



SIOS DataKeeper Cluster Edition

8.6.5 — Last update: 2019/11/15

SIOS TECHNOLOGY CORP.

Table of Contents

- SIOS DataKeeper Cluster Edition 7
- DataKeeper Cluster Edition Quick Start Guide 8
- Deploying DataKeeper Cluster Edition in AWS 12
- Deploying DataKeeper Cluster Edition in Azure 37
- DataKeeper Cluster Edition Release Notes..... 69
- DataKeeper Cluster Edition Installation Guide 75
 - Installation 76
 - Core Software..... 77
 - Installing Core 78
 - Third Party Product Files 80
 - Application Directory Anomaly 82
 - Localized Language Supplement 83
 - Silent Installation 84
 - Removing a Clustered DataKeeper Volume 86
 - Licensing 90
 - Uninstalling SIOS DataKeeper Cluster Edition 94
 - Upgrading SIOS DataKeeper Cluster Edition..... 96
- DataKeeper Cluster Edition Technical Documentation 98
 - User Interface 100
 - Components..... 102
 - DataKeeper Service Log On ID and Password Selection..... 104
 - Understanding Replication 110
 - SIOS DataKeeper Intent Log 111
 - Relocation of Intent Log 113
 - Resynchronization 115
 - Synchronous and Asynchronous Mirroring 117
 - Read and Write Operations..... 124
 - Volume Considerations..... 126
 - Specifying Network Cards for Mirroring..... 127
 - Performance Monitor Counters 128
 - Configuration 135
 - Sector Size 136
 - Network Bandwidth 137

Network Adapter Settings	139
DataKeeper Service Log On ID and Password Selection	141
Firewall Configurations	148
High-Speed Storage Best Practices	154
Configuration of Data Replication From a Cluster Node to External DR Site	156
Performance Tuning	157
Disable “Automatically manage paging file size for all drives”	158
WAN Considerations.....	162
Initial Synchronization of Data Across the LAN or WAN	163
Compression	167
Bandwidth Throttle	168
Administration	169
DataKeeper Event Log Notification	170
Primary Server Shutdown	172
Secondary Server Failures.....	173
Extensive Write Considerations	174
CHKDSK Considerations	175
DKHEALTHCHECK	176
DKSUPPORT	177
Event Log Considerations	178
Using Disk Management	179
Registry Entries	180
Using EMCMD with SIOS DataKeeper	192
Mirror State Definitions	196
Using the -proxy option with EMCMD.....	197
BREAKMIRROR	198
CHANGEMIRRORENDPOINTS	199
CHANGEMIRRORTYPE	205
CLEARBLOCKTARGET.....	208
CLEARSNAPSHOTLOCATION	209
CLEARSWITCHOVER	210
CONTINUEMIRROR	211
CREATEJOB	212
CREATEMIRROR	213
DELETEJOB	215
DELETELOCALMIRRORONLY	216
DELETEMIRROR	217
DROPSNAPSHOT	218
GETBLOCKTARGET	219

GETCOMPLETEVOLUMELIST	220
GETCONFIGURATION	221
GETEXTENDEDVOLUMEINFO	222
GETJOBINFO	223
GETJOBINFOFORVOL	224
GETMIRRORTYPE	225
GETMIRRORVOLINFO	226
GETREMOTEBITMAP	228
GETRESYNCSTATUS	229
GETSERVICEINFO	231
GETSNAPSHOTLOCATION	232
GETSOURCEMIRROREDVOLUMES	233
GETTARGETMIRROREDVOLUMES	234
GETVOLUMEDRVSTATE	235
GETVOLUMEINFO	236
ISBREAKUSERREQUESTED	238
ISPOTENTIALMIRRORVOL	239
LOCKVOLUME	241
MERGETARGETBITMAP	242
PAUSEMIRROR	243
PREPARETOBECOMETARGET	244
READREGISTRY	245
REGISTERCLUSTERVOLUME	247
RESTARTVOLUMEPIPE	248
RESYNCMIRROR	249
SETBLOCKTARGET	250
SETCONFIGURATION	251
SETSNAPSHOTLOCATION	253
STOPSERVICE	254
SWITCHOVERVOLUME	255
TAKESNAPSHOT	256
UNLOCKVOLUME	257
UPDATECLUSTERTARGET STATEPROPERTIES	258
UPDATEJOB	259
UPDATEVOLUMEINFO	260
Using DKPwrShell with SIOS DataKeeper	261
New-DataKeeperMirror	262
New-DataKeeperJob	264
Remove-DataKeeperMirror	266

Remove-DataKeeperJob	267
Add-DataKeeperJobPair	268
Get-DataKeeperVolumeInfo	270
User Guide.....	271
Getting Started	272
Disk-to-Disk.....	274
One-to-One	276
One-to-Many (Multiple Targets)	278
Many-to-One	280
N-Shared-Disk Replicated to One.....	282
N-Shared-Disk Replicated to N-Shared-Disk	284
N-Shared-Disk Replicated to Multiple N-Shared-Disk Targets.....	286
Setting Up SIOS DataKeeper	288
Connecting to a Server	289
Disconnecting from a Server	290
Creating a Job	291
Configuring Mirrors	292
Creating a Mirror	293
Creating Mirrors With Shared Volumes	295
Safe Creation of a Shared-Storage Volume Resource	299
Creating Mirrors With Multiple Targets.....	301
Switchover and Failover with Multiple Targets	303
Working With Jobs	307
Jobs	308
Renaming a Job	310
Deleting a Job	311
Reassigning a Job	312
Switching Over a Mirror	313
Working With Mirrors	315
Managing Mirrors	316
Pause and Unlock	317
Continue and Lock.....	318
Partial Resync	319
Break	320
Resync	321
Deleting a Mirror.....	322
Replacing a Target	323
DataKeeper Volume Resize.....	324
Mirror Properties	327

Changing the Compression Level of an Existing Mirror	329
Working With Shared Volumes	331
Managing Shared Volumes.....	332
Adding a Shared System	334
Removing a Shared System	335
Using Microsoft iSCSI Target With DataKeeper on Windows 2012.....	336
Installation of the iSCSI Target.....	338
Mirror Creation and Cluster Configuration.....	340
Creation of iSCSI Virtual Disks	344
Setup of iSCSI Initiator on Windows 2012	347
DataKeeper Notification Icon	349
DataKeeper Target Snapshot.....	352
Using SIOS DataKeeper Standard Edition To Provide Disaster Recovery For Hyper-V Virtual Machines.....	365
Clustering	378
Running chkdsk on Cluster Volumes during Cluster Volume Online.....	379
Creating a DataKeeper Volume Resource in WSFC	381
Manual Creation of a Mirror in WSFC	383
DataKeeper Volume Resource Health Check	385
DataKeeper Volume Resource Private Properties.....	386
Extending a Clustered DataKeeper Volume to a Node Outside the Cluster.....	389
Extending a Single SQL Server Node to a Cluster	394
Extending a Traditional 2-Node WSFC Cluster to a Third Node via DataKeeper	397
Extending a Traditional 2-Node WSFC SQL Server Cluster to a Third Node via DataKeeper	410
Extending a Traditional 2-Node Cluster to a Shared-Replicated Configuration.....	424
Using DataKeeper Cluster Edition to Enable Multi-Site Hyper-V Clusters	426
Split-Brain Issue and Recovery.....	437
Switchover in an N-Shared x N-Shared Configuration.....	439
Installing and Using DataKeeper Cluster Edition on Windows Server 2008 R2 / 2012 Core	
Platforms.....	443
Non-mirrored Volume Resource	445
Using DKCE to Enable Multi-Site File Share Resources with Windows Server 2008R2 WSFC .	448
Creating Other Server Resource in WSFC.....	456
FAQs	458
Awareness of Windows Filenames and Directory Names	459
AWS Issues and Workarounds.....	460
Change Mirror Endpoints	461
Change Mirror Type	462
Create a Mirror and Rename Job and Delete Job Actions Grayed Out	463

Data Transfer Network Protocols	464
Delete and Switchover Actions Grayed Out.....	465
Deleting a Mirror	466
Error Messages Log.....	467
Inability to Create a Mirror	468
Network Disconnect.....	469
Reclaim Full Capacity of Target Drive	471
Resize or Grow Mirrored Volumes	472
Server 2012: Server Manager “File and Storage Services” Disk Status	473
Split-Brain FAQs	474
Stop Replication Between Source and Target	477
Using Volume Shadow Copy.....	478
Volumes Unavailable for Mirroring	479
DataKeeper Troubleshooting.....	480
Known Issues and Workarounds.....	483
Access to Designated Volume Denied	485
DataKeeper Volume cannot come Online after Network failure with clustered IP Address on	
Replication network	486
DataKeeper Volume Not Available as Cluster Resource Type	487
Failed to Create Mirror.....	488
Hyper-V Host Cluster Error	489
Live Migration Failure	492
MaxResyncPasses Value	493
Mirroring with Dynamic Disks.....	494
New Resources Offline But Unlocked	495
Server Login Accounts and Passwords Must Be Same on Each Server in the Cluster	496
System Event Log – Create Mirror Failed in the GUI	498
Unable to Determine Previous Install Path.....	499
User Interface – Failed to Create Mirror.....	500
User Interface – Shows Only One Side of the Mirror.....	501
WSFC – MS DTC Resource Failure.....	502
WSFC 2008 R2 SP1 Procedure Change.....	503
Windows Server 2012 Specific Issues	504
Windows Server 2012 MMC Snap-in Crash.....	505
Windows Server 2012 — Simultaneous Move of Multiple Clustered File Server Roles Can	
Result in DataKeeper Switchover Failures.....	507
Windows Server 2012 iSCSI Target Role Does Not Support Dynamic Disks	508
Windows Server 2012 Default Information Missing During Mirror Creation	509
Windows Server 2012 NIC Teaming Issue	511

WSFC 2012 Cluster Creation Default Setting Issue	512
WSFC 2012 Failover Cluster Manager UI Defect (Delete Action Missing).....	514
WSFC 2012 File Server Resource Manager Event Log Errors	516
WSFC 2012 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager	517
WSFC 2012 New File Server Type Not Supported	519
WSFC 2012 Server Manager — Incorrect Volume Display	521
WSFC 2012 Server Manager — DataKeeper “Disk” Not Shown as Clustered	523
Windows 2012 File Share	524
Windows Server 2016 Specific Issues	525
Occasional Job Creation Failure.....	526
WSFC 2016 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager	527
Restrictions.....	529
Bitlocker Does Not Support DataKeeper	530
CHANGEMIRRORENDPOINTS	531
CHKDSK	532
DataKeeper Volume Resize Restriction	533
Directory for Bitmap Must Be Created Prior to Relocation.....	534
Duplicate IP Addresses Disallowed Within a Job	535
Intensive I-O with Synchronous Replication	536
Resource Tag Name Restrictions	537
DKCE Support Matrix	538
Product Support Schedule.....	542

SIOS DataKeeper Cluster Edition

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](#).

DataKeeper Cluster Edition Quick Start Guide

This topic provides step-by-step instructions for installing and configuring DataKeeper Cluster Edition. The series of steps includes links to the documentation that describe each step in detail.

Prerequisites and Installation

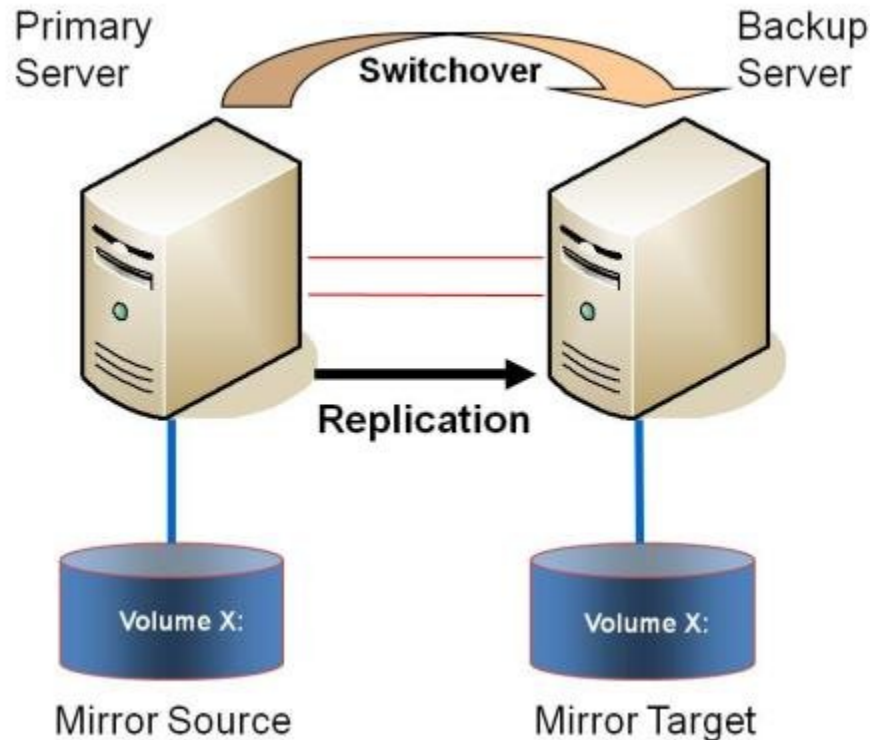
1. Read the [DataKeeper Cluster Edition Release Notes](#) for late breaking information.
2. [Firewall Configurations](#) – Make sure you understand what ports must be opened on any firewalls.
3. [Network Bandwidth](#) – If replicating across a WAN, it is critical that a [rate of change](#) analysis be done to ensure there is adequate bandwidth.
4. DataKeeper is a block-level volume replication solution and requires that each node in the cluster have additional volume(s) (other than the system drive) that are the same size and same drive letters. Please review [Volume Considerations](#) for additional information regarding storage requirements.
5. Configure your Cluster – It is important to have Windows Server configured as a cluster using either a node majority quorum (if there is an odd number of nodes) or a node and file share majority quorum (if there is an even number of nodes). Consult the Microsoft documentation on clustering or this article on the [Clustering for Mere Mortals](#) blog for step-by-step instructions. Microsoft released a [hotfix](#) that allows disabling of a node's vote which may help achieve a higher level of availability in certain multi-site cluster configurations. This hotfix and when it should be used is described in this article in [Clustering for Mere Mortals](#).
6. After the basic cluster is configured but prior to any cluster resources being created, install and license DataKeeper Cluster Edition on all cluster nodes. See the [DataKeeper Cluster Edition Installation Guide](#) for detailed instructions.

 **Note – If installing DataKeeper Cluster Edition on Windows “Core” (GUI-less Windows), make sure to read this section for detailed instructions – [Installing and Using DataKeeper on Windows 2008R2 Server Core Platforms](#).**

Configuration

The following sections describe the most common cluster configurations. Follow the instructions in the section that most likely matches your environment.

2-Node Replicated Cluster



1. The initial configuration must be done from the [DataKeeper UI](#) running on one of the cluster nodes. If it is not possible to run the DataKeeper UI on a cluster node, such as when running DataKeeper on a Windows Core only server, install the DataKeeper UI on any computer running Windows XP or higher and follow the instruction in the [Core Only](#) section for creating a mirror and registering the cluster resources via the CLI.
2. Once the DataKeeper UI is running, [connect to each of the nodes](#) in the cluster.
3. [Create a Job](#) using the DataKeeper UI. This process creates a mirror and adds the DataKeeper Volume resource to the Available Storage.

Note – If clustering Hyper-V VMs, do not add the DataKeeper Volume Resource to Available Storage at the end of the mirror creation process. Instead, allow the mirror to create but do not choose to register the DataKeeper Volume in Available Storage at the end of the Mirror Creation Wizard, then follow the instructions on this page, [Using DataKeeper Cluster Edition to Enable Multi-Site Hyper-V Clusters](#). Make sure that Virtual Network Names for NIC connections are identical on all cluster nodes.

4. If additional mirrors are required, you can [Add a Mirror to a Job](#).
5. With the DataKeeper Volume(s) now in Available Storage, you are able to create cluster resources (SQL, File Server, etc.) in the same way as if there were a shared disk resource in the cluster. Refer to Microsoft documentation for additional information or view this article in [Clustering for Mere Mortals](#)

for step-by-step cluster configuration instructions.

3- or 4-Node Multi-Site Cluster with Mixed Shared/Replicated Storage



1. The initial configuration must be done from the [DataKeeper UI](#) running on one of the cluster nodes. If it is not possible to run the DataKeeper UI on a cluster node, such as when running DataKeeper on a Windows Core only server, install the DataKeeper UI on any computer running Windows XP or higher and follow the instruction in the [Core Only](#) section for creating a mirror and registering the cluster resources via the CLI.
2. Once the DataKeeper UI is running, [connect to each of the nodes](#) in the cluster. Important – In order for DataKeeper to detect that a disk is shared, ALL of the nodes of the cluster must be connected to through the DataKeeper UI.
3. Prior to creating the DataKeeper Job, the storage must be configured such that the nodes located in the same location each have access to the shared storage. The instructions for the [Safe Creation of a Shared-Storage Volume](#) contain the information needed to safely give both servers access to shared storage once the storage has been provisioned and the same LUN has been handed off to each of the shared cluster nodes. The process of provisioning the storage and handing it off to two or more servers at the same time will be dependent upon the storage array. Please refer to your storage documentation for instructions on provisioning storage for clustered environments.
4. [Create a job](#) using the instructions in “[Creating Mirrors With Shared Volumes](#).” This process creates a mirror as well as collects information about the shared disks and then adds the DataKeeper Volume resource to the Available Storage.

Note – If clustering Hyper-V VMs, do not add the DataKeeper Volume Resource to Available Storage at the end of the mirror creation process. Instead, allow the mirror to create but do not choose to register the DataKeeper Volume in Available Storage at the end of the Mirror Creation Wizard, then follow the instructions on this page, [Using DataKeeper Cluster Edition to Enable Multi-Site Hyper-V Clusters](#). Make sure that Virtual Network Names for NIC connections are identical on all cluster nodes.

5. If additional mirrors are required, you can [Add a Mirror to a Job](#).

6. With the DataKeeper Volume(s) now in Available Storage, you are able to create cluster resources (SQL, File Server, etc.) in the same way as if there were a shared disk resource in the cluster. Refer to Microsoft documentation for additional information or view this article in [Clustering for Mere Mortals](#) for step-by-step cluster configuration instructions.

Management

Once a DataKeeper volume is registered with Windows Server Failover Clustering, all of the management of that volume will be done through the Windows Server Failover Clustering interface. All of the management functions normally available in DataKeeper [will be disabled](#) on any volume that is under cluster control. Instead, the DataKeeper Volume cluster resource will control the mirror direction, so when a DataKeeper Volume comes online on a node, that node becomes the source of the mirror. The properties of the DataKeeper Volume cluster resource also display basic mirroring information such as the source, target, type and state of the mirror.

For more information, please refer to the [DataKeeper Cluster Edition Technical Documentation](#).

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](#)

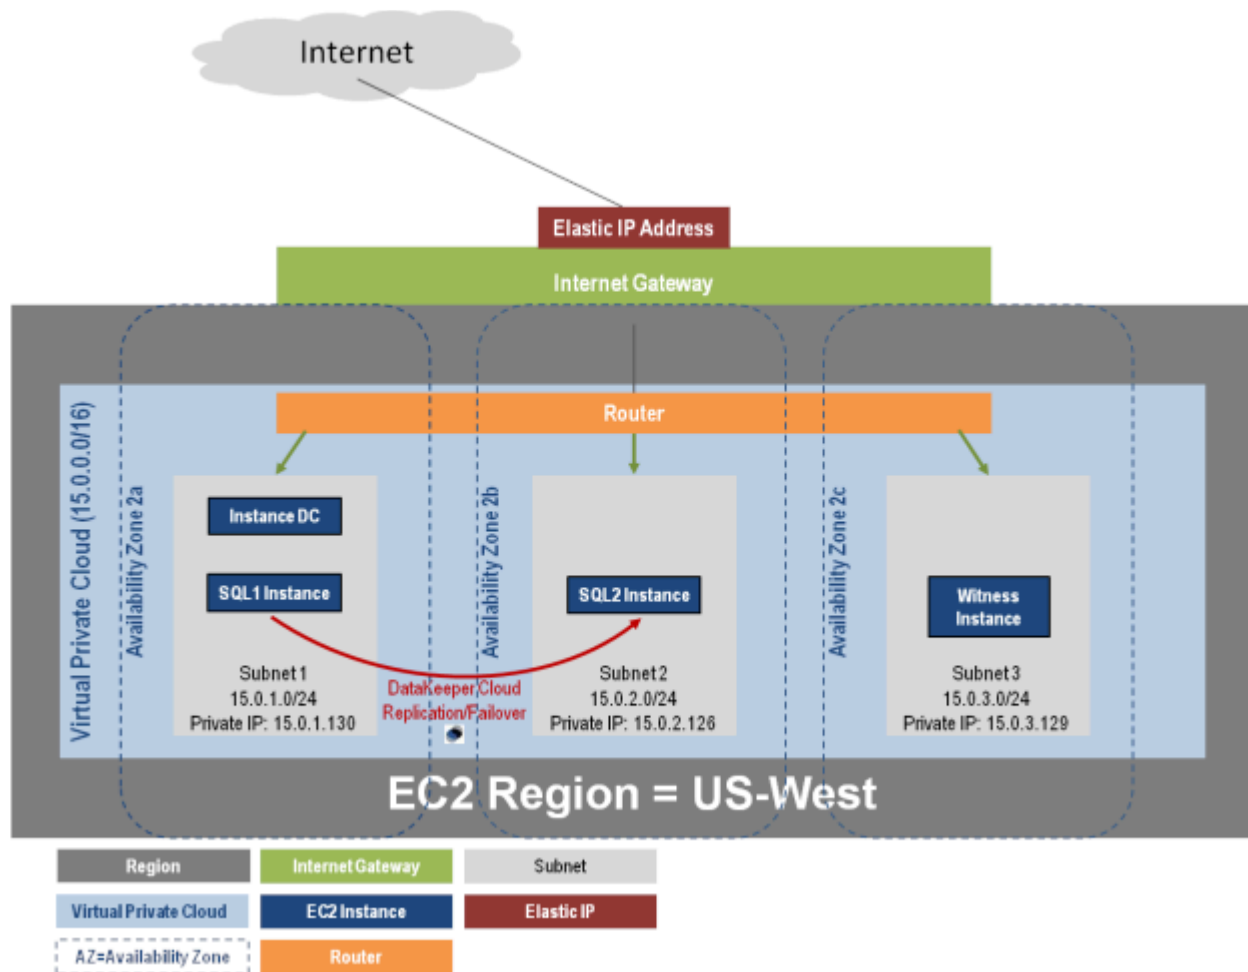
Deploying DataKeeper Cluster Edition in AWS

* **DISCLAIMER:** While the following completely covers the high availability portion within the scope of our product, this is a setup “guide” only and should be adapted to your own configuration.

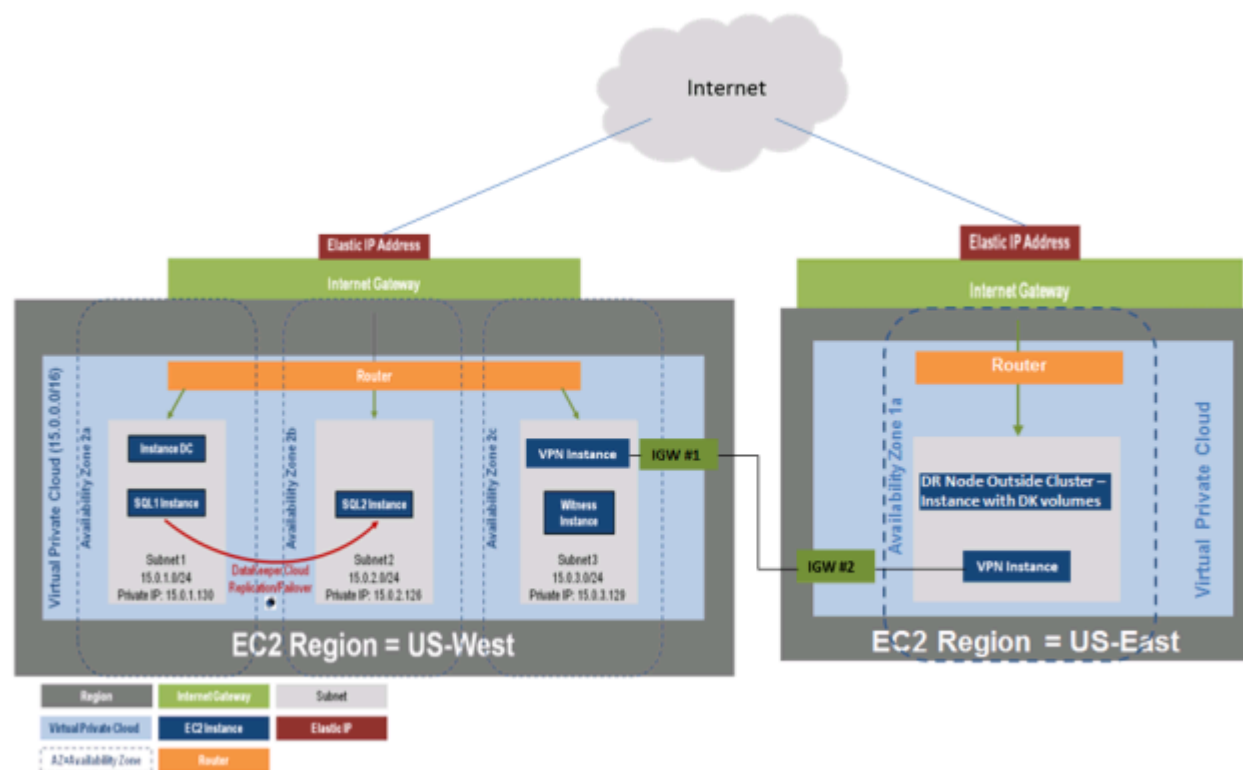
* Please refer to the [AWS Marketplace](#) for a list of supported instance sizes.

Overview

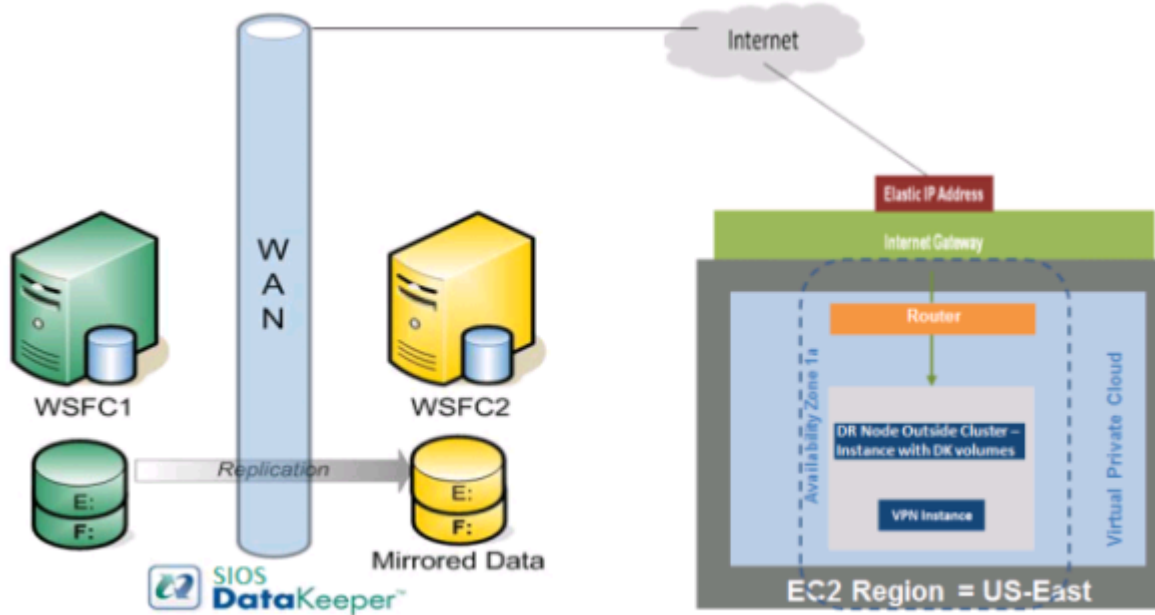
DataKeeper Cluster Edition provides replication in a virtual private cloud (VPC) within a single region across availability zones. In this particular SQL Server clustering example, we will launch four instances (one domain controller instance, two SQL Server instances and a quorum/witness instance) into three availability zones.



DataKeeper Cluster Edition provides support for a data replication node outside of the cluster with all nodes in AWS. In this particular SQL Server clustering example, four instances are launched (one domain controller instance, two SQL Server instances and a quorum/witness instance) into three availability zones. Then an additional DataKeeper instance is launched in a second region including a VPN instance in both regions. Please see [Configuration of Data Replication From a Cluster Node to External DR Site](#) for more information. For additional information on using multiple regions please see [Connecting Multiple VPCs with EC2 Instances](#).



DataKeeper Cluster Edition also provides support for a data replication node outside of the cluster with only the node outside of the cluster in AWS. In this particular SQL Server clustering example, WSFC1 and WSFC2 are in an on-site cluster replicating to an AWS instance. Then an additional DataKeeper instance is launched in a region in AWS. Please see [Configuration of Data Replication From a Cluster Node to External DR Site](#) for more information.



Requirements

Description	Requirement
Virtual Private Cloud	In a single region with three availability zones
Instance Type	Minimum recommended instance type: M1 Medium
Operating System	See the DKCE Support Matrix
Elastic IP	One elastic IP address connected to the domain controller
Four instances	One domain controller instance, two SQL Server instances and one quorum/witness instance
Each SQL Server	ENI (Elastic Network Interface) with 4 IPs <ul style="list-style-type: none"> Primary ENI IP statically defined in Windows and used by DataKeeper Cluster Edition Three IPs maintained by EC2 while used by Windows Failover Clustering, DTC and SQLFC
Volumes	Three volumes (EBS and NTFS only) <ul style="list-style-type: none"> One primary volume (C drive) Two additional volumes <ul style="list-style-type: none"> One for Failover Clustering One for MSDTC

Release Notes

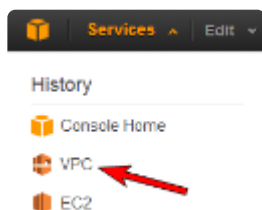
Before beginning, make sure you read the [DataKeeper Cluster Edition Release Notes](#) for the latest information. It is highly recommended that you read and understand the [DataKeeper Cluster Edition Installation Guide](#).

Create a Virtual Private Cloud (VPC)

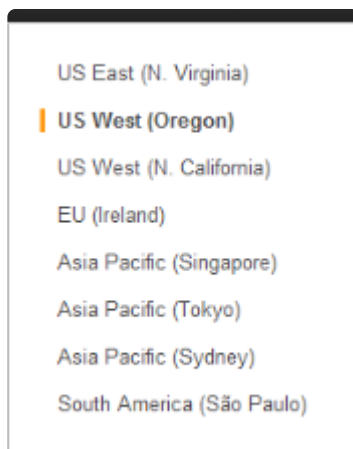
A virtual private cloud is the first object you create when using DataKeeper Cluster Edition.

* A virtual Private Cloud (VPC) is an isolated private cloud consisting of a configurable pool of shared computing resources in a public cloud.

1. Using the email address and password specified when signing up for **Amazon Web Services (AWS)**, sign in to the [AWS Management Console](#).
2. From the **Services** dropdown, select **VPC**.



3. On the right side of the top navigation bar, select the region for your virtual private cloud.



4. On the **VPC Dashboard**, select **Your VPCs** from the left navigation pane.
5. Select **Create VPC**.
6. Define your virtual private cloud subnet by entering your **CIDR (Classless Inter-Domain Routing)** as described below, then click **Yes, Create**.

Create VPC Cancel

A VPC is an isolated portion of the AWS cloud populated by AWS objects, such as Amazon EC2 instances. Please use the Classless Inter-Domain Routing (CIDR) block format to specify your VPC's contiguous IP address range, for example, 10.0.0.0/16. Please note that you can create a VPC no larger than /16.

CIDR Block: (e.g. 10.0.0.0/16)

Tenancy: Default

Cancel Yes, Create

7. Once your virtual private cloud has been successfully created, click **Close** to return to the **VPC Dashboard**.

✿ **HELPFUL LINK:**
[Amazon's Creating a Virtual Private Cloud \(VPC\)](#)

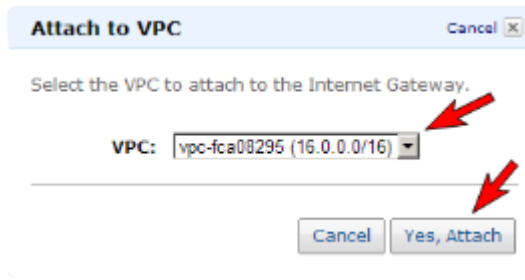
Create Internet Gateway and Attach to Virtual Private Cloud

Create and attach an **Internet Gateway** which provides access to your virtual private cloud from the Internet (from outside the virtual private cloud).

✿ An Internet Gateway connects your VPC directly to the Internet and provides access to other AWS resources.

1. Select **Internet Gateways** from the left navigation pane.
2. Select your **Gateway ID** from the list of **Internet Gateways**.
3. Click **Attach to VPC**.

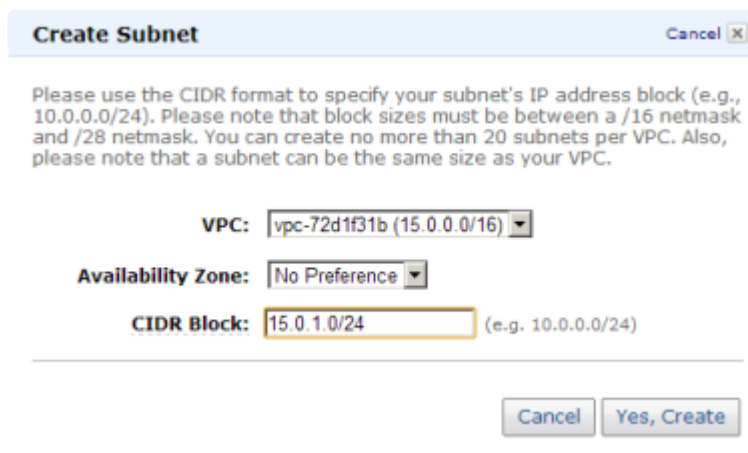
4. Select your **virtual private cloud** from the dropdown list and click **Yes, Attach**.



Configure Availability Zones

Availability zones enable you to group instances based on your security and high availability requirements. You will launch your instances into these availability zones, so for a DataKeeper configuration, you will want to configure at least two availability zones. In this example, we will configure a third availability zone for the quorum witness server.

1. From the left navigation pane of the **VPC Dashboard**, select **Subnets**.
2. Click **Create Subnet**.
3. On the **Create Subnet** dialog, select your **virtual private cloud** and choose an **availability zone**, then enter a **CIDR** based on the instructions in the dialog (shown below). Click **Yes, Create**.

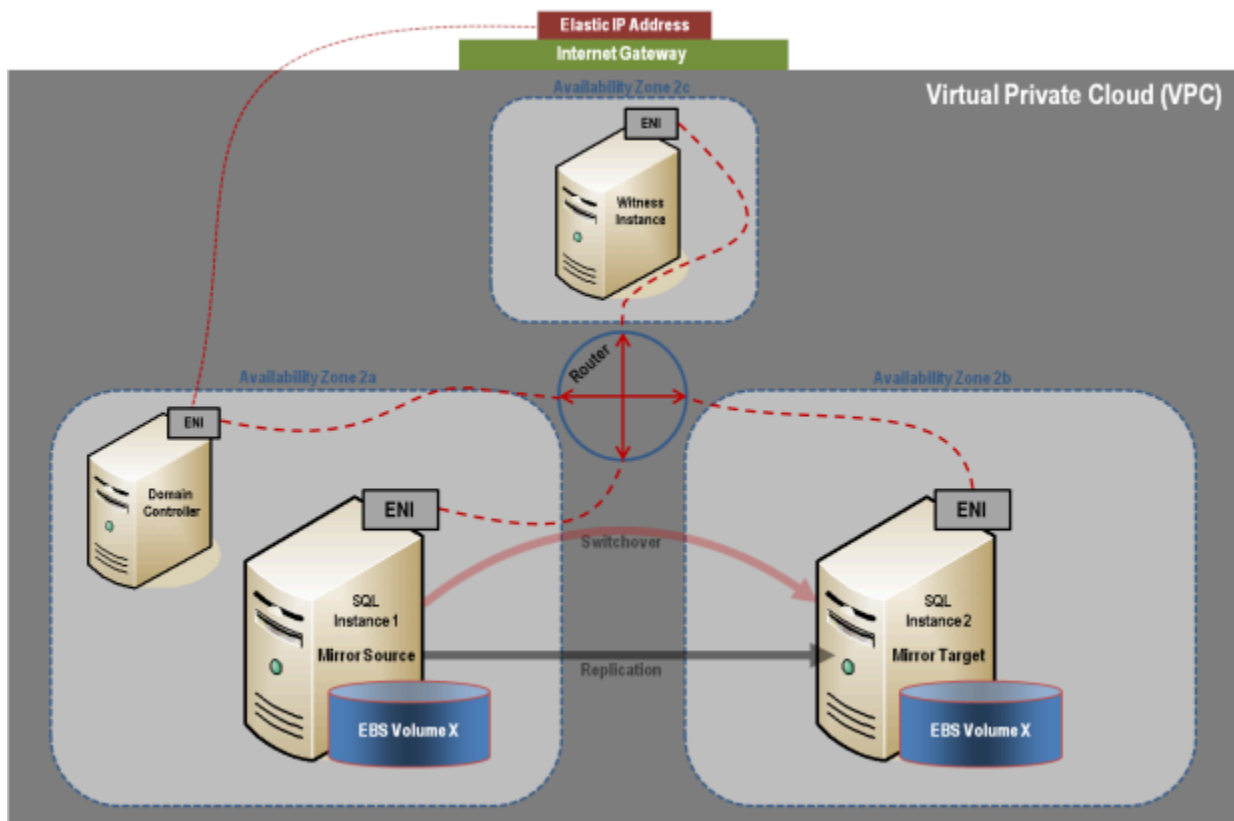


Virtual Network Topology

An elastic IP address, connected to the domain controller, is the entry point into your “virtual lab”. This elastic IP address is associated with the domain controller’s primary elastic network interface (ENI). Elastic

network interfaces are created and assigned to each instance at launch time. It is through these elastic network interfaces that the instances are able to communicate with one another.

The domain controller's elastic network interface is attached to a public subnet, and through a rule that you will create, routes 0.0.0.0/0 (all traffic) to the virtual private cloud's Internet Gateway. You will also create a rule to allow for utilizing Remote Desktop Connection to connect to your instances. You will initially connect to your domain controller through the elastic IP via Remote Desktop Connection. Once connected to your domain controller, a router allows you to Remote Desktop into your other instances.



Set Up Routing and Security

You will set up routing and security to control the flow of traffic in and out of the availability zones.

Set Up Route Tables

Each subnet in your virtual private cloud must be associated with a route table to determine how the traffic between availability zones flows.

1. From the left navigation pane of the **VPC Dashboard**, select **Route Tables**.
2. Select **Create Route Table**.
3. Select your **virtual private cloud** and click **Yes, Create**.
4. Select your new **route table**.
5. Under the **Routes** tab in the bottom pane, the first row is the local route. This enables communication within the virtual private cloud. Associate your **Internet Gateway** with **0.0.0.0/0**. This will appear in the second row and provides access into the virtual private cloud (0.0.0.0/0). This subnet is referred to as public because all traffic from the subnet goes to the **Internet Gateway**.



Destination	Target	Status	Propagated	Actions
15.0.0.0/16	local	active	No	Remove
0.0.0.0/0	igw-4fd7f526	active	No	Remove
<input type="text"/>	<input type="text" value="select a target"/>			Add

To associate your Internet Gateway with 0.0.0.0/0:

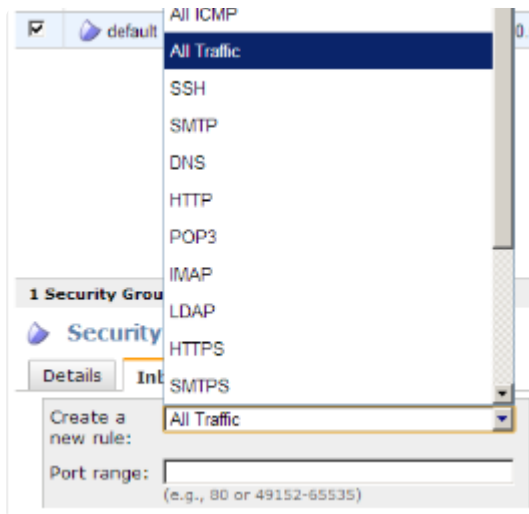
- a. In the box under **Destination**, enter **0.0.0.0/0**, then under **Target**, select your **Internet Gateway** and click **Add**.
- b. A dialog will appear asking for confirmation. Select **Yes, Create**.
- c. Under the **Associations** tab, click a subnet, then select **Associate**.
- d. A dialog will appear asking for confirmation. Select **Yes, Associate**.
- e. Repeat this for all subnets.

Create Network Security Group

In order to control inbound traffic as well as traffic between availability zones within the virtual private cloud, a security group should be set up.

1. From the left navigation pane of the **VPC Dashboard**, select **Security Groups**.
2. Click **Create Security Group**.
3. Enter **Name** and **Description**, then select your **virtual private cloud** and click **Yes, Create**.

4. Select your **Security Group**.
5. Under the **Inbound** tab, select **All Traffic** from the **Create a new rule** dropdown.



6. Enter your **private IP address** as **Source** and select **Add Rule**.
7. To enable a **Remote Desktop Connection**, select **RDP** from the **Create a new rule** dropdown.
8. Enter **3389** for the **Port** and **0.0.0.0/0** as **Source** and select **Add Rule**.
9. Under the **Outbound** tab, select **All Traffic** from the **Create a new rule** dropdown.
10. Enter **0.0.0.0/0** for the **Destination** and select **Add Rule**.

✿ **IMPORTANT:** Make sure you click **Apply Rule Changes** or your changes will not be saved.

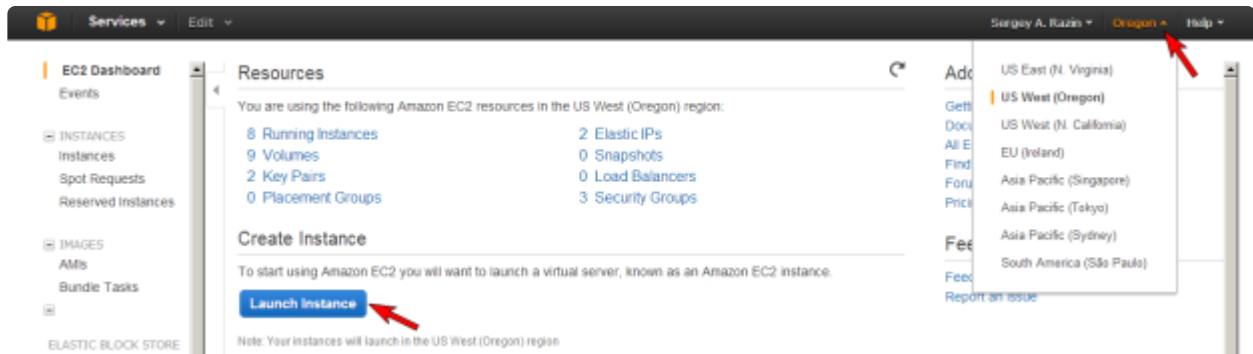
✿ **HELPFUL LINKS:**
[Amazon EC2 Security Groups](#)
[Create and Configure Your Amazon EC2 Security Group](#)

Launch an Instance

The following walks you through launching an instance into your subnet. You will want to launch two instances into one availability zone, one for your domain controller instance and one for your SQL instance.

Then you will launch another SQL instance into another availability zone and a quorum witness instance into yet another availability zone.

1. Using the email address and password specified when signing up for Amazon Web Service (AWS), sign in to the [Amazon EC2 Console](#).
2. From the top right of the navigation bar, select the region for your instance from the dropdown selection.



3. Select **Instances** from the left navigation pane, then click the **Launch Instance** button.
4. Select the **Classic Wizard** from the **Create a New Instance** dialog and select **Continue**.
5. Choose an AMI.
 - Select the **Microsoft Windows Server 2008 R2 Base AMI (2008 R2 SP1 Datacenter Edition)**.
6. Configure instance details.
 - a. Select your **Instance Type**. (Note: Select M1 Small or larger.)



HELPFUL LINKS:

[Amazon EC2 Instances](#)
[Available Instance Types](#)

- b. Under **Launch Instances**, select **Launch into: VPC**, then select your **availability zone**. Click **Continue**.
- c. Accept defaults on the **Advanced Instance Options** and the **Storage Device**

Configuration dialogs by clicking **Continue**.

d. Add a **Tag** to name your instance and select **Continue**.

7. Create **Key Pair**.

a. Unless choosing an existing Key Pair, you'll select **Create a New Key Pair**.

b. Enter a name and then select the **Create & Download your Key Pair** box.

c. Save the **Key Pair** file in a place you'll remember. **Note:** You can use this key pair to launch other instances in the future or visit the **Key Pairs** page to create or manage existing ones.

8. Configure **firewall** (refer to [Firewall Configurations](#) for further information).

- Choose a **Security Group**, then select **Continue**.

9. A review page will display. Select **Launch**. Click **Close** to close the confirmation page and return to the **Amazon Management Console**.

Get Windows Admin Password

You'll need an administrator password to connect to your instance with Remote Desktop. **Note:** You'll need the private key file that you created when you launched your instance.

1. Click **Instances** in the left navigation pane to view the status of your new instance. The status should be **pending** while it is launching, but status will change to **running**.
2. Select your new instance.
3. From the **Actions** dropdown menu, select **Get Windows Admin Password**. **Note:** It may take a few minutes before you are able to retrieve your password.

Actions ▼

Instance Management

- Connect
- Get System Log
- Get Windows Admin Password
- Create Image (EBS AMI)
- Add/Edit Tags
- Change Security Group



IMPORTANT: Make a note of this initial administrator password. It will be needed to log on to your instance.

Repeat the above steps for all instances.

Assign a Virtual Private Cloud Elastic IP Address to the Domain Controller Instance

For an instance in your virtual private cloud to be reachable from the Internet, it must have a virtual private cloud elastic IP (EIP) address assigned to it. This is your entry point into your “virtual lab.”

1. From the left navigation pane of the **VPC Dashboard**, select **Elastic IPs**.
2. Click **Allocate New Address**.
3. From the **EIP used in:** list, select **VPC** and click **Yes, Allocate**.
4. Select the **new IP address** from the list and click **Associate Address**.
5. In the **Associate Address** dialog box, select the **domain controller instance** to associate the address with and click **Yes, Associate**.

Associate Address Cancel

Select the instance or network interface to which you wish to associate this IP address (50.112.177.125).

Instance: i-d62456e4 - Doc_DC

Private IP address: 15.0.5.155*
* denotes the primary private IP address

or

Network Interface: Select a network interface

Private IP address:
* denotes the primary private IP address

☐ Allow Reassociation

Cancel Yes, Associate

Connect to Instances

Once you've retrieved your initial administrator password and set up your Remote Desktop Connection (RDP) "Rule", you can connect to your domain controller instance via Remote Desktop Connection. Once you are connected to your domain controller instance, you can Remote Desktop into your other instances from there.

1. Open a **Remote Desktop Connection** and enter the **elastic IP address** of your **domain controller instance**.
2. Enter your **administrator password**.

✿ **BEST PRACTICE:** Once logged on, it is best practice to change your password.

Create the Domain Controller Instance

Now that the instances have been created, we started with setting up the Domain Service instance.

This guide is not a tutorial on how to setup an Active Domain Service. We recommend reading [articles](#) on the web on how to setup an Active Directory Service specifically needed in an AWS cloud.

It is very important to understand that even though the instance is running in an AWS cloud, this is a regular installation of Active Directory.

Static IP Addresses

Configure Static IP Addresses for your Instances

1. Connect to your domain controller instance.
2. Click **Start / Control Panel**.
3. Click **Network and Sharing Center**.
4. Select your network interface.
5. Click **Properties**.
6. Click **Internet Protocol Version 4 (TCP/IPv4)**, then **Properties**.
7. Obtain your current **IPv4 address**, **default gateway** and **DNS server** for the network interface from **Amazon**.
8. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, under **Use the following IP address**, enter your **IPv4 address**.
9. In the **Subnet mask** box, type the subnet mask associated with your virtual private cloud subnet.
10. In the **Default Gateway** box, type the **IP address** of the default gateway and then click **OK**.
11. For the **Preferred DNS Server**, enter the **Primary IP Address of Your Domain Controller** (ex. 15.0.1.72).
12. Click **Okay**, then select **Close**. Exit **Network and Sharing Center**.
13. Repeat the above steps on your other instances.

Join the Two SQL Instances and the Witness Instance to Domain

1. On each instance, click **Start**, then right-click **Computer** and select **Properties**.
2. On the far right, select **Change Settings**.

3. Click on **Change**.
4. Enter **Computer Name**.
5. Select **Domain**.
6. Enter **Domain Name** – (ex. docs.aws.com).

Assign Secondary Private IPs to the Two SQL Instances

In addition to the Primary IP, you will need to add three additional IPs (Secondary IPs) to the elastic network interface for each SQL instance.

1. From the left navigation pane of the **EC2 Dashboard**, select **Instances**.
2. Right-click the instance for which you want to add a secondary private IP address for.
3. Select **Manage Private IP Addresses**.
4. Select **Assign a secondary private address** and enter an IP address that is within the subnet range for the instance (ex. For 15.0.1.25, enter 15.0.1.26). Repeat to add two additional IP addresses.
5. Select **Yes, Update**.
6. Select **Close**.
7. Perform the above on **both SQL Instances**.



HELPFUL LINKS:

[Configuring a Secondary Private IP Address for Your Windows Instance](#)
[Private IP Addresses Per ENI Per Instance Type](#)
[Multiple IP Addresses](#)

Create and Attach Volumes

DataKeeper is a block-level volume replication solution and requires that each node in the cluster have additional volume(s) (other than the system drive) that are the same size and same drive letters. Please review [Volume Considerations](#) for additional information regarding storage requirements.

Create Volumes

Create two volumes in each availability zone for each SQL instance.

1. From the left navigation pane of the **EC2 Dashboard**, select **Instances** to display your instances.
2. Select your instance. In the **Description** tab in the bottom pane, note the **Zone** for the instance.
3. From the left navigation pane, select **Volumes** under **Elastic Block Store** (EBS).

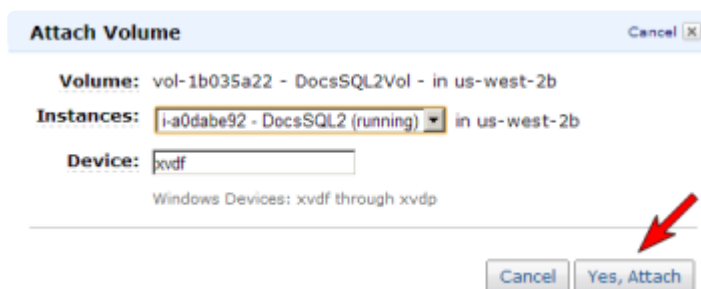
The console displays a list of current elastic block store volumes in that region. (You should see the elastic block store volume that serves as the root device volume for your instance.)

4. Click **Create Volume**.
5. In the **Create Volume** dialog box, select the **Standard** volume type, enter the desired size and select the correct zone (noted above), then click **Yes, Create**.

Attach Volumes

Once your volumes are created, attach them to your two SQL instances.

1. From the left navigation pane of the **EC2 Dashboard**, select **Volumes**.
2. Right-click on a volume, then select **Attach Volume**.
3. Select the **instance** to attach the volume to, then select **Yes, Attach**. You will see the status go from **available** to **in-use**.



HELPFUL LINKS:

[Creating an Amazon EBS Volume](#)

[Making an Amazon EBS Volume Available for Use / Make the Volume Available on Windows](#)

Configure the Cluster

Prior to installing DataKeeper Cluster Edition, it is important to have Windows Server configured as a cluster using either a node majority quorum (if there is an odd number of nodes) or a node and file share majority quorum (if there is an even number of nodes). Consult the Microsoft documentation on clustering in addition to this topic for step-by-step instructions. Microsoft released a [hotfix](#) for Windows 2008R2 that allows disabling of a node's vote which may help achieve a higher level of availability in certain multi-site cluster configurations.

Add Failover Clustering

Add the Failover Clustering feature to both SQL instances.

1. Launch **Server Manager**.
2. Select **Features** in the left pane and click **Add Features** in the **Features** pane. This starts the **Add Features Wizard**.
3. Select **Failover Clustering**.
4. Select **Install**.

Validate a Configuration

1. Open **Failover Cluster Manager**.
2. Click on **Validate a Configuration**.
3. Click **Next**, then add your two **SQL instances**.

Note: To search, select **Browse**, then click on **Advanced** and **Find Now**. This will list available instances.

4. Click **Next**.
5. Select **Run Only Tests I Select** and click **Next**.

6. In the **Test Selection** screen, deselect **Storage** and click **Next**.
7. At the resulting confirmation screen, click **Next**.
8. Review **Validation Summary Report**, then click **Finish**.

Create Cluster

1. In **Failover Cluster Manager**, click on **Create a Cluster** then click **Next**.
2. Enter your two **SQL instances**.
3. On the **Validation Warning** page, select **No** then click **Next**.
4. On the **Access Point for Administering the Cluster** page, enter a unique name for your WSFC Cluster. Then enter the **Failover Clustering IP address** for each node involved in the cluster. This is the first of the three **secondary IP addresses** added previously to each instance. Click **Next**.
5. Click **Next** on the **Confirmation** page.
6. On **Summary page**, review any warnings, then select **Finish**.

Configure Quorum/Witness

1. Create a folder on your quorum/witness instance (witness).
2. Share the folder.
 - a. Right-click folder and select **Share With / Specific People....**
 - b. From the dropdown, select **Everyone** and click **Add**.
 - c. Under **Permission Level**, select **Read/Write**.
 - d. Click **Share**, then **Done**. (Make note of the path of this file share to be used below.)
3. In **Failover Cluster Manager**, right-click cluster and choose **More Actions** and **Configure Cluster Quorum Settings**. Click **Next**.
4. On the **Select Quorum Configuration**, choose **Node and File Share Majority** and click **Next**.

5. On the **Configure File Share Witness** screen, enter the path to the file share previously created and click **Next**.
6. On the **Confirmation** page, click **Next**.
7. On the **Summary** page, click **Finish**.

Install and Configure DataKeeper

After the basic cluster is configured but prior to any cluster resources being created, install and license **DataKeeper Cluster Edition** on all cluster nodes. See the [DataKeeper Cluster Edition Installation Guide](#) for detailed instructions.

1. Run **DataKeeper setup** to install **DataKeeper Cluster Edition** on both SQL instances.
2. Enter your **license key** and reboot when prompted.
3. Launch the **DataKeeper GUI** and **connect to server**.

✿ **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. Refer to [DataKeeper Service Log On ID and Password Selection](#) for additional information.

4. [Create a Job](#).
5. When asked if you would like to auto-register the volume as a cluster volume, select **Yes**.

✿ **Note:** If installing DataKeeper Cluster Edition on Windows “Core” (GUI-less Windows), make sure to read [Installing and Using DataKeeper on Windows 2008R2/2012 Server Core Platforms](#) for detailed instructions.

Configure MSDTC

1. For Windows Server 2012 and 2016, in the **Failover Cluster Manager GUI**, select **Roles**, then select **Configure Role**.
2. Select **Distributed Transaction Coordinator (DTC)**, and click **Next**.



For Windows Server 2008, in the **Failover Cluster Manager GUI**, select **Services and Applications**, then select **Configure a Service or Application** and click **Next**.

3. On the **Client Access Point** screen, enter a name, then enter the **MSDTC IP address** for each node involved in the cluster. This is the second of the three **secondary IP addresses** added previously to each instance. Click **Next**.
4. Select the **MSDTC volume** and click **Next**.
5. On the **Confirmation** page, click **Next**.
6. Once the **Summary** page displays, click **Finish**.

Install SQL on the First SQL Instance

1. Via **Map Network Drive**, map the **SQL IMG file** from the **domain controller** to the **two SQL servers**.
2. Mount the IMG file (via **MagicDisc** or other virtual CD tool).
3. Once the IMG file is mounted, **launch** the **SQL setup**. To launch the SQL setup, open a **Command** window, browse to your **SQL install directory** and type the following command:

```
F:\>Setup /SkipRules=Cluster_VerifyForErrors /Action=InstallFailoverCluster
```

4. On **Setup Support Rules**, click **OK**.
5. On the **Product Key** dialog, enter your **product key** and click **Next**.
6. On the **License Terms** dialog, accept the **license agreement** and click **Next**.
7. On the **Product Updates** dialog, click **Next**.
8. On the **Setup Support Files** dialog, click **Install**.
9. On the **Setup Support Rules** dialog, you will receive a warning. Click **Next**, ignoring this message, since it is expected in a multi-site or non-shared storage cluster.
10. Verify **Cluster Node Configuration** and click **Next**.

11. Configure your **Cluster Network** by adding the “third” secondary IP address for your SQL instance and click **Next**. Click **Yes** to proceed with multi-subnet configuration.
12. Enter **passwords** for service accounts and click **Next**.
13. On the **Error Reporting** dialog, click **Next**.
14. On the **Add Node Rules** dialog, skipped operation warnings can be ignored. Click **Next**.
15. Verify features and click **Install**.
16. Click **Close** to complete the installation process.

Install SQL on the Second SQL Instance

Installing the second SQL instance is similar to the first one.

1. Mount the IMG file (again, via **MagicDisc** or other virtual CD tool).
2. Once IMG file is mounted, run **SQL setup** once again from the command line in order to skip the **Validate** process. Open a **Command** window, browse to your **SQL install directory** and type the following command:

```
Setup /SkipRules=Cluster_VerifyForErrors /Action=AddNode  
/INSTANCENAME="MSSQLSERVER"
```

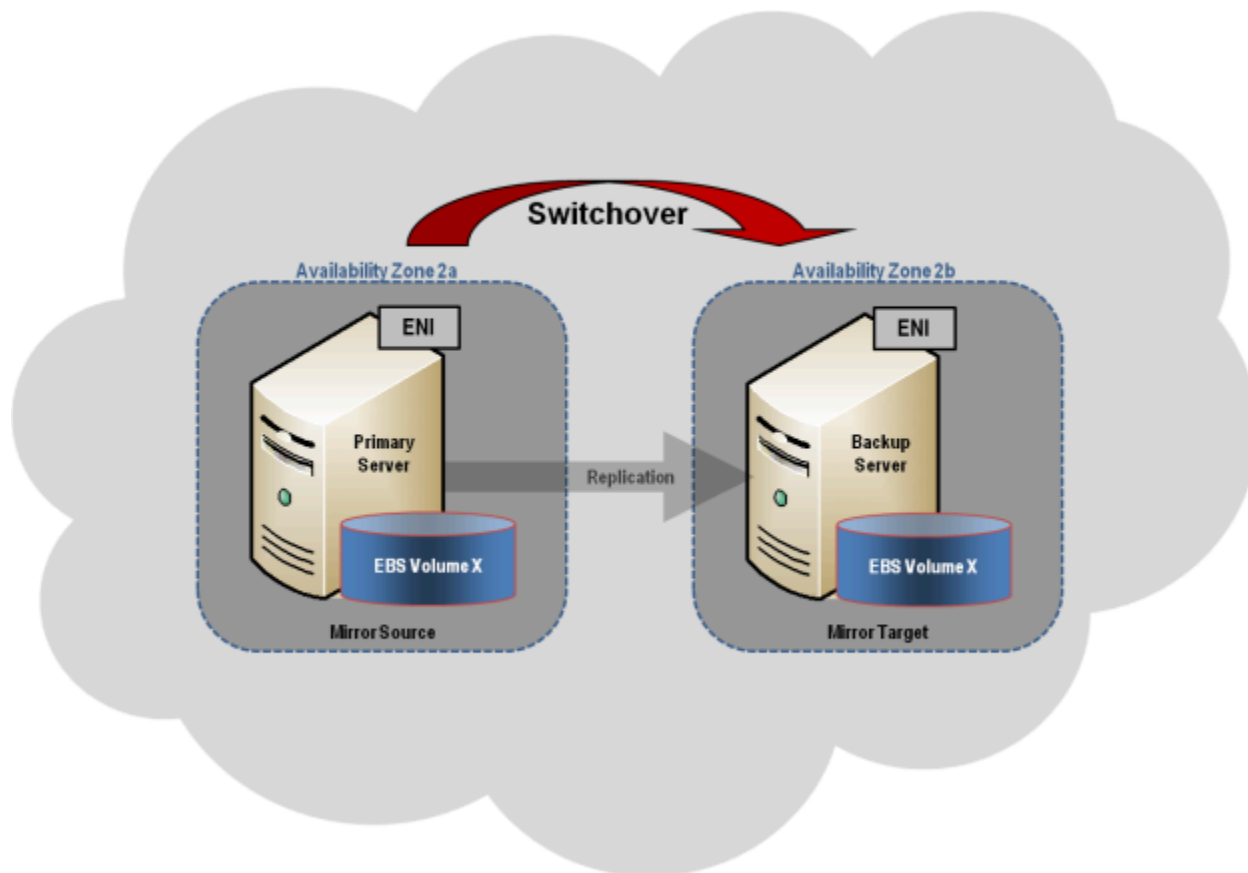
(**Note:** This assumes you installed the default instance on the first node)

3. On **Setup Support Rules**, click **OK**.
4. On the **Product Key** dialog, enter your **product key** and click **Next**.
5. On the **License Terms** dialog, accept the **license agreement** and click **Next**.
6. On the **Product Updates** dialog, click **Next**.
7. On the **Setup Support Files** dialog, click **Install**.
8. On the **Setup Support Rules** dialog, you will receive a warning. Click **Next**, ignoring this message, since it is expected in a multi-site or non-shared storage cluster.

9. Verify **Cluster Node Configuration** and click **Next**.
10. Configure your **Cluster Network** by adding the “third” secondary IP address for your SQL Instance and click **Next**. Click **Yes** to proceed with multi-subnet configuration.
11. Enter **passwords** for service accounts and click **Next**.
12. On the **Error Reporting** dialog, click **Next**.
13. On the **Add Node Rules** dialog, skipped operation warnings can be ignored. Click **Next**.
14. Verify features and click **Install**.
15. Click **Close** to complete the installation process.

Common Cluster Configuration

This section describes a **common 2-node replicated cluster configuration**.



1. The initial configuration must be done from the [DataKeeper UI](#) running on one of the cluster nodes. If it is not possible to run the DataKeeper UI on a cluster node, such as when running DataKeeper on a Windows Core only server, install the DataKeeper UI on any computer running Windows XP or higher and follow the instruction in the [Core Only](#) section for creating a mirror and registering the cluster resources via the command line.
2. Once the DataKeeper UI is running, [connect to each of the nodes](#) in the cluster.
3. [Create a Job](#) using the DataKeeper UI. This process creates a mirror and adds the DataKeeper Volume resource to the Available Storage.



IMPORTANT: Make sure that **Virtual Network Names for NIC connections** are identical on all cluster nodes.

4. If additional mirrors are required, you can [Add a Mirror to a Job](#).
5. With the **DataKeeper Volume(s)** now in **Available Storage**, you are able to create cluster resources (SQL, File Server, etc.) in the same way as if there were a shared disk resource in the cluster. Refer to Microsoft documentation for additional information in addition to the above for step-by-step cluster configuration instructions.

Management

Once a DataKeeper volume is registered with Windows Server Failover Clustering, all of the management of that volume will be done through the Windows Server Failover Clustering interface. All of the management functions normally available in DataKeeper [will be disabled](#) on any volume that is under cluster control. Instead, the DataKeeper Volume cluster resource will control the mirror direction, so when a DataKeeper Volume comes online on a node, that node becomes the source of the mirror. The properties of the DataKeeper Volume cluster resource also display basic mirroring information such as the source, target, type and state of the mirror.

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](#)

Additional Resources:

Step-by-Step: Configuring a 2-Node Multi-Site Cluster on Windows Server 2008 R2 – Part 1 —

<http://clusteringformeremortals.com/2009/09/15/step-by-step-configuring-a-2-node-multi-site-cluster-on-windows-server-2008-r2-%E2%80%93-part-1/>

Step-by-Step: Configuring a 2-Node Multi-Site Cluster on Windows Server 2008 R2 – Part 3 —

<http://clusteringformeremortals.com/2009/10/07/step-by-step-configuring-a-2-node-multi-site-cluster-on-windows-server-2008-r2-%E2%80%93-part-3/>

Deploying DataKeeper Cluster Edition in Azure

DEPLOYING MICROSOFT SQL SERVER 2014 FAILOVER CLUSTERS IN AZURE RESOURCE MANAGER (ARM)

Before beginning, make sure you read the [DataKeeper Cluster Edition Release Notes](#) for the latest information. It is highly recommended that you read and understand the [DataKeeper Cluster Edition Installation Guide](#).

Steps required to deploy a 2-node SQL Server Failover Cluster in a single region using Azure Resource Manager

Note: This guide does not apply to the Azure Classic portal.

DataKeeper Cluster Edition allows you to take locally attached storage, whether it uses Premium or Standard Disks, and replicate those disks synchronously, asynchronous, or a mix of both, between two or more cluster nodes. In addition, a DataKeeper Volume resource is registered in Windows Server Failover Cluster which takes the place of a Physical Disk resource. Instead of controlling SCSI-3 reservations like a Physical Disk Resource, the DataKeeper Volume controls the mirror direction, ensuring the active node is always the source of the mirror. For SQL Server and Failover Cluster the DataKeeper volume is similar to a Physical Disk and is used the same way Physical Disk Resources would be used.

PRE-REQUISITES

- You have used the Azure Portal (<http://portal.azure.com>) before and are comfortable deploying virtual machines in Azure IaaS
- You have obtained a full license or [eval license of SIOS DataKeeper](#)

THE EASY WAY TO DO A PROOF-OF-CONCEPT

The Azure Resource Manager has the ability to use Deployment Templates to rapidly deploy applications consisting of interrelated Azure resources. Many of these templates are developed by Microsoft and are readily available in their community on Github as [Quickstart Templates](#). Community members are also free to extend templates or to publish their own templates on GitHub. One such template entitled “SQL Server

2014 AlwaysOn Failover Cluster Instance with SIOS DataKeeper Azure Deployment Template” published by SIOS Technology completely automates the process of deploying a 2-node SQL Server FCI into a new Active Directory Domain.

To deploy this template click on **Deploy to Azure** in the template.



Refer to <https://github.com/SIOSDataKeeper/SIOSDataKeeper-SQL-Cluster> to rapidly provision a 2-node SQL cluster.

DEPLOYING A SQL SERVER FAILOVER CLUSTER INSTANCE USING THE AZURE PORTAL

While the automated Azure deployment template is a quick and easy way to get a 2-node SQL Server FCI up and running, there are some limitations. For one, it uses a 180 Day evaluation version of SQL Server, so it cannot be used in production unless you [upgrade the SQL eval licenses](#). Also, it builds an entirely new AD domain so if you plan to integrate it with your existing domain it will have to be rebuilt manually.

PROVISIONING THE DOMAIN CONTROLLER (DC1)

To build a 2-node SQL Server Failover Cluster Instance in Azure, you will need a basic Virtual Network based on Azure Resource Manager (not Azure Classic) and at least one virtual machine up and running configured as a Domain Controller. This guide does not cover these steps. We will refer to the domain controller as DC1 for the rest of this guide. When creating DC1 you may choose either Windows Server 2008R2 or Windows Server 2012R2. The only other requirements for DC1 are that it be the same type (Premium or Standard) as the cluster nodes, SQL1 and SQL2, and be in the same Availability Set. Once the Virtual Network and Domain Controller are configured, you will provision two more virtual machines which will act as the two nodes in the cluster.

Example:

DC1 – Our Domain Controller and File Share Witness

SQL1 and SQL2 – The two nodes of our SQL Server Cluster

PROVISIONING THE TWO CLUSTER NODES (SQL1 AND

SQL2)

Using the Azure Portal, provision both SQL1 and SQL2 exactly the same way. There are numerous options to choose from including instance size, storage options, etc. This guide is not meant to be an exhaustive guide to deploying SQL Server in Azure. There are a few key things to keep in mind when creating your instances, especially in a clustered environment.

Availability Set – It is important that both SQL1, SQL2, and DC1 reside in the same Availability Set. By putting them in the same Availability Set we are ensuring that each cluster node and the File Share Witness reside in different Fault and Update Domains. This helps guarantee that during both planned maintenance and unplanned maintenance the cluster will continue to be able to maintain quorum and avoid downtime.

Create virtual machine

Settings

1 Basics Done ✓

2 Size Done ✓

3 Settings Configure optional features >

4 Summary Windows Server 2012 R2 Data... >

Storage account ⓘ
(new) test226033 >

Network

* **Virtual network** ⓘ
(new) test22 >

* **Subnet** ⓘ
default (10.1.0.0/24) >

* **Public IP address** ⓘ
(new) Dave >

* **Network security group** ⓘ
(new) Dave >

Extensions

Extensions ⓘ
No extensions >

Monitoring

Diagnostics ⓘ
Disabled Enabled

* **Diagnostics storage account** ⓘ
(new) test226033 >

Availability

* **Availability set** ⓘ
None >

OK

Make sure all of your cluster nodes and your file share witness resides in the same Availability Set. You will initially have to create the Availability Set but then you can add the other servers to the Availability Set once it is created.

STATIC IP ADDRESS

Once each VM is provisioned, change the settings of the IP address to Static so that the IP addresses of the

cluster nodes will not change.

IP addresses
sios-0-nic

Save Discard

Public IP address settings

Public IP address
Disabled Enabled

* IP address
sios0RdpIP (13.68.17.112)

Private IP address settings

Virtual network
ergergergfdswVNET

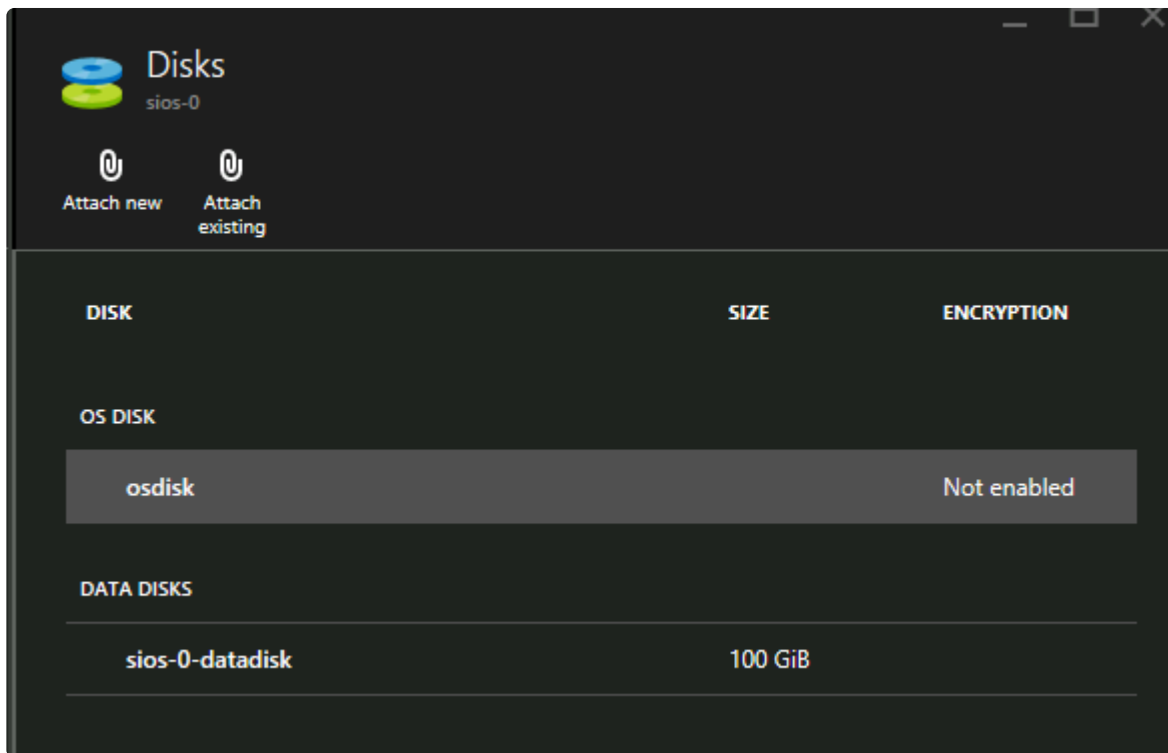
Subnet
staticSubnet (10.0.0.0/24)

Assignment
Dynamic Static

* IP address
10.0.0.5

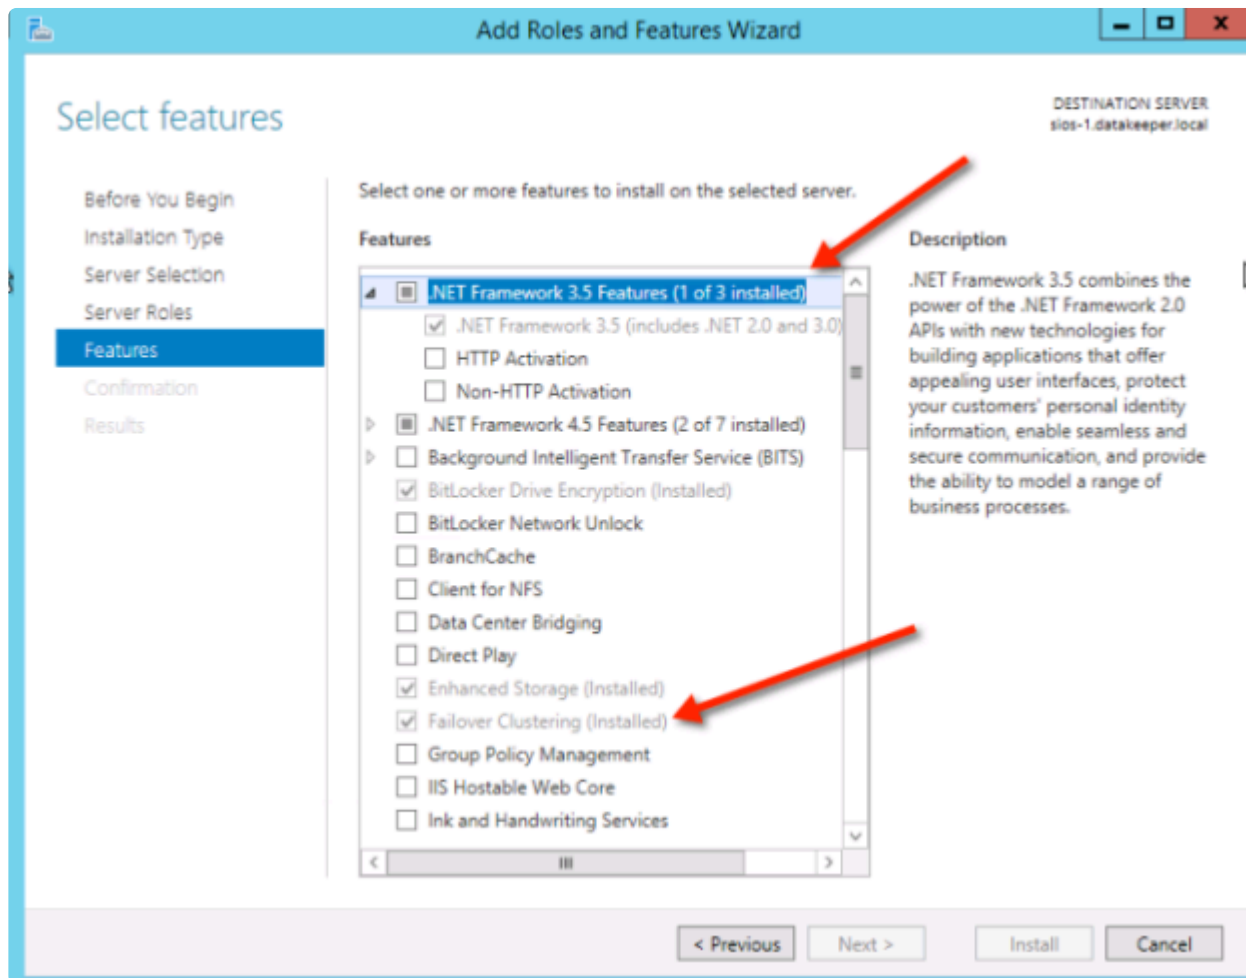
STORAGE

For Storage information refer to [Performance best practices for SQL Server in Azure Virtual Machines](#). At a minimum add at least one additional disk to each of your cluster nodes. DataKeeper can use Premium or Standard disks, but Azure requires you to configure data disks to use the same type as the OS disk. If you created VMs that reside on Premium disks, then you must attach Premium data disks as well. DataKeeper is compatible with Storage Pools so you may attach multiple data disks if your chosen VM size allows it.



CREATE THE CLUSTER

After both cluster nodes (SQL1 and SQL2) have been provisioned as described above and added to your existing domain, the cluster can be created. Before creating the cluster, both the .Net 3.5 Framework and Failover Clustering Features must be enabled on both cluster nodes.



Once these features have been enabled, you are ready to build your cluster. The following steps can be performed both via PowerShell and the WSFC GUI. It is recommended that PowerShell be used to create your cluster. **Note:** If the Failover Cluster Manager GUI is used, a duplicate IP address will be issued to the cluster that is not attached.

Azure VMs are required to use DHCP. By specifying a “Static IP” in the Azure portal when the VM was created, something similar to a DHCP reservation was established. It is not exactly a DHCP reservation because a true DHCP reservation removes the IP address from the DHCP pool. Instead, specifying a Static IP in the Azure portal means that if that IP address is still available when the VM requests it, Azure will issue that IP to it. However, if your VM is offline and another host comes online in that same subnet it could be issued that same IP address.

There is another side effect to the way Azure has implemented DHCP. When creating a cluster with the Windows Server Failover Cluster GUI, when a host uses DHCP (which is required), there is not an option to specify a cluster IP address. Instead it relies on DHCP to obtain an address. DHCP will issue a duplicate IP address, usually the same IP address as the host requesting it. The cluster creation will usually complete,

but you may encounter errors and need to run the Windows Server Failover Cluster GUI from a different node in order to get it to run. Once it is running, change the cluster IP address to an address that is not currently in use on the network.

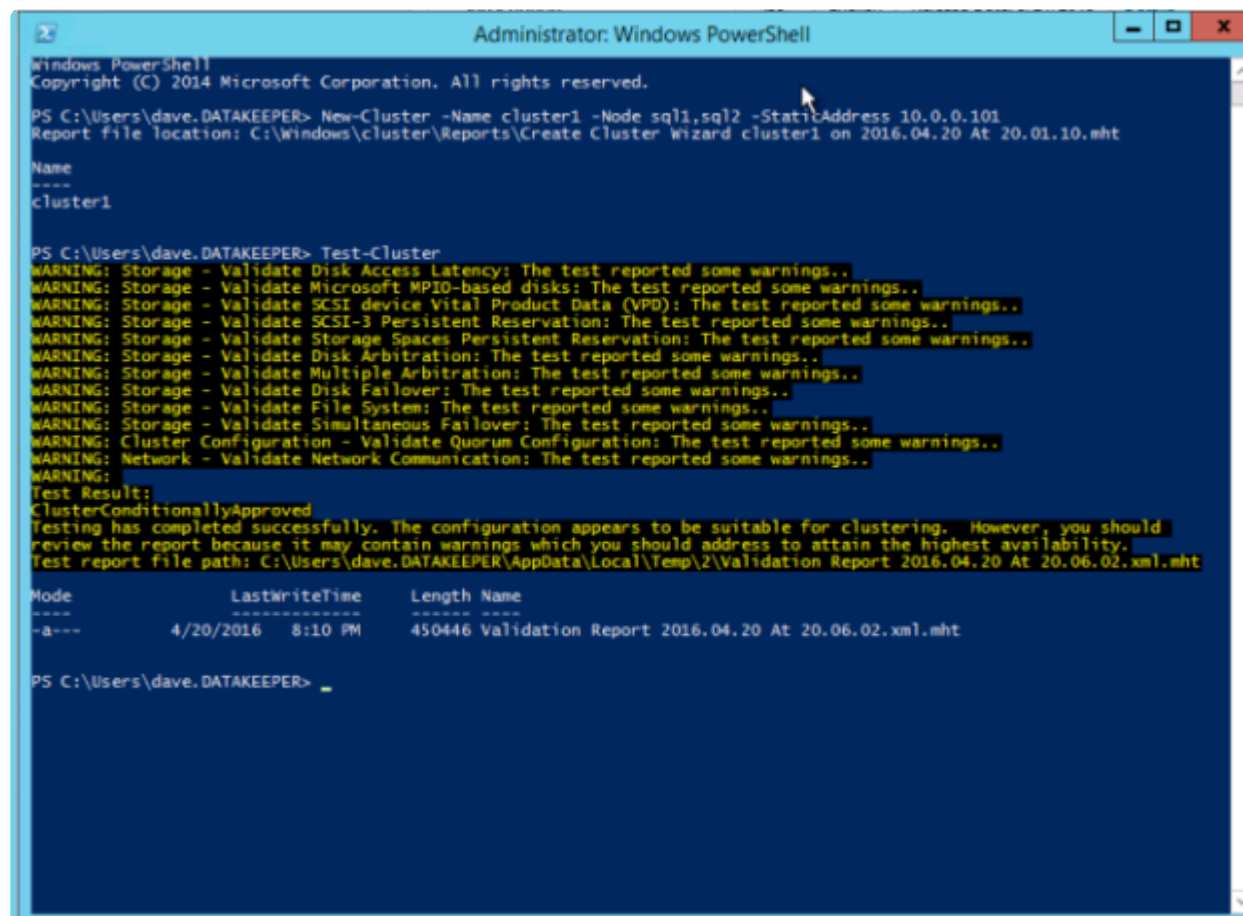
To avoid this scenario, create the cluster via PowerShell by specifying the cluster IP address as part of the PowerShell command.

To create the cluster run the following New-Cluster command:

```
New-Cluster -Name cluster1 -Node sql1,sql2 -StaticAddress 10.0.0.101
```

After the cluster is created, run the following cluster validation:

```
Test-Cluster
```



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2014 Microsoft Corporation. All rights reserved.

PS C:\Users\dave.DATAKEEPER> New-Cluster -Name cluster1 -Node sql1,sql2 -StaticAddress 10.0.0.101
Report file location: C:\Windows\cluster\Reports\Create Cluster Wizard cluster1 on 2016.04.20 At 20.01.10.mht

Name
----
cluster1

PS C:\Users\dave.DATAKEEPER> Test-Cluster
WARNING: Storage - Validate Disk Access Latency: The test reported some warnings..
WARNING: Storage - Validate Microsoft MPIIO-based disks: The test reported some warnings..
WARNING: Storage - Validate SCSI device Vital Product Data (VPD): The test reported some warnings..
WARNING: Storage - Validate SCSI-3 Persistent Reservation: The test reported some warnings..
WARNING: Storage - Validate Storage Spaces Persistent Reservation: The test reported some warnings..
WARNING: Storage - Validate Disk Arbitration: The test reported some warnings..
WARNING: Storage - Validate Multiple Arbitration: The test reported some warnings..
WARNING: Storage - Validate Disk Failover: The test reported some warnings..
WARNING: Storage - Validate File System: The test reported some warnings..
WARNING: Storage - Validate Simultaneous Failover: The test reported some warnings..
WARNING: Cluster Configuration - Validate Quorum Configuration: The test reported some warnings..
WARNING: Network - Validate Network Communication: The test reported some warnings..
WARNING:
Test Result:
ClusterConditionallyApproved
Testing has completed successfully. The configuration appears to be suitable for clustering. However, you should
review the report because it may contain warnings which you should address to attain the highest availability.
Test report file path: C:\Users\dave.DATAKEEPER\AppData\Local\Temp\2\Validation Report 2016.04.20 At 20.06.02.xml.mht

Mode                LastWriteTime         Length Name
----                -
-a---            4/20/2016   8:10 PM         450446 Validation Report 2016.04.20 At 20.06.02.xml.mht

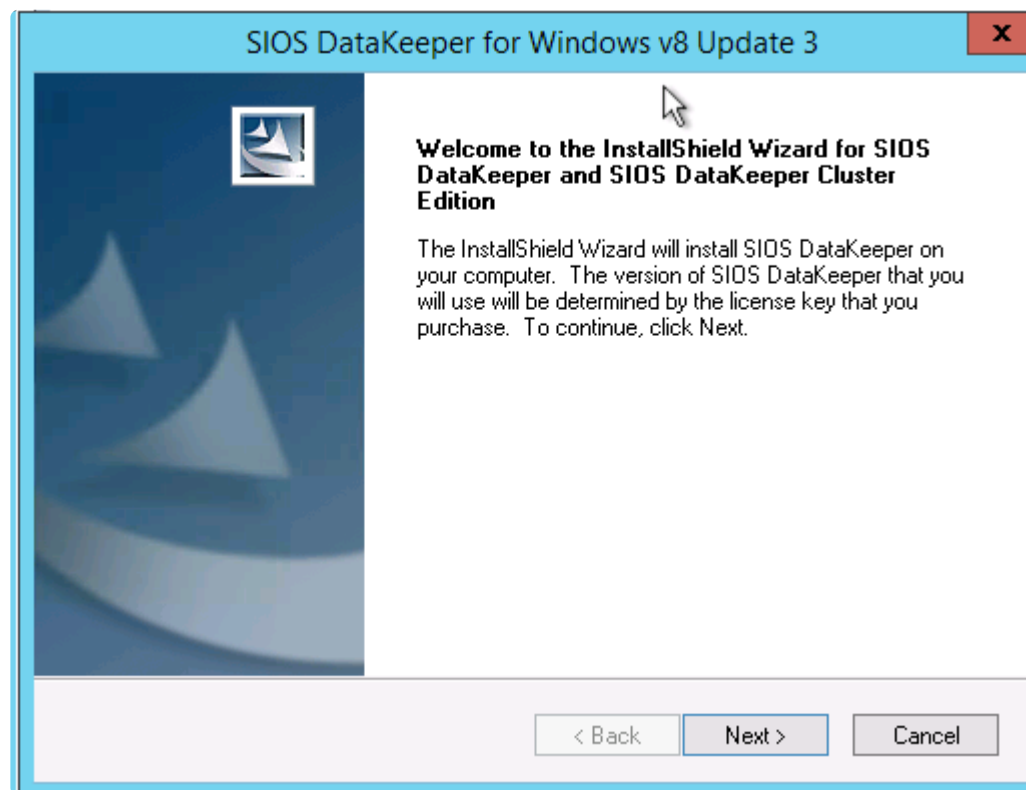
PS C:\Users\dave.DATAKEEPER> _
```

CREATE FILE SHARE WITNESS

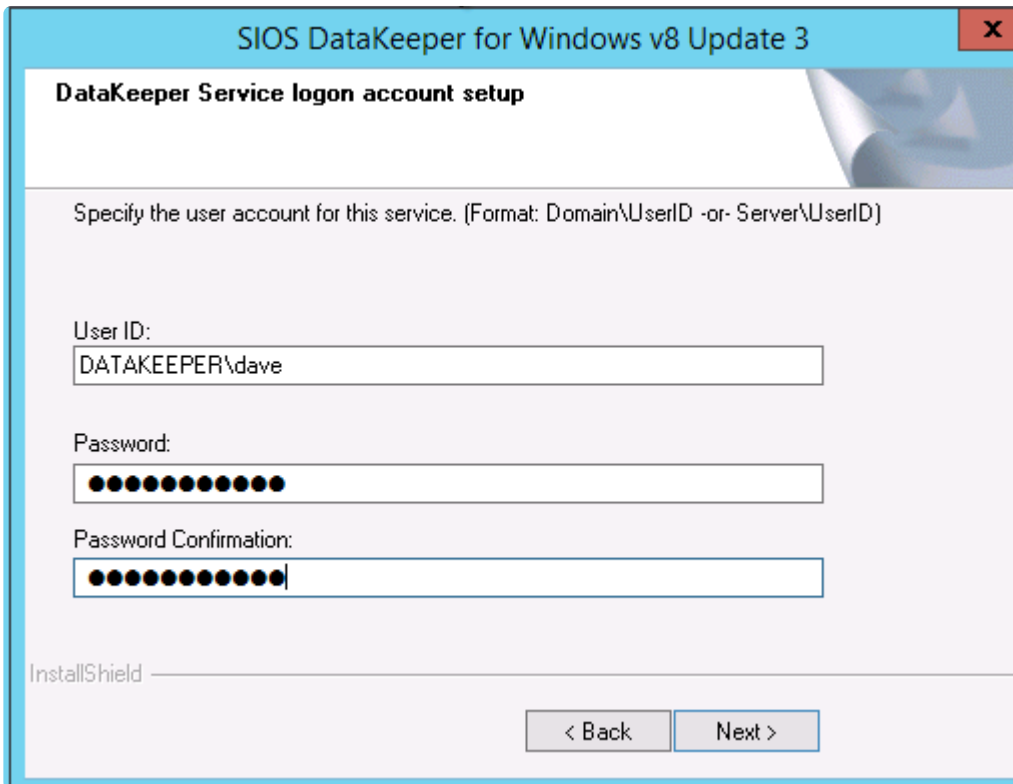
Since there is no shared storage, you must create a file share witness on another server in the same Availability Set as the two cluster nodes. Putting it in the same availability set ensures that you only lose one vote from your quorum at any given time. Refer to <http://www.howtonetworking.com/server/cluster12.htm> for information on how to create a File Share Witness. For this example, the file share witness was put on the domain controller, DC1. For more information on cluster quorums refer to https://blogs.msdn.microsoft.com/microsoft_press/2014/04/28/from-the-mvps-understanding-the-windows-server-failover-cluster-quorum-in-windows-server-2012-r2/.

INSTALL DATAKEEPER

During the installation use all of the default options.



The service account used must be a domain account, and must also be in the Local Administrators group on each node in the cluster.



SIOS DataKeeper for Windows v8 Update 3

DataKeeper Service logon account setup

Specify the user account for this service. (Format: Domain\UserID -or- Server\UserID)

User ID:
DATAKEEPER\dave

Password:
●●●●●●●●●●

Password Confirmation:
●●●●●●●●●●

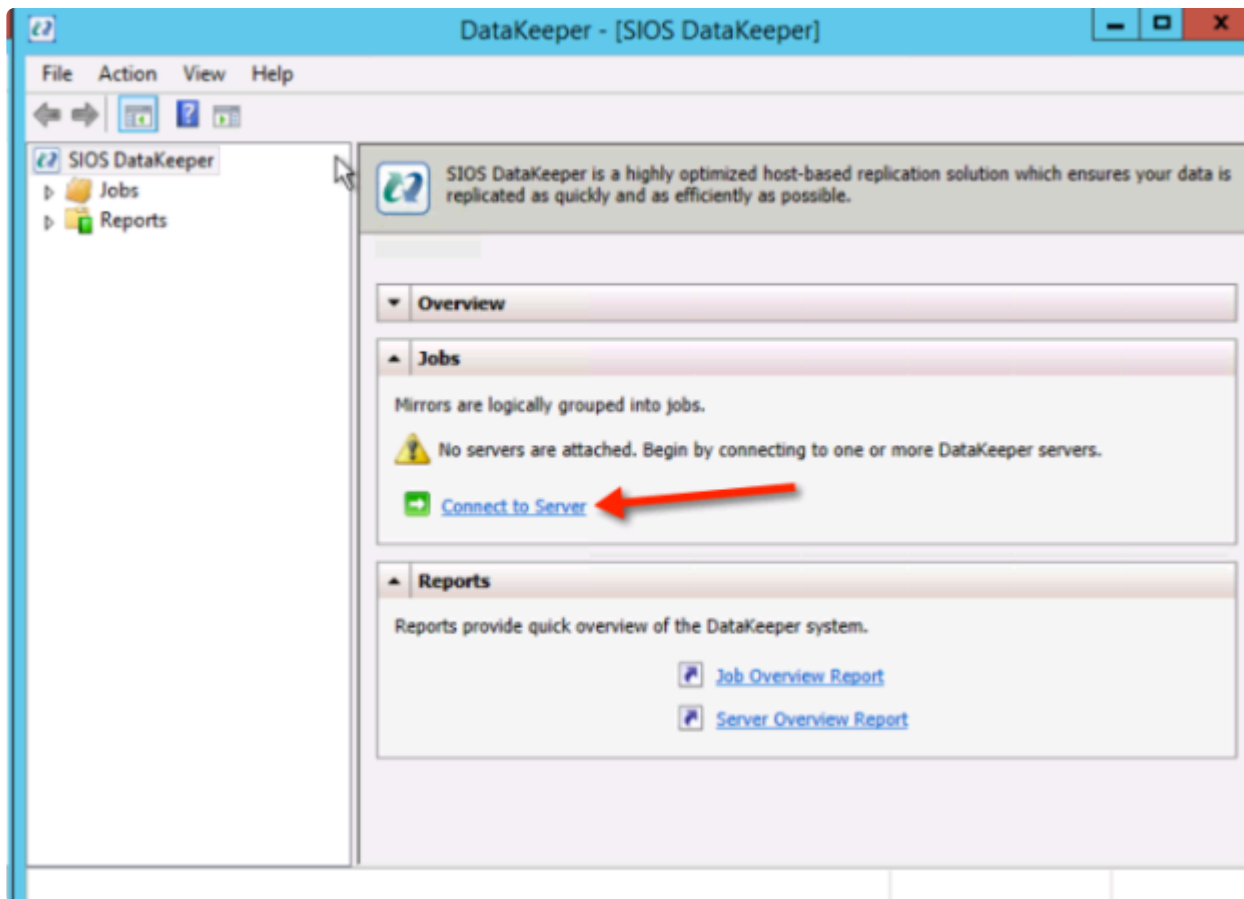
InstallShield

< Back Next >

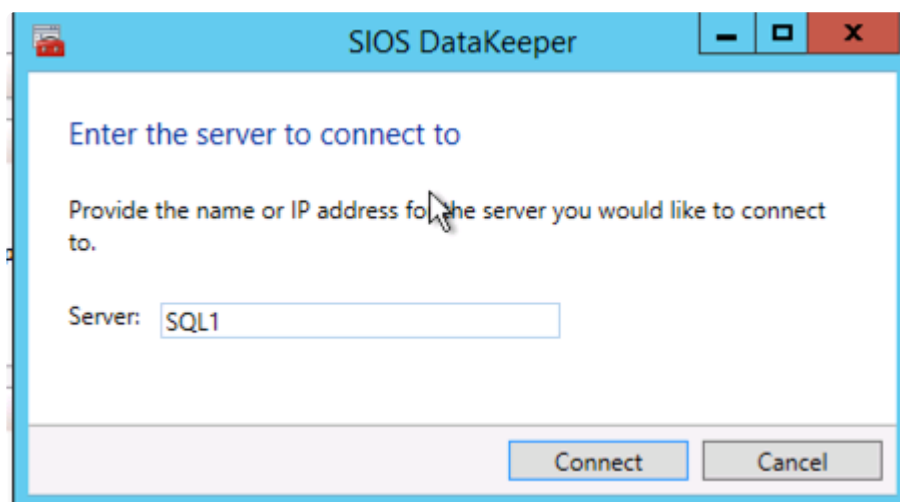
Reboot the servers once DataKeeper is installed and licensed on each node.

CREATE THE DATAKEEPER VOLUME RESOURCE

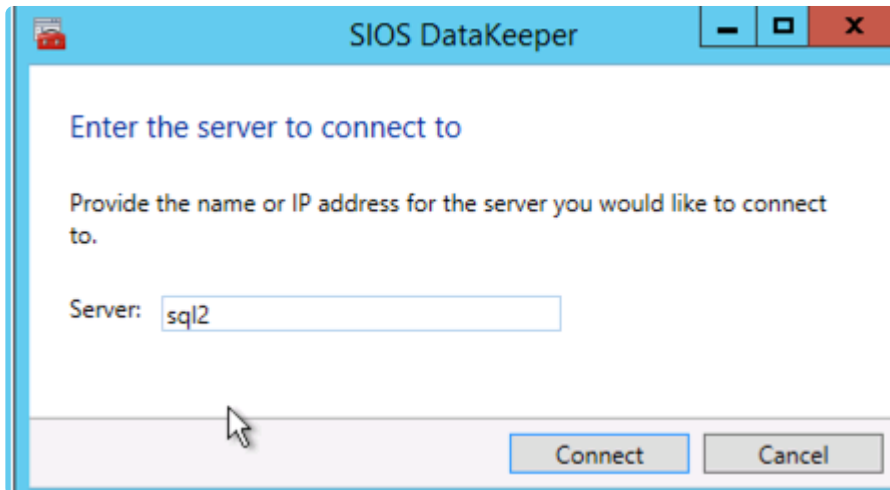
To create the DataKeeper Volume Resource start the DataKeeper UI and connect to both of the servers.



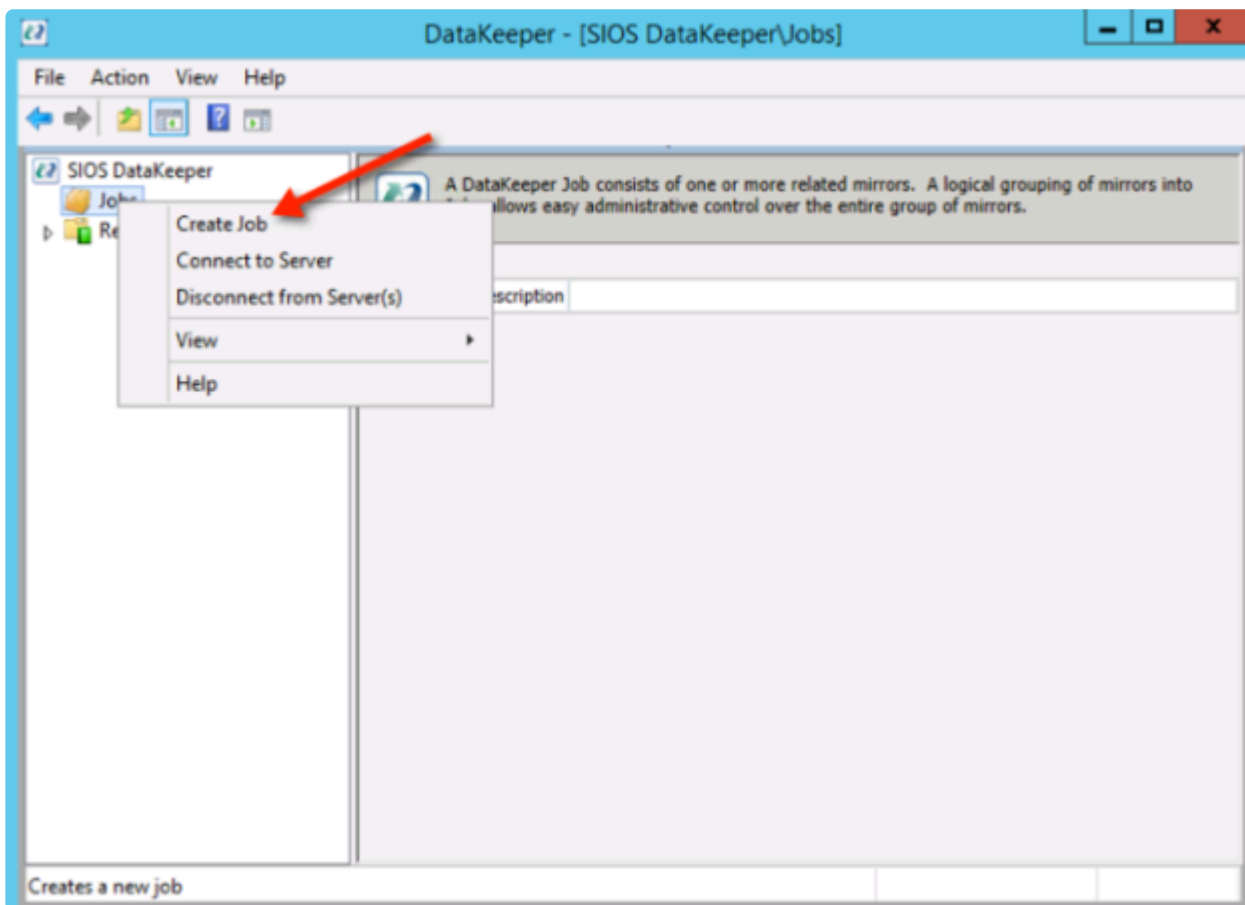
Connect to SQL1



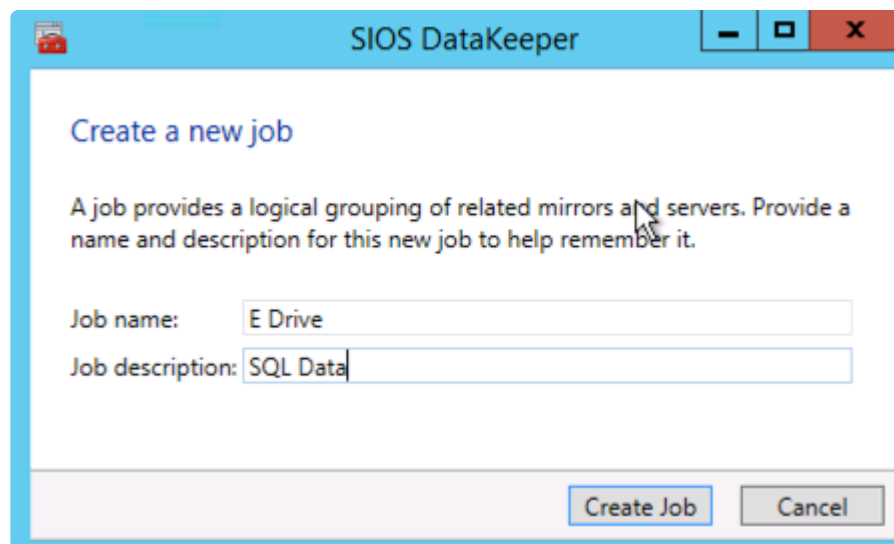
Connect to SQL2



Once you are connected to each server, create your DataKeeper Volume in the Navigation Pane, right click on **Jobs** and choose **Create Job**.

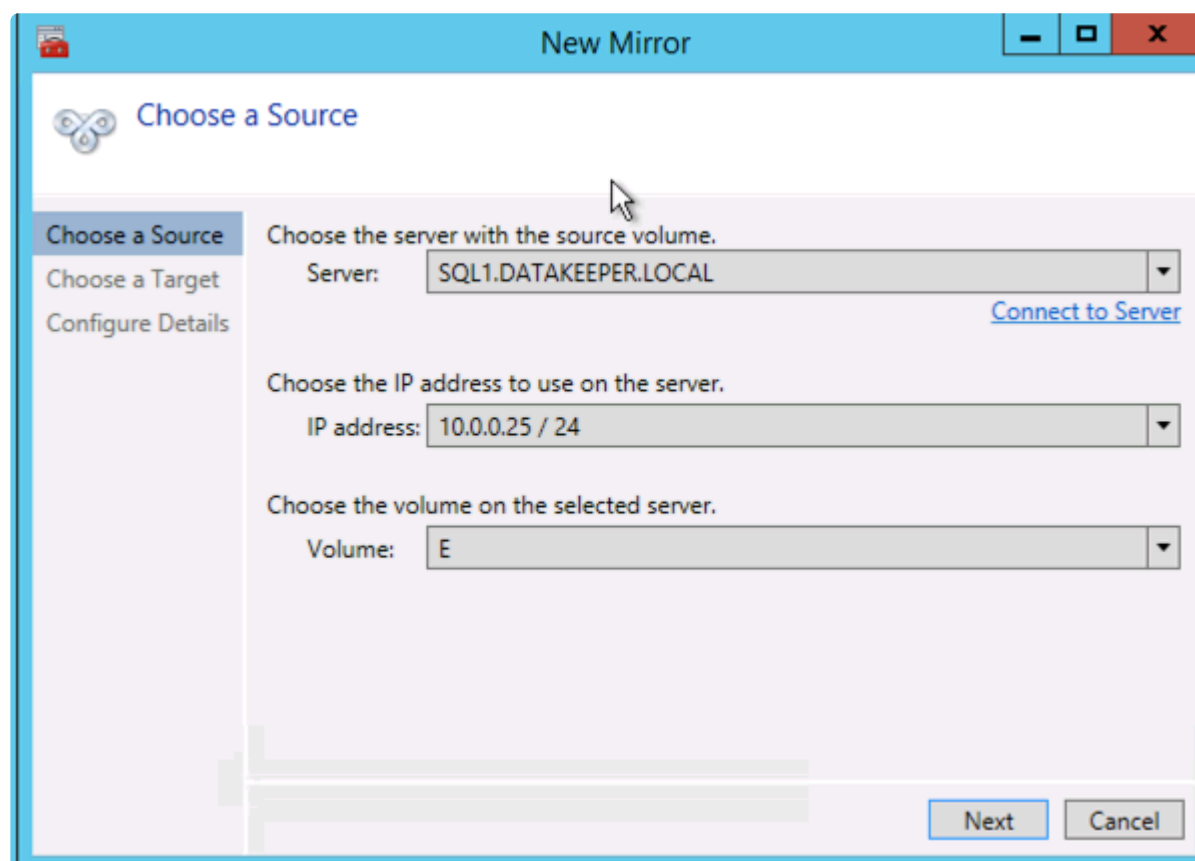


Give the Job a name and description.



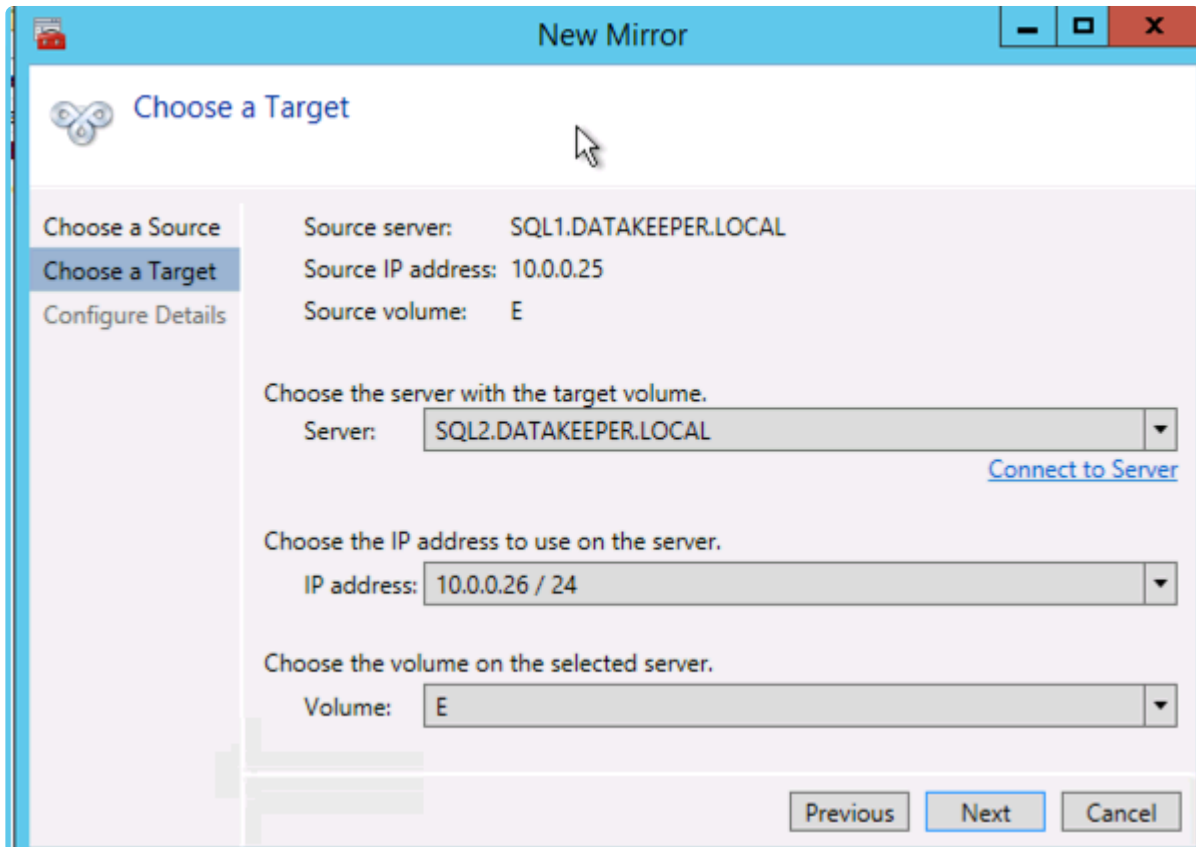
The screenshot shows a window titled "SIOS DataKeeper" with a standard Windows title bar. Inside, the heading "Create a new job" is followed by a descriptive paragraph: "A job provides a logical grouping of related mirrors and servers. Provide a name and description for this new job to help remember it." Below this, there are two input fields: "Job name:" with the text "E Drive" and "Job description:" with the text "SQL Data". At the bottom right, there are two buttons: "Create Job" and "Cancel".

Choose your source Server, IP and Volume. The IP address chosen will determine the replication network.



The screenshot shows a window titled "New Mirror" with a standard Windows title bar. The "Choose a Source" tab is selected, indicated by a blue header and a left sidebar with three options: "Choose a Source", "Choose a Target", and "Configure Details". The main area contains three sections: "Choose the server with the source volume." with a dropdown menu showing "SQL1.DATAKEEPER.LOCAL" and a "Connect to Server" link; "Choose the IP address to use on the server." with a dropdown menu showing "10.0.0.25 / 24"; and "Choose the volume on the selected server." with a dropdown menu showing "E". At the bottom right, there are "Next" and "Cancel" buttons.

Choose your target server.

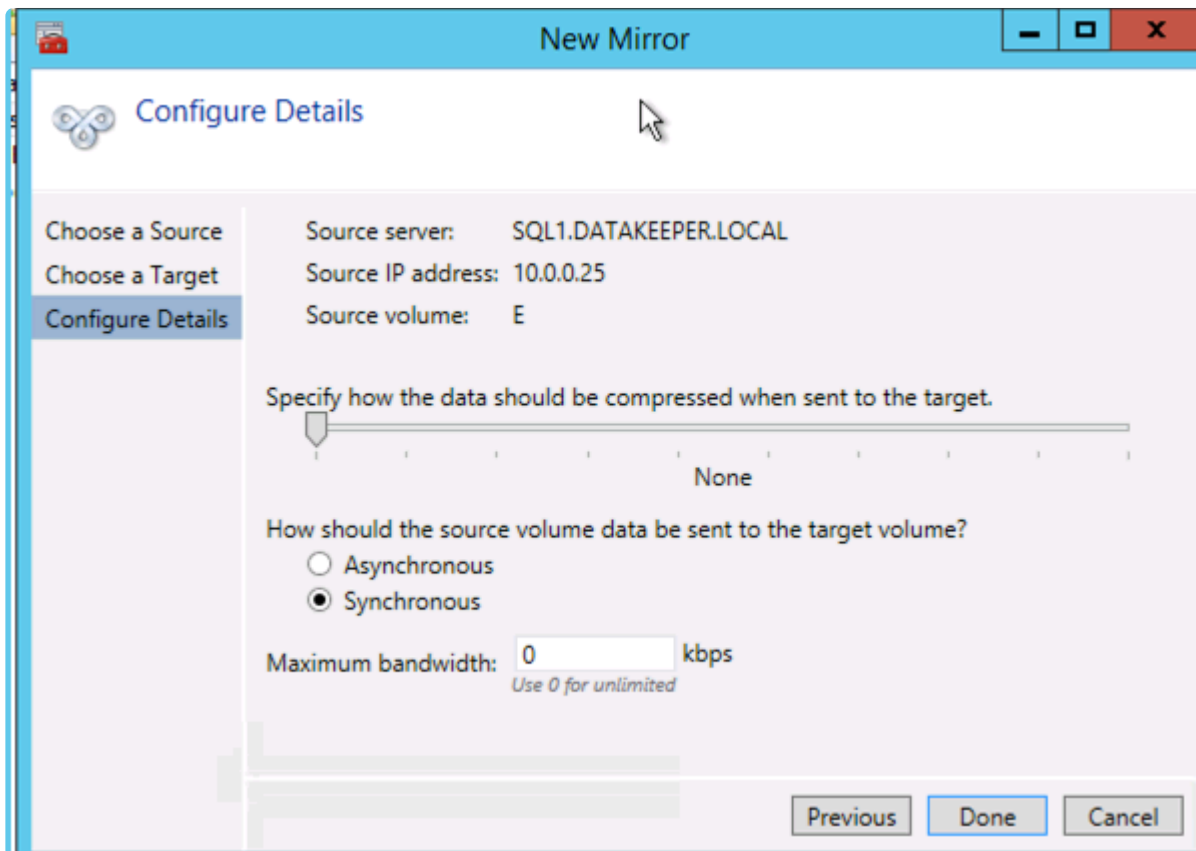


The screenshot shows the 'New Mirror' wizard window with the 'Choose a Target' step selected. The window has a blue title bar and a sidebar on the left with three options: 'Choose a Source', 'Choose a Target' (highlighted), and 'Configure Details'. The main area contains the following fields and controls:

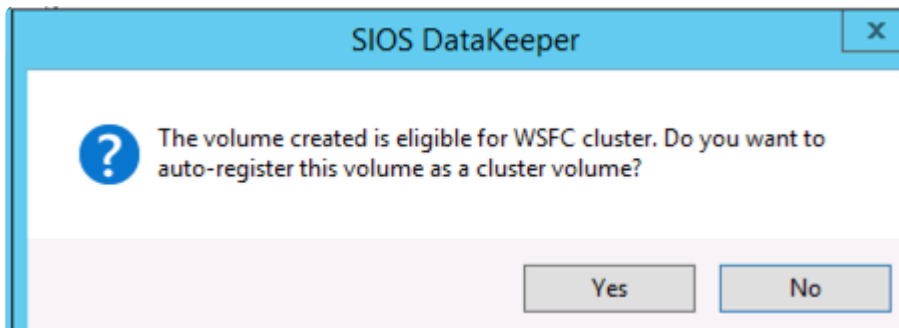
- Source server:** SQL1.DATAKEEPER.LOCAL
- Source IP address:** 10.0.0.25
- Source volume:** E
- Choose the server with the target volume.**
 - Server:** A dropdown menu showing 'SQL2.DATAKEEPER.LOCAL'.
 - [Connect to Server](#) (blue link)
- Choose the IP address to use on the server.**
 - IP address:** A dropdown menu showing '10.0.0.26 / 24'.
- Choose the volume on the selected server.**
 - Volume:** A dropdown menu showing 'E'.

At the bottom right, there are three buttons: 'Previous', 'Next' (highlighted with a blue border), and 'Cancel'.

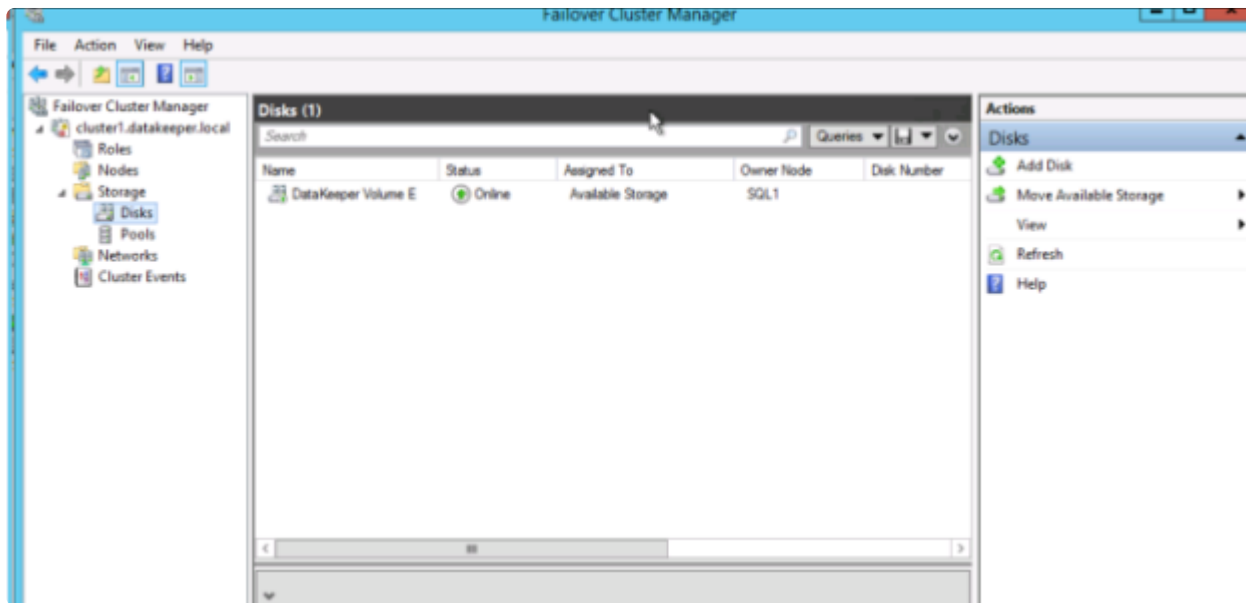
Choose your options. If two VMs are in the same geographic region we recommend using synchronous replication. For long distance replication we recommend using asynchronous replication with some level of compression. Since both SQL1 and SQL2 are in the same region, select Synchronous here.



Click **Yes** to register a new DataKeeper Volume Resource in Available Storage in Failover Clustering.

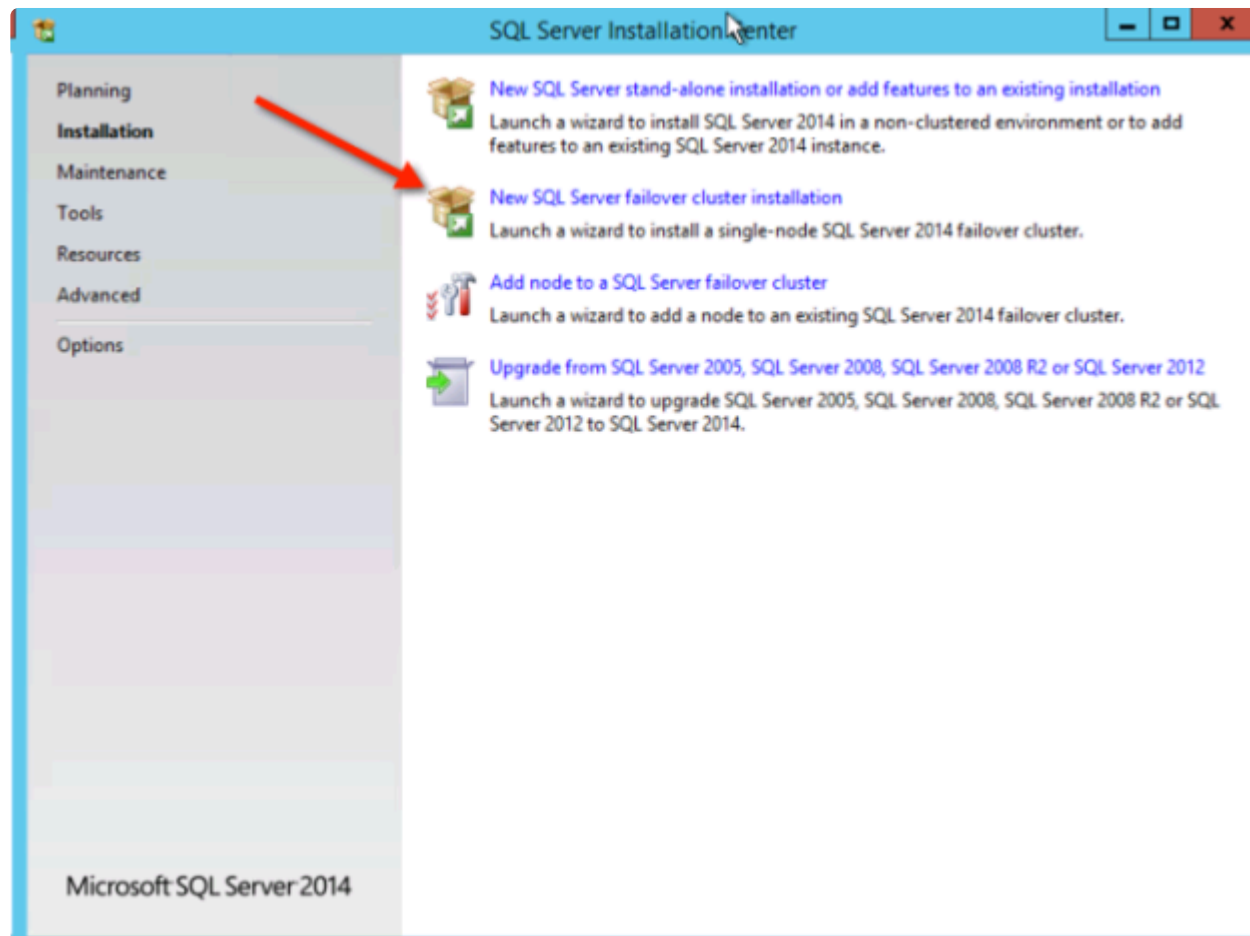


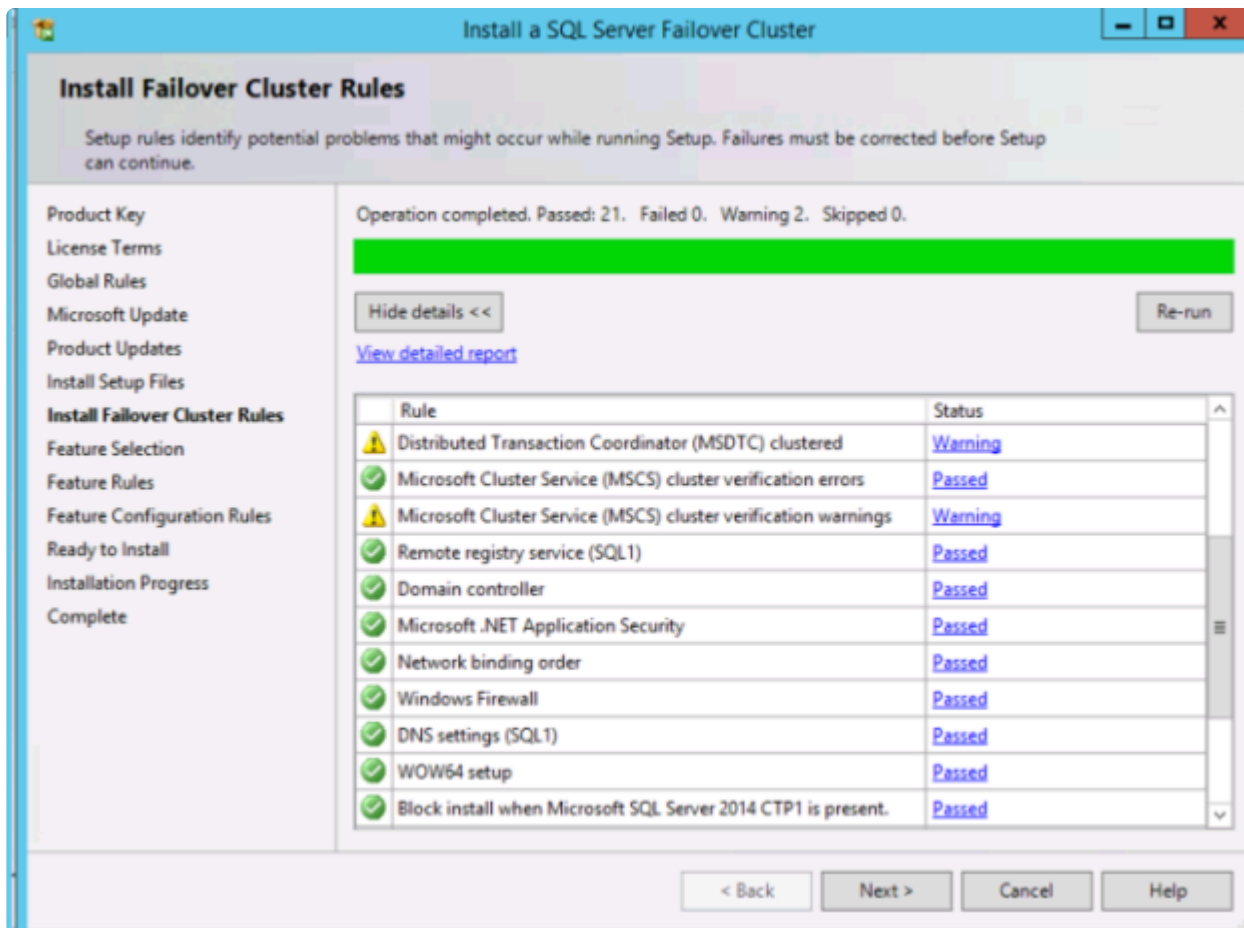
The new DataKeeper Volume Resource will appear in the Available Storage cluster group.

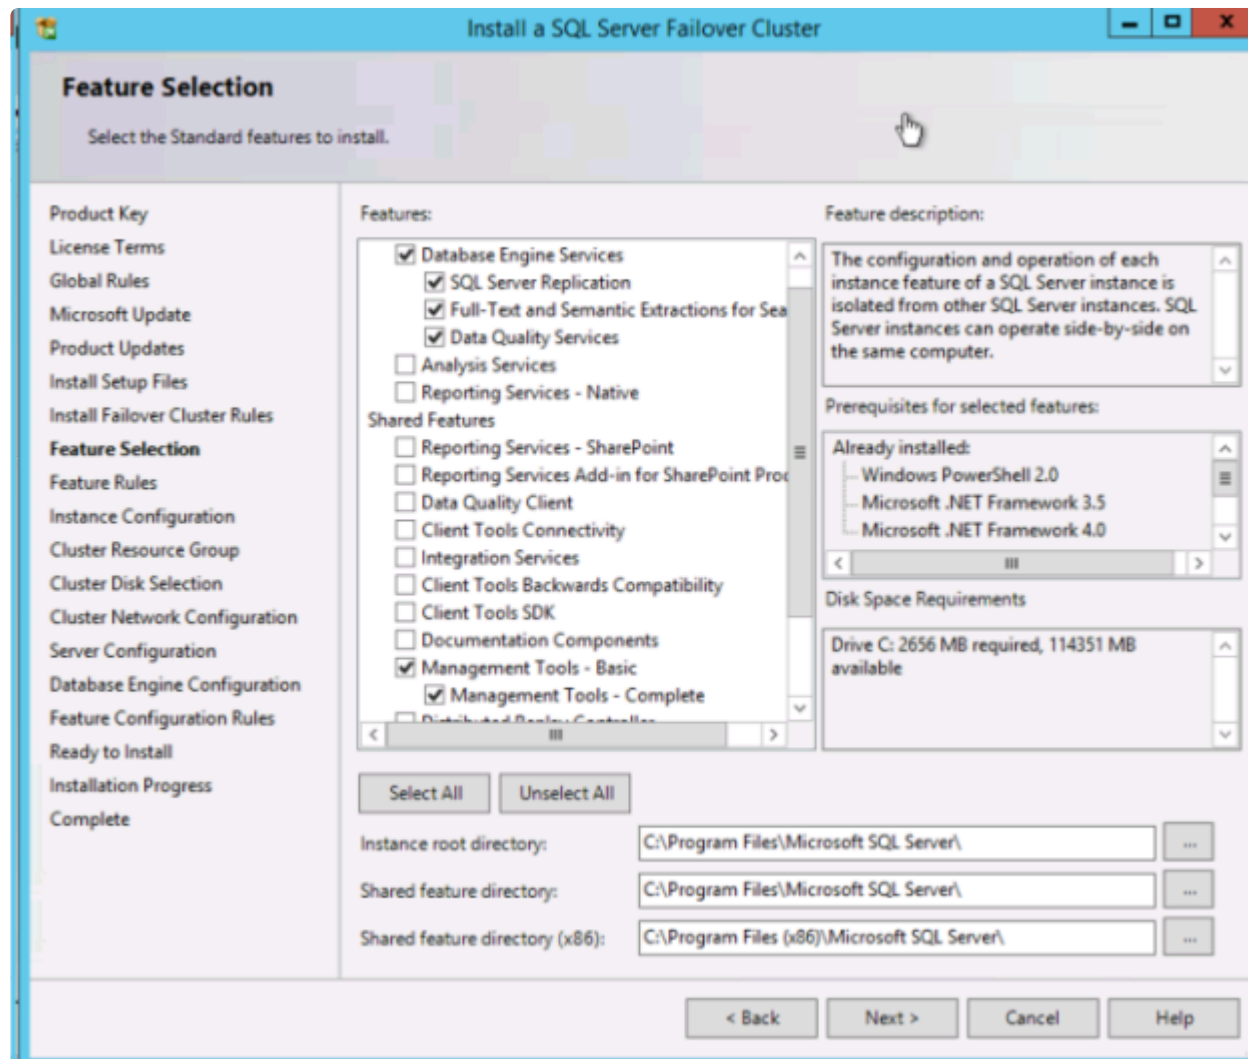


INSTALL THE FIRST CLUSTER NODE

You are now ready to install your first node. The cluster installation will proceed the same as any other SQL cluster. Start the installation on the first cluster node using the **New SQL Server failover cluster installation** option.







The screenshot shows the 'Install a SQL Server Failover Cluster' wizard, specifically the 'Instance Configuration' step. The left sidebar lists various installation steps, with 'Instance Configuration' currently selected. The main area contains fields for configuring the new SQL Server failover cluster. A mouse cursor is pointing at the 'SQL Server Network Name' field, which contains the text 'sqlcluster'. Below this, the 'Default instance' radio button is selected, and the 'Instance ID' field contains 'MSSQLSERVER'. The 'SQL Server directory' is set to 'C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER'. At the bottom, there is a table titled 'Detected SQL Server instances and features on this computer:' which is currently empty. Navigation buttons at the bottom include '< Back', 'Next >', 'Cancel', and 'Help'.

Install a SQL Server Failover Cluster

Instance Configuration

Specify the name and instance ID for the instance of SQL Server. Instance ID becomes part of the installation path.

Specify a network name for the new SQL Server failover cluster. This will be the name used to identify your failover cluster on the network.

SQL Server Network Name:

☒ Default instance

☐ Named instance:

Instance ID:

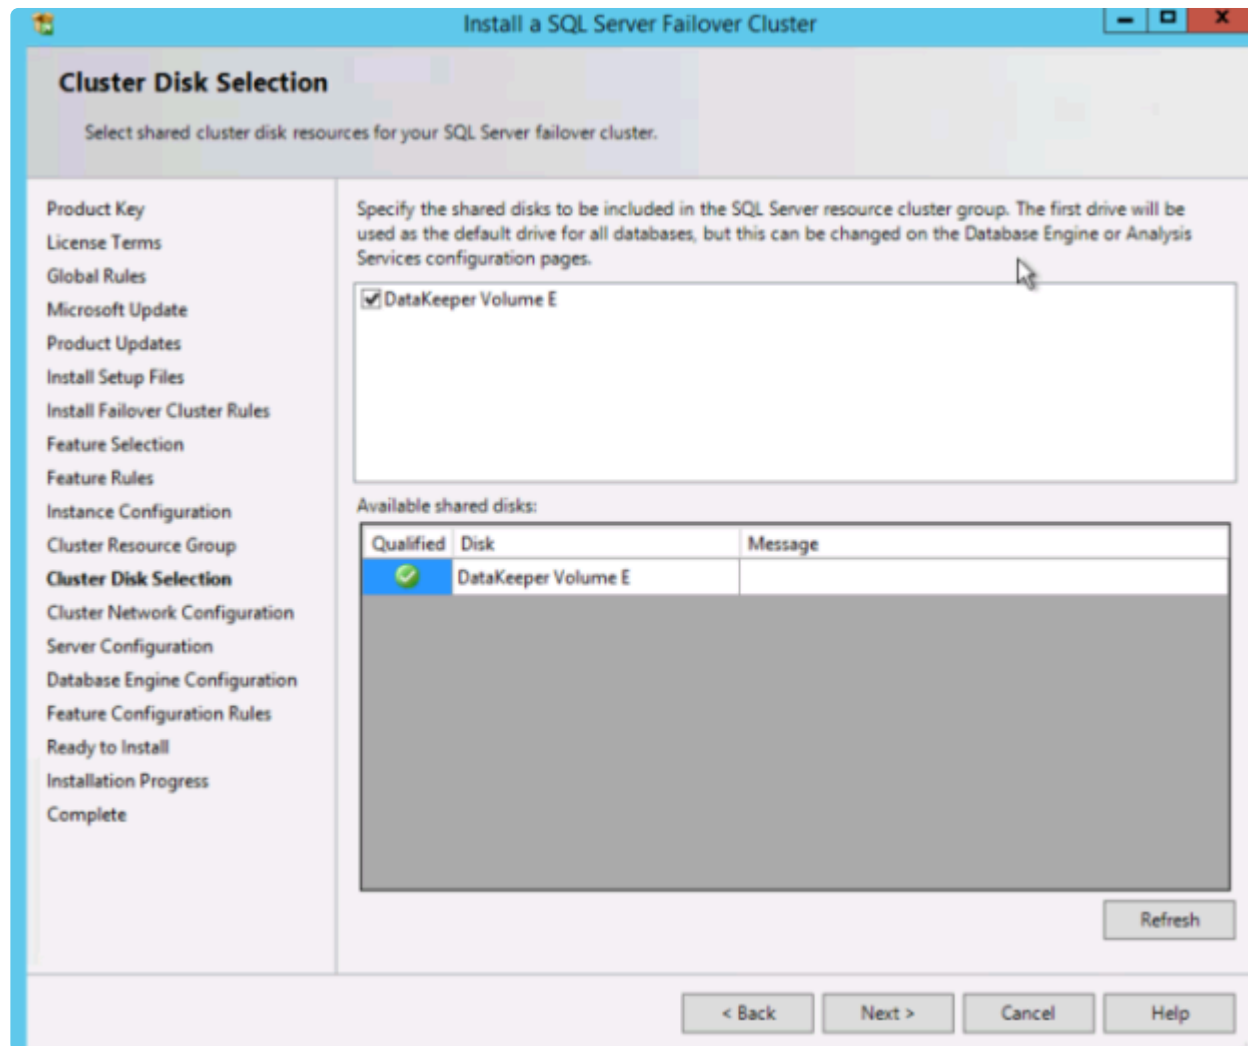
SQL Server directory: C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER

Detected SQL Server instances and features on this computer:

Instance	Cluster Network Name	Features	Edition	Version	Inst
----------	----------------------	----------	---------	---------	------

< Back Next > Cancel Help

The DataKeeper Volume Resource is recognized as an available disk resource, just as if it were a shared disk.



Make a note of the IP address you select here. It must be a unique IP address on your network. This IP address will be used later when creating the Internal Load Balancer.

Install a SQL Server Failover Cluster

Cluster Network Configuration

Select network resources for your SQL Server failover cluster.

Product Key
License Terms
Global Rules
Microsoft Update
Product Updates
Install Setup Files
Install Failover Cluster Rules
Feature Selection
Feature Rules
Instance Configuration
Cluster Resource Group
Cluster Disk Selection
Cluster Network Configuration
Server Configuration
Database Engine Configuration
Feature Configuration Rules
Ready to Install
Installation Progress
Complete

Specify the network settings for this failover cluster:

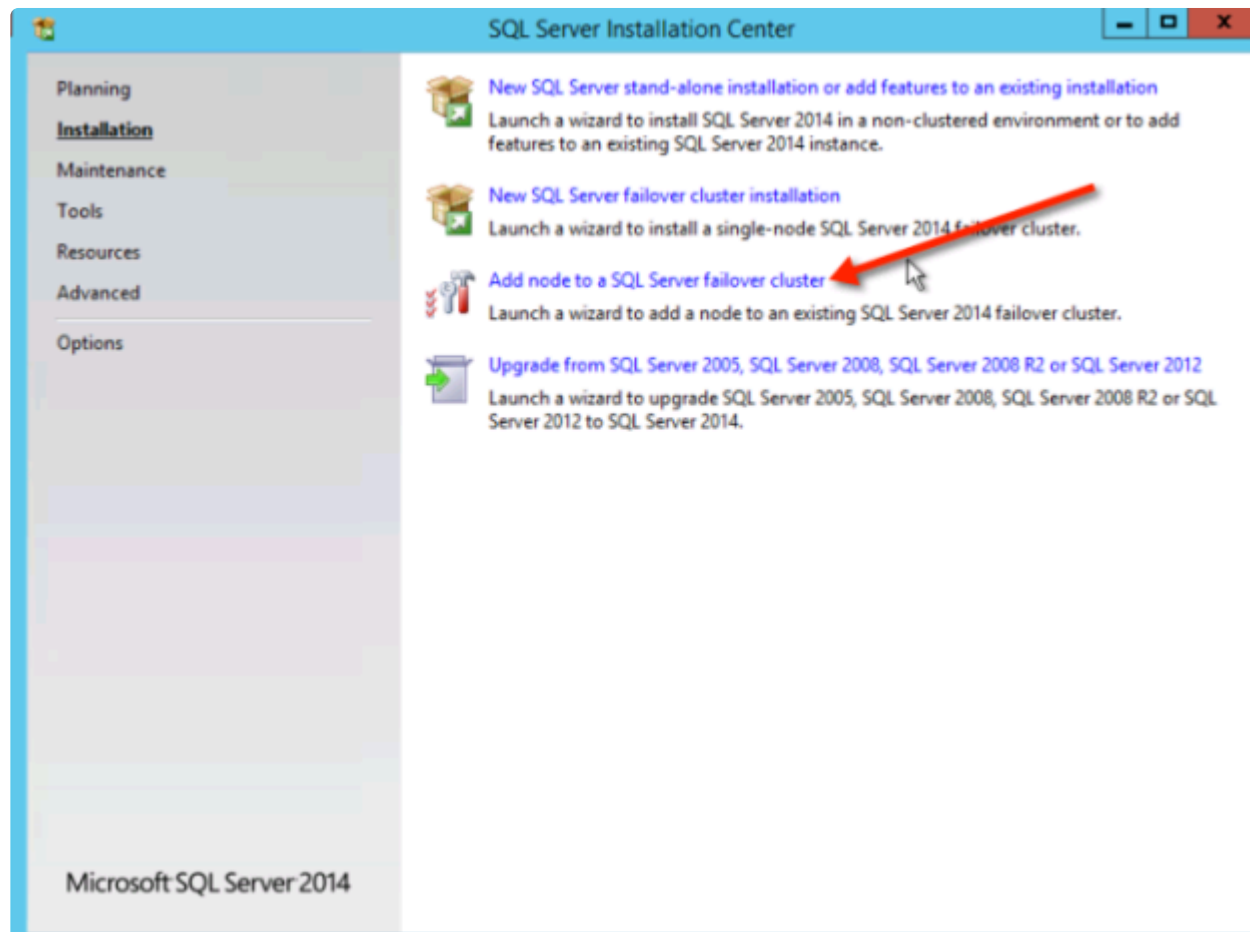
<input checked="" type="checkbox"/>	IP Type	DHCP	Address	Subnet Mask	Subnet(s)	Network
<input checked="" type="checkbox"/>	IPv4	<input type="checkbox"/>	10.0.0.201	255.255.255.0	10.0.0.0/24	Cluster Network 1

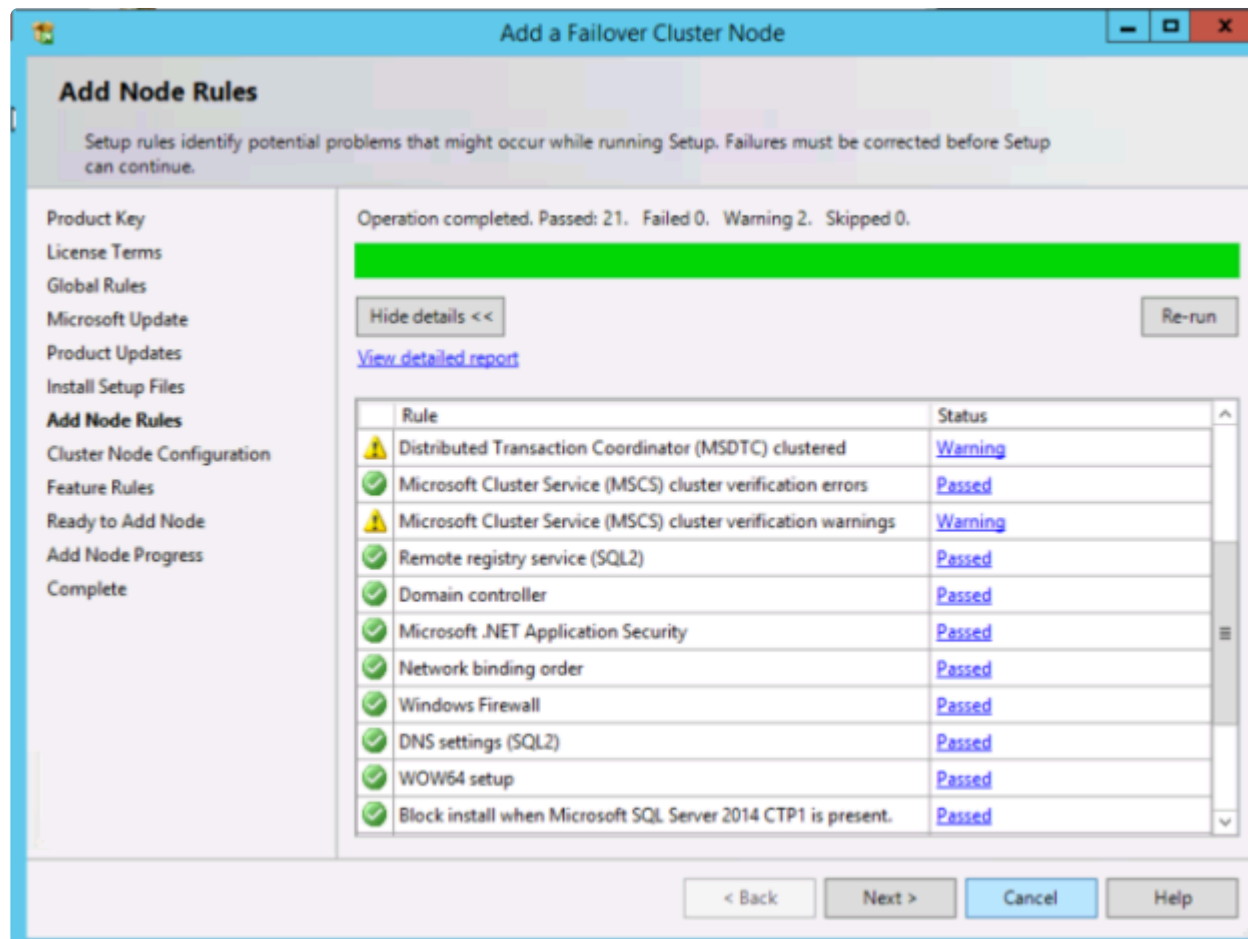
Refresh

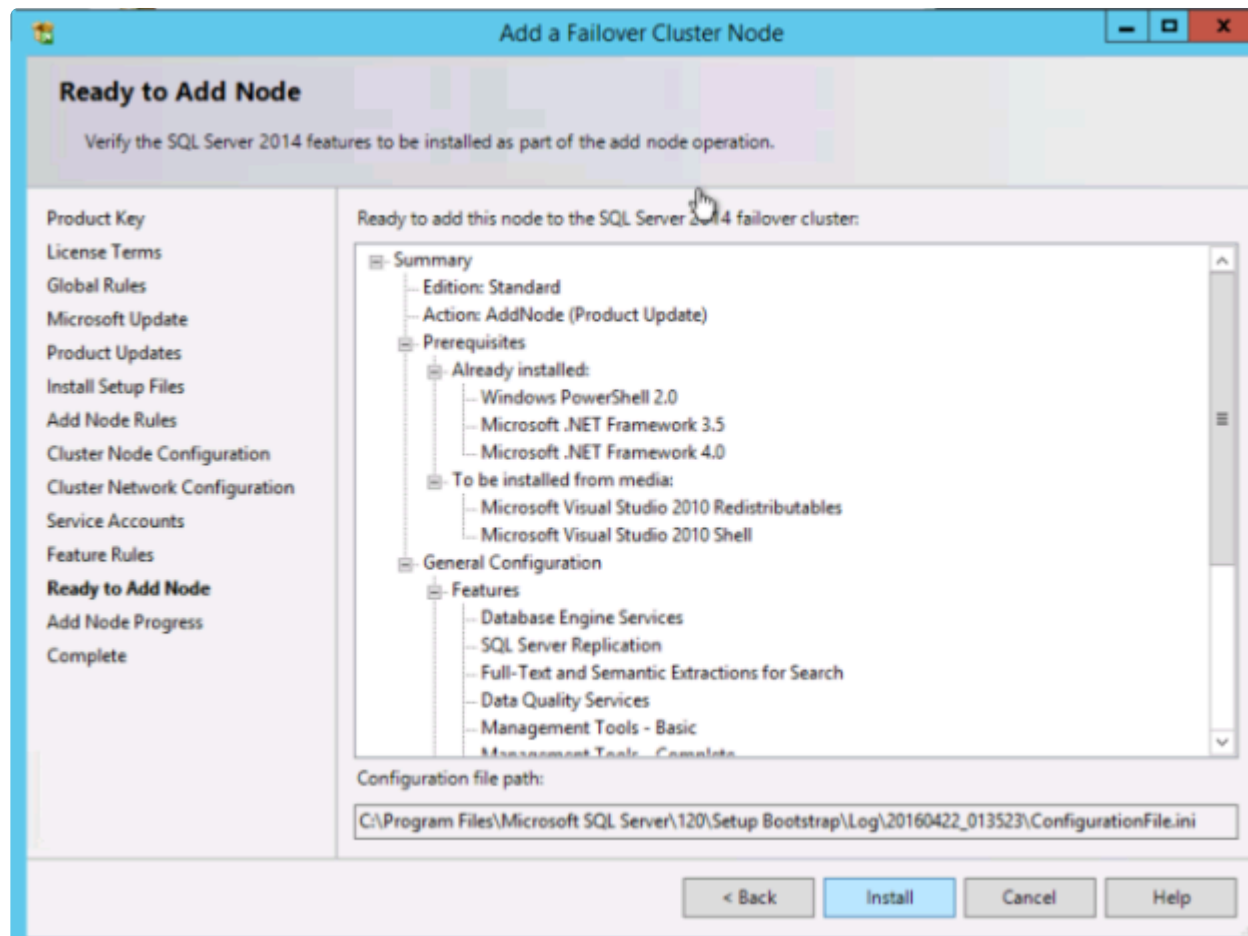
< Back Next > Cancel Help

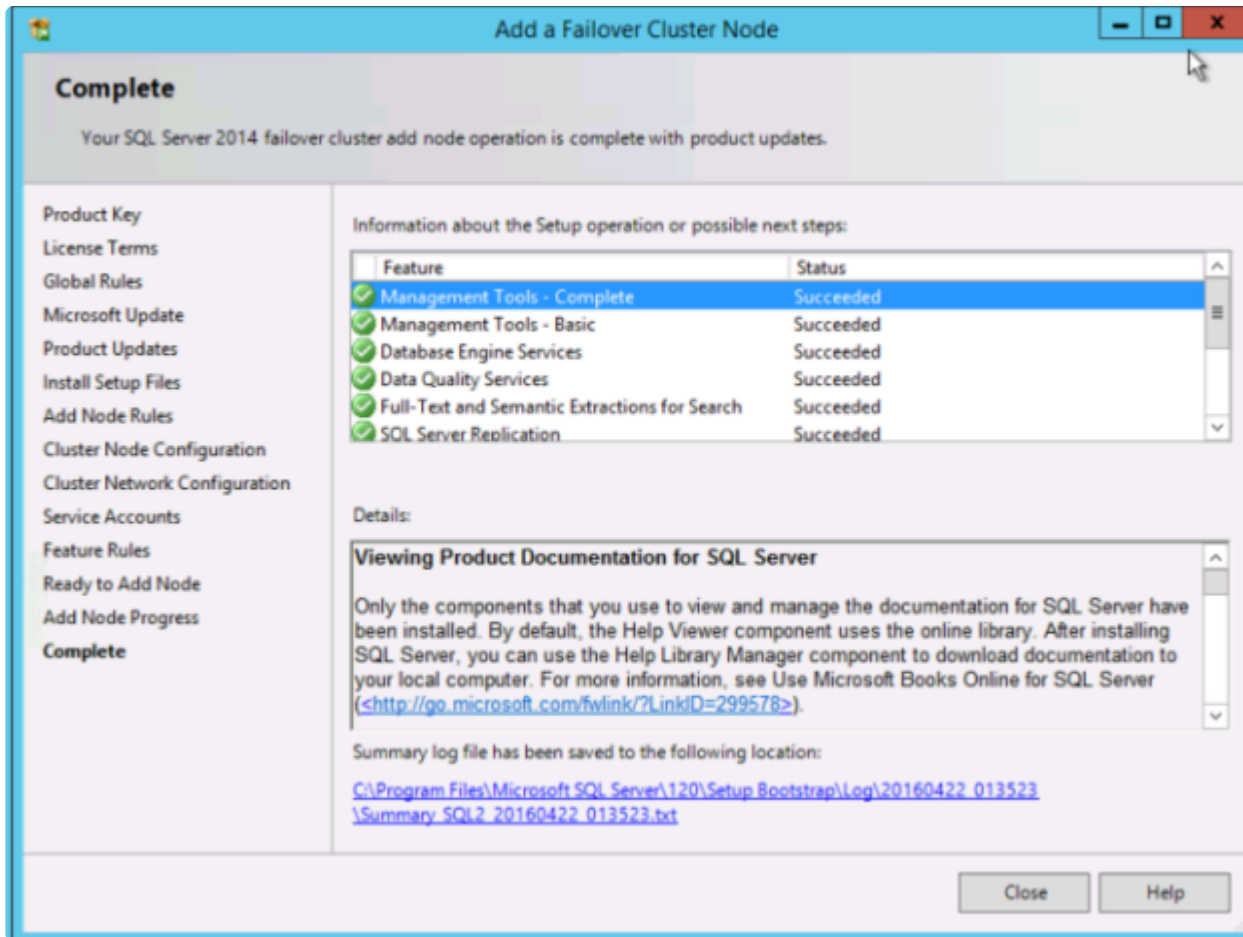
ADD THE SECOND NODE

After the first node is installed successfully, start the installation on the second node using the **Add node to a SQL Server failover cluster** option.



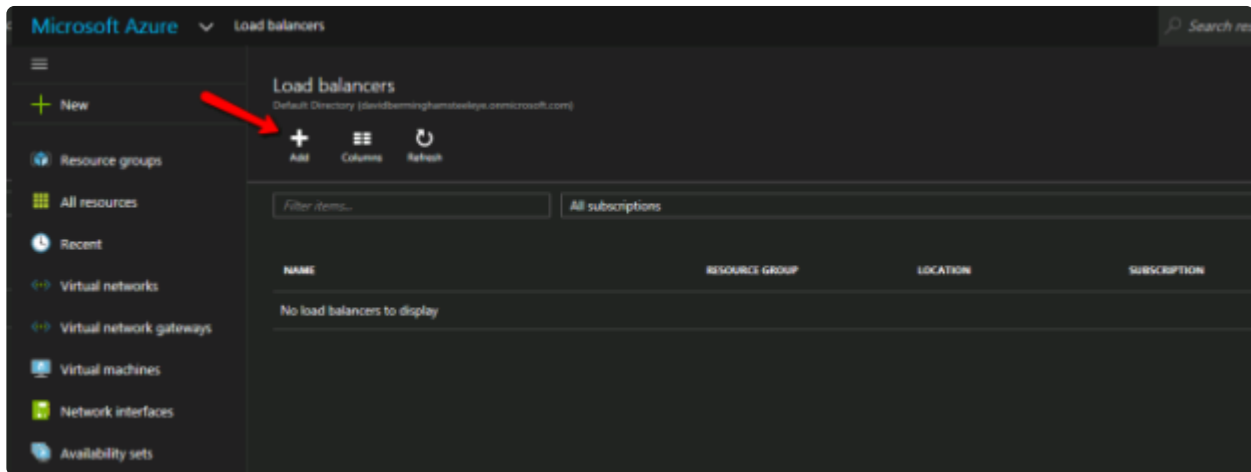






CREATE THE INTERNAL LOAD BALANCER

The failover clustering in Azure is different than traditional infrastructures. The Azure network stack does not support gratuitous ARPS, so clients cannot connect directly to the cluster IP address. Instead, clients connect via a load balancer resource which redirects them to the active cluster node. Thus an Internal Load Balancer must be created. This can all be done through the Azure Portal as shown below.



A Public Load Balancer can be used if your client connects over the public Internet. Since our clients reside in the same vNet, we will create an Internal Load Balancer. It is important that the Virtual Network is the same as the network where your cluster nodes reside. The Private IP address that you specify must be EXACTLY the same as the address you used to create the SQL Cluster Resource.

Microsoft Azure

Load balancers > Create load balancer

New

Resource groups

All resources

Recent

Virtual networks

Virtual network gateways

Virtual machines

Network interfaces

Availability sets

Load balancers

Storage accounts

Subscriptions

Network security groups

What's new

Marketplace

Storage accounts (class...

Virtual networks (classic)

Virtual machines (classic)

Browse >

Create load balancer

* Name

SQLILB

* Scheme ⓘ

PublicInternal

* Virtual network

ergergergfdswVNET

* Subnet

staticSubnet (10.0.0.0/24)

* IP address assignment

DynamicStatic

* Private IP address

10.0.0.201

Subscription

Windows Azure MSDN - Visual Studio Prei

* Resource group

SQLCluster

Location

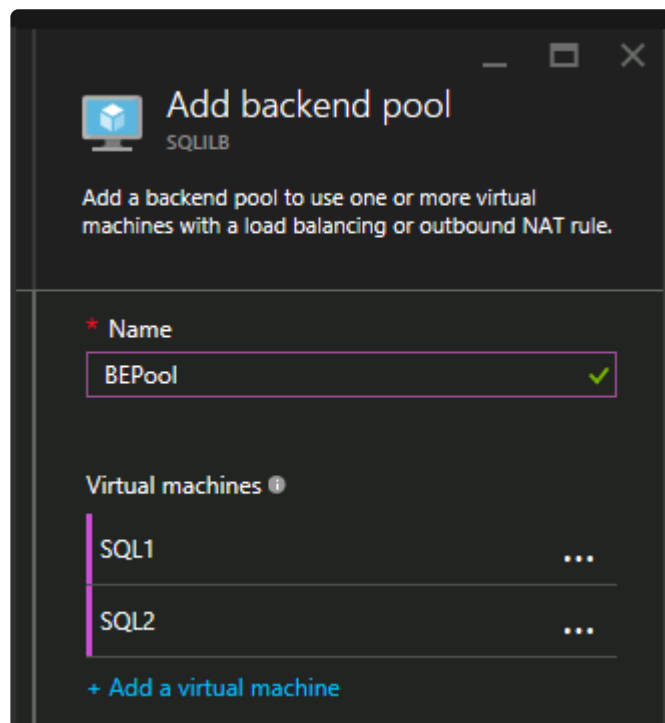
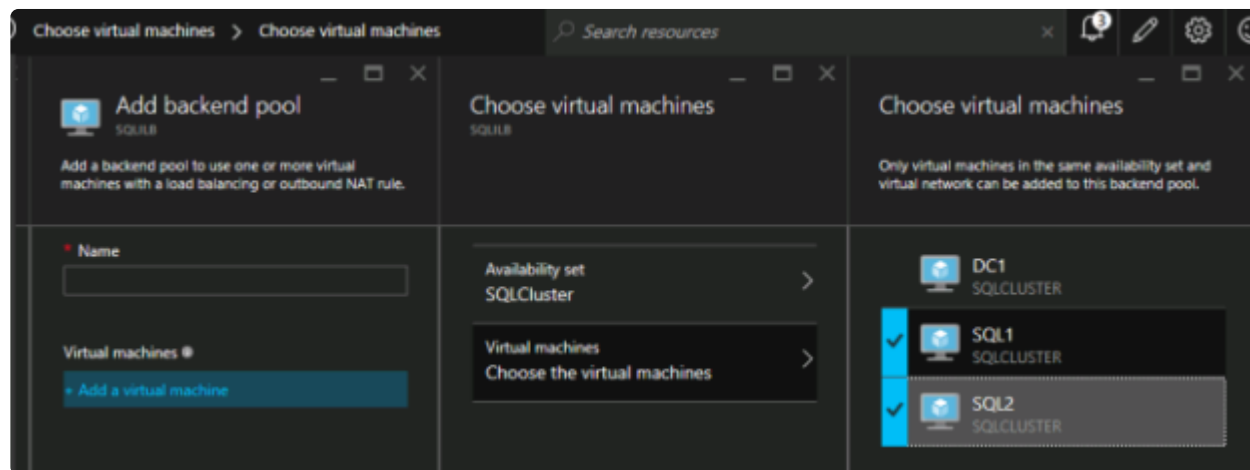
East US 2

☐ Pin to dashboard

Create

Page 64 of 544

After the Internal Load Balancer (ILB) is created add a backend pool. Through this process you will choose the Availability Set where your SQL Cluster VMs reside. However, when you choose the actual VMs to add to the Backend Pool, be sure not to choose the VM hosting your file share witness, DC1. You do **not** want to redirect SQL traffic to your file share witness.



The next step is to add a Probe. The probe we add will probe Port 59999. This probe determines which node is active in our cluster.

Add probe

SQLILB

*

Name

SQLProbe

✓

Protocol

HTTP

TCP

*

Port

59999

*

Interval

ⓘ

5

seconds

*

Unhealthy threshold

ⓘ

2

consecutive failures

Finally, a load balancing rule is needed to redirect the SQL Server traffic. A Default Instance of SQL uses port 1433. You can add rules for 1434 or others depending upon your application requirements. It is very important that Floating IP (direct server return) is Enabled.

Add load balancing rule

SQULB

* Name

SQL1433 ✓

Protocol

☒ TCP ☐ UDP

* Port

1433 ✓

* Backend port ⓘ

1433

Backend pool ⓘ

BEPool (2 virtual machines) ▼

Probe ⓘ

SQLProbe (TCP:59999) ▼

Session persistence ⓘ

None ▼

Idle timeout (minutes) ⓘ

4

Floating IP (direct server return) ⓘ

☐ Disabled ☒ Enabled

OK

FIX THE SQL SERVER IP RESOURCE

The final step in the configuration is to run the following PowerShell script on one of your cluster nodes. This will allow the Cluster IP Address to respond to the ILB probes and ensure that there is no IP address conflict between the Cluster IP Address and the ILB. **Note:** You will need to edit this script to fit your environment. The subnet mask is set to 255.255.255.255, this is not a mistake, leave it as is. This creates a host specific route to avoid IP address conflicts with the ILB.

```
# Define variables

$ClusterNetworkName = ""

# the cluster network name (Use Get-ClusterNetwork on Windows Server 2012
of higher to find the name)

$IPResourceName = ""

# the IP Address resource name

$ILBIP = ""

# the IP Address of the Internal Load Balancer (ILB)

Import-Module FailoverClusters

# If you are using Windows Server 2012 or higher:

Get-ClusterResource $IPResourceName | Set-ClusterParameter -Multiple
@{Address=$ILBIP;ProbePort=59999;SubnetMask="255.255.255.255";Network=$ClusterNetworkName}

# If you are using Windows Server 2008 R2 use this:

#cluster res $IPResourceName /priv enabledhcp=0 address=$ILBIP
probeport=59999 subnetmask=255.255.255.255
```

DataKeeper Cluster Edition Release Notes

SIOS DataKeeper Cluster Edition

Release Notes

Version 8.6.5

(Version 8 Update 6 Maintenance 5)

Important!!

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

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

Introduction

SIOS DataKeeper Cluster Edition is a highly optimized host-based replication solution which integrates seamlessly with Windows Server 2012, Windows Server 2012 R2, and Windows Server 2008 R2/2008 R2 SP1 Failover Clustering. Features of Windows Server Failover Clustering such as cross-subnet failover and tunable heartbeat parameters make it easy for administrators to deploy geographically dispersed clusters. SIOS DataKeeper provides the data replication mechanism which extends both versions of Windows Clustering, allowing administrators to take advantage of these advanced features to support non-shared disk high availability configurations.

Once SIOS DataKeeper Cluster Edition is installed, a new storage class resource type called DataKeeper Volume is available. This new SIOS DataKeeper Volume resource can be used in place of the traditional Physical Disk shared storage resource to enable geographically dispersed clusters, also known as multi-site clusters.

New Features of SIOS DataKeeper Cluster Edition v8

Feature	Description
New in This Release	

General maintenance	Bug Fix.
New in Version 8.6.4	
Windows 2019 Support	DataKeeper 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.
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	DataKeeper supports Microsoft SQLServer 2017.
New in Version 8.6	
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 8.5.1	
Windows 2016 Support	DataKeeper now supports Windows 2016.
VSS Provider	SIOS VSS Provider has been disabled by default.
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.
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	
DataKeeper Volume	DataKeeper Volume Resource Health Check determines if the underlying volume

Resource Health Check	device becomes unreachable.
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.
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.
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).
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	
DataKeeper Non-Mirrored Volume Cluster Resource	The DataKeeper Non-Mirrored Volume Cluster Resource allows users to use a local volume in a failover cluster without requiring that it be part of a mirror. Some of the common use cases for this feature include enabling rolling cluster OS upgrades on existing hardware and moving tempdb to local storage in SQL 2008 R2 clusters and earlier.
General Maintenance	Bug Fixes.
New in Version 8.1	
General Maintenance	Bug Fixes.
New in Version 8.0.1	
General Maintenance	Bug Fixes.
New in Version 8.0	
Replicate to Node Outside of Cluster	DataKeeper now allows you to have a replication target which resides outside of the failover cluster.
Operating System Support	DataKeeper now only supports Windows 2008R2 and later 64-bit Operating Systems. To run on an earlier version of Windows or on 32-bit systems, you must use DataKeeper v7.
Windows 2012 R2 Support	DataKeeper now supports Windows 2012 R2.

General Maintenance	Bug Fixes.
---------------------	------------

Bug Fixes

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

	Description
4463	Prevent unexpected lock of cluster volume resource

Product Definitions and Platforms

Product Requirements

Product	Operating Systems	Additional Software
Server Components	See the DKCE Support Matrix	<p>Hotfix – KB 951308</p> <p>https://support.microsoft.com/en-us/help/951308</p> <p>If protecting Hyper-V resources, Hotfix KB 958065</p> <p>https://support.microsoft.com/en-us/help/958065</p> <p>Note: These hotfixes are not required for Windows Server 2008 R2/2008 R2 SP1.</p> <p>Microsoft Hotfix KB 2741477 is now available that will allow NICs to be added to a Virtual Machine after the VM has been placed into a Failover Cluster. (For more information, see Hyper-V Host Cluster Error.)</p>
User Interface	See the DKCE Support Matrix	<p>MMC 3.0 – download from:</p> <p>https://download.cnet.com/Microsoft-Management-Console-3-0-for-Windows-XP-KB907265/3000-2206_4-10741230.html</p>

Note: All servers should run the same version of Windows and the same version of DataKeeper software.

Make sure you verify the following settings prior to installing and configuring SIOS DataKeeper Cluster

Edition.

- **Important:** SIOS Technology Corp. recommends that users use Domain accounts that have local Admin privileges on all servers running DataKeeper. If you are using Local accounts, the username and passwords must match on all servers running DataKeeper. This recommendation is for all editions and all platforms.
- Follow Microsoft best practices for deploying geographically dispersed clusters, including changing the quorum mode majority node setting with a file share witness.
- DataKeeper Failover Cluster registration is automatic and occurs 60 seconds after the following events have occurred on each cluster node:
 - A DataKeeper Cluster Edition license has been installed on each cluster node.
 - The Windows server Failover Clustering feature has been installed on each server.
 - A Windows server Cluster configuration has been formed.

Local Security Policy Requirement

If your Windows servers are not in a domain and you are going to run the DataKeeper Service as the local system account, the Local Security policy setting “**Network Access: Let Everyone permissions apply to anonymous users**” must be enabled.

Known Issues

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.

Windows 2016

- [Occasional job creation failure](#)

SCVMM 2012

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

Windows Server 2012

For issues and enhancements related to **Windows Server 2012**, see the following topics:

[WSFC 2012 Failover Cluster Manager UI Defect](#)

[WSFC 2012 New File Server Type Not Supported](#)

[Manual Creation of a Mirror in WSFC](#)

[WSFC 2012 Cluster Creation Default Setting Issue](#)

[WSFC 2012 File Shares Cannot be Created for File Server Resource](#)

[WSFC 2012 Server Manager — Incorrect Volume Display](#)

[WSFC 2012 Server Manager — DataKeeper “Disk” Not Shown as Clustered](#)

[Windows Server 2012 Default Information Missing During Mirror Creation](#)

[Windows Server 2012 DataKeeper 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](#)

[Using iSCSI Target with DataKeeper](#)

Also see the [Known Issues and Workarounds](#) and [Restrictions](#) sections of DataKeeper Cluster Edition Technical Documentation.

DataKeeper Cluster Edition Quick Start Guide

To get started using SteelEye DataKeeper Cluster Edition, refer to the [DataKeeper Cluster Edition Quick Start Guide](#).

DataKeeper Cluster Edition Installation Guide

The DataKeeper Cluster Edition Installation Guide contains information on how to install and license your Cluster Edition software.

Once you have completed the steps in this guide, you will be ready to configure your Cluster resources. The [DataKeeper Cluster Edition Technical Documentation](#) provides the information necessary to complete your DataKeeper Cluster Edition configuration.

DataKeeper Cluster Edition uses the Flexera InstallShield product to provide a standard installation interface.

Installation

Once you have downloaded the DataKeeper Cluster Edition software, the topics that follow explain the installation process.

[Core Software](#)

[Installing Core](#)

[Third Party Product Files](#)

[Application Directory Anomaly](#)

[Localized Language Supplement](#)

[Silent Installation](#)

Core Software

DataKeeper Cluster Edition Core Software

- 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 DataKeeper Cluster Edition Core Software

SIOS DataKeeper Cluster Edition 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 [DataKeeper Cluster Edition for Windows Release Notes](#) before installing and configuring DataKeeper Cluster Edition.

To install DataKeeper Cluster Edition software, run the setup program delivered with the DataKeeper Cluster Edition for Windows product. Follow the setup instructions on each screen. Some explanatory notes are included below.

Installation Notes

Once installation begins, 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 Cluster Edition 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.

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 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 contains “ExtMirr” when using DataKeeper.

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

Localized Language Supplement

For information regarding the Localized Language Supplement, please refer to [Installing LifeKeeper for Windows Localized Language Supplement](#) in the SIOS Protection Suite Documentation.

Silent Installation

Silent Installation of DataKeeper Cluster Edition

You can install DataKeeper Cluster Edition 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.

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 DataKeeper Cluster Edition 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 to install the license key.

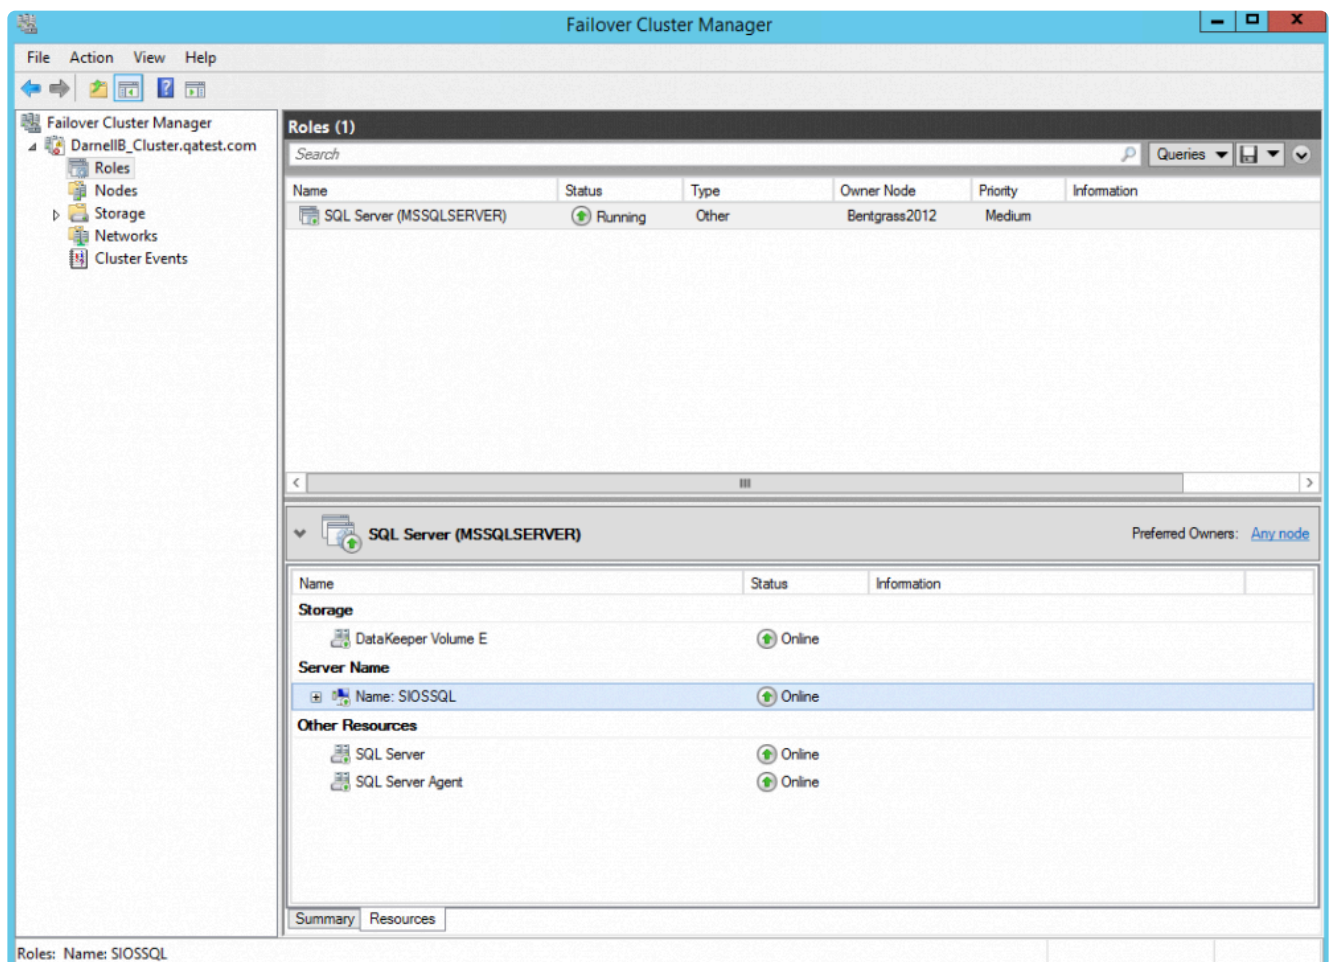
Start->All Programs->SIOS->DataKeeper->License Key Installer.

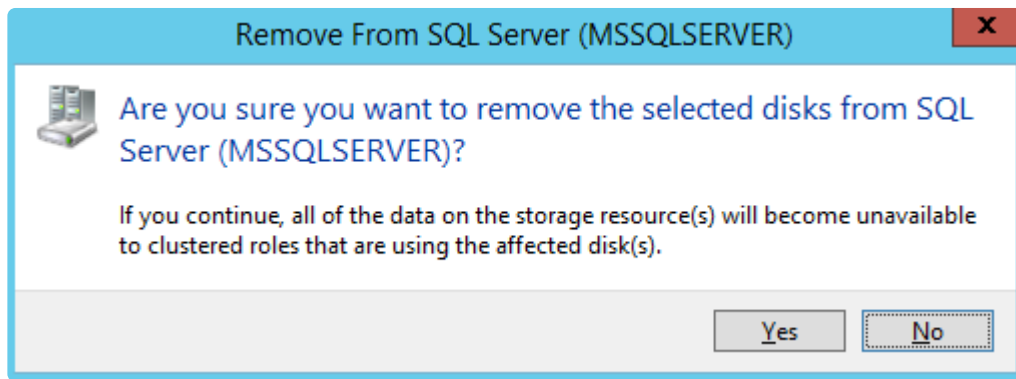
Reboot the server.

Removing a Clustered DataKeeper Volume

To remove a Clustered DataKeeper Volume follow these steps:

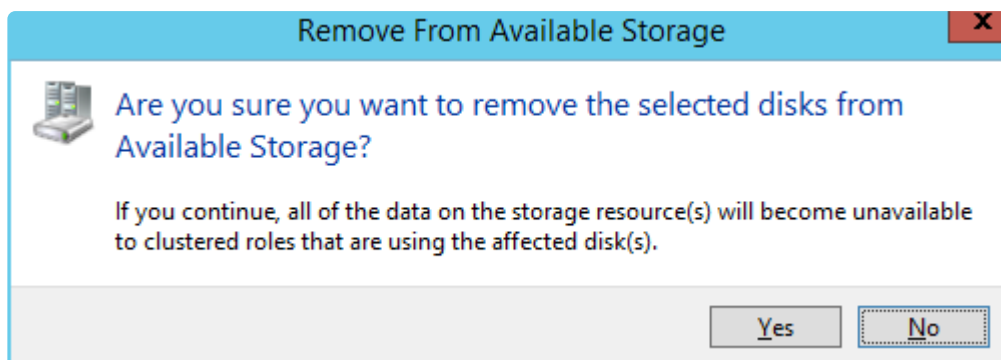
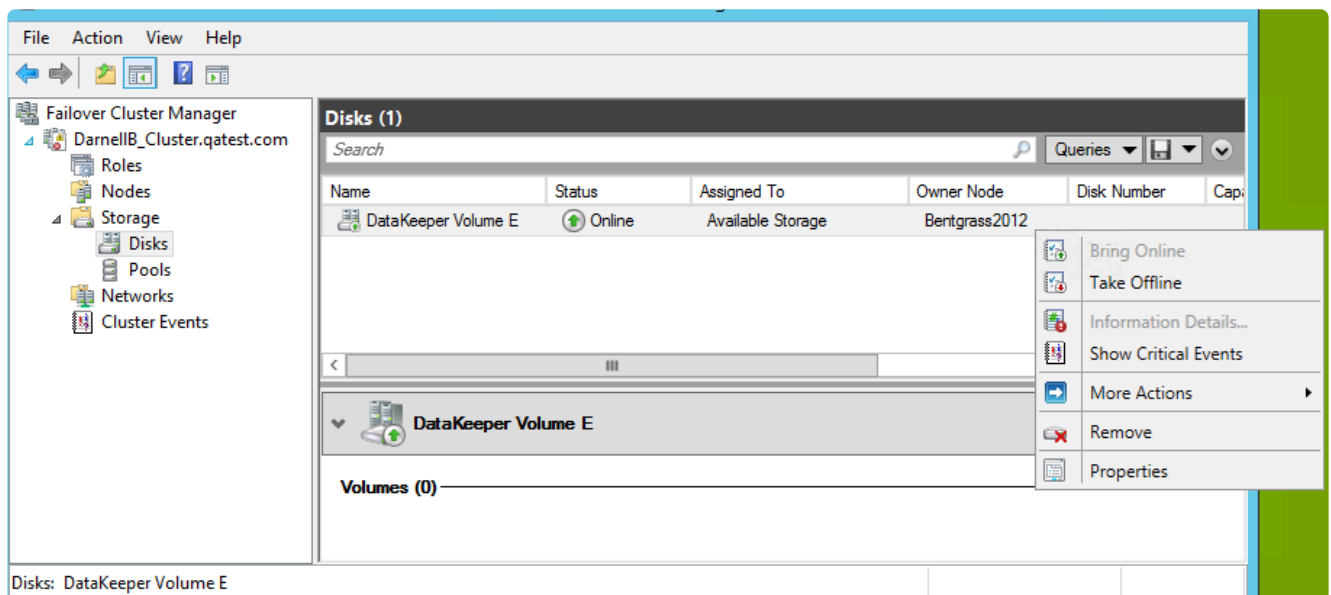
1. Launch the **Failover Cluster Manager**.
2. Select the **Role** and related **DataKeeper Storage**.
3. Remove the **DataKeeper Volume** resource from the **Role**. This will move the DataKeeper Volume resource to the “Available Storage” group.



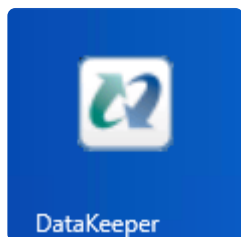


Once the DataKeeper Resource is returned to Storage/Disk and is listed as Available Storage it can be removed.

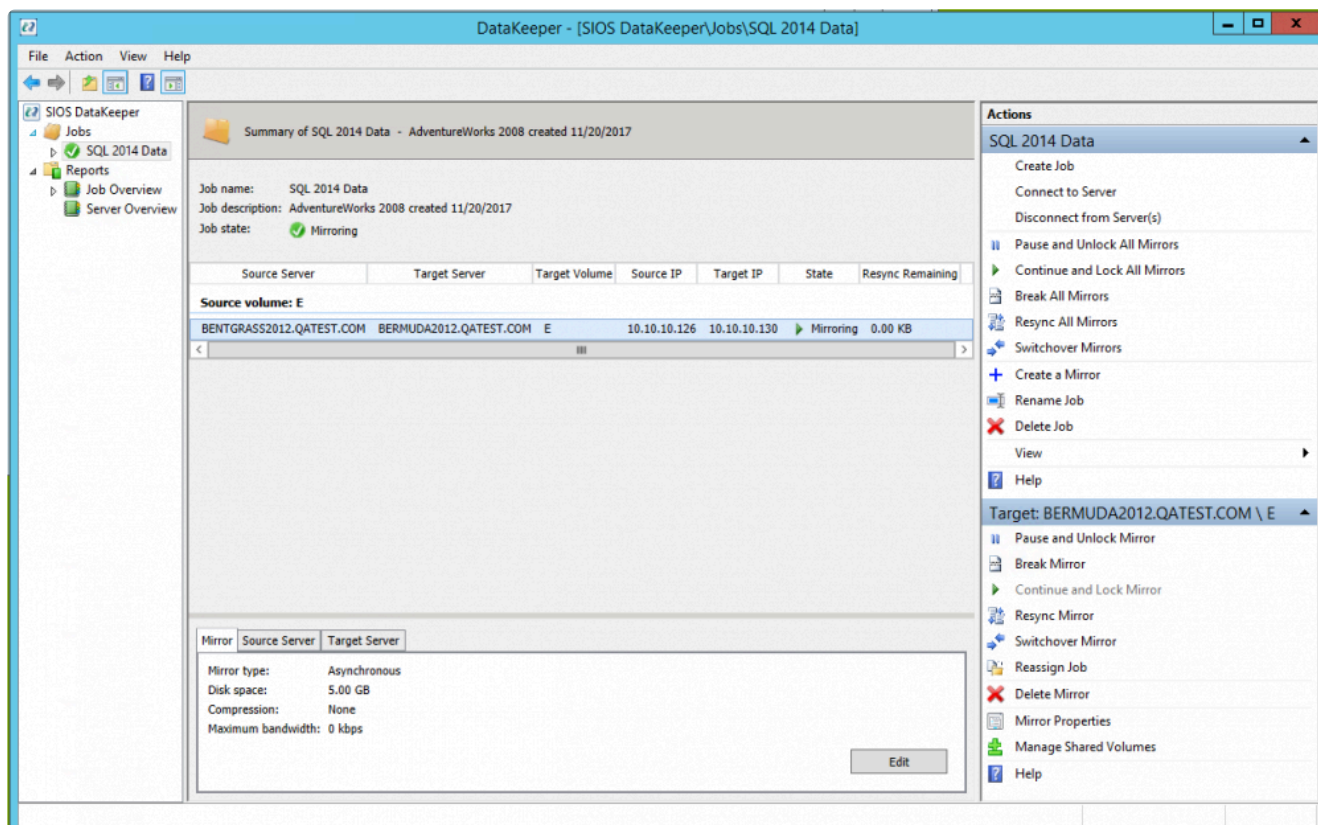
4. Select the Storage then select **Remove**.

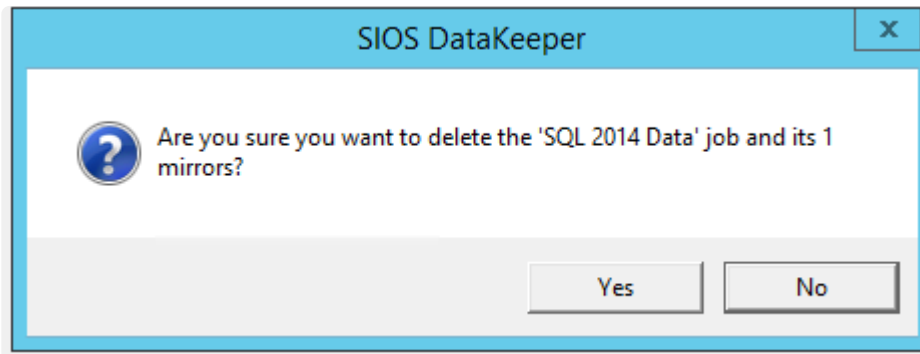


5. Launch DataKeeper.

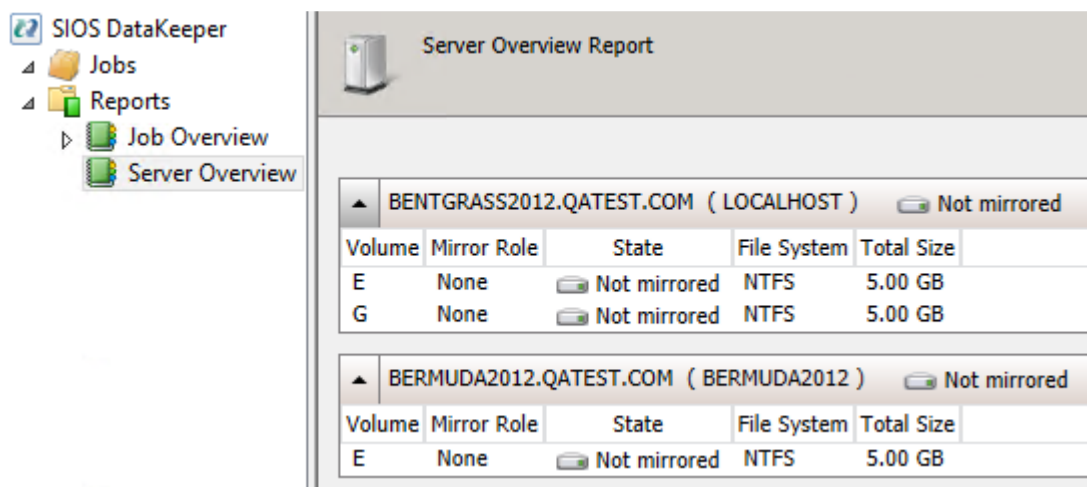
6. In the Action Pane select **Delete Job**.

✿ If the Job contains multiple mirrors use the **Delete Mirror** option to remove the mirrors only for the volume that is being deleted. Leave the other volumes in the job.





The Server Overview will show a **Not mirrored** state once the mirror has been successfully removed.

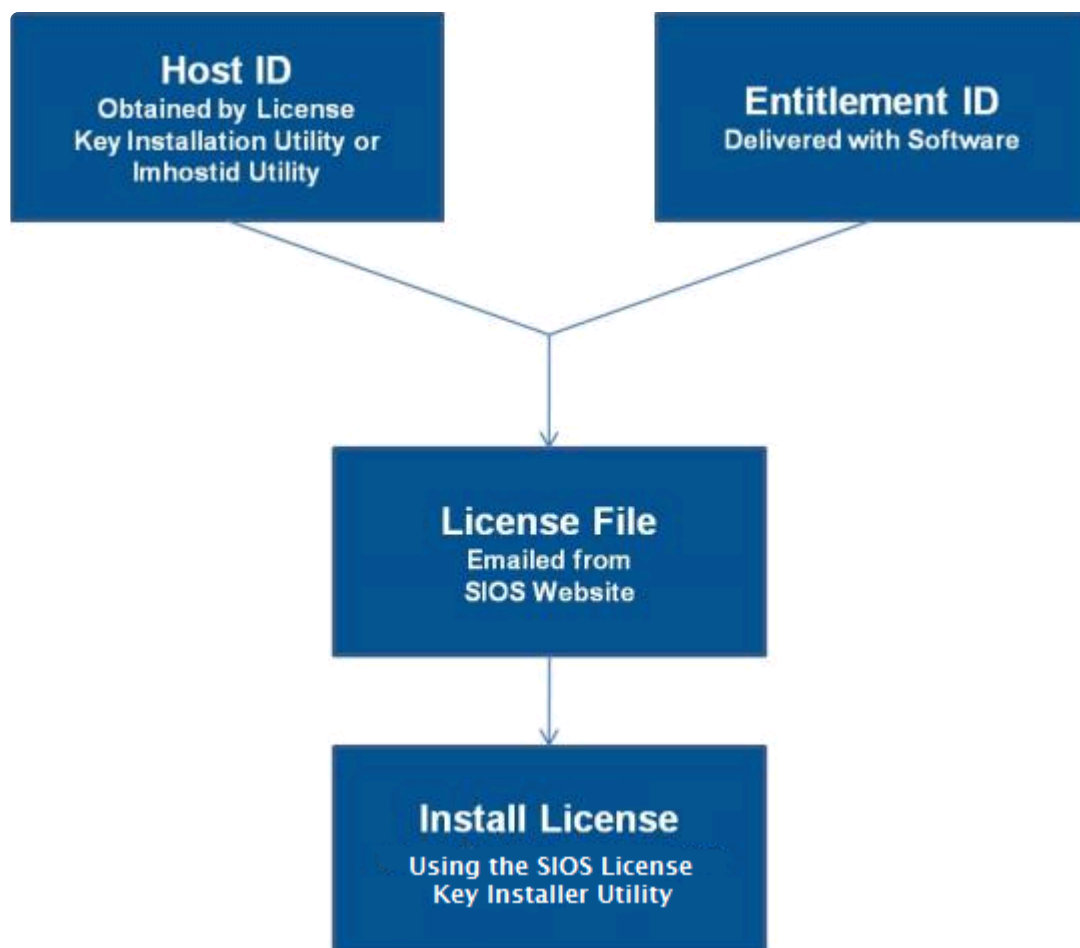


Licensing

Obtaining and Installing the License

DataKeeper Cluster Edition 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 DataKeeper Cluster Edition.

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 DataKeeper Cluster Edition software, is used to obtain the license required to run DataKeeper Cluster Edition. The process is illustrated below.



License Key Manager

In addition to installing DataKeeper Cluster Edition 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 in the DataKeeper Cluster Edition cluster:

1. Get your **Host ID**. At the end of the DataKeeper Cluster Edition 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 DataKeeper installation path, by default `C:\Program Files (x86)\SIOS\DataKeeper`) 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-DataKeeper-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 DataKeeper Cluster Edition **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-DataKeeper-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 DataKeeper Cluster Edition server(s) using the unique Host ID for each server.
7. Restart DataKeeper Cluster Edition.

Primary Network Interface Change May Require a License Rehost

The Host ID used by the License Key Installer utility is obtained from the DataKeeper Cluster Edition server's primary network interface card (NIC). DataKeeper Cluster Edition will check for a valid license each time it starts. If your DataKeeper Cluster Edition server should require a NIC replacement in the future that would cause the Host ID to change, then the next time DataKeeper Cluster Edition 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**.

Uninstalling SIOS DataKeeper Cluster Edition

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 DataKeeper Cluster Edition

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

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

Note: Uninstalling automatically stops the DataKeeper Cluster Edition 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/</code>	This file describes the DataKeeper performance monitor counters . This file can be

<code>ExtMirr/0011/ ExtMirrCounters.ini</code>	removed or left alone. It is not an executable.
----------------------------------------------------	-------------------------------------------------

Notes

- **Important:** Uninstallation of DataKeeper Cluster Edition software requires that the Microsoft Visual C++ 2008 Redistributable package be installed. Do not remove this package until DataKeeper Cluster Edition has been uninstalled.
- **Modify** or **Repair** must be run from the DataKeeper Cluster Edition setup program.
- Removal of DataKeeper Cluster Edition may NOT delete the DataKeeper Cluster Edition directory. This directory can be deleted manually after the **Add/Remove** operation is complete.
- A reboot of the system is required to completely remove DataKeeper Cluster Edition remnants.

Upgrading SIOS DataKeeper Cluster Edition

Upgrading DataKeeper Cluster Edition from previous versions of DataKeeper Cluster Edition is very straightforward. Simply run the installation process outlined below on all systems. The upgrade process stops the DataKeeper Service, copies the new files to the DataKeeper directory and requires a reboot at the end to load the new DataKeeper driver.

The following information applies to a DataKeeper Cluster Edition upgrade:

- Existing mirrors are not affected by the upgrade and will remain in place.
- It is not necessary to pause or manipulate the mirror(s) in any way before upgrading.
- DataKeeper Cluster Edition licensing is not affected and does not need to be redone after the upgrade.
- **IMPORTANT: Before rebooting a WSFC node that is currently the owner of DataKeeper volume resources which are online, it is recommended that all DataKeeper volume resources either be taken offline or moved to a different node/cluster owner.**

The DataKeeper Cluster Edition upgrade will be performed on the target systems first. Cluster resources will then be switched over to allow for the upgrade of the original source system.

Upgrading the Target Server

1. Using Microsoft Cluster Manager, move all resources to one node/cluster owner so that only one node is the source server.
2. Close the DataKeeper UI if it is currently running.
3. On each target system, run the setup.exe program distributed with the DataKeeper Cluster Edition product. Setup will detect that you are upgrading your existing DataKeeper product and display a confirmation dialog. Click **Yes** to continue the upgrade.
4. The DataKeeper service will be stopped during the upgrade process. After setup completes, you will be prompted to enter a new DataKeeper license key. When upgrading from a previous version of DataKeeper Cluster Edition, applying new licenses is not required and you may exit the License Manager.

5. Reboot your server.
6. Bring up the target systems and allow the mirror(s) to resync and return to the **Mirroring** state.
7. Repeat Steps 2-6 for each target system.

Upgrading the Original Source Server

1. Using Microsoft Cluster Manager, move all resources to a DataKeeper node that has been upgraded so the source server can be upgraded.
2. Once all resources are online on another node and in the mirroring state, repeat the above procedure on the original source system and reboot the server.
3. Run the DataKeeper UI to view your existing mirrors.

Reinstalling SIOS DataKeeper Cluster Edition

To reinstall DataKeeper Cluster Edition, 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 DataKeeper Cluster Edition 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.

Considerations

Upgrading from a previous release to DataKeeper v7.6 or later [read the considerations](#) for chkdisk.

DataKeeper Cluster Edition Technical Documentation

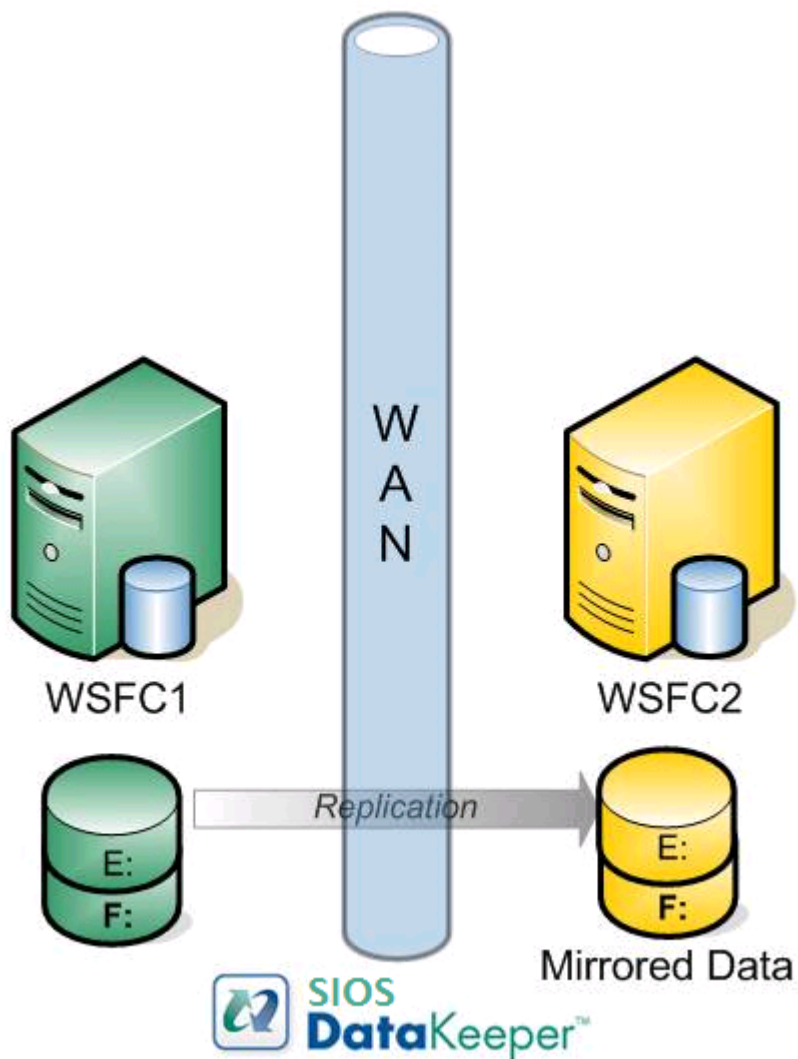
DataKeeper with High Availability – Cluster Edition Overview

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.

SIOS DataKeeper Cluster Edition is a highly optimized host-based replication solution which integrates seamlessly with Windows Failover Clustering. Windows Failover Clustering features, such as cross-subnet failover and tunable heartbeat parameters, make it possible for administrators to deploy geographically dispersed clusters. SIOS DataKeeper provides the data replication mechanism which extends Windows Server Failover Clustering, allowing administrators to take advantage of these advanced features to support high availability and disaster recovery configurations.

SIOS DataKeeper Cluster Edition is a separately licensed product. Once installed, a new storage resource type called **DataKeeper Volume** is available in Microsoft Windows Server Failover Clustering. This new SIOS DataKeeper Volume resource can be used in place of the traditional Physical Disk shared storage resource to enable geographically dispersed clusters.

Important Consideration: Prior to installing SIOS DataKeeper Cluster Edition, your Microsoft Windows Server Failover Cluster environment should be installed and created. This product requires a SIOS DataKeeper Cluster Edition License. SIOS DataKeeper resource type registration occurs 60 seconds after detecting a Failover Cluster configuration.



Features

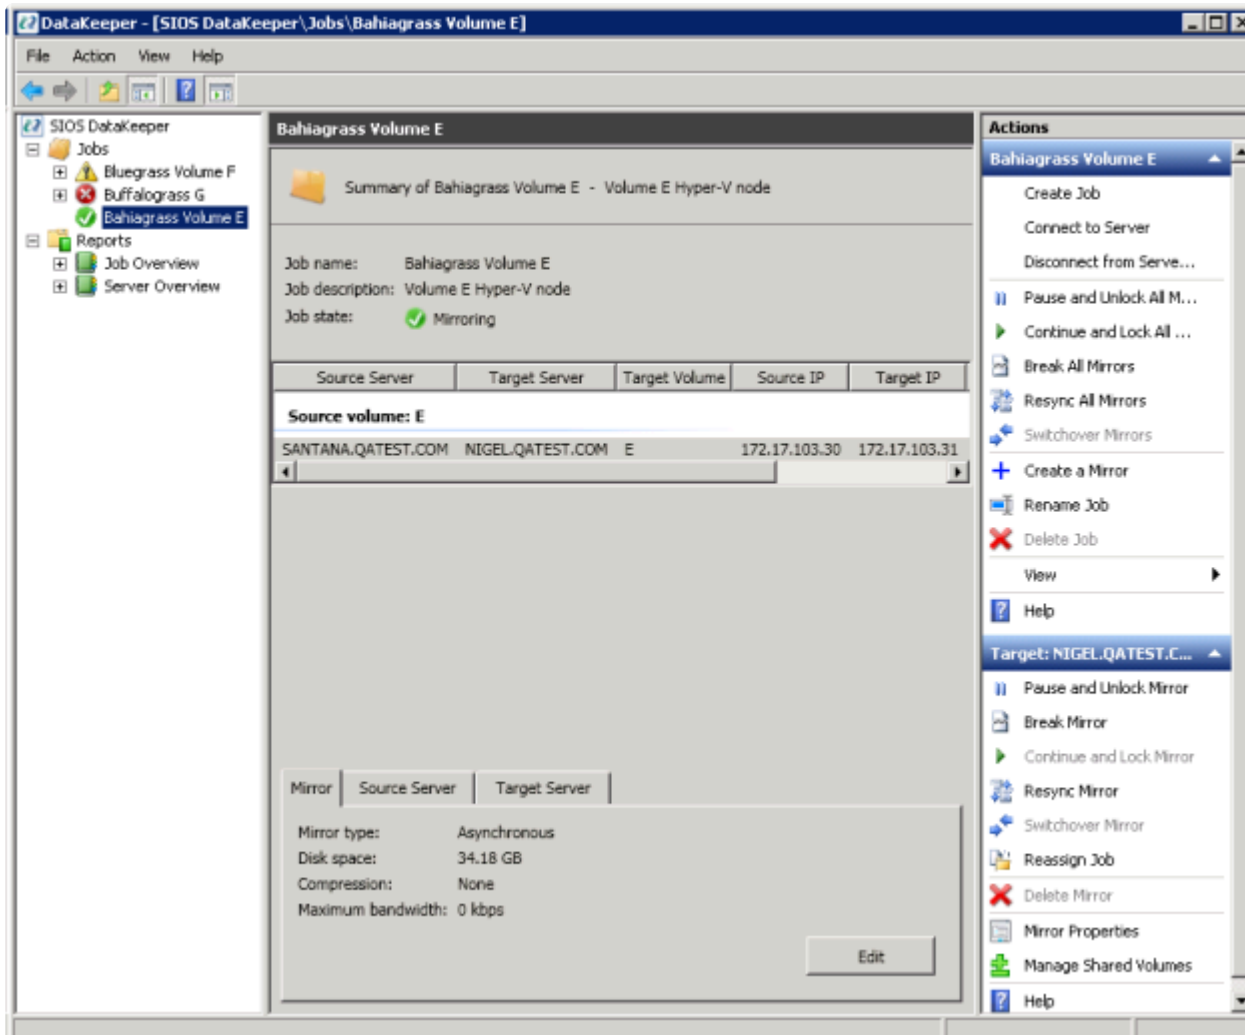
Some of the features include the following:

- 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.

Components

DataKeeper Components

SIOS DataKeeper for Windows is comprised of the following components:

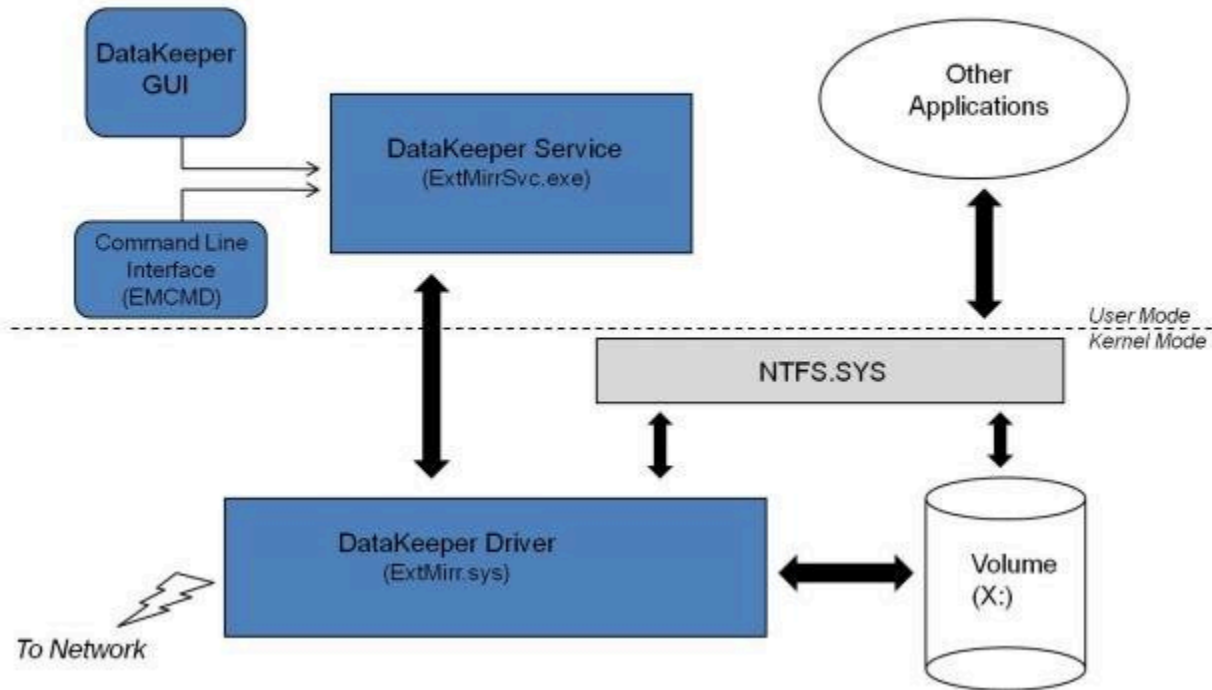
- **DataKeeper Driver (ExtMirr.sys)** – The DataKeeper Driver is a kernel mode driver and is responsible for all mirroring activity between the mirror endpoints.
- **DataKeeper Service (ExtMirrSvc.exe)** – The DataKeeper Service links the DataKeeper GUI and Command Line Interface to the DataKeeper Driver. 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



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:

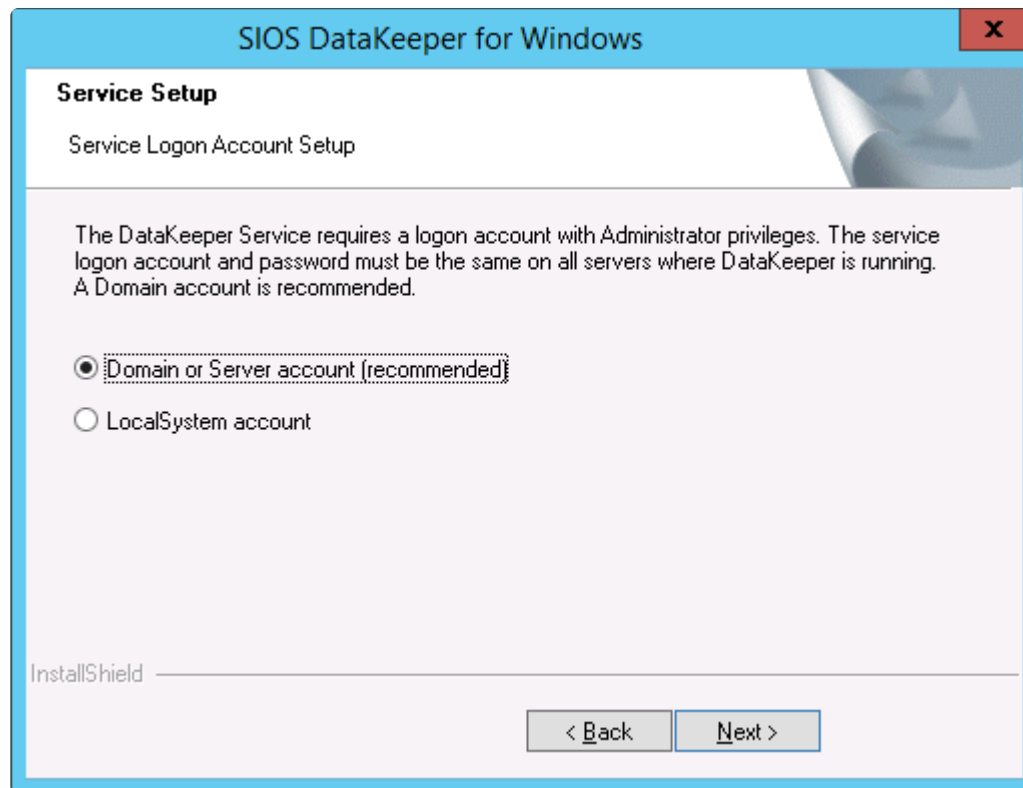
- A **Domain Account** with administrator privileges, valid on all connected servers in the domain (*recommended*)
- A **Server Account** with administrator privileges, valid on all connected servers
- The **Local System Account** (*not recommended*)

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](#)).

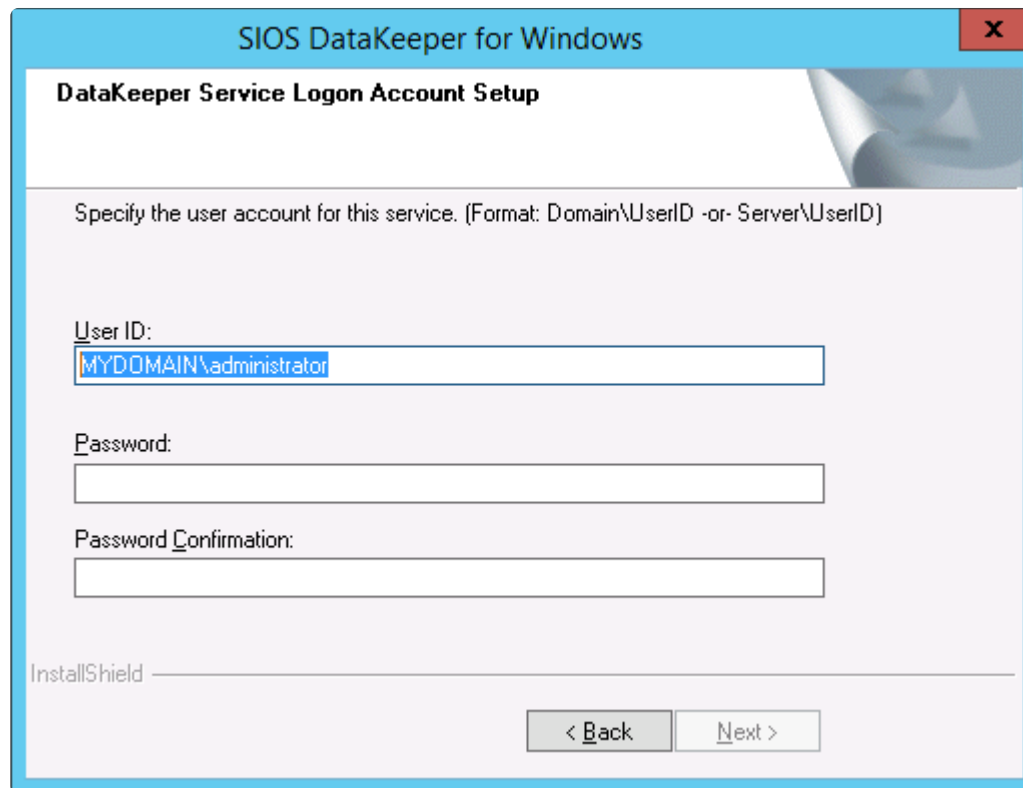
Note: The domain or server account used must be added to the Local System Administrators Group. The account must have administrator privileges on each server that DataKeeper is installed on.

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:



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 red close button. The main heading inside 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

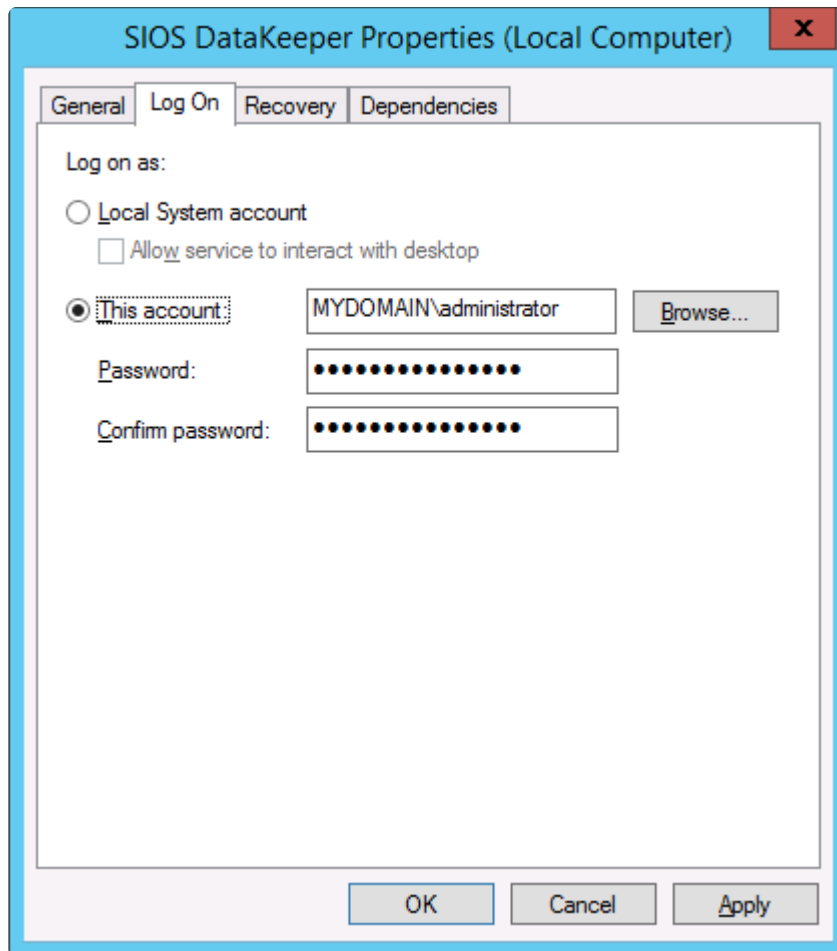
Password:
●●●●●●●●

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 GUI Or 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>
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 • Add this local account to the Administrator Group • Run the DK Service on all systems with this local administrator account 	<ul style="list-style-type: none"> • Log in using the local administrator account you created to run the DK Service • Run the DK GUI

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.

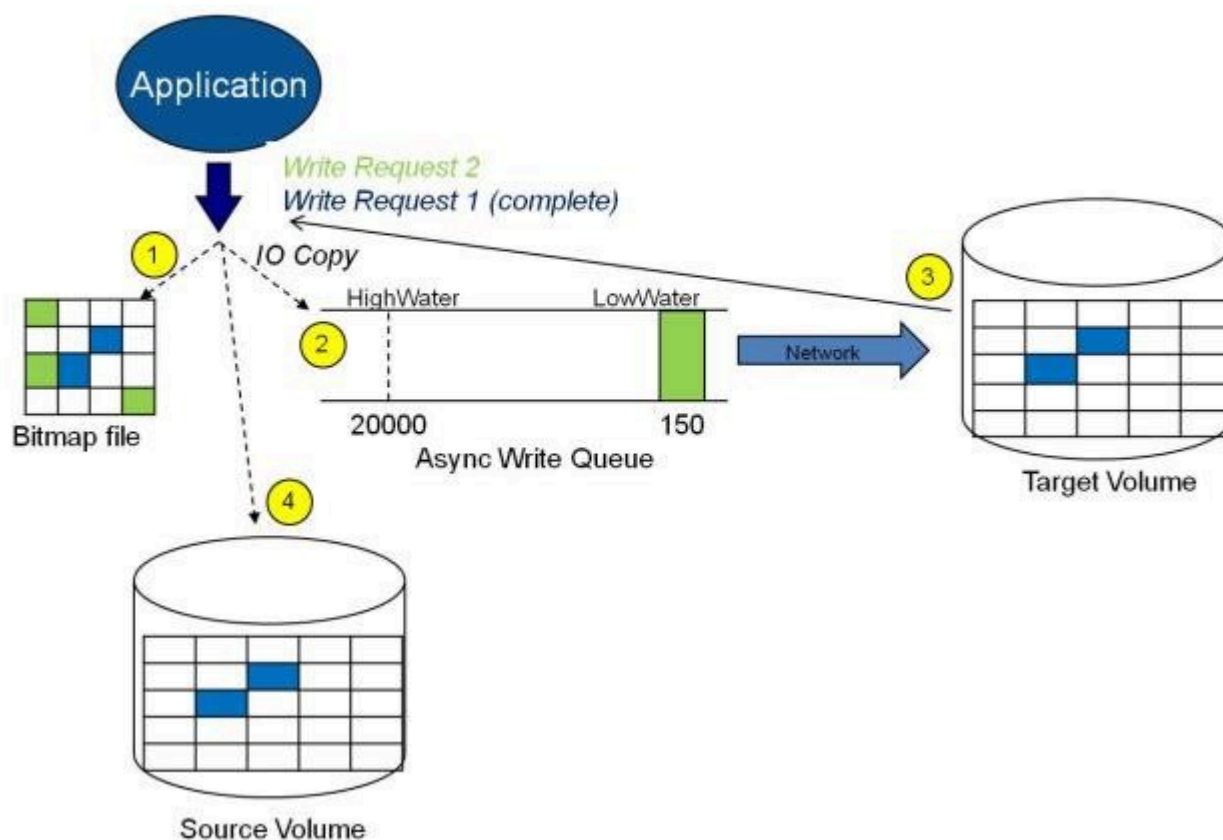
Synchronous Mirroring

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.

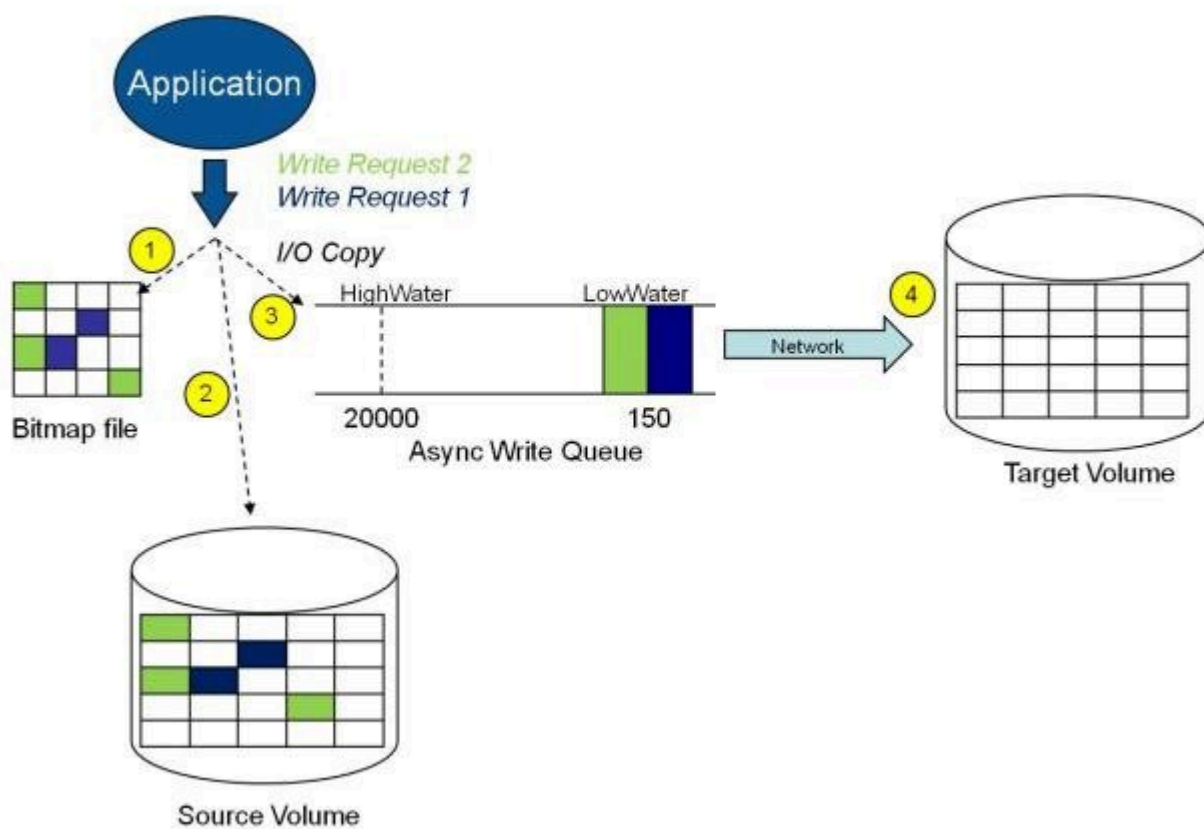
The following sequence of events describes what happens when a write request is made to the source volume of an asynchronous mirror.

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

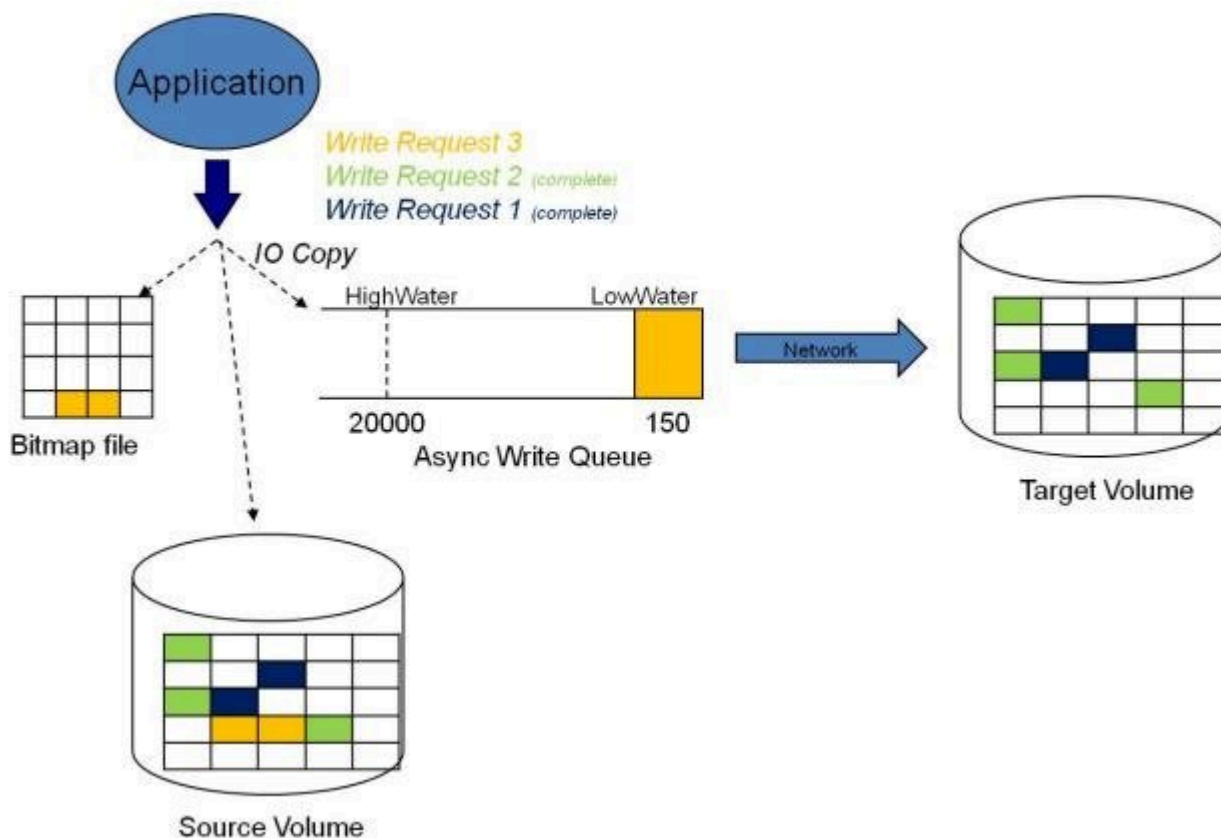


In the diagram above, the two write requests have been written to the source volume and are in the queue to be sent to the target system. However, control has already returned back to the application who initiated the writes.

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.

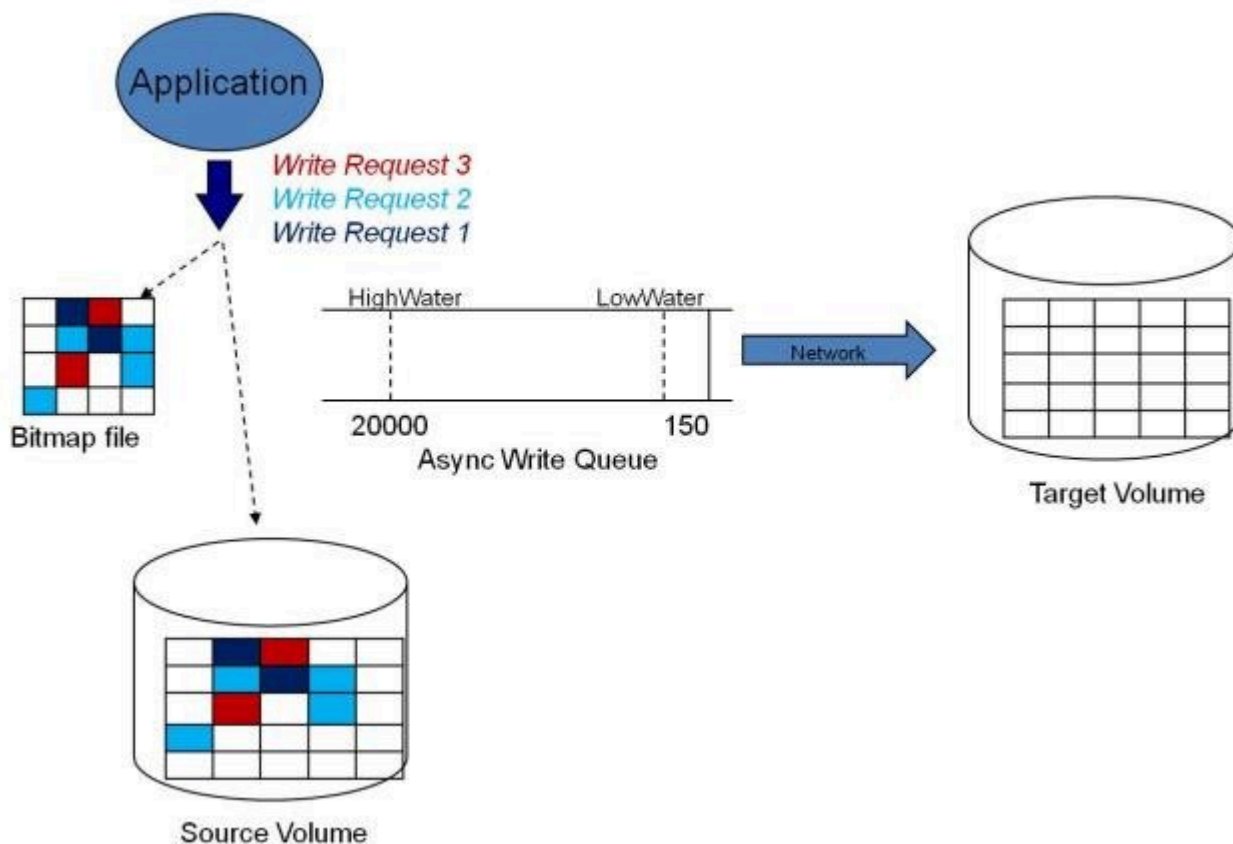
Asynchronous Replication: Mirroring



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.

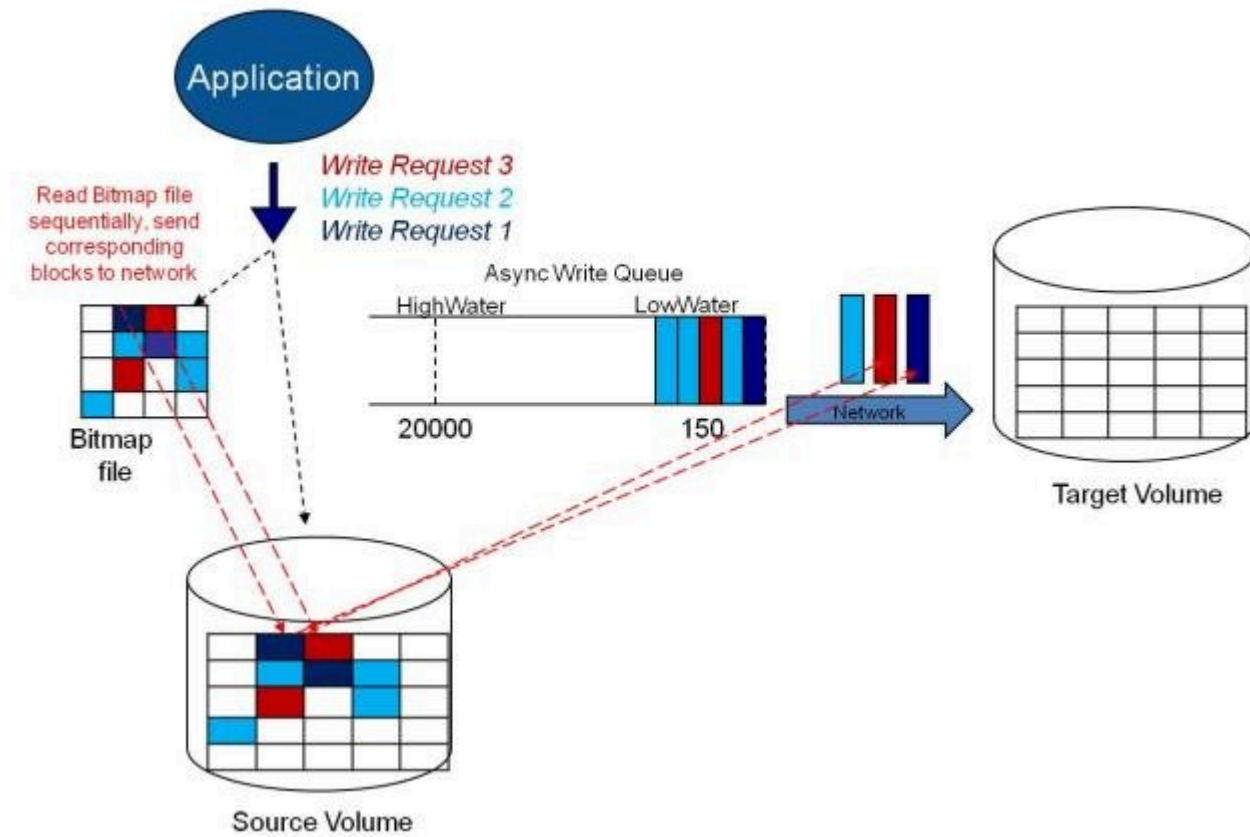
DataKeeper reads sequentially through the persistent bitmap file to determine what blocks have changed on the source volume while the mirror was **PAUSED** and then resynchronizes only those blocks to the target volume. This procedure is known as a partial resync of the data.

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.

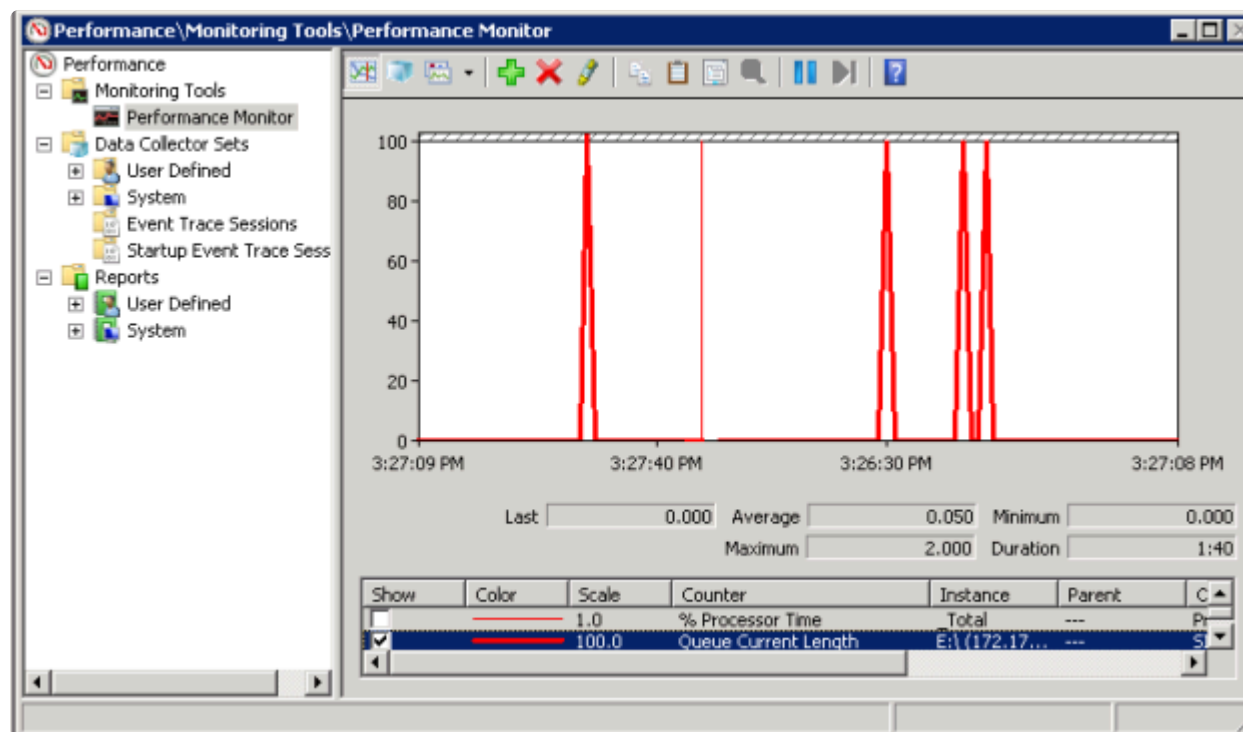
On a **Microsoft Windows 2012** system, start the **Windows Performance Monitor** through the **Performance Monitor** option in the **Administrative tools**.

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) and DataKeeper Cluster Edition (DKCE) 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.

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](#)

[WAN Considerations](#)


[Initial Synchronization of Data Across the LAN/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*)

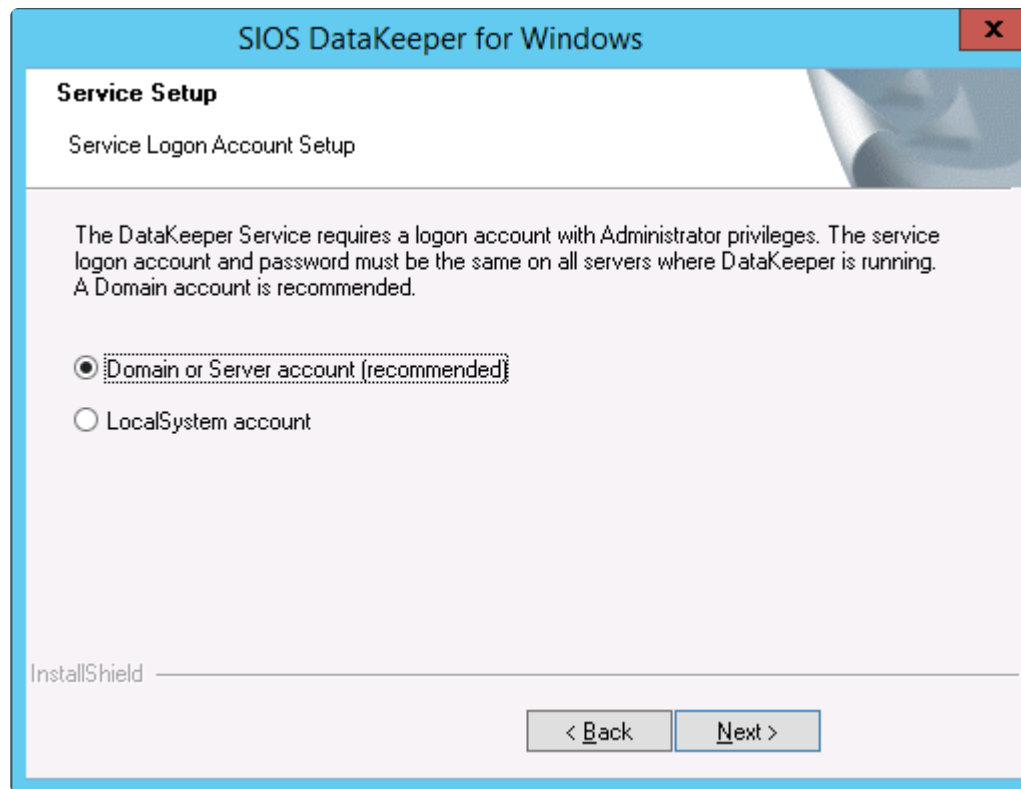
Note: For Workgroups, use the **Server Account** option and use the server name \ administrator on each system as the Service Account for DataKeeper. **You should also log on to all servers using this same Log On ID and Password** (see related [Known Issue](#)).

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.

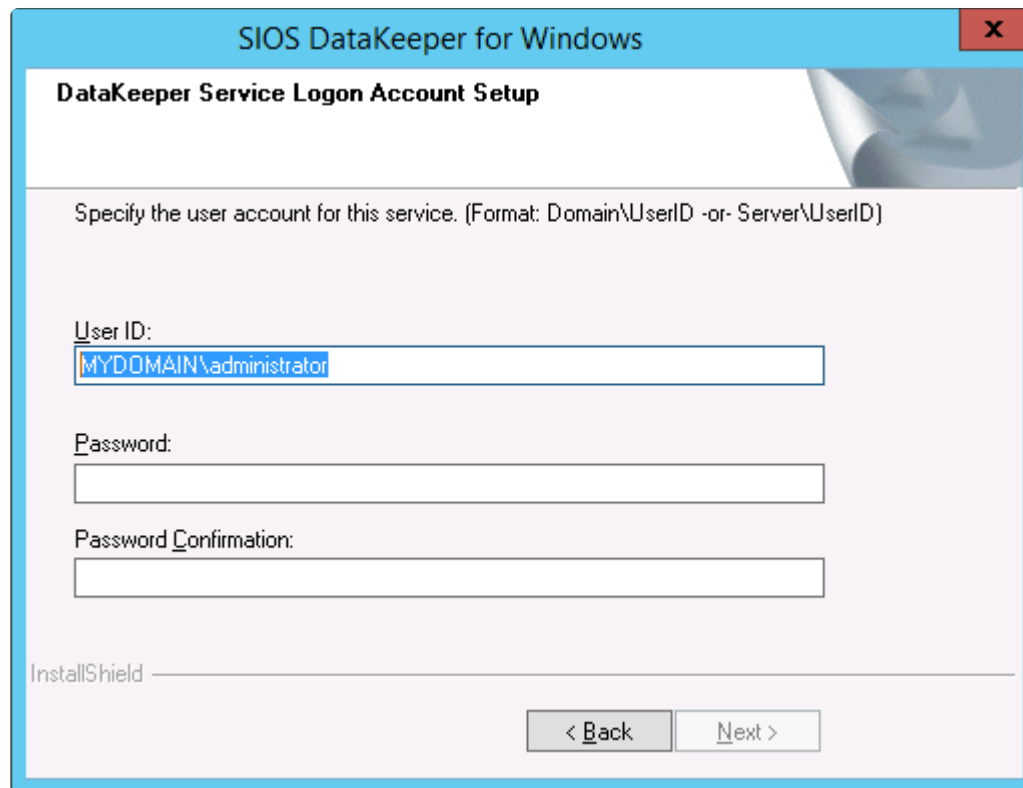
Please note that the Local System account cannot be authenticated properly in a domain when network connectivity with Active Directory is lost. In that situation, connections between servers cannot be established with the Local System account causing DataKeeper volume switchover commands, via the network, 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:

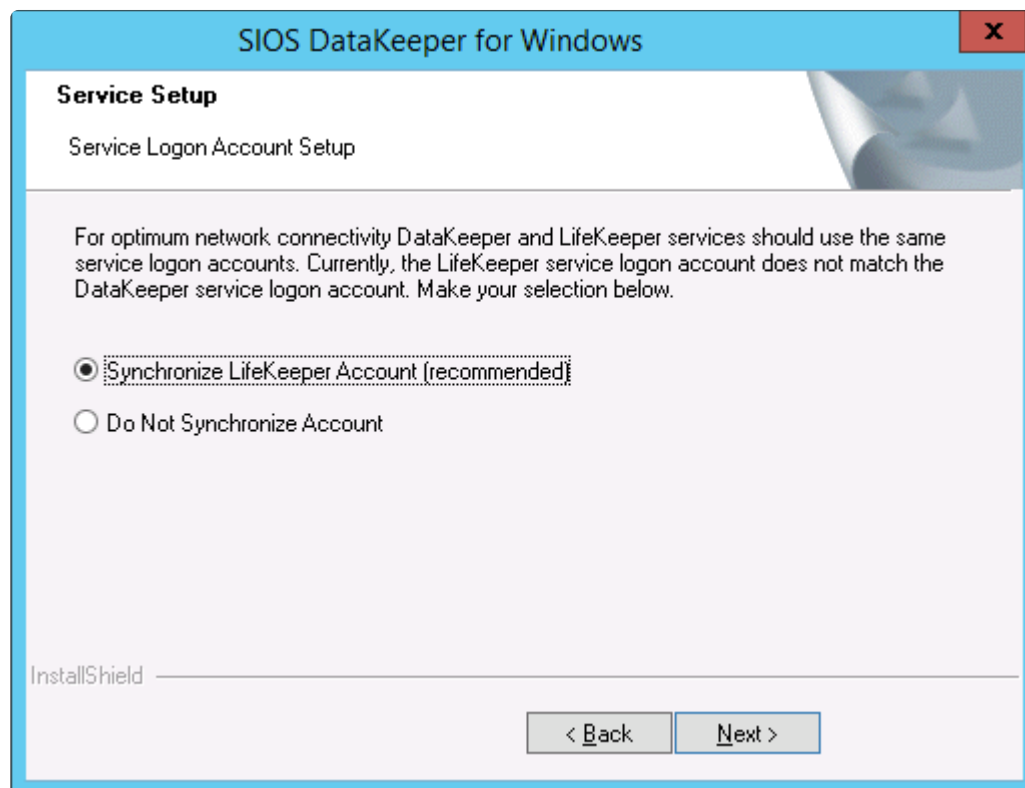


If a Domain or Server account is selected above, the DataKeeper Service Log On ID and Password Entry Form is displayed to enter that information.



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 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: "User ID:" containing "MYDOMAIN\administrator", "Password:" which is empty, and "Password Confirmation:" which 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.

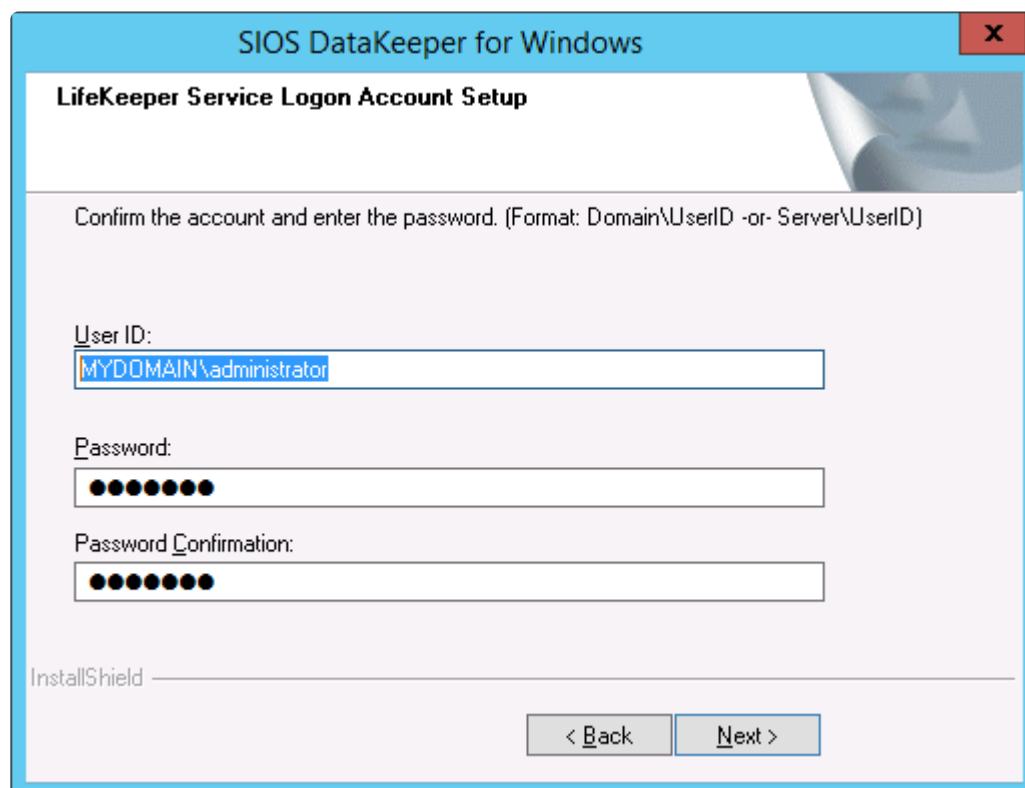
☒ Synchronize LifeKeeper Account (recommended)

☐ 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

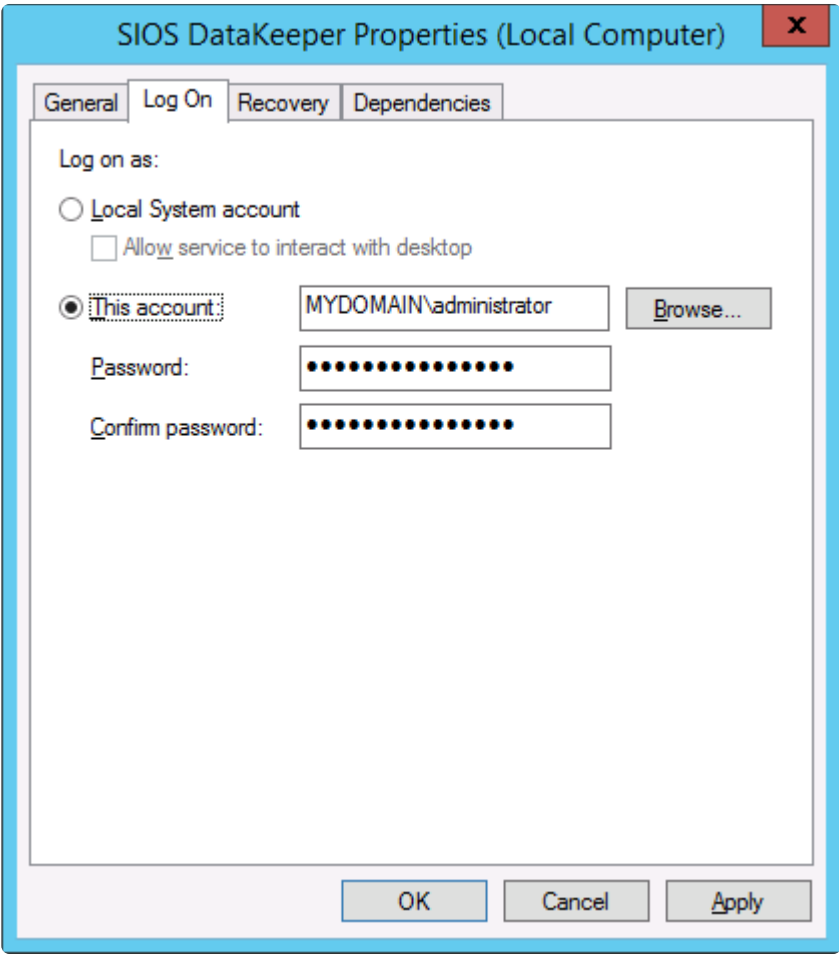
Password:
●●●●●●●●

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	<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.

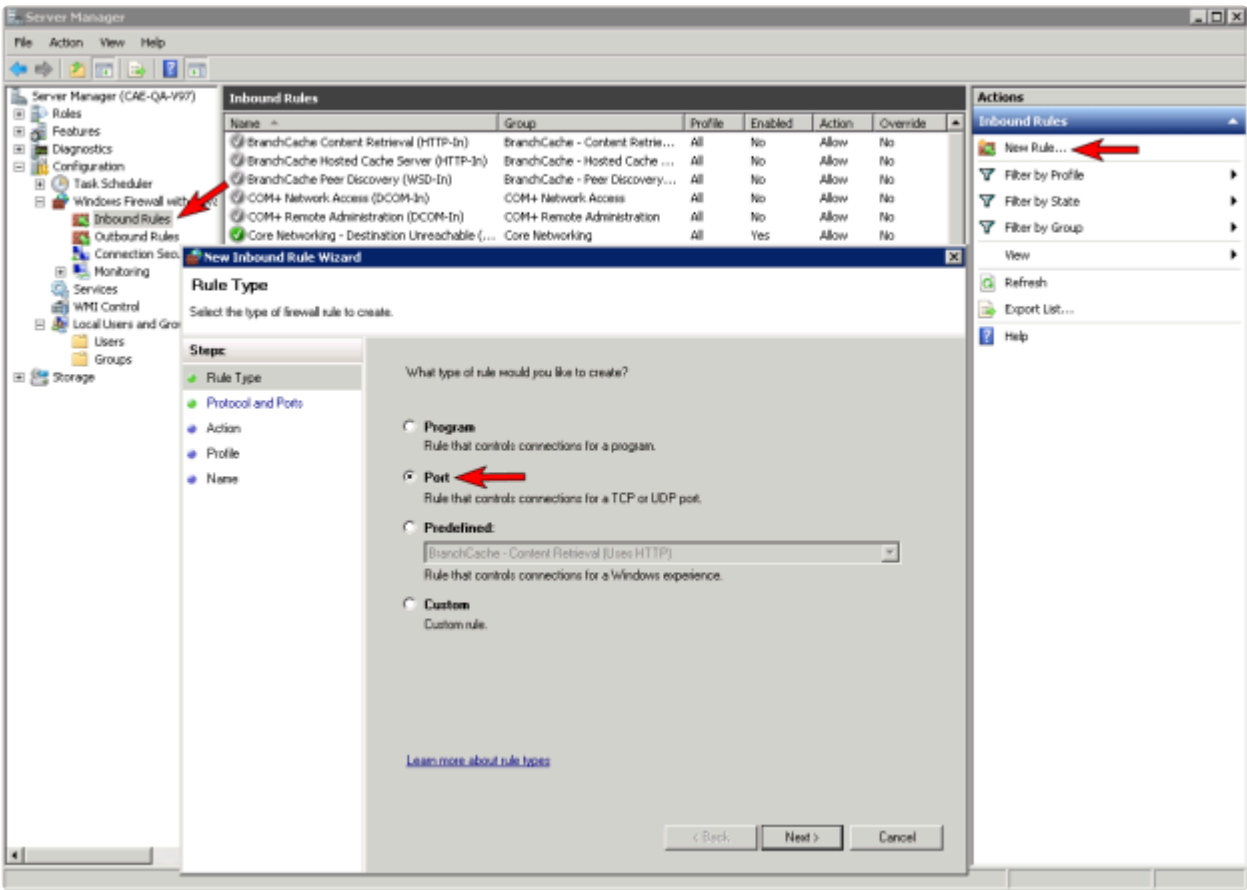
Port #	Volume Letter	Port #	Volume Letter
10000	A	10013	N
10001	B	10014	O
10002	C	10015	P
10003	D	10016	Q
10004	E	10017	R
10005	F	10018	S
10006	G	10019	T
10007	H	10020	U
10008	I	10021	V
10009	J	10022	W
10010	K	10023	X
10011	L	10024	Y

10012	M	10025	Z
-------	---	-------	---

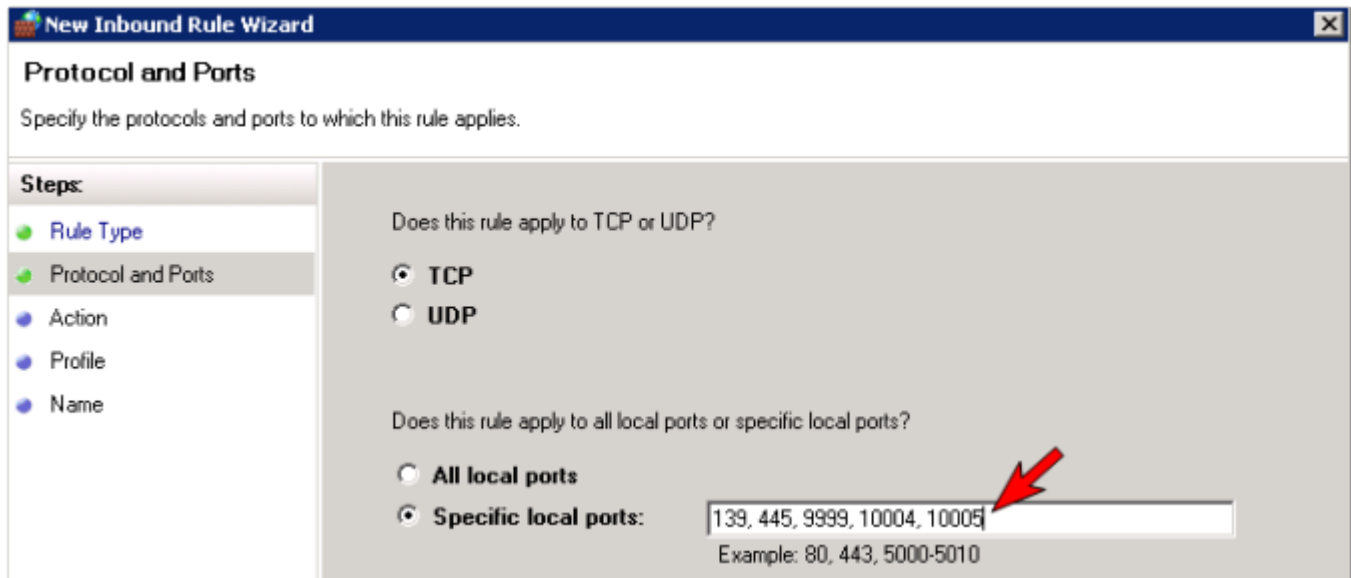
Configuring Microsoft's Windows Firewall with Advanced Security - Example

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**.



New Inbound Rule Wizard

Protocol and Ports

Specify the protocols and ports to which this rule applies.

Steps:

- Rule Type
- Protocol and Ports**
- Action
- Profile
- Name


Does this rule apply to TCP or UDP?

☒ **TCP**

☐ **UDP**

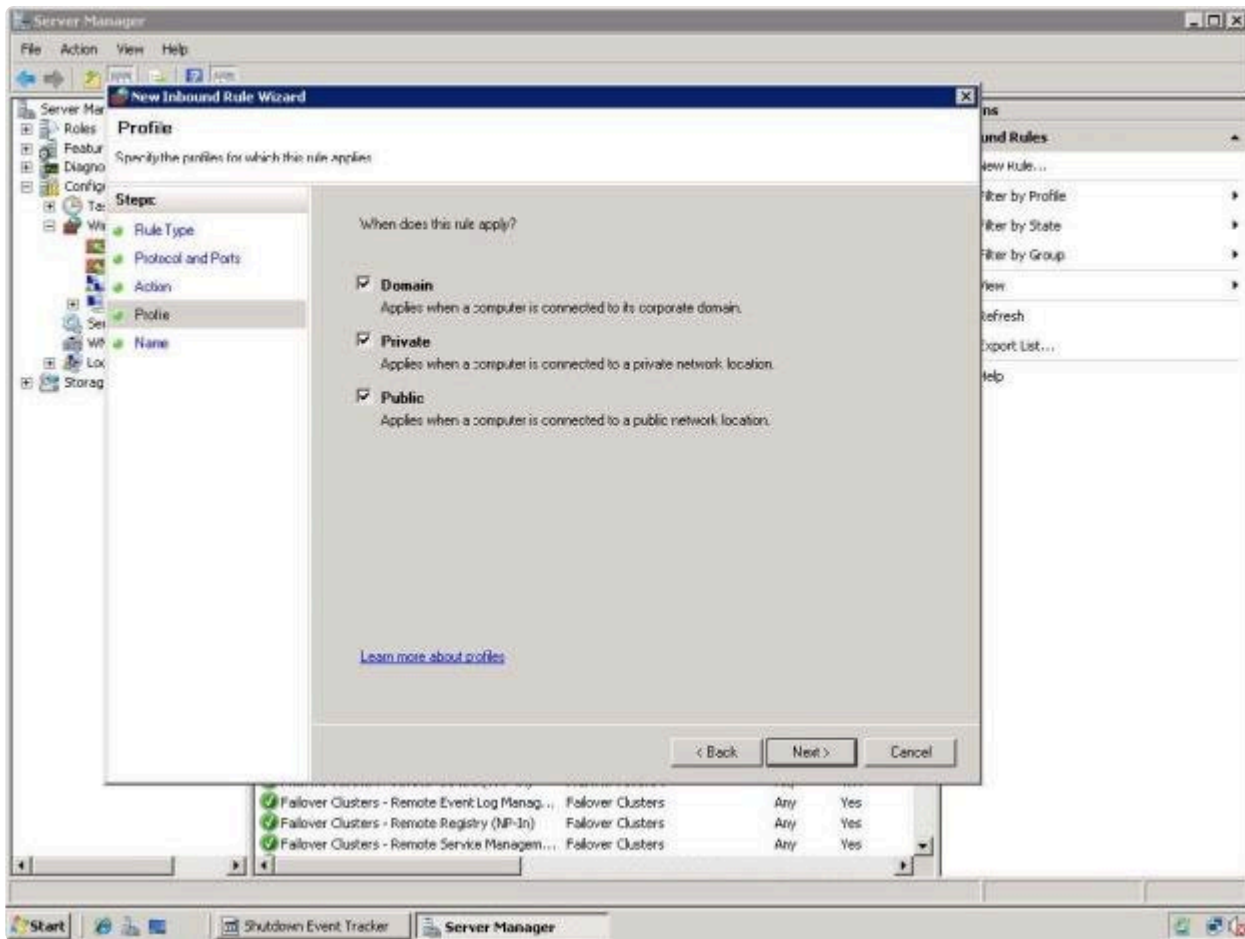
Does this rule apply to all local ports or specific local ports?

☐ **All local ports**

☒ **Specific local ports:** 

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**

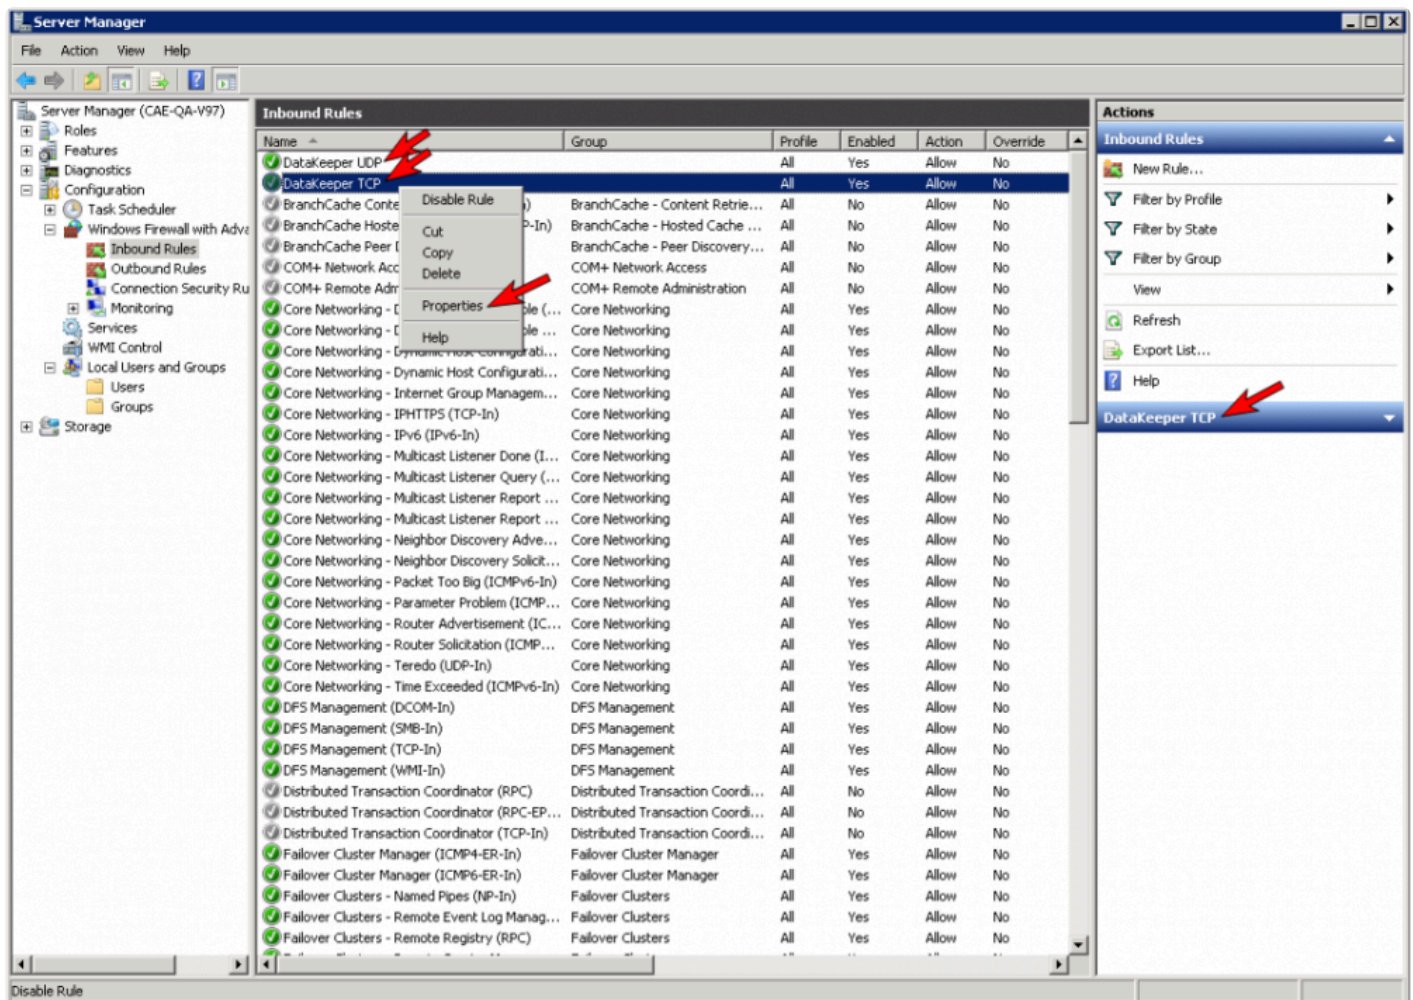
Name:
DataKeeperTCP

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.

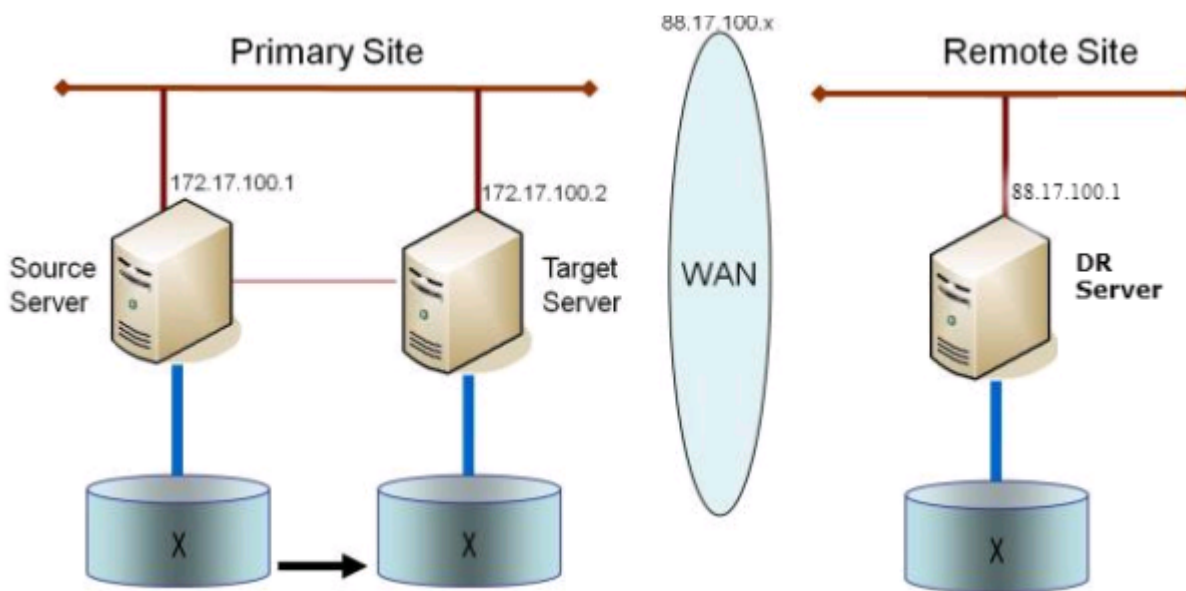
Handling Unmanaged Shutdown Issues

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"

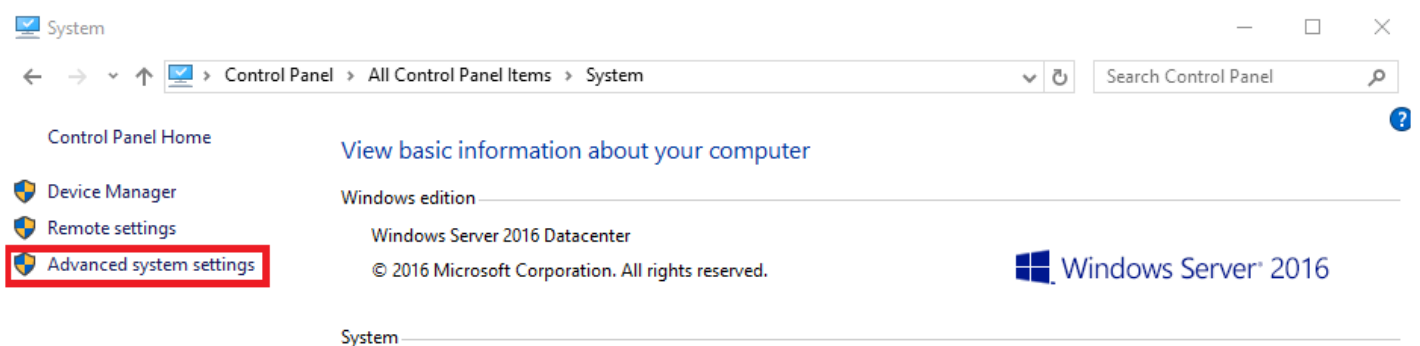
By default, Windows configures virtual memory so that page files are automatically created on volumes as the Operating System determines is best. This Virtual Memory setting is called "Automatically manage paging file size for all drives".

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.

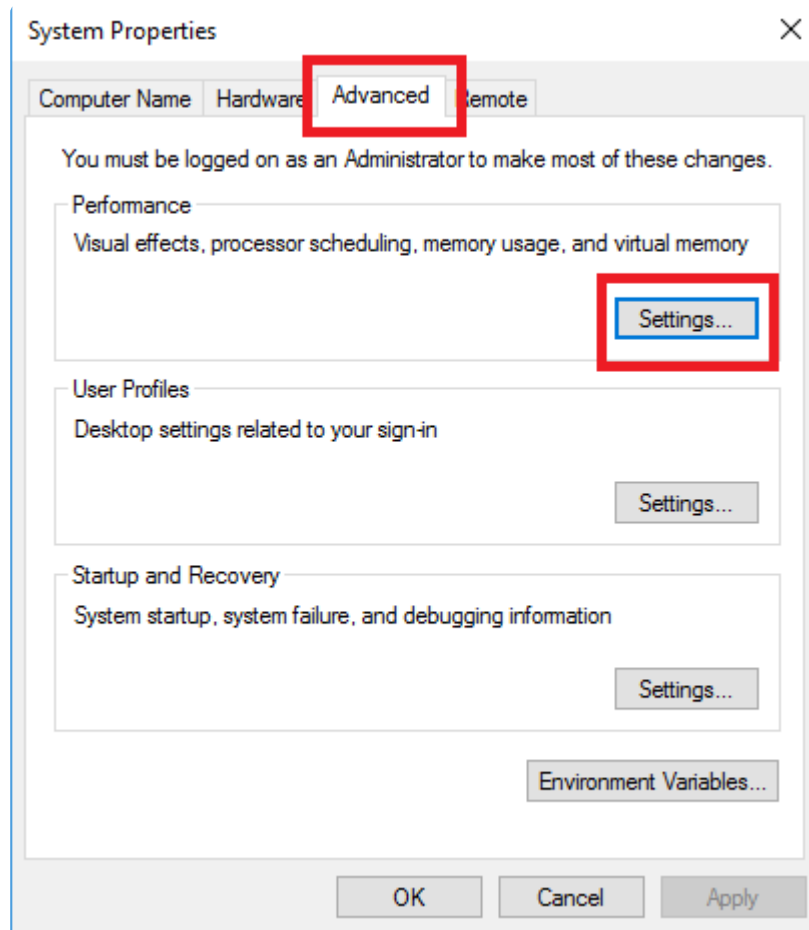
How to disable "Automatically manage paging file size for all drives"

This setting can be found in Control Panel "System" dialog.

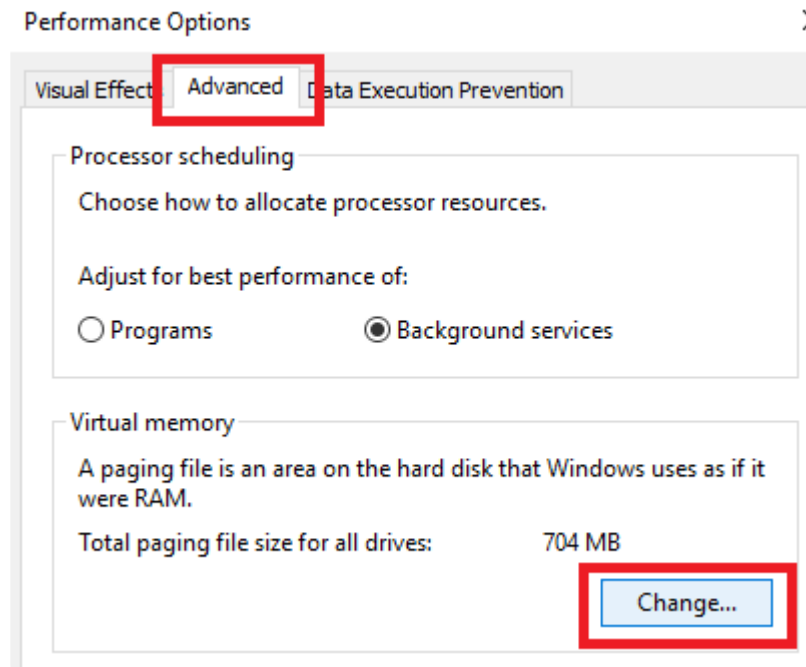
First, click the **Advanced system settings** option:



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.

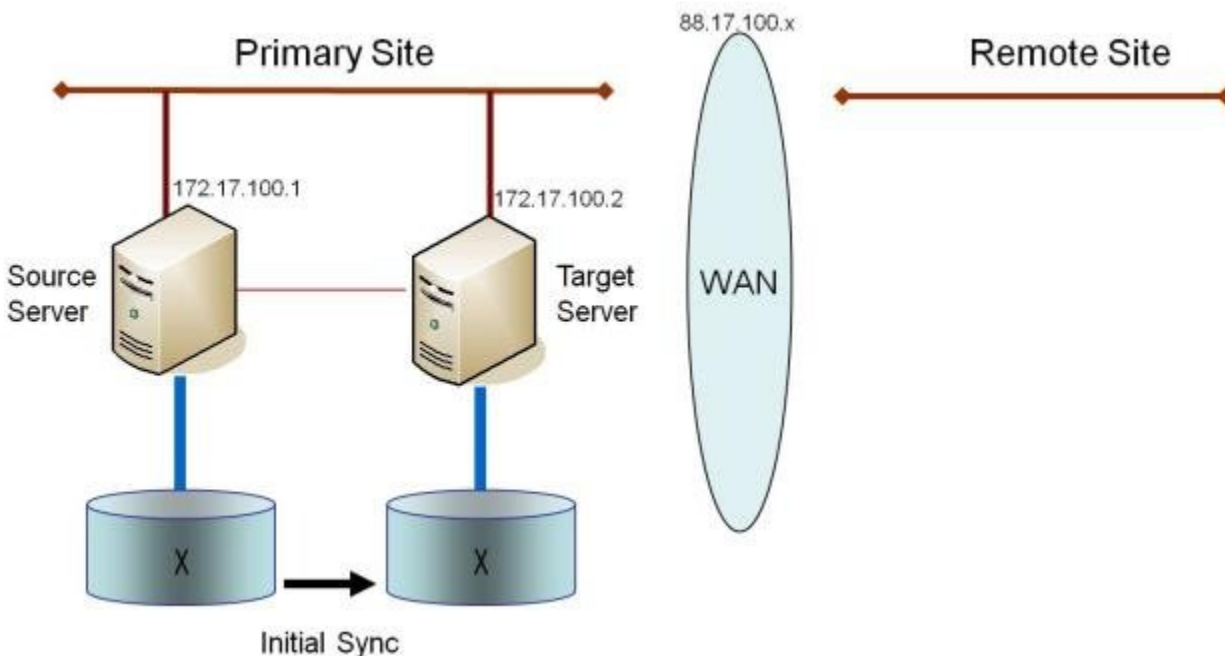
✳ **Note:** This command supports changing the endpoints of a mirrored volume that is configured on 3 nodes or fewer. For configurations greater than three nodes, create mirrors with the final endpoint at the local site and use route adds to get the mirrors created and resynced before moving the server 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:

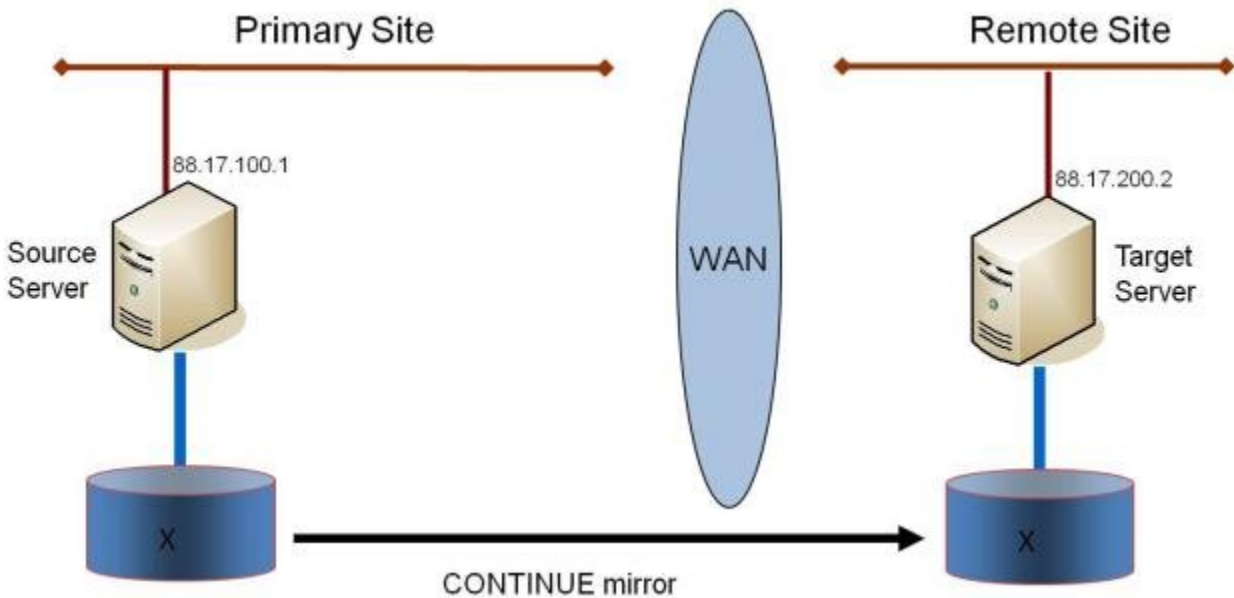
```
cd %EXTMIRRBASE%
```

- 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](#)

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](#)

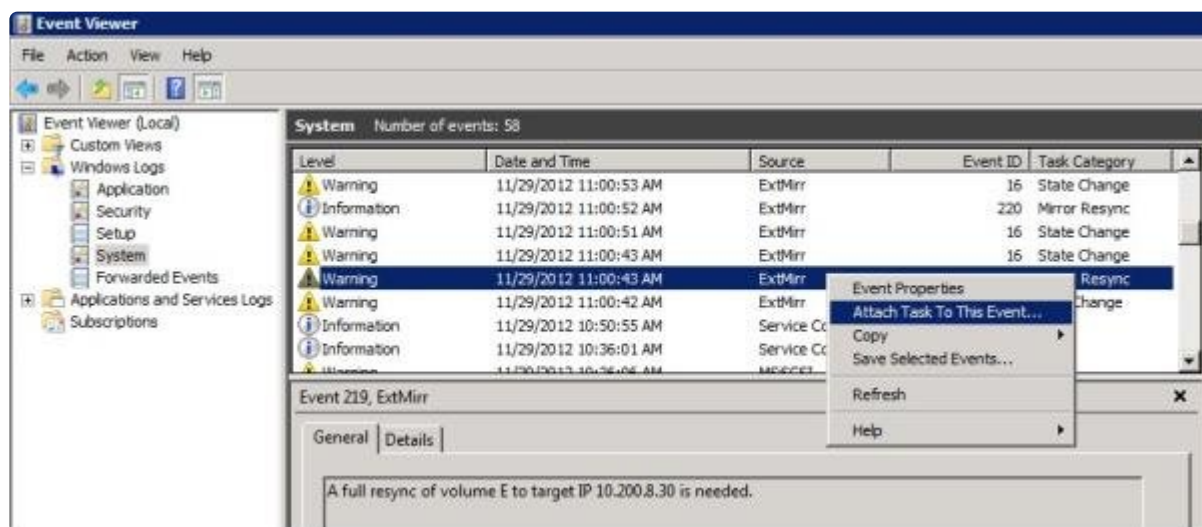
DataKeeper Event Log Notification

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.

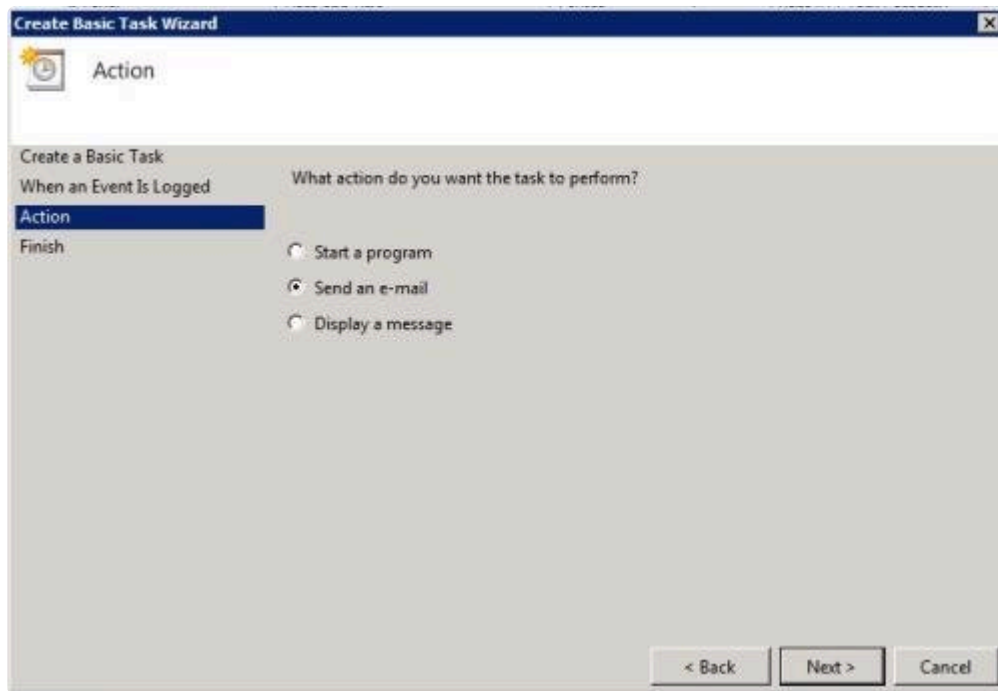
✿ **Note:** This option is only available for Windows Server 2008 R2.

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).

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.

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.

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)
<p>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.</p>		

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)
<p>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.</p> <p>Note: This tunable applies to synchronous and asynchronous mirrors.</p>		

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>
<p>Specifies the folder where the target snapshot file for this volume will be stored.</p>		

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](#)

[GETMIRRORTYPE](#)

[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

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`.

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 1x1 mirror (1 source, 1 target), only one command is required. For a 2x1 mirror (2 nodes with a shared volume with one target node) or a 1x1x1 (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.

1x1 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.

1x1x1 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 [relocked](#) 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 output is one of the following mirror states:

- 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> <ul style="list-style-type: none"> 1 – 0x01: All Net Alerts is enabled 2 – 0x02: Broken State Alert is enabled 4 – 0x04: Resync Done Alert is enabled 8 – 0x08: Failover Alert is enabled 16 – 0x10: Net Failure Alert is enabled 32 – 0x20: LifeKeeper Config is enabled 64 – 0x40: Auto Resync is enabled 128 – 0x80: MS Failover Cluster Config is enabled 256 – 0x100: Shared Volume is enabled

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 snapshotted.
<volume letter>	This is the drive letter(s) of the volume(s) to be snapshotted on the target server. If multiple volumes are to be snapshotted, 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.

UPDATECLUSTERTARGET STATEPROPERTIES

EMCMD <system> UPDATECLUSTERTARGET STATEPROPERTIES

This command will update the TargetState private properties for all clustered DataKeeper volumes for which the given system is the source.

<system>	This is the system whose volume states will be checked.
----------	---------------------------------------------------------

For more information on the TargetState properties, please see [DataKeeper Volume Resource Private Properties](#).

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 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
```

User Guide

The topics in this section are designed to be a reference for you as you get started using SIOS DataKeeper Cluster Edition, helping you identify the type of configuration you are interested in implementing and providing detailed instructions for effectively using your SIOS DataKeeper Cluster Edition software.

[Getting Started](#)

[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](#)

[Clustering](#)

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](#)

[Setting Up SIOS DataKeeper](#)

[Connecting to a Server](#)

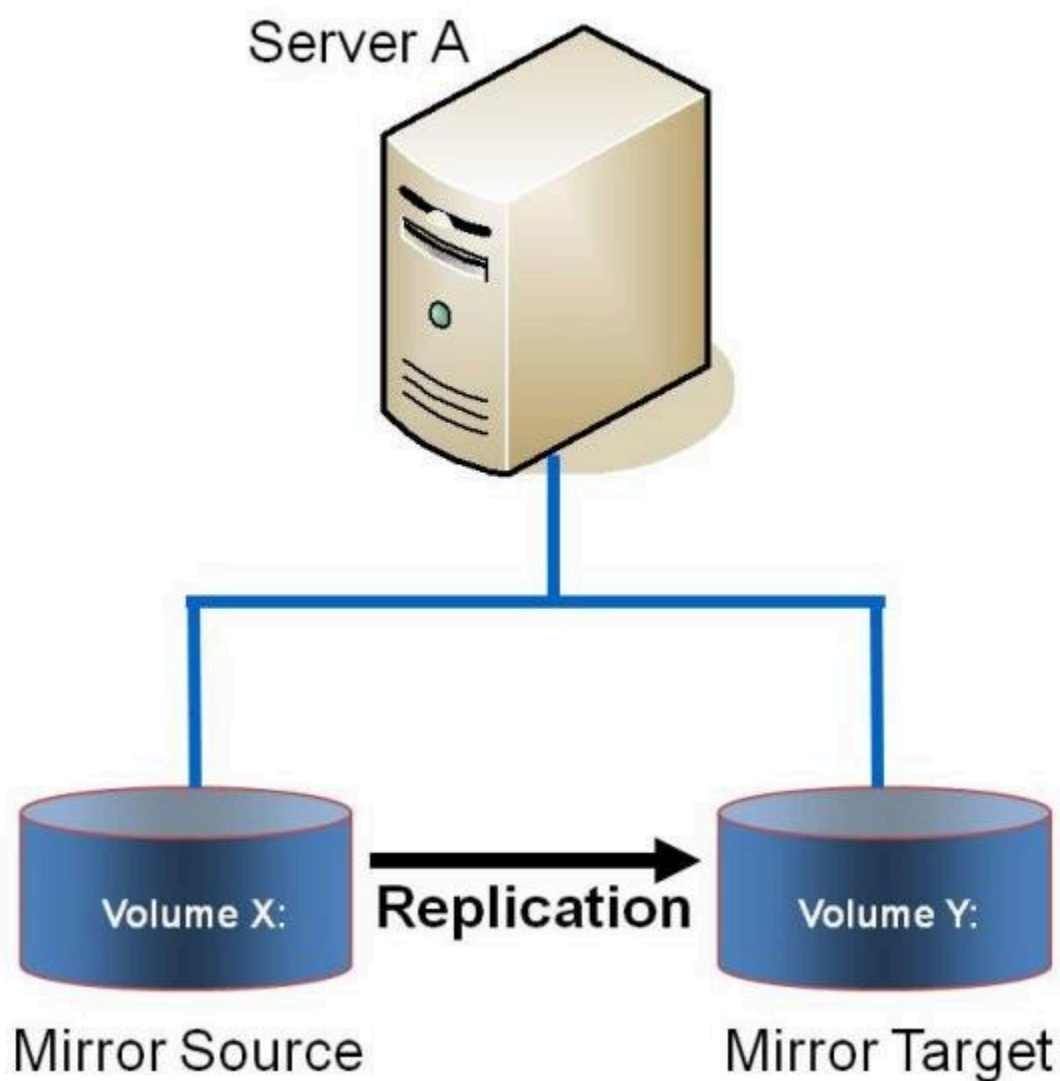
[Disconnecting from a Server](#)

[Creating a Job](#)

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

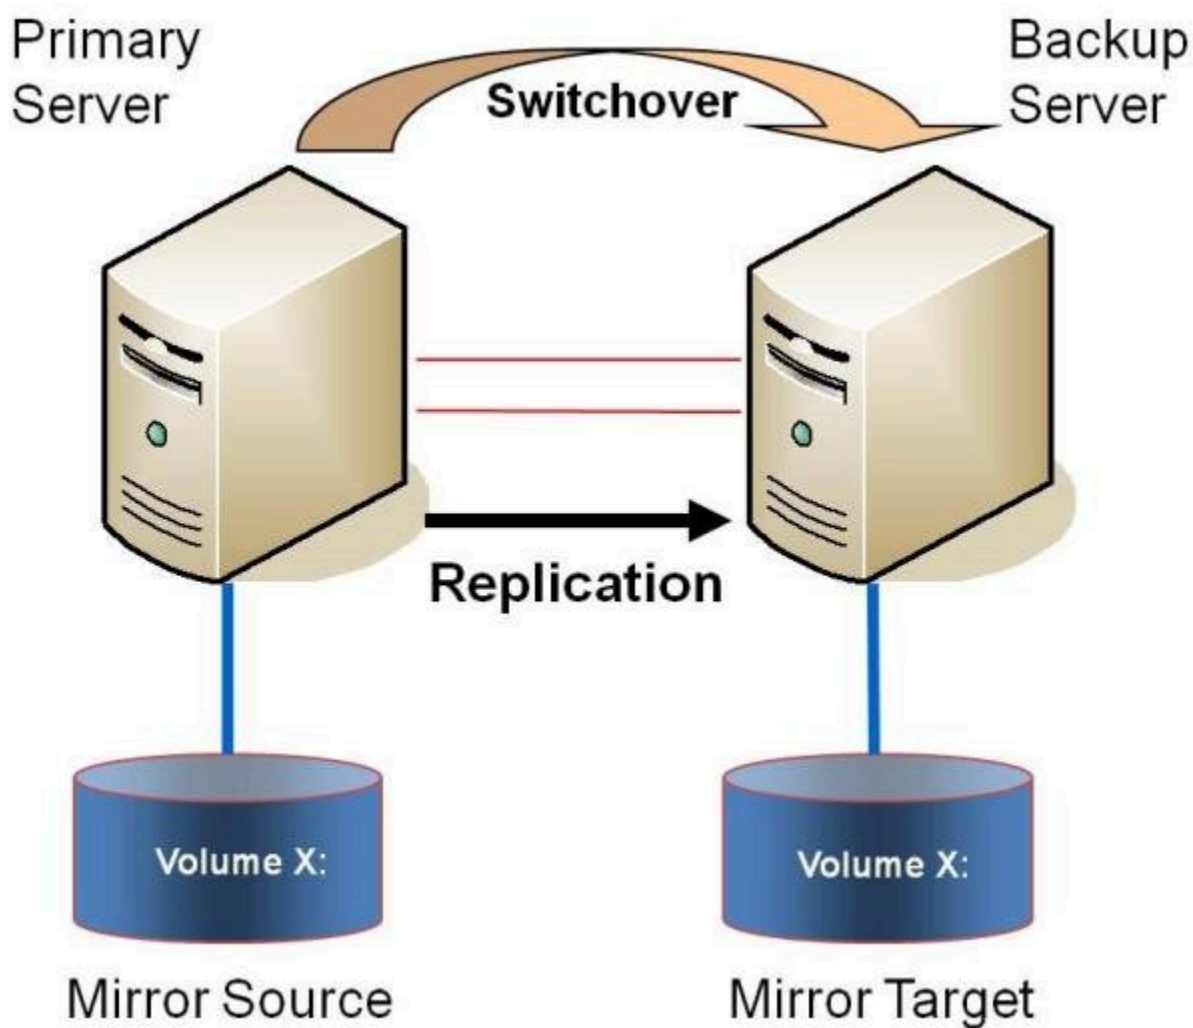
server. These disks can be different storage arrays, protecting against data loss should the primary SAN fail.

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.



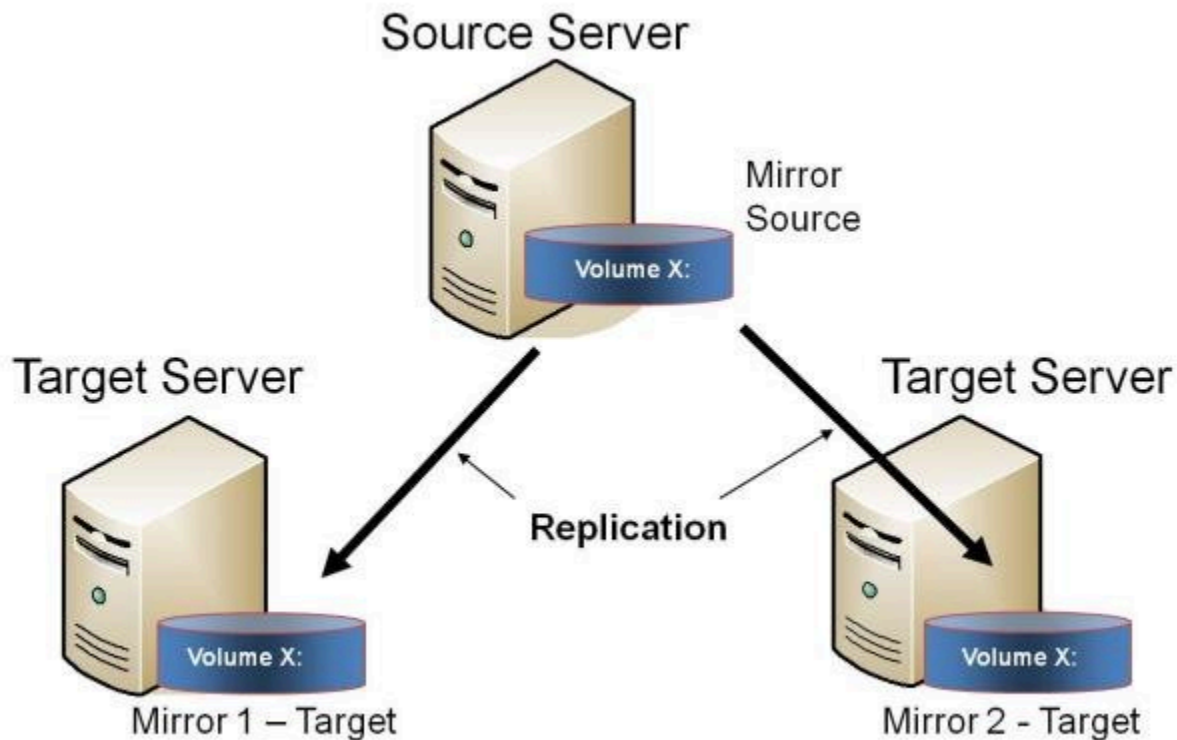
Example:	Replicate data on one or more volumes from a server in one city to another server in another city.
USE CASE	

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).

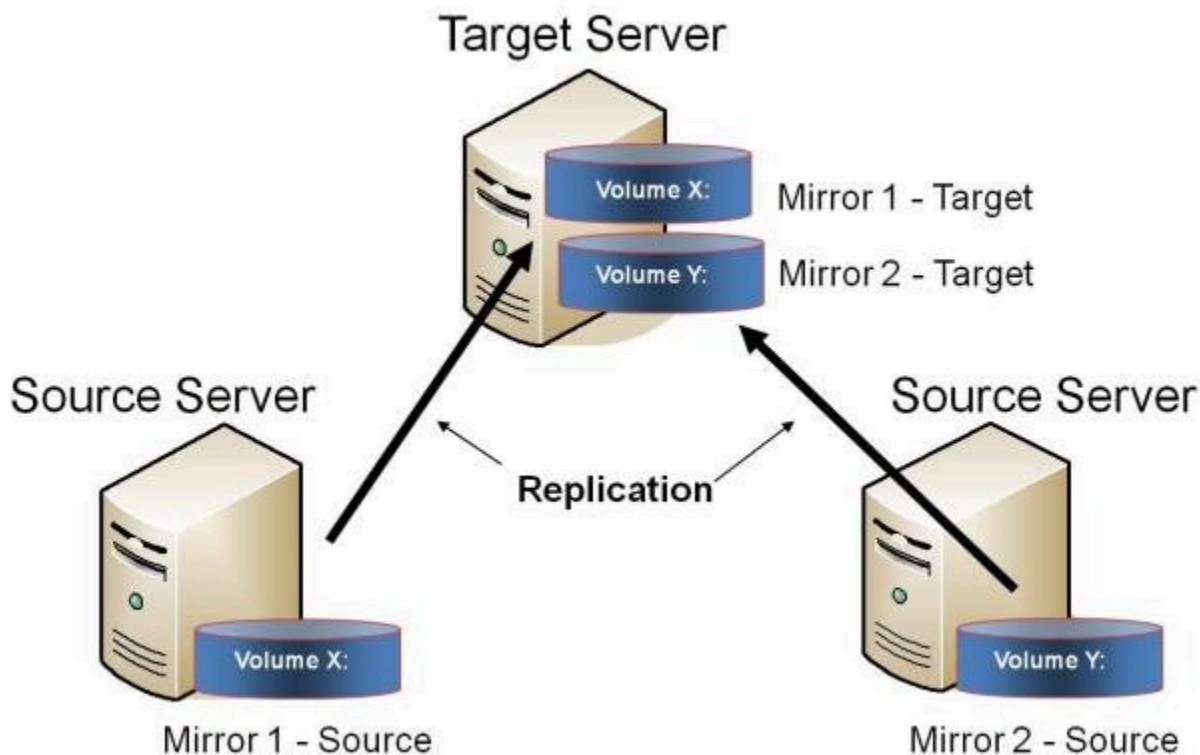
Example: USE CASE	Replicate data to one target server that resides locally in the same site with the primary server and replicate another copy of the data to a remote site for disaster recovery purposes should something happen to the first site.
Example: USE CASE	To periodically replicate or “push” data 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:

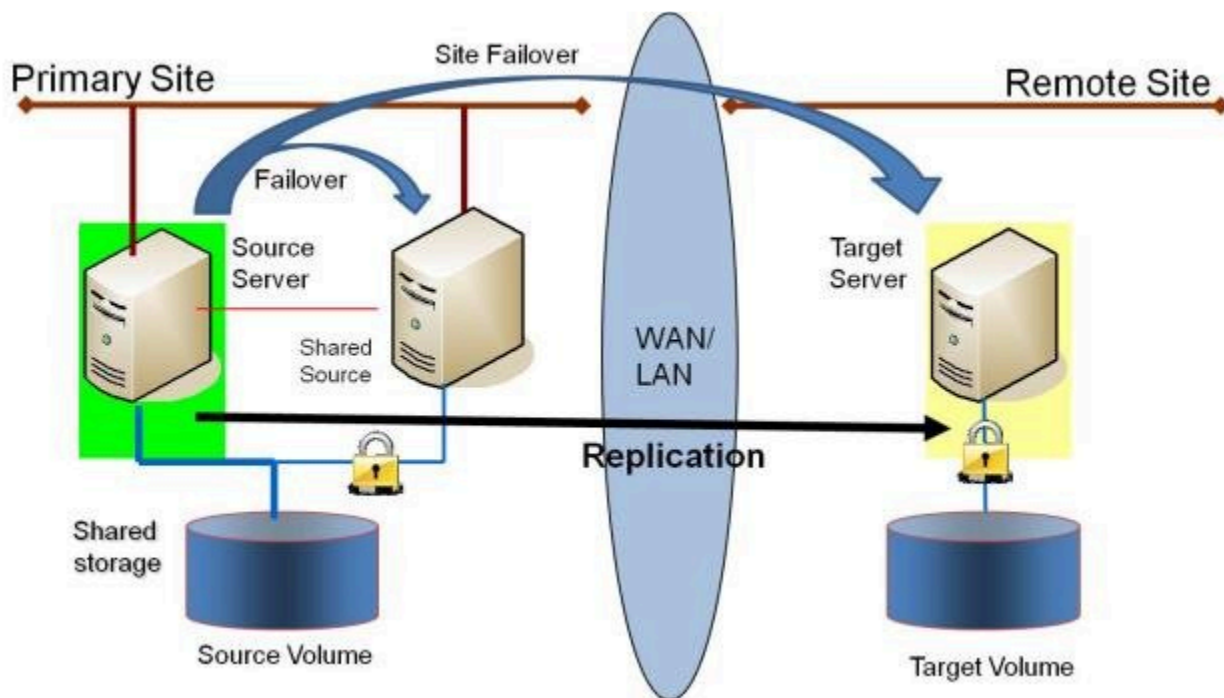
- [Primary Server Shutdown](#)
- [Secondary Server Failures](#)
- [Using DataKeeper Standard To Provide Disaster Recovery For Hyper-V Virtual](#)

[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](#)

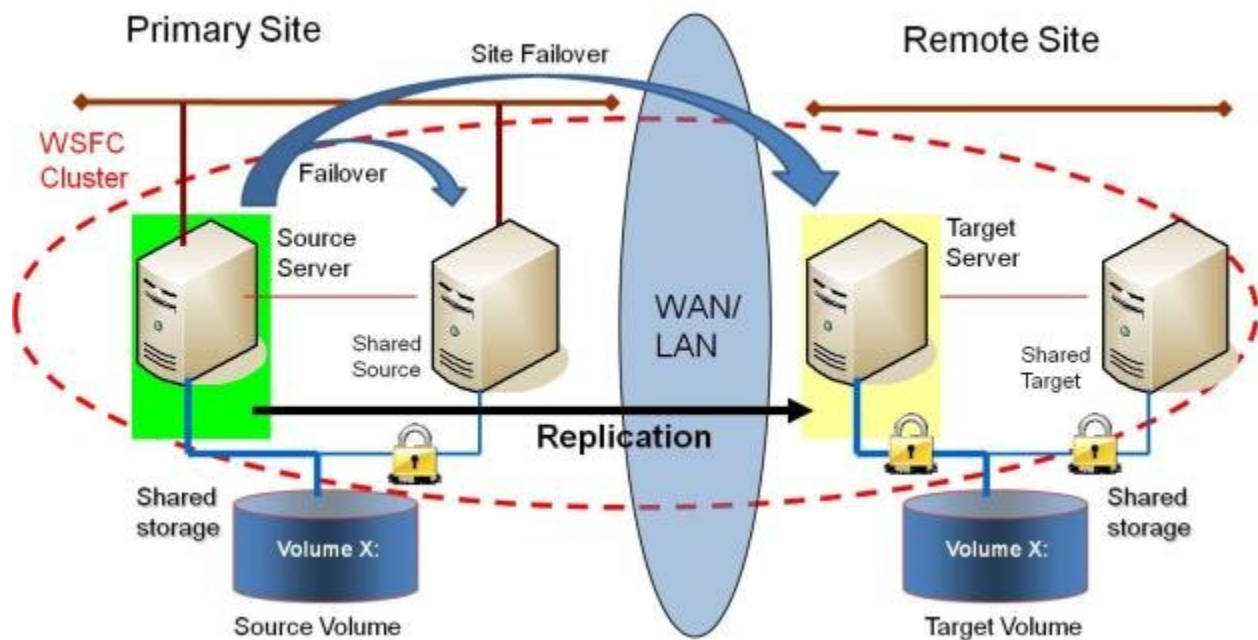
- [Removing a Shared System](#)
- [Frequently Asked Questions](#)

DataKeeper & Failover Clustering

- [DataKeeper Cluster Edition Overview](#)
- [Creating a DataKeeper Volume Resource in WSFC](#)
- [Switchover in an N-Shared x N-Shared Configuration](#)
- [Split Brain Issue and Recovery](#)
- [Using DataKeeper Cluster Edition to Enable Multi-Site Hyper V Clusters](#)

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

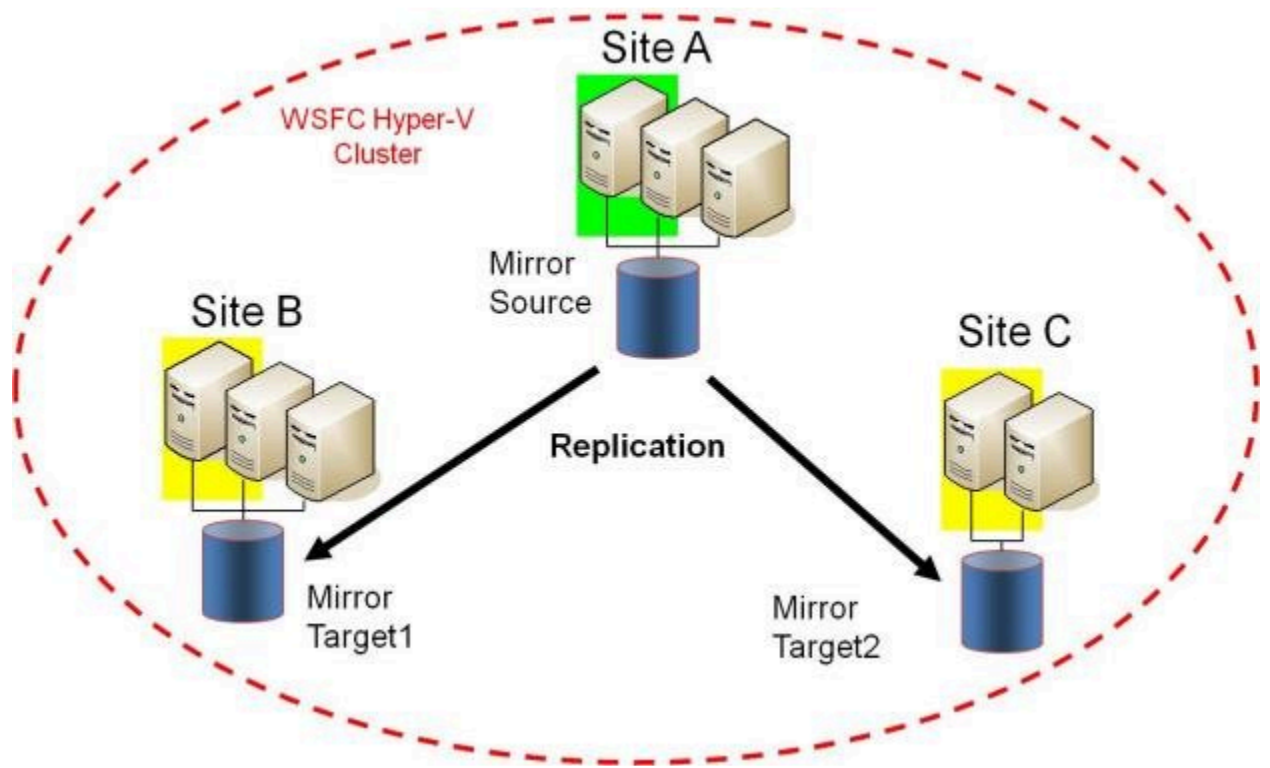
- [Creating Mirrors with Shared Volumes](#)
- [Managing Shared Volumes](#)
- [Adding a Shared System](#)
- [Removing a Shared System](#)
- [Frequently Asked Questions](#)

DataKeeper & Failover Clustering

- [DataKeeper Cluster Edition Overview](#)
- [Creating a DataKeeper Volume Resource in WSFC](#)
- [Switchover in an N-Shared x N-Shared Configuration](#)
- [Split Brain Issue and Recovery](#)
- [Using DataKeeper Cluster Edition to Enable Multi-Site Hyper V Clusters](#)

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](#)

DataKeeper & Failover Clustering

- [DataKeeper Cluster Edition Overview](#)
- [Creating a DataKeeper Volume Resource in WSFC](#)
- [Switchover in an N-Shared x N-Shared Configuration](#)
- [Split Brain Issue and Recovery](#)
- [Using DataKeeper Cluster Edition to Enable Multi-Site Hyper V Clusters](#)

Setting Up SIOS DataKeeper

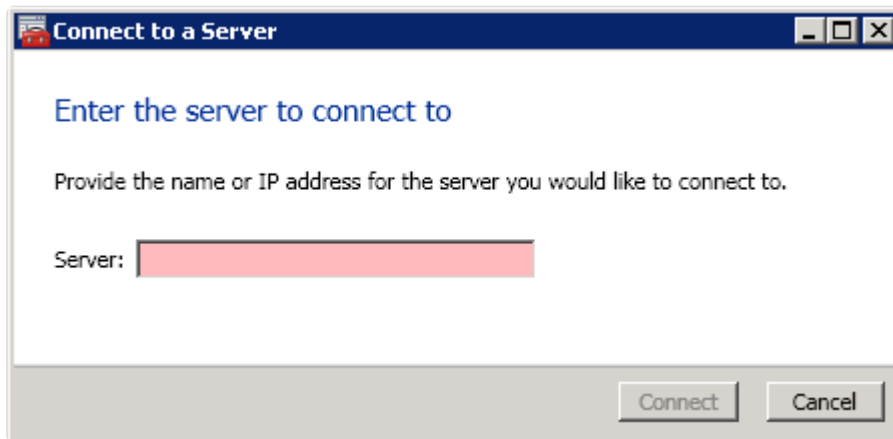
Setting Up SIOS DataKeeper

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

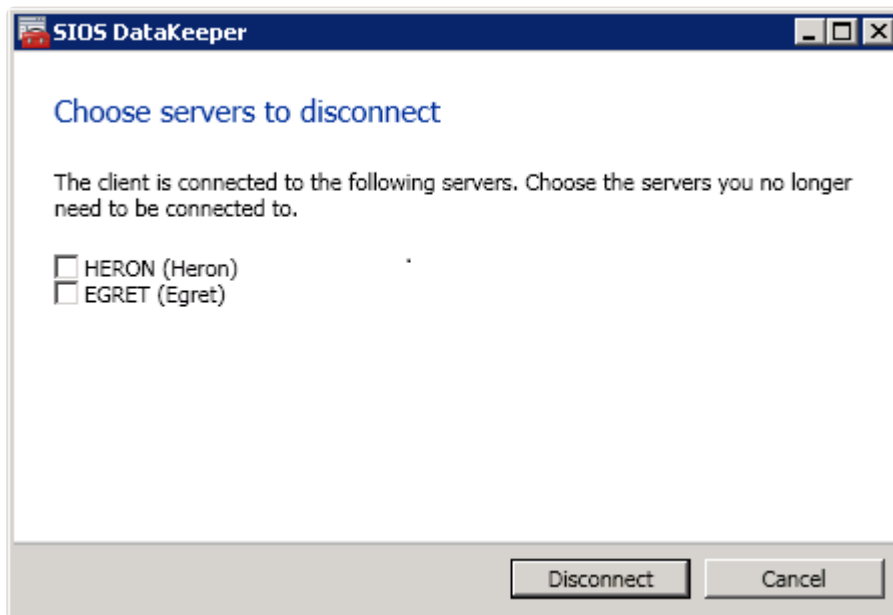
Use this dialog to connect to the server of your choice. You may enter the IP address, system NetBIOS name or the full system domain name for the server. Click **Connect** to select it.



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

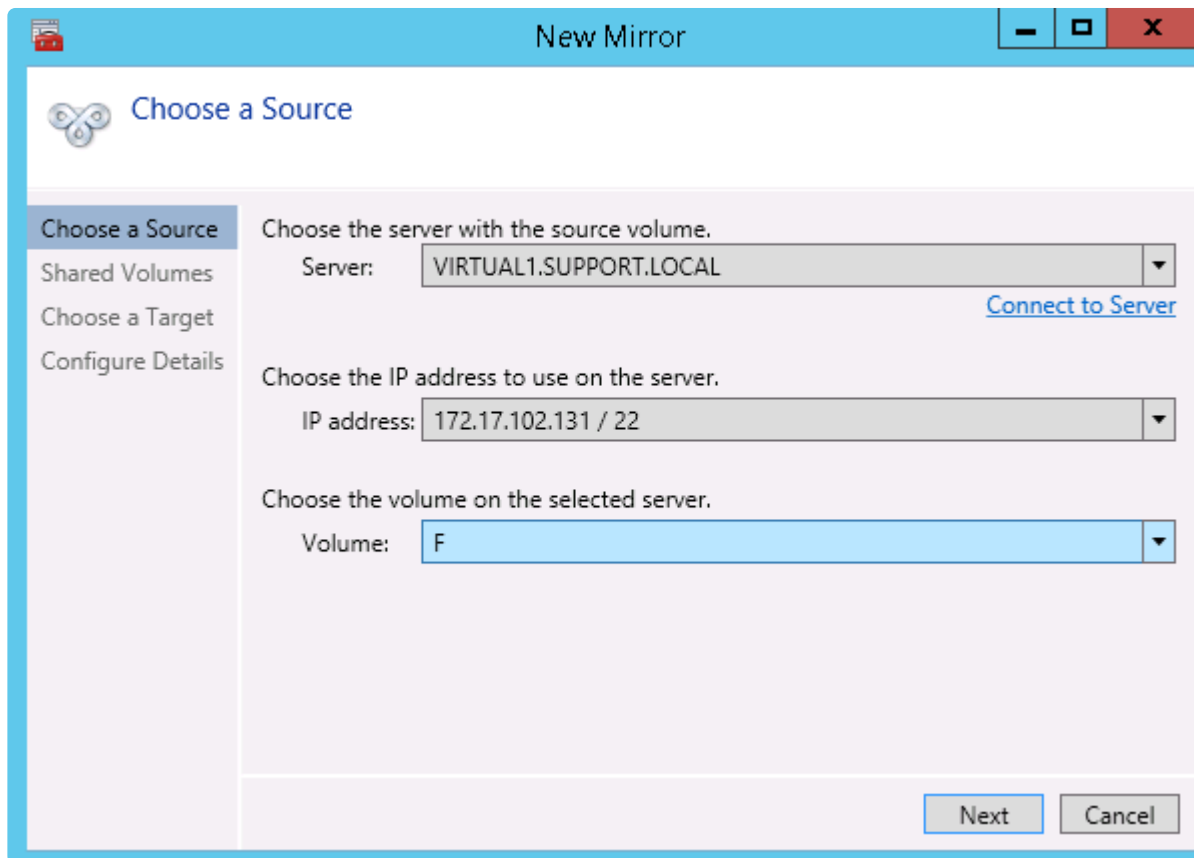
In order to properly configure DataKeeper in a shared volume configuration, use the **DataKeeper GUI** to connect to all systems where the shared volumes are configured. When connected, the DataKeeper GUI uses hardware signatures to automatically detect which volumes are shared and which are not.

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 of the dialog, 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.

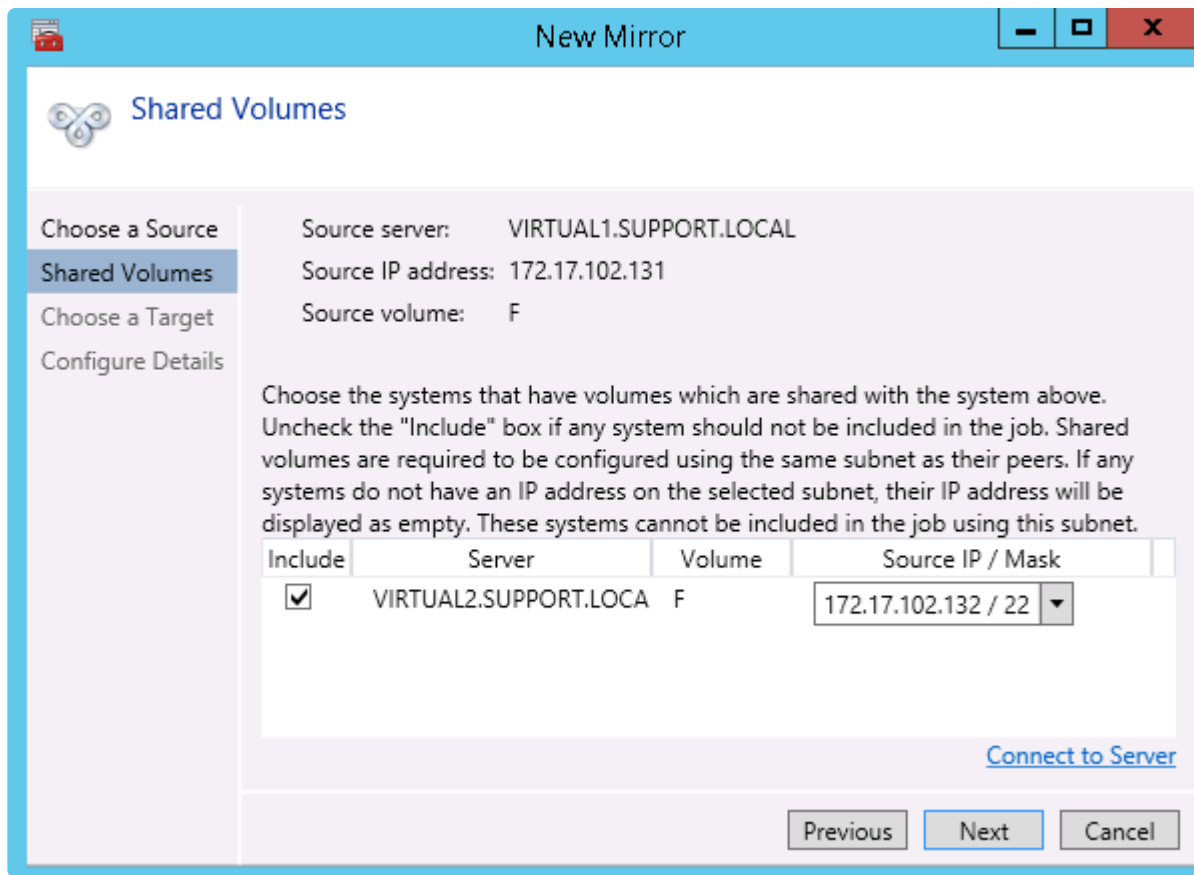
IP address:

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.



6. 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.

7. Select **Next**. The **Choose a Target** dialog box appears.
8. Choose a **Target System**, **IP Address** and **Volume**.
9. 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

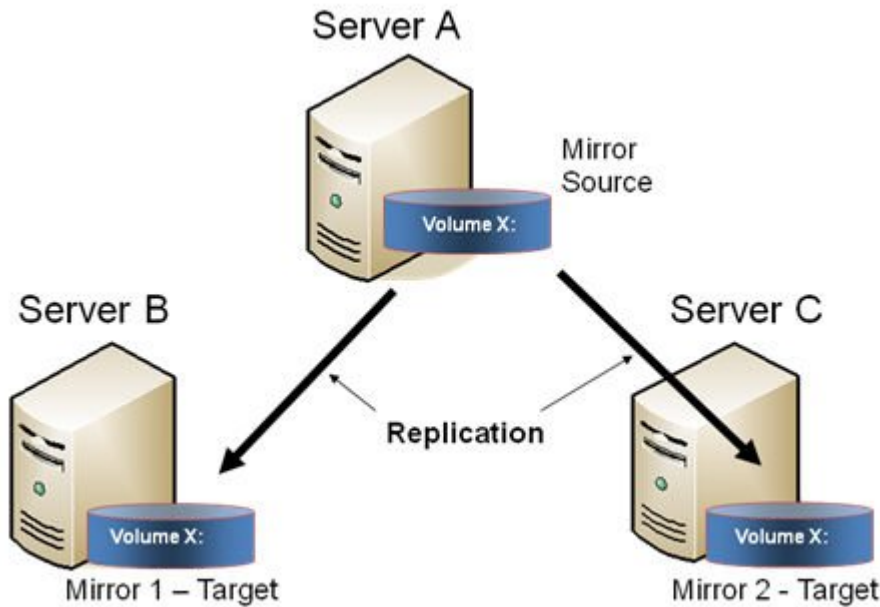
Example:

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



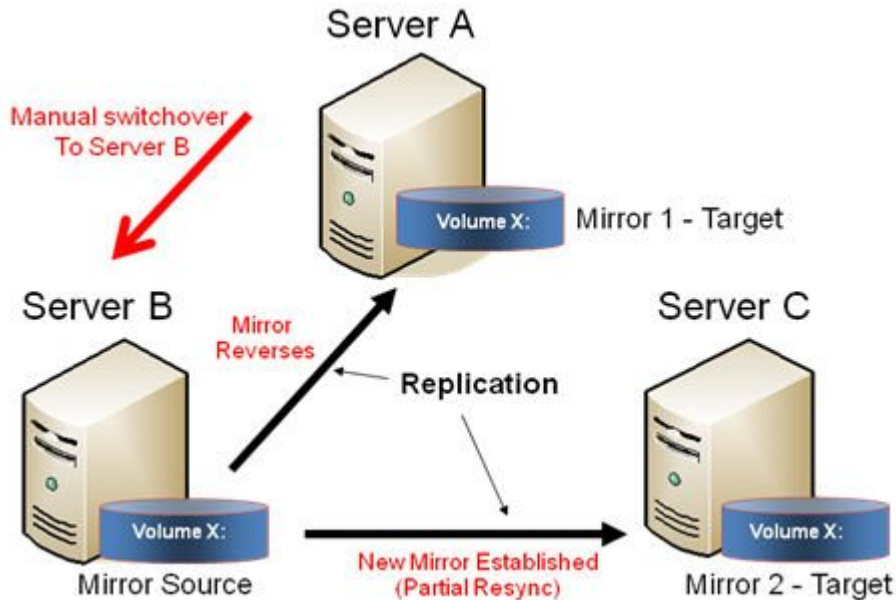
Manual Switchover to a Target Server

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

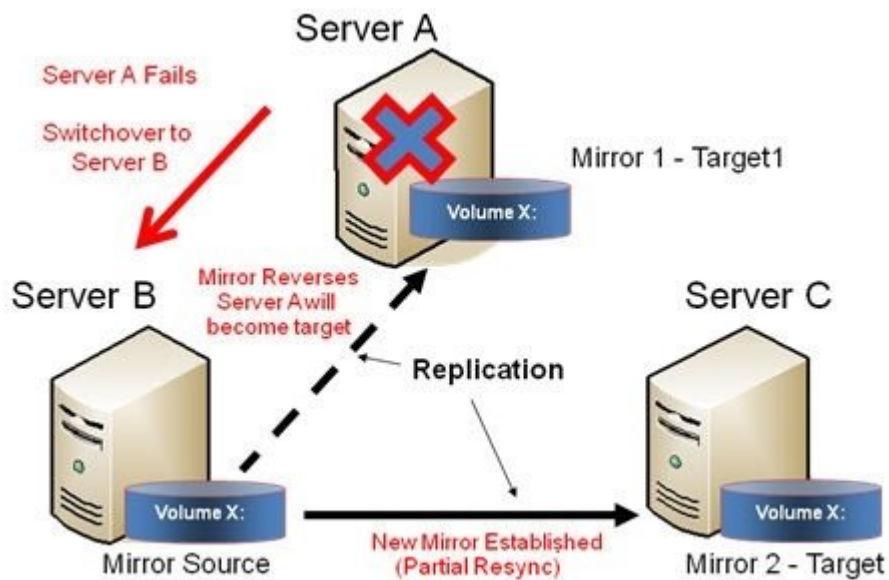
In the event the active (source) server fails, DataKeeper allows you to make Server B become the active (source) server. The following actions will occur:

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.

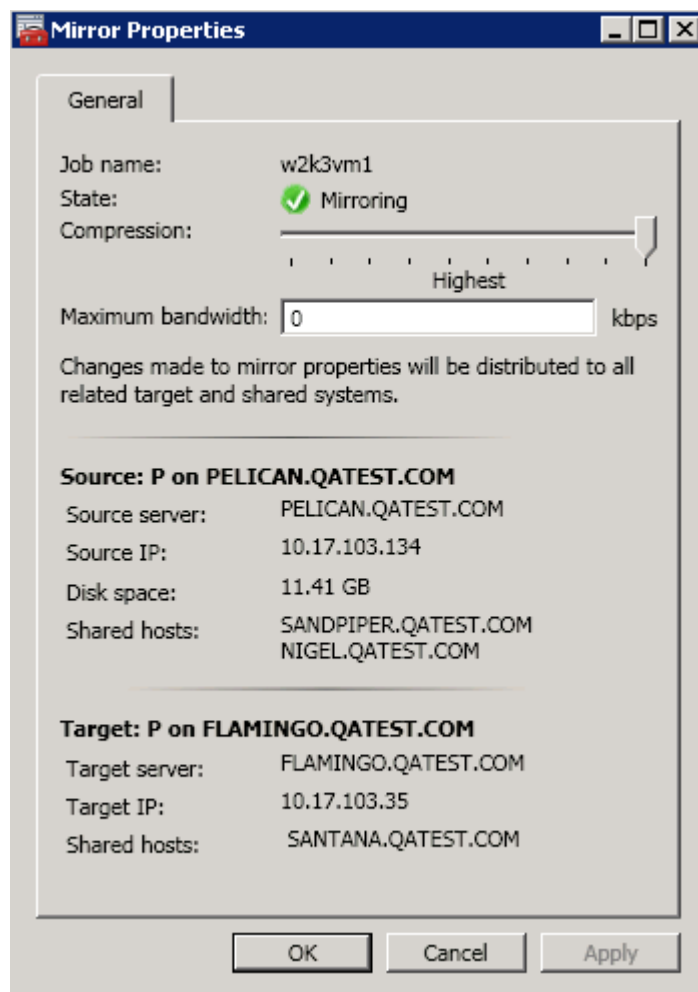
2. Shut down (power off) all shared source and/or shared target systems. (**Note:** Current source and current target systems should not be shut down.)
3. Change the volume sizes as noted above in the Non-Shared Volume procedure.
4. [Continue and Lock](#) the mirrors after resizing has completed.
5. 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

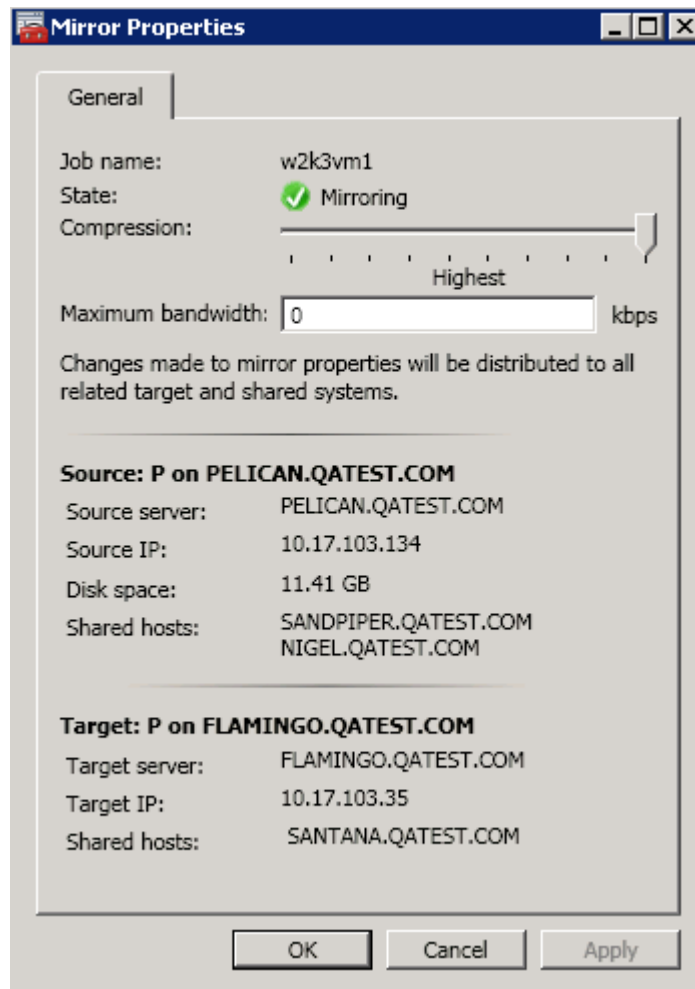
The compression level of the mirror is set during mirror creation and applies to that specific mirror only.

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 screenshot shows the 'Mirror Properties' dialog box with the 'General' tab selected. The dialog contains fields for job name, state, compression, and bandwidth, as well as sections for source and target server details.

Mirror Properties

General

Job name: w2k3vm1

State: ☒ Mirroring

Compression: 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](#)

Next **Cancel**

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

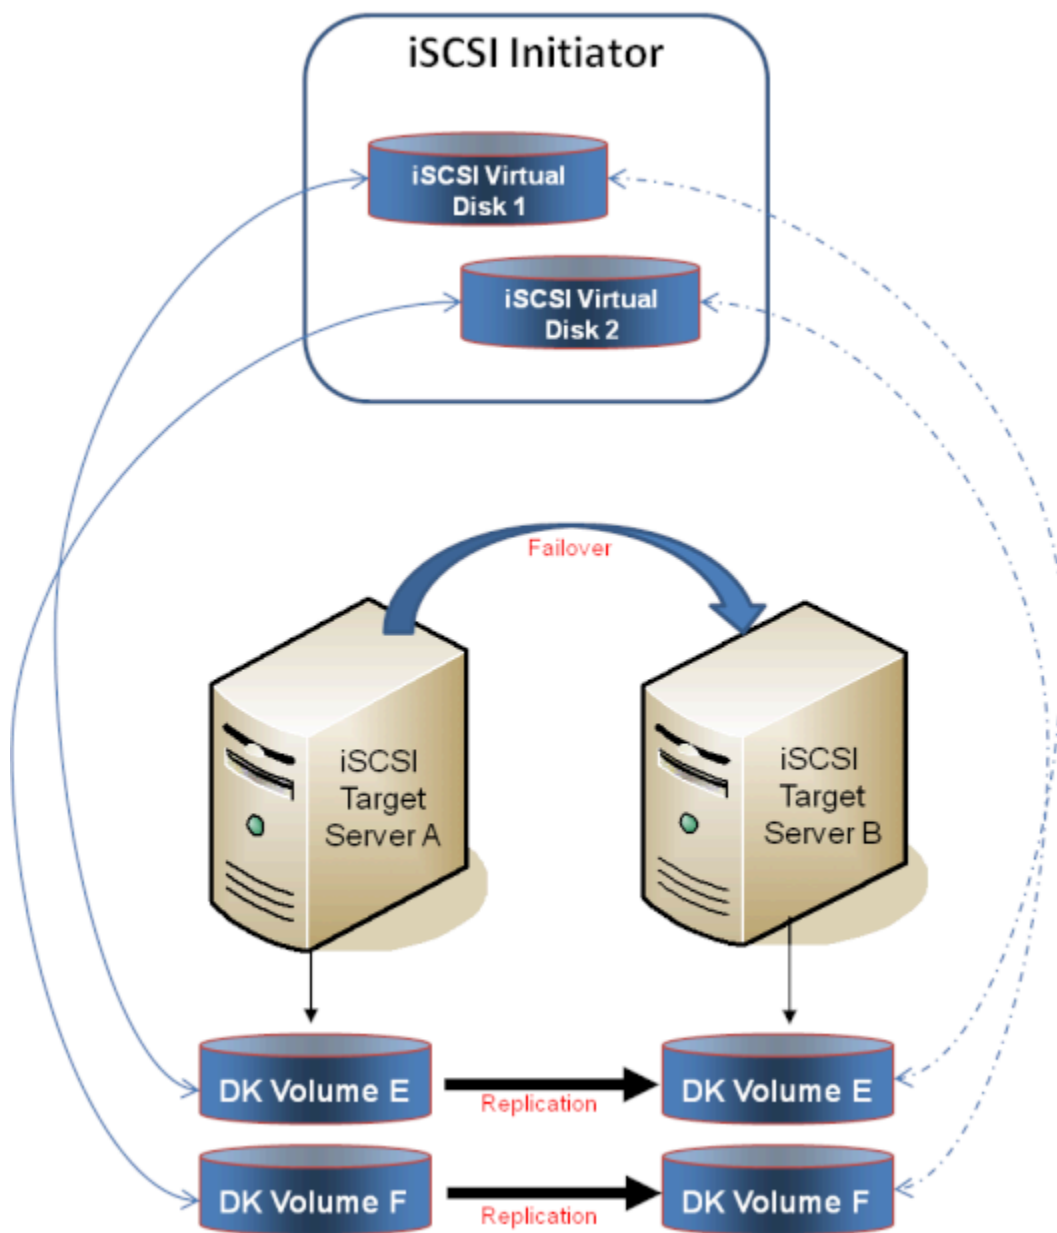
To remove a shared system from either side of the mirror, bring up the **Manage Shared Volumes** dialog and uncheck the **Include** box for the system to be removed. When you click **Done**, the job will be updated so that the system is not part of the job.

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](#)

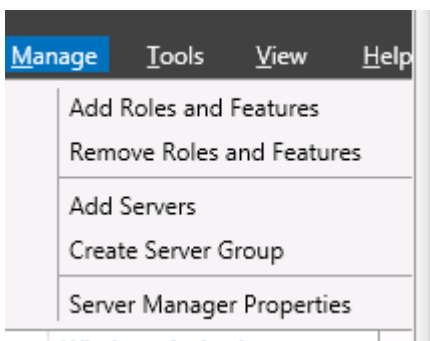
[Creation of Mirror and Configuration of Cluster](#)

[Creation of iSCSI Virtual Disks](#)

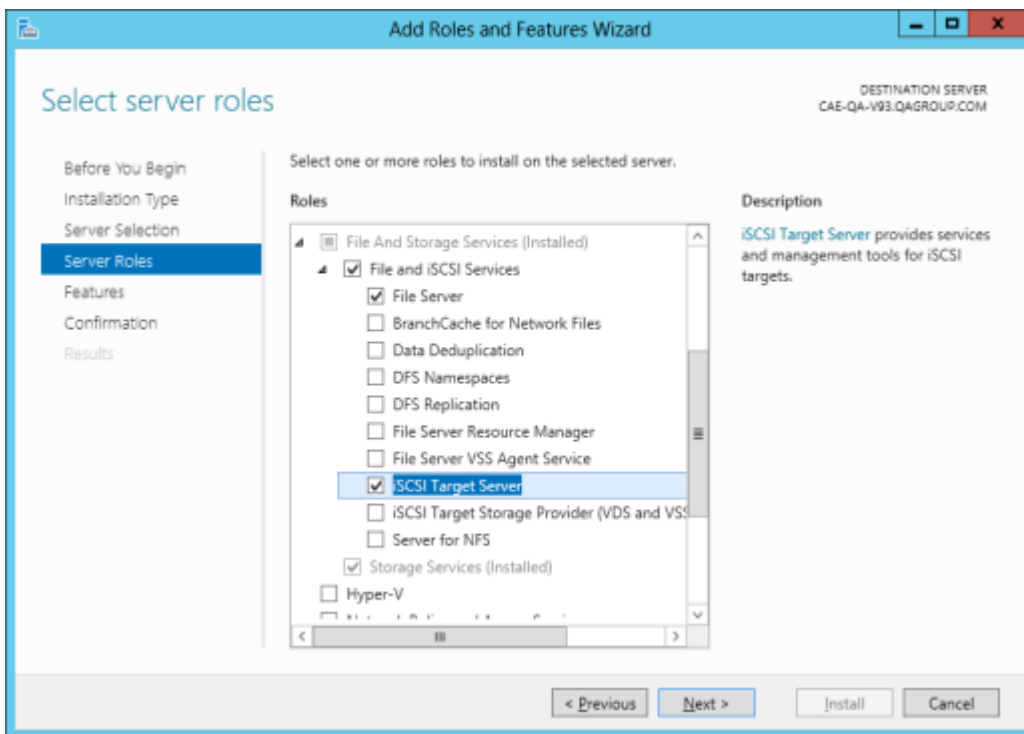
[Setup of iSCSI Initiator on Windows 2012](#)

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.



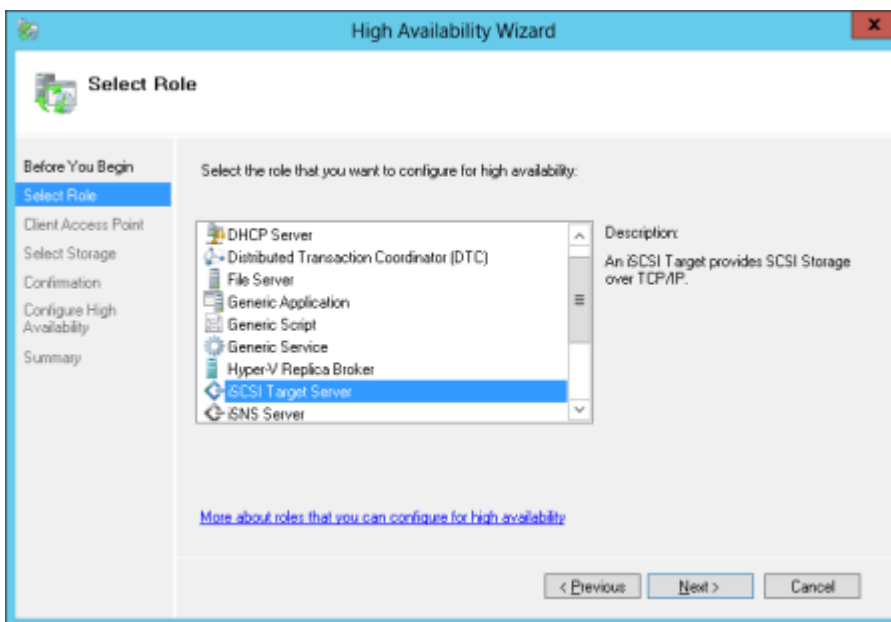
5. Click "**Next**" twice to get to the "**Install**" button to be able to install the role.
6. The feature will install and the progress will be shown.
7. Upon completion, the message "**Installation succeeded**" will be displayed.
8. Repeat these steps for all servers in the cluster.

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**.

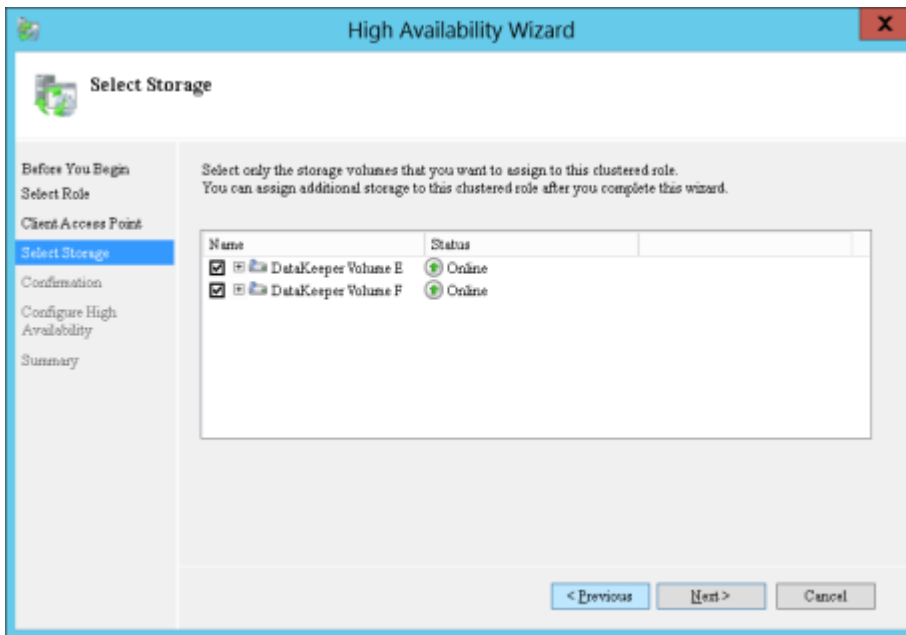


3. Select **iSCSI Target Server** role and select **Next**.
4. The **Client Access Point** page appears. Type the **Client Access Point name** and **IP address** for the iSCSI Target Server instance.

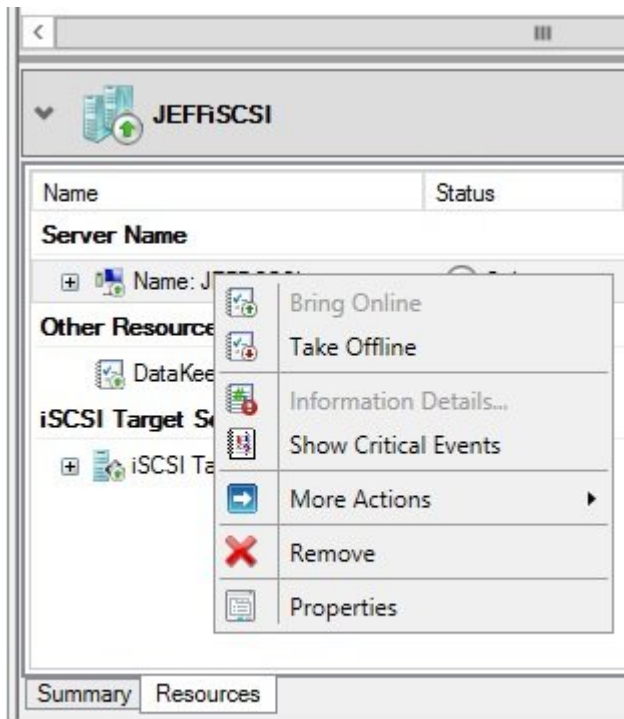
***IMPORTANT:** This name and IP address will be used later by clients to access the server address, so it should be recorded in DNS. This is very

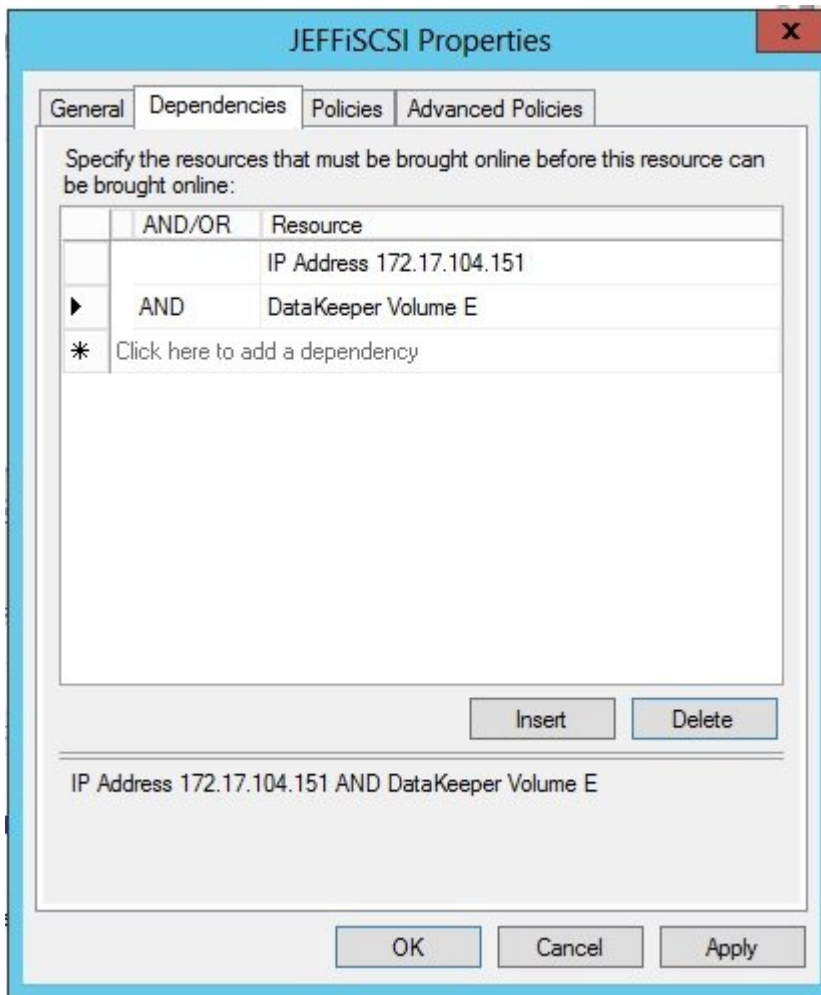
important for the servers to be able to resolve these names.

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.



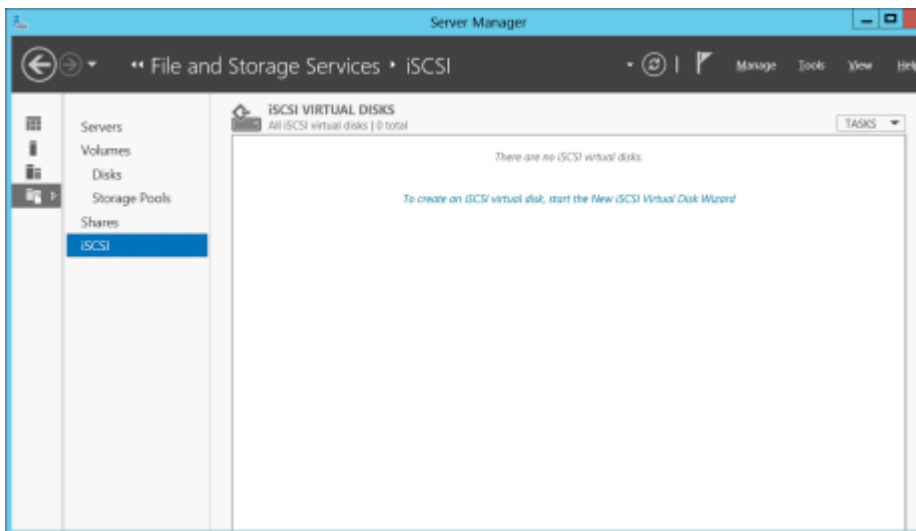


8. Setup is complete. Proceed to the [iSCSI Virtual Disks](#) configuration.

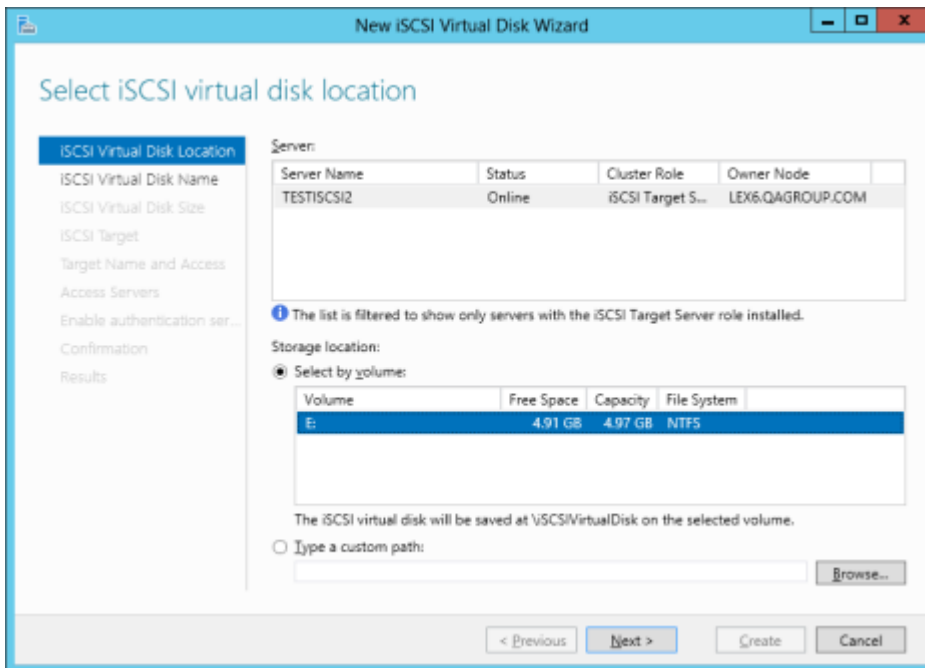
Creation of iSCSI Virtual Disks

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.

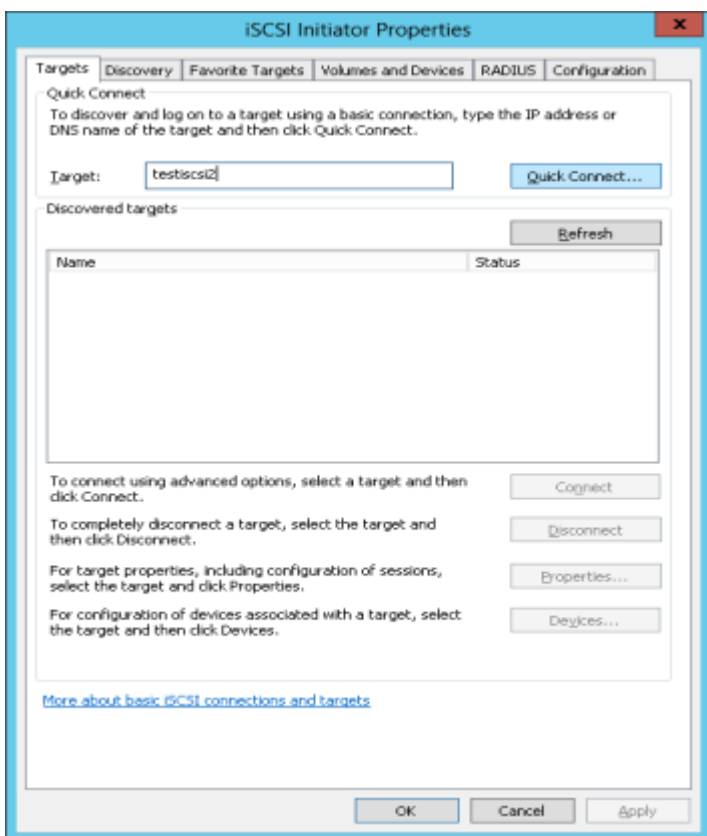
You need to have a plan ahead of time that describes which files you want to create and whether those files should all be accessed simultaneously or if the disks need to be accessed separate from one another.

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).

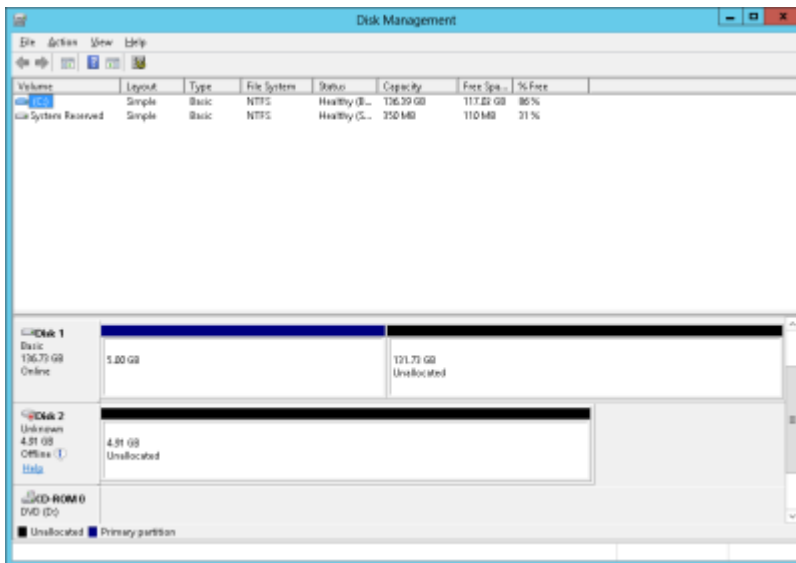
Setup of iSCSI Initiator on Windows 2012

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.



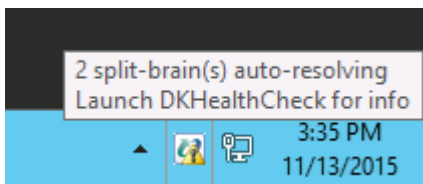
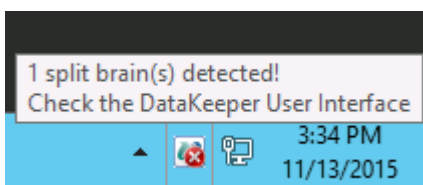
5. Right-click on the disk(s) to bring it online.
6. Initialize the disk(s).
7. Create the new volume and assign the drive letter.
8. Configuration is now complete.

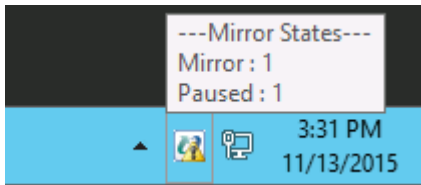
DataKeeper Notification Icon

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

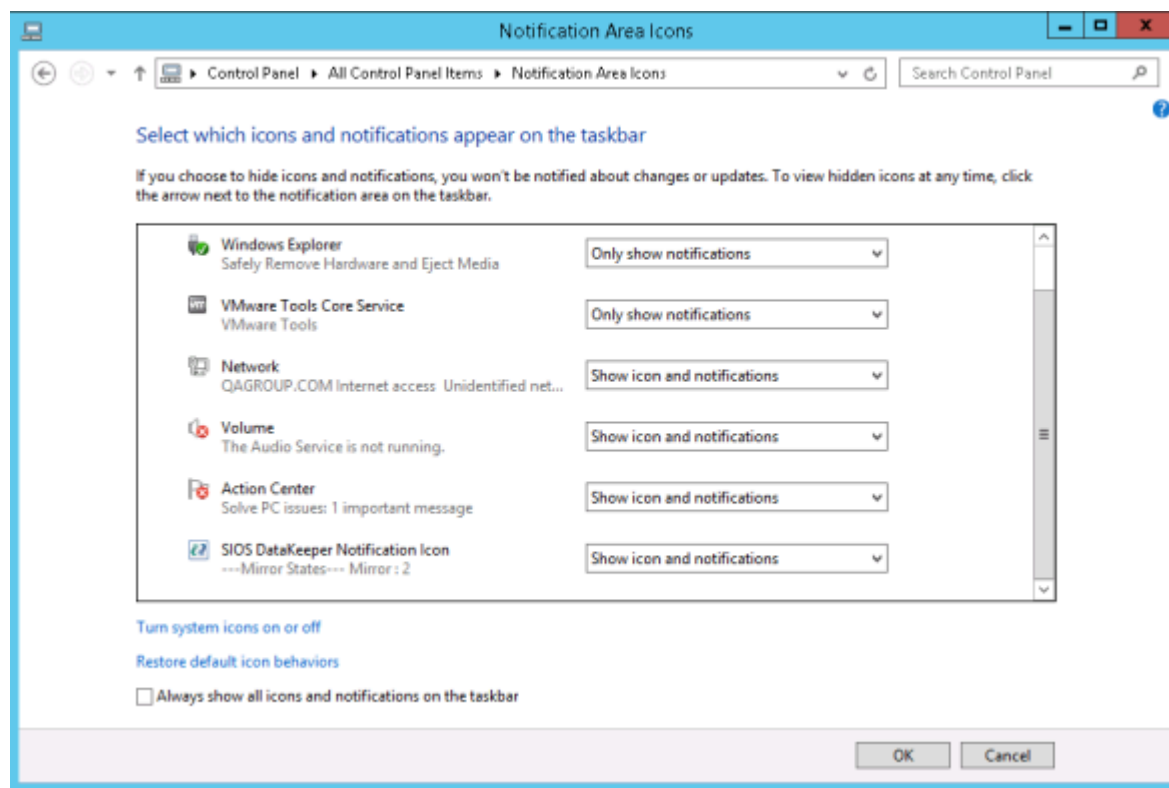
The Notification Icon should automatically appear in the Windows Notification Tray upon logging in to the node.

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

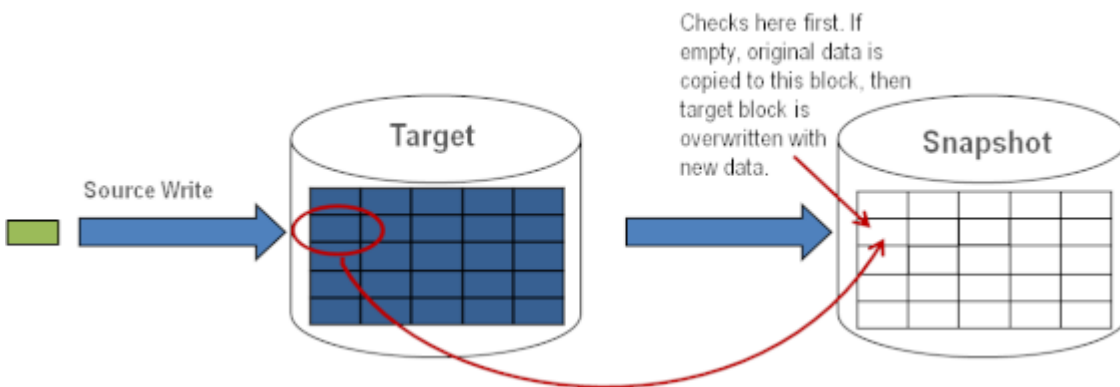
The snapshot exists in parallel with the live copy of the volume to be backed up, so except for the brief period of the snapshot's preparation and creation, an application can continue its work. Writes to the target, however, will now be processed differently while the target is in this

state.

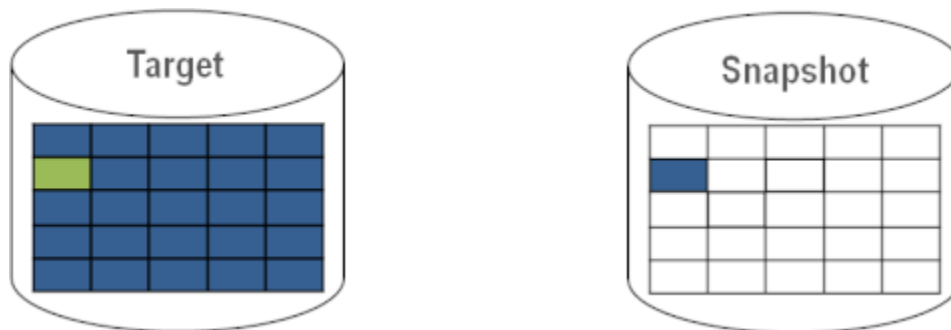
Data mirroring from the source system will continue uninterrupted, but any new data from the source that is received after the snapshot is taken will not be visible on the target system until the snapshot is dropped. This allows an application on the target system to run, using (and updating) data that 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.

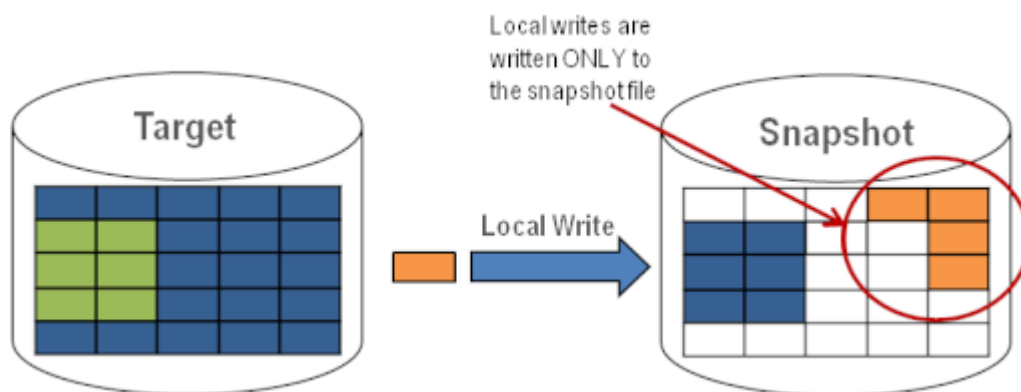


If DataKeeper determines that this block has already been written to the snapshot file, then this step is skipped and the block is

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.

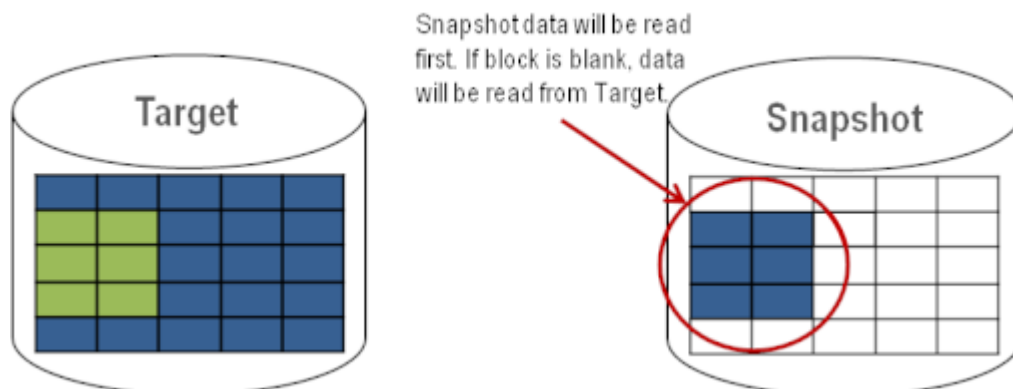
Local Write

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

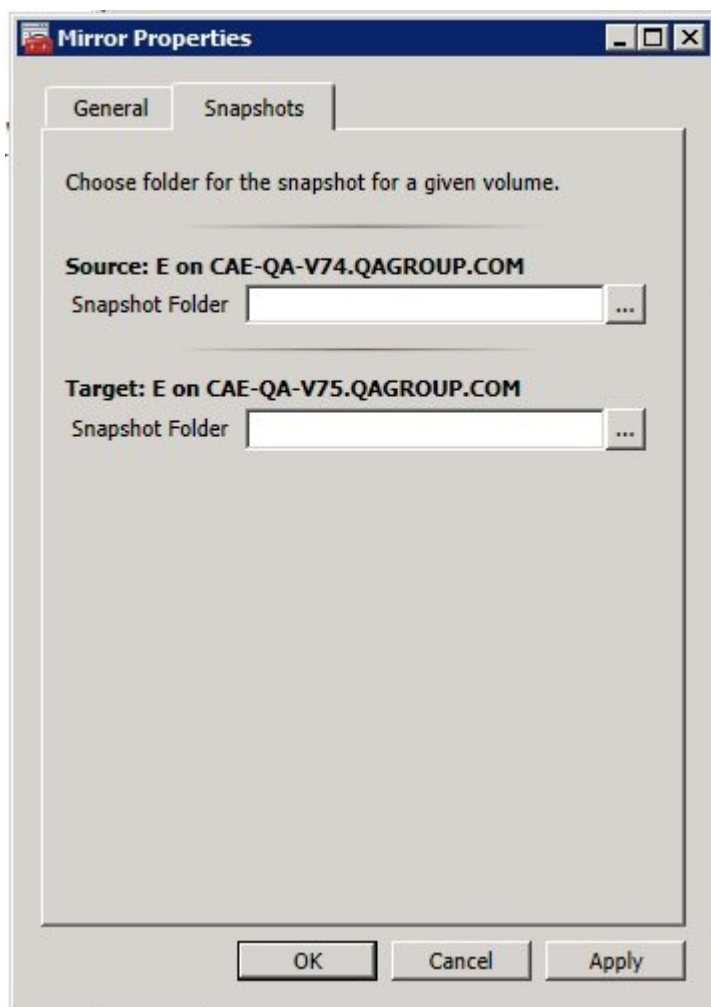
The size of the snapshot location should be determined on an individual basis based on several criteria. In practice, the size required for the snapshot file will be far less than the size of the volume being snapshot. The storage required needs to be big enough to contain any data that changes on the 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

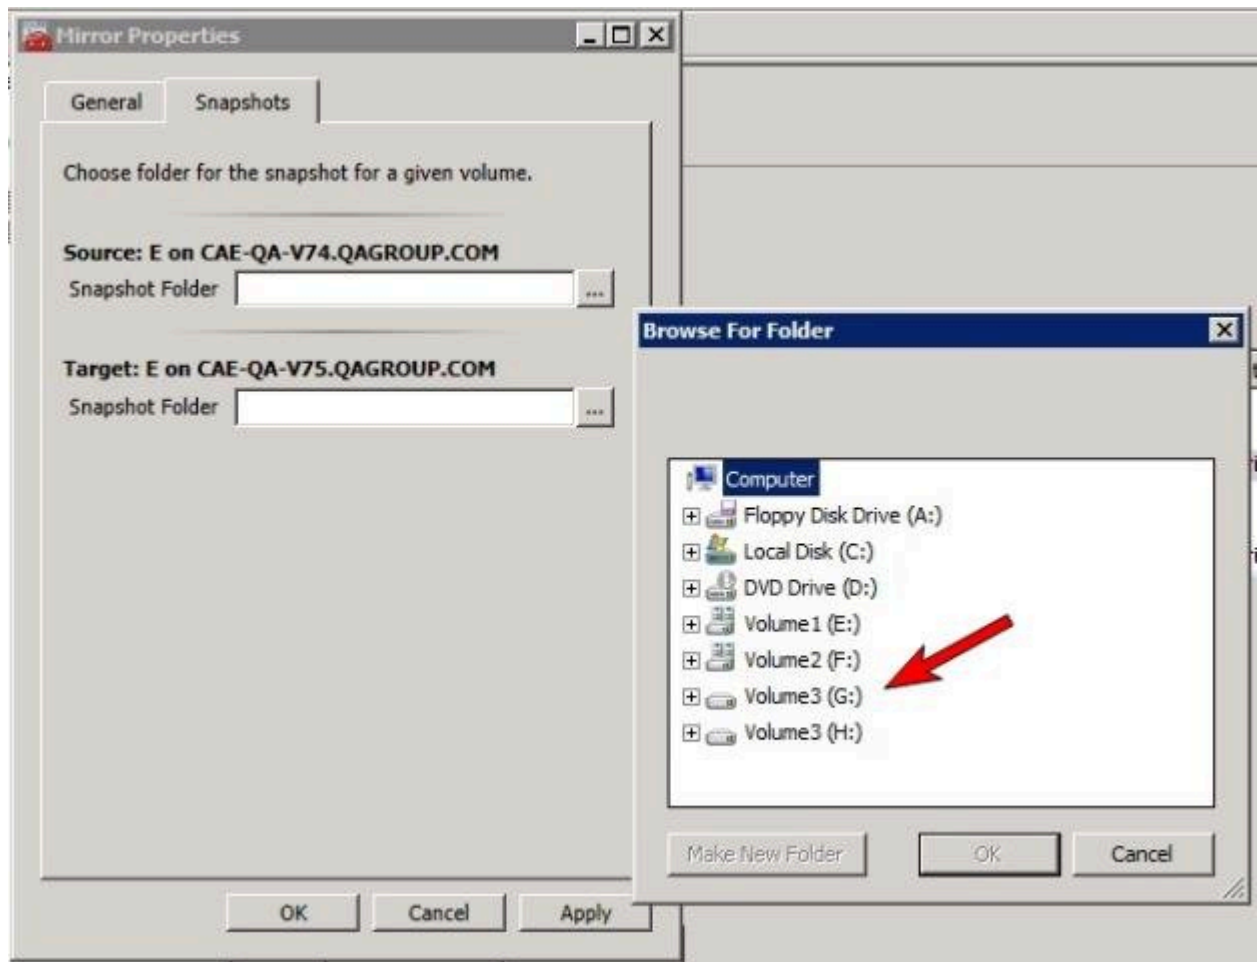
1. Right-click on the appropriate mirror and select **Mirror Properties**.
2. From the **Mirror Properties** dialog, select the **Snapshots** tab.



***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 DKCE.

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.

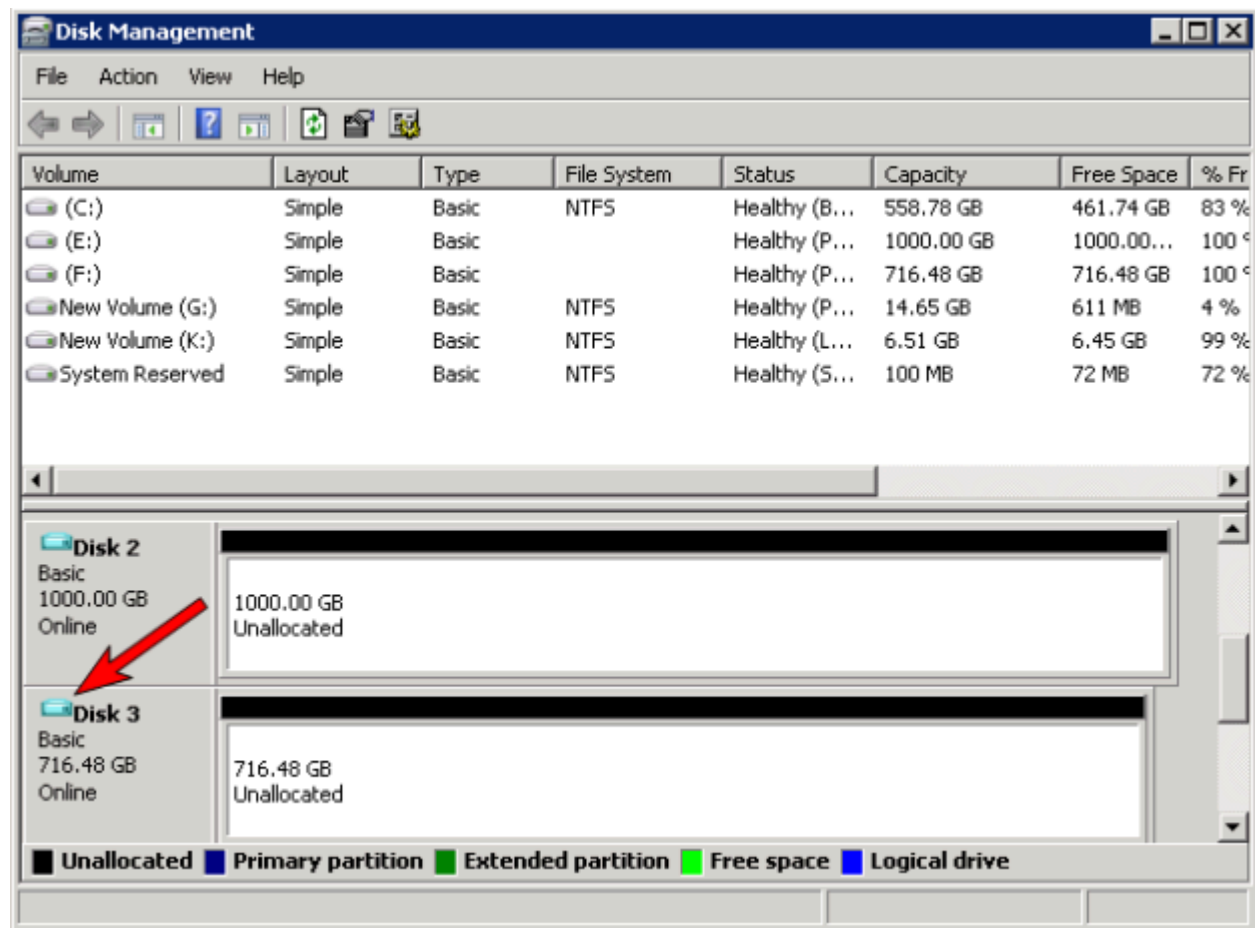
Note: To provide protection during SQL Server recovery we provide a generic script that needs to be added to stop the reporting SQL instance on the target node. Instructions are located in DKSnapshotCleanup.vbs script located in "\support". Please review the script code on how to add to your WSFC hierarchy.

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.

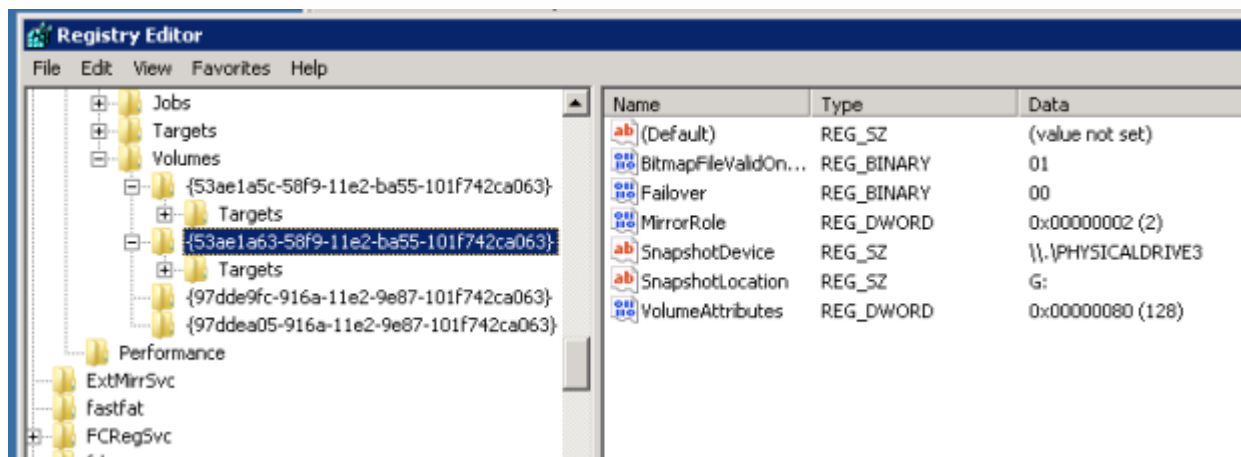
***NOTE:** The colored icon next to the disk number represents this disk as a

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

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.

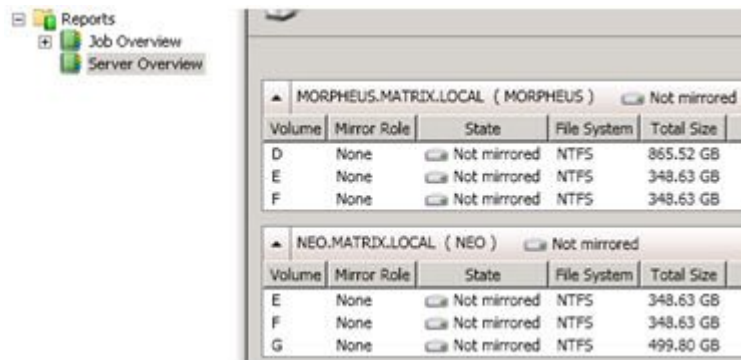
Preparing the Environment

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

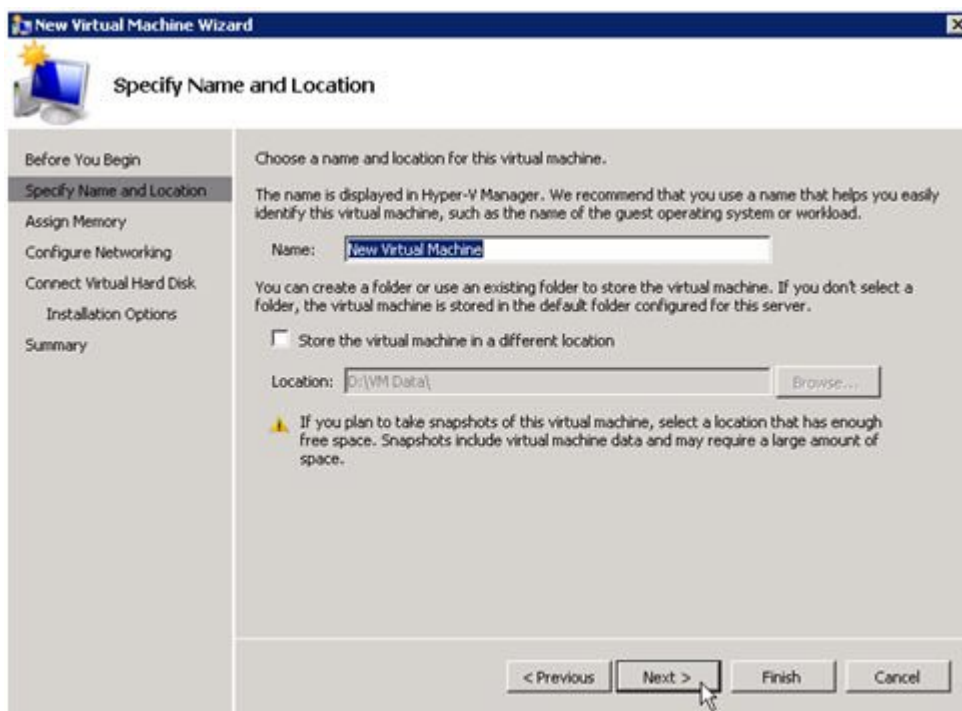
NEO.MATRIX.LOCAL (NEO) <input type="checkbox"/> Not mirrored				
Volume	Mirror Role	State	File System	Total Size
E	None	<input type="checkbox"/> Not mirrored	NTFS	348.63 GB
F	None	<input type="checkbox"/> Not mirrored	NTFS	348.63 GB
G	None	<input type="checkbox"/> Not mirrored	NTFS	499.80 GB

4. [Create a Job](#) / [Mirrored Volume](#).

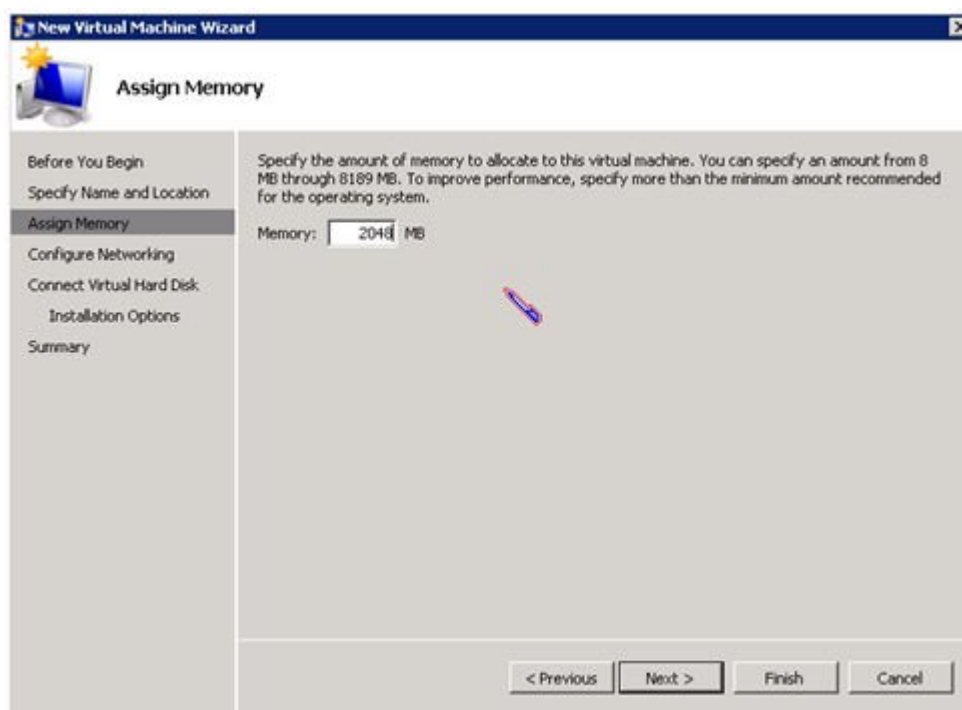
Note: When you select your source server, ensure you select the server whose volume you want to replicate from. Reversing the source and target in these steps will completely overwrite your source volume with whatever is on the target server's volume, even if it is empty, causing you to lose any and all data stored on the source volume.

Create and Configure a Hyper-V Virtual Machine

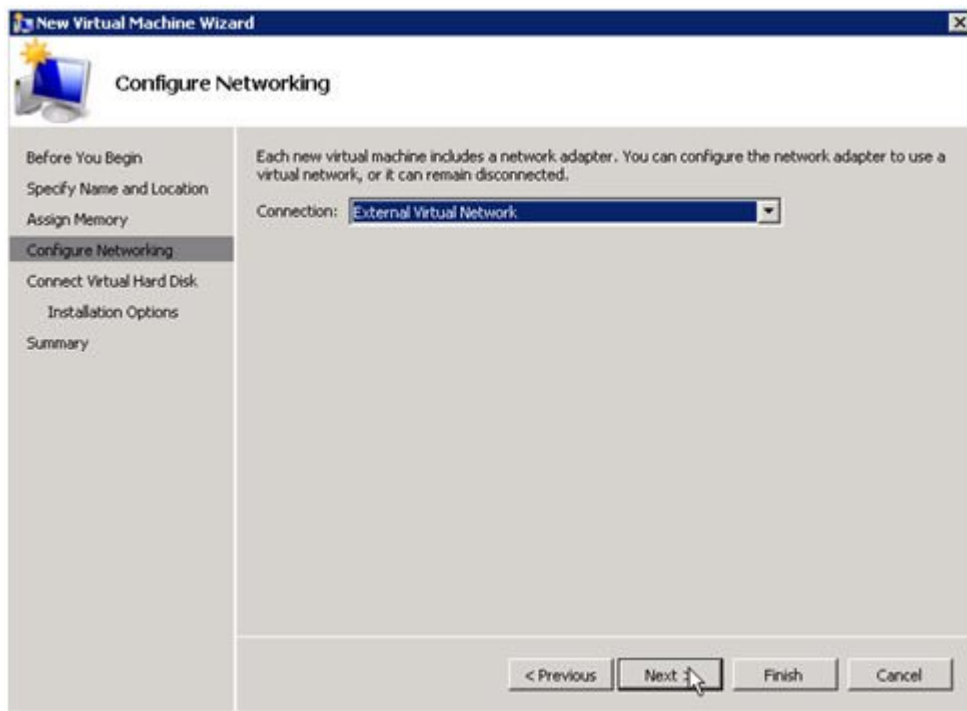
1. Launch the **Hyper-V Console** from **Start - Administrative Tools - Hyper-V Manager**.
2. Start the **New Virtual Machine Wizard**.



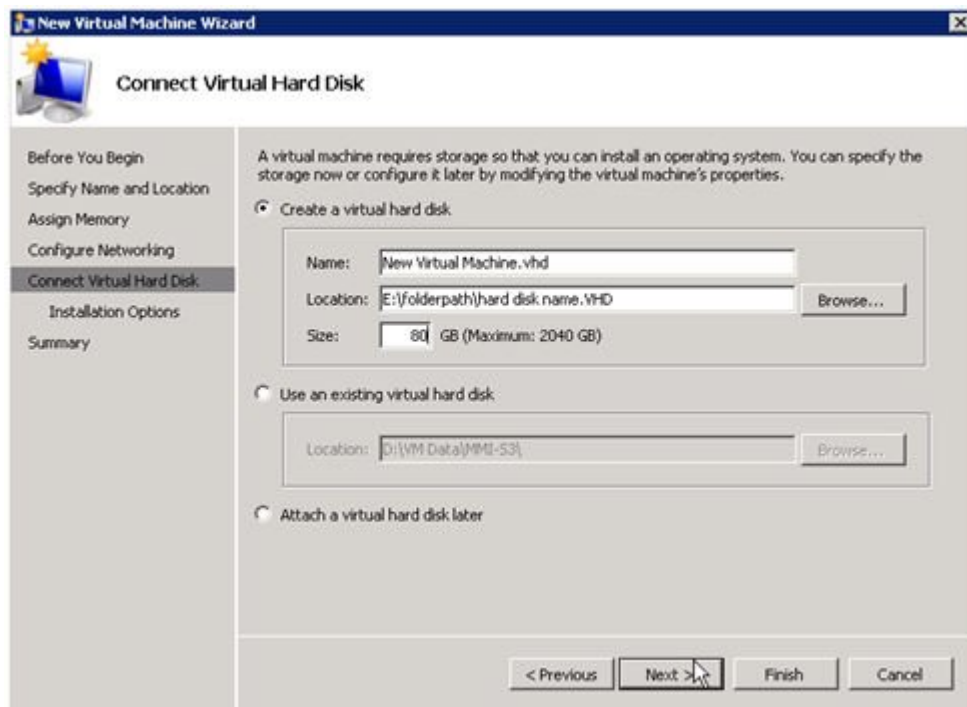
3. Specify the amount of **RAM** to use.



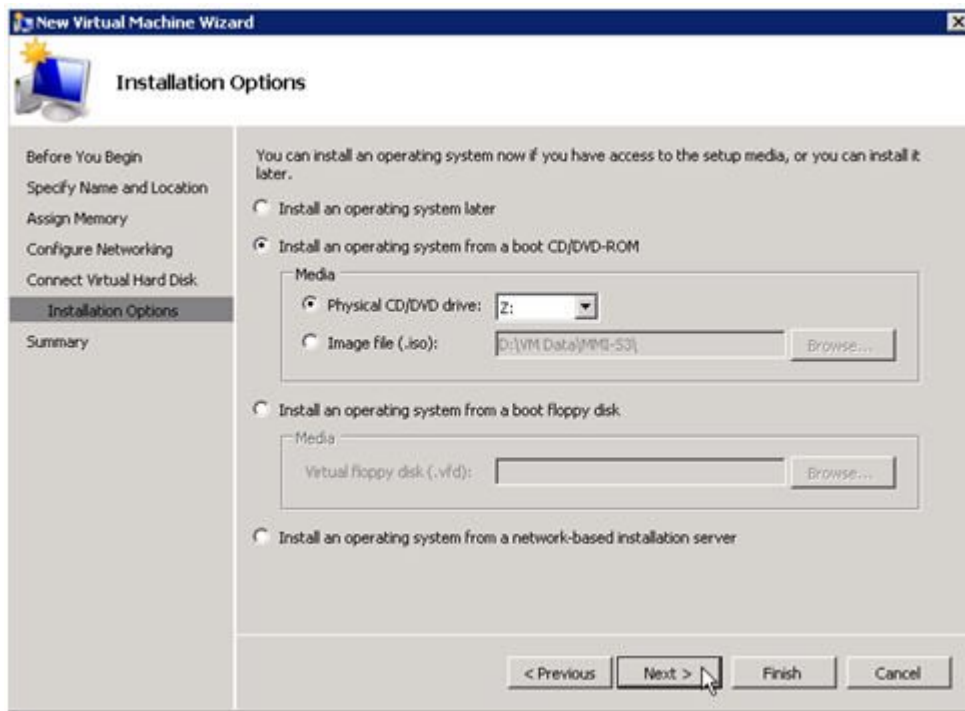
4. Select a **network adapter** to use.



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**.

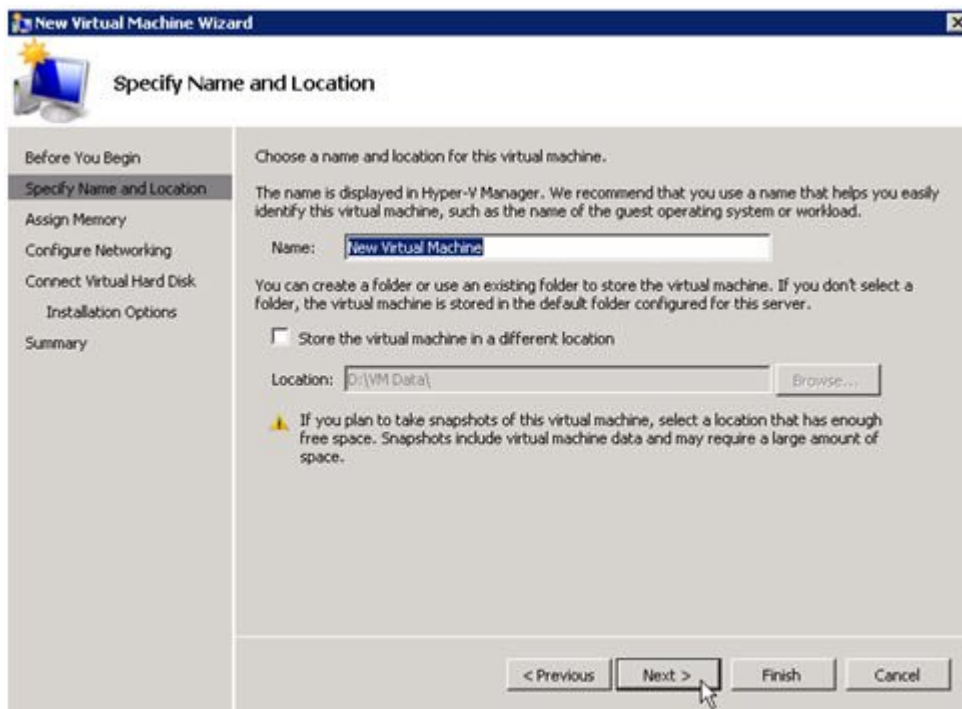
State	Resync Remaining
Mirroring	0.00 KB

4. Select the mirror and click **Switchover** in the **Actions** pane.

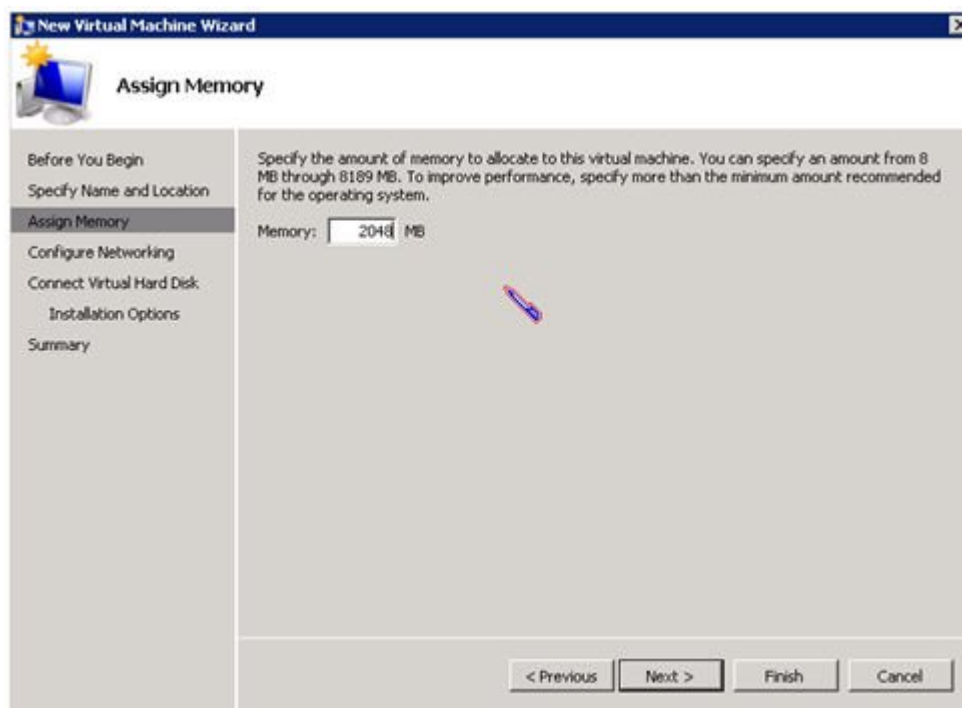


This will reverse the source and target and allow you to provision the virtual machine on the target server.

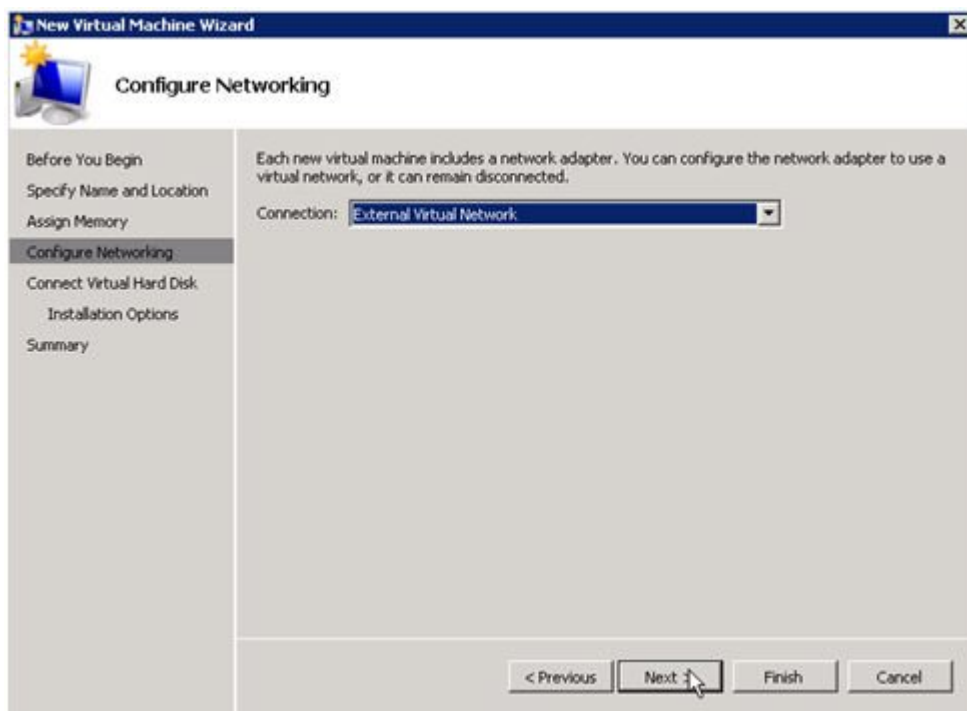
5. On the target server, start the **Hyper-V Manager**.
6. Start the **New Virtual Machine Wizard**.



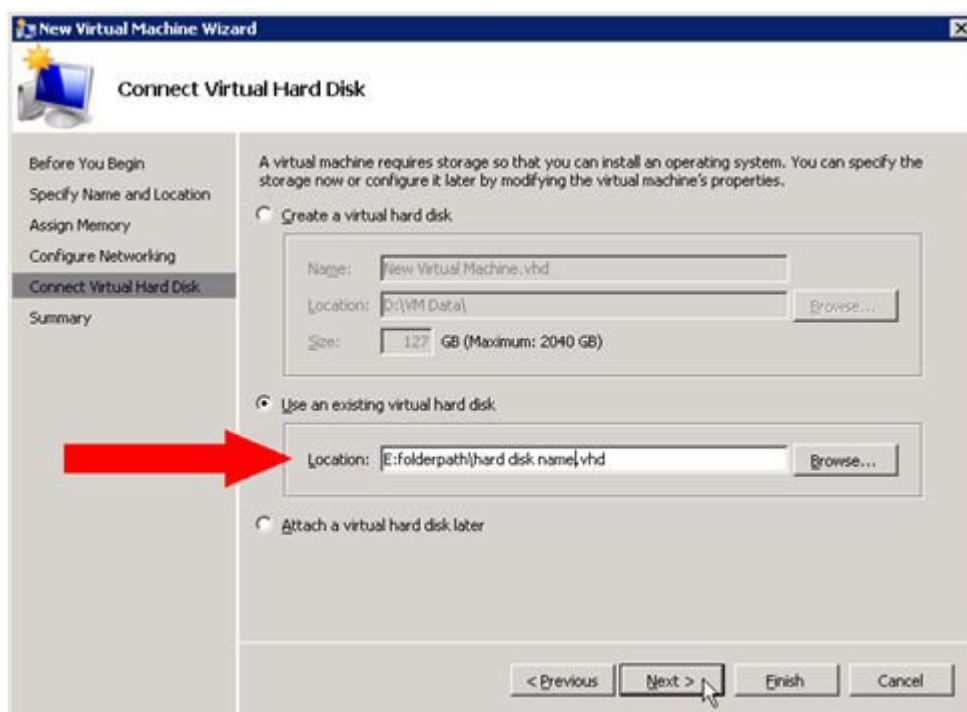
7. Specify the amount of **RAM** to use.



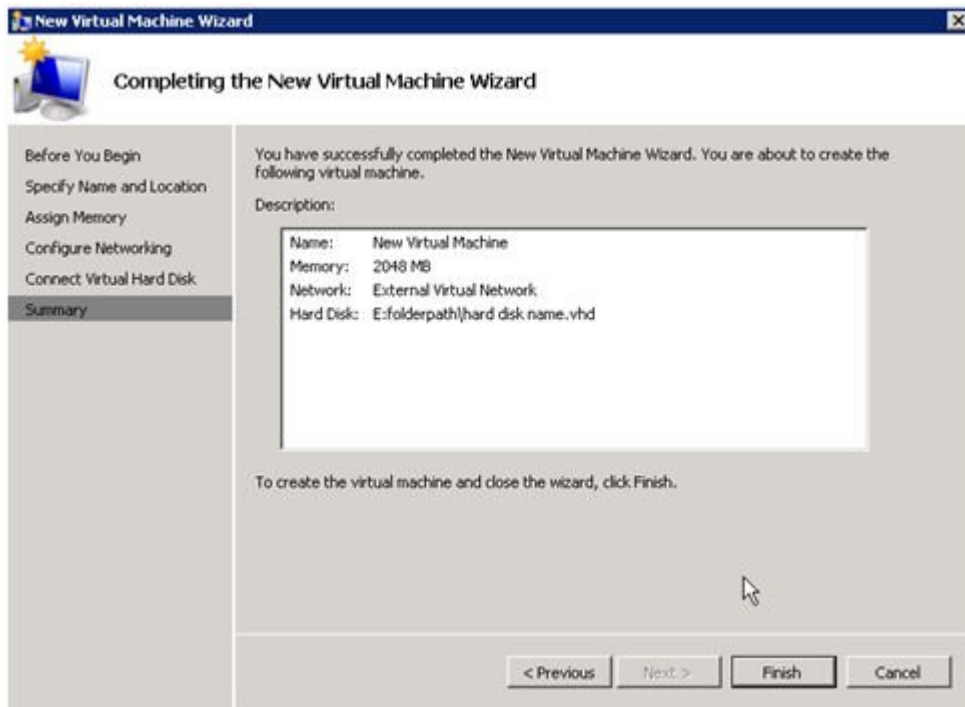
8. Select a **network adapter** to use.



***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

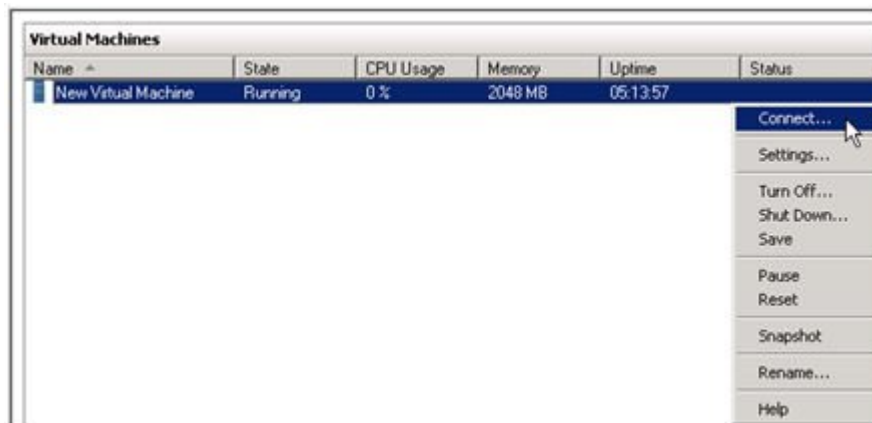
Initiate a **Planned Switchover** to migrate the virtual machine back to your source server.

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**.

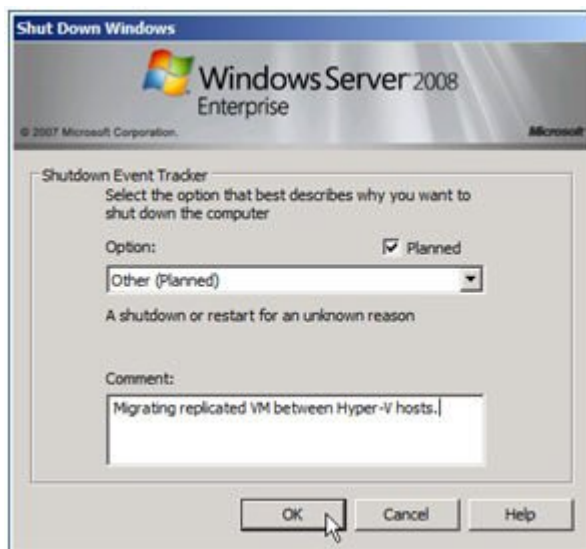
Planned Switchover

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.

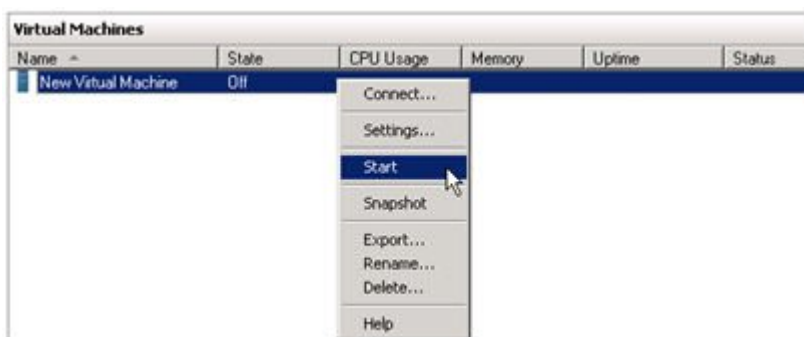
State	Resync Remaining
Mirroring	0.00 KB

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.

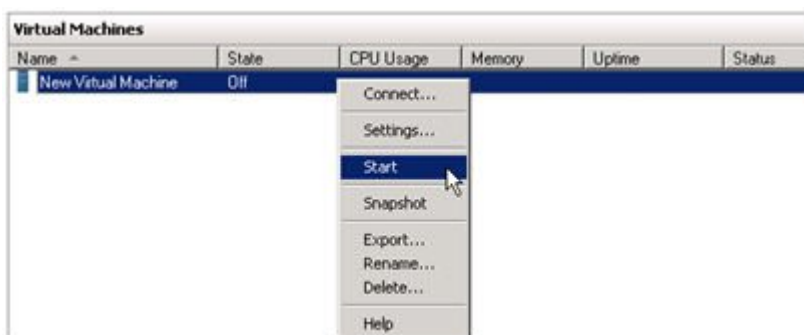
1. On the target server, start the **DataKeeper console** as described previously.
2. Select the mirror and click **Switchover** in the **Actions** pane.



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.

Clustering

[Running chkdsk on Cluster Volumes during Cluster Volume Online](#)

[Creating a DataKeeper Volume Resource in WSFC](#)

[Manual Creation of a Mirror in WSFC](#)

[DataKeeper Volume Resource Health Check](#)

[DataKeeper Volume Resource Private Properties](#)

[Extending a Clustered DataKeeper Volume to a Node Outside the Cluster](#)

[Extending a Single SQL Server Node to a Cluster](#)

[Extending a Traditional 2-Node WSFC Cluster to a Third Node via DataKeeper](#)

[Extending a Traditional 2-Node WSFC SQL Server Cluster to a Third Node via DataKeeper](#)

[Extending a Traditional 2-Node Cluster to a Shared-Replicated Configuration](#)

[Using DataKeeper Cluster Edition to Enable Multi-Site Hyper-V Clusters](#)

[Split-Brain Issue and Recovery](#)

[Switchover in an N-Shared x N-Shared Configuration](#)

[Installing and Using DataKeeper Cluster Edition on Windows Server 2008 R2 / 2012 Core Platforms](#)

[Non-mirrored Volume Resource](#)

[Using DKCE to Enable Multi-Site File Share Resources with Windows Server 2008R2 WSFC](#)

[Creating Other Server Resource in WSFC](#)

Running chkdsk on Cluster Volumes during Cluster Volume Online

As of version 7.6 DataKeeper now runs chkdsk on all mirrored volumes prior to the volume being available for use. DataKeeper now creates a new flag DiskRunChkDsk in Windows Failover Clustering for each DataKeeper volume. The flag determines how chkdsk runs on each volume during disk check operations.

Volume Commands:

Flag value 0 means chkdsk is enabled and the entire volume will be checked

Flag value 4 means chkdsk is skipped for the entire volume

The volume command performed depends on the argument located in the DiskRunChkDsk flag for each DataKeeper volume. The location of this flag is in HKEY_LOCAL_MACHINE/Cluster/Resources/{DataKeeper Volume GUID}/Parameters.

The complete list of values are numeric and are documented in the Microsoft article [http://msdn.microsoft.com/en-us/library/windows/desktop/bb309232\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/bb309232(v=vs.85).aspx)

Running chkdsk before a volume is brought online is recommended to ensure that the volume is healthy.

*Large volumes may take a long time to complete the check; therefore, it is recommended that the chkdsk be done during planned maintenance by setting the flag to 0 for each DataKeeper volume.

It is highly recommended to change the flag on all nodes in such circumstances. Please refer to the Microsoft blog on this specific issue <http://blogs.technet.com/b/askcore/archive/2012/03/05/understanding-diskskipchkdsk-in-2008.aspx>

The following messages will be logged in the Application Event log:

Message that disk is being checked

Event ID 101: Checking the dirty flag on volume <volume>

Message that volume is dirty and needs to be checked

Event ID 102: Volume <volume> dirty flags is <error number>

Message that an error occurred during volume online

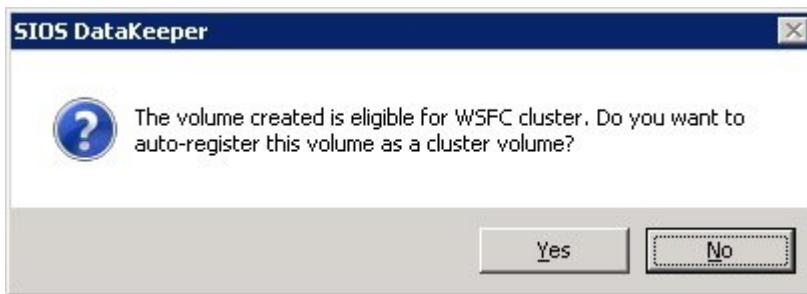
Event ID 80: A failure occurred during the check of volume <volume>.

Error: <error number>. The volume <volume> may be marked as failed.

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.

DataKeeper Volume Resource Private Properties

A DataKeeper Volume resource includes several private properties that are used by DataKeeper Cluster Edition. Among these properties are:

- VolumeLetter (REG_SZ) - the volume letter that is replicated by DataKeeper and associated with this DataKeeper Volume Resource.
- LastSource (REG_SZ) - name of the cluster node that was most recently the source of the mirror.
- NonMirrored (REG_DWORD) - This optional property can be used to configure a non-mirrored storage location for things like SQL tempdb. This private property does not normally exist - it must be manually configured. See [Non-mirrored Volume Resource](#) for more information.
- TargetState_<node> (REG_DWORD) - For each <node> in the system that has at some point been a mirror target, this private property will exist. Its value is the current state (the internal DataKeeper mirror state) of the Owner node's mirror to the given node. Values include:

0 = Node is not currently a target (may be source or shared with a source or target)

1 = Mirroring

2 = Resync

3 = Broken

4 = Paused

5 = Resync Pending

Below is a sample screenshot of the Private Properties of a DataKeeper Volume.

Use the following command in powershell to generate the output:

```
Get-ClusterResource "<DataKeeper Resource Name>" | Get-ClusterParameter
```

Object	Name	Value	Type
DataKeeper Volume E	VolumeLetter	E	String
DataKeeper Volume E	LastSource	CAE-QA-U46.QAGROUP.COM	String
DataKeeper Volume E	DiskSignature	0x45	UInt32
DataKeeper Volume E	DiskRunChkDsk	0	UInt32
DataKeeper Volume E	NonMirrored	0	UInt32
DataKeeper Volume E	TargetState_CAE-QA-U47	1	UInt32

Implications of the TargetState_ <node> value on Failover

DataKeeper Cluster Edition maintains the TargetState_<node> property value, and updates it whenever the mirror state changes. The cluster networks provide multiple paths for a mirror's state to be available on a target system. This enhances the reliability of DataKeeper, and reduces the chance of Split Brain occurrences in the cluster for both synchronous and asynchronous mirrors. The DataKeeper Volume Online method tests to make sure that a node is in a Mirror State which allows Online to proceed. For a mirror target node, that state is "1" (Mirroring). All other states will cause the Online to fail and the DataKeeper Volume resource to be marked as Failed on this node.

In some cases, for instance the catastrophic failure of a previous Owner node that is not going to be recovered for a long period of time, it may be necessary to remove the TargetState_<node> private property in order to force a DataKeeper Volume resource to come online on a node. However, because the node may not have been in a mirroring state with the previous Owner, some Data Loss could occur.

The TargetState_<node> property can be removed using the following command in powershell:

```
Get-ClusterResource "<DataKeeper Resource Name>" | Set-ClusterParameter
-Name "TargetState_<node>" -Delete
```

Note: Servers running Windows 2008 R2 will need to import the 'failoverclusters' module before running these commands. Use the following command in powershell:

```
import-module failoverclusters
```

Behavior of a Synchronous Mirror after Cluster Integration

The purpose of a synchronous mirror is to ensure data consistency between the source and the target at all times. When a synchronous mirror is integrated as a cluster resource, DataKeeper will begin using the cluster to further ensure data consistency.

When the state of a synchronous mirror changes from a mirroring to a non-mirroring state (such as mirrored to paused), DataKeeper will attempt to set the `TargetState_<node>` private property to ensure data consistency in the event of a failover.

DataKeeper, to ensure the consistency of the source's local volume with the targets, will hold all writes that come down to that volume until it has set the `TargetState_<node>` private property to the appropriate value. The writes will then be allowed to continue.

Should DataKeeper be unable to properly set the `TargetState_<node>` private property, it will fail those writes and lock the volume. This ensures that the mirrored volume has the same data consistency as when it was last in the mirroring state.

Extending a Clustered DataKeeper Volume to a Node Outside the Cluster

DataKeeper Volume cluster resources can be extended to a DR Node for Disaster Recovery purposes. In the event of a complete failure of all systems in the cluster, data will be accessible on the DR Node (referred to as the "DR Node"). Here you will find instructions on how to set up this configuration, how to access your data on the DR Node, and how to bring your data back into service in the cluster after the cluster nodes have been restored.

Configuration Tasks

Configuring a non-clustered DataKeeper target node

Recommended configuration for the DR Node

- If possible, the DR Node should be a member of the same domain that the clustered nodes are a member of. Refer to [DataKeeper Service Log On ID and Password Selection](#) for more information about configuring the DataKeeper Service account settings.
- Firewalls (Windows as well as any other firewall devices / software on the DR or cluster site) must allow access to DataKeeper-specified ports on the DR Node from all cluster nodes, and vice-versa. See [Firewall Configurations](#) for more information.
- Configure a volume on the DR Node for each clustered DataKeeper Volume that is to be extended to the DR Node. The volume should be at least as big as the clustered volume.

Scenario 1 – Extending existing DataKeeper Volume resources

If you have already configured DataKeeper Volume resources in your cluster, you can extend these volumes to a DR Node using the DataKeeper MMC GUI by following these steps:

1. Connect the DataKeeper GUI to the DR Node using the "Action / Connect To Server" option.
2. Connect the DataKeeper GUI to the cluster node where the DataKeeper Volume resource is online.
3. For each DataKeeper volume that you are extending to the DR Node:
 - a. In the Jobs view, choose the job that contains the volume to be extended.
 - b. Choose "Create a Mirror".
 - c. Select the mirror Source Node, Volume, and source IP address.
 - d. Select the DR Node as Target, along with the Volume and IP Address.
 - e. Choose the mirror parameters and click "OK" to create the mirror.
 - f. Configure any additional mirror information needed

See [Creating Mirrors with Multiple Targets](#) for more information.

Scenario 2 – Creating a new DataKeeper Volume resource, and extending it to the DR Node.

If you do not have a DataKeeper Volume resources in your cluster that represents the volume that you want to extend to the DR Node, first create the clustered resource, then use the steps in "Scenario 1" above to extend it to the DR Node.

Scenario 3 – Extending a traditional shared-volume cluster to a DR Node using DataKeeper

See [Extending a Traditional 2-node WSFC cluster to a third node using DataKeeper](#) for detailed steps that will guide you through extending a shared-volume Microsoft clustered volume to another cluster node.

In this case, we are extending to a **non-clustered** node. This makes it unnecessary to perform step 2 (configure cluster Quorum settings) and

step 7 (add the node to the cluster). Otherwise, the steps remain the same.

Configuration Summary

After you have extended your clustered volume(s) to a DR Node, you will be able to bring the volumes Online and Offline in the cluster as before. The DR Node will remain the mirror target under normal operating conditions.

If you wish, you can check the data on the non-mirrored system by using the "Pause and Unlock Target" option for the mirror whose target is the DR Node. See [Pause and Unlock](#) for more information.

Accessing Data on the Non-Clustered Disaster Recovery Node

In the event that all of your clustered nodes are unavailable (possibly due to a disaster of some sort at your primary cluster site), you may need to be able to access the data that has been replicated to a DR Node. Use the following procedure to accomplish this.

Note: Please refer to [Switching Over a Mirror](#) guidelines.

Option 1 - using the DataKeeper GUI

1. Start the DataKeeper GUI and connect to the DR Node.
2. Choose the Job that contains the mirror to be made accessible on the DR Node.
3. Choose "Switchover Mirror" to make the DR Node the mirror source, and to make the data on that node accessible.

***NOTE:** If any of the cluster nodes are still running and accessible over the network from the DR Node, the "Switchover Mirror" option will not be available. The DataKeeper GUI will see that the volume is part of a cluster that is still operational, and will prevent a switchover from being selected.

Option 2 - using EMCMD

On the DR Node, start a command prompt and run the commands:

1. `cd ExtMirrBase`
2. `EMCMD . SWITCHOVERVOLUME`

***NOTE:** Use this command with caution; if any cluster nodes are still operational and accessible over the network from the DR Node, EMCMD will NOT prevent the switchover from occurring. This will cause resource failures in the cluster, and the results will be undetermined.

Repeat these actions for all volumes that you need to access on the DR Node. DataKeeper will keep track of all changes that occur while the volume is accessible on this node, and will automatically resync these changes to the cluster nodes when they are brought back up and are accessible from the DR Node. However, the volume resources will not automatically come Online in the cluster - manual steps as outlined in the next section must be taken to move the DataKeeper volumes back into the cluster.

Restoring Data Access to the Cluster

When a cluster node is powered back up after a failure, there are several states that its mirrors can be put into, depending on the mirror state at the time of the original failure, the current network conditions, and the state of other nodes in the cluster. The volume may be in the Source, Target, or None role after all cluster nodes have been restored. You should use the DataKeeper GUI on any of the cluster nodes to determine the mirror role, and to resolve any possible Split-Brain conditions that may exist. See [Split Brain Issue and Recovery](#) for more information. If you are resolving Split-Brain, you should choose the DR Node to be the node that remains source, since it is the one that has the most up-to-date data.

As long as the DR Node is accessible from the cluster node, and the non-cluster node's mirror is in the Source role, any Online request on that cluster node will fail.

Steps to bring the clustered DataKeeper Volume resource back Online

In order to bring the DataKeeper Volume resource Online on a cluster node, the mirror must be switched over to the cluster node where that volume was last Online before the outage (the "Last Source" node for that volume), and the DR Node's volume must be made a Target of the clustered

volume. At that point, the DataKeeper Volume resource can be brought Online on the cluster node.

To determine which cluster node is the Last Source node for a particular volume, run either of the following commands on any of the cluster nodes:

- (to use cluster.exe) - cluster res "<DataKeeper Volume Resource name>" -priv
- (to use powershell) - get-clusterresource -Name "<DataKeeper Volume Resource name>" | get-clusterparameter
The output produced should include a line with the value "LastSource" listed. The name of the Last Source node is given on that line.

Follow these steps to bring the resource Online:

1. If any DataKeeper Volume resources are already Online in the cluster, take them Offline. This is necessary in order to resolve Split Brain conditions in the next step.
2. Start the DataKeeper GUI on one of the cluster nodes. Resolve any Split Brain conditions, choosing the DR Node as the mirror source.
3. Monitor the state of the mirrors that have been created from the DR Node (Mirror Source) to the clustered nodes (Targets). If any clustered nodes are shared, only one of them will be a mirror target.
4. Once the mirror from the DR Node (Mirror Source) to the LastSource cluster node has reached the Mirroring state, the LastSource node can be made Source.
 - a. Open a command prompt on that cluster node
 - b.. Run the command `cd ExtMirrBase`
 - c. Run the command `EMCMD . SWITCHOVERVOLUME`

Repeat these steps for each volume. If multiple volumes are part of the same resource group, be sure to switch each of them over to their Last Source node.

Then, using Failover Cluster Manager, bring the volumes and associated applications / roles Online.

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.



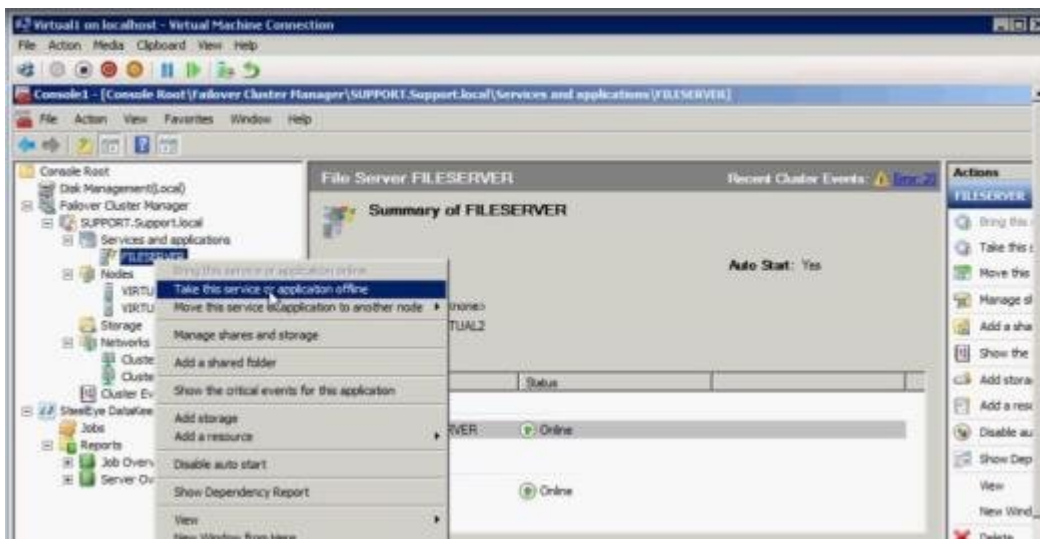
1. Remove physical disk resources from WSFC.

This also removes any dependencies on these physical disk resources. These dependencies will need to be reestablished to the new

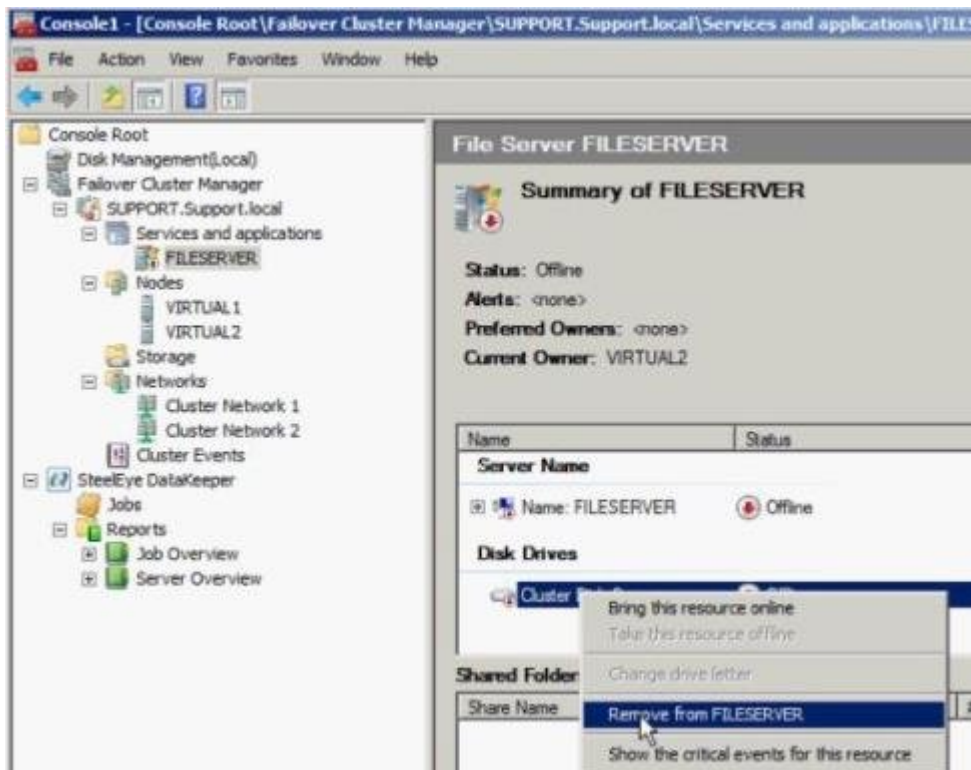
DataKeeper volume resource, so please take note of what these dependencies are before completing this first step by viewing the **Dependency Report**. Highlight your resource and select **Show Dependency Report**.

Use **Failover Cluster Manager MMC** to perform the following steps:

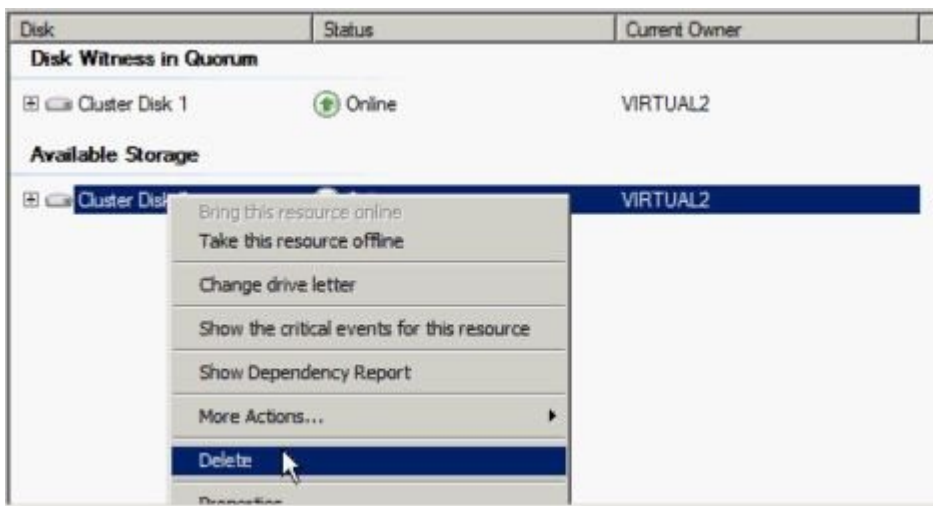
- a. Offline the cluster resource group by right-clicking and selecting **Take this service or application offline**.



- 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...**



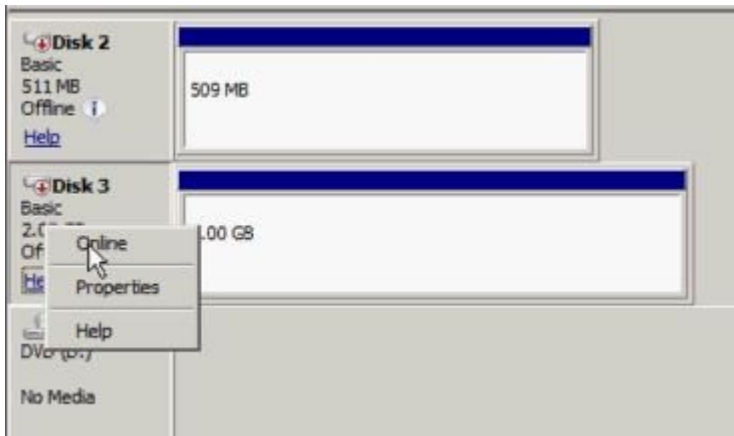
- b. Select **Node Majority**. **Note:** Change the quorum type to **Node Majority** if the final number of nodes is odd as in a 3-node cluster, or **Node and File Share Witness Majority** if the final number of nodes will be even as in a 4-node cluster.

- c. Delete disk witness from **Available Storage**.

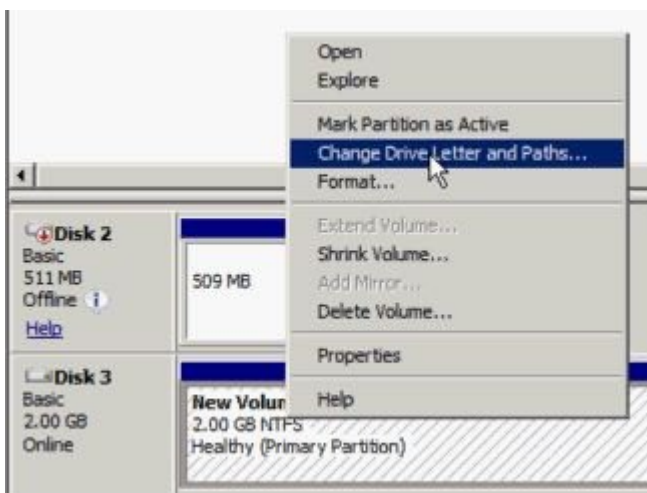
3. Online Disk.

Use **Disk Management MMC** to perform the following steps:

- a. Online physical disk.



b. Change drive letter to match previous configuration (if needed) .

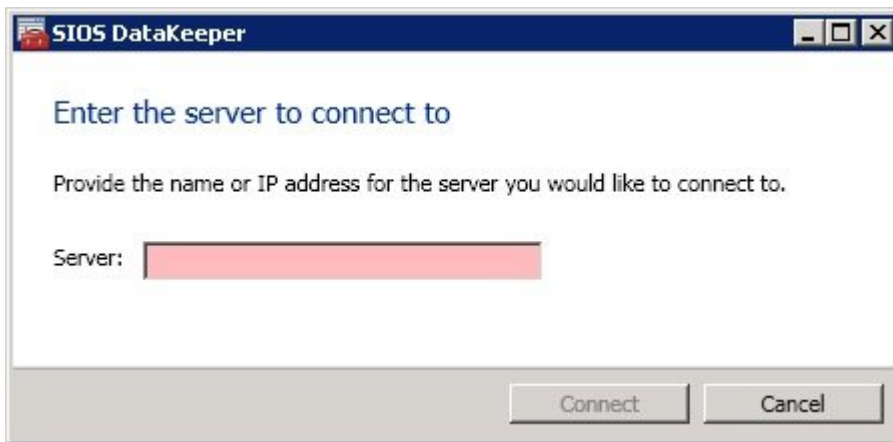


c. Repeat the above steps, 2a and 2b, on all shared nodes in the cluster.

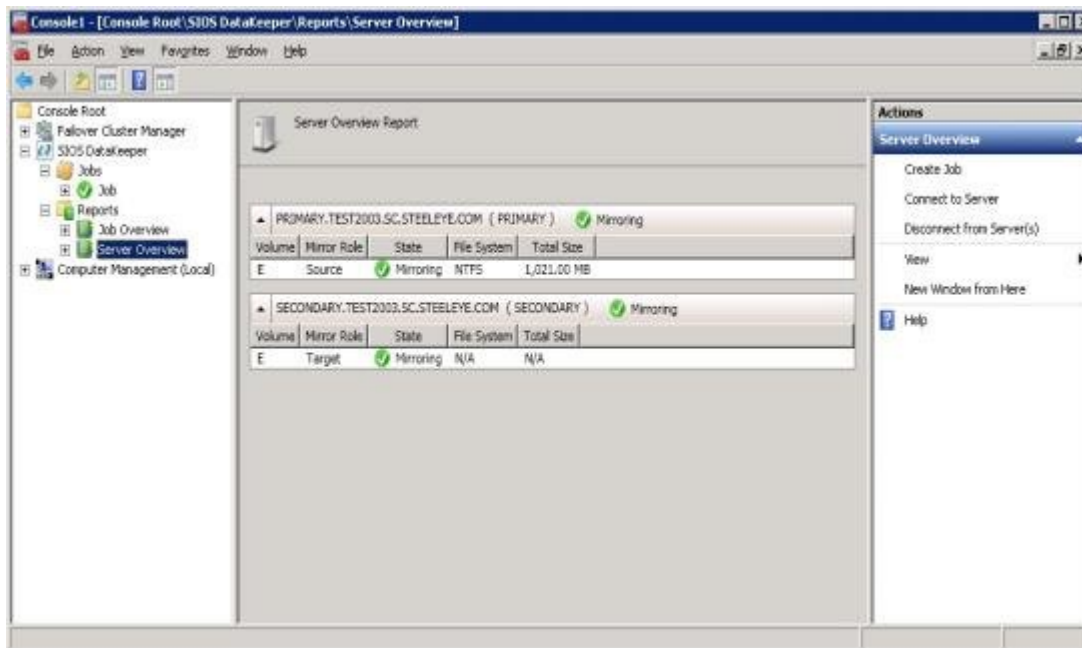
4. Make sure the volume of the third node is available to become a target. It should be formatted and online, and it should be at least as big or bigger than the source volume.
5. [Create Mirror](#).

Use the DataKeeper UI to perform the following steps (make sure DataKeeper Service is running on all servers):

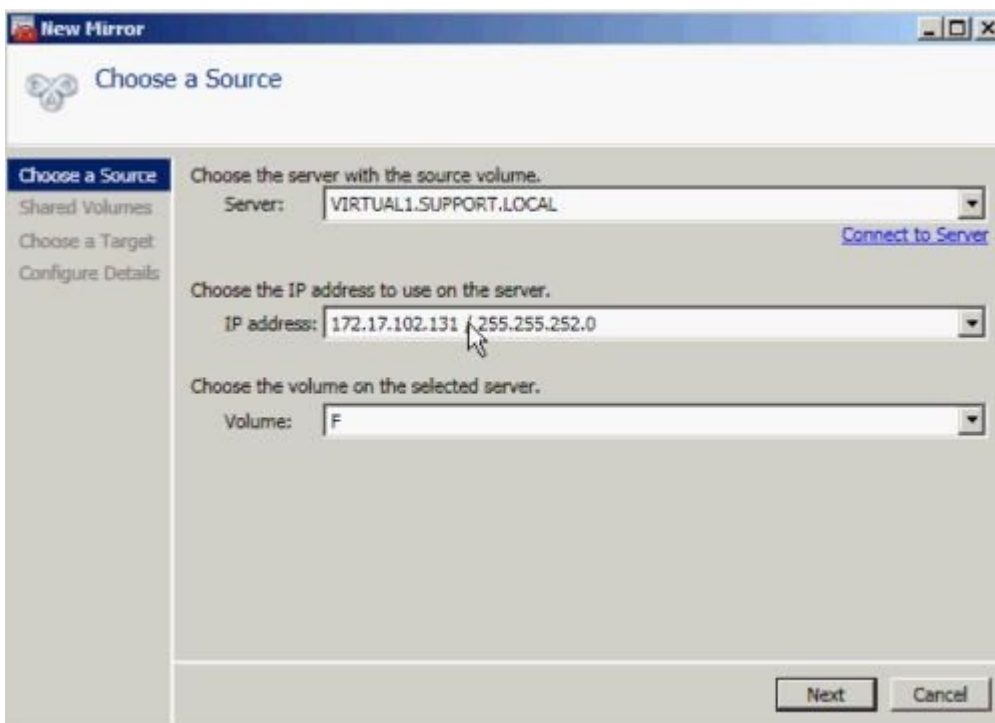
a. Connect to all shared nodes and the third node.



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](#)

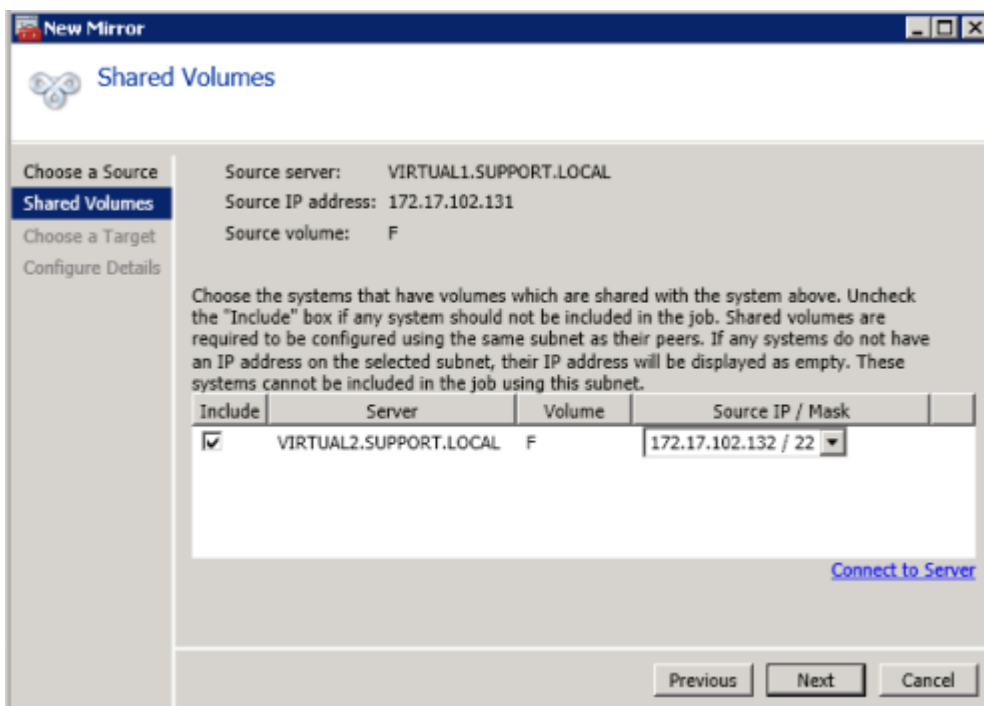
Choose the IP address to use on the server.

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 block explains the requirements for shared volumes. A table lists the systems to be included in the job. 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

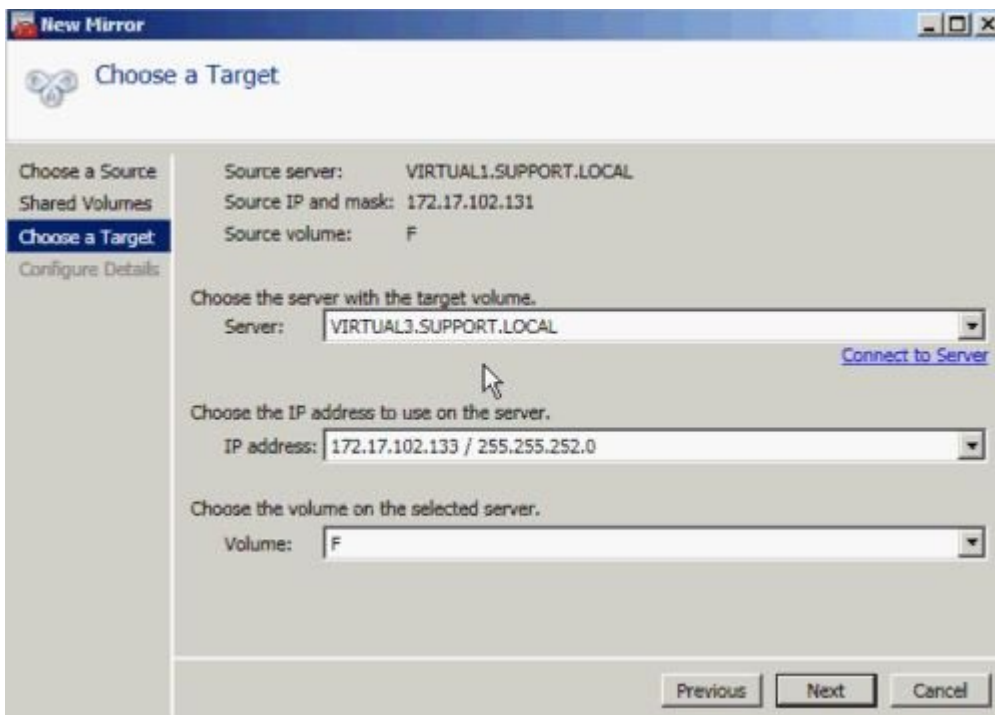
Source volume: F

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.

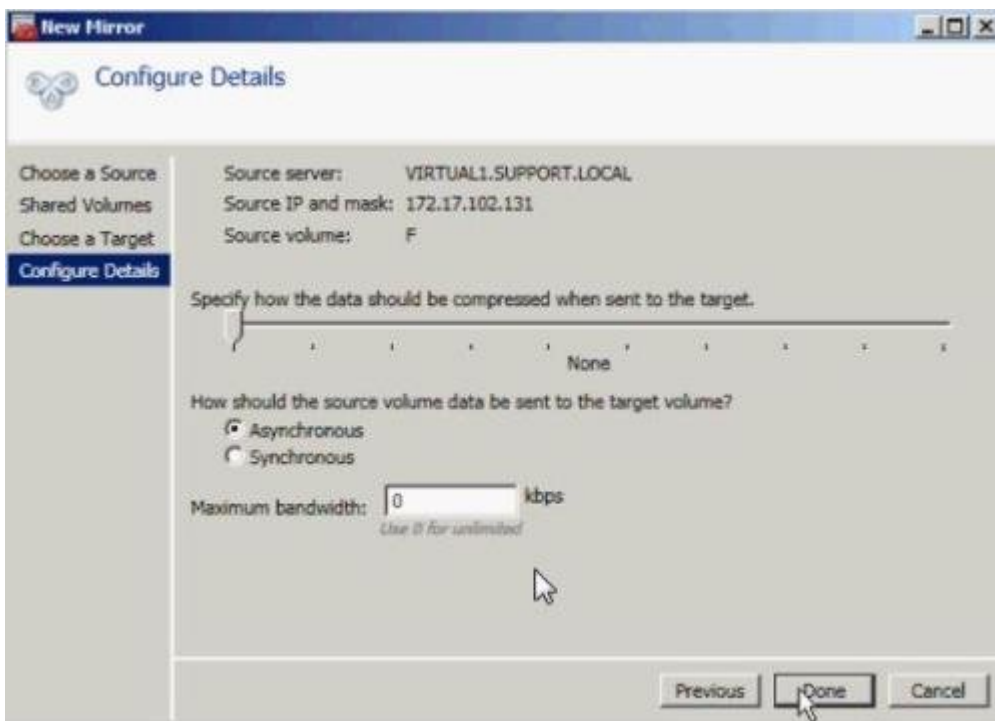
Include	Server	Volume	Source IP / Mask
<input checked="" type="checkbox"/>	VIRTUAL2.SUPPORT.LOCAL	F	172.17.102.132 / 22

[Connect to Server](#)

c. Select a **Target**.

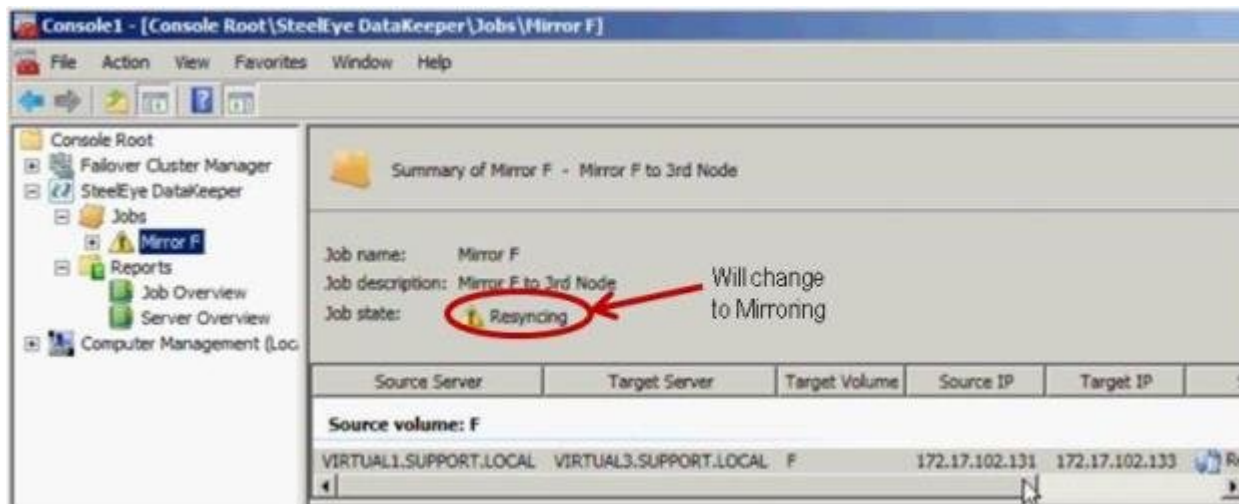


f. **Configure Details.** Choose compression settings, if applicable, and choose mirror type, [asynchronous or synchronous](#). Select **Done**.



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:

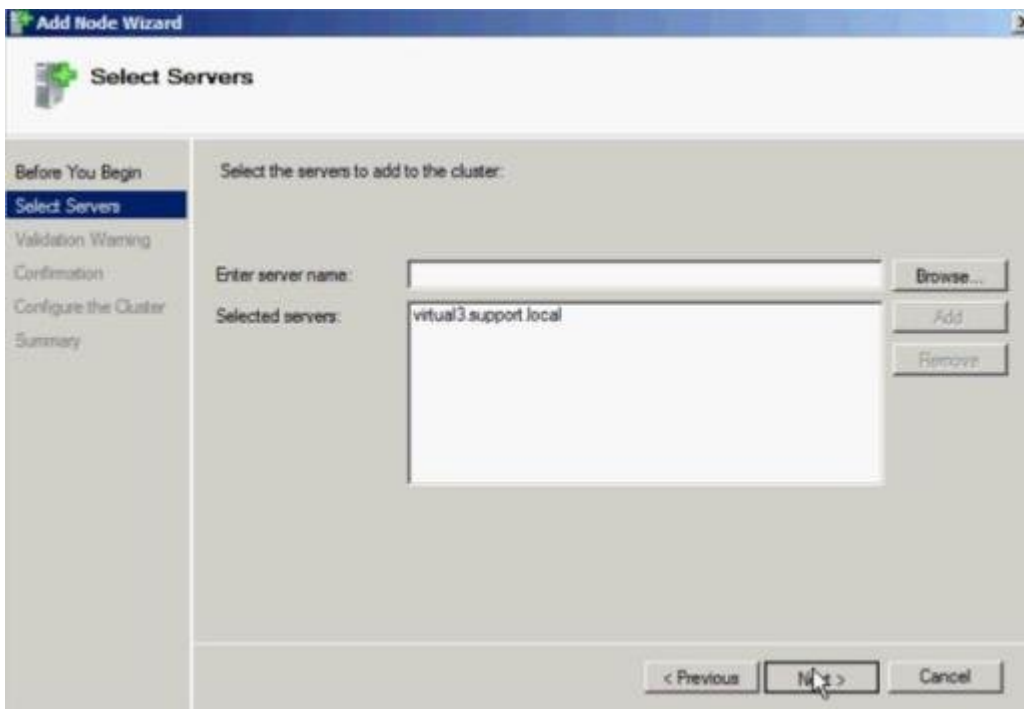
[Automatic Creation of a Mirror in WSFC](#)

[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):

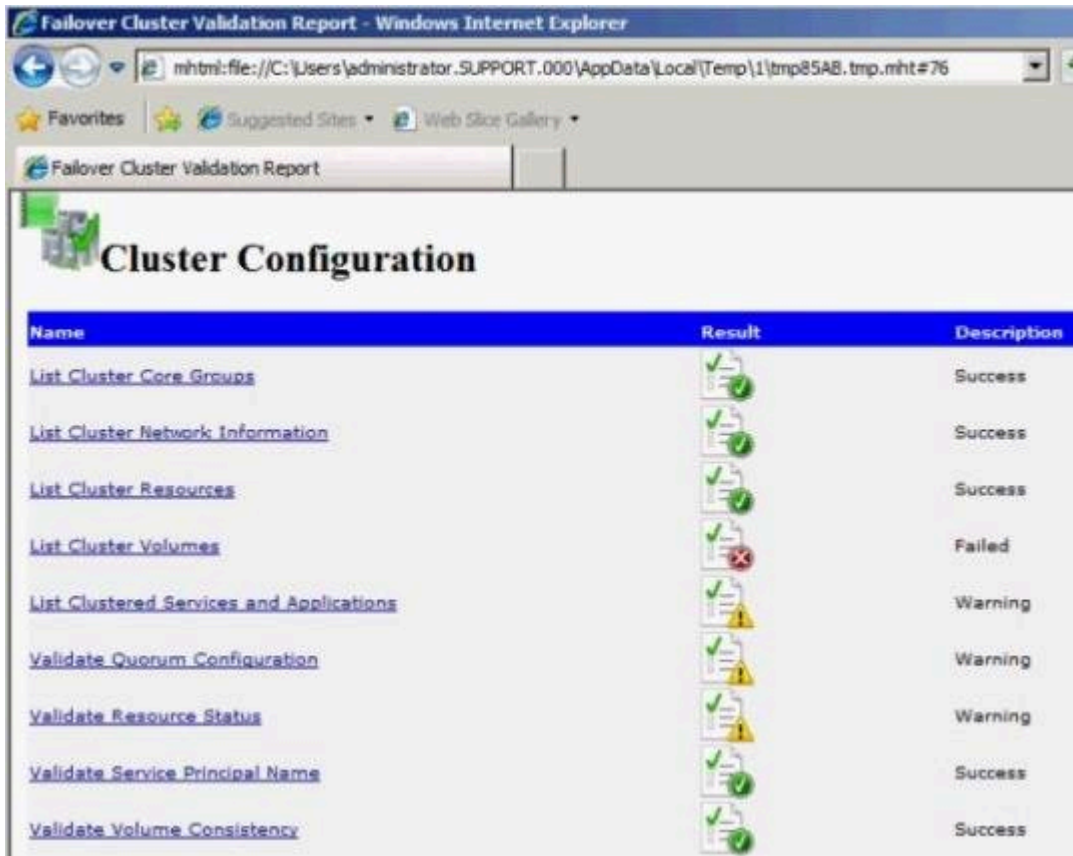
- a. Right-click **Nodes**.
- b. Select **Add Node**.
- c. In the **Add Node Wizard**, enter name of server to be added, click **Add** and select **Next**.



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.



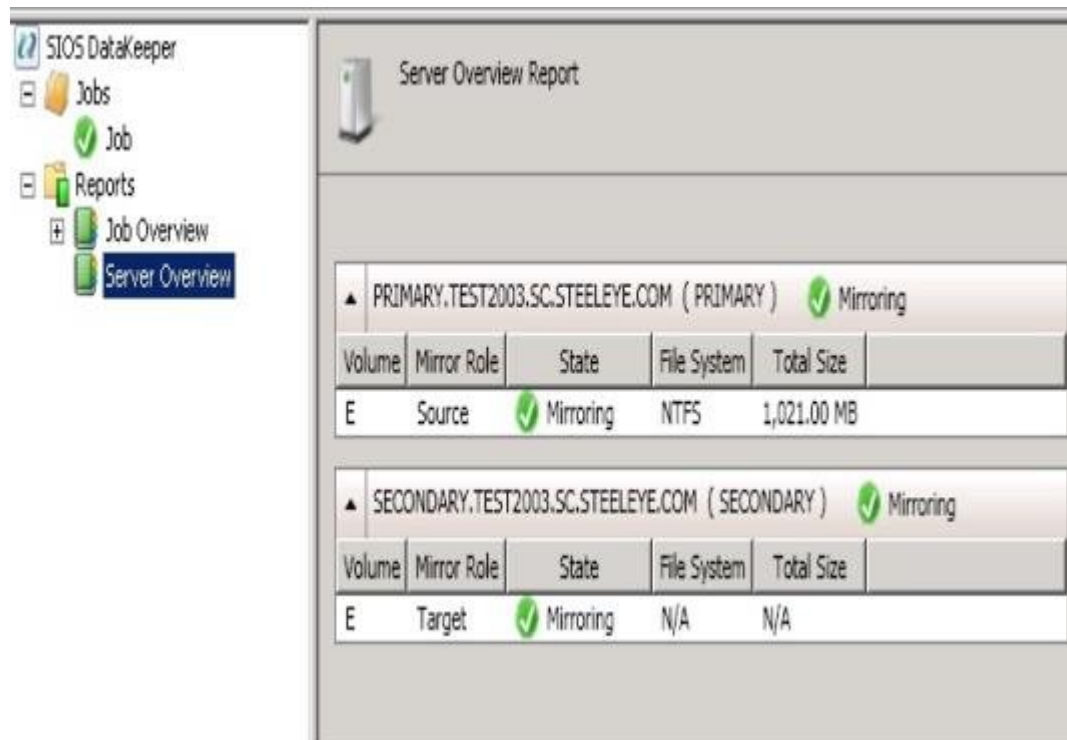
Name	Result	Description
List Cluster Core Groups		Success
List Cluster Network Information		Success
List Cluster Resources		Success
List Cluster Volumes		Failed
List Clustered Services and Applications		Warning
Validate Quorum Configuration		Warning
Validate Resource Status		Warning
Validate Service 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.

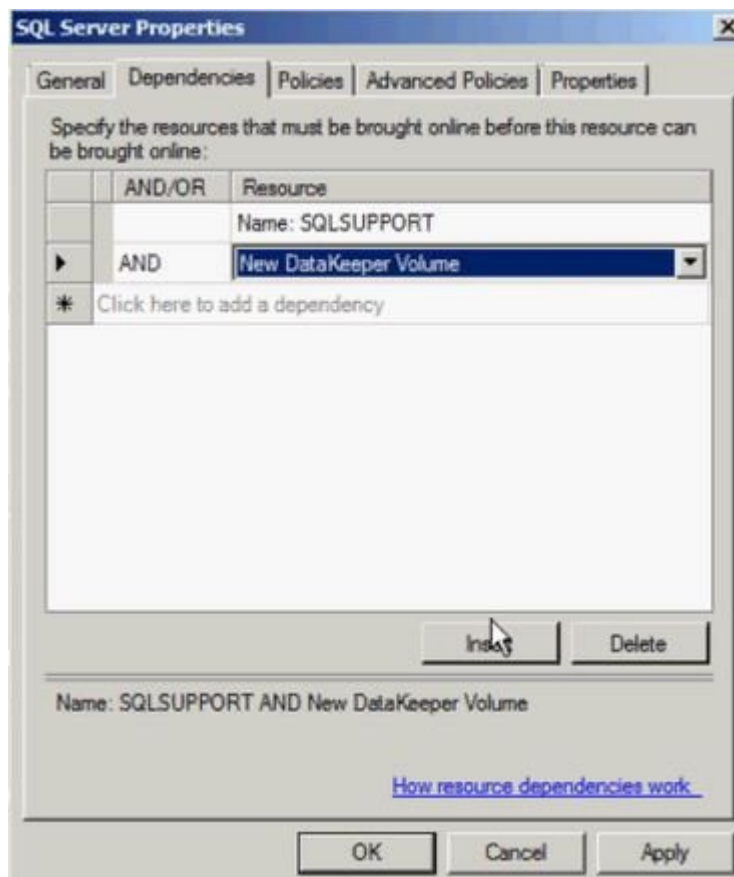
f. Select **Next** on the **Confirmation** screen.

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.

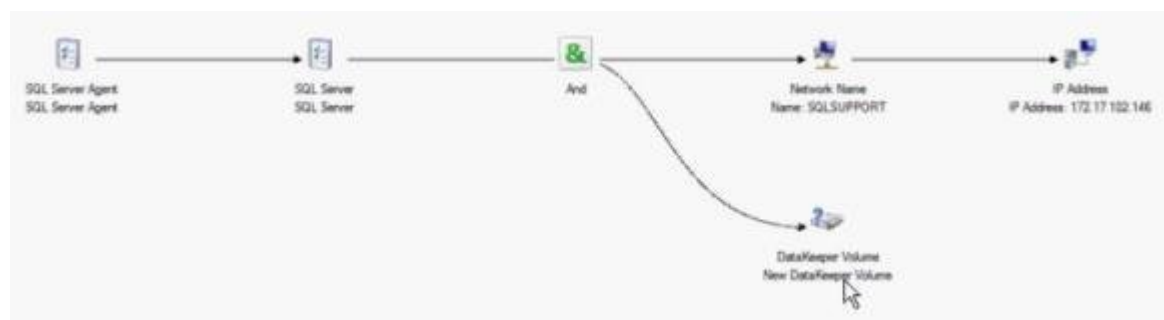


8. Reestablish any/all resource dependencies to the new DataKeeper volume resource.
 - a. Right-click application resource and select **Properties**.



- Select the **Dependencies** tab.
- Click on the **Click here** to add a dependency tab.
- Enter the new DataKeeper volume
- 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

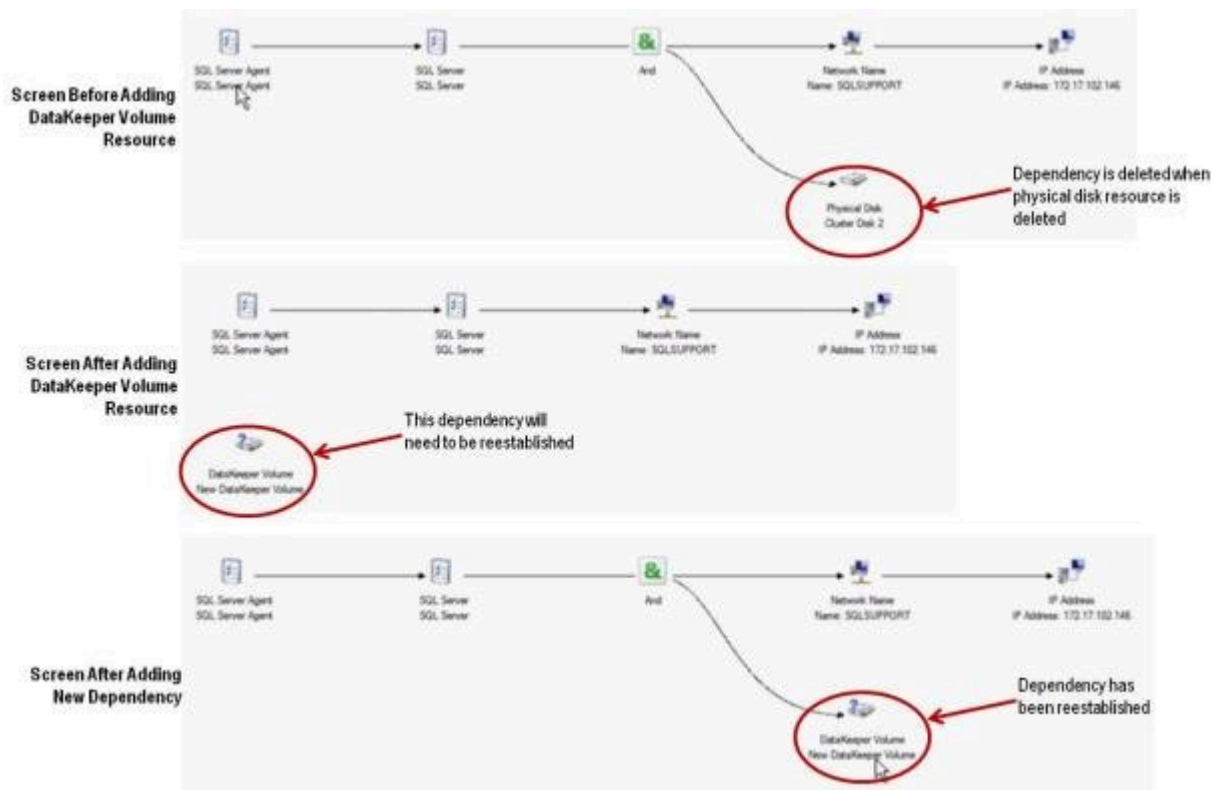
When replicating from a WSFC SQL Server 2-node cluster to a third node via DataKeeper, the following tasks are required:

- 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.

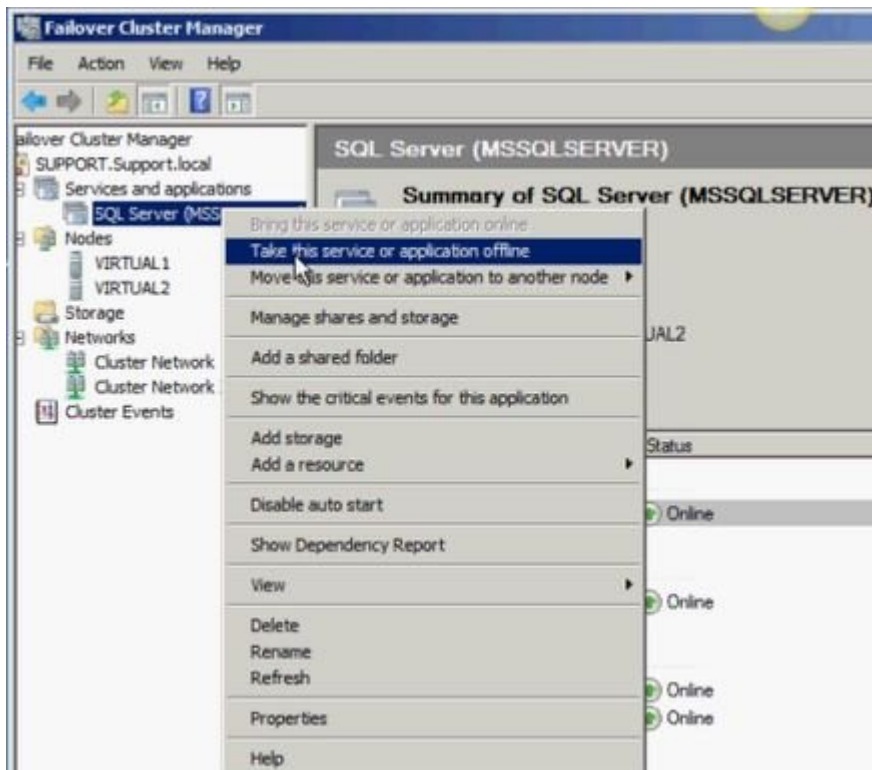
1. Remove physical disk resource from WSFC.

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**.

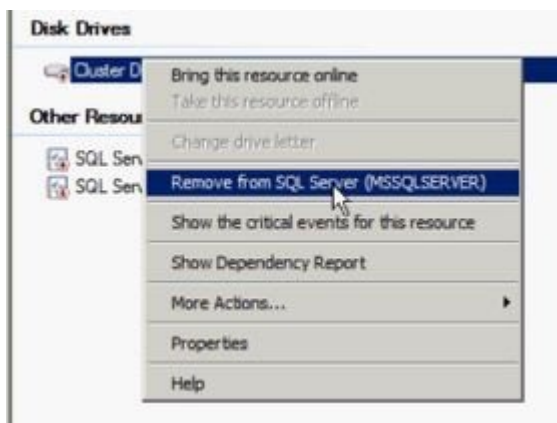


Use Failover Cluster Manager MMC to perform the following steps:

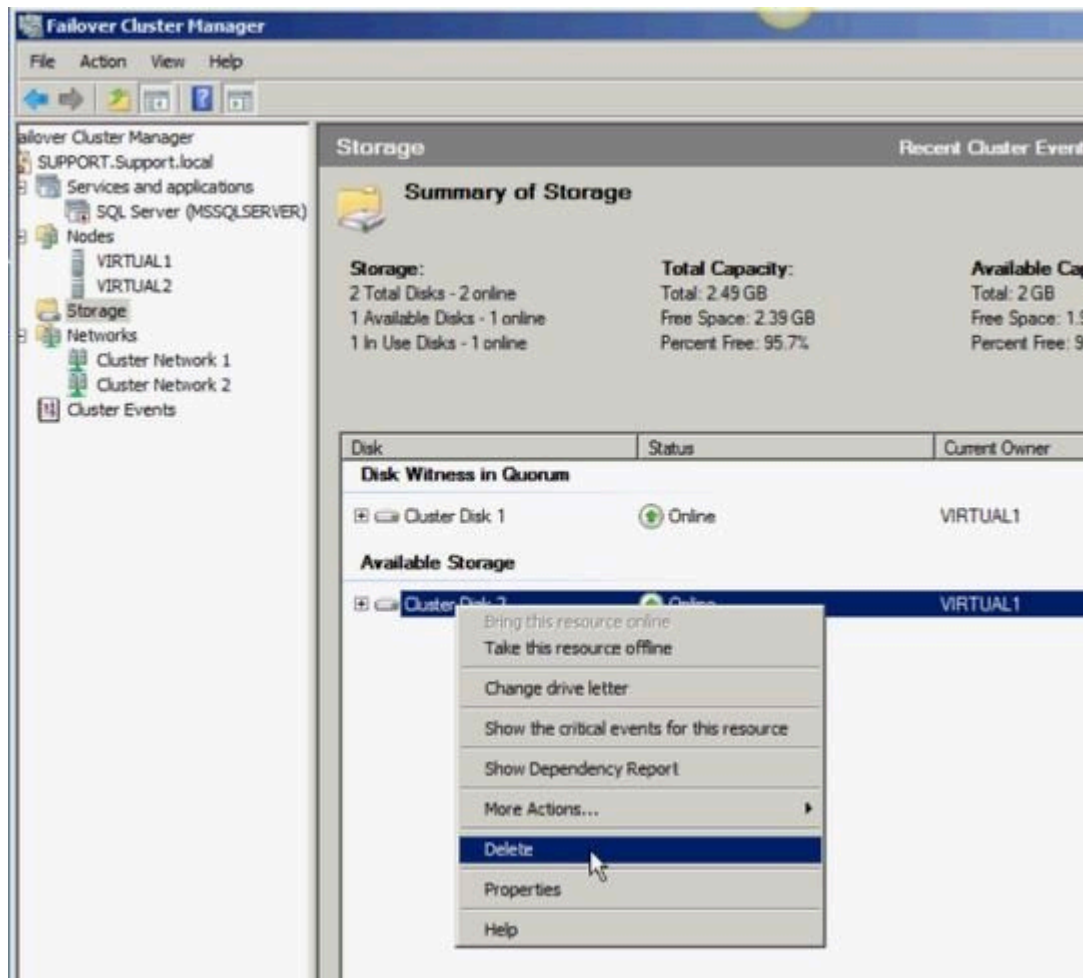
- a. Offline the SQL Server cluster resource group by right-clicking and selecting **Take this service or application offline**.



b. Remove the physical disk from the SQL Server cluster resource group (*moves to Available Storage*).



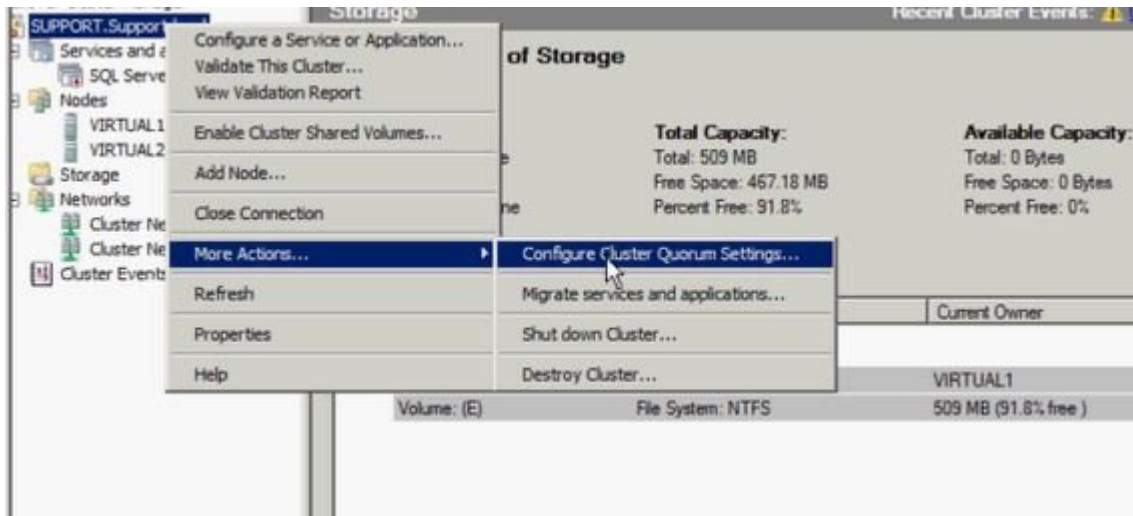
c. Remove the physical disk resource from the SQL Server cluster configuration by deleting the resource from the **Available Storage** group.



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...**



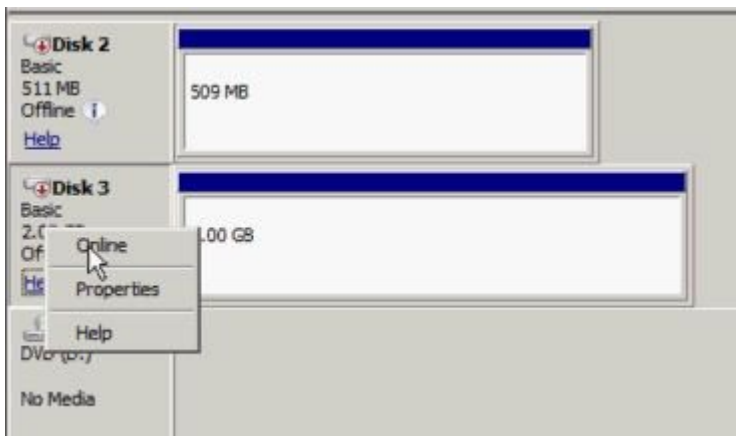
b. Select **Node Majority**. **Note:** Change the quorum type to **Node Majority** if the final number of nodes is odd as in a 3-node cluster, or **Node and File Share Witness Majority** if the final number of nodes will be even as in a 4-node cluster.

c. Delete disk witness from **Available Storage**.

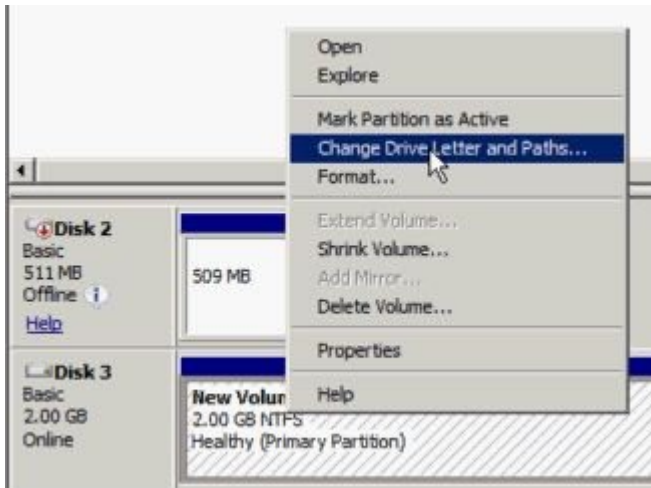
3. Online Disk.

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).



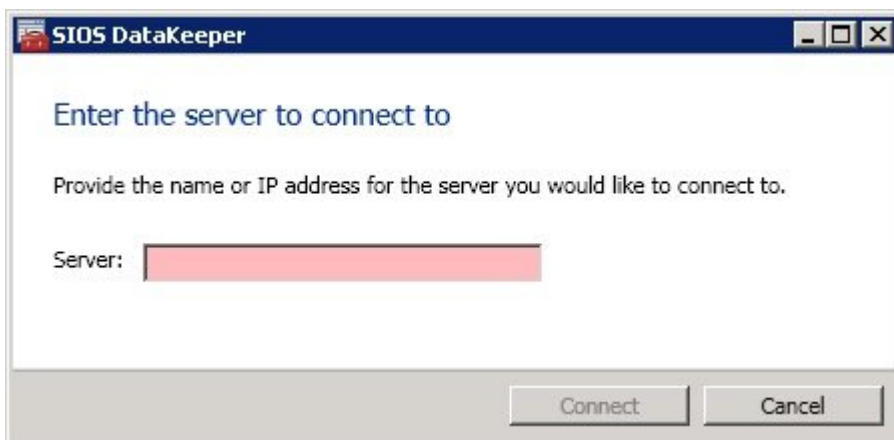
c. Repeat the above steps, 2a and 2b, on all shared nodes in the cluster.

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.

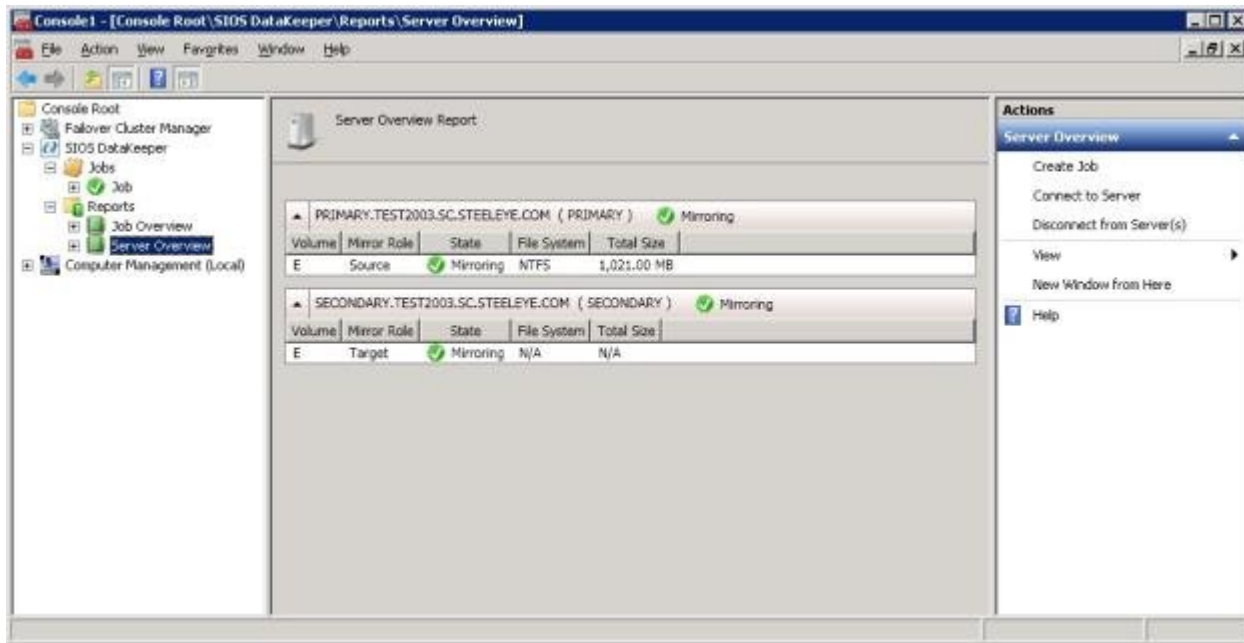
5. **Create Mirror.**

Use the DataKeeper UI to perform the following steps (**Note:** Make sure DataKeeper is running on all servers):

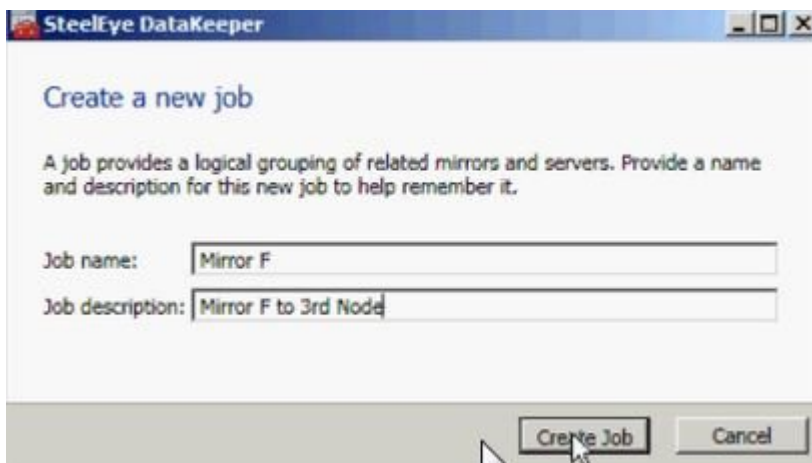
- a. Connect to all shared nodes and the third node.



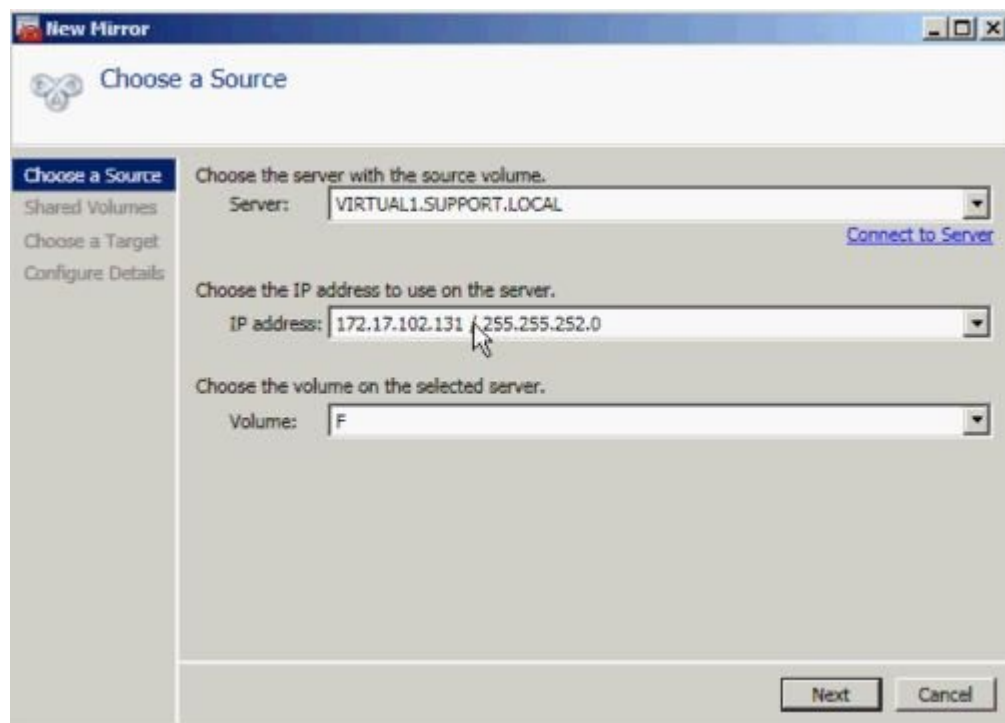
The Server Overview Report will show connection to all three nodes:



b. Create a job containing a mirror to the third node.

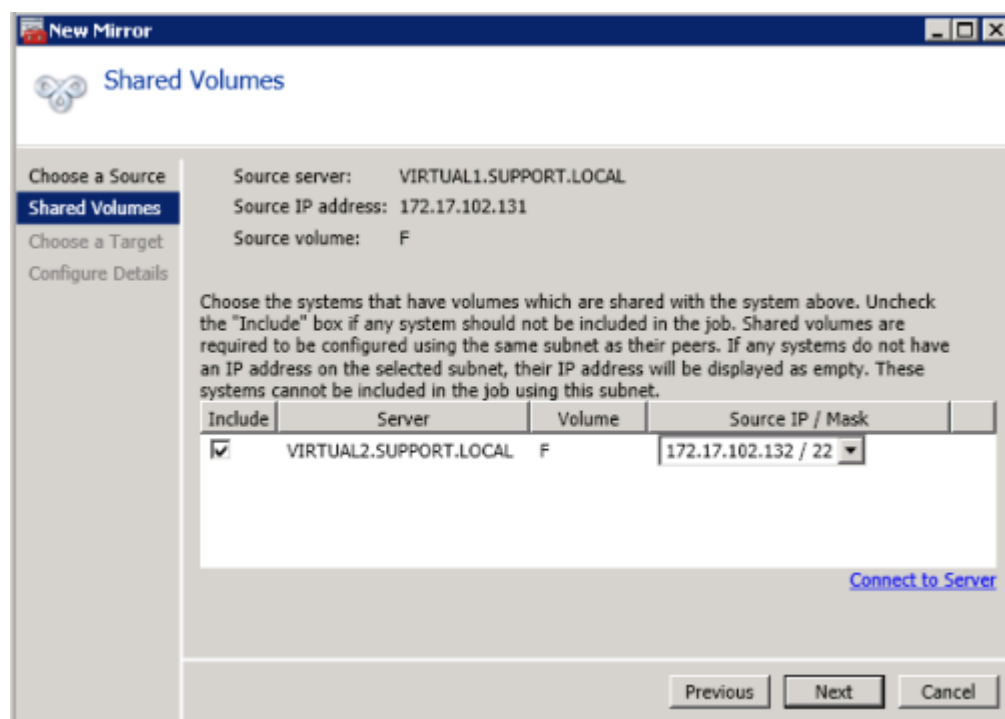


c. Select a **Source**.



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**.

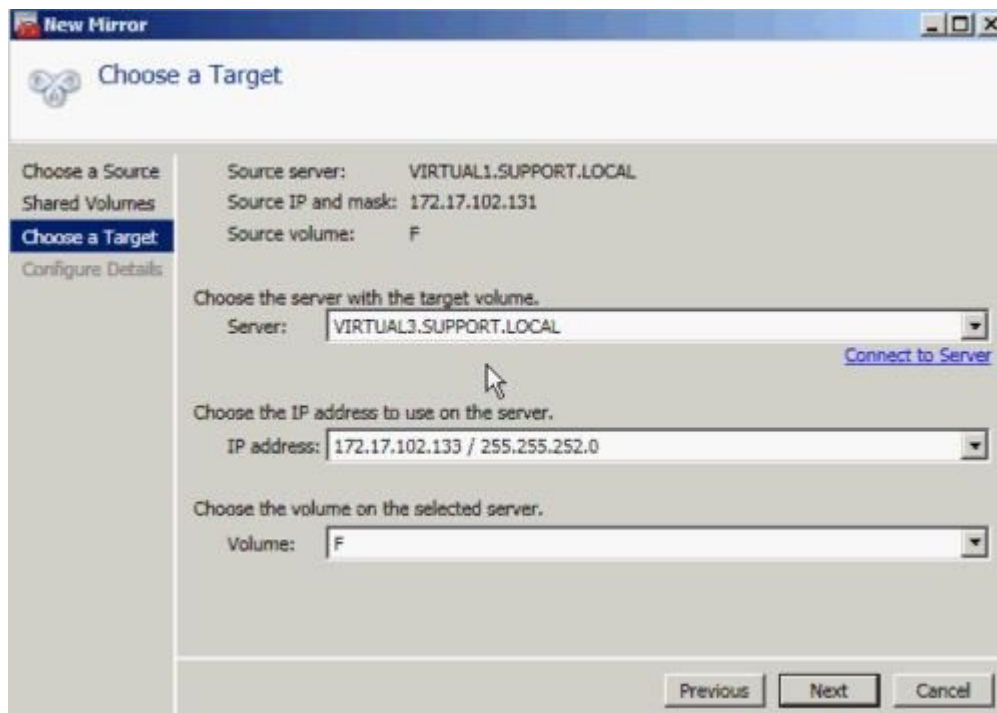


The "New Mirror" dialog box is shown with the "Shared Volumes" tab selected. The left sidebar contains "Choose a Source", "Shared Volumes", "Choose a Target", and "Configure Details". The main area displays the source configuration: "Source server: VIRTUAL1.SUPPORT.LOCAL", "Source IP address: 172.17.102.131", and "Source volume: F". Below this is a text block explaining the "Include" checkbox and subnet requirements. A table lists systems to include:

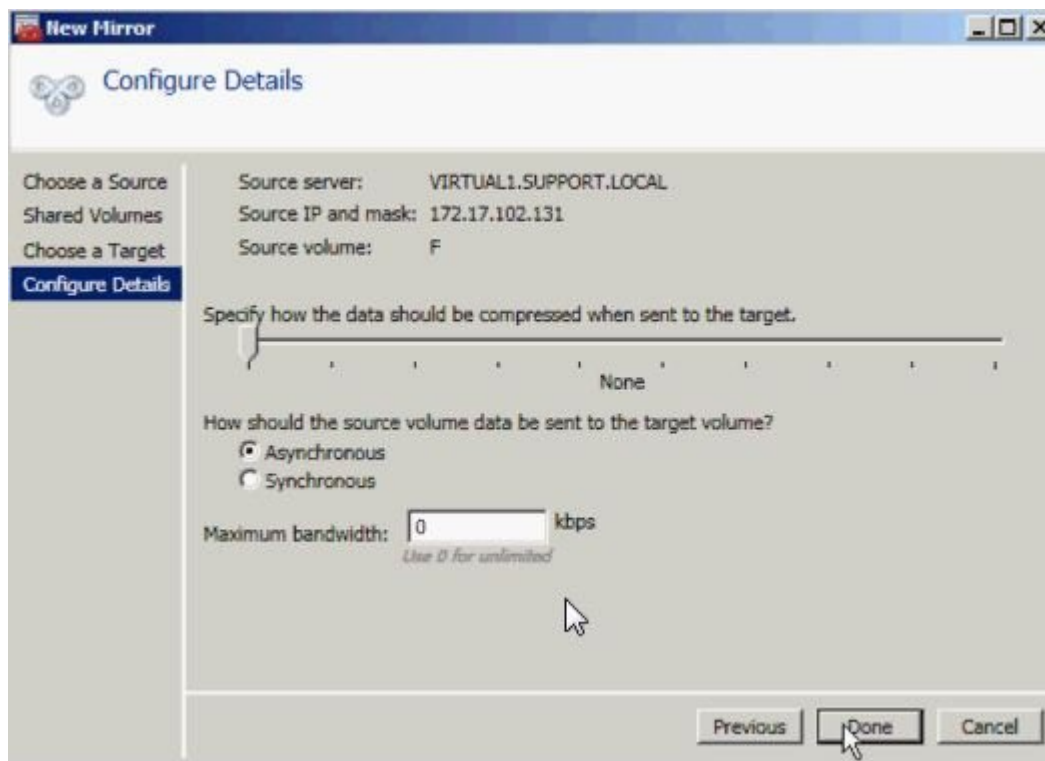
Include	Server	Volume	Source IP / Mask
<input checked="" type="checkbox"/>	VIRTUAL2.SUPPORT.LOCAL	F	172.17.102.132 / 22

Below the table is a "Connect to Server" link. "Previous", "Next", and "Cancel" buttons are at the bottom right.

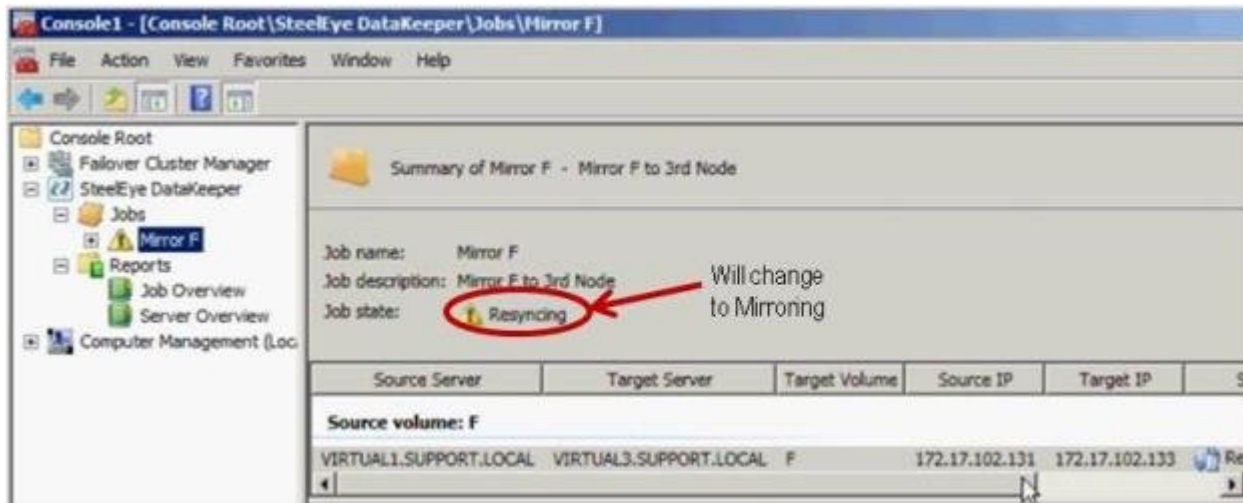
e. Select a **Target**.



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 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](#)

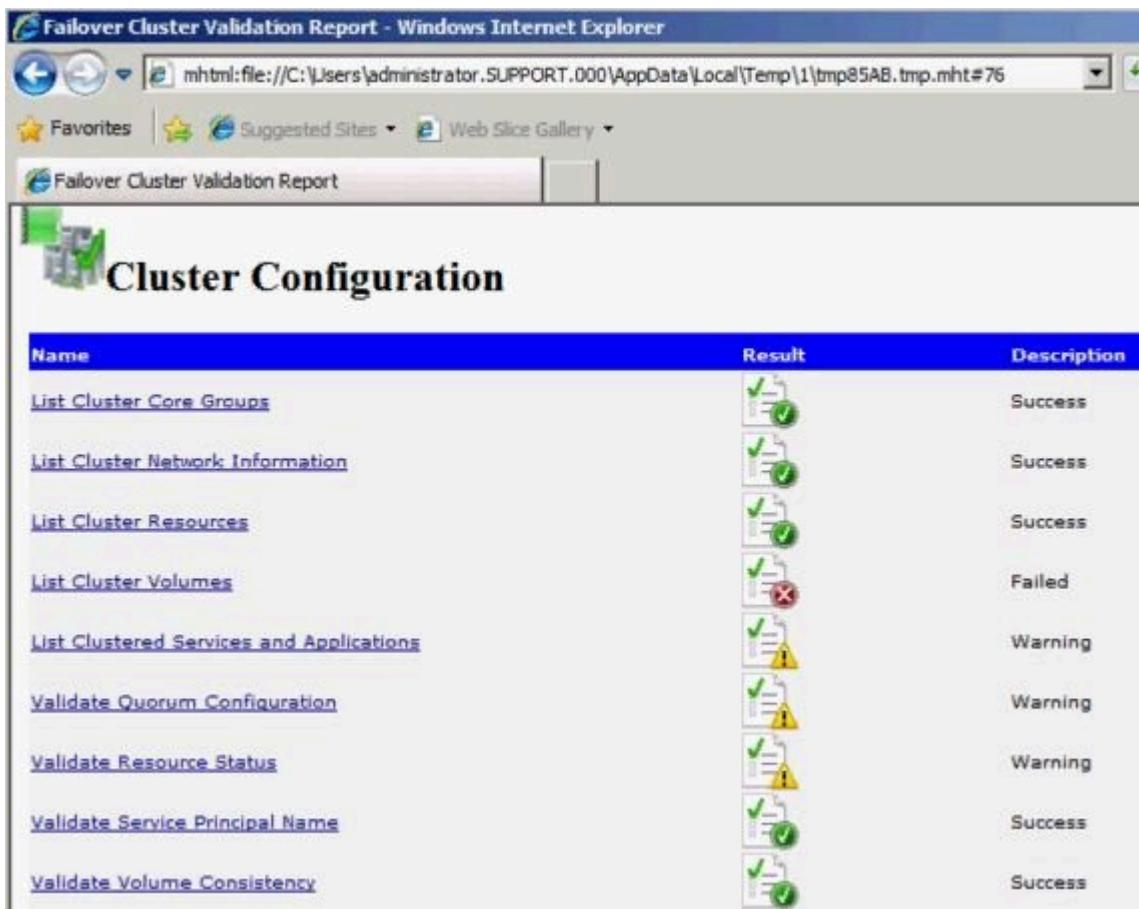
7. Add third node into the cluster.

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.

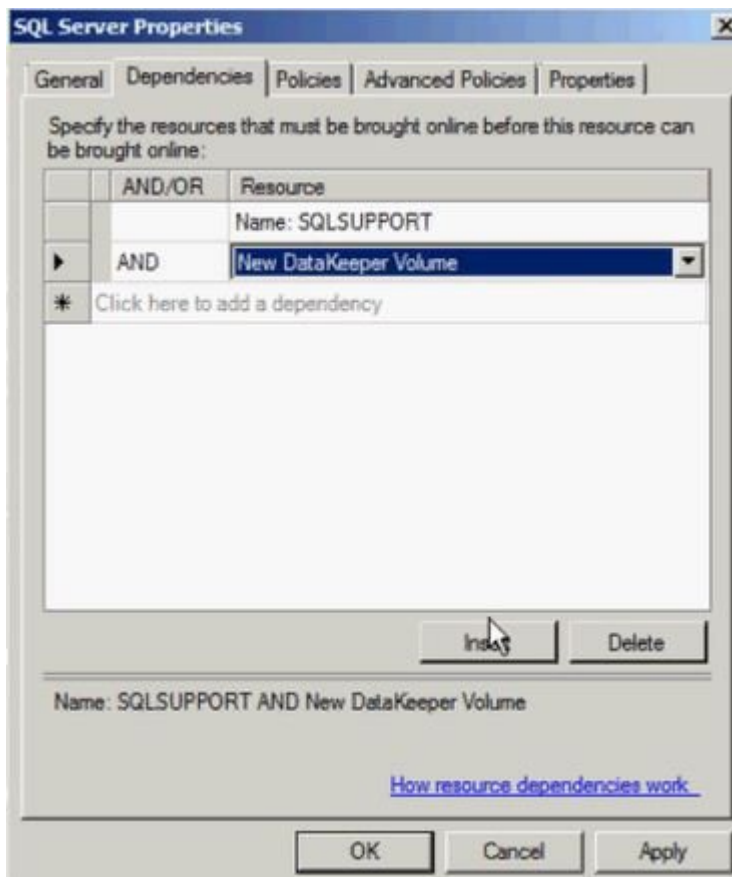
Service	Account Name	Password	Startup Type
SQL Full-text Filter Daemon Launcher	NT AUTHORITY\LOCALSER...		Manual
SQL Server Database Engine	support\administrator	*****	Manual
SQL Server Browser	NT AUTHORITY\LOCAL SER...		Automatic
SQL Server Agent	support\administrator	*****	Manual

Use the same account for all SQL Server services

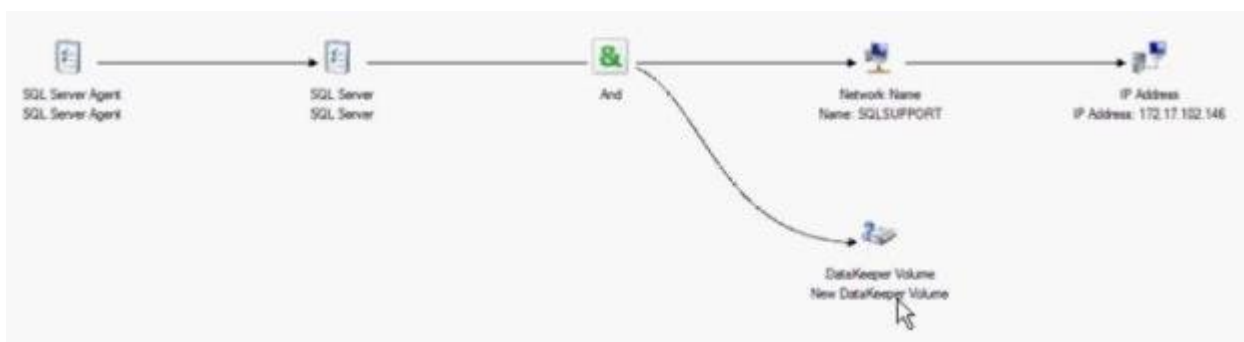
< Back Next > Cancel Help

- 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>"


Using DataKeeper Cluster Edition to Enable Multi-Site Hyper-V Clusters

Prerequisite

Familiarity with Microsoft Windows Server, Microsoft Windows Failover Cluster Server Management and Hyper-V Virtual Management procedures and commands is highly recommended.

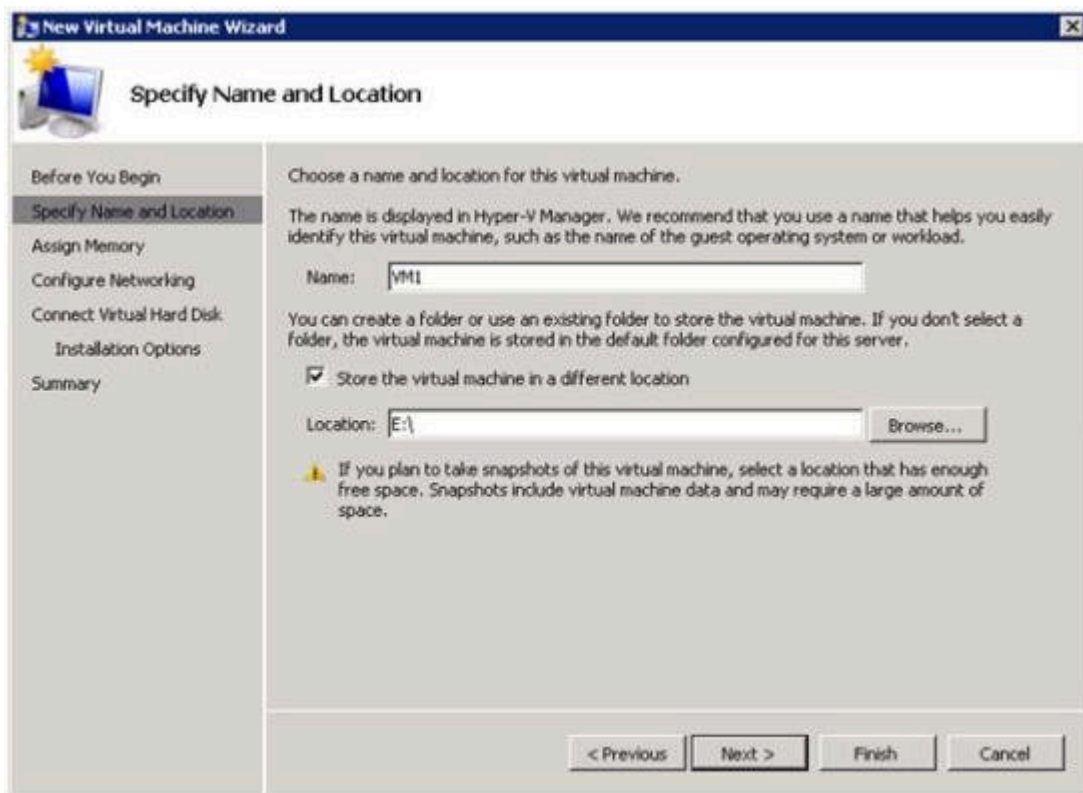
Note: DataKeeper ClusterCluster Edition does not support Cluster Shared Volumes; therefore, when configuring Hyper-V with DataKeeper Cluster Edition, you must have one volume per virtual machine.

1. Before you begin, all the Microsoft Windows Servers in your cluster must be installed and configured as outlined in other topics for this product (refer to "[Creating a DataKeeper Volume in WSFC](#)"). In addition, the following products must be installed and configured on each server in your cluster.

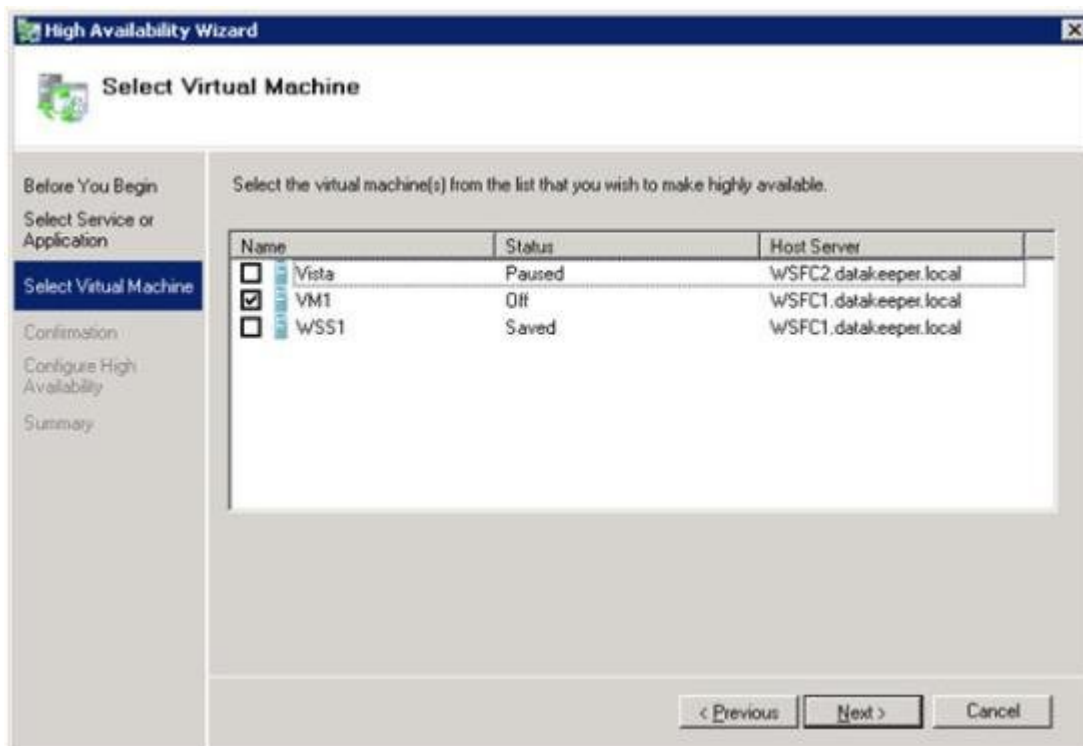
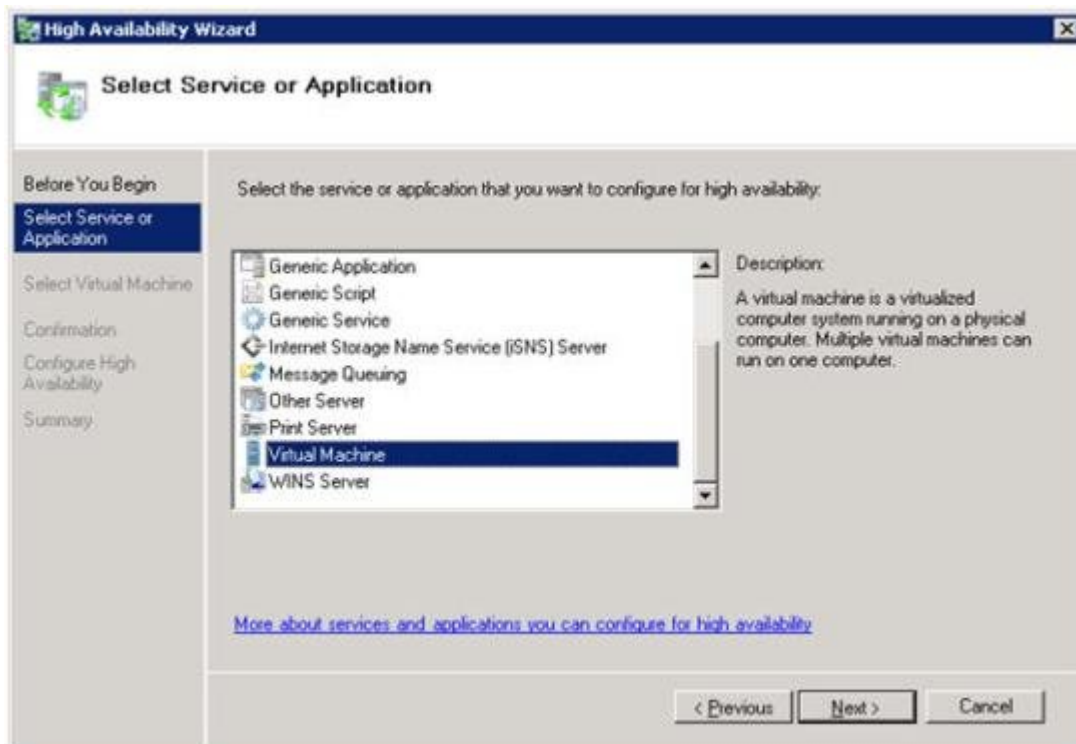
Operating Systems	Microsoft Windows Server 2008 R2 or Later
Hardware Platform	64-bit environment and supports hardware-assisted virtualization (Intel VT) technology
Virtual Management Software	Hyper-V role and all Hyper-V updates
SIOS DataKeeper License	A DataKeeper ClusterCluster Edition license key is required for every server on which DataKeeper runs. This applies to both physical and virtual servers
Network Configuration	<div><div>A high speed WAN link between the primary data center and the DR site (see performance benchmark on bandwidth configuration)</div><div> IMPORTANT: Make sure that Virtual Network Names for NIC connections are identical on all cluster nodes.</div></div>
Cluster Server Software	Windows Server 2008 R2, Server 2012, or later Failover Cluster Management

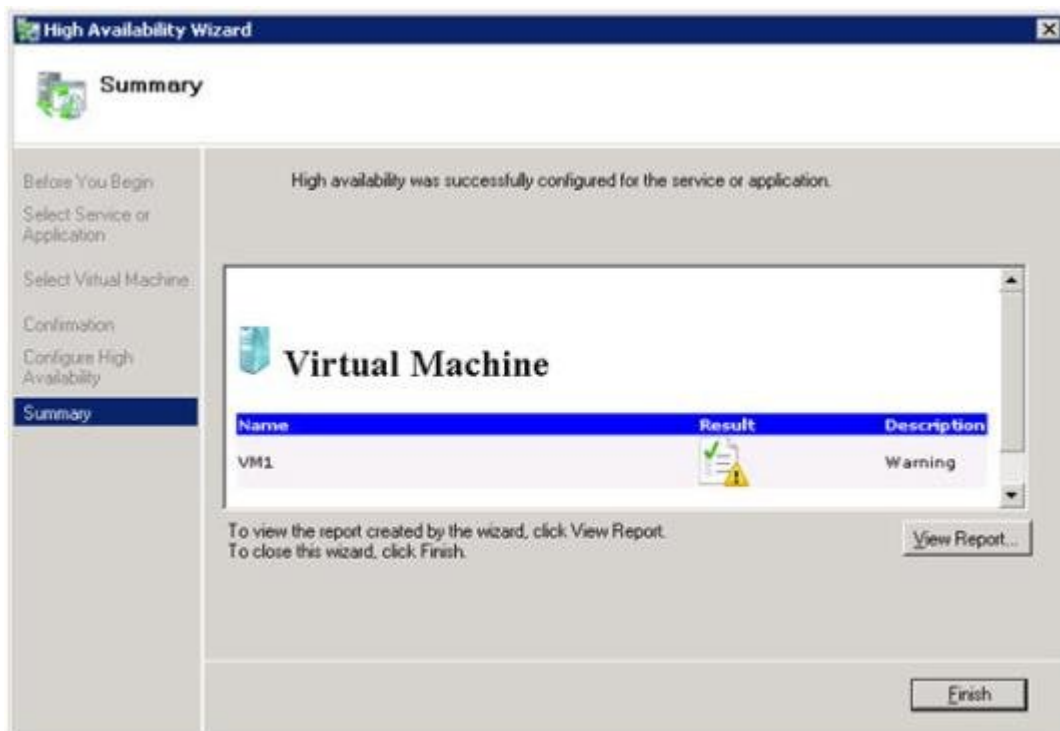
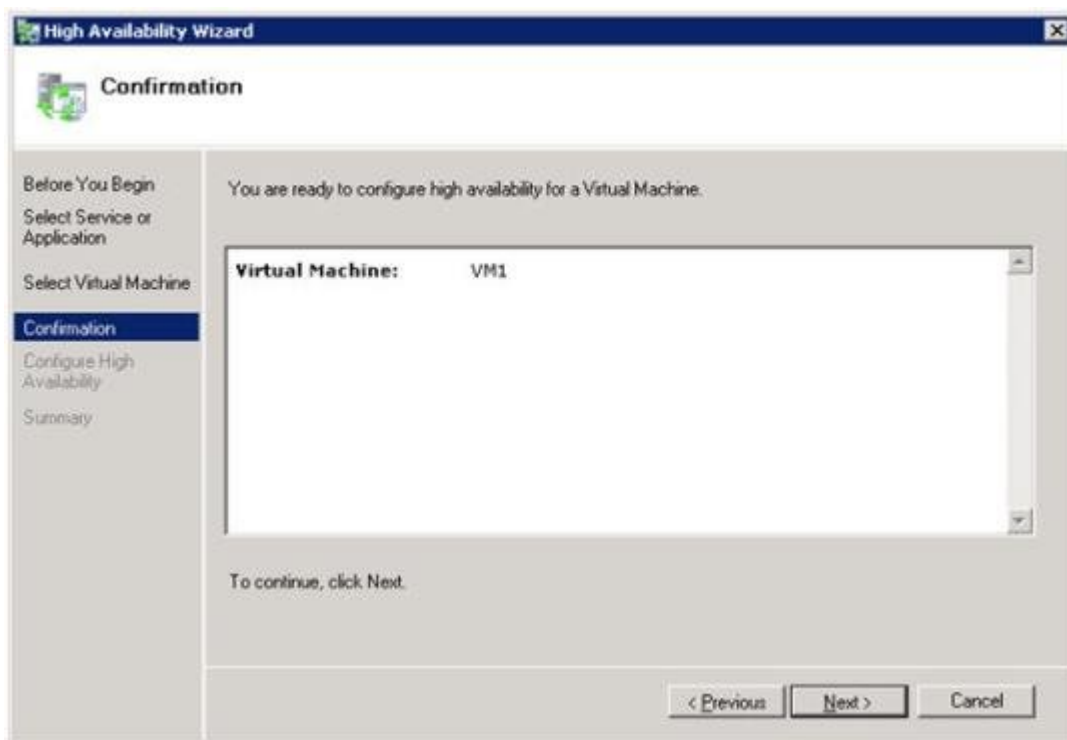
***IMPORTANT:** DataKeeper Cluster Edition registration with Failover Cluster is automatic and occurs 60 seconds after detecting a Failover Cluster configuration on each node.

2. Use the SIOS DataKeeper GUI to configure a job that includes all mirrors and shared relationships for the nodes that are part of the cluster. Wait until the mirrors are in the **Mirroring** state. For more information about setting up your SIOS DataKeeper environment, see the [Getting Started](#) topic.
3. Using Hyper-V, configure the first virtual machine and make sure to specify E:\ in the **Store the virtual machine in a different location** text box. This setting is in the **Specify Name and Location** window.

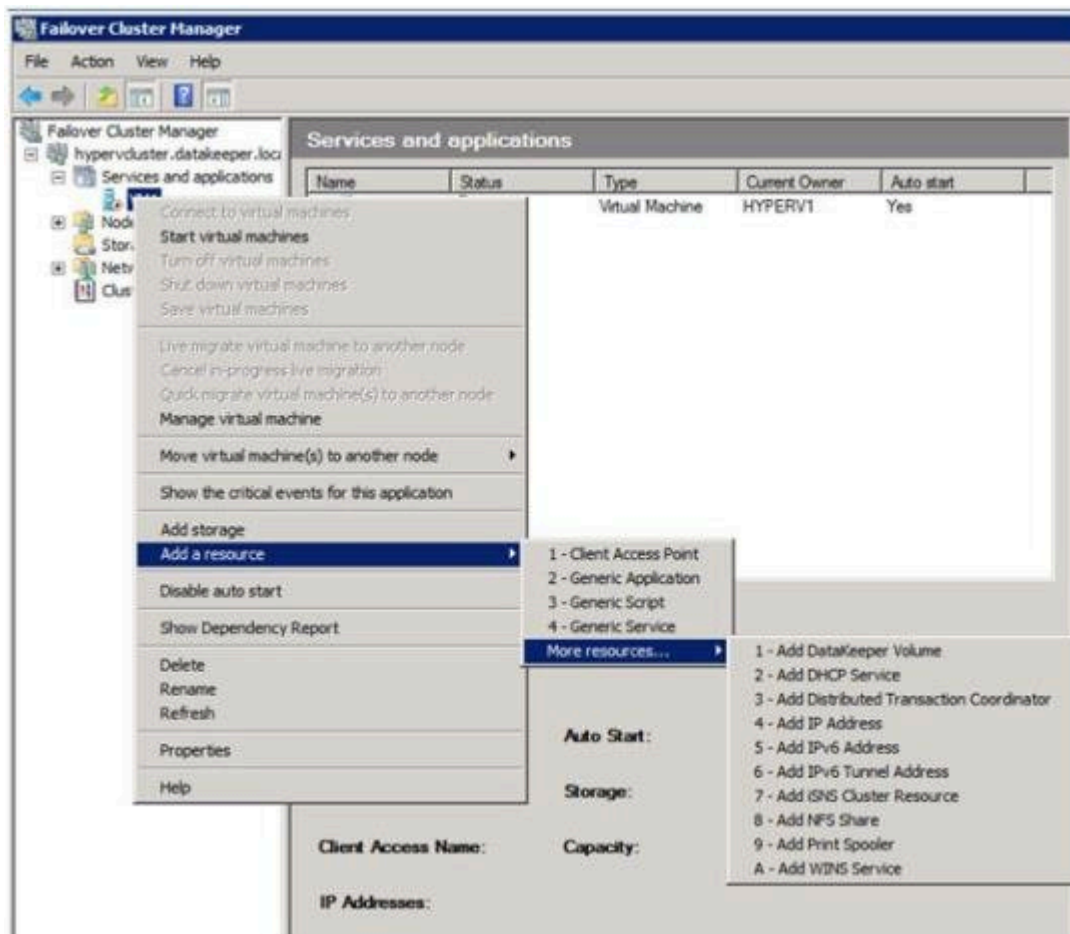


4. Select **Finish** and then shut down the virtual machine.
5. Use **WSFC** to create a Virtual Server resource choosing to protect the first virtual machine. This option is available under **Configure a Service or Application**.





6. Add a **SIOS DataKeeper volume resource**. Right-click on the virtual server resource created in Step #5 and choose **Add a Resource**, then **More Resources -> Add DataKeeper Volume** from the context menu.




7. Right-click on the **SIOS DataKeeper volume resource** and choose **Properties – DataKeeper Volume Parameters**. Denote which drive letter it is associated with the DK resource (e.g. Volume E)

New DataKeeper Volume Properties [X]

General | Dependencies | Policies | Advanced Policies

Shadow Copies | DataKeeper Volume Parameters

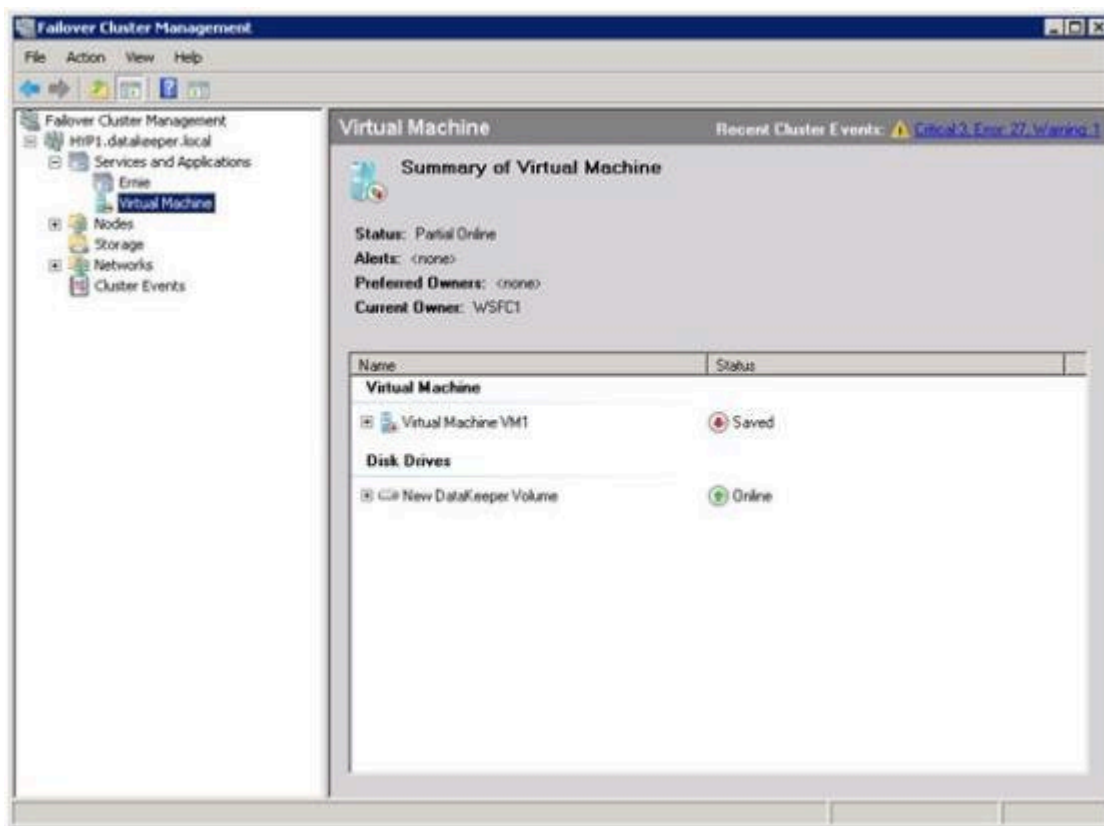
 DataKeeper volume not yet assigned. Please assign a DataKeeper Volume for this resource before proceeding.

DataKeeper Volume Parameters

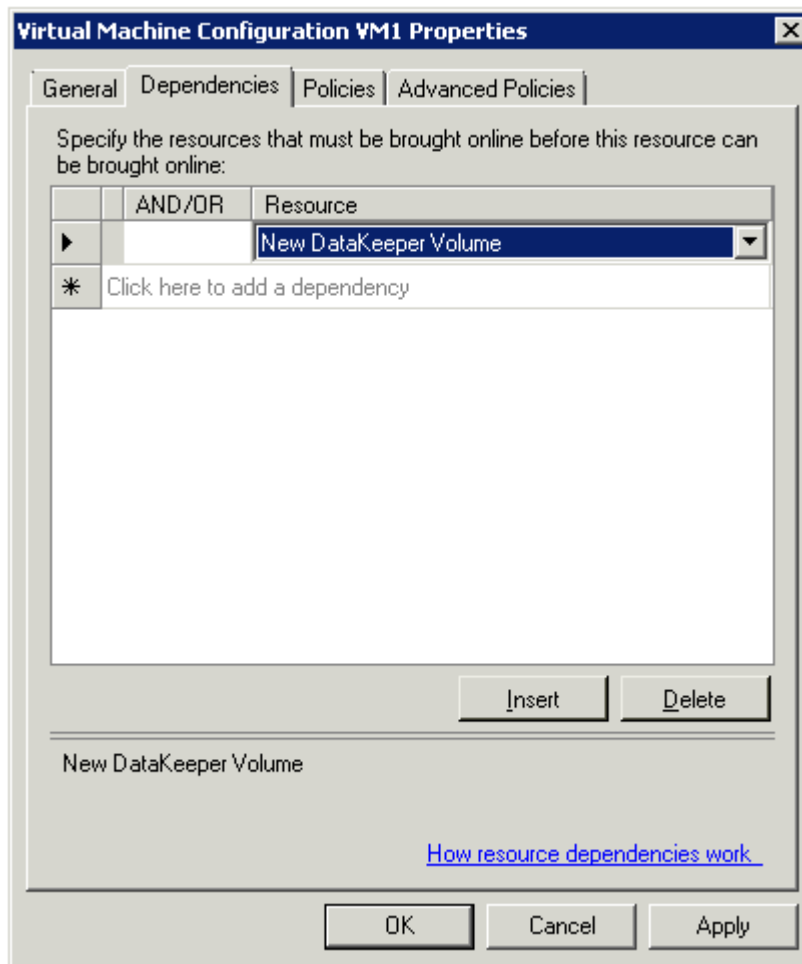
Volume:

Total Size:

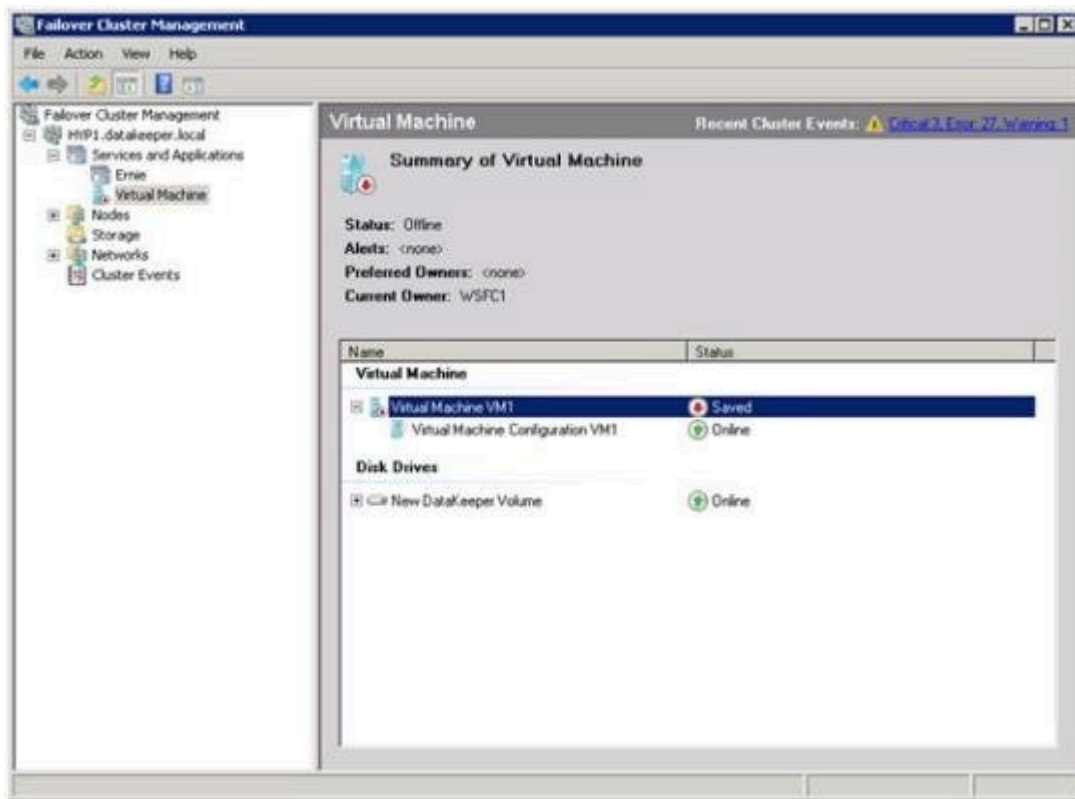
Source System:



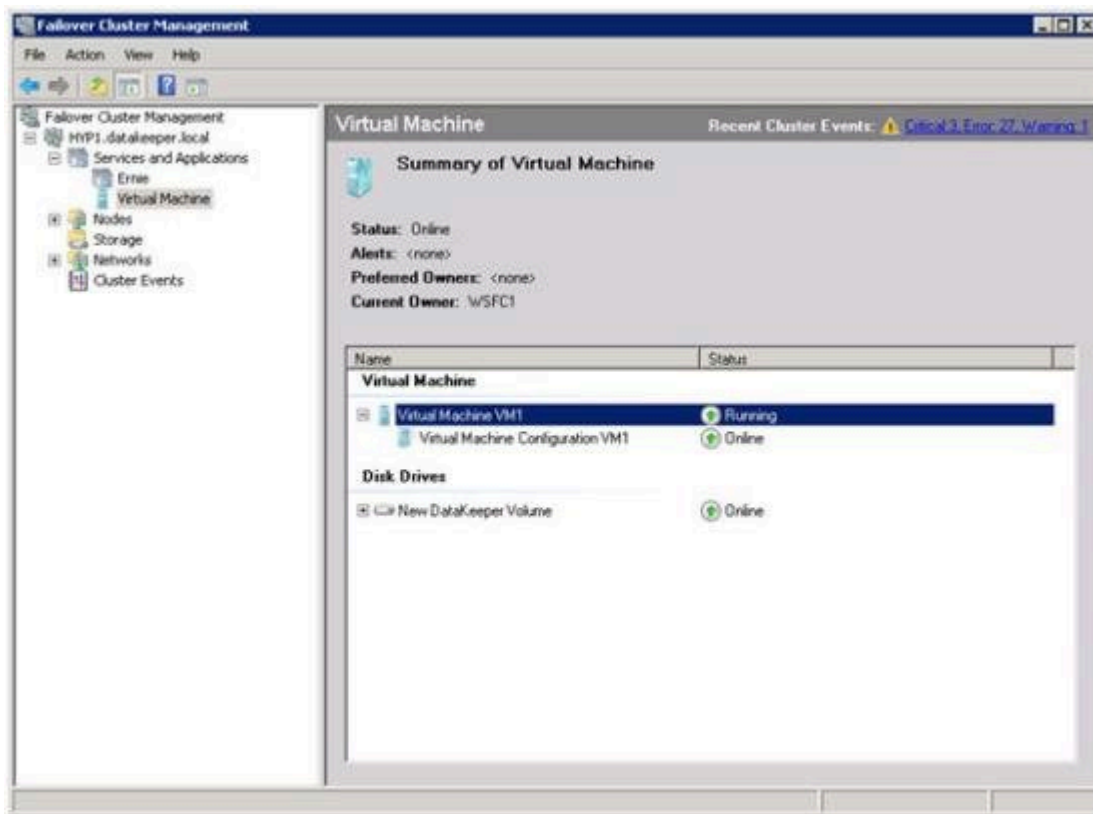
8. Use **Microsoft WSFC Manager**, right-click on the **Virtual Machine Configuration VM1** and choose **Properties**. In the **Properties** window, choose the **Dependencies** tab and add the **New DataKeeper Volume** as a dependency. Click **OK**.



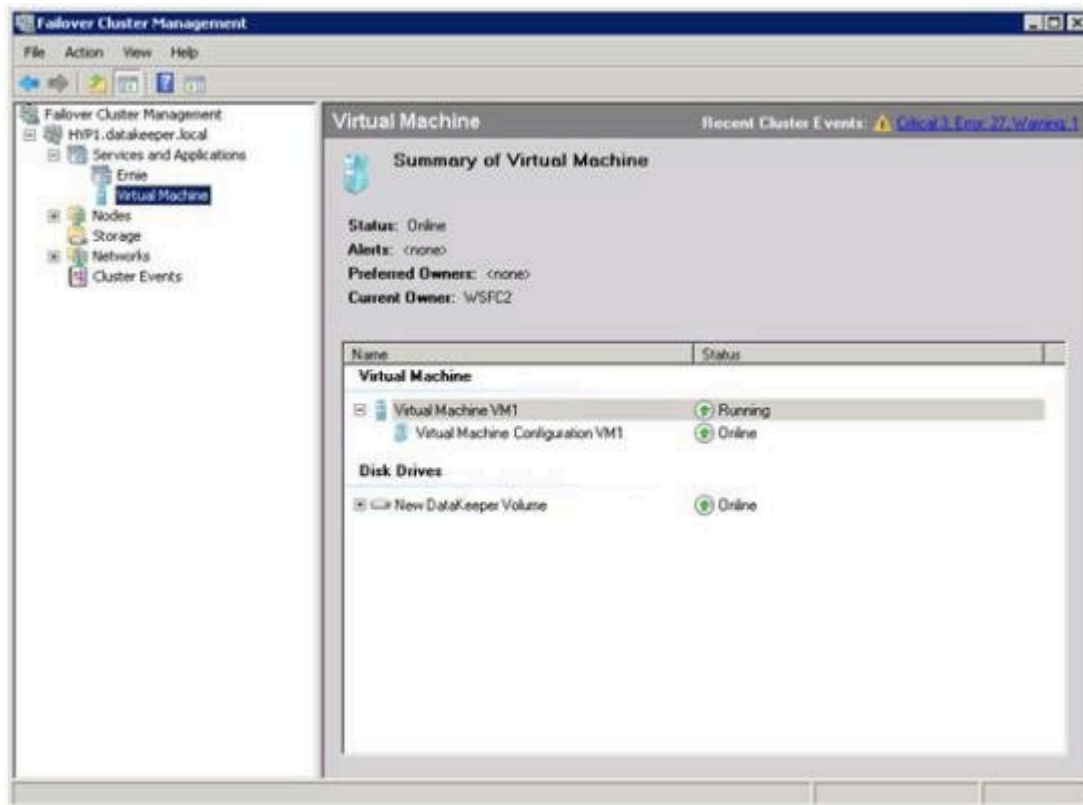
9. Right-click on the **Virtual Machine VM1** resource and choose **Start**.



10. The virtual machine is now online and highly available.



11. Verify that **Quick Migration** works by right-clicking on the virtual machine resource and choose **Move Virtual Machine(s) to Another Node**.
12. Verify that the virtual machine is now running on the secondary server.

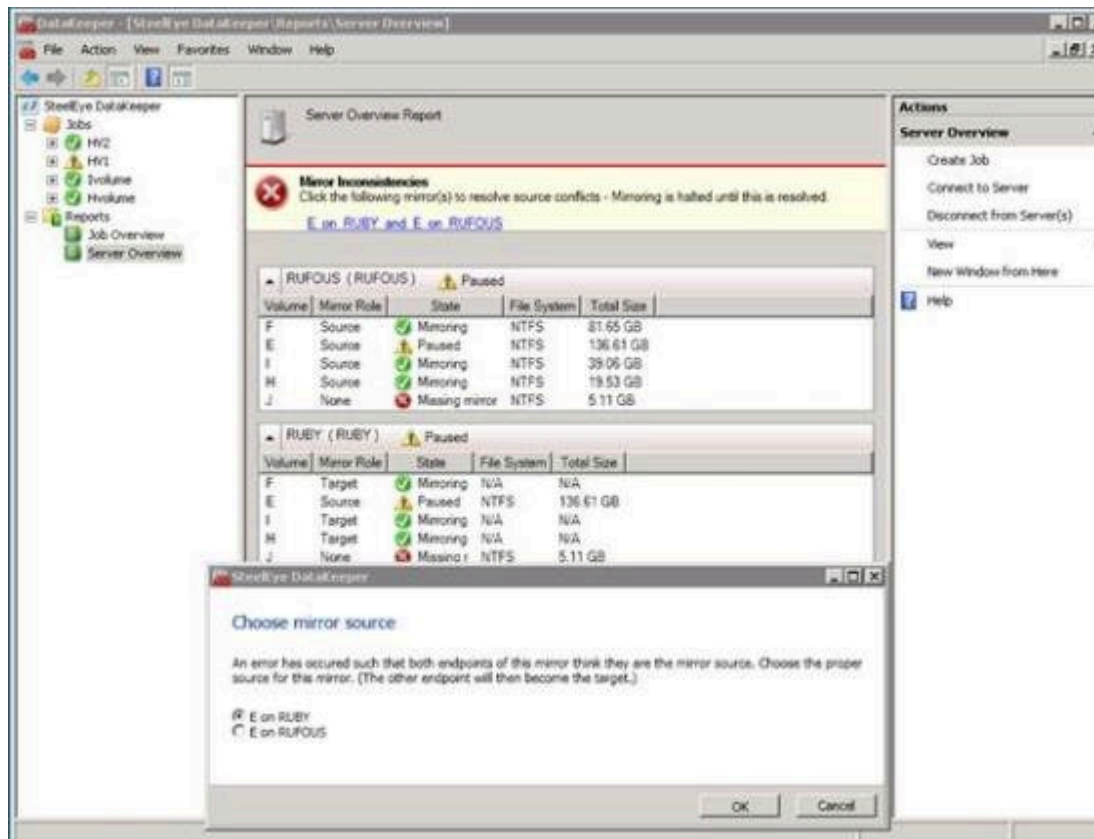


13. Simulate a catastrophic failure by pulling the power cord on secondary server and verify that the virtual machine automatically restarts on primary server.

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
```

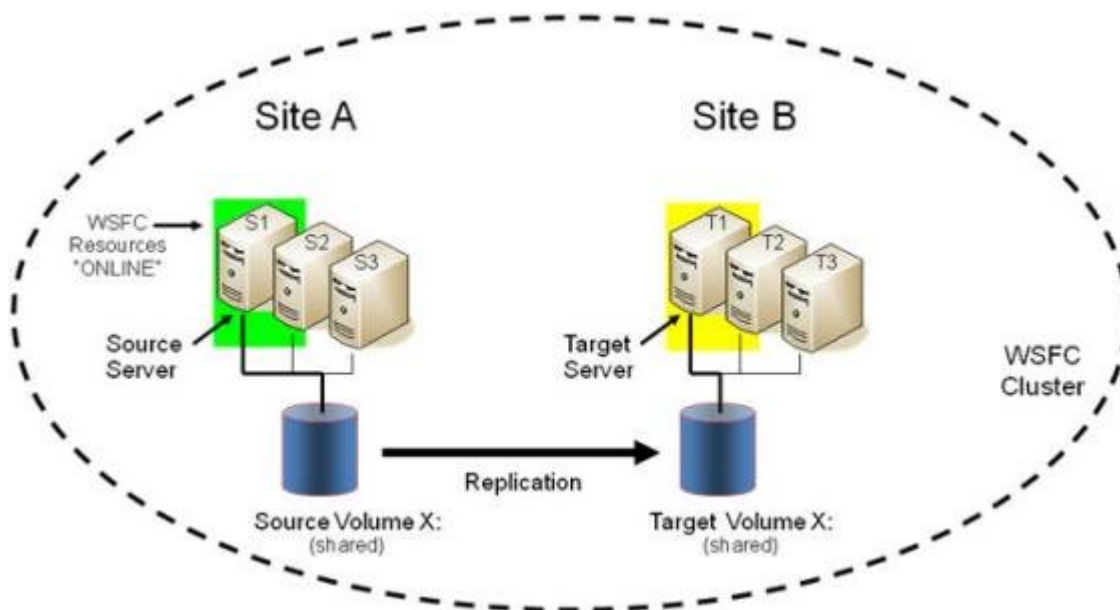
Mirroring of the F: drive from SYSA to SYSB will be established, a partial resync will occur (overwriting any changes that had been made on SYSB), and the mirror will reach the **Mirroring** state.

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.

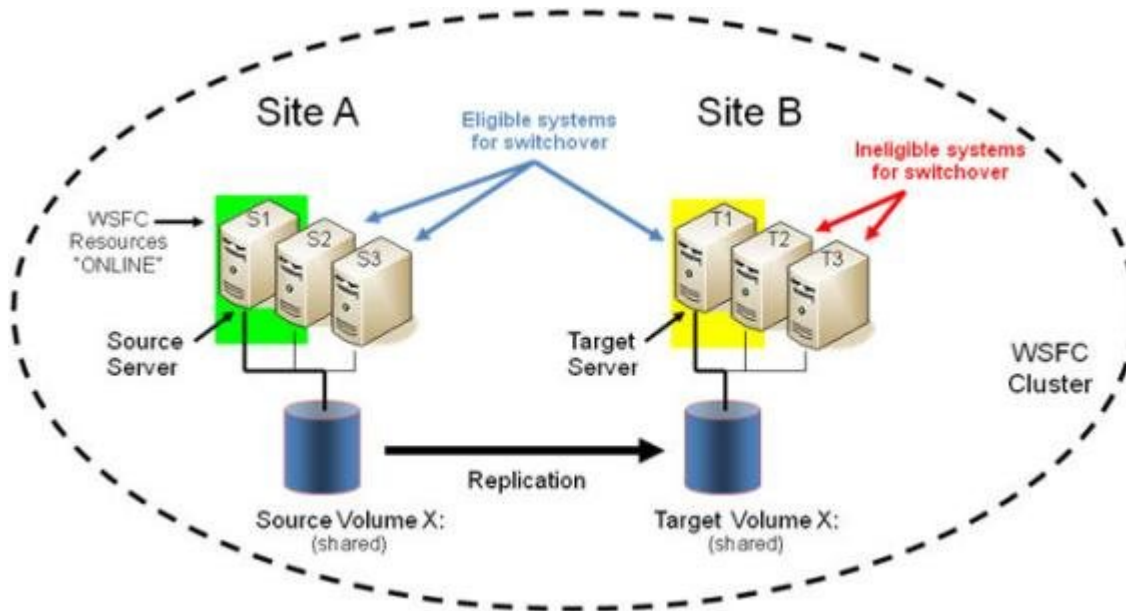
- T1 - Currently the TARGET of the mirror (locked as the primary target)
- T2 - Shared Target (locked)
- T3 - Shared Target (locked)

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.

Installing and Using DataKeeper Cluster Edition on Windows Server 2008 R2 / 2012 Core Platforms

1. Prepare the servers.

Install Windows Server 2008 R2/2012 Core on two servers, configure IP, join a domain, configure firewall for remote administration, configure the server for remote access and install the Failover Cluster feature.

[Server Core Installation Option Getting Started Guide](#)

[Configuring a Server Core Installation of Windows Server 2008 R2 with Sconfig.cmd](#)

[Install and Deploy Windows Server 2012](#)

2. Install SIOS DataKeeper.

Run the SIOS DataKeeper (DK) setup.exe from the command line on both servers. Make sure you choose to install the DataKeeper core services only, not the DataKeeper GUI.

3. An email should have been received containing the license file(s). It is recommended that this file be renamed to the YYYYMMDD.lic format to distinguish the day the license was activated. Once received, copy the license file(s) to the appropriate directory.

- On each system, copy the file(s) to:

`%windir%\sysWOW64\LKLicense` (ex. `c:\windows\SysWOW64\LKLicense`)

Note: If the LKLicense directory does not already exist, it will need to be created prior to copying the files.

- ## 4. On a management server (or Vista workstation with "Remote Server Admin Tools" installed), install the DataKeeper GUI console only.
- ## 5. Using the management server's DataKeeper GUI, connect to the new

core only servers and create a "job" with a mirror of the data volumes.

6. Using the management server, create the WSFC cluster.
7. Change the quorum mode to **Node Majority** if deploying an odd number of nodes or **Node and File Share Majority** if deploying an even number of nodes. Refer to the appropriate Microsoft documentation for official guidance on quorum configurations in a multi-site environment.

<http://download.microsoft.com/download/...Clustering.doc>

Configuring the Quorum in a Failover Cluster

8. Add the DataKeeper ClusterCluster Edition resource to the cluster.

Note: The example below assumes you want to add the E drive mirror to your Hyper-V resource.

Start Powershell on the primary core only server and enter the following commands:

- Import-Module FailoverClusters
- Add-ClusterResource -Name "DataKeeper Volume E" -ResourceType "DataKeeper Volume" -Group "<name of Hyper-V resource>"
- Get-ClusterResource "DataKeeper Volume E" | Set-ClusterParameter VolumeLetter E

9. Add your resource to Windows Server Failover Clustering to make it highly available (HA)

To make a Hyper-V VM HA - Hyper-V: Using Hyper-V and Failover Clustering

10. Finally, if not already done in the previous step, add the DataKeeper replicated storage to your resource to complete the dependency creation and setup.

Non-mirrored Volume Resource

A non-mirrored volume resource is a DataKeeper resource **where no data is replicated to any node in the cluster**. This resource type should only be used where the data is temporary and/or non-critical such as MS SQL Server tempdb space. In this case, when MS SQL restarts on another node after a failover or switchover, the tempdb space is automatically recreated so replication of the data is not necessary.

This non-mirrored volume resource will be able to come Online and go Offline on all cluster node without ever affecting the configured volume. Additionally, the volume will remain unlocked and writable at all times on all nodes in the cluster.

To configure a non-mirrored volume resource:

- Configure a volume on all cluster nodes using the same drive letter (all nodes must use the same drive letter).
- Create any directories that are required for the volume on all cluster nodes.
- Create a DataKeeper Volume resource using the Failover Clustering UI. Provide a name that best describes its intended use - (Example: "DataKeeper Volume F (Non-Mirrored)"). The following steps will set the Properties needed for the non-mirrored resource:
 - In the Failover Cluster Manager, right-click the **Cluster Group** or the **Role** that will contain the non-mirrored DataKeeper Volume Resource. If a role has **NOT** been defined, an [Empty Role](#) must be created.
 - Select **Add a Resource, More Resources**, then select **DataKeeper Volume**.
 - Right-click on the new **DataKeeper volume resource** and select **Properties**.
 - Enter the **Resource Name** you selected earlier (Example: "DataKeeper Volume F (Non-Mirrored)") then click **OK**. You do not need to change any other properties at this time.

The following steps will set the Properties needed for the non-mirrored

resource:

- 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)
```

- Add the 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
```

If this non-mirrored volume resource is to be used with MS SQL Server for tempdb space the following configurations steps are needed:

- Ensure that the volume security settings for the user account that is chosen to run SQL Server services has full access to the volume on all nodes in the cluster.
- Ensure that the "SQL Server" resource in the Failover Clustering group has a dependency on the new DataKeeper Volume resource.

SQL Server Properties [X]

General Dependencies Policies Advanced Policies Properties

Specify the resources that must be brought online before this resource can be brought online:

	AND/OR	Resource
▶		DataKeeper Volume F (Non-Mirrored)
	AND	Name: SQLEDMNET
	AND	DataKeeper Volume E
*	Click here to add a dependency	

[Insert] [Delete]

DataKeeper Volume F (Non-Mirrored) AND Name: SQLEDMNET AND DataKeeper Volume E

[How resource dependencies work](#)

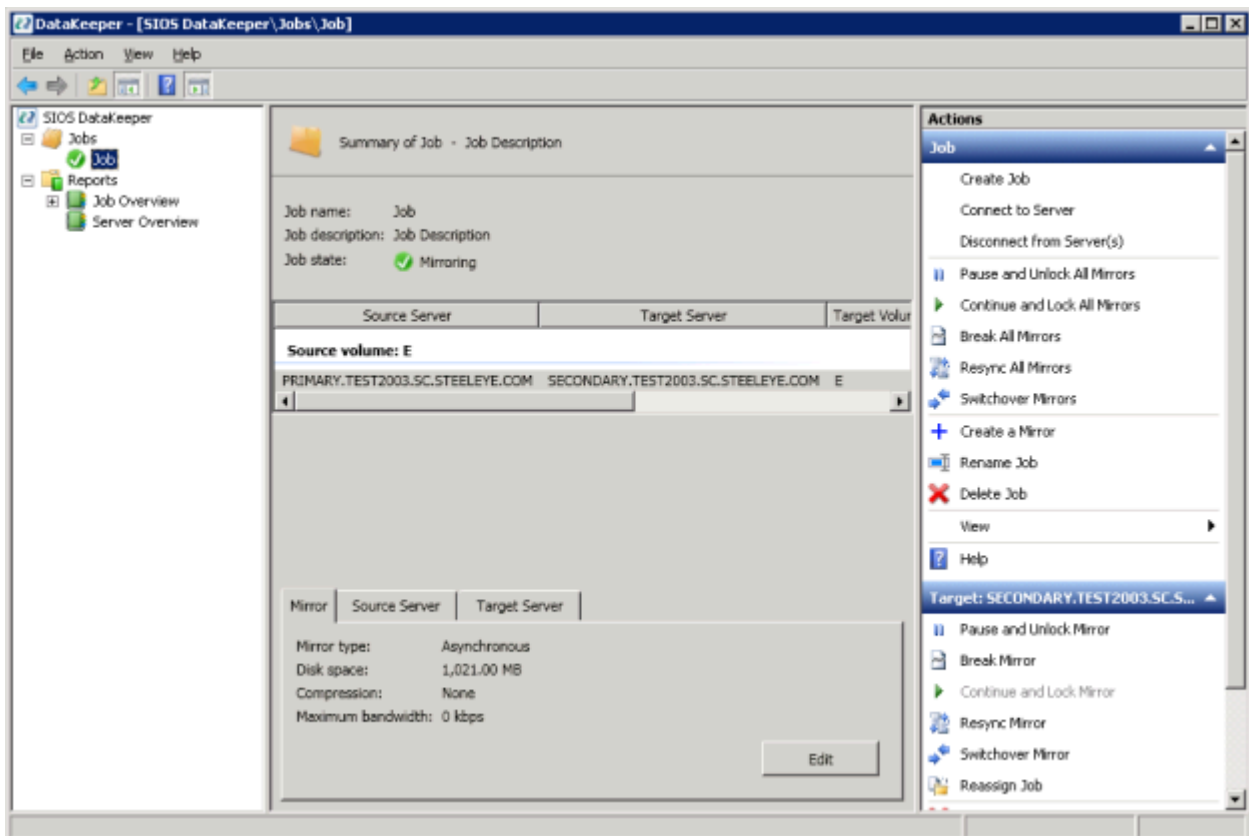
[OK] [Cancel] [Apply]

Using DKCE to Enable Multi-Site File Share Resources with Windows Server 2008R2 WSFC

Use the following procedure to protect File Share resources using SIOS DataKeeper and Microsoft Failover Cluster (WSFC).

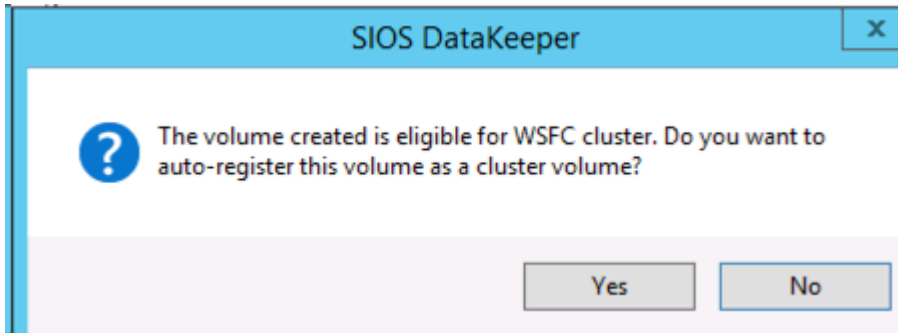
***IMPORTANT NOTE:** DataKeeper Cluster Edition registration with Failover Cluster is automatic and occurs 60 seconds after detecting a Failover Cluster configuration on each node.

1. [Create a mirror](#) for the volume using the SIOS DataKeeper user interface. Make sure the mirror is in "**Mirroring**" state before protecting in WSFC.



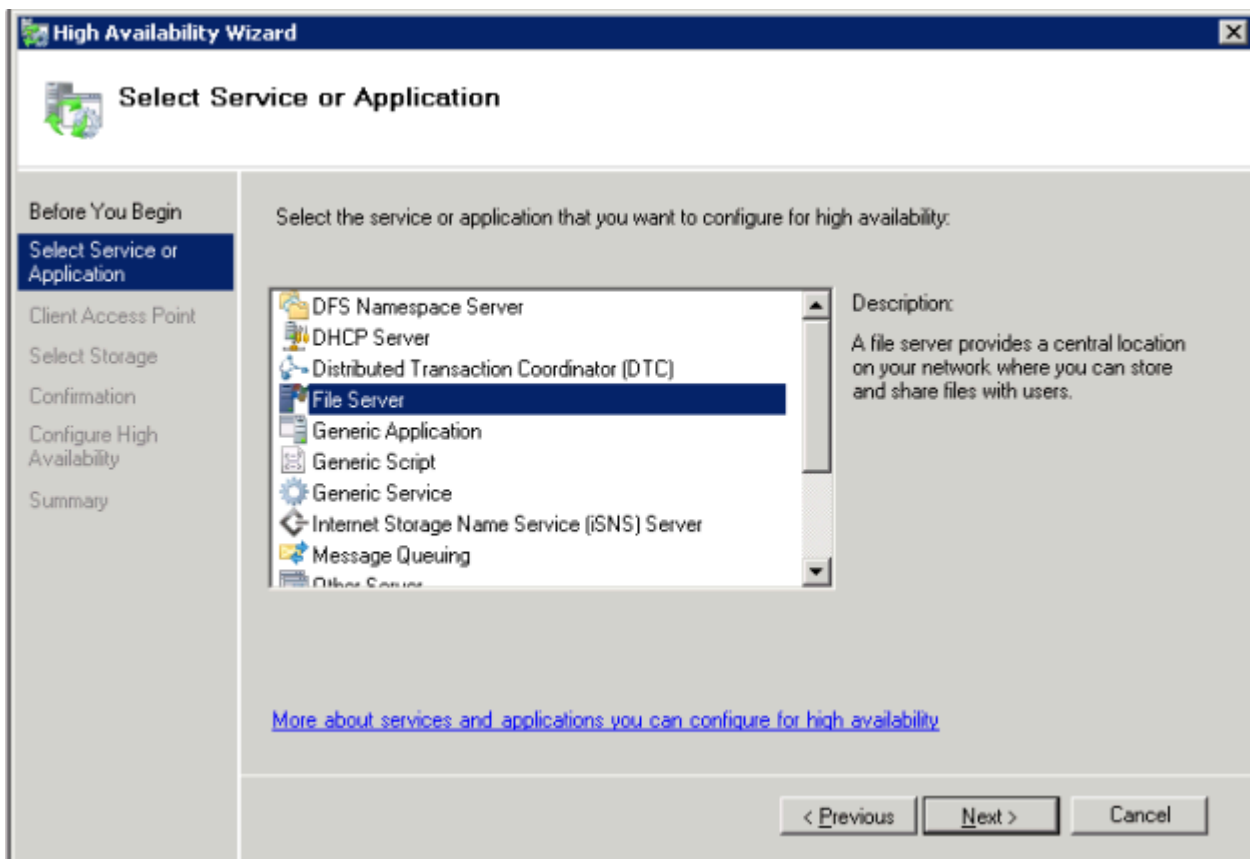
2. Use the DataKeeper Auto-register Pop-up to quickly register the replicated volume with WSFC. Select **YES** to auto-registration. The DataKeeper replicated volume will automatically be added to cluster

Available Storage.



Open the **WSFC Failover Cluster Manager**

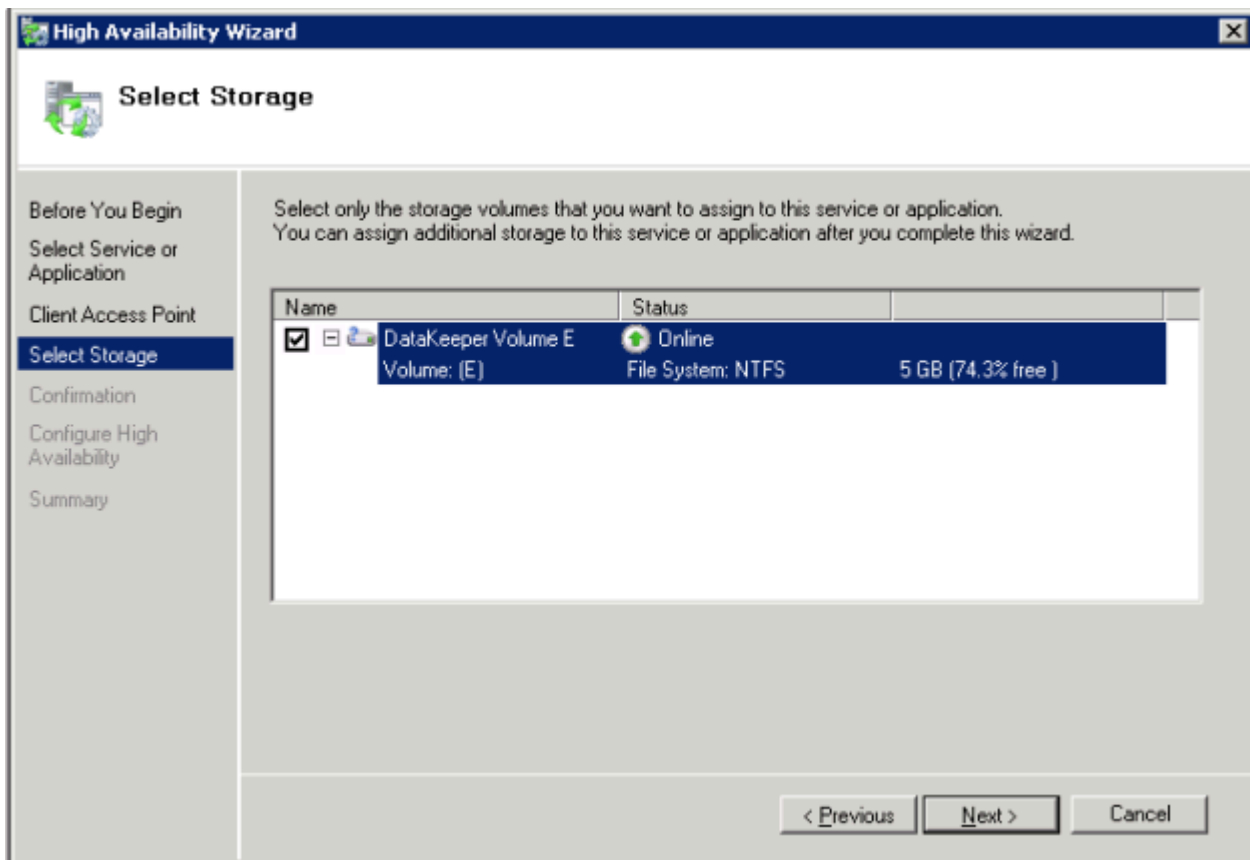
3. Right click on "Services or Application" and create a File Server resource. The File Server Role must already be installed on all the nodes in your cluster. If not, an error will be displayed.



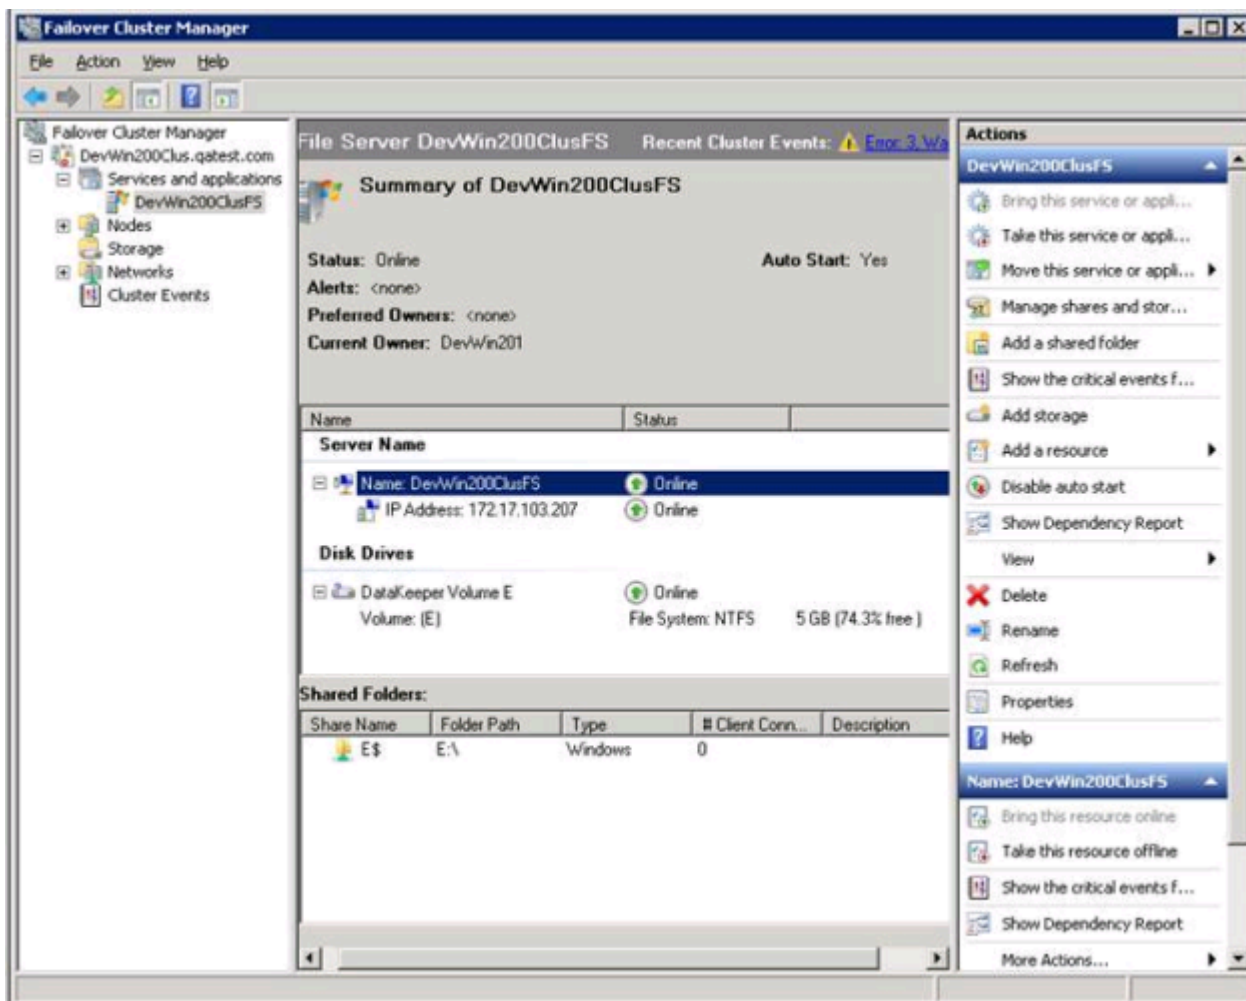
a. Click **Next**.

b. On the next screen fill in the **File Server resource name** and provide an **IP address**. Click **Next**.

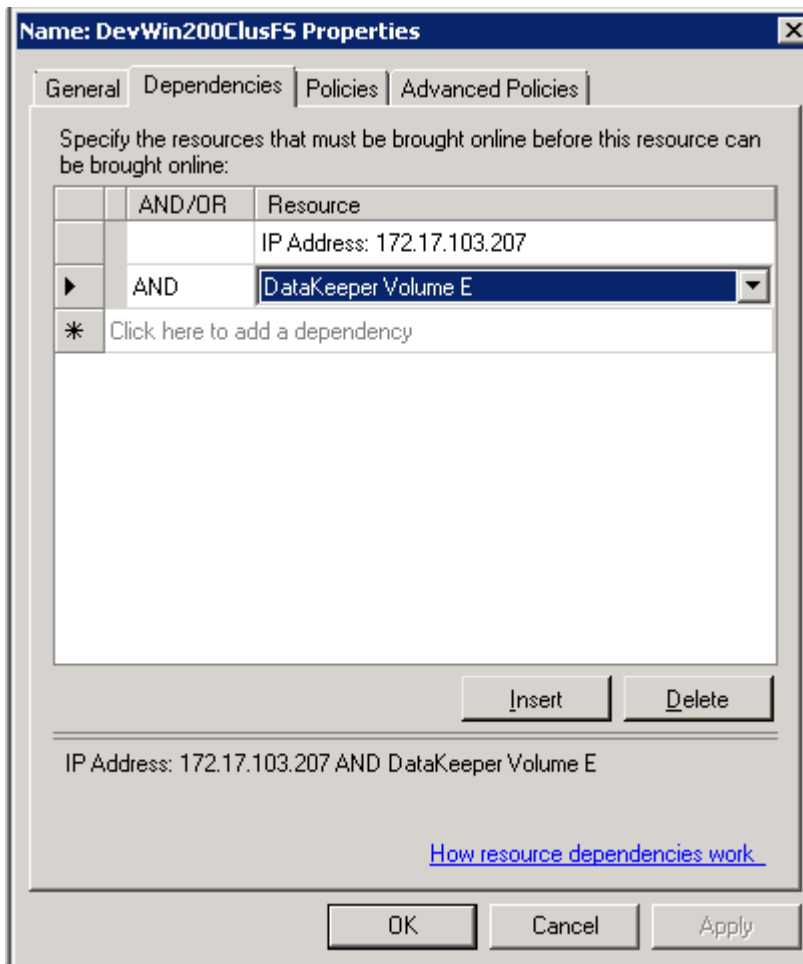
4. Add DataKeeper Storage to the File Server resource.



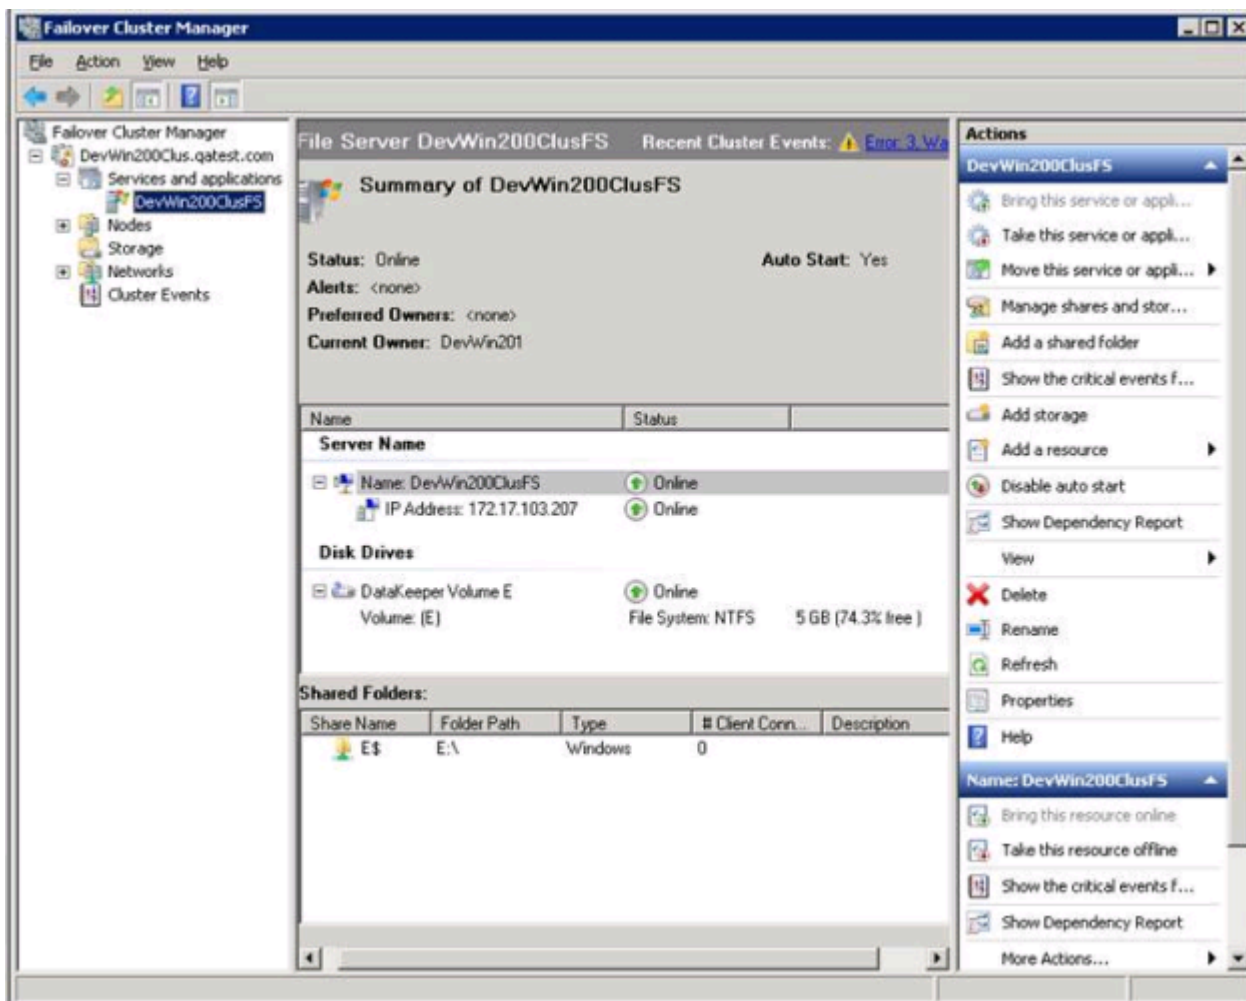
- Select the **DataKeeper Volume** check box above and click **Next**.
- The **File Server** Confirmation screen will be displayed, confirm it and click **Next**.
- The **File Server** Configuration Summary will be displayed, Click **Finish**.
- The **File Server** resource will be created with the DataKeeper volume root path as shown below.



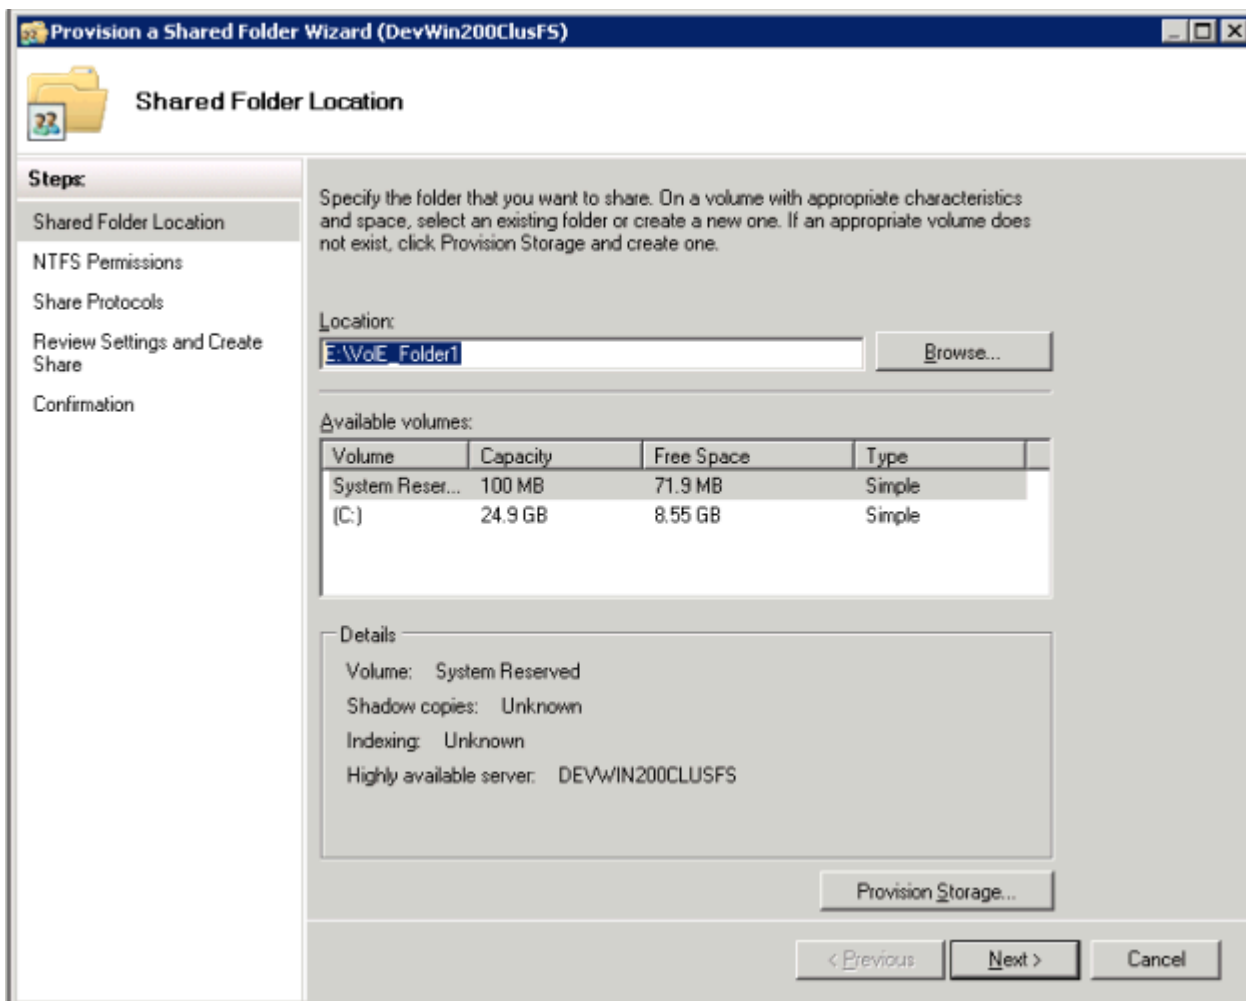
5. Using the **Failover Cluster Manager**, create a resource dependency from the File Server resource to the DataKeeper Volume resource. This dependency relationship is very important to coordinate in-service and out-of-service operations of the file share and volume resources when moving the File Share resource from one node to another node.
 - a. Right-click on the **File Server** resource name (center panel, selected above), then select **Properties**.
 - b. Add the **DataKeeper Volume** resource as a dependency for the **File Server** resource, and click **OK**.



6. The File Server resource is now ready for file shares to be added. File Shares can be added several ways.
 - a. Right-click on the **File Server resource** (highlighted below) and select **"Add a shared folder"**. Or, use the "Add a shared folder" action on the right panel shown below.



b. The **Shared Folder Location** screen shown below includes a browse button to locate the folder to share on the DataKeeper Volume. Then click **Next**.



c. The **NTFS Permissions** screen provides alternative **NTFS permission level**, and click **Next**.

d. The **Share Protocols** screen provides **protocol** selection, **share name**, confirm **Share Path**, and click **Next**.

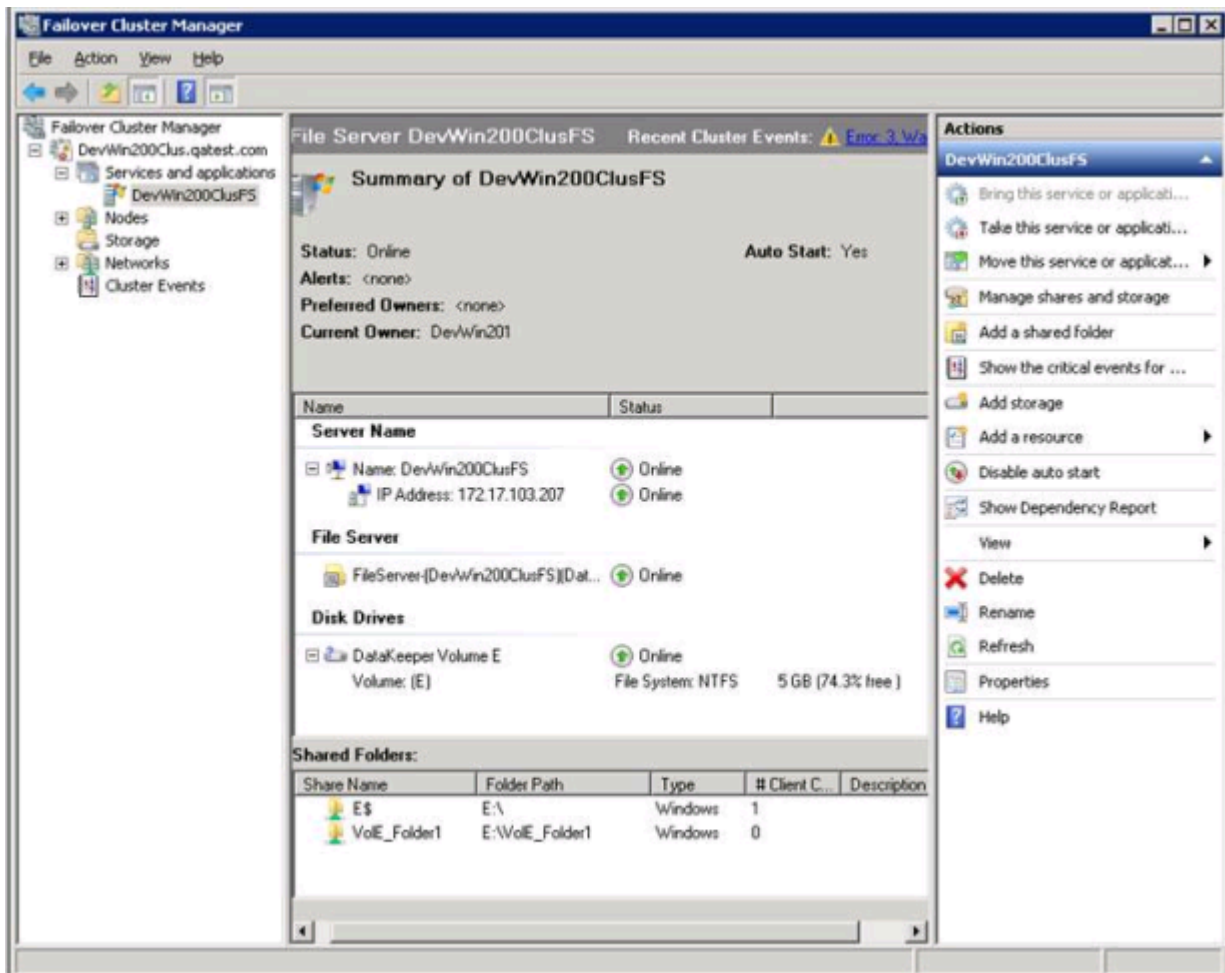
e. The **SMB Settings** screen provides a **Description** for the Share, and click **Next**.

f. The **SMB Permissions** screen provides alternative SMB permissions, click **Next**.

g. The **DFS Namespace** screen provides a publish to DFS namespace option, click **Next**.

h. The **Review** screen provides a configuration summary, click **Next**.

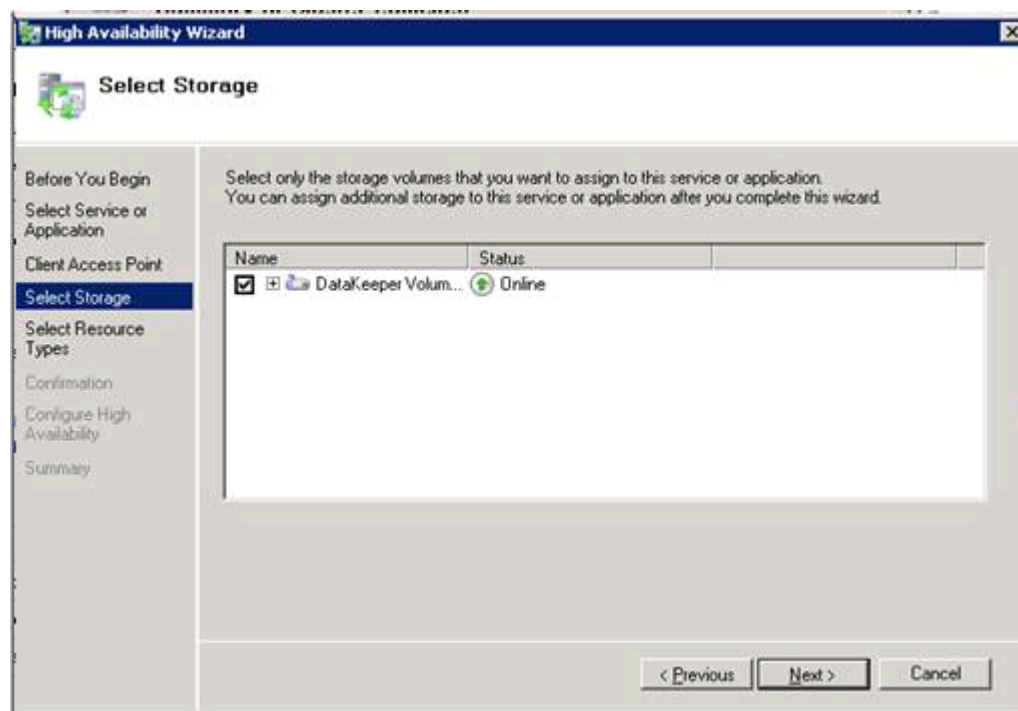
- i. The **Confirmation** screen displays success or failure to create the file share, click **Close**.
7. Add a file share using a DataKeeper replicated volume resource is now completed. Shares are displayed in the lower Shared Folders panel shown below.



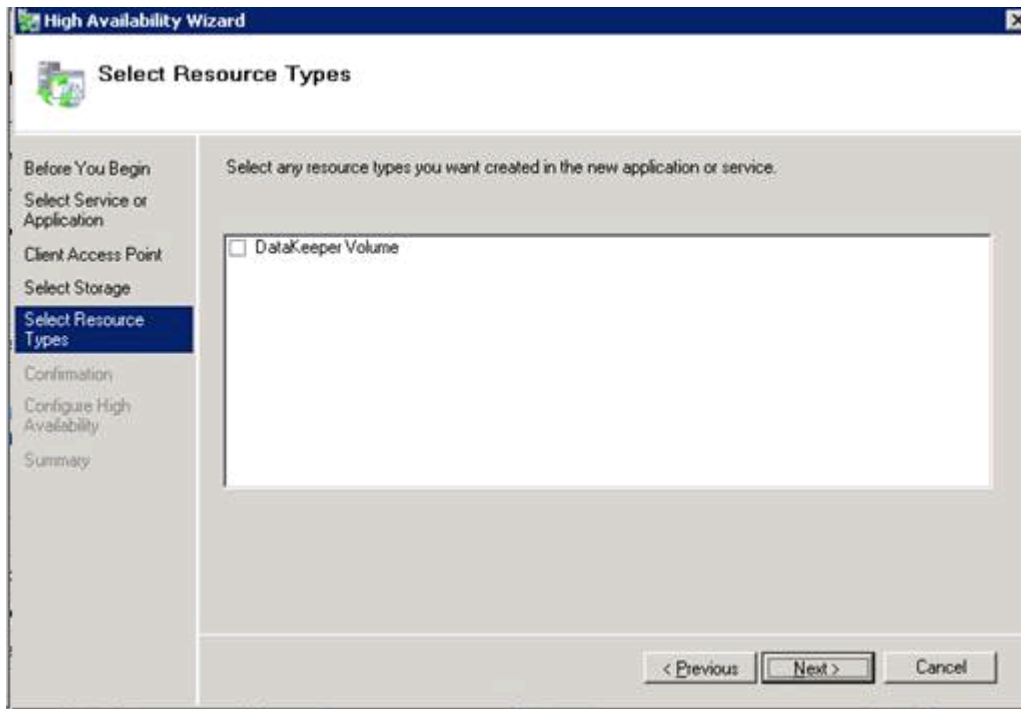
Creating Other Server Resource in WSFC

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.



FAQs

Refer to this section for answers to the most frequently asked questions about SIOS DataKeeper and SIOS DataKeeper Cluster Edition.

[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 FAQ](#)

[Error Messages Log](#)

[Inability to Create a Mirror](#)

[Network Disconnect](#)

[Reclaim Full Capacity of Target Drive](#)

[Resize or Grow Mirrored Volumes](#)

[Server 2012: Server Manager "File and Storage Services" Disk Status](#)

[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.

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.

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](#)

[WSFC 2016 File Shares Cannot be Created for File Server Role Using Server Manager or Failover Cluster Manager](#)

[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 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.

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

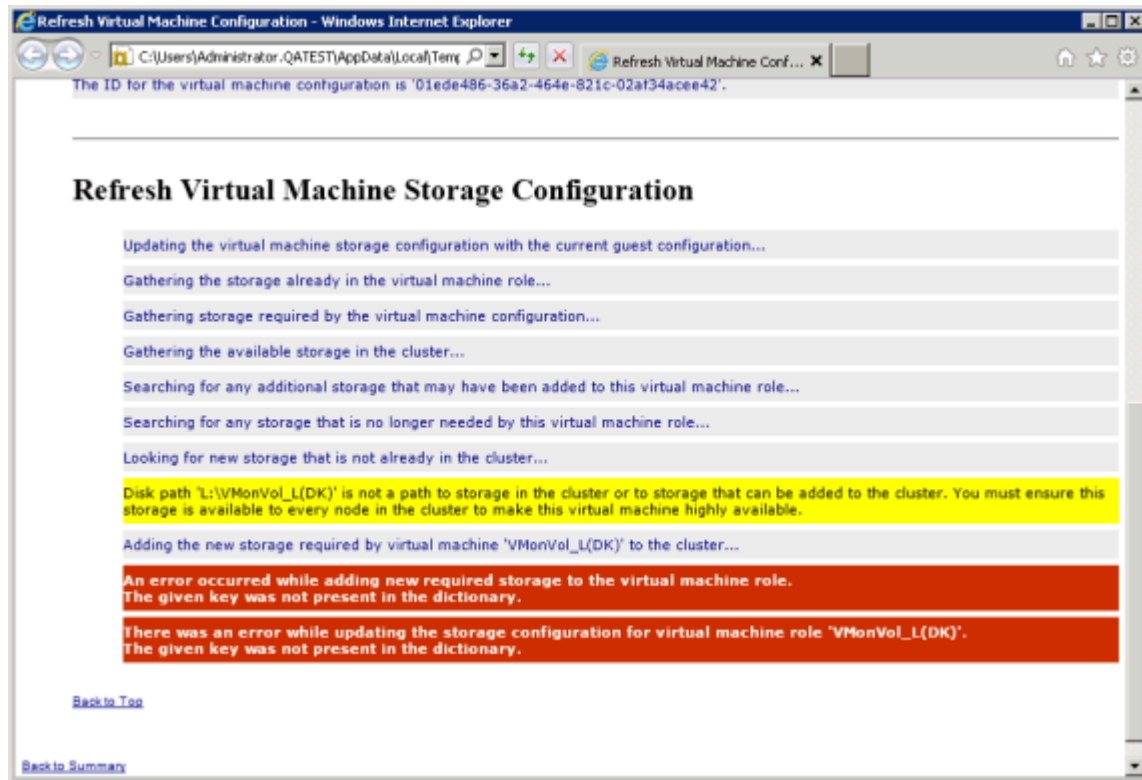
If Failover Cluster Manager is used to modify the VM configuration while the VM is clustered such as adding a Network Interface to the VM, **"Refresh Virtual Machine Storage Configuration"** errors may be generated and the VM will fail Quick Migration and/or Live Migration to another cluster node.

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:

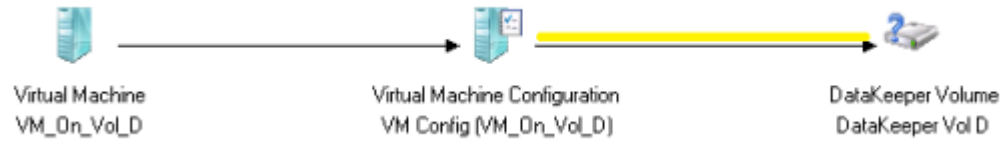


Suggested Action

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.

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 or download and install the Microsoft Hotfix described in the following KB article:
<http://support.microsoft.com/kb/978476>. 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 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 – 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.

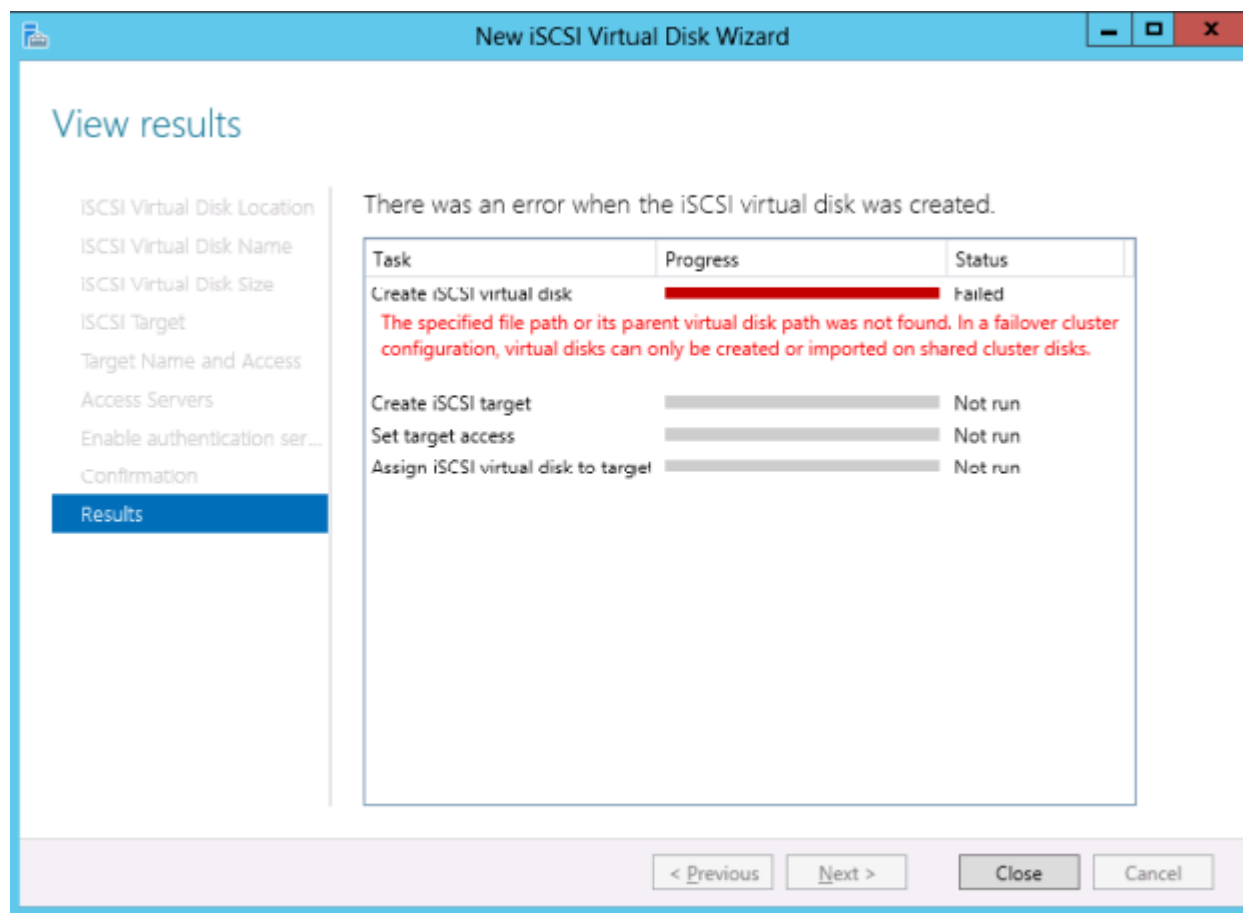
Suggested Action

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 iSCSI Target Role Does Not Support Dynamic Disks

Description

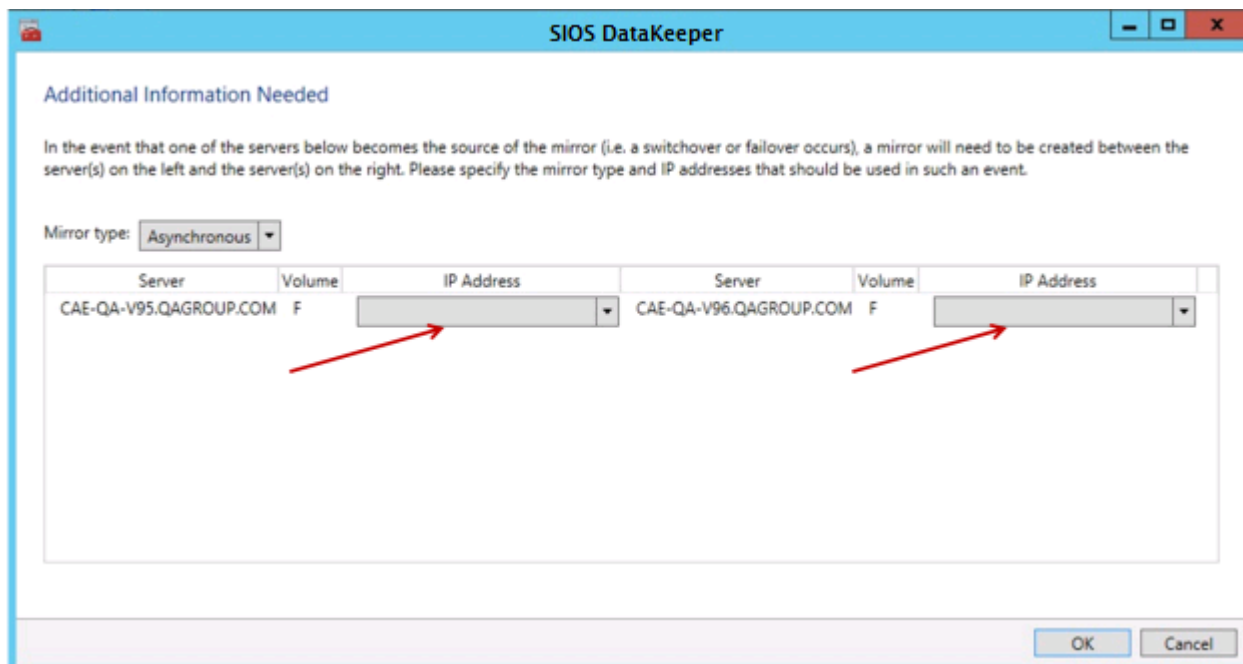
The iSCSI Target Role only supports DataKeeper Volumes that are mirrors of **Simple Volumes** placed on **Basic Disks**. If any of your mirrors are using volumes such as Striped or Spanned volumes on a Dynamic Disk on either the source or target system, then you cannot create an iSCSI Target role that uses those DataKeeper Volume resources for storage.



Windows Server 2012 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.

New Mirror

Shared Volumes

Choose a Source
Shared Volumes
Choose a Target
Configure Details

Source server: CAE-QA-V94.QAGROUP.COM
Source IP and mask: 10.200.8.94
Source volume: H

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.

Include	Server	Volume	Source IP / Mask
<input checked="" type="checkbox"/>	CAE-QA-V95.QAGROUP.COM	H	No Valid IP Selection Found

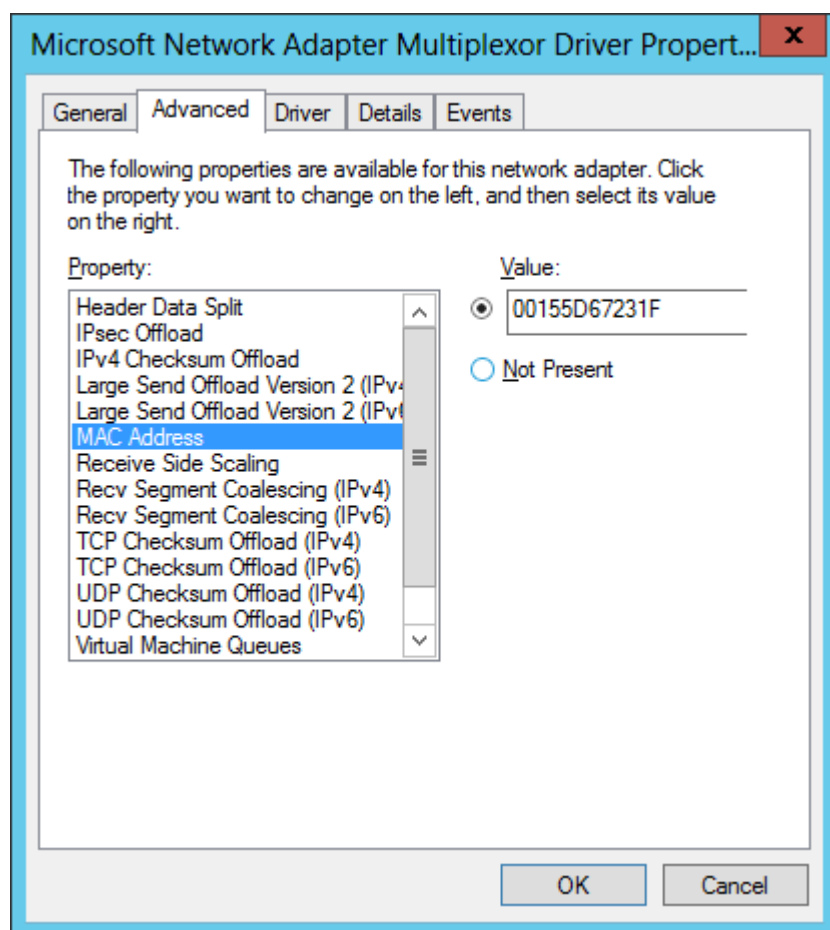
[Connect to Server](#)

Previous Next Cancel

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:



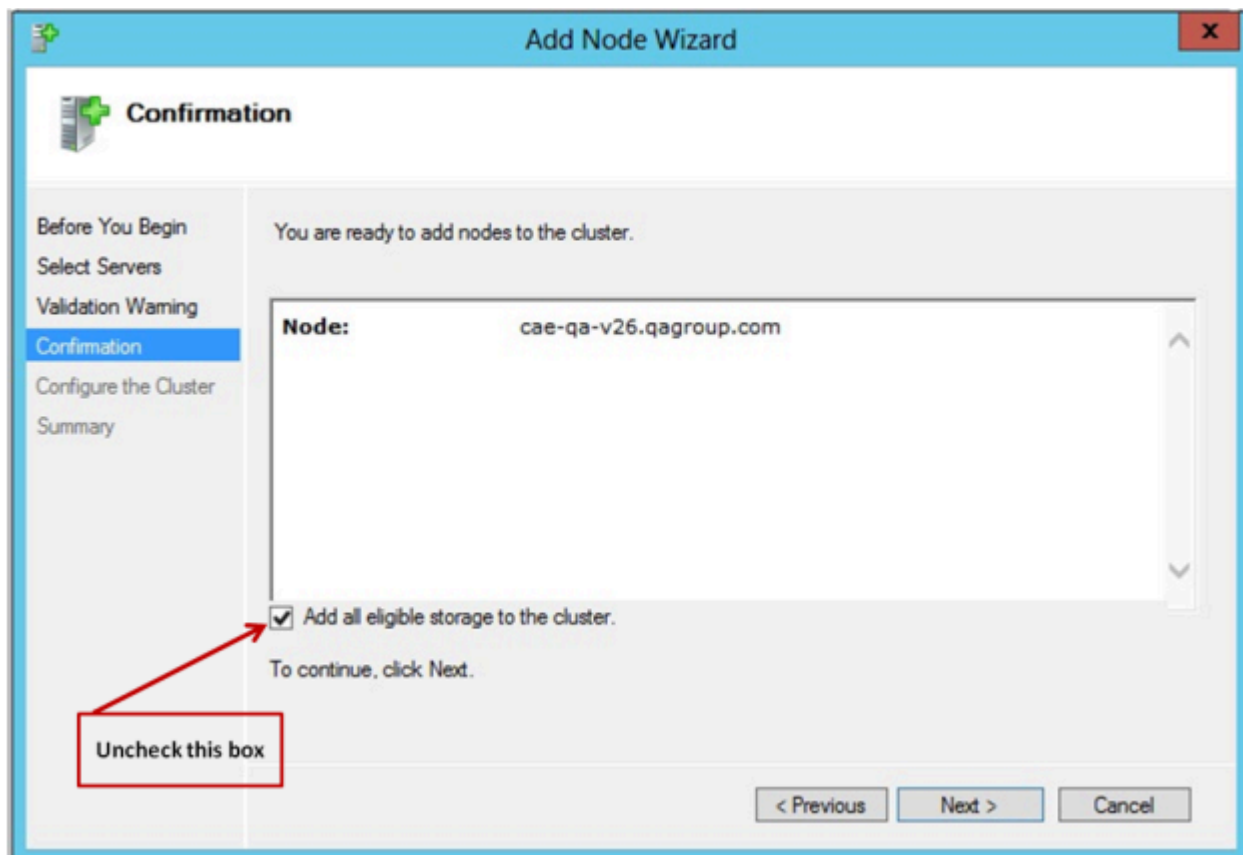
WSFC 2012 Cluster Creation Default Setting Issue

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**.

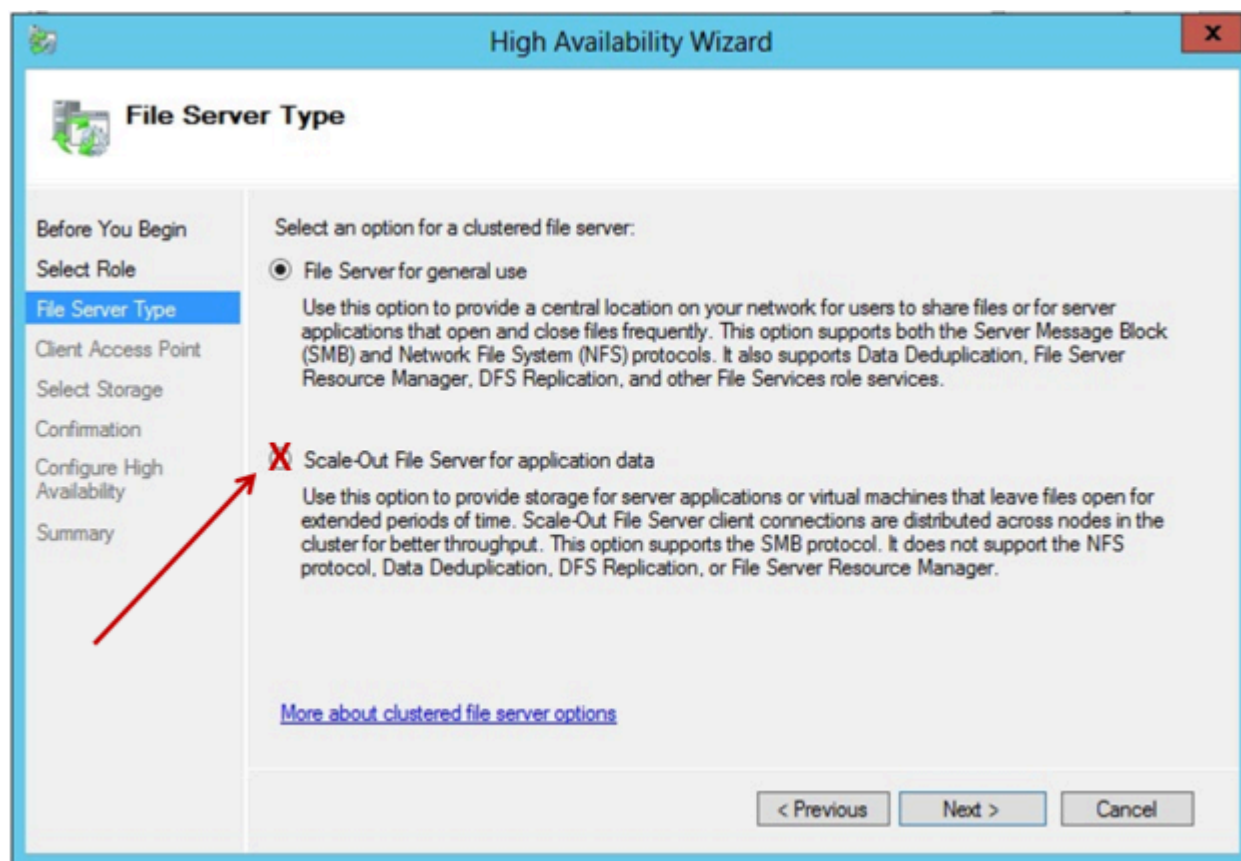
WSFC 2012 New File Server Type Not Supported

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.

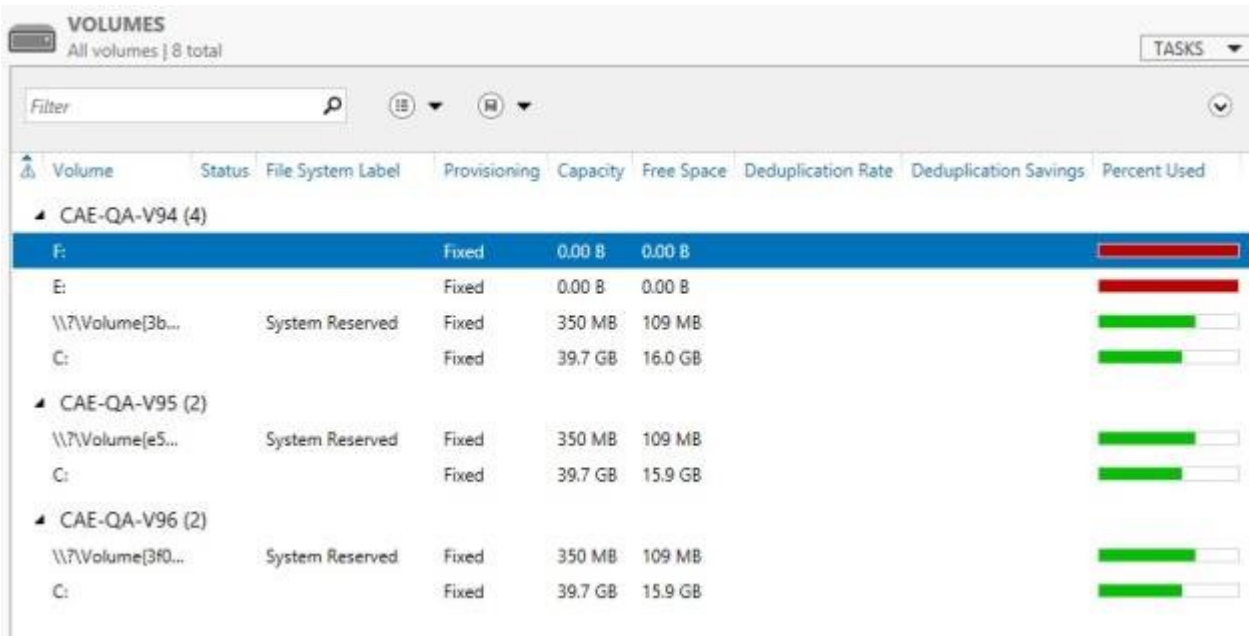
Note: Windows Server 2012 ReFS (Resilient File System) is also not currently supported.

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			
C:			Fixed	39.7 GB	16.0 GB			
CAE-QA-V95 (2)								
\\?\Volume{e5...	System Reserved		Fixed	350 MB	109 MB			
C:			Fixed	39.7 GB	15.9 GB			
CAE-QA-V96 (4)								
\\?\Volume{3f0...	System Reserved		Fixed	350 MB	109 MB			
C:			Fixed	39.7 GB	15.9 GB			
E:	Volume1		Fixed	4.97 GB	3.95 GB			
F:	Volume2		Fixed	4.97 GB	3.95 GB			

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**.

Server Manager

File and Storage Services > Volumes > Disks

DISKS
All disks | 8 total

Number	Virtual Disk	Status	Capacity	Unallocated	Partition	Read Only	Clustered	Subsystem	Bus Type	Name
CAE-QA-V95 (4)										
0		Online	40.0 GB	0.00 B	MBR				SAS	VMware
1		Online	5.00 GB	0.00 B	GPT	?	?		SAS	VMware
2		Online	5.00 GB	0.00 B	GPT	?	?		SAS	VMware
3		Online	10.7 GB	5.74 GB	MBR		✓		iSCSI	ROCKET
CAE-QA-V96 (4)										
0		Online	40.0 GB	0.00 B	MBR				SAS	VMware
1		Online	5.00 GB	0.00 B	GPT	?	?		SAS	VMware
2		Online	5.00 GB	0.00 B	GPT	?	?		SAS	VMware
3		Offline	10.7 GB	5.74 GB	MBR	✓	✓		iSCSI	ROCKET

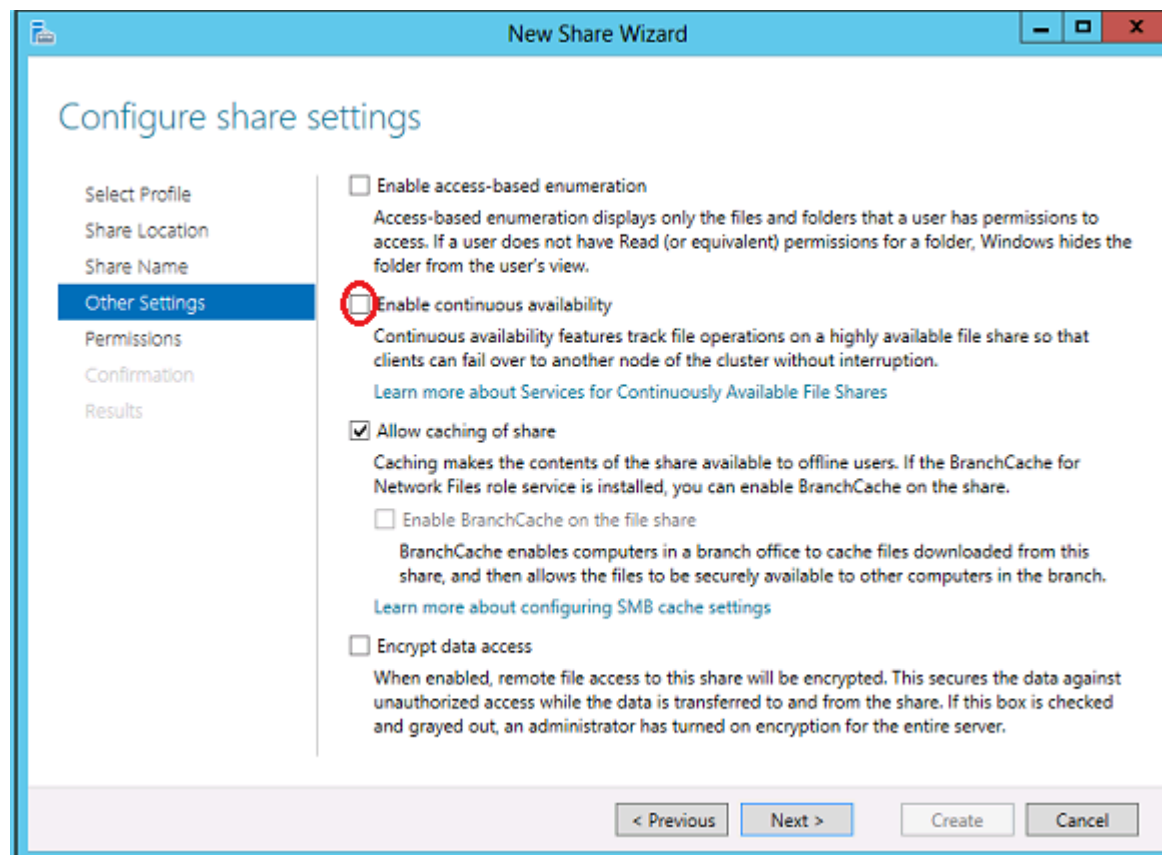
DK Clustered Volumes

iSCSI Shared Disk

Last refreshed on 10/15/2012 10:25:38 AM

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.



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"`

DKCE 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	v8.5.1	v8.6.0	v8.6.1
Microsoft Windows Server 2008 Enterprise and DataCenter Editions											
Microsoft Windows Server 2008 R2 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

Note: The operating system versions listed above are supported for guests running on the following virtual platforms:

- 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 are supported on XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later
- Google Cloud

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	v8.5.1	v8.6.0	v8.6.1
Microsoft Windows Server 2012 R2 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 2016 Standard and DataCenter Editions									64-bit	64-bit	64-bit
Microsoft Windows Server 2019 Standard and DataCenter Editions											

Note: The operating system versions listed above are supported for guests running on the following virtual platforms:

- 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 are supported on XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later
- Google Cloud

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	v8.5.1	v8.6.0	v8.6.1	v8.6.2
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	64-bit	64-bit
Microsoft Windows Server 2016									64-bit	64-bit	64-bit	64-bit
Microsoft Windows Server 2019												
Windows XP												
Windows Vista												

Note: The operating system versions listed above are supported for guests running on the following virtual platforms:

- 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 are supported on XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later
- Google Cloud

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	v8.5.1	v8.6.0	v8.6.1	v8.6.2
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
Windows 8.1	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: The operating system versions listed above are supported for guests running on the following virtual platforms:

- 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 are supported on XenServer 6.5 or later)
- KVM with Kernel 2.6.32 or later
- Google Cloud

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 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

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.