has not already been done, configure two systems following the steps given in [Installation and Configuration - Adding PostgreSQL to an Existing LifeKeeper Configuration](#).
2. Use the **Disk Management** utility to configure your disk resources and define the shared or replicated volumes that you want to use. If you are using shared storage assign the same drive letter to the shared volume as assigned on the primary server. For replicated storage be sure the volume size is adequate.
3. Install the PostgreSQL Server software EXACTLY as you did on the primary server (program files in the same directory on the local disk and data files in the same location on the protected volume) using the following guidelines:
 1. Using the `-extract-only` argument to the PostgreSQL installer is not recommend as it does not configure all of the information required by the PostgreSQL Recovery Kit.
 2. By default the installation of the PostgreSQL software creates a single database cluster (a database cluster is a collection of databases that is managed by a single instance of a running PostgreSQL database server). During installation the Data Directory prompt determines where the database cluster will be created. Select any location as the data directory will be removed in a later step. If replicated storage is being used it can be the volume created above in step 2. If shared storage is being used do not use the volume located above to prevent overwriting the current PostgreSQL database cluster.
 3. The database service created during install does not configure the postmaster process to start with the port argument (`-p port`) which is required by the SIOS Protection Suite PostgreSQL Server Recovery Kit to properly manage the instance. The Windows service (e.g. `postgresql-x64-9.6` with

PostgreSQL v9.6) created for the default database cluster will need to be updated to include this option if it will be protected by the SIOS Protection Suite. See [Configuring the Postmaster Port Argument](#) for more information. While updating the Postmaster Port Argument you will also need to update the data directory path (-D argument value) to match the install location on the primary server.

4. Follow the steps outlined in [Configure for Unattended Connections](#) if the pgpass.conf will be used for authentication (if setting up a trust relationship via the pg_hba.conf file this step can be skipped as that will have already been done as part of the configuration on the primary server).
5. The default database cluster instance created during installation is not required and can be removed.
 1. Stop the database cluster instance created during installation.
 2. Open **Explorer** and migrate to the location of the data directory entered during installation of the PostgreSQL software.
 3. Delete the PostgreSQL database cluster directory.
 4. Delete the PostgreSQL service that was created during installation. You can use the Windows "sc delete <servicename>" command to do this.
6. If no comm paths exist to the new LifeKeeper cluster node, use the LifeKeeper GUI on the Primary Server to create comm paths between the primary and the new LifeKeeper cluster node and to create comm paths between the backup server and the new LifeKeeper cluster node.
7. Extend the PostgreSQL hierarchy to the new LifeKeeper cluster node. See [Creating the PostgreSQL Hierarchy](#) for more information. Test the new PostgreSQL Server hierarchy by performing a manual switchover to the new LifeKeeper cluster node.

Server Configuration Considerations

PostgreSQL Server Configuration Considerations

Before you install and configure your PostgreSQL database clusters, it is important to understand how to configure them. It is also important to understand the concepts of Active/Standby and Active/Active configurations, and how they can be set up in a PostgreSQL configuration.

PostgreSQL Database Cluster Configuration Considerations

The SIOS Protection Suite for PostgreSQL uses Windows services for administration of the PostgreSQL database cluster. If a Windows Service does not already exist for the PostgreSQL database cluster, then one will be setup when the PostgreSQL hierarchy is created.

For SIOS Protection Suite to protect a PostgreSQL database cluster the following conditions must exist:

- The PostgreSQL Server must be running
- The PostgreSQL postmaster process must be running with the port option: `-p port`
- The PostgreSQL database cluster data directory must reside on a protected volume
- The PostgreSQL database cluster data directory, sub-directories, and all files must be accessible by the Windows Service account on all servers

The following configuration limitations currently exist in the kit:


- Does not automatically include IP resource instances as part of the hierarchy. If a user connects remotely, then an IP resource will need to be created and added as a child resource in the PostgreSQL hierarchy.
- Only the location of the database cluster data directory is taken into consideration when determining which volume resources need to

be included as part of the resource hierarchy. If any database table space is not located on the same protected volume as the data directory, then the volume containing the table space will need to be protected and added as a child resource in the PostgreSQL hierarchy.

PostgreSQL Active/Standby Configuration

A configuration is Active/Standby when there is only one PostgreSQL database cluster, located on a shared or replicated volume. The PostgreSQL database cluster services run on only one system at a time. The servers are assigned priorities within SIOS Protection Suite which determine the order of failover for a particular hierarchy.

The figure below depicts a single PostgreSQL instance installed on a pair of servers. The instance contains one database cluster, PGSQL1 residing on a single volume.



When you create the PostgreSQL hierarchy within SIOS Protection Suite, you are asked to specify the PostgreSQL data directory (database cluster location). If remote connections to the database cluster will be made then a protected IP resource will need to be configured and added to the


resource hierarchy. SIOS Protection Suite then reads the configuration data for that instance and pulls the associated volumes into the hierarchy.

Once the hierarchy is created, it will appear as follows in the LifeKeeper GUI.



Active/Standby Failover

In the event of failure, SIOS Protection Suite brings the PostgreSQL Server hierarchy In Service on the backup Server. PostgreSQL Server is started on the backup server and it takes over protection of the database cluster as depicted in the figure below.




PostgreSQL Active/Active Configuration

Multiple PostgreSQL Server database clusters can be configured on any of

the servers using initdb. SIOS Protection Suite can protect the multiple PostgreSQL database clusters in what is called an Active/Active configuration. SIOS Protection Suite identifies each instance by the port used for connections.

Each database cluster is protected in a single SIOS Protection Suite hierarchy.

The figure below depicts two PostgreSQL database clusters: PGSQL1 and PGSQL2.



Notes:

- In this configuration Server1 is the primary server for the PGSQL1 database cluster and Server2 is the primary server for the PGSQL2 database cluster.

- Each server can be the primary and backup server for multiple instances.
- It would be possible for Server1 or Server2 to be the primary server for both database clusters.

Creating the PostgreSQL Hierarchy

After you have completed the necessary setup tasks outlined in the [SIOS Protection Suite for Windows Installation Guide](#), use the steps listed below to create the PostgreSQL Server hierarchy to protect your database cluster.

Important



If you have an existing PostgreSQL database cluster installed, you may need to close any client applications (local or remote) that are accessing the PostgreSQL database cluster prior to completing this procedure. Closing all client connections is required if any of the following conditions exist:

- The PostgreSQL database cluster data directory does not reside on a protected volume. To be highly available the PostgreSQL data directory must reside on a protected volume that can be switched between nodes in the LifeKeeper cluster. This will require manually moving the data directory prior to creating the resource hierarchy. Once the move is complete, restart the PostgreSQL database cluster services.
- The PostgreSQL Server instance is not controlled by an existing Windows service. To facilitate the administration of the protected PostgreSQL Server, the hierarchy create will create a Windows service if one does not already exist. This requires stopping and restarting the PostgreSQL database cluster services.
- The postmaster process is not running with the “-p port” option. See [Configuring the Postmaster Port Argument](#) for more information on how to verify and start PostgreSQL with the postmaster port argument.

1. From the LifeKeeper GUI menu, select **Edit** then **Server**. From the menu, select **Create Resource Hierarchy**.

2. The **Create Protected Application** dialog box will display. Select the **Primary** and **Backup** servers from the pull-down list. Select **Next** to continue.
3. The dialog box will appear with a drop down list box displaying all recognized recovery kits installed within the cluster. Select **PostgreSQL Server** and click **Next**.
4. 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
PostgreSQL Service Name	<p>Enter the Windows Service Name for this instance.</p> <p>If the name entered is an existing PostgreSQL service, then the PostgreSQL Recovery Kit will use the service information to pull information for the resource creation.</p> <p>If the name entered is not an existing PostgreSQL service, then the create process will require additional inputs in order to create a Windows Service using the entered name.</p> <p>Note: If a Windows Service is created during the PostgreSQL resource create, then a delete of the PostgreSQL resource will remove the PostgreSQL service. If the PostgreSQL service already exists at create time, then a delete of PostgreSQL resource will not remove the service.</p>
PostgreSQL Executable Location	<p>Select the location of the PostgreSQL executables (directory where pg_ctl.exe and psql.exe reside). If multiple versions of PostgreSQL are installed on the system, then the path selected must match the PostgreSQL version used to create the protected database cluster. Example: If intidb from version 10 was used to create the database cluster, then the path selected for pg_ctl.exe and psql.exe must be the version 10 path.</p>
PostgreSQL	Enter the path to the data directory for the

Data Directory	PostgreSQL database cluster to be protected. If the PostgreSQL Service name entered exists, then it will default to the data directory defined for that service.
PostgreSQL Port	Select the port to be used for the PostgreSQL database cluster. This field is used to specify the TCP/IP port number on which the postmaster daemon is listening for connections from client applications. The default choice is obtained from the running PostgreSQL Server.
Enter Database Administrator User	Enter the name of the PostgreSQL database cluster administrative user. This user must have connection and administrative privileges.
PostgreSQL Database Tag	Enter a unique tag name, or you can accept the default tag name offered by SPS. Note: The tag name must consist of printable ASCII characters.
The following fields only appear if the PostgreSQL Service Name entered does not already exist as a Windows service.	
PostgreSQL Service Logon Account	Enter the user account to be used for the logon credentials for the new PostgreSQL Windows Service. This account must have privileges to start and stop the PostgreSQL server instance being protected. This user account must also have full control file permissions for all files in the PostgreSQL Data Directory.
PostgreSQL Service Logon Password	Enter the user password to be used for the logon credentials for the new PostgreSQL Windows Service. Note: If built-in system accounts (e.g. "Local System" or "Network Service") are used, enter any non-blank value into this field (blank is not allowed). For built-in accounts, the password is ignored.

- After you click **Create**, the **Wizard** will create your PostgreSQL resource. SIOS Protection Suite will validate the data entered. If SIOS Protection Suite detects a problem, an error message will appear in the information box.

6. Another information box will appear indicating that you have successfully created a PostgreSQL resource hierarchy, and you must **Extend** that hierarchy to another server in your cluster in order to achieve failover protection. Click **Next**.
7. After you click **Next**, SIOS Protection Suite will launch the **Pre-Extend Wizard**.

Extending a PostgreSQL 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. 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**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information. **Note:** These first two fields appear only if you initiated the **Extend** from the **Edit** menu.

Field	Tips
Primary Server	Select a server where a resource hierarchy to be extended is currently defined and in service.
Resource Hierarchy to Extend	Select the resource hierarchy to extend.
Backup Server	Select a server to be the backup server for the resource hierarchy.

3. After receiving the message that the pre-extend checks were successful, click **Next**.
4. The **Extend Wizard** will prompt you to enter the following information. **Note:** The first two fields appear only if a Windows Service for the PostgreSQL Server does not exist on the backup server (the service name entered on the primary server during create is used for checking on the backup server).

Field	Tips
PostgreSQL Service Logon Account	Enter the user account to be used for the logon credentials for the PostgreSQL Windows Service on the backup server. This account must have privileges to start and stop the PostgreSQL server instance being protected.
PostgreSQL Service Logon	Enter the user password to be used for the logon credentials for the new PostgreSQL Windows Service..

Password	
PostgreSQL Executable Location	Select the location of the PostgreSQL executables (directory where pg_ctl.exe and psql.exe reside). If multiple versions of PostgreSQL are installed on the system, then the path selected must match the PostgreSQL version used to create the protected database cluster. Example: If intidb from version 10 was used to create the database cluster, then the path selected for pg_ctl.exe and psql.exe must be the version 10 path.
Backup Priority	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. SIOS Protection Suite 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.

5. Click **Extend**.

Unextending a PostgreSQL Hierarchy

To remove a resource hierarchy from a single server in the SIOS Protection Suite 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 SQL resource. It cannot be the server where the PostgreSQL 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 SQL 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.

Deleting a PostgreSQL Hierarchy

Before deleting a PostgreSQL hierarchy or instance, make sure that the hierarchy is active (green) on its primary server. You may also wish to remove the dependencies before deleting the hierarchy; otherwise, the dependencies will also be deleted.

To delete a resource hierarchy from all the servers in your SIOS Protection Suite 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 PostgreSQL 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 PostgreSQL resource was deleted successfully.
6. Click **Done** to exit.

Testing Your PostgreSQL Resource Hierarchy

You can test your PostgreSQL resource hierarchy by initiating a manual switchover. This will simulate a failover of a resource instance from the primary server to the backup server.

Select **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 the 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.

PostgreSQL Server Hierarchy Administration

Follow these guidelines when administering your PostgreSQL Server.

[Access Via Protected Communication Paths](#)

[Reserve Volumes For Exclusive PostgreSQL Use](#)

[Start and Stop PostgreSQL Server Only Through SIOS Protection Suite](#)

[Creating and Protecting Additional PostgreSQL Database Clusters](#)

[PostgreSQL Administrative Login](#)

[Configuration for Unattended Connections](#)

[Configuring the Postmaster Port Argument](#)

[Monitoring Your PostgreSQL Hierarchy](#)

Access Via Protected Communication Paths

All remote access of the service should be done through the hierarchy's LifeKeeper network resources. This will ensure that users can access the PostgreSQL database cluster regardless of which server it is currently running on.

*Note: Currently the PostgreSQL recovery kit does not automatically include an IP communication resource as a dependent resource in the hierarchy. To ensure remote access regardless of which server is currently running the PostgreSQL database cluster, an IP resource will need to be created and added as a child to the PostgreSQL resource. Additionally, the PostgreSQL database cluster must be configured to listen on this address.

Reserve Volumes For Exclusive PostgreSQL Use

The volumes containing the protected PostgreSQL files should be reserved for use by PostgreSQL exclusively.

A SIOS Protection Suite protected volume may fail to switch over if it is accessed by an another application, process or remote user.

Start and Stop PostgreSQL Server Only Through SIOS Protection Suite

Although most of the administrative tasks for the PostgreSQL Server are done through the PostgreSQL tools, starting and stopping of the PostgreSQL Server should not be one of them:

1. **Consistent state** - When SIOS Protection Suite stops and starts the PostgreSQL Server, it maintains a consistent state for the protected Microsoft Service. Performing start and stop requests via the command line using the PostgreSQL tools such as pg_ctl.exe creates an inconsistent state for the Microsoft Service, as it is unable to detect the state as running or stopped. This can result in failures to detect and correct issues for the PostgreSQL Server.
2. **Protected PostgreSQL services** should be set to **Manual** startup mode through the **Control Panel "Services"** tool. **Note:** When creating a PostgreSQL resource hierarchy the protected service will automatically be set to **Manual** startup mode.

Creating and Protecting Additional PostgreSQL Database Clusters

As your environment grows, you may need to add new PostgreSQL Server database clusters on existing or new shared or replicated volumes.

To add and protect a new database cluster and the associated volume follow these steps:

1. **Create the volume resource.** On the server where the PostgreSQL database cluster will be placed, create and extend a volume resource.
2. **Create the database cluster.** Run `initdb` to create the new database cluster. Be sure to locate the data directory on the volume resource created above.
3. **Set access permissions.** In order to start and stop the PostgreSQL database cluster, the data directory and all files and sub-directories must have access rights that are recognized by all nodes in the cluster. By default the data directory and all files and sub-directories will have access rights based on the user running `initdb`. If the user is the local administrator, then attempting to start the database cluster on any other server will fail as that server will not have access rights. Either running `initdb` while logged on as a domain user or adding access rights for **NT AUTHORITY\NetworkService** is recommended. The account is used to provide access across all servers and should be the same account used for the Windows service logon credentials when creating the PostgreSQL resource.
4. **Configure for unattended connections.** The database cluster must be configured to allow connections without requiring a password. See [Configuration for Unattended Connections](#) for more information.
5. **Start the database cluster.** The database cluster must be running to create the PostgreSQL resource. Start the database cluster via `pg_ctl.exe` with the `-o "-p port"` argument. See [Configuring the Postmaster Port Argument](#) for more information on how to verify and start PostgreSQL with the postmaster port argument.
6. **Create the PostgreSQL resource.** On the server that the new database

cluster was created on, create and extend a PostgreSQL resource. During the resource create you will be prompted for a Windows service name along with the logon credentials (user account and password). The Windows service name can be anything and will be used to create a Service account for administration (starting, stopping ...) of the database cluster. The logon credentials used should be the ones setup in step 3 to ensure access rights on all servers.

PostgreSQL Administrative Login

During the creation of a SIOS Protection Suite PostgreSQL resource the user must enter a PostgreSQL administrative username for that database cluster. This administrative username is used for client connections through the psql utility.

The username:

1. Must have the ability to connect to the database (template1), as well as obtain the listing of defined databases for the instance.
2. Must have the ability to view system tables and make generalized queries.
3. Must allow unattended (non-terminal or scripted) connections. See [Configuration for Unattended Connections](#) for details.

Configuration for Unattended Connections

During the creation of a SIOS Protection Suite PostgreSQL resource, the user must enter a PostgreSQL administrative username for that database cluster. No password is requested for this username by the PostgreSQL recovery kit. Therefore, for the connection tests performed during health checking to be successful one of two configuration methods must be used:

- Trust configuration (no password)
- Credentials supplied via pgpass.conf

Trust Configuration

To configure trusted connections for the PostgreSQL database cluster modifications to the `pg_hba.conf` file are required. The authentication method for the administrative user for the database cluster must be set to `'trust'`. The following is an example entry for the administrative user `pgsql`:

#	DATABASE	USER	ADDRESS	METHOD
TYPE				
# IPv4 local connections:				
host	all	pgsql	127.0.0.1/32	trust
# IPv6 local connections:				
host	all	pgsql	:::1/128	trust

Credentials via pgpass.conf

To supply credentials for the PostgreSQL database cluster administrative user and password, creation of a `pgpass.conf` file for the logon account specified for the LifeKeeper service is required. This file must be created in the user's APPDATA folder. Login as the user account used to start LifeKeeper, and follow these steps:

- Change directories to `appdata`

- Create the directory postgresql if it does not exist
- Change directories to postgresql
- Create the file pgpass.conf with the following format:

```
hostname:port:database:user:password
```

If the PostgreSQL database cluster administrative user password changes, then the pgpass.conf file will need to be updated with the password setting.

If a non-login user account (such as the built-in accounts "Network Service" or "Local System") is used to start LifeKeeper on a node, the APPDATA folder can be determined by running the following LifeKeeper command on a different node in the LifeKeeper cluster.

```
lcdremexec -d <node> - echo %APPDATA%
```

The pgpass.conf file can be created by a system administrator in that folder. Be sure to add read permissions for the LifeKeeper login account user to the pgpass.conf file that is created.

Configuring the Postmaster Port Argument

To properly manage the PostgreSQL database cluster the PostgreSQL Recovery Kit requires the postmaster process to be running with the port argument: `"-p port"`. The port is required to create a PostgreSQL resource hierarchy and for monitoring once the hierarchy is created. To view the current argument list for the postmaster process see the `postmaster.opts` file located in the data directory for PostgreSQL database cluster. The following is an example of the contents of this file:

```
C:/Program Files/PostgreSQL/9.6/bin/postgres.exe "-D" "E:\PGSQL1" "-p" "5432"
```

In this example the postmaster process is running with the port argument. If the postmaster process is not running with the port argument it will need to be added via one of the two methods listed below:

- Adding the port argument to an existing Windows Service Instance
- Adding the port argument to a non-Windows Service Instance

Adding the port argument to an existing Windows Service Instance

If the PostgreSQL database cluster to be protected by the PostgreSQL recovery kit is running via a Windows Service, such as the `postgresql-x64-9.6` service created by the initial install of the PostgreSQL software v9.6, then the startup command line for the service will need to be modified. This can be done via the following steps:

1. Stop the existing PostgreSQL database cluster instance. This can be done using several different methods:
 1. Use the Windows Service Interface (`services.msc`) and select to stop the service.
 2. Use the command line utility `"sc"`, e.g. `sc stop service name`.
 3. Use the command `"net"`, e.g. `net stop service_name`.
2. Modify the startup command line for the service. This can be done

via the command line using the "sc" utility or by editing the startup command line in the registry. **Note:** The startup command line for the PostgreSQL service uses the pg_ctl.exe utility. To pass arguments to the postmaster process from the pg_ctl utility requires the use of the "-o" argument. The "-o" argument takes a quoted list of postmaster startup options such as the "-p port" as used in the modifications shown below.

To edit the startup command line using the "sc" command line utility follow these steps. This example uses the postgresql-x64-9.6 service:

1. Retrieve the binPath (startup command line) for the service, e.g "sc qc postgresql-x64-9.6" it will return something like the following (for this example only the output line with the startup command line is shown):

```
BINARY_PATH_NAME : "C:\Program Files\PostgreSQL\9.6\bin\pg_ctl.exe" runservice -N "postgresql-x64-9.6" -D "E:\PGSQL1" -w
```

2. Update the binPath. **Note:** There is a space after the "=" for the binPath argument and the value must be in double quotes therefore it requires escaping the imbedded double quotes.

```
sc config postgresql-x64-9.6 binPath= "\"C:\Program Files\PostgreSQL\9.6\bin\pg_ctl.exe\" runservice -N \"postgresql-x64-9.6\" -D \"E:\PGSQL1\" -w -o \"-p 5432\" "
```

To edit the startup command line in registry requires the use of a registry edit tool such as regedit.

- a. Using regedit find the service ImagePath registry value:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\postgresql-x86-9.6\ImagePath
 - b. Modify the ImagePath value by adding the port argument at the end of the existing command line: -o "-p 5432"
3. Restart the existing PostgreSQL database cluster instance. This can be done using several different methods:
 1. Use the Windows Service Interface (services.msc) and select to start the service.

2. Use the command line utility "sc", e.g. `sc start service_name`.
3. Use the command "net", e.g. `net start service_name`.
4. Verify the postmaster process is running with port argument by checking the `postmaster.opts` file as described above.

Adding the port argument to a non- Windows Service Instance

If a Windows Service does not exist for the PostgreSQL database cluster then stopping and restarting the PostgreSQL database cluster via the PostgreSQL utility `pg_ctl.exe` will be required. This can be done via the following steps:

1. Stop the PostgreSQL database cluster.

```
pg_ctl stop -D F:\PGSQL2
```

2. Start the PostgreSQL database cluster.

```
pg_ctl start -D F:\PGSQL2 -w -o "-p 5433"
```

3. Verify the postmaster process is running with port argument by checking the `postmaster.opts` file as described above.

Monitoring Your PostgreSQL Hierarchy

SIOS Protection Suite monitors the PostgreSQL database cluster services for every PostgreSQL resource. If the services stop, the monitoring process associated with the PostgreSQL resource will detect this and attempt to restart the services on the local server if Local Recovery is enabled. If Local Recovery is disabled, the resource will fail over to the backup server.

SIOS Protection Suite will also perform a query to test the connection to the protected PostgreSQL database cluster. If the query fails, the monitoring process associated with the PostgreSQL resource will detect this and attempt to restart the services if Local Recovery is enabled. If Local Recovery is disabled, the resource will fail over to the backup server.

Troubleshooting

This section provides suggestions and insights into occurrences that are not specifically related to the SIOS Protection Suite software but have a relationship with the total environment.

[Create Fails](#)

[Restore Fails](#)

Create Fails

Symptom:

The Create of a PostgreSQL Server resource will fail if the database cluster data directory is located on a volume and that volume is not already protected by SIOS Protection Suite.

Suggested Action:

Create the volume resource for the data directory.

Symptom:

The Create of a PostgreSQL Server resource will fail if the postmaster process for PostgreSQL database cluster is not running with the "-p port" option.

Suggested Action:

To ensure access to the correct PostgreSQL database cluster the postmaster process must be running with the "-p port" argument. See [Configuring the Postmaster Port Argument](#) for more information on how to verify and start PostgreSQL with the postmaster port argument.

Symptom:

The Create of a PostgreSQL Server resource will hang after entering the administrative user if the unattended connection configuration is incorrect or non-existent.

Suggested Action:

Correctly configure unattended connections. See [Configuration for Unattended Connections](#).

Restore Fails

Symptom:

The Restore of a PostgreSQL Server resource will fail after switchover if the database cluster data directory access permissions are configured incorrectly.

Suggested Action:

Verify the access permissions are configured correctly. See step 3 in [Creating and Protecting Additional PostgreSQL Database Clusters](#).

Symptom:

The Restore of a PostgreSQL Server resource will fail if the database cluster instance is started via **pg_ctl.exe start** and not via an in service action in LifeKeeper or via a service start via Windows APIs. Using pg_ctl.exe to start the database cluster creates an inconsistency in the Windows Service state causing a LifeKeeper restore to fail on the attempt to start an already running instance.

When attempting to start an already running instance, PostgreSQL will log the following messages:

```
FATAL: lock file "postmaster.pid" already exists
```

```
HINT: Is another postmaster (PID 3488) running in data directory  
"E:/PGSQL1"?
```

Suggested Action:

To correct this condition the database cluster must be stopped via **pg_ctl stop**. Once the stop completes the LifeKeeper in service action should be successful.

Symptom:

The Restore of a PostgreSQL Server resource can fail if the database cluster did not shut down cleanly because of server crash or the PostgreSQL service was hung when the shutdown occurred (windbg was used to simulate a hang). The inability to shutdown cleanly will force a database cluster recovery action on the next startup. This recovery action can cause the Window's Service start action to fail placing the service in an inconsistent state with the database cluster state. During startup after a unclean shutdown, PostgreSQL may log the following messages (along with a number of others):

```
Waiting for server start up
```

```
LOG: database system was interrupted; last known up at 2017-07-25  
16:12:10 EDT
```

```
FATAL: the database system is starting up
```

```
LOG: database system was not properly shut down; automatic recovery  
in progress
```

Once the recovery is complete the PostgreSQL database cluster processes are running but the Window's Service state is "Stopped" and the LifeKeeper PostgreSQL resource is in the failed state. If a LifeKeeper restore action is attempted with the database cluster up and running, PostgreSQL will log the following messages:

```
FATAL: lock file "postmaster.pid" already exists
```

```
HINT: Is another postmaster (PID 3488) running in data directory  
"E:/PGSQL1"?
```

Suggested Action:

To correct this condition the database cluster must be stopped via **pg_ctl stop** once the recovery is complete. Once the stop completes the LifeKeeper in service action should be successful.

Tunable Settings for the PostgreSQL Recovery Kit

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 %LKROOT%\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_START_RETRIES**

This tunable controls the number of times the PostgreSQL Recovery Kit will loop waiting for the database cluster to start before giving up with a failed status. There is a 5 second wait between each retry. By default the Recovery Kit will retry 12 times resulting in a 60 second wait for the database cluster to startup. The minimum value for this tunable setting is 12.

- **LKPGSQL_STOP_RETRIES**

This tunable controls the number of times the PostgreSQL Recovery Kit will loop waiting for the database cluster to stop before giving up with a failed status. There is a 5 second wait between each retry. By default the Recovery Kit will retry 12 times resulting in a 60 second wait for the database cluster to stop. The minimum value for this tunable setting is 12.

- **LKPGSQL_RESTORE_CONNECT_RETRIES**

This tunable controls the number of times during a restore action that the PostgreSQL Recovery Kit will loop waiting for the database cluster to respond to a connection request. If the number of retries is reached the restore action will be failed. There is a 5 second wait between each retry. By default the Recovery Kit will retry 2 times resulting in a 10 second wait for the database cluster to respond to the connection request. The minimum value for this tunable setting is 2.

Tunable Settings for the PostgreSQL Database Cluster

To check the health of the PostgreSQL database cluster, the Recovery Kit's `deepchk` script will attempt to connect to the database cluster's `template1` db. This connection operation can fail if all the available connections are in use. This can be prevented via the PostgreSQL super user connection tunable.

- `superuser_reserved_connections`

This tunable in the database cluster `postgresql.conf` file controls the number of super user connections allowed. By default this setting is commented out. This setting should be uncommented. For the change to take effect, the PostgreSQL database cluster will need to be stopped and restarted.

SIOS Protection Suite Oracle Recovery Kit Introduction

LifeKeeper Oracle

The LifeKeeper Oracle Recovery Kit provides a way to recover an Oracle database instance (version 10g, 11g, 12c or 18c) from a failed server to a backup server. You can also extend the protection of the database instance to other servers. Using the LifeKeeper GUI, you can easily create a complete resource hierarchy so that the recovery operation includes all the disk resources used by the Oracle System Identifier (SID) as well as the Named Pipe and/or IP socket resources used to access the database.

Oracle Overview

The LifeKeeper Oracle Recovery Kit provides the ability to concurrently run Oracle database instances on other servers, and to optionally place these instances under LifeKeeper protection. Such a configuration is known as Active/Active and allows LifeKeeper servers to be fully utilized under normal operating conditions.

The LifeKeeper Oracle Recovery Kit includes the ability to recover the database instance locally (local recovery) before trying to fail over the database instance to a standby server.

The LifeKeeper Oracle Recovery Kit protects the following Core/Standard Oracle services:

- Oracle Service
- Oracle TNS Listener

The LifeKeeper Oracle Recovery Kit protects the following Optional Services for each release:

Optional 11g Services	Optional 12c Services	Optional 18c Services
Oracle DB Console	Oracle DB Console	Oracle DB Console
Oracle Job Scheduler	Oracle Job Scheduler	Oracle Job Scheduler
Oracle Cluster Service	Oracle Cluster Service	Oracle Cluster Service
Note: This service is for Automatic Storage Management (ASM) and is not available for protection under LifeKeeper because LifeKeeper does not currently support ASM.	Note: This service is for Automatic Storage Management (ASM) and is not available for protection under LifeKeeper because LifeKeeper does not currently support ASM.	Note: This service is for Automatic Storage Management (ASM) and is not available for protection under LifeKeeper because LifeKeeper does not currently support ASM.

The typical Oracle resource hierarchy consists of the following resources:


- Oracle

- Shared communication resource (IP or LAN Manager alias name)
- Volume(s)

All Oracle data, log, and trace (core database) files for the protected SID are stored on shared or replicated volumes. Upon detecting a failure, LifeKeeper switches the core database files, along with its associated data volumes and communication resources, to a backup server. The recovery can be completely transparent to database users. Once LifeKeeper switches all dependent resources to the backup server, it starts the Oracle service on that server.

The LifeKeeper GUI display shown below depicts a typical resource hierarchy. The Oracle resource is the topmost resource in the hierarchy tree. It is responsible for starting and stopping the dependent resources (communication and volume resources) in the correct order.

This particular Oracle hierarchy uses only IP for its communication/Listener resource.



Oracle Services

The LifeKeeper Oracle Recovery Kit protects the following Core/Standard Oracle services:

- Oracle Service
- Oracle TNS Listener

The LifeKeeper Oracle Recovery Kit protects the following Optional Services for each release:

Optional 11g Services	Optional 12c Services	Optional 18c Services
Oracle DB Console	Oracle DB Console	Oracle DB Console
Oracle Job Scheduler	Oracle Job Scheduler	Oracle Job Scheduler
Oracle Cluster Service	Oracle Cluster Service	Oracle Cluster Service
Note: This service is for Automatic Storage Management (ASM) and is not available for protection under LifeKeeper because LifeKeeper does not currently support ASM.	Note: This service is for Automatic Storage Management (ASM) and is not available for protection under LifeKeeper because LifeKeeper does not currently support ASM.	Note: This service is for Automatic Storage Management (ASM) and is not available for protection under LifeKeeper because LifeKeeper does not currently support ASM.

Resource Hierarchy for Oracle


The typical Oracle resource hierarchy consists of the following resources:

- Oracle
- Shared communication resource (IP or LAN Manager alias name)
- Volume(s)

All Oracle data, log, and trace (core database) files for the protected SID are stored on shared or replicated volumes. Upon detecting a failure, LifeKeeper switches the core database files, along with its associated data volumes and communication resources, to a backup server. The recovery can be completely transparent to database users. Once LifeKeeper switches all dependent resources to the backup server, it starts the Oracle service on that server.

The LifeKeeper GUI display shown below depicts a typical resource hierarchy. The Oracle resource is the topmost resource in the hierarchy tree. It is responsible for starting and stopping the dependent resources (communication and volume resources) in the correct order.

This particular Oracle hierarchy uses only IP for its communication/Listener resource.



Restriction: If you are configuring multiple ORACLE instances, then you must configure each listener to listen on a unique virtual IP address.

LifeKeeper will not allow you to configure multiple listeners to listen on the same virtual IP address with different ports.

Recovery Kit Requirements

Before installing and configuring the LifeKeeper Oracle Recovery Kit, be sure that your configuration meets the following requirements:

Operating System software. See the [SIOS Protection Suite for Windows Support Matrix](#) for a list supported versions of Windows operating systems.

***LifeKeeper software.** You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [Release Notes](#) for specific LifeKeeper requirements.

SIOS DataKeeper software (optional). If you plan to use Oracle with replicated volumes rather than shared storage, you should install the SIOS DataKeeper for Windows software on each server.

Oracle Database. The recovery kit supports Oracle versions 10g or 11g.

Installation

Oracle Installation

The LifeKeeper Oracle Recovery Kit is available via ftp download. Installation is simple and quick using InstallShield to provide a standard installation interface.

Recovery Kit Installation

Before installing the LifeKeeper Oracle Recovery Kit, be sure you are familiar with the product prerequisites listed above, as well as the installation/configuration procedure outlined in the section [Installing and Configuring LifeKeeper with Oracle](#).

Upgrading Recovery Kit from Previous Version

You may upgrade from the previous version of the LifeKeeper Oracle Recovery Kit software while preserving your resource hierarchies.

***Note:** You must close and restart the LifeKeeper GUI after upgrading the LifeKeeper Oracle Recovery Kit.

Kit Removal

To remove the LifeKeeper Oracle Recovery Kit software, choose the **"LifeKeeper Oracle Recovery Kit vX.X"** in the **Program and Features** applet in the control panel.

Installing and Configuring LifeKeeper with Oracle

The LifeKeeper with Oracle installation and configuration procedure differs slightly for replicated storage systems and shared stored systems. We have created separate sections in this document for each type of storage configuration in an effort to clarify the installation and configuration process.

For the most efficient setup, perform the following tasks to create an Oracle database instance first on the primary server and then on the secondary server.

Important Considerations

- When working on a particular server, switch the communication resource on that particular server, e.g. if working on the backup server, the communication resource should be switched to the backup server.
- The resource hierarchy will be preserved on an upgrade
- Only one SID is supported per Oracle Home directory
- Special consideration is required when the Oracle Home is not installed to the same LifeKeeper protected volume as the database (SID). If Oracle Home is installed on a different shared or replicated volume, that volume must be LifeKeeper protected and manually added as a dependent resource in the Oracle resource hierarchy. If Oracle Home is not installed on a LifeKeeper protected volume no changes are necessary.
- Use Windows disk management tools to configure the disk resources and the volumes. Use Oracle tools, e.g. Net Manager, to configure network protocols

Before Installing Oracle

Before you install the Oracle software, the servers and storage must be configured and LifeKeeper must be installed on each server in the cluster. By doing so, you can then install Oracle onto a volume that is already LifeKeeper-protected.

Replicated Storage Systems

On the Primary Server

1. Use the **Windows Disk Management** tool to configure your disk resources and define the replicated volumes that you want to use. (Be sure the volume size is adequate.)
2. It is recommended that you use **Windows Explorer** to unshare from the network all volumes to be used by the Oracle SID.
3. Configure your networking to support the LifeKeeper TCP/IP comm path(s) and, if applicable, the switchable IP address.
4. Install the LifeKeeper Core software on a local disk followed by the LifeKeeper Oracle Recovery Kit.
5. If you have a very large Oracle database, you should review the MAXWAIT value and consider increasing it.
6. Install the SIOS DataKeeper software to the local disk now. Refer to the [SIOS Protection Suite Installation Guide](#) for more details.

On the Backup Server

1. Bring up the backup server and use the **Disk Management** utility to assign the same drive letter to the replicated volume as assigned on the primary server.
2. Install the LifeKeeper Core software on a local disk followed by the LifeKeeper Oracle Recovery Kit.
3. If you have a very large Oracle database, you should review the MAXWAIT value and consider increasing it.

4. Install the [SIOS DataKeeper](#) software.

On the Primary Server

Now that you have LifeKeeper installed on both servers, go back to the primary server and do the following:

1. Using the LifeKeeper GUI, create comm paths between the primary and backup servers.
2. In LifeKeeper, create your communication resources (including either IP, LAN Manager or both) and extend them to the backup server. Later, when you create your Oracle resource hierarchy, LifeKeeper will automatically bring these resources into the hierarchy as dependencies.

***Note:** When the Oracle Hierarchy is created, the SIOS DataKeeper resource will automatically be created and brought into the Oracle resource hierarchy as a dependency.

Installing Oracle

Once you have installed LifeKeeper and configured the volume and communication resources, you are ready to install Oracle to the protected volume(s).

Oracle 12c Installation: Oracle12c introduces some new installation options. The following options for software and sample database installation are recommended.

- Oracle Services Account - Select the "Windows Built-In Account". This is referring to the Local System Account.
- Container Database - Uncheck the "Container Database" checkbox option for the sample database.

*If you are installing **Oracle 12c R2** go to **Advanced Options** to change the passwords for SYS and SYSTEM usernames.

- Password Management - After the sample database is installed the "Password Management" button is presented. Select it and set the password for the SYS and SYSTEM administrator accounts. One of these accounts (your choice) will be needed for creating the Oracle resource hierarchy.

On the Primary Server

1. Install the Oracle software to the protected shared volume. This creates the Oracle SID. Note that all files related to this Oracle SID (including log, trace, control, and data files) must be located on protected volumes.
2. Stop the default TNSListener service OracleTNSListener and set the startup mode to **Manual**. (You will create a new Listener for the SID to be protected in a later step.)
3. Using Oracle Net Manager, configure Oracle to use the LifeKeeper-protected communication resource(s) as follows:
 1. Create a new **TNSListener Service** using the SID name. Configure **Listening Locations** designating the LifeKeeper-protected IP address and/or named pipe (LAN Manager alias name). Then configure the Database Services specifying the **Oracle Home**

directory and SID.

2. Modify the **Oracle Service** for your SID. For TCP/IP, change the host name to the protected IP address. For **Named Pipes**, change the machine name to the LAN Manager alias.
4. Create a separate **TNSListener Service** instance for the SID to be protected under LifeKeeper. The service should be created using the `lsnrctl Start <SID>` command. This will create a service with the name `OracleTNSListener<SID>`.
5. Use the **Services** tool to test your Oracle services as follows:
 1. Verify that the new **TNSListener service** can be stopped and started successfully.
 2. Ensure that the `OracleService<SID>` service has been created by Oracle.
 3. Stop all Oracle services.

On the Backup Server

1. In LifeKeeper, bring the protected volume in service on the backup server.
2. Remove the Oracle inventory directory and rename the directory or directories that contain the Oracle data files. If this is a new installation you can delete the data files.
3. Install the Oracle software to the protected volume. Use EXACTLY the same installation options as on the primary server (the Oracle Home, SID name and paths must be identical). If prompted, choose to overwrite the existing Oracle configuration. **Note:** Ignore errors regarding moving files to `*.bak`.
4. Stop the default TNSListener service `Oracle<OraHome>TNSListener`, and set the startup mode to **Manual**.
5. Using **Oracle Net Manager**, configure Oracle to use the LifeKeeper-protected communication resource(s) as follows, if required:
 1. Create a new **TNSListener Service** using the SID name. Configure **Listening Locations**, designating the LifeKeeper-protected IP address and/or named pipe (LAN Manager alias name). Then

configure the **Database Services**, specifying the Oracle Home directory and SID.

2. Modify the **Oracle Service** for your SID. For TCP/IP, change the host name to the protected IP address. For **Named Pipes**, change the machine name to the LAN Manager alias.
6. Create a separate **TNSListener Service** instance for the SID to be protected under LifeKeeper. The service should be created using the `lsnrctl Start <SID>` command. This will create service with the name *Oracle<OraHome>TNSListener<SID>*.
7. Use the **Services** tool to test your Oracle services as follows:
 1. Verify that the new **TNSListener Service** can be stopped and started successfully.
 2. Ensure that the *OracleService<SID>* service has been created by Oracle.
 3. Stop all Oracle services on the backup server.

On the Primary Server

1. Bring the volume resource back in service on the primary server.
2. Start the *OracleService<SID>* service on the primary server.
3. Create the Oracle hierarchy on the primary server and extend it to the backup server. See [Creating an Oracle Hierarchy](#) for details.
4. Test the new Oracle hierarchy by performing a manual failover.

Oracle Resource Configuration Tasks

Once you have completed the setup tasks as described in the previous section, you are ready to create and extend your Oracle resource hierarchies.

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.
- **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 Administrative Tasks section within the LifeKeeper Online Product Manual, 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 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.

Creating an Oracle Hierarchy

After you have completed the necessary setup tasks, use the following steps to define the Oracle Server hierarchy to protect your database(s).

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the menu, select **Create Resource Hierarchy**.



2. The **Create Protected Application** dialog box will display. Select the appropriate **Primary** and **Backup Servers** from the pull-down list. Select **Next** to continue. A window will display with a list of all the recognized Recovery Kits installed within the cluster
3. Select **Oracle** and click **NEXT**.
4. 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
Select the Oracle Home directory	Select the appropriate Oracle Home directory for this hierarchy.
Select the Oracle SID	Select the Oracle SID that you wish to place under LifeKeeper protection.
Enter the Oracle User Name	Enter the administrative user name for Oracle. This user account must have system permissions to the database.
Enter Password	Enter the system password for the Oracle administrative user.

Optional Services	Select the optional services to be protected with this hierarchy. The list includes only those services that are eligible for LifeKeeper protection.
Oracle Tag Name	Enter a unique tag name, or you can accept the default tag name offered by LifeKeeper.

5. After you click **Next**, the **Create Resource Wizard** will create your Oracle resource. LifeKeeper will validate the data entered. If LifeKeeper detects a problem, an error message will appear in the information box.
6. Another information box will appear indicating that you have successfully created an Oracle resource hierarchy, and you must extend that hierarchy to another server in your cluster in order to achieve failover protection. Click **Next**.
7. After you click **Next**, LifeKeeper will launch the **Pre-Extend Wizard**. Refer to [Extending An Oracle Hierarchy](#) for details on how to extend your resource hierarchy to another server.

Extending an Oracle 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**.
- 2. The **Pre-Extend Wizard** will prompt you to enter the following information.

Field	Tips
Backup Priority	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.

- 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**.

Unextending an Oracle 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 Oracle resource. It cannot be the server where the Oracle 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 Oracle 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 Oracle resource hierarchy you have chosen to unextend. Click **Unextend**.
5. Another information box appears confirming that the Oracle resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

Deleting an Oracle Hierarchy

Before deleting an Oracle hierarchy or instance, make sure that the hierarchy is active (green) on its primary server. You may also wish to remove the dependencies before deleting the hierarchy; otherwise, the dependencies will be deleted also.

Deleting an Oracle hierarchy accomplishes the following:

- Stops the Oracle services.
- Deletes the Oracle hierarchy and all dependencies.

Notes:

- Make sure both servers are active when a delete is initiated for LifeKeeper to properly withdraw the databases from the backup server.
- If you want the IP address and volume to remain under LifeKeeper protection, you should delete volume and TCP/IP dependencies prior to deletion.

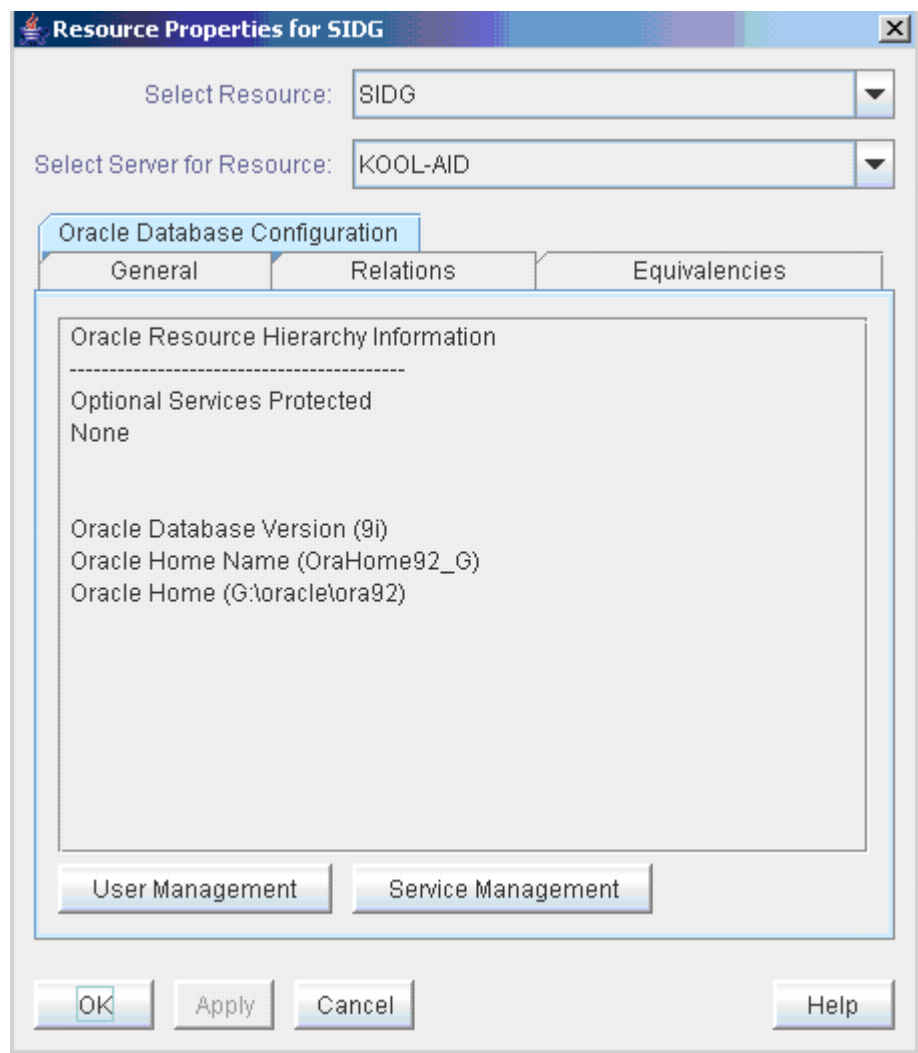
To delete a resource hierarchy from all the servers in your LifeKeeper environment, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the **Target Server** where you will be deleting your Oracle 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 **Next**.
5. Another information box appears confirming that the Oracle resource was deleted successfully.

6. Click **Done** to exit.

Manage Oracle Database Configuration

To administer a protected Oracle resource from the LifeKeeper GUI, right-click on the Oracle resource (on the right-hand side of the LifeKeeper GUI) and select **properties**, then select the **Oracle Database Configuration** tab. Use the **Oracle Database Configuration** page to view or change information about your Oracle resource.



User Management:

This menu allows users to manage the Oracle DBA user that will be used during LifeKeeper operations.

Select Management Action:

- **Show Current User** - Display the current user name used by the protected resource hierarchy.
- **Change Password** - Update the user password for the current user associated with the protected resource hierarchy.
- **Change User and Password** - Update both the Oracle DBA user and password to be used during LifeKeeper operations to administer and monitor the Oracle instance. The user must have DBA privileges for all databases under protection.

Field	Tips
Enter User Name	Enter the administrative user name. This user account must include DBA permissions to all databases under LifeKeeper protection.
Enter Password	Enter the administrative password for the user account being updated.

Service Management:

This menu allows users to modify the list of optional Oracle Services that are protected under the resource hierarchy. LifeKeeper will monitor all protected optional services.

Select Service Action:

- **Add Service** - Add an additional service to the protected configuration. LifeKeeper will start monitoring added optional Oracle service.
- **Delete Service** - Remove a service from the protected configuration. LifeKeeper will stop monitoring optional Oracle service.

Field	Tips
Service Name	Enter the service name for the service to Add or Delete from the protected configuration. For Add operations, enter the service name. For Delete operations, choose the service name to remove from the list provided.
Update Cluster	Select Yes to update all systems in this cluster. Otherwise, select No to only update the current system. If you chose No , you must manually add the service to the backup servers.

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.

You can select **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, and the Oracle services are stopped.

Oracle Hierarchy Administration

The topics in this section assist in the administration of Oracle hierarchies.

Oracle Hierarchy Administration Guidelines

Follow these guidelines for administering your Oracle hierarchy:

- **Access via protected communication paths.** To ensure that users can access the Oracle SID, regardless of the physical system on which it is running, all remote access of the database should be done through the protected Named Pipe (LAN Manager alias) or IP addresses, which are part of the Oracle hierarchy. LifeKeeper automatically makes protected communication paths available on the backup system in case of a switchover.
- **Reserve volumes for exclusive Oracle use.** Reserve volumes containing the Oracle database files for use exclusively by Oracle. They should not be shared for users to access via LAN Manager, and should not be accessed by any other local applications. This is because LifeKeeper operations that remove a volume resource from service, for example in a failover, can fail if a remote user is accessing one of the volumes over the network or if a local process has done an open for write access on the volume.

Local processes that have read-only access to volumes do not prevent removal of a resource from service, but the read-only access may cause a restore to fail when you attempt to switch the resource back. Examples of processes with read-only access are the Performance Monitor, which periodically polls each volume, or any running process which is installed on the shared volume.

- **Start and Stop Oracle Through LifeKeeper.** Although much of your administration of Oracle databases is done through the Oracle tools, use the LifeKeeper Out of Service function to stop the Oracle SID and use the In Service function to start the Oracle SID. When LifeKeeper stops and starts the SID, it maintains a consistent view of the server on all nodes in the configuration.
- **Protect volume resources before adding to Oracle SID.** As your environment grows, if you need to add new volumes to the Oracle SID already under LifeKeeper protection, you should do the following:
 1. Protect the volume first (create a volume resource).

2. Add the volume to the SID.
3. Manually create a dependency between the Oracle resource and the volume resource.

LifeKeeper Oracle Recovery Kit

Recovery Variables

The LifeKeeper Oracle Recovery Kit installation creates 3 registry entry variables stored in the following registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\SIOS\LifeKeeper\RK\ORAapp

MAXWAIT is a decimal integer that specifies the number of seconds that the recovery kit will wait for a single Oracle service to start or stop. If the service has not started within the specified time frame, LifeKeeper will mark the resource as failed. The default value for *MAXWAIT* is **300**; however, it is possible that for extremely large databases, 300 seconds might not be enough time for the database services to reach the *STARTED* or *STOPPED* state. If this is the case, change this registry entry to a reasonable value.

RESTORE_DEEPCHK_MAX_RETRY is a decimal integer that allows multiple attempts to verify the Oracle service state during a restore or local recovery operation. On a server that is unexpectedly heavily loaded, the default service state check time may not always be sufficient to verify that protected Oracle services are in the *RUNNING* state. The default value for this variable is **0** and normally only 1 Oracle service state check attempt is performed for each service. This value can be changed if extra attempts may be needed to verify the Oracle service state.

RESTORE_DEEPCHK_SLEEP is a decimal integer, measured in seconds, to insert sleep intervals between each extra attempt to verify the Oracle service state during a restore or local recovery operation. This option is enabled if the *RESTORE_DEEPCHK_MAX_RETRY* option described above is used. The default value for this variable is **0** and normally no sleep times are inserted between extra Oracle service state check attempts. If the *RESTORE_DEEPCHK_MAX_RETRY* variable is set, it is highly recommended that the *RESTORE_DEEPCHK_SLEEP* variable be set as well to improve the reliability and performance of Oracle service state checks.

Updating Oracle Username and Password for LifeKeeper

During the creation of a LifeKeeper Oracle resource, the user must enter an Oracle user name and password for that instance of Oracle. Should the password of this user name change at some point in the future, the LifeKeeper Oracle resource must be updated on all systems in the cluster with this new password. Failure to do so will leave the Oracle resource out of sync and will prevent it from coming in and out of service properly. LifeKeeper will log an error message to the **Application Event Log** stating that LifeKeeper cannot remove or restore the resource during any subsequent failover or manual switchover.

The LifeKeeper GUI provides an interface to manage the user account associated with the Oracle resource. See the section on [Manage Oracle Database Configuration](#) for more information.

Manually Configure Oracle 11g DB Console

You must manually configure Oracle 11g DB Console to use the protected virtual IP address or LAN Manager Alias so that the service will start successfully on all systems in the cluster.

On the primary, edit `<ORACLE_HOME>\<SYSTEM_NAME>_<ORACLE_SID>\sysman\config\emoms.properties`

Set the property `oracle.sysman.emSDK.svlt.ConsoleServerHost=<virtual IPaddress>` or `<LAN Manager Alias>`.

Oracle Troubleshooting Tips

This section is intended to provide suggestions and insights into occurrences that are not specifically related to the LifeKeeper software, but have a relationship with the total environment.

Create Hierarchy Failed

Suggestion

Check the following:

- All volumes and communications resources (IP and/or LAN Manager) associated with the SID should already be under LifeKeeper protection.
- All the shared or replicated volumes are available to the primary server and all volumes are mapped to the same drive letters on each server.

Bring In Service Failed

Suggestion

Check to see if any other Oracle resource is already in service on the system you are trying to bring in service and/or that shared volumes can be accessed from this system.

If additional time is required for the Oracle service to reach the running state, consider using the LifeKeeper Oracle Recovery Kit Variables in the registry to extend the allowed service startup time interval. [LifeKeeper Oracle Recovery Kit Variables](#)

Oracle TNSListener Service is Not Started or Stopped as it Should Be

Suggestion

Do the following:

- Check if at least one TCP/IP or LAN Manager resource is part of the Oracle hierarchy.
- Take hierarchy out of service and bring back in service to start *Oracle<OraHome>TNSListener<SID>*.

Server Not Responding

TCP/IP Client Cannot Access Server (Server Not Responding) After a Successful Switchover by LifeKeeper

Insight

The client system has old information in its IP-to-Physical address translation table used by address resolution protocol (arp).

Suggestion

The IP address being used to access the server must be reset. To reset this address, issue the command `arp -d server_ip_address`. This deletes the address from the translation table. On the next request of that IP address, the table entry will be filled.

Remote Users Cannot Log In ORA 12504 or ORA 12514 or ORA 12541

Insight

Oracle connect issues indicate Oracle TNS setup issues are not caused by LifeKeeper.

Suggestion

Do the following:

- Read the specific Oracle error and message
- Using **Oracle Net Manager**, make certain that the listener and the TNS names are configured correctly.
- Make certain the service is configured using the IP or LAN Manager alias created.
- Use TNSPING to reach the service.

SIOS Protection Suite Microsoft Internet Information Services Recovery Kit Introduction

The SIOS Protection Suite Microsoft IIS Recovery Kit extends your SIOS Protection Suite product by adding specific protection to Internet servers. The SIOS Protection Suite Microsoft IIS Recovery Kit continuously monitors the health of your Internet servers, and if a problem arises, provides automatic failover of the affected sites to a standby system. The Recovery Kit protects Web, FTP and SMTP sites. When multiple IIS resources are configured, Web, FTP or SMTP sites can be stopped/started independently.

IIS Overview















The SIOS Protection Suite Microsoft IIS Recovery Kit protects Internet servers from the following problems:

- System failure or server shutdown
- Network Interface Card (NIC) failures
- Communication failures (Web server is running but stops responding)
- Startup failures (Web server aborts on startup)

The SIOS Protection Suite Microsoft IIS Recovery Kit has two recovery procedures. For system or NIC failures, the Recovery Kit transfers the affected web server's IP address to a standby system, and then starts up the standby web server. If there is a communication or startup failure, and local recovery is enabled, the SIOS Protection Suite Microsoft IIS Recovery Kit will first stop and restart the affected web server locally to see if that corrects the problem. If the restart is unsuccessful, then the Recovery Kit transfers service to the backup web server.

The SIOS Protection Suite Microsoft IIS Recovery Kit manages the dependencies between the IIS application, IP and volume resources. First, you create the IP and volume resources to be used by your web servers. Then, when you create the IIS resource, the SIOS Protection Suite Microsoft IIS Recovery Kit reads the Microsoft IIS configuration and automatically creates the required dependencies between the IIS resource and the IP and volume resources.

The following is a sample IIS hierarchy as shown in the LifeKeeper GUI. The Web site has dependencies on both the IP address "Switchable113", and on the volume "WEB.Vol.X", where the home directory containing the Web site's content resides. Both the IP and volume resources were created prior to the web site creation.

		MyFTPSite		Active	1
		FTP.Vol.Y		Active	1
		Switchable113		Active	1
		MyWebSite		Active	1
		Switchable113		Active	1
		Web.Vol.X		Active	1

Installation

The topics in this section will assist in installing the LifeKeeper Microsoft IIS Recovery Kit.

[Hardware and Software Requirements](#)

[Kit Installation](#)

[Kit Removal](#)

[Configuration Definitions and Restrictions](#)

[IIS Configuration Considerations](#)

[Installing and Configuring IIS with LifeKeeper](#)

[Install and Configure SIOS Protection Suite and Recovery Kits](#)

[Install and Configure Microsoft IIS on All Servers](#)

[Install and Configure SMTP Virtual Server](#)

Hardware and Software Requirements

Before attempting to install or remove the SIOS Protection Suite Microsoft IIS Recovery Kit, be sure that your configuration meets the following requirements:

- **Operating System software.** Refer to the [SIOS Protection Suite for Windows Support Matrix](#) for a list of supported operating systems.
- **SIOS Protection Suite software.** You must install the same version of SIOS Protection Suite software and any patches on each server. Please refer to the [Release Notes](#) for specific requirements.
- **SIOS DataKeeper software (optional).** If you plan to use IIS with replicated volumes rather than shared storage, you should install the SIOS DataKeeper for Windows software on each server.
- **SIOS Protection Suite IP Recovery Kit.** You must have the SIOS Protection Suite IP Recovery Kit installed on each server. All TCP/IP configuration requirements for the IP Recovery Kit also apply to the SIOS Protection Suite Microsoft IIS Recovery Kit.
- **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: throughput requirements, elimination of single points of failure, network segmentation, and so forth.
- **TCP/IP protocol.** Each server requires TCP/IP to be installed and configured properly.
 - The two servers must be on the same LAN segment (that is, no routers between the two systems).
 - Free IP addresses to create SIOS Protection Suite hierarchies:
 - For each active Microsoft IIS site, you will need one switchable IP address to be shared between the active and standby IIS site.

- If you plan to protect multiple Microsoft IIS sites, you will need additional IP addresses for each protected IIS resource.
- **IIS Software.** The IIS Recovery Kit supports the native version of IIS on any supported version of Microsoft Windows.

Microsoft IIS software, IIS server roles, FTP feature or SMTP feature must be installed and configured prior to configuring SIOS Protection Suite and the SIOS Protection Suite Microsoft IIS Recovery Kit. The same version should be installed on each server. Please refer to [IIS Required Roles and Role Services and Features](#) for requirements that must be installed on each server.

Kit Installation

The SIOS Protection Suite Microsoft IIS Recovery Kit is included with the SIOS Protection Suite for Windows core product which is available on CD-ROM or via ftp download.

Kit Removal

!CAUTION: Be sure to remove all SIOS Protection Suite IIS resources from service and delete them prior to removing the recovery kit. Once the kit is removed these resources will be unusable.

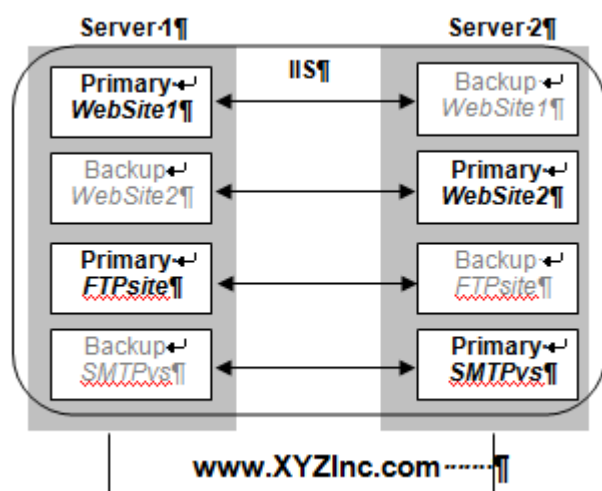
The SIOS Protection Suite Microsoft IIS Recovery Kit is included with the SIOS Protection Suite for Windows core product and will be uninstalled when removing the core product.

Configuration Definitions and Restrictions

Active/Active Configuration

IIS allows multiple Web and FTP sites and SMTP virtual servers to run on each server in a cluster. Only a single instance of the IIS software is allowed (or required) on a given system. The Web sites, FTP sites and SMTP virtual servers can be protected and managed individually by LifeKeeper. Subsequent descriptions may use the term "site" to refer generically to a Website, FTP site or SMTP virtual server.

The figure below illustrates a typical configuration of web servers.



In this configuration, each server has two Web sites: one primary and one backup Web site. Server 1 has the primary instance of *WebSite1* and the backup instance of *WebSite2*. Server 2 has the reciprocal configuration: a primary instance of *WebSite2* and a backup instance of *WebSite1*. Only the primary instances of the Web sites actually service incoming user requests on any given server.

In addition, Server 1 has the primary instance of an FTP site named *FTPsite* and the backup instance of the SMTP virtual server named *SMTPvs*, while Server2 has the backup instance of the FTP site and the primary instance of the SMTP virtual server.

If a primary Web site stops servicing user requests, LifeKeeper activates the backup instance on the backup server to resume service there. Thus, if *WebSite1* on Server 1 fails, then LifeKeeper activates the backup instance of *WebSite1* on Server 2. After the switchover, there will be two active instances running on Server 2. Once the problem with the failed web server is corrected, you may switch service back to Server 1. The LifeKeeper Microsoft IIS Recovery Kit allows you to manually switch service back or take advantage of the LifeKeeper automatic switchback feature.

IIS Configuration Considerations

LifeKeeper places certain restraints on your Internet server configurations. These restrictions will ensure that the standby web server/site can successfully and completely replace the active web server/site.

Default Web Site or New Web Site

The Default Web Site created by the IIS installation process may be protected by LifeKeeper with one minor configuration change. The Default WebSite must be reconfigured to use a LifeKeeper protected IP address for the site.

LifeKeeper can also protect new Web Sites that have been configured to use a LifeKeeper protected IP address for the site.

Primary and Backup Designations

The server where the active site is created will be the primary LifeKeeper server for this Web site. The server where the standby site is created will be the backup server for this site. Keep in mind that the designations "primary" and "backup" server change for each site you are configuring.

Naming Restrictions

In order to receive LifeKeeper protection, you should adhere to the following rules for site name (which is entered in the Description field of the IIS console):

- Use only alphanumeric characters and dashes (should NOT contain spaces).
- If you need to change the name (description) of a protected Web Site, first delete the LifeKeeper IIS resource, then change the description, and recreate the resource.

Identical Primary/Backup Web Sites

For each primary IIS site, you must create an identical backup IIS site on the other server. These two servers must be connected by a LifeKeeper heartbeat. In order for the primary and backup sites to be identical, the following criteria must be met:

- The site names entered in the *Description* field of the Properties form must be identical, including using the same case.
- The switchable IP addresses, port, and header assigned to the sites in the Properties form must be identical.
- If using a shared or replicated volume for your web or FTP content, the drive letter and folder of the volume you assign in the *Home Directory Path* must be identical.
- If you configure multiple backup sites for a particular Web site, then you must configure the other Web sites with the same identities; that is, the primary and backup Web sites must contain the same IP addresses, ports, and headers.
- If you configure one Web site as a secure Web site, then you must configure the other Web site as a secure Web site. Additional limitations apply to secure Web sites. See the following section for details.

Configuring Secure Servers

A secure server is a web server that uses Secure Socket Layers (SSL) for communication. Security is improved because the data sent and received are encrypted, and because the web client and the web server can identify one another. Secure servers use *https:* rather than *http:* in their URL. The default port number for a secure server is 443.

With regards to the LifeKeeper Microsoft IIS Recovery Kit and LifeKeeper, there is no difference in running a secure IIS Web site. In fact, IIS allows the same Web site to have both a TCP port and SSL port. There is no change in the startup or operational procedures. Therefore, after a key is generated and a corresponding digital certificate is installed in IIS, you may configure and run with SSL ports.

IIS Configuration

The following configuration rules must be followed to ensure LifeKeeper protection:

- IIS sites that do not have IP addresses specified in the "IPAddress" field of the Properties or Bindings form cannot be protected.
- If using a shared or replicated volume for your web content, the Home Directory should be specified as "A directory located on this computer". LifeKeeper will not be able to protect the Home Directory if specified as either of the following:
 - a share located on another computer
 - a redirection to a URL
 - a volume that is not protected by LifeKeeper

Document Content Location

Shared and Replicated Content Storage

If the content volume is on a shared or replicated volume, both Web sites must point to the same shared or replicated volume and folder. The primary and backup servers must contain the same content files for the active and standby web servers/sites to be identical. However, if the content volume is not shared or replicated, the content may come from any location on either system.

To ensure data availability on a failover we suggest that you configure the Home Directory on the primary server as a folder on a shared or replicated disk and configure the Home Directory on the backup server identical to the primary server. You then have only one copy of the content files to maintain.

Non-Shared Storage


If your configuration does not utilize shared storage, then the content must be synchronized between local volumes on each server. While the LifeKeeper Microsoft IIS Recovery Kit does not contain any specific features to synchronize the content between two servers, the following are a few suggestions:

- Use SIOS DataKeeper to automatically replicate the data volumes on each active server to the standby server(s).
- Use a content replication tool such as Microsoft Site Server 3.0. You can also use the utility *Robocopy* as a content replication tool. Microsoft Site Server is the preferred solution.
- If you have a tape backup system, make a tape backup of the files on the primary server, and then restore them to the backup server, as needed.

Use Different Volume for Multiple IIS Sites

When the LifeKeeper Microsoft IIS Recovery Kit creates an IIS resource hierarchy, it creates dependencies associated with the IP address and content volume using the home directory path designated in the IIS configuration. We commend that if you protect multiple sites, then you should designate DIFFERENT IP addresses and volumes for each site.

For example, the hierarchy shown below shows both *MyFTPSite* and *MyWebSite* utilizing the same IP address and different volume resources. Any maintenance done on one site will affect the other site since these have common IP resource dependency.



Bringing *MyFTPSite* In Service on the backup server will also move its dependencies to the backup server. This causes *MyWebSite* to be taken out of service on the primary server. You would then need to manually bring *MyWebSite* In Service on the backup server.

Assigning DIFFERENT IP addresses and volumes to each protected IIS site will give you more flexibility in managing your resources by NOT tying their recovery actions together. However, you may prefer to have them grouped as shown above.

Installing and Configuring IIS with LifeKeeper

Before proceeding, you should have already configured your storage and networking according to the recommendations in the previous chapters of this guide.

Installation Checklist

The installation and setup sequence should be performed in the following order (more detailed instructions for each of these steps are provided in other topics):

1. If using replicated volumes, install SIOS DataKeeper software on each server and create your mirrors.
2. Install and configure the LifeKeeper Core, which includes the LifeKeeper IP Recovery Kit and LifeKeeper Microsoft IIS Recovery Kit, on each server.
3. Install and configure Microsoft IIS on all servers.

Install SIOS DataKeeper and Create Mirrors

If you will be using replicated volumes, you should now install the SIOS DataKeeper for Windows software and create your mirrors. Refer to the [SIOS Protection Suite Installation Guide](#) for more details.

Install and Configure SIOS Protection Suite and Recovery Kits

The next step is to install the SIOS Protection Suite Core, which includes the IP Recovery Kit and Microsoft IIS Recovery Kit, on all servers. See the [SIOS Protection Suite Installation](#) topics for details on installing SIOS Protection Suite. You must have the same version of SIOS Protection Suite on all servers.

After you have installed LifeKeeper, you will need to reboot all the servers. After rebooting, make sure LifeKeeper is up and running on all servers.

You can now configure LifeKeeper. The LifeKeeper setup tasks are given in the proper sequence below. For detailed instructions on the LifeKeeper configuration tasks, use the **Help** button or refer to [LifeKeeper Configuration](#).

*Switchover on Shutdown with DataKeeper is not supported. See [Set Server Shutdown Strategy](#) for more information.

1. Set up your LifeKeeper communication paths. For the most reliable communication path configuration, we recommend creating two separate TCP/IP comm paths, and if possible, a third comm path. For TCP/IP comm paths, the best results are obtained when you use a private network between the two servers.
2. Create switchable IP addresses for each web server/site pair as required. On the **Edit** menu, select **Server**, then **Create Resource Hierarchy**. From the drop down list, select **IP Address**, and then fill in the information required by the **Protected Application Wizard**. Repeat for each switchable IP address needed.

When creating LifeKeeper switchable IP addresses, consider the following:

- a. The primary server is the one that normally runs the active Webserver/site. It should be set to priority 1 so that the IP resource matches the IIS resource to be
- b. If desired, change the Switchback Strategy from *Intelligent* (the default) to *Automatic**

c. If you have two NICs on the same subnet, you can set the IP Local Recovery* option to have LifeKeeper transfer service of the switchable IP address between the two cards for increased availability.

*See related topics for additional information on [Switchback Strategy](#) and [IP Local Recovery](#).

4. Test your switchable IP addresses for switchover and response.
 1. To test switchover, open the LifeKeeper GUI. Your switchable IP resources should display as green (Active) on the primary server and blue(Standby) on the backup. Right click on the IP instance in the hierarchy tree.From the pop-up menu, select In Service, and select the backup server from the list box. This switchable IP resource will turn from blue to green on the backup server. Repeat this test on any remaining switchable IP addresses. When you are finished testing all the switchable IP addresses, bring them back In Service on their primary servers.
 2. To test response, open an MS-DOS window and use the ping command on each switchable IP address. Your switchable IP addresses should return a response time and packet loss value for each ping.
 3. Do not proceed until your switchable IP addresses pass both the switchover and ping tests successfully.
5. Create your Volume resource(s) which will contain the home directories for your Web/FTP/SMTP services. Perform Volume resource switchovers to ensure that that your volume(s) can be placed in service on primary and backup servers. Also make sure that the priorities you assign to protected IP and Volume resources match on each server.

Install and Configure Microsoft IIS on All Servers

The next step is to configure Microsoft IIS on all the servers in your cluster.

Install and Configure Microsoft IIS Web or FTP Site

To create a Web site or FTP site pair, you need to add a new site on each server and configure the pair to be identical, or you can choose to protect the Default Web and FTP sites:

1. Use the LifeKeeper GUI to bring the switchable IP address and protected Volume resource, if any, in-service on the primary server.
2. Launch the Internet Information Services (IIS) Manager (from Administrative Tools in the Program menu). You will use this interface to create the primary and backup sites.
3. Launch the New Web Site Wizard and enter the Site Name or Description, remembering that the name must be identical on both primary and backup servers.
4. **Web site only:** Enter the TCP port and Host Header fields. They will need to be identical on primary and backup servers.
5. Select or enter the Switchable IP Address for the Web site Note that a protected switchable IP address may not appear in the drop-down list if it is out-of-service, but you can type it in. The IP address for the web site must be identical on primary and backup servers.
6. **Web site only:** If you plan on configuring your Web sites with multiple identities, enter the switchable IP address in the IP Address field for each multiple identity.
7. The Home Directory or Physical Path for the site can be local (not protected by LifeKeeper) or LifeKeeper shared or replicated storage may be used. If using shared or replicated storage, the volume must

be in-service and this directory must already exist.

8. **Web site only:** Select your Web Site Access Permissions.
9. Ensure the Microsoft IIS site is started, and then test it to ensure it is accessible and working properly before proceeding with web site setup on the backup server.
10. Use the LifeKeeper GUI to bring the switchable IP address(es) and protected volume(s), if any, in-service on the backup server.
11. Repeat steps 2-10 to create the site on the backup server. Remember that the Web or FTP site pairs must be identical on the primary and backup servers.

Once you have created and configured all your IIS sites, do the following:

1. Use the LifeKeeper GUI to bring the switchable IP address and protected volume, if any, in-service on their primary server(s).
2. Use the Internet Information Services (IIS) Manager to start the primary site(s) on their primary server(s), and stop the backup sites on the backup servers.

***Note:** The Microsoft Management Console IIS screen may show all available IP addresses for a site even if the LifeKeeper-protected IP address is NOT in service on that server. This is a bug in the Microsoft Management Console snap-in. It does NOT affect operations.

If your protected FTP site will not permit anonymous logins, the LifeKeeper deep check process must be configured to perform non-anonymous logins. See [Protecting FTP Sites with Non-Anonymous Login](#) for more information.

If you wish to create an SMTP Virtual Server, then proceed to [Install and Configure SMTP Virtual Server](#). Otherwise, you are ready to add LifeKeeper protection to the Internet servers by configuring one or more IIS resource hierarchies.

Install and Configure SMTP Virtual Server

Follow the steps below to create an SMTP Virtual Server. You need to add a new virtual server on each system and configure the pair to be identical.

1. Use the LifeKeeper GUI to bring the switchable IP address and protected volume, if any, in-service on the primary server.
2. Launch the Internet Information Services (IIS) Manager for SMTP sites (from Administrative Tools in the Program menu). You will use this interface to create the primary and backup SMTP Virtual Server.
3. Launch the New Web Site Wizard by selecting the server name, then on the Action menu, select New, and then SMTP Virtual Server.
4. Enter the site description, remembering that the description should be identical on both primary and backup servers.
5. Select the switchable IP address for the SMTP Virtual Server when prompted to "Select the IP address to be used for this Web site." It must be the same on primary and backup servers.
6. When prompted for the home directory path, choose a file share or volume on your shared or replicated storage device.
7. Enter the default domain for this virtual server.
8. Ensure the SMTP virtual server is started, and then test it to ensure it's accessible and working properly before proceeding to setup on the backup server.
9. Use the LifeKeeper GUI to bring the switchable IP address and protected volume, if any, in-service on the backup server.
10. Repeat steps 2-9 to create the site on the backup server. Remember that site pairs must be identical on the primary and backup servers.

Once you have created and configured all your SMTP virtual servers, start

the primary sites and stop the backup Web sites as follows:

1. Use the LifeKeeper GUI to bring the switchable IP address and protected volume, if any, in- service on the primary server(s).
2. Use the Internet Information Services (IIS) Manager to start the primary SMTP site on the primary server.

Now you are ready to add LifeKeeper protection to the Internet servers by creating one or more IIS resource hierarchies.

IIS Configuration Definitions and Restrictions

The topics in this section will assist in configuring the SIOS Protection Suite Microsoft Internet Information Services Recovery Kit.

[IIS Required Roles and Role Services and Features](#)

[IIS Active Active](#)

[Primary and Backup Designations](#)

[Naming Restrictions](#)

[Identical Primary Backup Web Sites](#)

[Configuring Secure Servers](#)

[IIS Configuration](#)

[Document Content Location](#)

[Use Different Volume for Multiple IIS Sites](#)

IIS Required Roles and Role Services and Features

LifeKeeper interfaces to IIS require the following roles, role services and features to be installed on the server:

Roles:

- Web Server (IIS)

Role Services:

- IIS Management Console
- FTP Server (If protecting FTP Sites)
- FTP Service

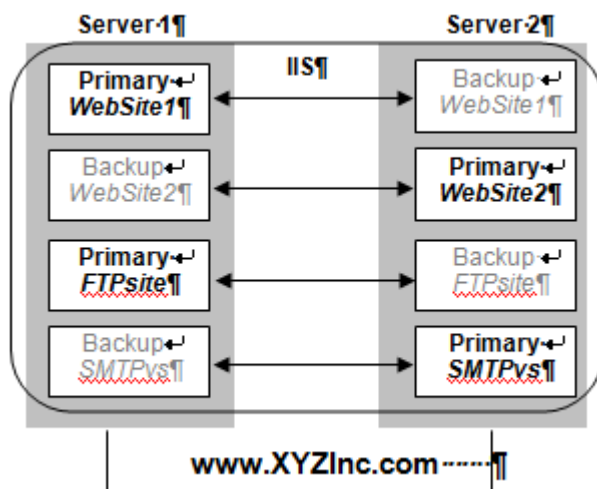
Features:

- SMTP Server (If protecting SMTP sites)

IIS Active Active

IIS allows multiple Web and FTP sites and SMTP virtual servers to run on each server in a cluster. Only a single instance of the IIS software is allowed (or required) on a given system. The Web sites, FTP sites and SMTP virtual servers can be protected and managed individually by LifeKeeper. Subsequent descriptions may use the term "site" to refer generically to a Website, FTP site or SMTP virtual server.

The figure below illustrates a typical configuration of web servers.



In this configuration, each server has two Web sites: one primary and one backup Web site. Server 1 has the primary instance of *WebSite1* and the backup instance of *WebSite2*. Server 2 has the reciprocal configuration: a primary instance of *WebSite2* and a backup instance of *WebSite1*. Only the primary instances of the Websites actually service incoming user requests on any given server.

In addition, Server 1 has the primary instance of an FTP site named FTP site and the backup instance of the SMTP virtual server named *SMTPvs*, while Server 2 has the backup instance of the FTP site and the primary instance of the SMTP virtual server.

If a primary Web site stops servicing user requests, LifeKeeper activates the backup instance on the backup server to resume service there. Thus, if *WebSite1* on Server 1 fails, then LifeKeeper activates the backup instance of *WebSite1* on Server 2. After the switchover, there will be two active instances running on Server 2. Once the problem with the failed

web server is corrected, you may switch service back to Server 1. The LifeKeeper Microsoft IIS Recovery Kit allows you to manually switch service back or take advantage of the LifeKeeper automatic switchback feature.

Primary and Backup Designations

The server where the active site is created will be the primary LifeKeeper server for this Web site. The server where the standby site is created will be the backup server for this site. Keep in mind that the designations "primary" and "backup" server change for each site you are configuring.

Naming Restrictions

In order to receive LifeKeeper protection, you should adhere to the following rules for site name (which is entered in the Description field of the IIS console):

- Use only alphanumeric characters and dashes.
- If you need to change the name (description) of a protected Web Site, first delete the LifeKeeper IIS resource, then change the description, and recreate the resource.

Identical Primary Backup Web Sites

For each primary IIS site, you must create an identical backup IIS site on the other server. These two servers must be connected by a LifeKeeper heartbeat. In order for the primary and backup sites to be identical, the following criteria must be met:

- The site names entered in the **Description** field of the **Properties** form must be identical, including using the same case.
- The switchable IP addresses, port and header assigned to the sites in the **Properties** form must be identical.
- If using a shared or replicated volume for your web or FTP content, the drive letter and folder of the volume you assign in the **Home Directory Path** must be identical.
- If you configure multiple backup sites for a particular Web site, you must configure the other Web sites with the same identities; that is, the primary and backup Web sites must contain the same IP addresses, ports and headers.
- If you configure one Web site as a secure Web site, then you must configure the other Web site as a secure Web site. Additional limitations apply to secure Web sites. See [Configuring Secure Servers](#) for details.

Configuring Secure Servers

A secure server is a web server that uses Secure Socket Layers (SSL) for communication. Security is improved because the data sent and received are encrypted and because the web client and the web server can identify one another. Secure servers use https: rather than http: in their URL. The default port number for a secure server is 443.

With regards to the LifeKeeper Microsoft IIS Recovery Kit and LifeKeeper, there is no difference in running a secure IIS Web site. In fact, IIS allows the same Web site to have both a TCP port and SSL port. There is no change in the startup or operational procedures. Therefore, after a key is generated and a corresponding digital certificate is installed in IIS, you may configure and run with SSL ports.

IIS Configuration

The following configuration rules must be followed to ensure LifeKeeper protection:

- IIS sites that do not have IP addresses specified in the "IPAddress" field of the **Properties** or **Bindings** form cannot be protected.
- If using a shared or replicated volume for your web content, the **Home Directory** should be specified as "**A directory located on this computer**". LifeKeeper will not be able to protect the **Home Directory** if specified as any of the following:
 - a share located on another computer
 - redirection to a URL
 - volume that is not protected by LifeKeeper

Document Content Location

Shared and Replicated Content Storage

If the content volume is on a shared or replicated volume, both Web sites must point to the same shared or replicated volume and folder. The primary and backup servers must contain the same content files for the active and standby web servers/sites to be identical. However, if the content volume is not shared or replicated, the content may come from any location on either system.

To ensure data availability on a failover we suggest that you configure the **Home Directory** on the primary server as a folder on a shared or replicated disk and configure the **Home Directory** on the backup server identical to the primary server. You then have only one copy of the content files to maintain.

Non-Shared Storage


Your configuration should utilize shared storage or replicated storage. If your configuration does not utilize shared or replicated storage, the content must be synchronized between local volumes on each server. While the LifeKeeper Microsoft IIS Recovery Kit does not contain any specific features to synchronize the content between two servers, the following are a few suggestions:

- Use [SIOS DataKeeper](#) to automatically replicate the data volumes on each active server to the standby server(s).
- Use a content replication tool such as Microsoft Site Server 3.0. You can also use the utility **Robocopy** as a content replication tool. Microsoft Site Server is the preferred solution.
- If you have a tape backup system, make a tape backup of the files on the primary server and then restore them to the backup server, as needed.

Use Different Volume for Multiple IIS Sites

When the LifeKeeper Microsoft IIS Recovery Kit creates an IIS resource hierarchy, it creates dependencies associated with the IP address and content volume using the home directory path designated in the IIS configuration. We recommend that if you protect multiple sites, then you should designate DIFFERENT IP addresses and volumes for each site.

The hierarchy shown below shows both MyFTPSite and MyWebSite utilizing the same IP address and different volume resources. Any maintenance done on one site will affect the other site since these have common IP resource dependency.



Bringing *MyFTPSite* In Service on the backup server will also move its dependencies to the backup server. This causes *MyWebSite* to be taken out of service on the primary server. You would then need to manually bring *MyWebSite* In Service on the backup server.

Assigning DIFFERENT IP addresses and volumes to each protected IIS site will give you more flexibility in managing your resources by NOT tying their recovery actions together. However, you may prefer to have them grouped as shown above.

IIS Resource Configuration Tasks

The topics in this section describe the tasks involved in configuring resources.

[Create an IIS Resource Hierarchy](#)

[Extend an IIS Resource Hierarchy](#)

[Delete an IIS Resource Hierarchy](#)

[Unextend Your IIS Hierarchy](#)

Create an IIS Resource Hierarchy

Before creating your IIS hierarchy, be sure you have created the associated IP and volume (if needed) resource hierarchies first.

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**. Click **Next** after selecting the serves for your protected application.

A dialog box will appear with a drop down list box with all the recognized applications you can protect within the cluster.

***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. The **Create Protected Application** window will appear. Select **IIS** from the drop down list for **Application to Protect** and click **NEXT**.
3. Select the Service Type (**WEB**, **FTP** or **SMTP**) and click **NEXT**.
4. Accept the Site Name listed or select enter a new one from the drop down list and click **NEXT**. LifeKeeper generates this list from the IIS configuration information.

Accept the Site Tag offered by LifeKeeper that is the same as the Site Name or enter a new **Site Tag** and click **NEXT** to create the IIS resource.

***Note:** The tag name must consist of printable ASCII characters.

LifeKeeper will validate that you have provided valid data to create your resource hierarchy. If LifeKeeper detects a problem, an error message will appear in the information box. If the validation is successful, your resource will be created.

Extend an IIS Resource Hierarchy

This operation can be started from the **Edit** menu or initiated automatically upon completing the **Create IIS 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.
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
Primary Server	Enter the server where your IIS resource is currently in service.
Resource Hierarchy to Extend	Select the IIS resource hierarchy to be extended.
Backup Server	Select a server from the list of connected servers for which you have Administrator permission to be the backup server for the IIS 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 which cannot be edited. Click **Next** to extend each of the dependencies.
5. Select a priority for the IIS resource on the backup server. Click **Extend** to extend the IIS resource to the backup server.
6. Click **Finish** to complete the extend process.

Delete an IIS 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 **Target Server** where you will be deleting your IIS resource hierarchy. Click **Next** to proceed to the next dialog box.
Note: This dialog will not appear if you selected the **Delete Resource** task by right-clicking on an individual resource instance in the right pane or on a global resource in the left pane where the resource is on only one server.
3. Select the hierarchy to delete. Remember that the list box displays every hierarchy on the target server (i.e. in-service and out-of-server). If you want to stop the IIS 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.
Note: This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a global resource in the left pane or an individual resource instance in the right pane.
4. Click **Next** to proceed to the next dialog box.
5. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete** to remove the IIS resource from LifeKeeper protection. The IIS resource and all its dependencies will be deleted from LifeKeeper protection. Another information box appears confirming that the IIS resource was deleted successfully.
6. Click **Done** to exit the **Delete Resource Hierarchy** menu selection.

Unextend Your IIS Hierarchy

To unextend your IIS hierarchy from a system:

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 IIS resource. It cannot be the server where the IIS resource is currently in service.

Note: The dialog to select the target server will not appear if you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance.

3. Click **Next** to proceed to the next dialog box.
4. Select the IIS resource hierarchy to unextend. Click **Next** to proceed to the next dialog box.

Note: This dialog will not appear if you selected the **Unextend** task by right-clicking on a global resource in the left pane or an individual resource instance in the right pane.

5. An information box appears confirming the target server and the IIS resource hierarchy you have chosen to unextend. Click **Unextend**.
6. Another information box appears confirming that the IIS resource was unextended successfully. Click **Done** to exit.

Testing Your IIS Resource Hierarchy

The topics in this section assist in testing your IIS resource hierarchy.

[Performing a Manual Switchover from the GUI](#)

[Recovery Operations](#)

Performing a Manual Switchover from the GUI

You can initiate a manual switchover from the LifeKeeper GUI by 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.

Recovery Operations

When the primary server fails, the LifeKeeper Microsoft IIS Recovery Kit software performs the following tasks:

- Brings the IIS resource hierarchy into service on the backup server by bringing In Service the IP address(s) on one of that server's physical network interfaces
- Unlocks the shared or replicated volume - if one is being used - for the backup server and locks it for the primary server
- Starts the IIS Web site on the backup server

After recovery, web server users may reconnect by clicking on the **Reload/Refresh** button of their browsers.

IIS Hierarchy Administration

The topics in this section assist in the administration of IIS hierarchies.

[Modifying Quick Check Interval, Deep Check Interval and Local Recovery](#)

[Manual Switchover](#)

[IIS Failover](#)

[Protecting FTP Sites with Non-Anonymous Login](#)

[Using an FTP Login Script](#)

[Disabling the FTP Deep Check Process](#)

[Changing LifeKeeper Microsoft IIS Recovery Kit Configuration](#)

[Removing Microsoft IIS](#)

[Changing the Network Interface Card](#)

Modifying Quick Check Interval, Deep Check Interval and Local Recovery

The default values for the quick check interval, deep check interval and local recovery may be modified after the resource has been placed under LifeKeeper protection by using the LifeKeeper GUI. To change the default values, right-click on the resource and select the entry from the menu list you wish to modify. After changing the value, click **Modify**. Results from the modify operation will be displayed in a dialog box. Click **Done** to complete the process.

Default values are as follows:

- Quick Check Interval is **180 seconds** (3 minutes)
- Deep Check Interval is **300 seconds** (5 minutes)
- Local Recovery is "**Enabled**"

Manual Switchover

Manual switchover can be performed from the LifeKeeper GUI using the **In Service** option. LifeKeeper will move the switchable IP addresses, volume and the IIS Web/FTP/SMTP site to the other server. You may wish to do this, for example, after a failover and you have fixed the primary and you want the primary to take over again.

IIS Failover

A failover occurs in two situations:

- The first situation is when the hardware or operating system has suffered a major failure and the server is no longer functioning. The LifeKeeper core on the backup server detects this when its heartbeat messages fail. At that time, the LifeKeeper core invokes the kit's recovery script. The recovery script ensures that the Internet server(s) is brought In Service on the backup server.
- The second situation is when the Recovery Kit's **Deepcheck** and **Quickcheck** scripts detect failures of the application. These scripts respond to the LifeKeeper core with code(s) indicating failure. The LifeKeeper core starts the failover process and invokes the kit's recovery scripts. The LifeKeeper core first stops the local server (if it is not stopped) and deactivates the switchable IP address and LifeKeeper volumes. It then continues the failover to the backup server.

When the primary server is repaired, the Internet server automatically returns to the primary server if switchback type is **Automatic**. You must perform a manual switchover if switchback type is **Intelligent**.

A relatively small number of web clients will experience a problem whenever a switchover or failover occurs. First, the process of moving the IP addresses and volumes to the backup system and starting the backup server there takes approximately 45 seconds (depending on the number of IP addresses and volumes), and users cannot connect to the server during that time. Second, the active server is stopped during the switchover, and users with open connections to the active server will be disconnected. In any case, if the IIS client retries the request, the request should succeed. Of course, once the switchable IP addresses and volumes have moved to the backup system and the backup site is running, service will be normal again.

Protecting FTP Sites with Non-Anonymous Login

The default procedure used to monitor protected FTP sites is to connect to the site and use an anonymous login. This feature is performed by the LifeKeeper deep check process assigned to monitor each FTP site. If your site does not permit an anonymous login, the default deep check operation will fail. Where anonymous logins are not permitted, you may either provide LifeKeeper with a small login script, or disable the deep check process for the LifeKeeper resource by setting the associated deep check interval to 0 seconds.

Using an FTP Login Script

The Microsoft FTP command provides the ability to use scripted FTP logins. LifeKeeper will create an empty login script file for each protected FTP site and the file names will match the FTP site names. The empty scripts are not used by LifeKeeper until written with FTP commands. They are created in the following folder:

```
<LifeKeeper Root Install Folder>\admin\kit\webapp
```

For instance, an empty login script for the "Default FTP Site" would be located at:

```
<LifeKeeper Root Install Folder>\admin\kit\webapp\Default  
FTPSite.txt
```

Spaces are permitted in the file name so it can match the FTP site name exactly with a .txt extension. The content of the script should include only 4 lines containing the FTP open command, the login ID, the login password, and the FTP bye command. For example, a login script might contain the following 4 lines:

```
open 192.168.1.10  
mytestloginID  
mytestloginPW  
bye
```

LifeKeeper will use a search mechanism for each response from the FTP client utility. A response starting with "230" indicates a successful login. A login failure will fail the deep check process. Other commands can be added to the script but they will be ignored by LifeKeeper.

Disabling the FTP Deep Check Process

If the protected FTP site does not permit anonymous logins and you prefer not to use a login script as described above, you may disable the deep check process for a particular LifeKeeper protected resource. Change directory to the LifeKeeper "bin" folder and use the following LifeKeeper command to disable deep check for the resource.

```
cd <LifeKeeper Root Install Folder>\bin
ins_setchkint -t <LifeKeeper Resource Tag Name> -c d -v 0
```

Changing LifeKeeper Microsoft IIS Recovery Kit Configuration

Any configuration change that affects the port number, IP address, hardware virtual servers or secure/non-secure setting, will affect the LifeKeeper configuration. As there is no direct linkage between the server and this kit, you should follow this procedure to synchronize the configuration.

1. Remove protection of the IIS resource by taking the hierarchy out of service and then deleting the hierarchy.
2. On the primary server, run the IIS Console as appropriate and apply the changes to the server.
3. On the backup server, run the IIS Console and apply the changes to the server.
4. Add LifeKeeper protection to the IIS resource by creating the IIS resource hierarchy and extending it to the backup server.

Removing Microsoft IIS

Remove LifeKeeper protection for the IIS Web/FTP/SMTP site before removing either the site itself or the entire software package. This is important so that LifeKeeper will not try to protect something that does not exist.

***Note:** This release of the kit was tested with Microsoft IIS 5.0 and Microsoft IIS 6.0. It will not necessarily be compatible with later releases, mainly because of dependencies on the location and content of certain registry keys and configuration files.

Changing the Network Interface Card (NIC)

The procedure for changing the Network Interface Card (NIC) will interrupt the web service. The basic reason for this is that the LifeKeeper IP hierarchy data includes the NIC model number. That piece of data cannot easily be changed, so the procedure is to remove and reinstall the Switchable IP addresses by doing the following:

1. Use the LifeKeeper GUI to remove protection from all IIS resources by taking them out of service, then deleting the hierarchies.
2. Stop all IIS services, including those on the backup server(s). All IIS services are now out of service.
3. Delete all IP hierarchies (be sure to save the configuration data in a notebook).
4. Shut down the server, replace the NIC, and then reboot the server.
5. Create all IP hierarchies again.
6. Bring the IP hierarchies In Service on each Web site's primary server.
7. On the primary server, start the IIS services again. The Web sites are now in-service again.
8. On each web server's primary server, use the LifeKeeper GUI to re-create the IIS resource hierarchy.

***Note1:** A quick way to identify the type of NIC currently installed is to run the ipconfig command from a DOS prompt. The type of NIC is displayed in the output (for example, "Ethernet adapter: DC21x41" describes an Ethernet NIC named DC21x41).

***Note2:** Changing the primary NIC could potentially affect LifeKeeper licensing. Refer to the Planning and Installation Guide for additional information on licensing.

IIS Troubleshooting

This section is intended to provide suggestions and insights into occurrences that are not specifically related to the LifeKeeper software, but have a relationship with the total environment.

Symptoms and Solutions

SYMPTOM

Identical sites are created on both the primary and the secondary servers, but the sites are not displayed in the **Site Name list** in the **Create Resource Wizard**.

SOLUTION

- Check that the same optional parameters are specified and their values are also the same (even the same letter case) on both systems.
- Also ensure that all sites have a valid description for LifeKeeper (alphanumeric characters and dashes only). If you need to change the name (description) of a protected Web Site, first delete the IIS resource, then change the description and recreate the resource.

SYMPTOM

When attempting to access the FTP site, the following error message is received:

"User x cannot login, home directory inaccessible".

SOLUTION

To allow users to log on to an FTP site as *ANONYMOUS*, the Internet Guest Account user IUSR_{machine name} must be given the right to logon

SYMPTOM

When configuring your IIS Web/FTP/SMTP in the Internet Information Services (IIS) Manager and you attempt to select the IP address from the drop-down list, your switchable IP address is not displayed.

SOLUTION

Manually enter the switchable IP address.

SIOS Protection Suite for Windows Support Matrix

Server Components

Supported Operating System	v8.0.0	v8.0.1	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0	v8.5.0	v8.6.0	v8.6.1	v8.6.2
Microsoft Windows Server 2008 R2 Standard, Enterprise, and DataCenter Editions	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2012 Standard and DataCenter Editions	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2012			64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit

Note: For all versions Microsoft .NET Framework 3.5 is required - download from <http://www.microsoft.com/net>

Note: The operating system versions listed above are supported for guests running on the following virtual platforms:

- Amazon EC2 (AWS)
- VMware vSphere 4.0 or later
- Microsoft Hyper-V Server 2008 R2 or later
- Citrix XenServer 5.5 or later (Microsoft Windows 2012 and later OS versions supported on Citrix XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later

Supported Operating System	v8.0.0	v8.0.1	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0	v8.5.0	v8.6.0	v8.6.1	v8.6.2
R2 Standard and DataCenter Editions											
Microsoft Windows Server 2016 Standard and DataCenter Editions									64-bit	64-bit	64-bit
Microsoft Windows Server 2019 Standard and DataCenter Editions											

Note: For all versions Microsoft .NET Framework 3.5 is required - download from <http://www.microsoft.com/net>

Note: The operating system versions listed above are supported for guests running on the following virtual platforms:

- Amazon EC2 (AWS)
- VMware vSphere 4.0 or later
- Microsoft Hyper-V Server 2008 R2 or later
- Citrix XenServer 5.5 or later (Microsoft Windows 2012 and later OS versions supported on Citrix XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later

User Interface Components

Supported Operating System	v8.0.0	v8.0.1	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0	v8.5.0	v8.6.0	v8.6.1	v8.6.2	v8.6.3
Microsoft Windows Server 2008	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2008 R2	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2012	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2012 R2			64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2016									64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2019												
Windows												

Note: For all versions Microsoft .NET Framework 3.5 is required - download from <http://www.microsoft.com/net>

Note: The operating system versions listed above are supported for guests running on the following virtual platforms:

- Amazon EC2 (AWS)
- VMware vSphere 4.0 or later
- Microsoft Hyper-V Server 2008 R2 or later
- Citrix XenServer 5.5 or later (Microsoft Windows 2012 and later OS versions supported on Citrix XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later

Supported Operating System	v8.0.0	v8.0.1	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0	v8.5.0	v8.6.0	v8.6.1	v8.6.2	v8.6.3
XP												
Windows Vista												
Windows 7	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit
Windows 8	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit	64-bit

Note: For all versions Microsoft .NET Framework 3.5 is required - download from <http://www.microsoft.com/net>

Note: The operating system versions listed above are supported for guests running on the following virtual platforms:

- Amazon EC2 (AWS)
- VMware vSphere 4.0 or later
- Microsoft Hyper-V Server 2008 R2 or later
- Citrix XenServer 5.5 or later (Microsoft Windows 2012 and later OS versions supported on Citrix XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later

SQL Server Recovery Kit

Supported Application	v8.0.0	v8.0.1	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0	v8.5.0	v8.6.0
Microsoft SQL Server 2000									
Microsoft SQL Server 2005									
Microsoft SQL Server 2008 R1	All editions	All editions	All editions	All editions	All editions	All editions	All editions	All editions	All editions
Microsoft SQL Server 2008 R2	All editions	All editions	All editions	All editions	All editions	All editions	All editions	All editions	All editions
Microsoft SQL Server 2012	All editions	All editions	All editions	All editions	All editions	All editions	All editions	All editions	All editions
Microsoft SQL Server 2014				All editions	All editions	All editions	All editions	All editions	All editions

Supported Application	v8.0.0	v8.0.1	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0	v8.5.0	v8.6.0
2014									
Microsoft SQL Server 2016									Al ed
Microsoft SQL Server 2017									

PostgreSQL Server Recovery Kit

Supported Application	v8.0.0	v8.0.1	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0	v8.5.0	v8.6.0	v8.6.1	v8.6.2
PostgreSQL									9.6.x	9.6.x	9.6.x

Oracle Recovery

Supported Application	v8.0.0*	v8.0.1*	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4.0
Oracle 11g							
Oracle 11g Release 2	Standard Edition,	Standard Edition,	Standard Edition,	Standard Edition,	Standard Edition,	Standard Edition,	Standard Edition,
	Standard Edition One,	Standard Edition One,	Standard Edition One,	Standard Edition One,	Standard Edition One,	Standard Edition One,	Standard Edition One,
	Enterprise Edition	Enterprise Edition	Enterprise Edition	Enterprise Edition	Enterprise Edition	Enterprise Edition	Enterprise Edition
Oracle 12c						Enterprise Edition	Standard Edition Two, Enterprise Edition

* For each of these versions the Oracle Recovery Kit version is 7.2.1

Supported Application	v8.0.0*	v8.0.1*	v8.1.0	v8.2.0	v8.2.1	v8.3.0	v8.4
Oracle 12c Release 2							
Oracle 18c							

* For each of these versions the Oracle Recovery Kit version is 7.2.1

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	End of Support
SIOS Protection Suite for Linux v9.1.1	January 31, 2020
SIOS Protection Suite for Linux v9.1.2	June 30, 2020
SIOS Protection Suite for Linux v9.2	October 31, 2020
SIOS Protection Suite for Linux v9.2.1	December 31, 2020
SIOS Protection Suite for Linux v9.2.2	March 31, 2021
SIOS Protection Suite for Linux v9.3.2	April 30, 2022
SIOS Protection Suite for Linux v9.4	October 31, 2022
SIOS Protection Suite for Windows v8.4	July 31, 2019
SIOS Protection Suite for Windows v8.5	December 31, 2019
SIOS Protection Suite for Windows v8.6	August 31, 2020
SIOS Protection Suite for Windows v8.6.1	March 31, 2021
SIOS Protection Suite for Windows v8.6.2	August 31, 2021
SIOS Protection Suite for Windows v8.6.3	December 31, 2021
SIOS Protection Suite for Windows v8.6.4	April 30, 2022
DataKeeper for Windows v8.4	July 31, 2019
DataKeeper for Windows v8.5	December 31, 2019
DataKeeper for Windows v8.5.1	March 31, 2020
DataKeeper for Windows v8.6	August 31, 2020
DataKeeper for Windows v8.6.1	March 31, 2021
DataKeeper for Windows v8.6.2	August 31, 2021
DataKeeper for Windows v8.6.3	December 31, 2021
DataKeeper for Windows v8.6.4	April 30, 2022
DataKeeper Cluster Edition for Windows v8.4	July 31, 2019
DataKeeper Cluster Edition for Windows v8.5	December 31, 2019

DataKeeper Cluster Edition for Windows v8.5.1	March 31, 2020
DataKeeper Cluster Edition for Windows v8.6.	August 31, 2020
DataKeeper Cluster Edition for Windows v8.6.1	March 31, 2021
DataKeeper Cluster Edition for Windows v8.6.2	August 31, 2021
DataKeeper Cluster Edition for Windows v8.6.3	December 31, 2021
DataKeeper Cluster Edition for Windows v8.6.4	April 30, 2022
DataKeeper Cluster Edition for Windows v8.6.5	July 31, 2022

Split-Brain Issue and Recovery

When protecting DataKeeper volume resources in Microsoft WSFC, if all nodes are included in the cluster split-brain recovery should occur automatically.

However, in a Node Outside the Cluster scenario a split-brain may occur if network connectivity is lost to the DR Node. The user might choose to manually switchover the volume to the DR node while WSFC maintains the source on the original cluster node. When network connectivity to the DR node is restored there will be a conflict known as a split-brain condition where both systems assume the ownership role over the volume. The SIOS DataKeeper user interface will display the error "**Mirror Inconsistencies - Click the following mirror(s) to resolve source conflicts - Mirroring is halted until this is resolved**" (as shown in the diagram below).



In addition, the following error is logged to the **System Event log**:

An invalid attempt to establish a mirror occurred. Both systems were

found to be Source.t

Local Volume: F

Remote system: 192.168.1.212

Remote Volume: F

The mirror has been paused or left in its current non-mirroring state.

Refer to Extending a Clustered DataKeeper Volume to a Node Outside the Cluster for the recovery procedure.

Resolving a Split-Brain Issue via the Command Line Interface

EMCMD <system> PREPARETOBECOMETARGET <volume letter>

This command should only be used to recover from a Split-Brain condition. It should be run on the system where the mirror is to become a target and is only valid on a mirror source. This command causes the mirror to be deleted and the volume to be locked.

To complete split-brain recovery, run [CONTINUEMIRROR](#) on the system that remains as the mirror source.

Example Scenario

If volume F: is a mirror source on both SYSA and SYSB, you can use emcmd to resolve this split-brain situation. Choose one of the systems to remain a source - for example, SYSA. Make sure there are no files or modifications on SYSB that you want to save - if so, these need to be copied manually to SYSA. To re-establish the mirror, perform the following steps:

```
EMCMD SYSB PREPARETOBECOMETARGET F
```

The mirror of F: on SYSB will be deleted and the F: drive will be locked.

```
EMCMD SYSA CONTINUEMIRROR F
```

Mirroring of the F: drive from SYSA to SYSB will be established, a partial resync will occur (overwriting any changes that had been made on SYSB), and the mirror will reach the **Mirroring** state.

DataKeeper Troubleshooting

The topics in this section contain important information about known issues and restrictions offering possible workarounds and/or solutions.

Known Issues and Workarounds

[Access to Designated Volume Denied](#)

[Failed to Create Mirror](#)

[MaxResyncPasses Value](#)

[Mirroring with Dynamic Disks](#)

[Server Login Accounts and Passwords Must Be Same on Each Server in the Cluster](#)

[System Event Log - Create Mirror Failed in the GUI](#)

[Unable to Determine Previous Install Path](#)

[User Interface - Failed to Create Mirror](#)

[User Interface - Shows Only One Side of the Mirror](#)

Windows Server 2012 Specific Issues

[Windows Server 2012 MMC Snap-in Crash](#)

[Windows Server 2012 iSCSI Target Role Does Not Support Dynamic Disks](#)

[Windows Server 2012 NIC Teaming Issue](#)

Restrictions

[Bitlocker Does Not Support DataKeeper](#)

[CHANGEMIRRORENDPOINTS Restriction](#)

[CHKDSK](#)

[DataKeeper Volume Resize Restriction](#)

[Directory for Bitmap Must Be Created Prior to Relocation](#)

[Duplicate IP Addresses Disallowed Within a Job](#)

[Intensive I-O with Synchronous Replication](#)

SIOS DataKeeper for Windows Quick Start Guide

This topic provides step-by-step instructions for installing and configuring DataKeeper Standard Edition. The series of steps includes links into the documentation that describe each step in detail.

Prerequisites and Installation

1. Read the [SIOS Protection Suite for Windows Release Notes](#) for late breaking information.
2. [Firewall Configurations](#) - Make sure you understand what ports must be opened on any firewalls.
3. [Network Bandwidth](#) - If replicating across a WAN, it is critical that a [rate of change](#) analysis be done to ensure there is adequate bandwidth.
4. DataKeeper is a block-level volume replication solution and requires that each server have additional volume(s) (other than the system drive) that are the same size. Please review [Volume Considerations](#) for additional information regarding storage requirements.
5. Review the section [Understanding Replication](#) to help understand how DataKeeper works and the difference between synchronous and asynchronous replication.
6. During installation/upgrade, you will be prompted to select and set up a user account that will be used to start the DataKeeper Service. SIOS recommends using a domain account on each server (same account with the same password). This account must be added to each server's local Administrator Group. Perform the following on **each server**:
 1. Select **Edit Local Users and Groups**
 2. Select **Groups, Administrators**
 3. Add the **DK Service Account** being used to this Administrator's Group.

The following topics provide further information about these accounts.

[DataKeeper Service Log On ID and Password Selection](#)

[Server Login Accounts and Passwords Must Be Same on Each Server in the Cluster](#)

7. Configure your [Network Adapter Settings](#).
8. Review [Sector Size](#) and [GUI Requirements, Platforms and Browsers](#).
9. Set your [Browser Security Parameters](#) to low.
10. Set the Event Log to [Overwrite Events as Needed](#).
11. Disable the **Distributed Link Tracking Client service** and set its startup policy to **Manual**.
12. If using High-Speed Storage, review [High-Speed Storage Best Practices](#).

Configuration

1. Choose the type of replication scheme that meets your data protection requirements. The links below describe the possible configuration options.
[Disk-to-Disk](#)
[One-to-One](#)
[One-to-Many](#)
[Many-to-One](#)
2. [Create a Job](#) using the DataKeeper UI. The Job creation process also adds your first mirror.
3. If additional mirrors are required, you can [Add a Mirror to a Job](#).
4. If you are creating a [One-to-Many](#) replication scheme, see [Creating Mirrors with Multiple Targets](#).

*Note: A "Mirror" consists of a source volume and target volume as well as a source IP and target IP. A "[Job](#)" consists of one or more related "Mirrors". For example, when replicating a database that consists of an E:\ volume for data and an F:\ volume for log files, you would place the E:\ and F:\ volumes into the same "Job". However, if you have three VMs, one on the G:\, H:\ and I:\ volumes, you would create three separate Jobs as the VMs themselves are not related to each other. A general rule of thumb is that mirrors should always be grouped into the same Job if they should move as a group.

Management

Some of the more common management functions are described below.

[Switching Over a Mirror](#)

[Managing Mirrors](#)

[Pause and Unlock](#)

[Continue and Lock](#)

[Break](#)

[Resync](#)

[Deleting a Mirror](#)

[Mirror Properties](#)

[Changing the Compression Level of an Existing Mirror](#)

[DataKeeper Volume Resize](#)

You may also want to review the [Frequently Asked Questions](#).

For more information, please refer to the DataKeeper section of the [DataKeeper for Windows Technical Documentation](#).

Replicating Hyper-V Virtual Machines

If you are using DataKeeper to replicate Hyper-V VHD files in a non-clustered environment, you need to understand how to access and re-provision virtual machines based on those replicated VHD files. Please follow the instructions in the topic [Providing Disaster Recovery for Hyper-V Virtual Machines](#).

Troubleshooting


Use the following resources to help troubleshoot issues:

- [Troubleshooting](#) issues section
- For customers with a support contract - <http://us.sios.com/support/overview/>
- For evaluation customers only - [Pre-sales support](#)

Creating a DataKeeper Volume Resource in WSFC

Automatic Creation of a Mirror in WSFC

1. Use the SIOS DataKeeper GUI to configure a mirror.
2. At the completion of the mirror creation operation, the following dialog box will pop up which will allow you to automatically add the mirror as a DataKeeper volume resource into WSFC.



If **Yes** is selected, then the DataKeeper mirror will be added as a **DataKeeper volume resource** to "**Available Storage**" in WSFC. However, if **No** is selected or an error is encountered, follow the steps in the section "[Manual Creation of a Mirror in WSFC](#)", to create the mirror manually.

Whether the mirror is created automatically or manually, the DataKeeper volume resource(s) is placed in "**Available Storage**." The DataKeeper volume resource(s) can now be utilized just as if it were a physical disk resource. While clustering some applications (SQL for example), the DataKeeper volume resource is recognized and treated exactly the same as a physical disk resource. This means that during the cluster create process, the DataKeeper volume resource will be pulled automatically into the cluster resource and the dependencies will be created automatically. Other applications, notably the Hyper-V "Virtual Machine" resource, are hard-coded to look for a physical disk resource and will not automatically pull the DataKeeper resource into the cluster configuration. In these circumstances, the DataKeeper volume resource can be easily added to the cluster resource and the dependencies can be created manually through the WSFC GUI.

Use the DataKeeper GUI to monitor the mirroring states. It provides more

information than the Failover Cluster manager UI.

Manual Creation of a Mirror in WSFC

1. Wait until the mirrors are in the **Mirroring** state.
2. Create an **Empty Application**.
 - Open the **Microsoft Windows Failover Cluster GUI**
 - Right-click on **Services and Applications** (*in Windows 2012, right-click on Roles*)
 - **Expand More Actions** (*not in Windows 2012*)
 - Select **Create Empty Service and Application** (*in Windows 2012, select Create Empty Role*)
 - Right-click on the newly created service and select **Properties**
 - Rename the empty application to **DataKeeper Service**
3. Add a resource to the new empty application.
 - Right-click on **DataKeeper Service**
 - Select **Add Resource**
 - Navigate to **More Resources/Add DataKeeper Volume**
4. Open the **Properties** of this newly added resource and change the **Resource Name** to something meaningful, like "**DataKeeper Volume E**" (which designates E drive). Select the **DataKeeper Volume Parameters** tab and associate a volume drive letter (E for the above example) with the resource. A list of all mirrored volumes will be presented. Choose the volume that you wish to control with this resource.
5. Right-click on the resource and choose **Bring Online**. Make sure it comes online. **Note:** Resources placed in online status are unlocked. Resources placed in the offline status are locked. (*In Windows 2012, select Start Role.*)
6. Repeat **Steps 3 and 4** for all the other DataKeeper volumes you wish to protect with WSFC.
7. Move the resource to available storage by doing the following:
 - Expand the **Services and Applications** tab
 - Select the DataKeeper service just created and note the resources
 - From the main pane, select the resource (i.e., DataKeeper Volume E)
 - Right-click and select **Remove from DataKeeper Service** – this will move the resource to **Available Storage**

- Repeat for all the other volumes

8. Delete the empty application created in Step 2.

DataKeeper Volume Resource Health Check

A DataKeeper Volume resource provides two functions that are used by the Microsoft Cluster service to check for availability and health of the DataKeeper Volume resource. A simple check LooksAlive and a more rigorous check IsAlive.

LooksAlive

The Cluster service calls the LooksAlive function based on the specified interval. The default is every 20 seconds on a freshly installed system, or 60 seconds after upgrading DataKeeper Cluster Edition from a version prior to 8.4.0. The LooksAlive function performs a quick check of the volume device. When the LooksAlive test fails, the cluster service will call the IsAlive test immediately.

IsAlive

Performs a thorough check to determine if the specified resource is online (available for use). The default is 120 seconds on a freshly installed system, or 300 seconds after upgrading DataKeeper Cluster Edition from a version prior to 8.4.0. If the device for the mirror becomes unreachable by DataKeeper, the IsAlive check will detect this condition and will mark the resource as Failed.

Extending a Single SQL Server Node to a Cluster

This guide explains how to install a single SQL Server node and extend it to a clustered node. Please read this document carefully prior to installation.

Planning Steps:

- Start with a Microsoft SQL Server standalone node, where data resides on the C drive.

Note: It is recommend that similar hardware be used for the standalone node (for the purpose of this guide, we will call this node the backup node).

- On the new backup node setup Windows failover cluster and setup a one node cluster (setup alongside a file share quorum).
- Setup Datakeeper Cluster Edition on the new node.
- Create empty DataKeeper Volume resource(s) using the Failover Clustering UI. Provide a name(s) that best describes its intended use (Example: "DataKeeper Volume F (NonMirrored)").

The following steps explain how to manually create this resource:

- a. In the Failover Cluster Manager, create an empty role, right-click on **Role** and select **Create Empty Role**.
- b. Right-click the empty role and select **Add a Resource, More Resources**, then select **Add DataKeeper Volume**.
- c. Right-click the **new DataKeeper Volume** resource and select **Properties**.
- d. Enter the Resource Name you chose earlier (Example: "DataKeeper Volume F (NonMirrored)") then select **Done**. No other properties changes are needed at this time. Follow the steps below for setting the Properties needed for the non-mirrored resource.
- e. Repeat the steps a - d for the E drive.

Assign the following properties using Powershell:

```
VolumeLetter = "F" (if the drive letter is F, otherwise whatever  
the drive letter is)
```

```
NonMirrored = 1 (there is no space between Non and Mirrored)
```

Assign the following properties using Powershell:

```
Get-ClusterResource "DataKeeper Volume F (Non-Mirrored)" | Set-  
ClusterParameter -Name VolumeLetter -Value "F"
```

```
Get-ClusterResource "DataKeeper Volume F (Non-Mirrored)" | Set-  
ClusterParameter -Name NonMirrored -Value 1
```

After the storage is created, right-click and select **Remove from empty role**. This will move the storage to **Available Storage** which can now be used by SQL Server installation in the following steps.

- Run SQL Server setup, choose the SQL Server Cluster edition to install. Select the same features that are installed on the primary node. (**Note:** Use the domain user ID and password to start SQL Server services instead of local account.)
- During installation setup will prompt you for failover cluster storage, select the storage you created earlier.
- Copy the database from the primary server to the new backup node the use the **BACKUP DATABASE** command to make a backup of the database.
- Connect to the Cluster database on the backup node and use the **RESTORE DATABASE** command to restore the database.

It is highly recommend using the **MOVE** option of the **RESTORE** option to move the files to separate drives (separate DATA and LOG file). The following example shows how to move a sample Sales database to different drives. (**Note:** The data and the log files are moved to different volumes.)

```
RESTORE DATABASE sales
```

```
FROM DISK = 'C:\Backup\Sales.bak'
```

```
WITH RECOVERY,
```

```
MOVE 'Sales_Dat' TO  
'E:\MSSQL11.MSSQLSERVER\MSSQL\Data\Sales.mdf',  
MOVE 'Sales_Log' TO  
'F:\MSSQL11.MSSQLSERVER\MSSQL\Data\Sales.ldf';
```

- Users can now be moved to the new database and the new backup server. Also, the original primary server can now be reconfigured. It is recommend that the entire Windows operating system be re-installed. Once you have re-installed the operating system, setup failover clustering and join the cluster created earlier. (**Note:** When joining the cluster, select **"No, I do not require support from Microsoft"**. The cluster can be verified later in the process.)
- Setup Datakeeper Cluster Edition on this server.
- Remove both non-mirrored DataKeeper volume resources from the MS SQL Server cluster group.
- Delete both non-mirrored DataKeeper Volume resources from Storage. (**Note:** It is important to remove the resource from both the cluster group and available storage. The resource should be deleted from failover clustering before proceeding.)
- Use the DataKeeper GUI to create jobs that contain the E and F volume mirrors for all nodes in the cluster. Choose the node that is currently running Microsoft SQL Server as the source. After creation choose **Yes** to add the volumes to Failover clustering
- Add the DataKeeper Volume E and DataKeeper Volume F storage to the SQL cluster group.
Right-click on the resource and select **Add Storage**.
- Setup SQL Server using the cluster edition **Add Node to a SQL Server Cluster** option.
- Finish by adding dependencies to the Microsoft SQL Server resource for both the DataKeeper volume resources.

The single node database has now been converted to a clustered highly available database and is ready for failover testing.

Extending a Traditional 2-Node WSFC Cluster to a Third Node via DataKeeper

When replicating from a WSFC 2-node cluster to a third node via DataKeeper, the following tasks are required:

- Replace the existing WSFC physical disk resource with a DataKeeper volume resource that supports data replication.
- Change the quorum type to Node Majority.
- Add the third node into the cluster for failover.
- Reestablish any/all resource dependencies to the new DataKeeper volume resource.

The following example details the steps necessary to extend a cluster resource group from a 2-node cluster to a third node via DataKeeper.



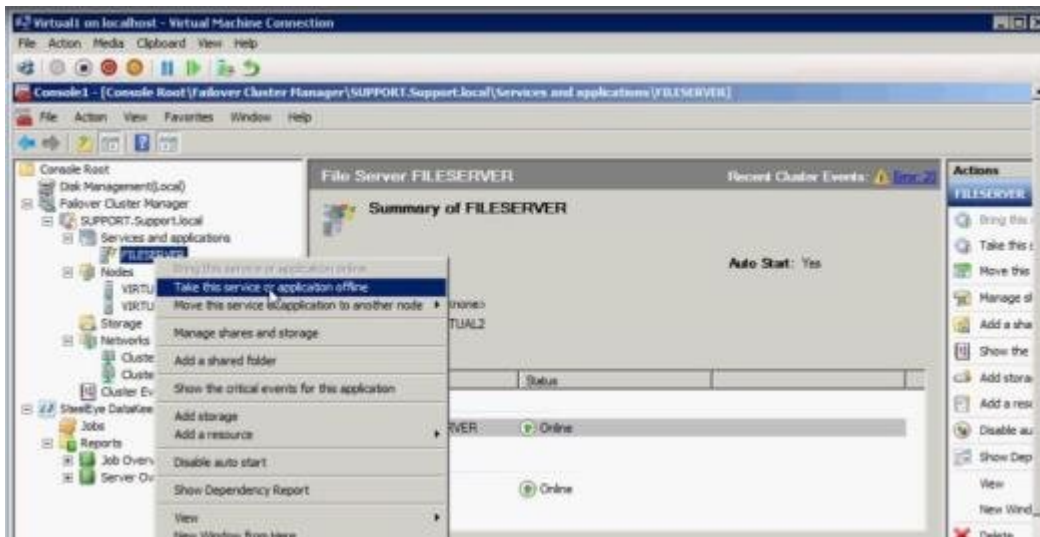
1. Remove physical disk resources from WSFC.

This also removes any dependencies on these physical disk resources. These dependencies will need to be reestablished to the new

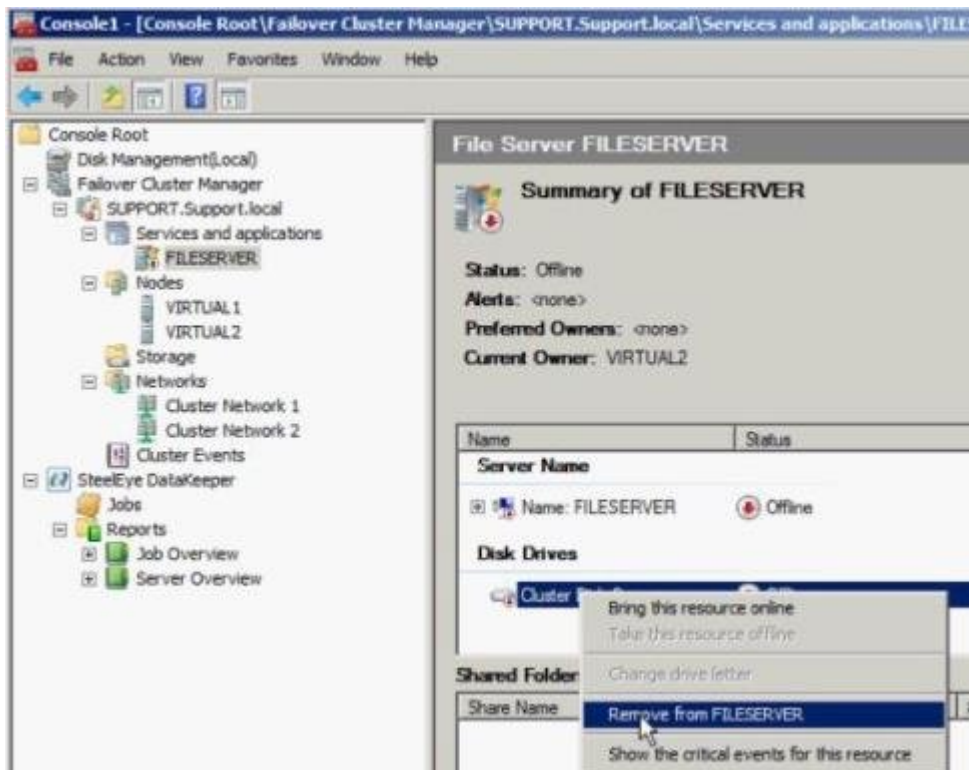
DataKeeper volume resource, so please take note of what these dependencies are before completing this first step by viewing the **Dependency Report**. Highlight your resource and select **Show Dependency Report**.

Use **Failover Cluster Manager MMC** to perform the following steps:


- a. Offline the cluster resource group by right-clicking and selecting **Take this service or application offline**.



- b. Remove the physical disk from cluster resource group (moves to *Available Storage*).



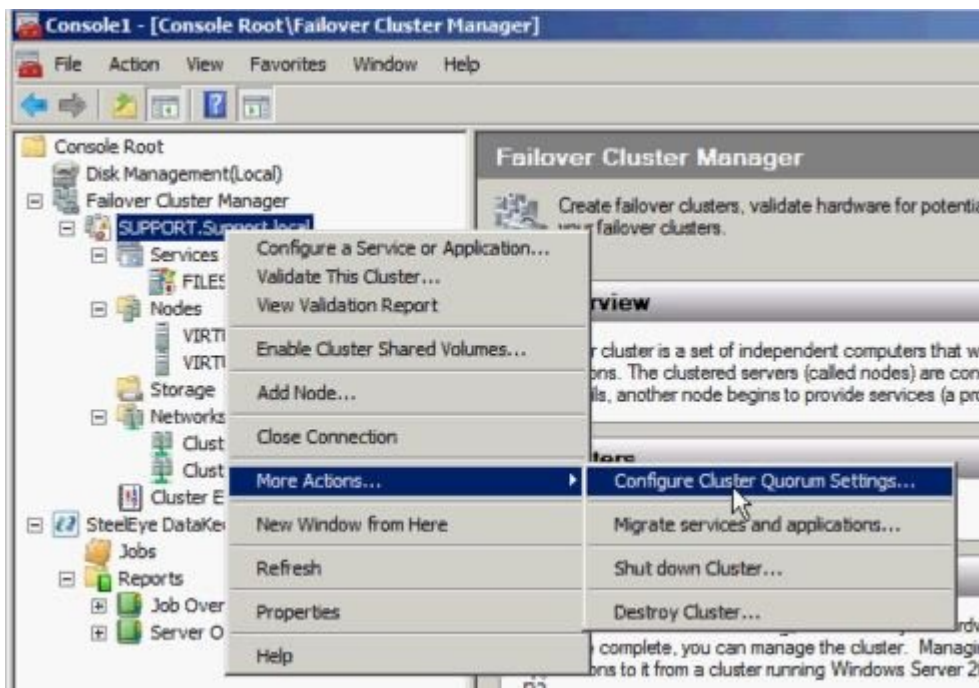
c. Remove the physical disk resource from the cluster configuration by deleting the resource from the **Available Storage** group.



2. Configure cluster quorum settings.

Because there will now be a third node in a remote site, the **Disk Witness in Quorum** is no longer valid, therefore, the **Node Majority** configuration should be selected.

- a. Right-click the cluster and select **More Actions / Configure Cluster Quorum Settings...**




- b. Select **Node Majority**. **Note:** Change the quorum type to **Node Majority** if the final number of nodes is odd as in a 3-node cluster, or **Node and File Share Witness Majority** if the final number of nodes will be even as in a 4-node cluster.

- c. Delete disk witness from **Available Storage**.

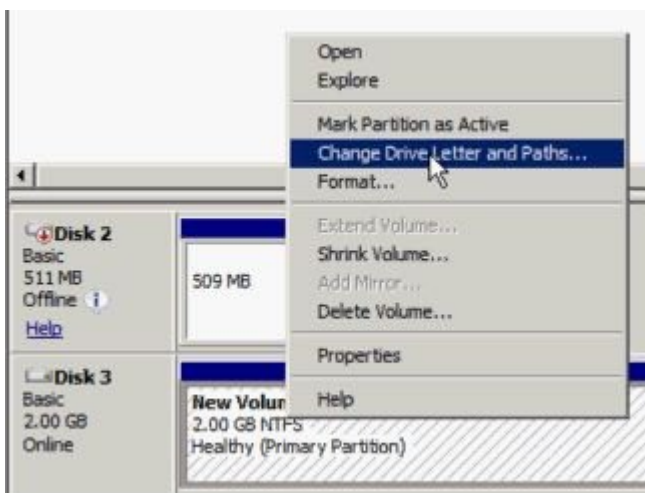
3. Online Disk.

Use **Disk Management MMC** to perform the following steps:

- a. Online physical disk.



b. Change drive letter to match previous configuration (if needed) .

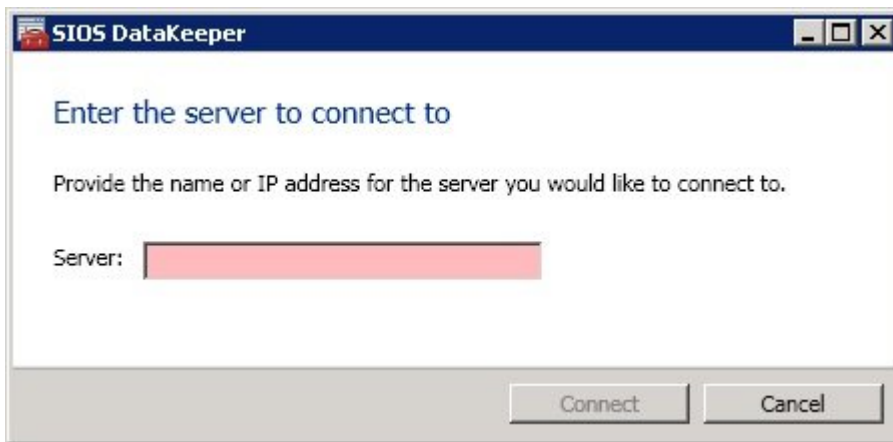


c. Repeat the above steps, 2a and 2b, on all shared nodes in the cluster.


4. Make sure the volume of the third node is available to become a target. It should be formatted and online, and it should be at least as big or bigger than the source volume.
5. [Create Mirror](#).

Use the DataKeeper UI to perform the following steps (make sure DataKeeper Service is running on all servers):

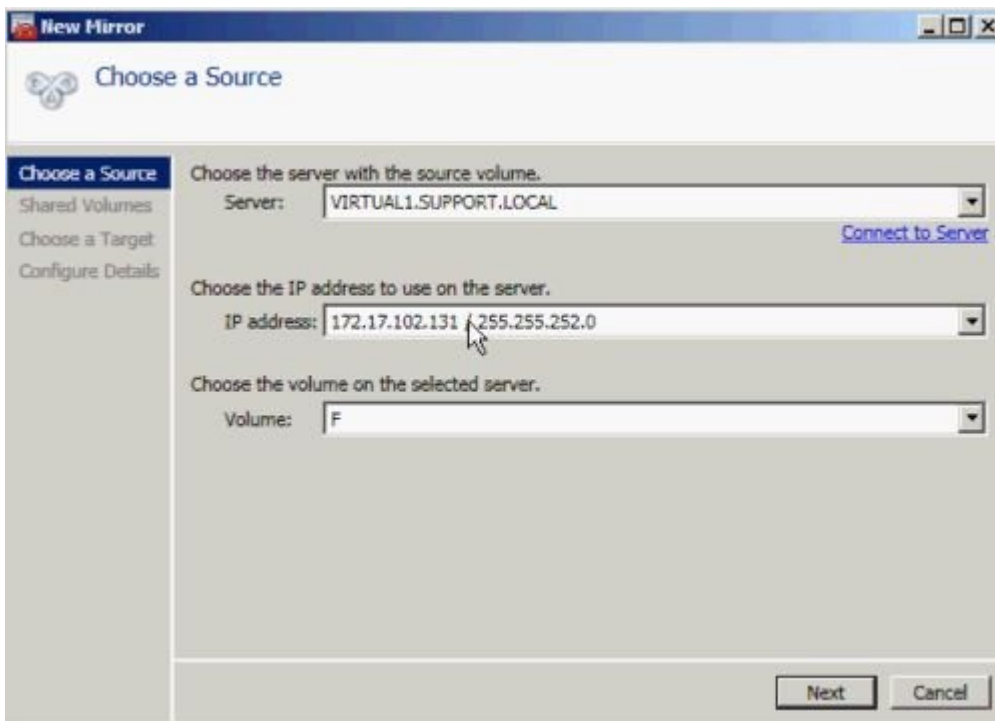
a. Connect to all shared nodes and the third node.



The **Server Overview Report** will show connection to all three nodes:



- b. Create a job containing a mirror to the third node.
- c. Select a **Source**.



The "New Mirror" dialog box is shown with the "Choose a Source" tab selected. The "Server" dropdown is set to "VIRTUAL1.SUPPORT.LOCAL". The "IP address" dropdown is set to "172.17.102.131 / 255.255.252.0". The "Volume" dropdown is set to "F". A "Connect to Server" link is visible. The "Next" and "Cancel" buttons are at the bottom right.

New Mirror

Choose a Source

Choose the server with the source volume.

Server:

[Connect to Server](#)

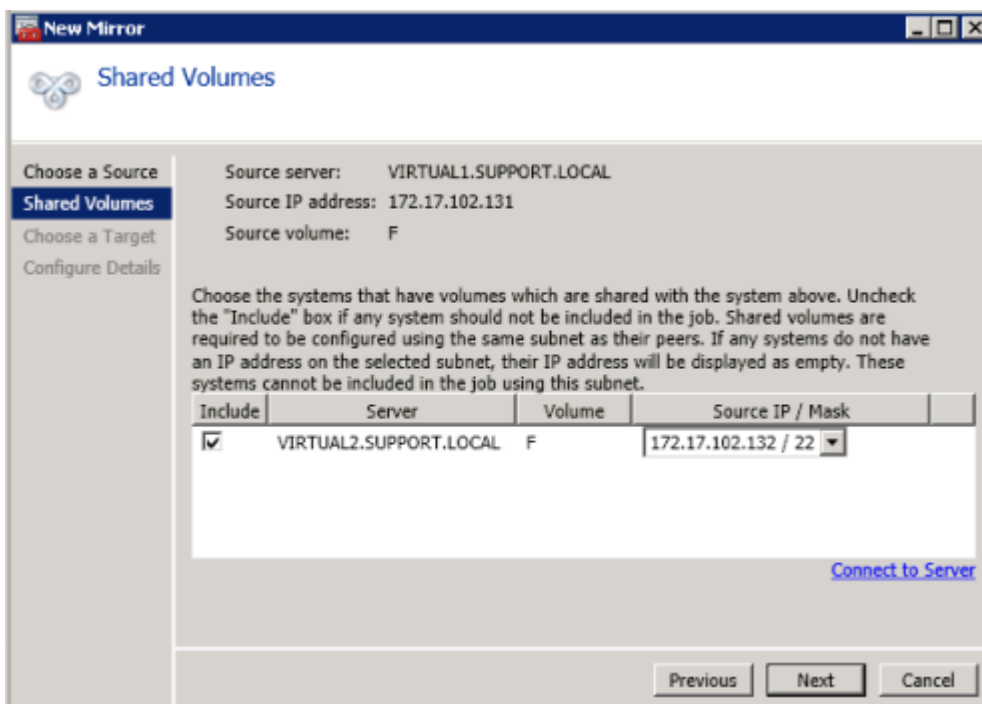
Choose the IP address to use on the server.

IP address:

Choose the volume on the selected server.

Volume:

d. Select **Shared Volumes**.



The "New Mirror" dialog box is shown with the "Shared Volumes" tab selected. The "Source server" is "VIRTUAL1.SUPPORT.LOCAL", the "Source IP address" is "172.17.102.131", and the "Source volume" is "F". A text box explains that systems should be included if they share volumes and are on the same subnet. A table lists the systems to include. The "Include" checkbox is checked for "VIRTUAL2.SUPPORT.LOCAL" with volume "F" and source IP "172.17.102.132 / 22". A "Connect to Server" link is visible. The "Previous", "Next", and "Cancel" buttons are at the bottom right.

New Mirror

Shared Volumes

Source server: VIRTUAL1.SUPPORT.LOCAL

Source IP address: 172.17.102.131

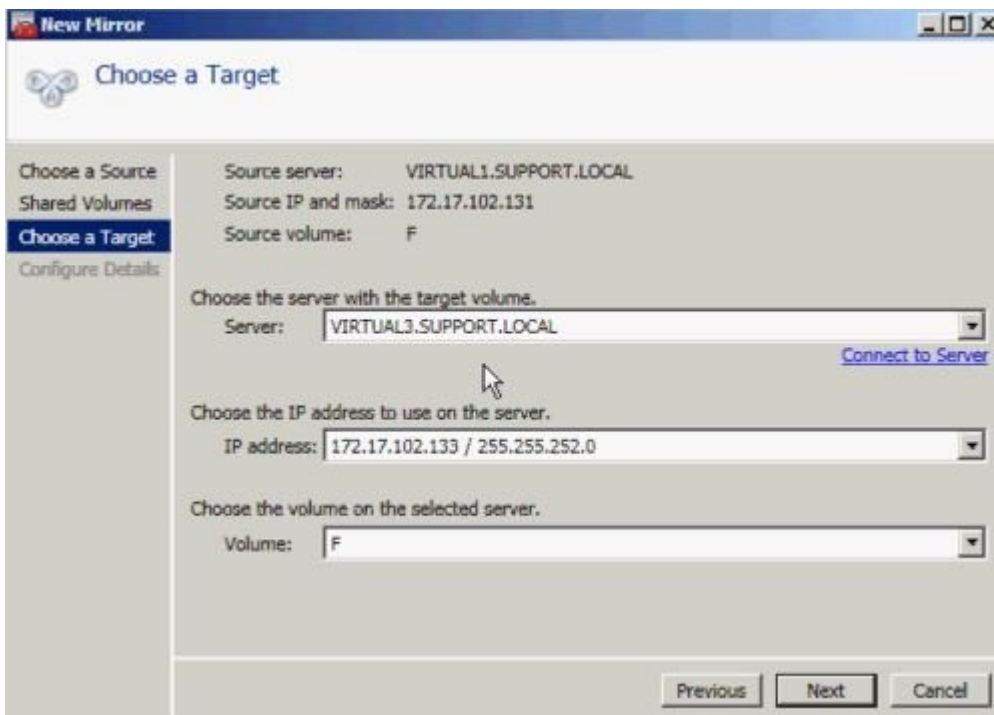
Source volume: F

Choose the systems that have volumes which are shared with the system above. Uncheck the "Include" box if any system should not be included in the job. Shared volumes are required to be configured using the same subnet as their peers. If any systems do not have an IP address on the selected subnet, their IP address will be displayed as empty. These systems cannot be included in the job using this subnet.

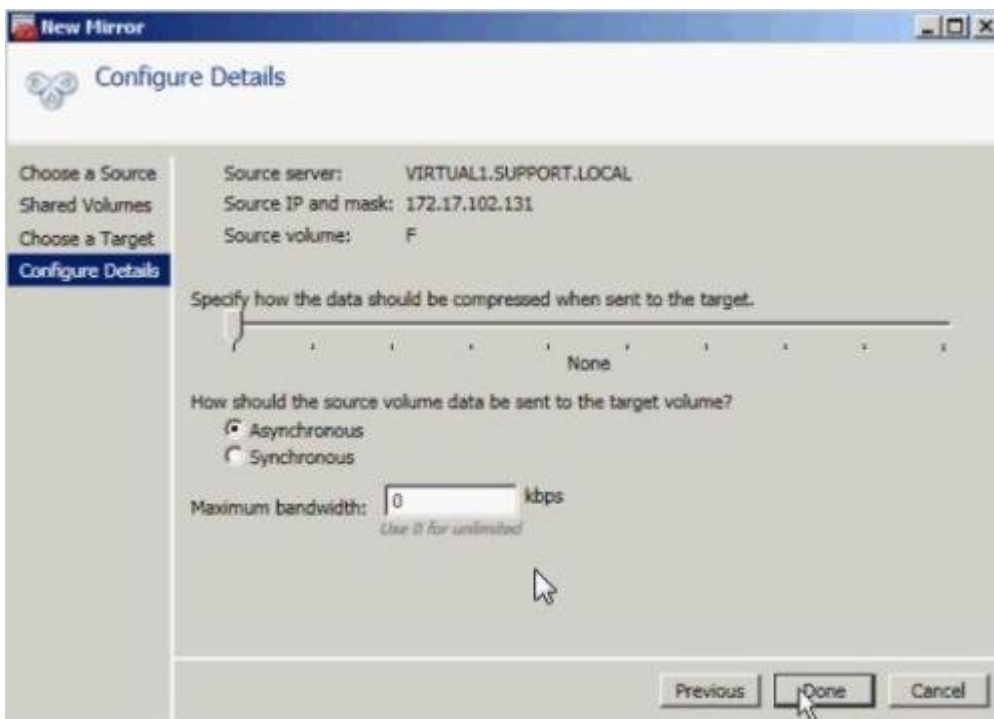
Include	Server	Volume	Source IP / Mask
<input checked="" type="checkbox"/>	VIRTUAL2.SUPPORT.LOCAL	F	172.17.102.132 / 22

[Connect to Server](#)

c. Select a **Target**.

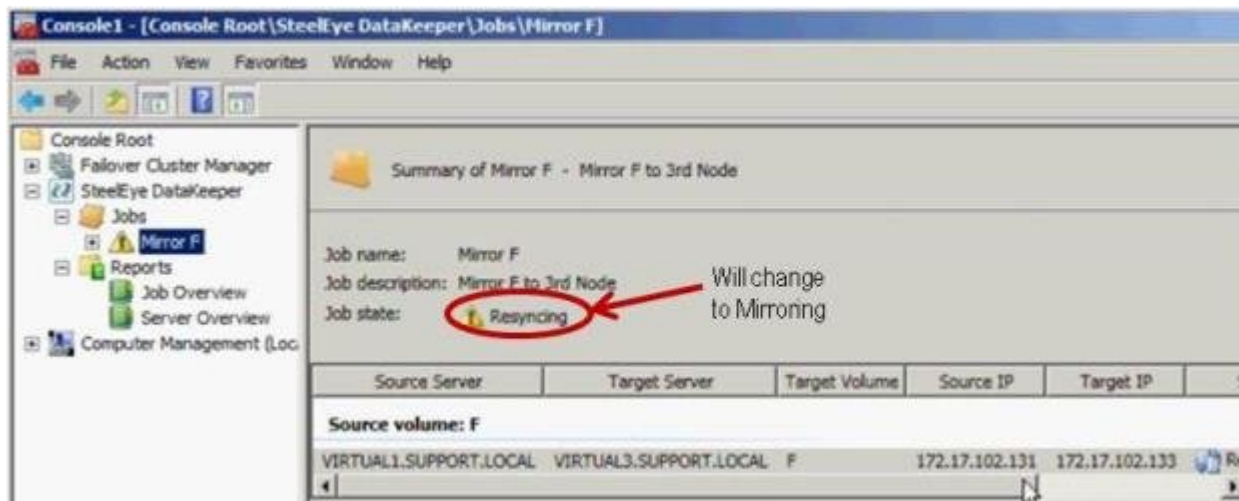


f. **Configure Details.** Choose compression settings, if applicable, and choose mirror type, [asynchronous or synchronous](#). Select **Done**.



The mirror will then begin resyncing to the third node. After resyncing

is complete, the **Job State** will change to **Mirroring**.



6. Add **DataKeeper Volume Resource**.

To add the DataKeeper Volume Resource in WSFC, perform one of the following:


[Automatic Creation of a Mirror in WSFC](#)

[Manual Creation of a Mirror in WSFC](#)

7. Add third node into the cluster.

Use Failover Cluster Manager MMC to perform the following steps (**Note:** Make sure the Failover Clustering feature is installed on the third node prior to adding into the cluster):

- a. Right-click **Nodes**.
- b. Select **Add Node**.
- c. In the **Add Node Wizard**, enter name of server to be added, click **Add** and select **Next**.



d. **Yes** may be chosen to perform **Validation Testing**, but be aware that some errors are expected due to the disk being locked on the target side. **Note:** When performing validation testing, select **Run only the tests I select**, then deselect the **Storage** tests from the **Test Selection** screen.



Even though the Storage tests were deselected, the **Failover Cluster Validation Report** will still report a **List Cluster Volumes failure** in the **Cluster Configuration Test**. This, along with a few other warnings, is expected due to the third node and its volume not being shared with the first two nodes.



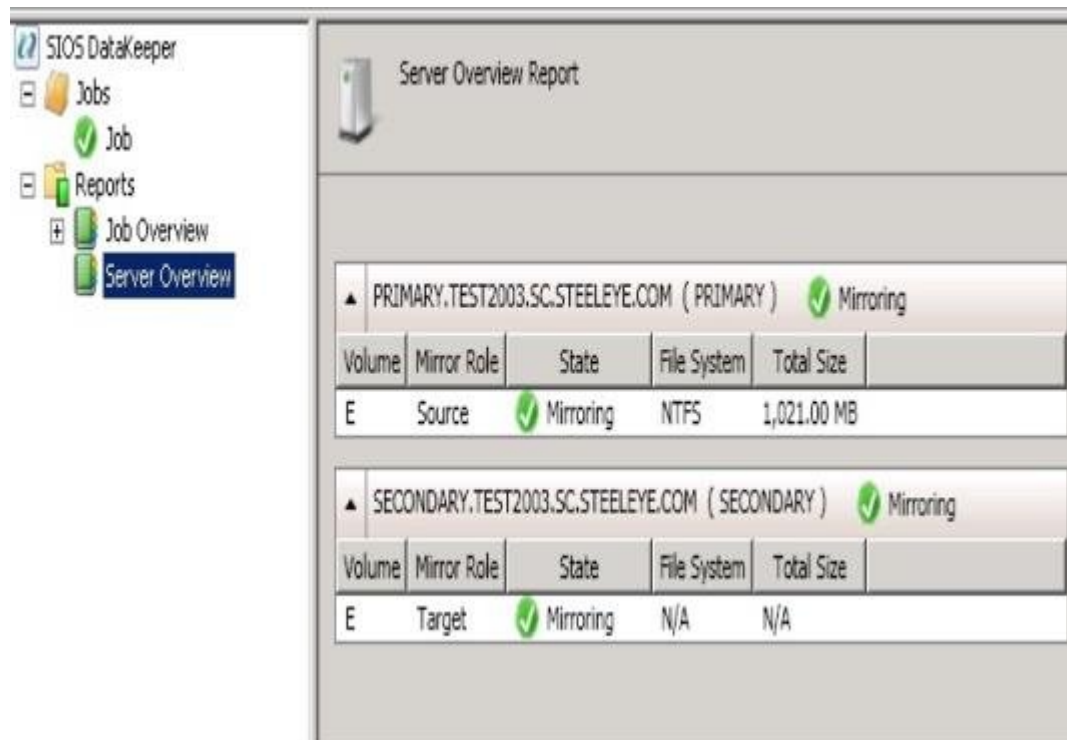
Name	Result	Description
List Cluster Core Groups		Success
List Cluster Network Information		Success
List Cluster Resources		Success
List Cluster Volumes		Failed
List Clustered Services and Applications		Warning
Validate Quorum Configuration		Warning
Validate Resource Status		Warning
Validate Service Principal Name		Success
Validate Volume Consistency		Success

e. After viewing the **Validation Report**, select **Finish**. The **Validation Warning** screen will reappear. Select **No** to add the node without performing the tests again.

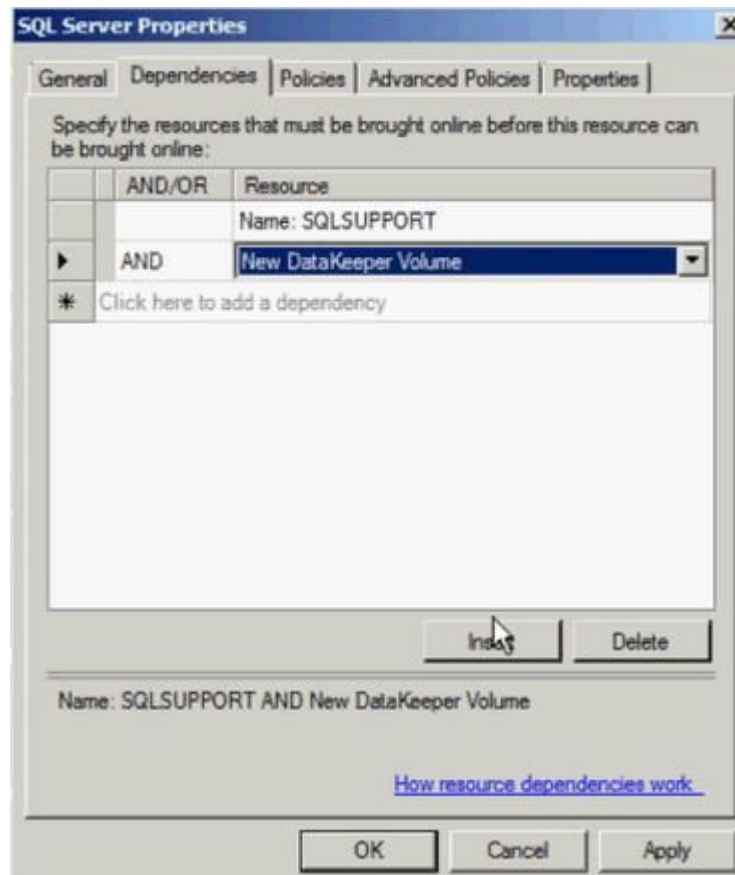
f. Select **Next** on the **Confirmation** screen.

g. The **Summary** screen will appear indicating that the node was added successfully.

At this point, the cluster resource group is defined across all three nodes in the cluster.




8. Reestablish any/all resource dependencies to the new DataKeeper volume resource.
 - a. Right-click application resource and select **Properties**.



- b. Select the **Dependencies** tab.
- c. Click on the **Click here** to add a dependency tab.
- d. Enter the new DataKeeper volume
- e. Select **Apply** and then **Okay**.

Dependencies should now be reestablished.



Extending a Traditional 2-Node WSFC SQL Server Cluster to a Third Node via DataKeeper


When replicating from a WSFC SQL Server 2-node cluster to a third node via DataKeeper, the following tasks are required:

- Replace the existing WSFC physical disk resource with a DataKeeper volume resource that supports data replication.
- Change the quorum type to **Node Majority**.
- Add the third node into the SQL Server cluster for failover.
- Reestablish any/all resource dependencies to the new DataKeeper volume resource.

The following example details the steps necessary to extend a SQL Server cluster resource group from a 2-node cluster to a third node via DataKeeper.

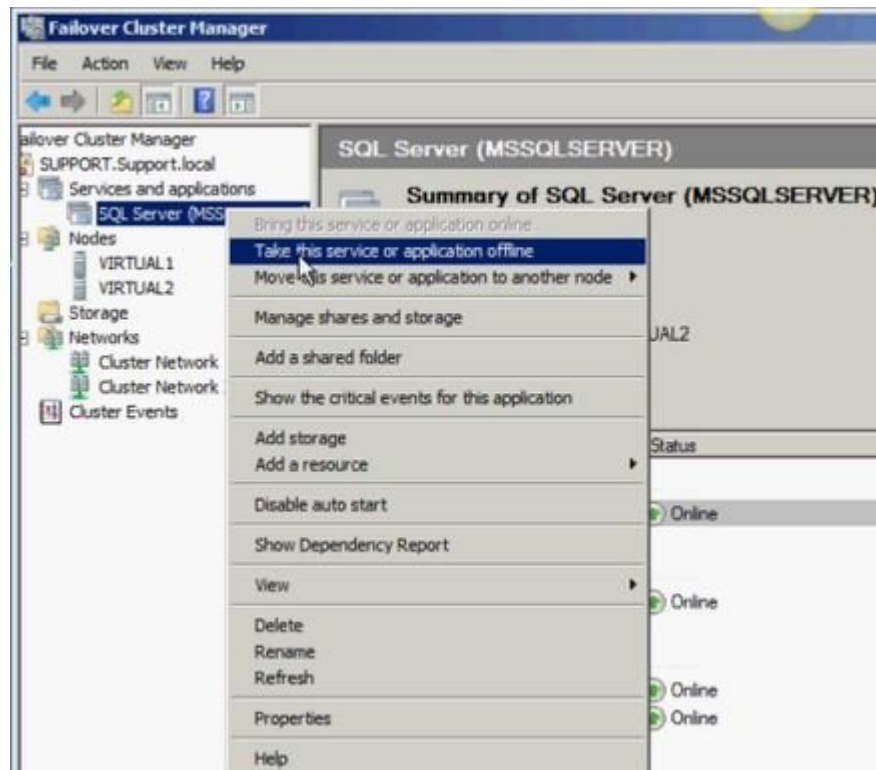
1. Remove physical disk resource from WSFC.

This also removes any dependencies on these physical disk resources. These dependencies will need to be reestablished to the new DataKeeper volume resource, so please take note of what these dependencies are before completing this first step by viewing the **Dependency Report**. Highlight your resource and select **Show Dependency Report**.




Use Failover Cluster Manager MMC to perform the following steps:

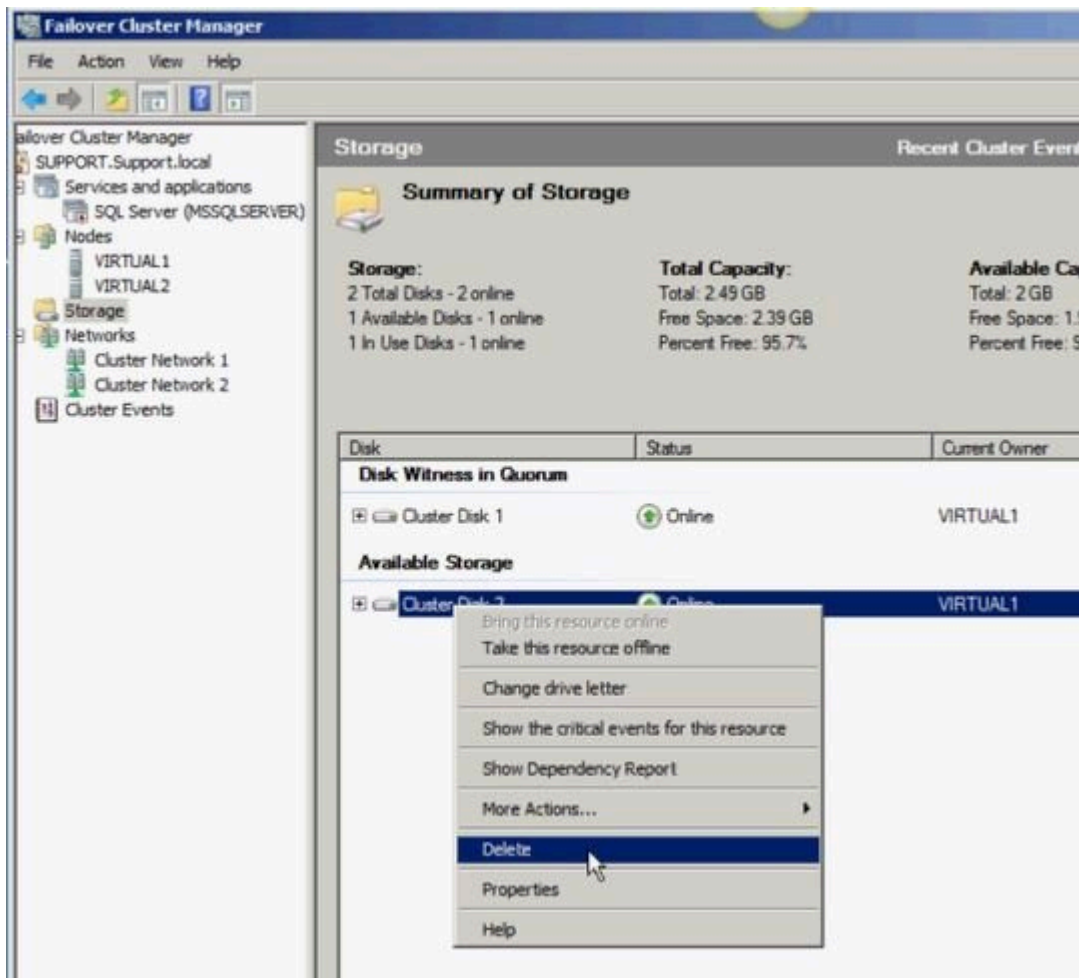
- a. Offline the SQL Server cluster resource group by right-clicking and selecting **Take this service or application offline**.



b. Remove the physical disk from the SQL Server cluster resource group (*moves to Available Storage*).



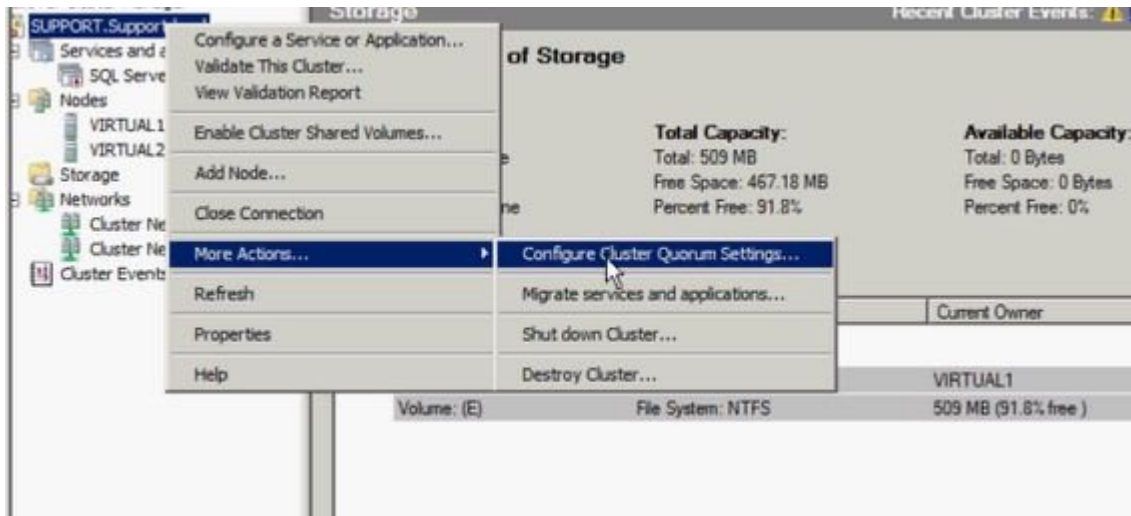
c. Remove the physical disk resource from the SQL Server cluster configuration by deleting the resource from the **Available Storage** group.



2. Configure cluster quorum settings.

Because there will now be a third node in a remote site, the **Disk Witness in Quorum** is no longer valid; therefore, the **Node Majority** configuration should be selected.

- a. Right-click the cluster and select **More Actions / Configure Cluster Quorum Settings...**




b. Select **Node Majority**. **Note:** Change the quorum type to **Node Majority** if the final number of nodes is odd as in a 3-node cluster, or **Node and File Share Witness Majority** if the final number of nodes will be even as in a 4-node cluster.

c. Delete disk witness from **Available Storage**.

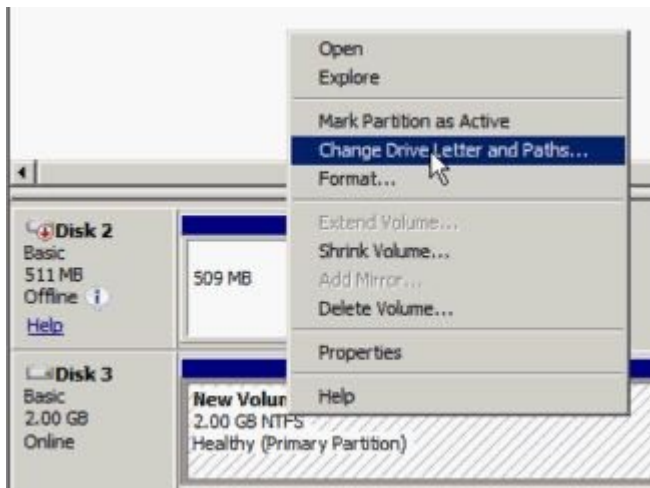
3. Online Disk.

Use Disk Management MMC to perform the following steps:

a. Online the physical disk.



b. Change the drive letter to match the previous SQL Server configuration (if needed).



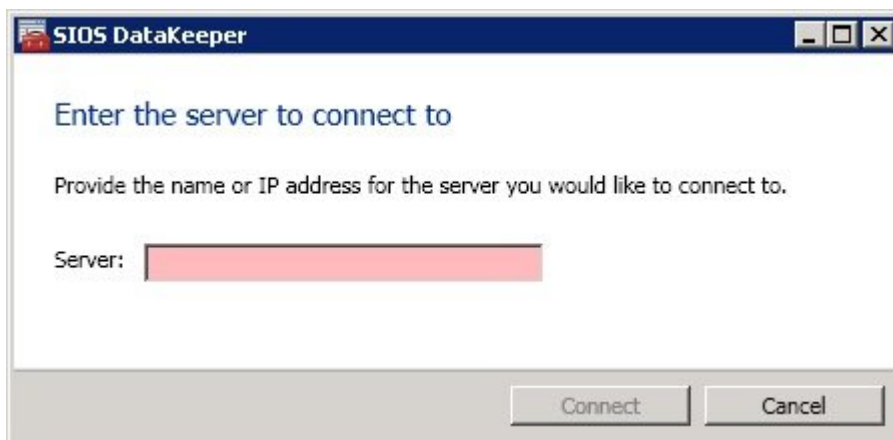
c. Repeat the above steps, 2a and 2b, on all shared nodes in the cluster.

4. Make sure the volume of the third node is available to become a target. It should be formatted and online, and it should be at least as big or bigger than the source volume.

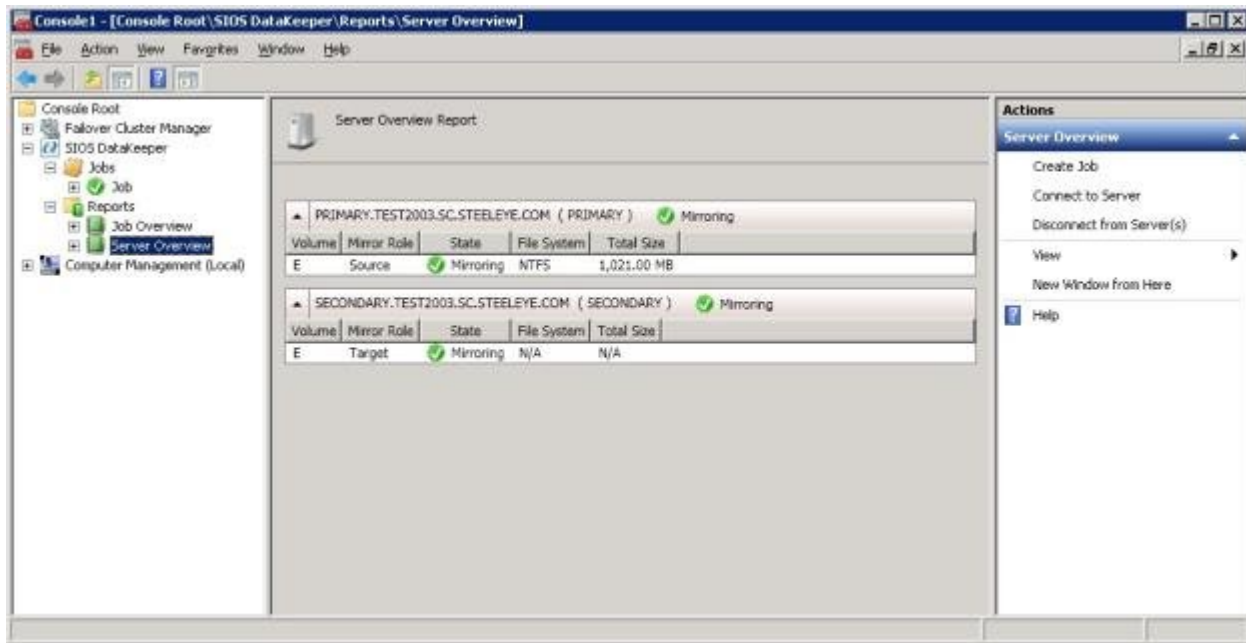
5. **Create Mirror.**

Use the DataKeeper UI to perform the following steps (**Note:** Make sure DataKeeper is running on all servers):


- a. Connect to all shared nodes and the third node.



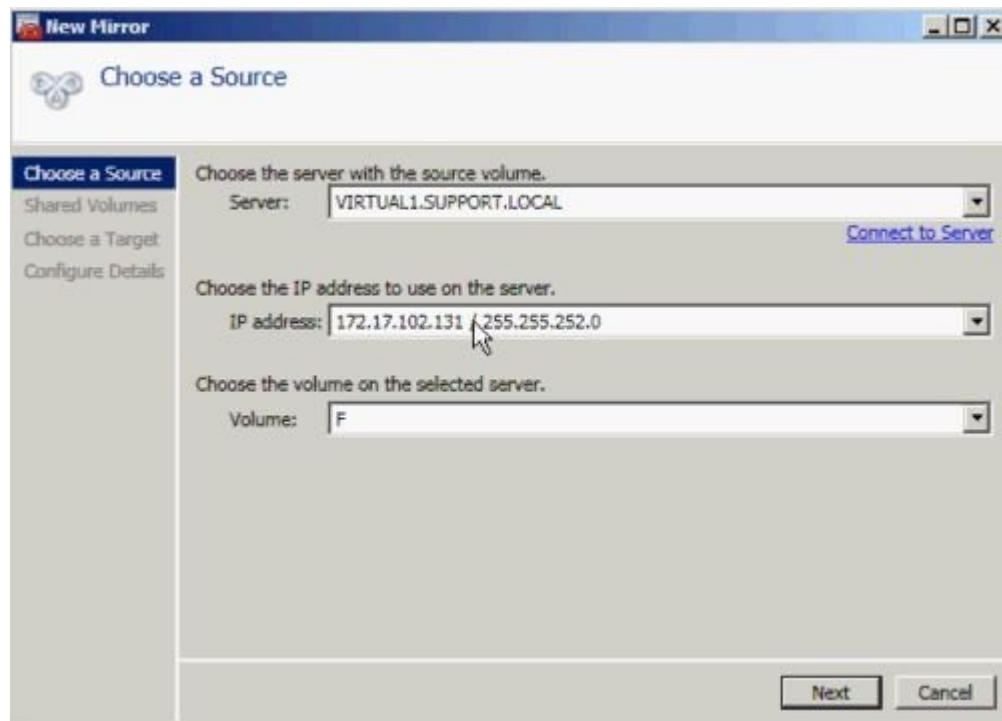
The Server Overview Report will show connection to all three nodes:



b. Create a job containing a mirror to the third node.



c. Select a **Source**.



The "New Mirror" dialog box is shown with the "Choose a Source" tab selected. The left sidebar contains "Choose a Source", "Shared Volumes", "Choose a Target", and "Configure Details". The main area has three sections: "Choose the server with the source volume." with a "Server" dropdown set to "VIRTUAL1.SUPPORT.LOCAL" and a "Connect to Server" link; "Choose the IP address to use on the server." with an "IP address" dropdown set to "172.17.102.131 / 255.255.252.0"; and "Choose the volume on the selected server." with a "Volume" dropdown set to "F". "Next" and "Cancel" buttons are at the bottom right.

d. Select **Shared Volumes**.

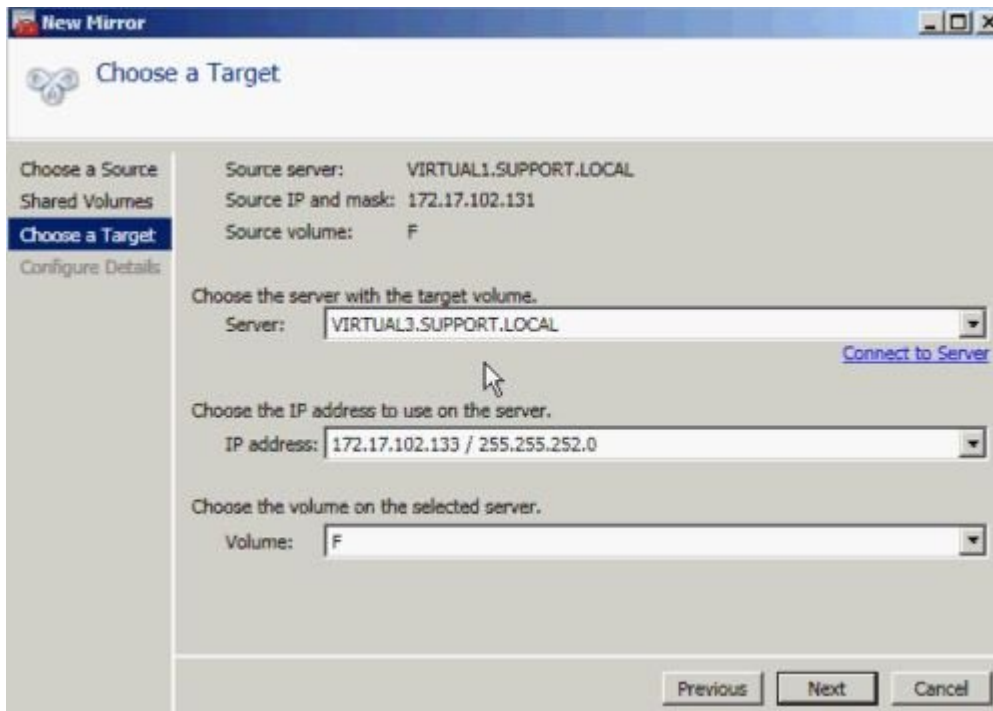


The "New Mirror" dialog box is shown with the "Shared Volumes" tab selected. The left sidebar contains "Choose a Source", "Shared Volumes", "Choose a Target", and "Configure Details". The main area displays the source configuration: "Source server: VIRTUAL1.SUPPORT.LOCAL", "Source IP address: 172.17.102.131", and "Source volume: F". Below this is a text block explaining the "Include" checkbox and subnet requirements. A table lists systems to include:

Include	Server	Volume	Source IP / Mask
<input checked="" type="checkbox"/>	VIRTUAL2.SUPPORT.LOCAL	F	172.17.102.132 / 22

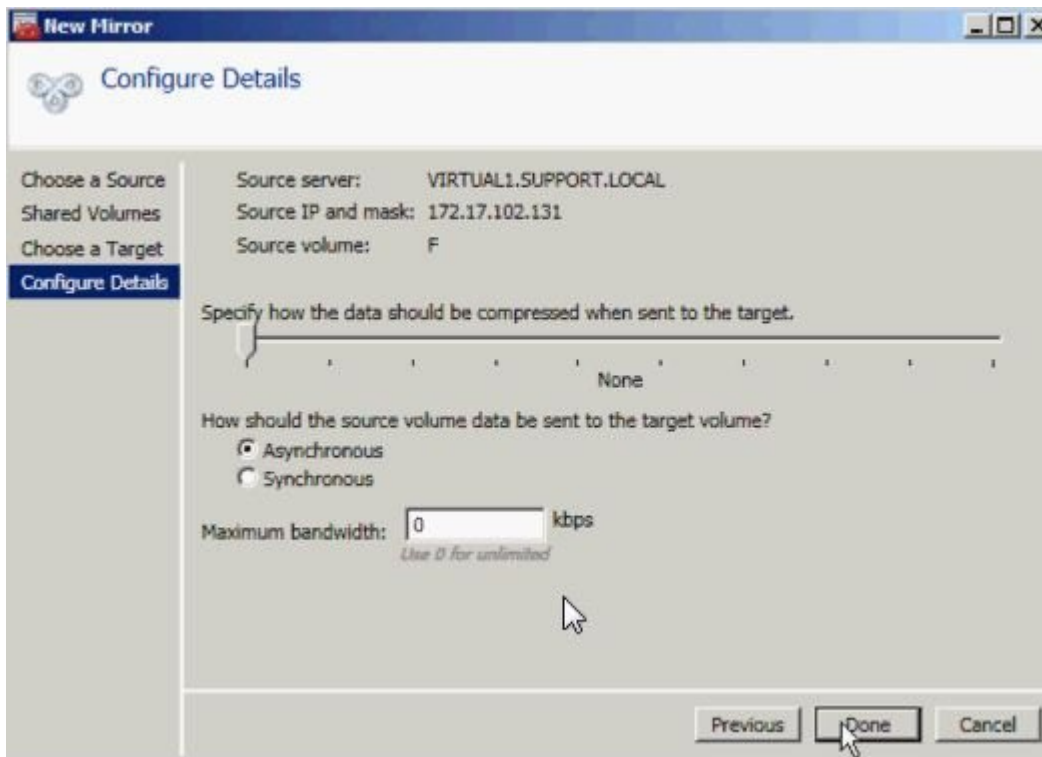
Below the table is a "Connect to Server" link. "Previous", "Next", and "Cancel" buttons are at the bottom right.

e. Select a **Target**.




The 'Choose a Target' dialog box in the 'New Mirror' wizard. It features a sidebar with 'Choose a Source', 'Shared Volumes', 'Choose a Target' (selected), and 'Configure Details'. The main area contains fields for 'Source server' (VIRTUAL1.SUPPORT.LOCAL), 'Source IP and mask' (172.17.102.131), and 'Source volume' (F). Below these are three sections: 'Choose the server with the target volume.' with a 'Server' dropdown (VIRTUAL3.SUPPORT.LOCAL) and a 'Connect to Server' link; 'Choose the IP address to use on the server.' with an 'IP address' dropdown (172.17.102.133 / 255.255.252.0); and 'Choose the volume on the selected server.' with a 'Volume' dropdown (F). At the bottom are 'Previous', 'Next', and 'Cancel' buttons.

f. **Configure Details.** Choose **compression settings**, if applicable, and choose **mirror type**, [Asynchronous or Synchronous](#). Select **Done**.



The 'Configure Details' dialog box in the 'New Mirror' wizard. The sidebar now highlights 'Configure Details'. The source information remains the same. The main area includes a slider for 'Specify how the data should be compressed when sent to the target.' with a 'None' label. Below is a section 'How should the source volume data be sent to the target volume?' with radio buttons for 'Asynchronous' (selected) and 'Synchronous'. A 'Maximum bandwidth' field is set to '0' kbps, with a note 'Use 0 for unlimited'. At the bottom are 'Previous', 'Done' (with a mouse cursor), and 'Cancel' buttons.

The mirror will then begin resyncing to the third node. After resyncing is complete, the **Job State** will change to **Mirroring**.



6. Add **DataKeeper Volume Resource**.

To add the DataKeeper Volume Resource to the SQL Server cluster resource group, perform one of the following:

[Automatic Creation of a Mirror in WSFC](#)

[Manual Creation of a Mirror in WSFC](#)

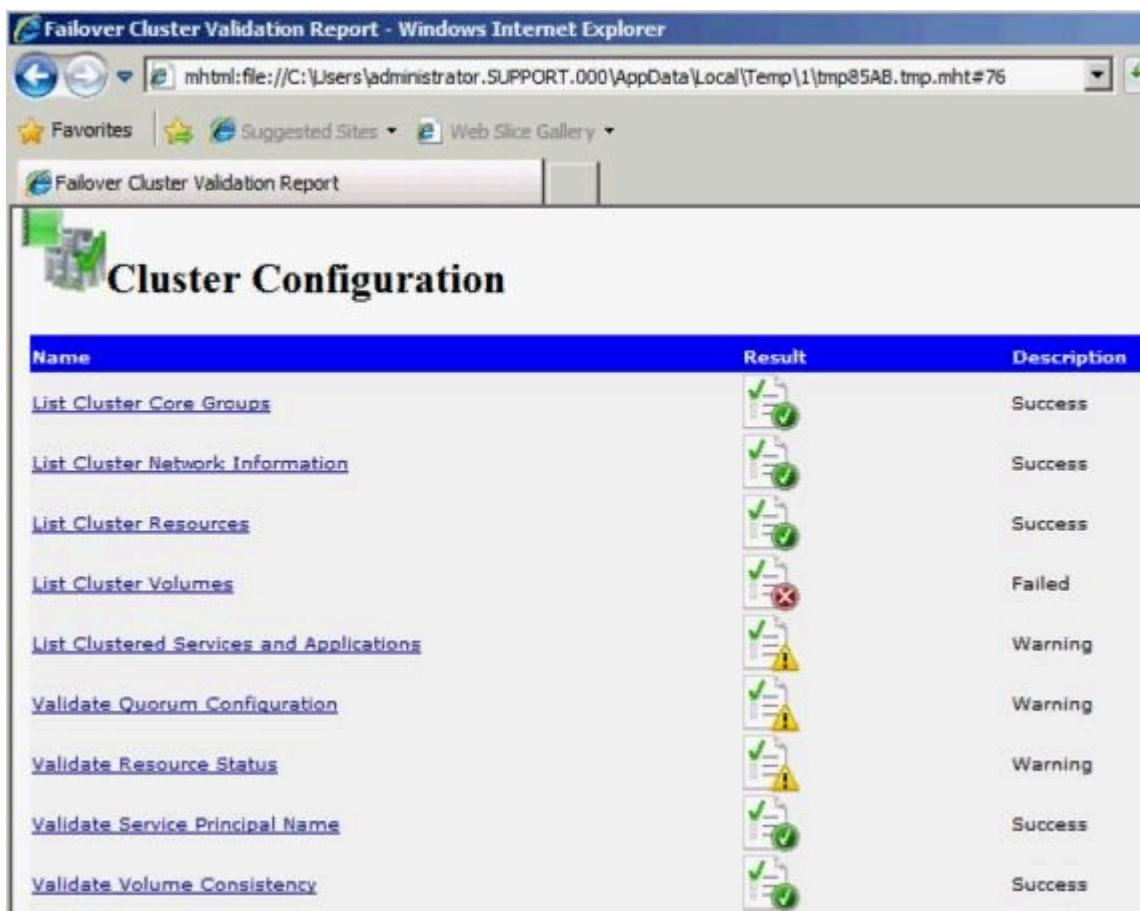
7. Add third node into the cluster.

Use **Failover Cluster Manager MMC** to perform the following steps (**Note:** Make sure the Failover Clustering feature is installed on the third node prior to adding into cluster):

- a. Right-click **Nodes**.
- b. Select **Add Node**.
- c. In the **Add Node Wizard**, enter name of server to be added, click **Add** and select **Next**.
- d. **Yes** may be chosen to perform **Validation Testing**, but be aware that some errors are expected due to the disk being locked on the target side. **Note:** When performing validation testing, select **Run only the tests I select**, then deselect the **Storage tests** from the **Test Selection** screen.



Even though the Storage tests were deselected, the **Failover Cluster Validation Report** will still report a **List Cluster Volumes failure** in the Cluster Configuration Test. This, along with a few other warnings, is expected due to the third node and its volume not being shared with the first two nodes.



e. After viewing the **Validation Report**, select **Finish**. The **Validation Warning screen** will reappear. Select **No** to add the node without performing the tests again.

f. Select **Next** on the **Confirmation** screen.

g. The **Summary** screen will appear indicating that the node was added successfully. A **Warning** may appear on this screen because SQL Server has not yet been installed on the third node. This

installation will be performed in the next step.

8. Install SQL Server on the third node.

a. On third node, run the following command:

```
Setup /SkipRules=Cluster_VerifyForErrors /Action=AddNode  
/INSTANCENAME="MSSQLSERVER"
```

b. On **Setup Support Rules** screen, select **Okay**.

c. When prompted, enter **Product Key** and select **Next**.

d. **Accept License Terms** and select **Next**.

e. On the **Setup Support Files** screen, select **Install**.

f. The **Setup Support Rules** screen will appear. Review the **System Configuration Check Report** for any failures, then select **Next**.

g. Select **Next** when the **Cluster Node Configuration** screen appears.

h. On the **Service Accounts** screen, enter **Passwords** for SQL Server services based on setup of the first nodes. Select **Next**.

Add a Failover Cluster Node

Service Accounts

Specify the service accounts and collation configuration.

Setup Support Rules
Cluster Node Configuration
Service Accounts
Error Reporting
Add Node Rules
Ready to Add Node
Add Node Progress
Complete

Microsoft recommends that you use a separate account for each SQL Server service.


Service	Account Name	Password	Startup Type
SQL Full-text Filter Daemon Launcher	NT AUTHORITY\LOCALSER...		Manual
SQL Server Database Engine	support\administrator	*****	Manual
SQL Server Browser	NT AUTHORITY\LOCAL SER...		Automatic
SQL Server Agent	support\administrator	*****	Manual

Use the same account for all SQL Server services


< Back Next > Cancel Help

- i. When the **Error Reporting** screen appears, select **Next**.
 - j. When the **Add Node Rules** screen appears, select **Next**.
 - k. Once the **Ready to Add Node** screen appears, select **Install**.
 - l. The **Complete** screen will appear notifying that the **Failover Cluster Add Node** operation is complete. Select **Close**.
9. The final step is to reestablish any/all resource dependencies between the SQL Server service resource and the new DataKeeper volume resource.
- a. Right-click the application resource and select **Properties**.
 - b. Select the **Dependencies** tab.
 - c. Click on the **Click here to add a dependency** tab.
 - d. Enter the new DataKeeper volume

e. Select **Apply** and then **Okay**.



Dependencies should now be reestablished. The **Dependency Report** may be viewed again to make sure.



Extending a Traditional 2-Node Cluster to a Shared-Replicated Configuration

Add a Shared Node Using Windows Server 2008R2 or 2012

When using Windows Server 2008R2 or 2012, a 1×1 2-Node replicated cluster using DataKeeper Volumes can be extended to a 2×1 shared and replicated 3-Node cluster by using either of the following methods:

- WSFC GUI
- WSFC command line tool: `"cluster /add /node:<standby node name>"`
- powershell -command : `"Add-ClusterNode -Name <host name>"`

In Windows Server 2008R2 or 2012, DataKeeper shared disks will remain DataKeeper Volume Resources in the cluster when additional nodes with shared disks are added. This is because the shared disk is never accessible on every node in the cluster – only two systems in a 3-node cluster. It is very important that the DataKeeper Volume Resource not be converted to a WSFC Physical Disk Resource.

Add a Shared Node Using Windows Server 2008R2 "SP1"

!WARNING – The WSFC mmc GUI behavior on Windows Server 2008R2 "SP1" has changed!

Starting with WSFC 2008R2SP1, Microsoft has changed the behavior of the WSFC mmc GUI when adding nodes to a cluster. If the new node is hosting a disk that is shared by one or more other systems already in the cluster, the shared disk on the new node and the existing DataKeeper Volume Resource will automatically be converted to a WSFC Physical Disk Resource when the node is added to the cluster! The transformation process includes a volume letter change that will break the DataKeeper Volume Resource and the associated replication mirror. Clustered applications will likely be impacted.

When using Windows Server 2008R2SP1, do not use the WSFC GUI to add a node to the cluster if the new node is hosting a DataKeeper shared disk.

When using Windows Server 2008R2SP1, additional nodes with DataKeeper shared volumes can be safely added to a WSFC cluster using only the WSFC command line tool as follows:


- WSFC command line tool: "cluster /add /node:<standby node name>"

Switchover in an N-Shared x N-Shared Configuration

In a multi-shared storage environment between two sites (as shown in the diagram below), each server in each site has access to the storage being shared between the servers at that site. When a DataKeeper mirror is created, one server in each site will be designated as the mirror endpoints of the mirror.

(**Note:** N represents a number from 1 to N; e.g. 4x1 represents 4 servers sharing a disk replication to another site with 1 server.)

In the following example (3x3), a DataKeeper mirror has been created to replicate the X: volume from Site A to Site B.



Note that there are three servers sharing storage in Site A. Those servers are:

- S1 - Currently the SOURCE of the mirror
- S2 - Shared Source (locked)
- S3 - Shared Source (locked)

Because S1 is the SOURCE of the mirror, we refer to S2 and S3 as shared

source systems. This implies that these servers currently share access to the volume on the SOURCE side of the mirror, but they cannot currently access the volume because they are not defined as the source of the mirror. (**Note:** A user on S2 will see "**access denied.**")

There are three other servers sharing storage in Site B.


- T1 - Currently the TARGET of the mirror (locked as the primary target)
- T2 - Shared Target (locked)
- T3 - Shared Target (locked)

They share access to the Target volume currently defined to T1. T2 and T3 are referred to as shared target systems. File system access to the target volume is locked on all three systems.

All six servers are part of the WSFC cluster, and all of the WSFC protected resources are currently active or "ONLINE" on S1.

Given this initial replicated configuration, it is important to understand which servers are eligible to "take over" and become the ACTIVE server. With a DataKeeper mirror in place, the following rules apply:

1. Switchover to a Shared Source server (S2, S3) is allowed.
2. Switchover to the current Target server (T1) is allowed.
3. Switchover to a Shared Target server (T2, T3) is not allowed; however, there is a two-step process to switch over to these servers.
 - First, switch over to the target T1.
 - Then you can switch over to either T2 or T3 server.



Switchover to a Shared Source Server

In our example, either S2 or S3 is eligible to become the ACTIVE server and source of the mirror. If we switch the protected resources to S2, S2 will become the new SOURCE of the mirror and T1 will remain the Target of the mirror.

1. Initial Mirror Configuration: S1 → T1
2. Action: Switchover to S2 (Bring resources ONLINE in WSFC)
3. Final Result: S2 → T1

Switchover to the Current Target System

In our example, switching the protected resources to T1 will effectively reverse the mirror direction making T1 the new SOURCE of the mirror and S1 will become the Target of the mirror.

1. Initial Mirror Configuration: S1 → T1
2. Action: Switchover to T1 (Bring resources ONLINE in WSFC)
3. Final Result: T1 → S1

Switchover to a Shared Target System

This is not allowed. The switchover operation will fail, but as noted

above, a two-step process can be used.

1. Initial Mirror Configuration: S1 → T1
2. Action: Switchover to T1 (Bring resources ONLINE in WSFC)
3. Intermediate Configuration: T1 → S1
4. Action: Switchover to T2 (Bring resources ONLINE in WSFC)
5. Final Result: T2 → S1

Failover

In a failover scenario where the current source system fails or a resource failure triggers a group failure, Failover Clustering attempts to fail over the resource group to another node in the cluster. Factors affecting the failover include the following:

- Node Failure vs Resource Failure
- Preferred Owner List is set for the Resource Group vs Preferred Owner List is not set
- Possible Owners

The following article describes how Failover Clustering determines which node to fail over to.


<http://blogs.msdn.com/b/clustering/archive/2009/08/11/9863688.aspx>

Important: In an N x N configuration, DataKeeper supports failover to any one of the shared source systems or to the target system. If Failover Clustering tries to online the group on an ineligible system (one of the shared target systems), that online will fail. Failover Clustering will continue trying to online the group on different nodes and eventually succeeds when a system that is eligible for failover is tried.

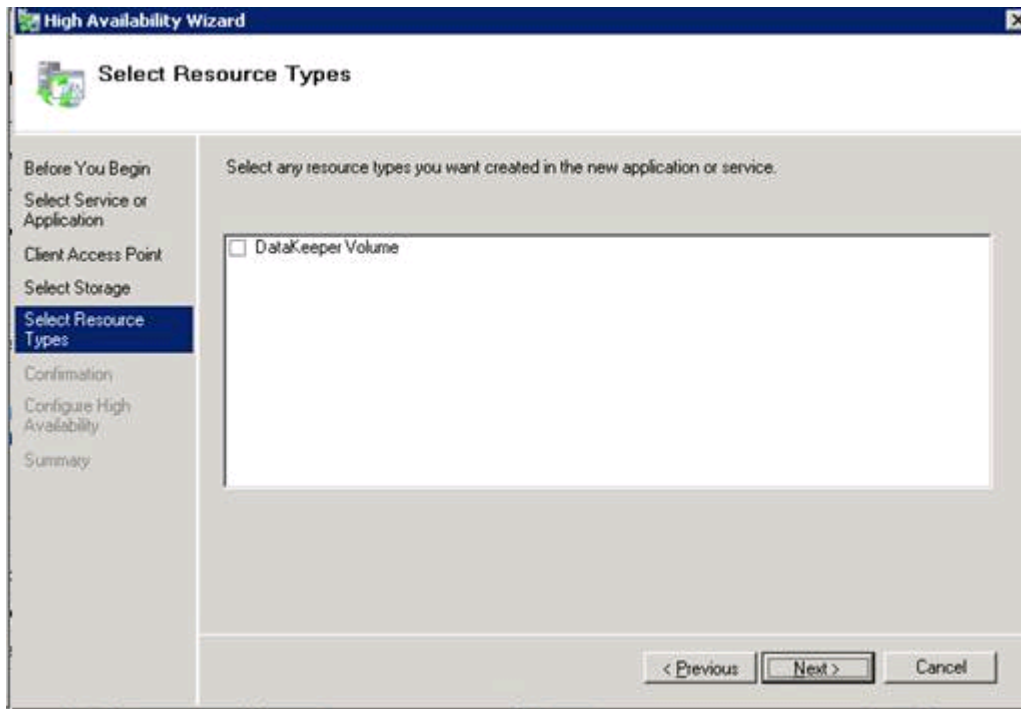
Creating Other Server Resource in WSFC

When creating an **Other Server** resource in WSFC, you will be prompted for a storage resource (DataKeeper resource) twice. You will check the first one, but do not need to check the second dialog.

Click the "checkbox" for the first dialog "Select Storage".



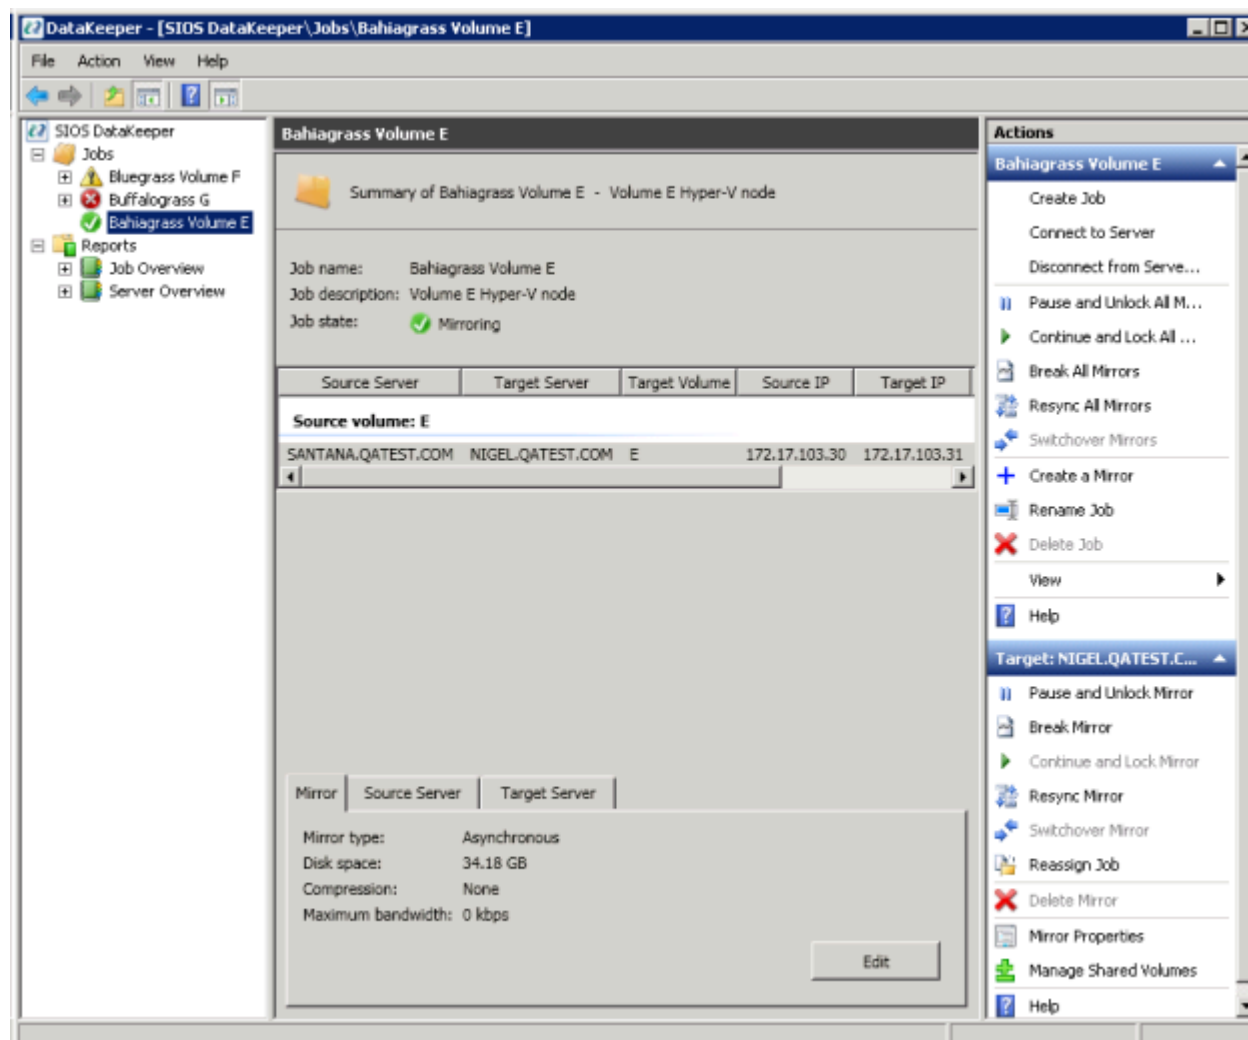
But **DO NOT** check the box in the 2nd dialog "Select Resource Types". This will attempt to create another resource, which is not required.



User Interface

SIOS DataKeeper User Interface

The SIOS DataKeeper User Interface uses a standard MMC snap-in interface.



- The left pane displays the Console Tree view. This includes the **Jobs** and **Reports**. Currently, there are two reports available - **Job Overview** and **Server Overview**. The **Job Overview** report provides a summary of all the jobs on the connected servers. The **Server Overview** report provides a summary of all the mirrors on the connected servers.
- The middle pane is the **Summary** view. This includes information

about the selected item.

- The right column is the **Actions** view. This pane appears when activated through the **View** menu. The options available from this pane are the same options available from the **Action** menu. This column is divided into two sections. The **Actions** in the top section apply to the job and every mirror within the job. The **Actions** in the bottom section apply only to the selected mirror.
- At the bottom of the main window, three tabs appear: **Mirror**, **Source Server** and **Target Server**. These tabs provide information on the mirror that has been selected.
- The icon shows the state of the mirror, which provides more information than the icons and states provided in the Failover cluster UI.

Administration

The topics in this section provide detailed instructions for performing DataKeeper administration tasks.

[DataKeeper Event Log Notification](#)

[Primary Server Shutdown](#)

[Secondary Server Failures](#)

[Extensive Write Considerations](#)

[CHKDSK Considerations](#)

[DKHEALTHCHECK](#)

[DKSUPPORT](#)

[Event Log Considerations](#)

[Using Disk Management](#)

[Registry Entries](#)

[Using EMCMD with SIOS DataKeeper](#)

[Using DKPwrShell with SIOS DataKeeper](#)

Components

DataKeeper Components

SIOS DataKeeper for Windows is comprised of the following components:

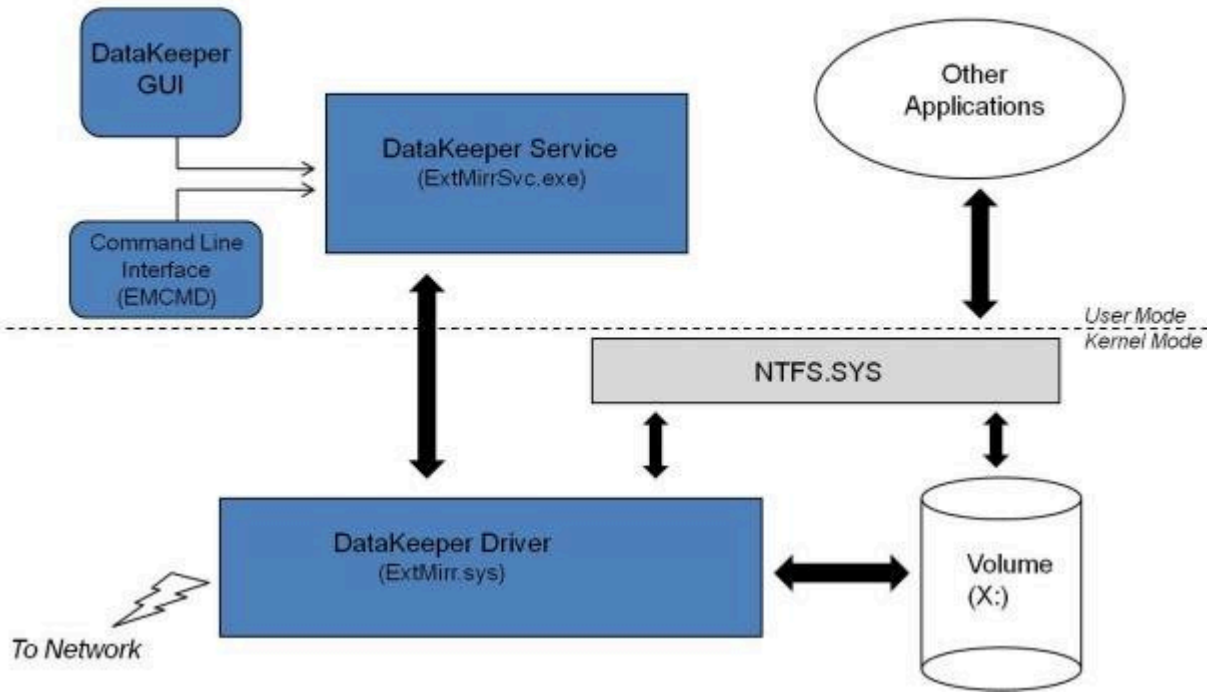
- **DataKeeper Driver (ExtMirr.sys)** - The DataKeeper Driver is a kernel mode driver and is responsible for all mirroring activity between the mirror endpoints.
- **DataKeeper Service (ExtMirrSvc.exe)** - The DataKeeper Service links the DataKeeper GUI and Command Line Interface to the DataKeeper Driver. All commands to manipulate the mirror are relayed through the DataKeeper Service to the DataKeeper Driver.

Important: Stopping the DataKeeper Service does not stop mirroring. Sending the driver a PAUSE mirror, BREAK mirror or DELETE mirror command is the only way to interrupt mirroring.

- **DataKeeper Service Log On ID and Password Selection** - The [DataKeeper Service Log On ID and Password Selection](#) allows you to select the type of account to be used to start the service. Domain and Server account IDs with administrator privileges allow improved disaster recovery when network disruptions occur.
- **Command Line Interface (EMCMD.exe)** - There is an entire suite of [EMCMD command options](#) that can be used to operate DataKeeper.
- **DataKeeper GUI (Datakeeper.msc)** - The [DataKeeper GUI](#) is an MMC 3.0 (Microsoft Management Console) based user interface which allows you to control mirroring activity and obtain mirror status.
- **Packaging files, SIOS Protection Suite scripts, help files, etc.**

The following diagram displays how the DataKeeper components interface with the NTFS file system and each other to perform data replication.

DataKeeper Architecture



Server 2012: Server Manager "File and Storage Services" Disk Status

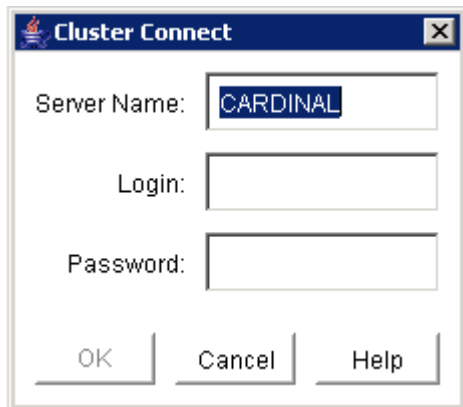
Question

Why are DataKeeper Volumes that are being used in a Microsoft Failover Cluster not shown by Server Manager as "Clustered"?

Answer

On Server 2012, the new Server Manager tool is capable of detecting if a complete "Disk" is being used by Failover Cluster, for example as a Microsoft Cluster Shared Disk. However, the tool cannot detect if one or more Volumes located on a "Disk" are being used as DataKeeper (replicated) Volumes in a Cluster.

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.

Setup

Setting Up SIOS DataKeeper

Follow these steps to start using SIOS DataKeeper:

1. [Connect to the servers](#) you wish to configure for replication. You can select **Connect to Server** from the **Action** pull down menu, right-click on the job folder in the left panel tree display and select **Connect to Server** or choose **Connect to Server** from the **Actions** pane.
2. [Create a Job](#). From the right **Actions** pane, select **Create Job** or you can right-click on the job folder in the left panel tree and select **Create Job**.
3. [Create a mirror](#) for the new job.

SIOS DataKeeper for Windows Technical Documentation

Localized Language Supplement

For information regarding the Localized Language Supplement, please refer to the topic [Installing LifeKeeper for Windows Localized Language Supplement](#) in the SIOS Protection Suite Documentation.

SIOS DataKeeper for Windows

Your information resource for SIOS DataKeeper Standard Edition

SIOS Technology Corp. maintains documentation for all supported versions of SIOS DataKeeper Standard Edition. We welcome your suggestions and feedback. To help us continue to improve our documentation, please complete our brief [Documentation Feedback Survey](#).