

Veritas Storage Foundation™ and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL

Windows Server 2003,
Windows Server 2008

5.1



Veritas Storage Foundation and HA Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL

Copyright © 2008 Symantec Corporation. All rights reserved.

Storage Foundation 5.1 for Windows HA

Symantec, the Symantec logo, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THIS DOCUMENTATION IS PROVIDED “AS IS” AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID, SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Windows is a registered trademark of Microsoft Corporation.

Licensing and registration

Storage Foundation for Windows and Storage Foundation HA for Windows are licensed products. See the *Storage Foundation and High Availability Solutions for Windows, Installation and Upgrade Guide* for license installation instructions.

Technical support

For technical assistance, visit <http://www.symantec.com/business/support/index.jsp> and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Section 1 Introduction

Chapter 1 Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL

About the solutions guides	19
Quick Recovery	20
Microsoft clustering	20

Chapter 2 Using the Solutions Configuration Center

About the Solutions Configuration Center	21
Starting the Configuration Center	22
Available options from the Configuration Center	22
About running the Configuration Center wizards	29
Following the workflow in the Configuration Center	30
Solutions wizard logs	32

Section 2 Quick Recovery

Chapter 3 Quick Recovery overview for SQL 2000

About Quick Recovery snapshot solutions	38
Advantages of Quick Recovery snapshot solutions	38
About the components used in Quick Recovery	39
FlashSnap and FastResync	39
Integration with SQL Server Virtual Device Interface (VDI)	40
Configuration requirements	40
Recommendations and best practices	41
VCS, MSCS and VVR considerations	42
VCS and MSCS consideration	42
VVR considerations	42
Vxsnapsql command syntax	43

Chapter 4 Implementing Quick Recovery for SQL 2000

Tasks for implementing Quick Recovery for SQL Server 2000	51
---	----

Reviewing the prerequisites	53
Supported software	53
Storage requirements	54
Configuration requirements	54
Reviewing the configuration	55
Configuring SQL Server storage with Veritas	
Storage Foundation for Windows	56
Creating dynamic disk groups	56
Creating dynamic volumes	58
Pointing the databases and log paths to the SFW volumes	60
Creating the initial snapshot set	61
Creating the snapshot mirrors	61
Creating the snapshot set	61
Manually refreshing the snapshot set	62
Reattaching the split-mirror snapshots	62
Creating the refreshed snapshot set	63
Automatically refreshing the snapshot set	64
Recovering a SQL Server 2000 database	65

Chapter 5 Quick Recovery overview for SQL 2005

About Quick Recovery snapshot solutions	69
About snapshot-assisted backups	70
Advantages of Quick Recovery snapshots	70
Quick Recovery process	71
Methods of implementing Quick Recovery snapshots for SQL 2005	72
About the Quick Recovery Configuration Wizard	72
About the VSS Snapshot Scheduler Wizard	73
About the VSS Snapshot and Snapback wizards and the vxsnap utility	74
About the components used in Quick Recovery	75
FlashSnap and FastResync	75
Integration with Microsoft Volume Shadow Copy Service	75
VCS, Microsoft clustering, and VVR considerations	77

Chapter 6 Preparing to implement Quick Recovery for SQL 2005

Tasks for preparing to implement Quick Recovery for SQL Server 2005 ...	79
Reviewing the prerequisites	80
Supported software	80
Storage requirements and best practices	82
Configuration requirements and best practices	82
Reviewing the configuration	83

	Configuring SQL Server storage with Veritas	
	Storage Foundation for Windows	84
	Creating dynamic disk groups	84
	Creating dynamic volumes	85
	Pointing the databases and log paths to the SFW volumes	88
Chapter 7	Implementing Quick Recovery for SQL 2005 with the configuration wizard	
	About the Quick Recovery Configuration Wizard	89
	Backup types for snapshot sets	91
	About snapshot templates	91
	Tasks for implementing snapshot sets with the configuration wizard	93
	Reviewing the prerequisites	94
	Scheduling and creating snapshot sets	95
	Viewing the status of scheduled snapshots	103
	Troubleshooting scheduled snapshots	104
	Deleting or modifying schedules	106
	Synchronizing schedules after adding a cluster node	107
Chapter 8	Scheduling or creating an individual snapshot set for SQL 2005	
	About scheduling or creating an individual snapshot set	109
	Tasks to schedule a new snapshot set	110
	Tasks to create a one-time snapshot set	111
	Reviewing the prerequisites	112
	Preparing the snapshot mirrors	112
	Scheduling a new snapshot set	114
	Viewing the status of scheduled snapshots	119
	Troubleshooting scheduled snapshots	120
	Deleting or modifying schedules	122
	Creating a one-time snapshot set	123
	Refreshing a snapshot set manually	126
	Reattaching the split-mirror snapshots	127
	Creating the refreshed snapshot set	128
Chapter 9	Recovering a SQL 2005 database	
	About recovering a SQL 2005 database	129
	Tasks for recovering a SQL 2005 database	130
	Prerequisites for recovering a SQL 2005 database	131
	Types of recovery	131
	Recovering using snapshots without log replay	132
	Recovering using snapshots and log replay	134

Restoring snapshots and manually applying logs	137
Recovering missing volumes	139
Preparing for the recovery	140
Performing the recovery	143
Post-recovery steps	145
Vxsnap restore command reference	146

Chapter 10 Vxsnap utility command line reference for SQL 2005

About the Vxsnap utility	149
Vxsnap keywords	150
vxsnap prepare	150
vxsnap create	151
vxsnap reattach	154
vxsnap restore	155

Section 3 Microsoft clustering solutions

About Microsoft clustering solutions	160
--	-----

Chapter 11 Deploying SFW with MSCS: New SQL 2000 installation

Tasks for a new SQL Server 2000 installation	
with SFW and MSCS (Windows Server 2003)	161
Reviewing the requirements	164
Requirements for Veritas Storage Foundation for Windows	164
Disk space requirements	165
System requirements	165
Reviewing the configuration	166
Sample configuration	167
Configuring the storage hardware and network	169
Establishing an MSCS cluster	171
Installing SFW with MSCS/Failover Cluster option	172
SFW installation tasks	172
Pre-installation tasks	172
Installing Veritas Storage Foundation for Windows	174
Post-installation tasks	180
Configuring SFW disk groups and volumes	181
Planning disk groups and volumes	181
Creating dynamic cluster disk groups	183
Creating dynamic volumes	185
Managing disk group and volumes	189

Importing a disk group and mounting a volume	189
Unmounting a volume and deporting a disk group	190
Creating the SQL virtual server group	191
Creating an IP address resource	191
Creating the SQL disk group resource	192
Creating the MSDTC resource	193
Installing SQL Server 2000	194
Applying the SQL SP4 patch	196
Verifying SQL installation	196
Implementing a dynamic quorum resource	196
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	197
Creating the quorum resource for the cluster group	197
Changing the quorum resource to a dynamic mirrored quorum resource	199
Testing the cluster	199

Chapter 12

Deploying SFW with MSCS: New SQL 2005 installation

Tasks for a new SQL Server 2005 installation with SFW and MSCS (Windows Server 2003)	201
Reviewing the requirements	204
Disk space requirements	205
System requirements	205
Reviewing the configuration	206
Sample configuration	207
Configuring the storage hardware and network	209
Establishing an MSCS cluster	210
Installing SFW with MSCS/Failover Cluster option	211
SFW installation tasks	211
Pre-installation tasks	211
Installing Veritas Storage Foundation for Windows	213
Post-installation tasks	219
Configuring SFW disk groups and volumes	220
Planning disk groups and volumes	220
Creating dynamic cluster disk groups	222
Creating dynamic volumes	224
Managing disk group and volumes	228
Importing a disk group and mounting a volume	228
Unmounting a volume and deporting a disk group	229
Creating the SQL virtual server group	230
Creating an IP address resource	231
Creating the SQL disk group resource	231

Creating the MSDTC resource	232
Installing SQL Server 2005	233
Verifying SQL installation	236
Implementing a dynamic mirrored quorum resource	237
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	237
Creating the quorum resource for the cluster group	238
Changing the quorum resource to a dynamic mirrored quorum resource	239
Verifying the cluster configuration	239

Chapter 13 Deploying SFW with Microsoft failover clustering: New SQL 2005 installation

Tasks for a new SQL Server 2005 installation with SFW and Microsoft failover clustering (Windows Server 2008)	242
Reviewing the requirements	244
System requirements	245
Reviewing the configuration	245
Sample configuration	247
Configuring the storage hardware and network	249
Establishing a Microsoft failover cluster	250
Installing SFW with MSCS/Failover Cluster option	253
SFW installation tasks	253
Pre-installation tasks	253
Installing Veritas Storage Foundation for Windows	254
Post-installation tasks	260
Configuring SFW disk groups and volumes	261
Planning disk groups and volumes	261
Creating dynamic cluster disk groups	263
Creating dynamic volumes	265
Managing disk group and volumes	269
Importing a disk group and mounting a volume	269
Unmounting a volume and deporting a disk group	270
Creating the SQL Server virtual server group	271
Creating a Volume Manager Disk Group resource for the application	271
Installing SQL Server 2005	272
Verifying SQL Server 2005 installation	275
Implementing a dynamic mirrored quorum resource	276
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	276
Creating the quorum resource for the cluster group	277

	Changing the quorum resource to a dynamic mirrored quorum resource	277
	Verifying the cluster configuration	278
Chapter 14	Deploying SFW with MSCS and SQL Server in a campus cluster	
	Tasks for a new SQL Server installation with SFW and MSCS in a campus cluster (Windows Server 2003)	282
	Reviewing the requirements	284
	Supported software	284
	System requirements	286
	Disk space requirements	287
	Reviewing the configuration	288
	Overview of campus clustering with MSCS	289
	MSCS campus cluster failure scenarios	290
	MSCS quorum and quorum arbitration	294
	Configuring the network and storage	296
	Establishing an MSCS cluster	298
	Installing and configuring the operating system and MSCS on Server A	298
	Configuring the shared storage and creating a partition for the Cluster quorum disk	299
	Creating the first node of the cluster on Server A	299
	Installing and configuring the operating system and MSCS on Server B	299
	Connecting the two nodes	300
	Creating the second node of the cluster on Server B	300
	Verifying the cluster configuration	300
	Creating the MSDTC resource	301
	Installing SFW	303
	SFW installation tasks	304
	Pre-installation tasks	304
	Installing Veritas Storage Foundation for Windows	306
	Post-installation tasks	311
	Creating disk groups and volumes	312
	Configuring the disks and volumes	313
	Creating a dynamic (cluster) disk group	314
	Creating a volume	316
	Implementing a dynamic quorum resource	321
	Creating a dynamic cluster disk group for the quorum, mirrored	321
	Making the quorum cluster disk group an MSCS resource	322
	Changing the quorum resource to the dynamic mirrored	

quorum resource	323
Setting up a group for SQL Server in MSCS	324
Installing the application on the cluster nodes	325
Verifying the cluster configuration	326

Chapter 15 Deploying SFW with Microsoft failover clustering and SQL Server in a campus cluster

Tasks for deploying SFW with Microsoft failover clustering in a campus cluster (Windows Server 2008)	330
Reviewing the requirements	331
Supported software	332
System requirements	332
Disk space requirements	333
Reviewing the configuration	334
Overview of campus clustering with Microsoft clustering	336
Campus cluster failure with Microsoft clustering scenarios	337
Microsoft clustering quorum and quorum arbitration	341
Configuring the network and storage	343
Establishing a Microsoft failover cluster	345
Connecting the two nodes	346
Installing SFW	347
SFW installation tasks	347
Pre-installation tasks	347
Installing Veritas Storage Foundation for Windows	348
Post-installation tasks	353
Creating disk groups and volumes	355
Configuring the disks and volumes	356
Considerations when creating new volumes	356
Creating a dynamic (cluster) disk group	357
Creating a volume	359
Implementing a dynamic quorum resource	365
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	365
Adding the volume manager disk group for the quorum	365
Changing the quorum resource to the dynamic mirrored quorum resource	366
Setting up a group for SQL Server in the failover cluster	368
Installing the application on the cluster nodes	369
Verifying the cluster configuration	371

Chapter 16	Deploying SFW and VVR with MSCS: New SQL 2000 installation	
	Tasks for a new SQL Server 2000 installation with SFW, VVR, and MSCS (Windows Server 2003)	374
	Reviewing the requirements	375
	Disk space requirements	376
	System requirements	377
	Reviewing the configuration	377
	Sample configuration	379
	Configuring the storage hardware and network	380
	Establishing an MSCS cluster	381
	Installing SFW with MSCS/Failover Cluster option	383
	SFW installation tasks	383
	Pre-installation tasks	383
	Installing Veritas Storage Foundation for Windows	385
	Post-installation tasks	390
	Configuring SFW disk groups and volumes	395
	Planning disk groups and volumes	395
	Creating dynamic disk groups	398
	Creating dynamic volumes	400
	Managing disk groups and volumes	405
	Importing a disk group and mounting a volume	405
	Unmounting a volume and deporting a disk group	406
	Creating the SQL virtual server group	407
	Creating an IP address resource	408
	Creating the disk group resource	408
	Creating the MSDTC resource	409
	Installing SQL Server 2000	410
	Applying the SQL SP4	413
	Verifying SQL installation	413
	Implementing a dynamic mirrored quorum resource	414
	Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	414
	Creating the quorum resource for the cluster group	415
	Changing the quorum resource to a dynamic mirrored quorum resource	416
	Verifying the cluster configuration	417
	Creating a parallel environment on the secondary site	418
	VVR components overview	419
	Creating resources for VVR	420
	Creating an IP address resource	420
	Creating a network name resource	421
	Configuring VVR: Setting up an RDS	422

Creating the RVG resource (primary and secondary sites)	433
Setting the SQL server resource dependency on the RVG resource	434
Working with a solution: Normal operations and recovery procedures ...	435
Monitoring the status of the replication	435
Performing planned migration	435
Replication recovery procedures	437

Chapter 17 Deploying SFW and VVR with MSCS: New SQL 2005 installation

Tasks for a new SQL Server 2005 installation with SFW, VVR, and MSCS (Windows Server 2003)	442
Reviewing the requirements	444
Disk space requirements	445
System requirements	445
Reviewing the configuration	446
Sample configuration	448
Configuring the storage hardware and network	449
Establishing an MSCS cluster	450
Installing SFW with MSCS/Failover Cluster option	452
SFW installation tasks	452
Pre-installation tasks	452
Installing Veritas Storage Foundation for Windows	454
Post-installation tasks	459
Configuring SFW disk groups and volumes	464
Planning disk groups and volumes	464
Creating dynamic cluster disk groups	466
Creating dynamic volumes	469
Managing disk groups and volumes	474
Importing a disk group and mounting a volume	474
Unmounting a volume and deporting a disk group	475
Creating the SQL virtual server group	476
Creating an IP address resource	477
Creating the disk group resource	477
Creating the MSDTC resource	478
Installing SQL Server 2005	479
Verifying SQL installation	483
Implementing a dynamic quorum resource	484
Creating a dynamic cluster disk group and a mirrored volume	484
Creating the quorum resource for the cluster group	485
Changing the quorum resource to a dynamic mirrored quorum resource	486
Verifying the cluster configuration	486
Creating a parallel environment on the secondary site	488

VVR components overview	489
Creating resources for VVR	490
Creating an IP address resource	490
Creating a network name resource	490
Configuring VVR: Setting up an RDS	492
Creating the RVG resource (primary and secondary sites)	503
Setting the SQL server resource dependency on the RVG resource	504
Working with the solution: Normal operations and recovery procedures	506
Normal operations	506
Performing planned migration	506
Replication recovery procedures	508

Chapter 18 Deploying SFW and VVR with Microsoft failover clustering: New SQL 2005 installation

Tasks for a new SQL 2005 installation with SFW, VVR, and Microsoft failover clustering (Windows Server 2008)	512
Reviewing the prerequisites	514
Supported software for Microsoft failover clusters with SFW	514
Disk space requirements	514
System requirements	515
Reviewing the configuration	515
Sample configuration	517
Configuring the storage hardware and network	518
Establishing an Microsoft failover cluster	520
Installing SFW with MSCS/Failover Cluster option	522
SFW installation tasks	522
Pre-installation tasks	522
Installing Veritas Storage Foundation for Windows	523
Post-installation tasks	528
Configuring SFW disk groups and volumes	533
Planning disk groups and volumes	533
Creating dynamic cluster disk groups	535
Creating dynamic volumes	537
Managing disk groups and volumes	542
Importing a disk group and mounting a volume	542
Unmounting a volume and deporting a disk group	543
Completing the primary site configuration	544
Creating a parallel environment on the secondary site	545
VVR components overview	546
Creating resources for VVR	547
Setting up the replicated data sets (RDS) for VVR	548
Creating the RVG resource (primary and secondary sites)	559

Setting the SQL Server resource dependency on the RVG resource	560
Working with the solution: Normal operations and recovery procedures	562
Normal operations	562
Performing planned migration	562
Replication recovery procedures	564
Index	567

Introduction

- [Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL](#)
- [Using the Solutions Configuration Center](#)

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL

This chapter covers the following topics:

- [About the solutions guides](#)
- [Quick Recovery](#)
- [Microsoft clustering](#)

About the solutions guides

The *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL* contains solutions using the following:

- [Quick Recovery](#)
- [Microsoft clustering](#)

Solutions for Microsoft SQL for High Availability and Disaster Recovery are in *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*.

Separate guides are available for Microsoft Exchange solutions and for other application solutions.

Quick Recovery

Quick Recovery is the process of creating and maintaining on-host point-in-time copies of production data that can be used to quickly recover SQL Server databases that have been subject to accidental or malicious updates or other data corruption.

A Quick Recovery solution serves as a first line of defense to recover SQL Server databases that have been subject to accidental or malicious updates. Quick Recovery is designed to augment your traditional backup methodology.

Microsoft clustering

Microsoft clustering solutions are covered in separate chapters according to operating system:

- Microsoft Cluster Server (MSCS) on Windows Server 2003
- Microsoft failover clustering on Windows Server 2008

Microsoft clustering may be used with Veritas Storage Foundation for Windows to provide high availability for SQL Server. Microsoft clustering may be used with Veritas Storage Foundation for Windows and Veritas Volume Replicator to provide replication support for SQL Server.

Using the Solutions Configuration Center

This chapter covers the following topics:

- [About the Solutions Configuration Center](#)
- [Starting the Configuration Center](#)
- [Available options from the Configuration Center](#)
- [About running the Configuration Center wizards](#)
- [Following the workflow in the Configuration Center](#)
- [Solutions wizard logs](#)

About the Solutions Configuration Center

The Storage Foundation and High Availability Solutions Configuration Center guides you through setting up your Veritas Storage Foundation for Windows (SFW) or SFW High Availability (HA) environment. The Configuration Center provides solutions for the following applications:

- Microsoft Exchange Server 2003 and 2007
- Microsoft SQL Server 2000 and 2005
- Enterprise Vault Server (high availability solution only)
- Additional applications

You can use the Configuration Center and its wizards to set up your environment for any combination of the following solutions:

- High availability at a single site for a new installation
- High availability at a single site for an existing server

- Campus cluster disaster recovery, including the following:
 - Campus cluster using Veritas Cluster Server (SFW HA)
 - Campus cluster using Microsoft clustering
- Wide area disaster recovery involving multiple sites
- Quick Recovery for on-host recovery from logical errors in application data (available for Microsoft Exchange 2003 and 2007 and for Microsoft SQL Server 2005)
- Fire drill to test the fault readiness of a disaster recovery environment that uses VVR replication

The Solutions Configuration Center provides two ways to access Solutions wizards:

- The Applications tab lists solutions by application. It provides step-by-step configuration instructions that include buttons to launch the appropriate wizard for each step.
- The Solutions tab, for advanced users, lists wizards by solution without additional instructions.

Starting the Configuration Center

You can start the Configuration Center in two ways:

- Click **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.
- Click **Start > Run** and type **scc**.

Available options from the Configuration Center

On the Applications tab, the Solutions Configuration Center is context-sensitive to the application. For example, the Solution Guides listed in the right pane match the selected application.

In addition, some choices can vary depending on the operating system of the node on which you launch the wizard. For example, since Microsoft Exchange 2003 runs only on 32-bit operating systems, on a 64-bit system only the Exchange 2007 configuration wizard is shown.

[Figure 2-1](#) shows the choices available on a 32-bit system when you click Solutions for Microsoft Exchange.

Figure 2-1 Solutions Configuration Center for Microsoft Exchange

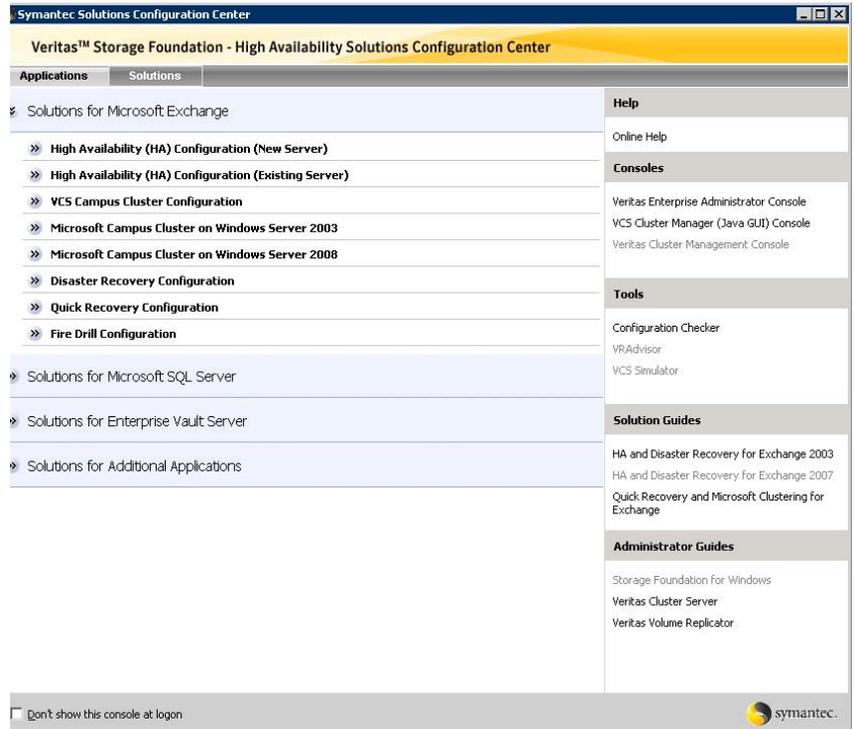


Figure 2-2 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 2-2 Solutions Configuration Center for Microsoft SQL Server

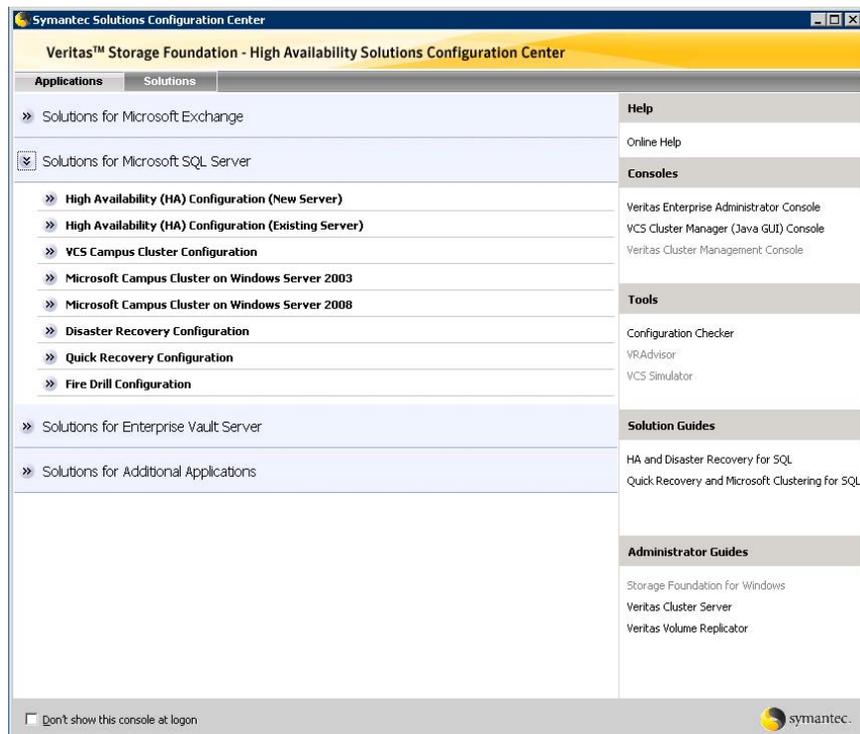


Figure 2-3 shows the choices available when you click Solutions for Enterprise Vault Server.

Figure 2-3 Solutions Configuration Center for Enterprise Vault Server

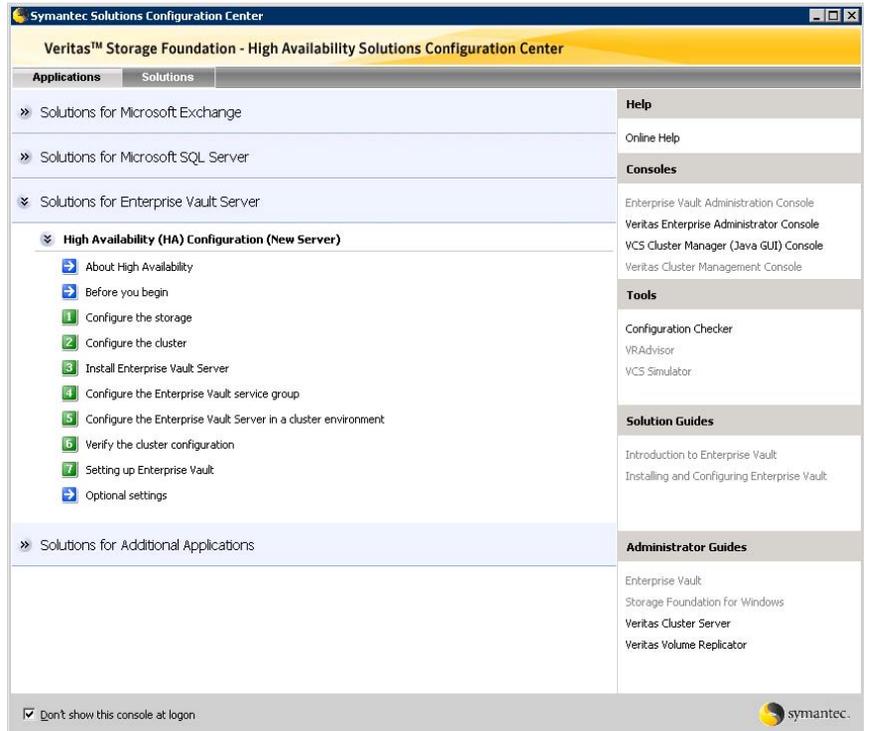
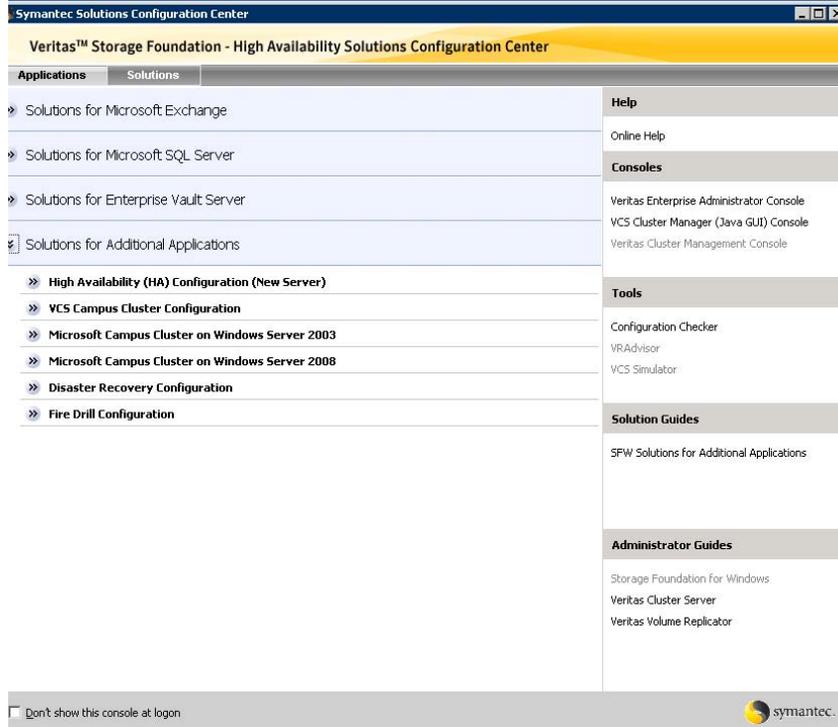


Figure 2-4 shows the choices available when you click Solutions for Additional Applications.

Figure 2-4 Solutions Configuration Center for additional applications



The submenu choices also vary by application. For example, different steps, information, or wizards are shown under High Availability (HA) Configuration for Exchange than those shown for SQL Server.

Figure 2-5 shows one of the steps for implementing high availability for Exchange.

Figure 2-5 Context-sensitive step for Exchange



Figure 2-6 shows one of the steps for implementing high availability for SQL Server.

Figure 2-6 Context-sensitive step for SQL Server



Figure 2-7 shows one of the steps for implementing high availability for Enterprise Vault Server.

Figure 2-7 Context-sensitive step for Enterprise Vault Server

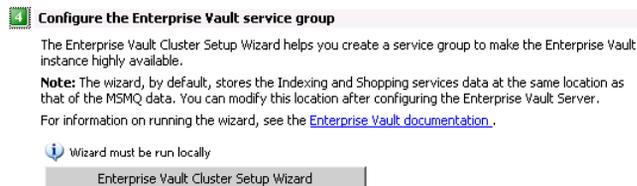


Figure 2-8 shows one of the steps for implementing high availability for additional applications.

Figure 2-8 Context-sensitive step for additional applications

4 Configure the service group

Create a service group to make your application or server role highly available.

[Read how to create a File Share service group](#)

Wizard must be run locally

File Share Configuration Wizard

[Read how to create a Print Share service group](#)

Wizard must be run locally

Print Share Configuration Wizard

[Read how to create an IIS Server service group](#)

Wizard must be run locally

IIS Configuration Wizard

[Read how to create a Microsoft Virtual Server virtual machine service group](#)

Wizard must be run locally

MSVirtual Machine Configuration Wizard

[Read how to create an Oracle service group](#)

Wizard must be run locally

Oracle Agent Configuration Wizard

[Read how to create a service group for an application, process or service](#)

Wizard must be run locally

Application Configuration Wizard

About running the Configuration Center wizards

You can run the wizards from the Applications tab if you are walking through the configuration steps on the Solutions Configuration Center. If you are already familiar with configuration, you can also go directly to a particular wizard by selecting the Solutions tab.

The Configuration Center and some wizards can be run from a remote system. Wizards that you can run remotely include the following:

VCS Configuration Wizard	Sets up the VCS cluster
Disaster Recovery Configuration Wizard	Configures wide area disaster recovery, including cloning storage, cloning service groups, and configuring the global cluster Also can configure Veritas Volume Replicator (VVR) replication or configure the VCS resource for EMC SRDF and Hitachi TrueCopy array-based hardware replication. Requires first configuring high availability on the primary site
Quick Recovery Configuration Wizard	Schedules preparation of snapshot mirrors and schedules the Quick Recovery snapshots
Fire Drill Wizard	Sets up a fire drill to test disaster recovery Requires configuring disaster recovery first

Wizards related to storage configuration and application installation must be run locally on the system where the process is occurring. Wizards that you must run locally include the following:

New Dynamic Disk Group Wizard	Launched from the Veritas Enterprise Administrator console
New Volume Wizard	Launched from the Veritas Enterprise Administrator console
Exchange Setup Wizard	Installs and configures Exchange for the high availability environment If Exchange is already installed, refer to the documentation for further instructions.
Exchange Configuration Wizard	Configures the service group for Exchange high availability
SQL Server Configuration Wizard	Configures the service group for SQL Server high availability You must first install SQL Server on each node according to the instructions in the documentation.

Enterprise Vault Cluster Setup Wizard Configures the service group for Enterprise Vault Server high availability.

In addition, the Additional Applications section of the Configuration Center provides wizards to be run locally for creating service groups for the following applications or server roles:

File Share Configuration Wizard Configures FileShare for high availability.

Print Share Configuration Wizard Configures PrintShare for high availability.

IIS Configuration Wizard Configures IIS for high availability.

MSVirtual Machine Configuration Wizard Configures MS Virtual Machine for high availability.

Oracle Agent Configuration Wizard Configures Oracle for high availability.

Application Configuration Wizard Configures any other application service group for which application-specific wizards have not been provided.

Following the workflow in the Configuration Center

During the multi-step High Availability Configuration workflow, you may find it helpful to run an SFW HA client on another system and leave the Configuration Center open on that system. In this way, you can see what step comes next, drill down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format.

When setting up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration.

[Figure 2-9](#) shows the high-level overview of the workflow steps for configuring high availability for Exchange from the Solutions Configuration Center.

Figure 2-9 Workflow for configuring Exchange high availability

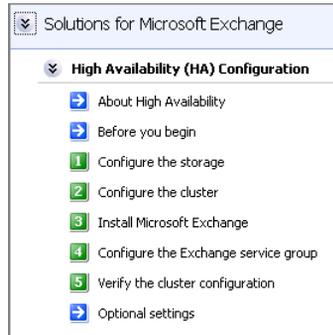


Figure 2-10 shows the high-level overview of the workflow steps for configuring high availability for SQL Server from the Solutions Configuration Center.

Figure 2-10 Workflow for configuring SQL Server high availability

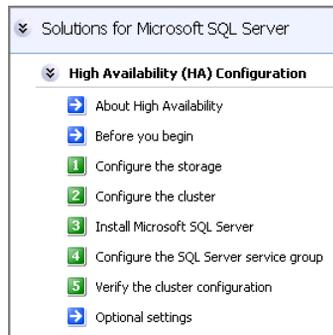


Figure 2-11 shows the high-level overview of the workflow steps for configuring high availability for Enterprise Vault Server from the Solutions Configuration Center.

Figure 2-11 Workflow for configuring high availability for Enterprise Vault Server

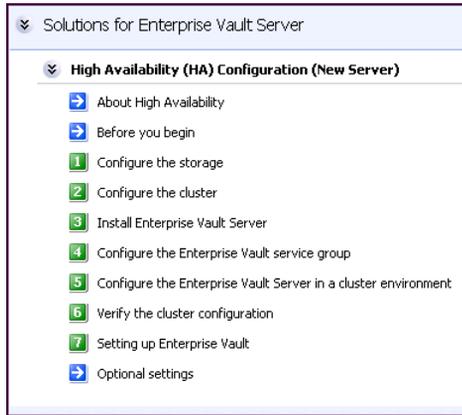
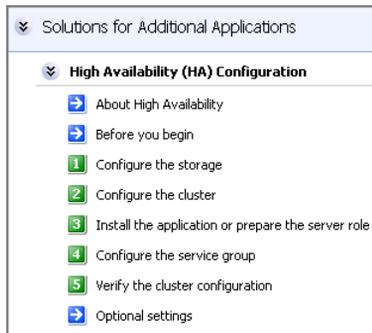


Figure 2-12 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

Figure 2-12 Workflow for configuring high availability for additional applications



Solutions wizard logs

The Solutions Configuration Center provides access to many wizards. However, three wizards are built in to the Solutions Configuration Center:

- Disaster Recovery Wizard
- Fire Drill Wizard
- Quick Recovery Configuration Wizard

These three Solutions wizards are launched only from the Solutions Configuration Center, whereas other wizards can be launched from product consoles or the Start menu.

Logs created by these three Solutions wizards are located in the following paths:

For Windows Server 2003:

C:\Documents and Settings\All Users\Application
Data\VERITAS\winsolutions\log

For Windows Server 2008:

C:\ProgramData\Veritas\winsolutions\log

Quick Recovery

This section includes the following chapters:

- [Quick Recovery overview for SQL 2000](#)
- [Implementing Quick Recovery for SQL 2000](#)
- [Quick Recovery overview for SQL 2005](#)
- [Preparing to implement Quick Recovery for SQL 2005](#)
- [Implementing Quick Recovery for SQL 2005 with the configuration wizard](#)
- [Scheduling or creating an individual snapshot set for SQL 2005](#)
- [Recovering a SQL 2005 database](#)
- [Vxsnap utility command line reference for SQL 2005](#)

Quick Recovery overview for SQL 2000

This chapter covers the following topics:

- [About Quick Recovery snapshot solutions](#)
- [Advantages of Quick Recovery snapshot solutions](#)
- [About the components used in Quick Recovery](#)
- [Configuration requirements](#)
- [Recommendations and best practices](#)
- [VCS, MSCS and VVR considerations](#)
- [Vxsnapsql command syntax](#)

About Quick Recovery snapshot solutions

Quick Recovery is the term used to describe the process of creating and maintaining on-host point-in-time copies of production data that can be used to quickly recover SQL Server databases that have been subject to accidental or malicious updates or other data corruption.

Veritas Storage Foundation for Windows (SFW) provides a Quick Recovery solution for SQL 2000. The Quick Recovery solution for SQL 2000 integrates with SQL Server Virtual Device Interface (VDI) to produce a snapshot set containing on-host, disk-based snapshots that provide fast recovery from logical errors and eliminate the time-consuming process of restoring data from tape. Databases can be restored to a specific point-in-time, recovered using current logs to the point-of-failure, or restored to the time when the snapshot set was created or refreshed.

`Vxsnapsql` integrates Veritas FlashSnap™ and SQL Server Virtual Device Interface (VDI) to create a snapshot set that provides a complete copy of the database at the time the `vxsnapsql create` command is issued. The snapshot set contains snapshots of the volumes associated with a database and its transaction log as well as metadata about the database.

SFW Quick Recovery solution uses a split-mirror snapshot method. A snapshot is a separate persistent volume that contains an exact duplicate of all the data on the original volume at the time the snapshot is taken. This type of persistent physical snapshot is also known as a Clone (HP) or a BCV (EMC). Copy-on-write snapshots, also known as metadata snapshots, only copy changed blocks to the snapshot and do not create a separate physical volume.

Veritas FlashSnap technology is also integrated into the Veritas NetBackup 6.0 Advanced Client Option and Symantec Backup Exec 10d Advanced Disk-based Backup Option. These products are the preferred solution for on and off host snapshot-assisted backup.

Advantages of Quick Recovery snapshot solutions

A Quick Recovery solution serves as a first line of defense to recover SQL Server databases that have been subject to accidental or malicious updates. Quick Recovery is designed to augment your traditional backup methodology. Maintaining a snapshot set requires just the few seconds it takes to detach a split-mirror snapshot from its original volume. On-host snapshot recovery is faster than restoring a full backup from tape or other media; on-host snapshot recovery reduces downtime and helps meet service-level agreements for application availability.

In addition to the primary benefit of recovery from logical errors, snapshot sets can be moved over a SAN to another server and used for other purposes including:

- Application tuning and testing--data can be updated and modified in a realistic environment without impacting users.
- Business reporting and decision analysis--up-to-date data is available with minimal impact on the production environment.

About the components used in Quick Recovery

The Quick Recovery solution uses Veritas FlashSnap and FastResync technology along with SQL Server VDI to quiesce a database and ensure a persistent snapshot of the data.

FlashSnap and FastResync

Veritas FlashSnap provides the ability to create and maintain the on-host point in time copies that are integral to the Quick Recovery solution. FastResync is a FlashSnap feature that optimizes the resynchronization of a snapshot volume and its original volume.

FlashSnap

FlashSnap is the multi-step process that is used to create and maintain split-mirror, persistent snapshots that are copies of the original volumes they mirror. Both the original and snapshot volume may consist of multiple physical devices, as in the case of RAID 0+1 (Mirrored Striped) volumes. FlashSnap cannot be used with software RAID-5 volumes. FlashSnap includes the following commands:

- **Prepare**
Creates a snapshot mirror and attaches it to the original volume.
- **Snap Shot**
Detaches the snapshot mirror from the original volume. This split-mirror snapshot volume is an exact duplicate of the original volume at the point in time the snapshot command is executed.
- **Snap Back**
Reattaches the snapshot volume to the original volume. The volumes can be resynchronized using either the original volume or the snapshot volume as the source. If a logical error has occurred on the original database volume, then the snapshot volume can be quickly restored to a consistent, point-in-time image.

- **Snap Clear**
Permanently removes the association between the snapshot volume and the original volume.
- **Snap Abort**
Aborts the snapshot operation after a Prepare or Snap Back command is issued. Snap Abort permanently removes the snapshot mirror from the volume and releases its space.

FastResync

The FastResync capability optimizes the resynchronization of a snapshot volume and its original volume. FlashSnap uses FastResync technology to track the changed blocks in an original volume after a snapshot is detached. A Disk Change Object (DCO) volume is automatically created to store a record of these changes. When the snapshot volume is resynchronized with the original volume using the Snap Back command, only the changed data blocks are written to the snapshot volume. This greatly reduces the time and performance impact of resynchronization which means that a snapshot set can be refreshed with minimal impact to production.

Integration with SQL Server Virtual Device Interface (VDI)

SFW integrates with VDI to perform snapshot operations on SQL Server database volumes while the database is online and available. VDI quiesces the database for the short period of time required to create the snapshot and then immediately thaws it. This quiescing allows SQL snapshots to be taken while the database application remains active.

FlashSnap integrates with VDI and uses the `vxsnapsql` command to provide the ability to detach multiple split-mirror snapshot volumes simultaneously. This allows you to snapshot all volumes associated with a SQL Server database at exactly the same point in time without taking the database offline. When viewed together with the metadata file that is also created, these snapshots form a snapshot set of the database. These persistent FlashSnap snapshots, taken through VDI, can later be used to quickly recover a database that has been subjected to accidental or malicious updates or otherwise corrupted.

Configuration requirements

- The system and boot volumes must reside on a separate disk (Harddisk0) from the dynamic volumes used for the SQL user-defined databases and split-mirror snapshots.

- User-defined database and transaction logs must be stored on disks within a single dynamic disk group.
- Each database must be in a separate volume, but the volumes may share the same dynamic disks.

Note: If multiple databases share a volume, then all the databases in the volume will be restored even if only one database is corrupt. This is by design, the Restore operation is done at a volume level.

- Disk groups must be of a Storage Foundation for Windows 4.0 or later version. Upgrade any disk groups created using an earlier version of Volume Manager for Windows before using the `vxsnapsql` command.
- User-defined database and transaction logs may not be stored in the same volume as the SQL Server program files or system data files.
- In order to perform a roll-forward recovery to the point of failure, database and transaction logs must be in separate volumes.

Recommendations and best practices

The following recommendations enable you to take advantage of SFW storage configuration functionality as you manage your SQL Server storage:

- The hardware for Quick Recovery should include sufficient storage to be able to create dynamic volumes on separate disks or LUNs for the following purposes:
 - Databases
 - Transaction logs
 - Split-mirror snapshots of the database and transaction log volumes
- For the split-mirror snapshots, select disks or LUNs that are not used for production data. However, you may create more than one snapshot volume on the same disk or LUN as long as there is sufficient space available.
- Database and transaction logs must be in separate volumes in order to perform a roll-forward recovery to the point of failure.
- Database and transaction logs should be on separate disks so that disk failure does not affect both the database and transaction logs.
- Transaction logs should always be configured in a redundant layout. The preferred software layout is RAID 0+1 (mirrored striped) volumes as this provides better read and write performance than RAID 1 (mirrored) alone.

The transaction log will generate the most I/O and thus should use the highest performance disks available.

- The preferred layout for the database is hardware RAID 5, software RAID 1 (mirrored with logging enabled) or software RAID 0+1 (mirrored striped).

Note: FlashSnap is not supported for software RAID 5 volumes.

VCS, MSCS and VVR considerations

Special precautions must be taken when implementing a snapshot solution with a VCS or MSCS cluster or with the Veritas Volume Replicator option (VVR).

VCS and MSCS consideration

Store the XML metadata file on a volume in the cluster disk group associated with the SQL database so that the metadata file is available from all nodes in the cluster. Specify the full path to the location of the XML metadata file for commands requiring the *filename* variable.

VVR considerations

- Store the XML metadata file on a volume that is included in the replicated data set so that the metadata file is available at the secondary site. Additionally, if VCS or MSCS is used, store the metadata file on a volume in the cluster disk group associated with the SQL database so that the metadata file is available from all nodes in the cluster. Specify the full path to the location of the XML metadata file for commands requiring the *filename* variable.
- During a point in time recovery, the volumes on the secondary site lose write-order fidelity. DCM automatically becomes active to ensure data consistency between the primary and secondary sites. While DCM is active, the volumes cannot be expanded by either manual or AutoGrow operations. Perform a manual resynchronization of the secondary to deactivate DCM.

To resynchronize the node at the secondary site with the node at the primary site:

- 1 Right-click on the primary RVG and click **Resynchronize Secondaries**.
- 2 Click **Yes** to resynchronize the nodes.

Vxsnapsql command syntax

The `vxsnapsql` utility is designed to simultaneously snapshot all the volumes associated with a SQL Server database.

`vxsnapsql` has the following keywords:

<code>start</code>	Creates snapstart mirrors on the specified disks.
<code>create</code>	Creates a simultaneous snapshot of all the volumes in the specified component.
<code>restore</code>	Restores a corrupted database using the snapshot set. Restore can be to a specified point-in-time, to the point-of-failure, or to the time of the snapshot.
<code>reattach</code>	Reattaches and resynchronizes an existing snapshot set to the original database volumes.

Typing the following sequence brings up a description of the command syntax:

```
vxsnapsql <Keyword> -?
```

Keywords or operands

```
vxsnapsql start
vxsnapsql [-b] start [server=<ServerName>]
[instance=<InstanceName>] component=<DBname>
[<Diskname>]...
```

This command creates snapstart mirrors of the volumes in the database (component) in preparation for creating a snapshot set using the `vxsnapsql create` command.

The following attributes apply:

<code>-b</code>	Runs the command in the background.
<code>server=</code> <code><ServerName></code>	Name of the local system (localhost) or of the SQL virtual server created in a clustered environment. The default, localhost, need not be specified.
<code>instance=</code> <code><InstanceName></code>	SQL Server instance name. This parameter is not required for a default instance.

<code>component=</code> <code><DBname></code>	SQL Server database name as it appears in the SQL Enterprise Manager. Snapstart mirrors are created for all the data and log volumes associated with the specified SQL Server database.
<code><Diskname></code>	Name of the disk or disks where the snapstart mirrors will be created, for example, <code>harddisk2</code> .

Example

```
vxsnapsql -b start instance=Accounting  
component=TestDB harddisk3
```

This command creates a snapstart mirror (snapplex) on `harddisk3` for each data or log volume associated with the SQL Server database `TestDB` in the `Accounting` instance. The snapstart mirrors remain synchronized with the original volumes until the `vxsnapsql create` command is issued.

```
vxsnapsql create  
vxsnapsql -x <Filename> [-o] create  
[server=<ServerName>][instance=<InstanceName>]  
component=<DBname> [<snapshot_tuple>...]
```

A `snapshot_tuple` consists of a number of attribute=value fields. Within each tuple, the source and snapshot volume attributes are paired by forward slashes (/). Although the entire `snapshot_tuple` is optional, if you choose to define the tuple for one volume, or you must define it for all volumes in the component. The tuple will be of the following form:

```
source=volume [/Newvol=SnapVol] [/plex=SnapPlexName]  
[/DriveLetter={A|B|...|Z}][/Label=<VolLabel>]  
[/DrivePath=<DrivePath>]
```

This command creates split-mirror snapshots of the volumes in the specified database (component) and a metadata file containing information about the database and snapshot volumes. Together the snapshots and metadata file form the snapshot set.

The following attributes apply:

<code>-x <Filename></code>	Indicates the name assigned to the XML metadata file that is created to store the snapshot information. The metadata file is used for restore operations. Symantec recommends that the file name include the ".xml" extension. The default path to the file is \Documents and Settings\All Users\Application Data\Veritas\SQLBACKUP. If you wish to place the file in another directory, specify a full path before the file name, for example J:\XML\Image1.xml. In a clustered environment, store the metadata file on shared storage.
<code>-o</code>	Overwrites an existing XML metadata file of the same name.
<code>server=<ServerName></code>	Name of the local system (localhost) or of the SQL virtual server created in a clustered environment. The default, local host, need not be specified.
<code>instance=<InstanceName></code>	SQL Server instance name. This parameter is not required for a default instance.
<code>component=<DBname></code>	SQL Server database name as it appears in the SQL Enterprise Manager. Snapshots are created for all the data and log volumes associated with the specified SQL Server database.
<code><snapshot_tuple></code>	Consists of a number of attribute=value fields. The source and snapshot volume attributes are paired by forward slashes (/). Although the entire snapshot_tuple is optional, if you choose to define the tuple for one volume, or you must define it for both the database and log volumes.

<code>source=<Volume></code>	Indicates the source volume for the split-mirror snapshot specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\". Repeat this parameter for each volume associated with the specified SQL Server database.
<code>[/Newvol=<SnapVol>]</code>	Specifies the name of the new snapshot volume that is created. If the name is not specified with this option, the form "SnapVolume01" is created. The full device path will be: \\Device\HarddiskDmVolumes\ <DiskGroupName>\<SnapVolName>
<code>[/plex=<SnapPlexName>]</code>	Specifies the name of the snapshot mirror (plex) to be detached. Use this parameter if there are multiple snap plexes available for the snapshot.
<code>[/DriveLetter={A B ... Z}]</code>	Specifies the drive letter assigned to the new snapshot volume.
<code>[/Label=<volLabel>]</code>	Specifies the volume label assigned to the new snapshot volume.
<code>[/DrivePath=<DrivePath>]</code>	Specifies the drive path assigned to the new snapshot volume. The drive path must reference an existing empty local NTFS folder. The path must include the drive letter and folder to be mounted, for example, C:\DB1VOL.

Examples

```
vxsnapsql -x TestDB.xml create component=TestDB  
source=M:/driveletter=R source=N:/driveletter=S
```

This command creates split-mirror snapshots of the data and log volumes associated with the TestDB database, in this example volumes M and N. TestDB is in a default instance running on the local host. The snapshots are assigned drive letters R and S respectively. Additionally, a metadata file, TestDB.xml, is created and stored in the default directory \Documents and Settings\All Users\Application Data\Veritas\SQLBACKUP.

```
vxsnapsql -x TestDB.xml create server=SQLVS  
instance=acctpay  
component=TestDB  
source=E:\DB1\Data/DrivePath=F:\DB1Snapshot\Data  
source=E:\DB1\TLog/DrivePath=F:\DB1Snapshot\TLog
```

This command creates snapshots of the volumes associated with the TestDB database in the acctpay instance on the SQL virtual server SQLVS. The original volumes are mounted on E:\DB1\Data and E:\DB1\TLog. The resulting snapshot volumes are assigned mount points F:\DB1Snapshot\Data and F:\DB1Snapshot\TLog respectively. The metadata involved in this operation is stored in TestDB.xml in the default directory \Documents and Settings\All Users\Application Data\Veritas\SQLBACKUP.

```
vxsnapsql restore  
vxsnapsql -x <Filename> [-b] [-f] [-s] restore  
{RestoreType=[RECOVERY|NO_RECOVERY]} [noLogs]  
[logFiles=<tlog1,tlog2,...>]
```

This command restores the snapshot volumes in the snapshot set and is used to recover a corrupted or missing SQL Server database. After a restore with the recovery option, the database is left in an online state. After a restore with the no_recovery option, the database is left in a loading state. The following attributes apply:

- | | |
|---------------|---|
| -x <Filename> | The metadata file created by the vxsnapsql create command. Each snapshot set must have a unique name for the metadata file. |
| -b | Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete. |
| -f | Forces the operation. Make sure the volume is not in use before using this option. |
| -s | Silent mode. Allows the restore operation to proceed without user interaction. |

<code>RestoreType=</code> <code>[RECOVERY NO_RECOVERY]</code>	Specifies the type of database recovery, either recovery or no recovery. <code>RECOVERY</code> leaves the database in an online state. <code>NO_RECOVERY</code> leaves the database in a loading state.
<code><noLogs></code>	Database and transaction log files are restored from the snapshot set. No other logs are applied. The database is left in an online state.
<code>logFiles=<tlg1,</code> <code>tlg2,...></code>	Transaction log backup files to be applied with the <code>RECOVERY</code> option to achieve a point of failure recovery and leave the database in an online state. Each transaction log must have a unique name and be created using the “overwrite existing media” option within SQL Server.

The options can be applied as follows:

Table 3-1 Recovery options

Selected options	Database state after recovery	Description
<code>RECOVERY,</code> <code>logFiles=tlg1,</code> <code>tlg2,...</code>	online	Database and transaction log volumes are restored and the specified backup transaction logs are applied.
<code>RECOVERY,</code> <code>noLogs</code>	online	Database and transaction log volumes are restored and mounted. No additional transaction logs are applied.
<code>NO_RECOVERY</code>	loading	Database and transaction logs are restored. The database is left in a loading state so that backup logs can be replayed to a specified point in time.

Exclusive access to the SQL Server database is required for this operation. Before using this command verify that the source volumes and the snapshot volumes are not in use. In the SQL Enterprise Manager, close the tree view to the server level.

Examples

Point in Time Restore

```
vxsnapsql -x TestDB.xml restore
RestoreType=NO_RECOVERY
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and leaves the database in a loading state so that backup logs can be manually restored to a specific point in time.

Point of Failure Restore

```
vxsnapsql -x TestDB.xml restore RestoreType=RECOVERY
logFiles=c:\backup\tLog1.bak, c:\tLog2.bak
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and then applies the specified transaction log backups (c:\backup\tLog1.bak and c:\tLog2.bak) and brings the database online.

Time of Snapshot Restore

```
vxsnapsql -x TestDB.xml restore RestoreType=RECOVERY
noLogs
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and brings the database online. The database is restored to the time the snapshot set was created or last refreshed.

```
vxsnapsql reattach
```

```
vxsnapsql -x <Filename> [-b] [-f] reattach
```

This command reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes. (This command is similar to a snapback operation.)

The following attributes apply:

- | | |
|---------------|---|
| -x <Filename> | The file created by the vxsnapsql create command. Each snapshot set must have a unique name for the metadata file.

Note: This file is deleted after the reattach operation has completed successfully |
| -b | Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete. |

-f Forces the operation. Make sure the volume is not in use before using this option.

Example

```
vxsnapsql -x TestDB.xml reattach
```

This command uses the information in the TestDB.xml file to reattach and resynchronize all the volumes in the snapshot set. This xml file is deleted after the reattach operation has completed successfully. The snapshot volumes remain synchronized with the original volumes until the vxsnapsql create command is issued.

Implementing Quick Recovery for SQL 2000

This chapter covers the following topics:

- [Tasks for implementing Quick Recovery for SQL Server 2000](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring SQL Server storage with Veritas Storage Foundation for Windows](#)
- [Creating the initial snapshot set](#)
- [Manually refreshing the snapshot set](#)
- [Automatically refreshing the snapshot set](#)
- [Recovering a SQL Server 2000 database](#)

Tasks for implementing Quick Recovery for SQL Server 2000

The Quick Recovery process creates and maintains snapshot sets (on-host copies of production volumes) that can be used to recover SQL databases in the event of corruption or an accidental or malicious update. The Quick Recovery process consists of three phases:

- Creating an initial snapshot set
- Refreshing a snapshot set
- Recovering a corrupted database

Table 4-1 outlines the high-level objectives and the tasks to complete each objective for implementing Quick Recovery for SQL Server 2000:

Table 4-1 Tasks for implementing SQL Server 2000 Quick Recovery

Objective	Tasks
“Reviewing the prerequisites” on page 53	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 55	<ul style="list-style-type: none"> ■ Reviewing the sample configuration
“Configuring SQL Server storage with Veritas Storage Foundation for Windows” on page 56	<ul style="list-style-type: none"> ■ Creating a dynamic disk group ■ Creating dynamic volumes ■ Creating a SQL Server database
“Creating the initial snapshot set” on page 61	<ul style="list-style-type: none"> ■ Creating snapshot mirrors ■ Creating the snapshot set
“Manually refreshing the snapshot set” on page 62	<ul style="list-style-type: none"> ■ Reattaching the snapshot mirrors ■ Creating the refreshed snapshot set
“Automatically refreshing the snapshot set” on page 64	<ul style="list-style-type: none"> ■ Using Veritas Net Backup, Veritas Backup Exec, or Windows task schedule to automatically refresh the snapshot set ■ Calling a refresh script from Veritas Backup Exec
“Recovering a SQL Server 2000 database” on page 65	<ul style="list-style-type: none"> ■ Restoring to a specified point in time ■ Restoring to the point of failure ■ Restoring to the point-in-time the snapshot set was created or refreshed

Reviewing the prerequisites

This solution assumes that the required software is already installed and configured. Refer to other solutions in this guide, the *Veritas Storage Foundation for Windows Administrator's Guide* and the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for installation and configuration information.

Supported software

- Veritas Storage Foundation 5.1 for Windows (SFW) with the FlashSnap option.
or
- Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the FlashSnap option.
- Any of the following Microsoft SQL 2000 servers and their operating systems:

Microsoft SQL Server 2000
 Standard Edition or Enterprise
 Edition (SP4 required)

- Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
- Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

Microsoft SQL Server 2000 (64-bit)
 Enterprise Edition

- Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)

Microsoft SQL Server 2000 (64-bit)
 Standard Edition or Enterprise
 Edition (SP4 required)

- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions)
- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions)

Storage requirements

- The hardware for Quick Recovery should include sufficient storage to be able to create dynamic volumes on separate disks or LUNs for the following purposes:
 - Databases
 - Transaction logs
 - Split-mirror snapshots of the database and transaction log volumes
- For the split-mirror snapshots, make sure to select disks or LUNs that are not used for production data. However, you may create more than one snapshot volume on the same disk or LUN as long as there is sufficient space available.

Configuration requirements

- The system and boot volumes must reside on a separate disk (Harddisk0) from the dynamic volumes used for the SQL user-defined databases and split-mirror snapshots.
- Database and transaction logs must be stored on disks within a single dynamic disk group.
- Disk groups must be of a Storage Foundation for Windows 4.0 or later version. Upgrade any disk groups created using an earlier version of Volume Manager for Windows before using the `vxsnapsql` command.
- Quick Recovery snapshots are supported only on volumes belonging to an SFW dynamic disk group. They are not supported on volumes belonging to a Microsoft Disk Management Disk Group. For more information on Microsoft Disk Management Disk Groups, see *Veritas Storage Foundation Administrator's Guide*.
- User-defined database and transaction logs may not be stored in the same volume as the SQL Server program files or system data files.
- In order to perform a roll-forward recovery to the point of failure, database and transaction logs must be in separate volumes.

Refer to “[Recommendations and best practices](#)” on page 41 for guidelines on planning your implementation.

Reviewing the configuration

This section describes how to create, refresh, and restore a snapshot set for a specified SQL Server 2000 database. A sample setup containing the following object names is used to illustrate the create, refresh, and restore tasks on a local server.

Name	Drive Letter	Object
billing		SQL Server instance name
billing_DB		component (SQL Server database name)
harddisk3		Disk where the snapstart mirrors will be created
billing_DG		disk group name for the Billing disk group
billing_data	L:	volume for Billing database user data
billing_log	M:	volume for Billing database log
billing_datasnap	S:	volume for Billing database user data snapshot
billing_logsnap	T:	volume for Billing database log snapshot
billing_DB.xml		name for the metadata file generated during the <code>create</code> command
tlog1.bak, tlog2.bak		log files saved at <code>c:\backup</code> prior to a <code>vxsnapsql restore</code> command

Configuring SQL Server storage with Veritas Storage Foundation for Windows

To use the SFW Quick Recovery snapshot functionality with SQL database components, you must store the database components on Veritas Storage Foundation for Windows (SFW) dynamic volumes. Configuring your SQL Server storage includes the following tasks:

- Creating one or more dynamic disk groups
- Creating volumes for the databases and transaction logs
- Using SQL Server Enterprise Manager to create a new database and set the appropriate paths to point to the new SFW volumes.

If your SQL Server environment is already configured with SFW, skip this section and proceed to “[Creating the initial snapshot set](#)” on page 61.

Creating dynamic disk groups

Create one or more dynamic disk groups. If your SQL Server production server is in a clustered environment, choose the cluster disk group option. See the appropriate chapter in the High Availability section of this guide for further configuration information.

To create a dynamic disk group from the VEA console

- 1 Click **Start > All Programs > Symantec > Veritas Enterprise Administrator** and if prompted to select a profile, select a profile (or Default).
- 2 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
For the local system you can specify **localhost**.
- 3 If prompted to do so, specify the user name, password, and domain for the system.
- 4 In the tree expand the system name and expand the storage agent.
- 5 Right-click **Disk Groups**, and click **New Dynamic Disk Group**.
- 6 On the **Welcome** screen of the New Dynamic Disk Group Wizard, click **Next**.
- 7 Enter a name for the disk group (for example, `billing_DG`).
- 8 For an off-host or cluster environment, choose from the following:
 - For a cluster environment, check the **Create cluster group** check box.
 - For an off-host environment, check the **Add private group protection** check box.

- 9 Select the appropriate disks in the Available disks list and click the **Add** button to move them to the Selected disks list.
- 10 Click **Next**.
- 11 Click **Next** to upgrade the selected disks.
- 12 Click **Finish** to create the new disk group.

To create a dynamic disk group from the command line

Type the following command:

For an standalone environment

```
> vxdbg -gbilling_DG init Harddisk1 Harddisk2 Harddisk3
```

For a cluster environment

```
> vxdbg -gbilling_DG -s init Harddisk1 Harddisk2 Harddisk3
```

where `-gbilling_DG` is the name of the dynamic disk group you want to create and `Harddisk1`, `Harddisk2`, and `Harddisk3` are the disks included in the dynamic disk group.

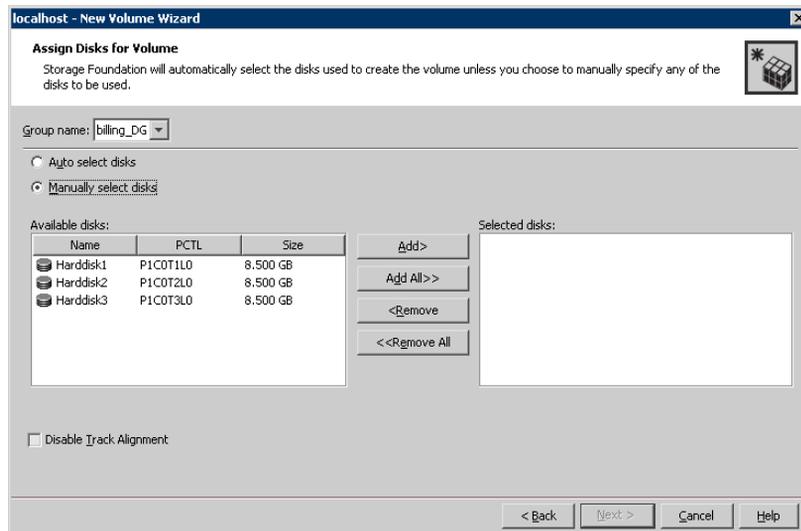
For the complete syntax for the `vxdbg init` command, see *Veritas Storage Foundation Administrator's Guide*.

Creating dynamic volumes

Create volumes for the database and transaction log.

To create a dynamic volume from the VEA console

- 1 Start the VEA and connect to the appropriate host.
- 2 In the tree, expand the system name, expand the storage agent, and then expand **Disk Groups**.
- 3 Right-click on the disk group in which to create the volumes (for example, billing_DG), and click **New Volume**.
- 4 In the Welcome panel of the New Volume Wizard, click **Next**.
- 5 Select the disks for the volume:

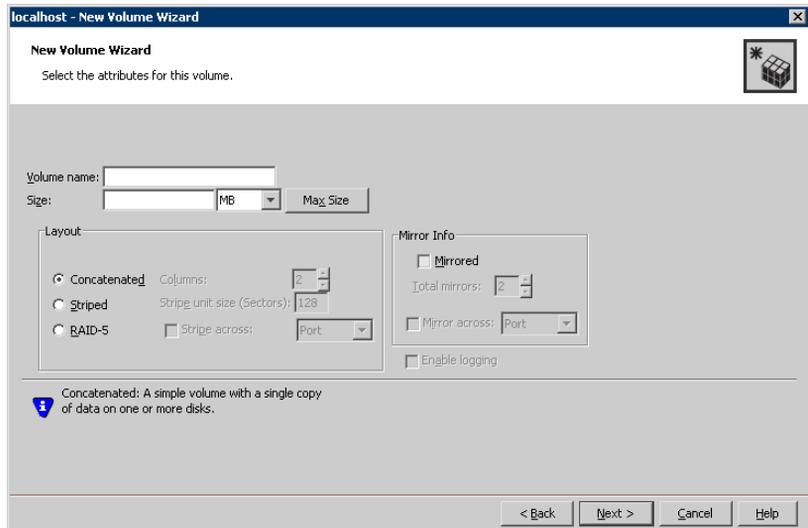


- Confirm that the **Group name** is the correct disk group (for example, billing_DG). If necessary, select the correct disk group name from the drop-down menu.
- Specify automatic or manual disk selection. Symantec recommends using the **Manually select disks** option.
- Select the appropriate disks in the Available disks list, and click the **Add** button to move them to the Selected disks list.
- You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does

not store blocks of data in alignment with the boundaries of the physical track of the disk.

- Click **Next**.

6 Specify the parameters of the volume:



- Enter the volume name (for example, `billing_data`).
- Enter the size.
- Select the layout.
- Select the appropriate mirror options.
- Click **Next**.

7 Assign a drive letter to the volume (for example, L: for the `billing_data` volume) and click **Next**.

8 Create an NTFS file system:

- Accept the default **Format this volume**.
- Click **NTFS**.
- Select an allocation size.
- Accept the default file system label, which is the same as the volume name you entered previously or enter a file system label.
- If desired, select **Perform a quick format**.
- Click **Next**.

9 Review the volume specifications, then click **Finish** to create the new volume.

- 10 Repeat the previous steps as necessary to create volumes for the transaction log (for example, `billing_log`) and any other database in your configuration.

To create a volume from the command line

- 1 Type the CLI commands:

```
> vxassist [-b] -gbilling_DG make billing_data L:  
> vxassist [-b] -gbilling_DG make billing_log M:
```

This will create volumes in the `billing_DG` disk group named `billing_data` and `billing_log` on drive letters L and M respectively.

- 2 Modify and repeat this command as necessary to create additional volumes. For the complete syntax of the `vxassist make` command, see *Veritas Storage Foundation Administrator's Guide*.

Pointing the databases and log paths to the SFW volumes

Create a new database and point the database and log paths to the SFW volumes.

To create a new SQL Server database

- 1 Click **Start > All Programs > Microsoft SQL Server > Enterprise Manager** to open the SQL Server Enterprise Manager.
- 2 Expand the Server icon located in the appropriate SQL Server Group.
- 3 Right-click the Databases folder and click **New Database**.
 - Enter a name for the new database, for example, `billing`.
 - On the **General** tab of the Database Properties, name the database "Test."
 - Click the **Data Files** tab.
 - Click the **Browse** button in the location column and set the Database file location to the drive letter or mount point of the volume created for the database, for example `L:`.
 - Click the **Transaction Log** tab.
 - Click the **Browse** button location column and set the transaction log location to the drive letter or mount point of the volume created for the database, for example `M:`.
 - Click **OK**.

Creating the initial snapshot set

Creating a snapshot is a two-step process. Use the `vxsnapsql start` command to create snapshot mirrors for the database and log volumes. These snapshots mirrors remain synchronized with the original volumes until the `vxsnapsql create` command is issued. The `create` command creates the snapshot set by detaching the snapshot mirrors from the original volumes and creating separate on-host snapshot volumes as well as an XML file for the metadata. The `vxsnapsql create` command integrates with VDI to quiesce the database and then snapshots the database and transaction log volumes simultaneously. This is done while the database is online and available. Once a snapshot set has been created, it can be refreshed quickly since the time-consuming `vxsnapsql start` step is not required.

Creating the snapshot mirrors

Create snapshot mirrors for the database and log volumes. Make sure to select disks or LUNs that are not used for production data. However, you may create more than one snapshot volume on the same disk or LUN as long as there is sufficient space available.

To create snapshot mirrors

- ◆ Type the command:

```
> vxsnapsql start instance=billing component=billing_DB  
harddisk3
```

The complete syntax of the `vxsnapsql start` command is:

```
vxsnapsql [-b] start [server=ServerName]  
[instance=InstanceName] component=DBname [diskname]...
```

Note: Make sure that the lower pane of the VEA console shows that the resynchronization process is complete before continuing with the `vxsnapsql create` command.

Creating the snapshot set

Create the snapshot set using the `vxsnapsql create` command. This provides a complete picture of the database at the time the snapshots are taken.

To create the snapshot set

- ◆ Type the command:

```
vxsnapsql -x billing_DB.xml create instance=billing  
component=billing_DB source=L:/driveletter=S  
source=M:/driveletter=T
```

This will create a snapshot set with the snapshot volume of the `billing_data` volume mounted on S: and the snapshot volume of the `billing_log` volume mounted on T:

The complete syntax of the `vxsnapsql create` command is:

```
vxsnapsql -x filename [-o] create  
[server=ServerName][instance=InstanceName]  
component=DBname [snapshot_tuple...]
```

The `snapshot_tuple` will be of the following form:

```
source=volume[/NewVol=SnapVol][/plex=SnapPlexName]  
[/DriveLetter={A/B/.../Z}][/Label=VolLabel]  
[/DrivePath=DrivePath]
```

The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes.

Manually refreshing the snapshot set

Periodically refresh or update your snapshot set so that it contains a current copy of the original volumes. Refreshing your snapshot set is a two-step process that can easily be incorporated into your regular backup routine.

Refreshing a snapshot set includes the following tasks:

- [“Reattaching the split-mirror snapshots”](#) on page 62
- [“Creating the refreshed snapshot set”](#) on page 63

Reattaching the split-mirror snapshots

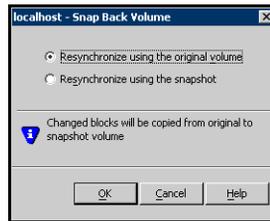
You can use the `vxsnapsql reattach` command to reattach all the split-mirror snapshots in the snapshot set with a single command.

Alternatively, from the VEA console, repeat the **snap back** operation for each volume in the snapshot set. Both commands reattach a snapshot volume to its original volume and use FastResync technology to automatically update the snapshot mirror and synchronize it with the original volume, applying only the changes tracked in the DCO volume.

To reattach the split-mirror snapshots to the original volumes from the VEA console

- 1 Close all open handles on the database and snapshot volumes.

- 2 Right-click the appropriate snapshot volume (for example, `billing_datasnap`), select **Snap** and then **Snap Back**.



- 3 Select the **Resynchronize using the original volume** option.
- 4 Click **OK**.
 You may get an error message asking if you want to force the command. This indicates there is an open handle to the volume. Click **No**, close any open handles and retry the command.
- 5 Repeat [step 2](#) through [step 4](#) for the other snapshot volumes in the snapshot set (for example `billing_logsnap`).

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close all open handles on the database and snapshot volumes.
- 2 Type the following command:
`> vxsnapsql -x billing_DB.xml reattach`

The complete syntax for the `reattach` command is:

```
vxsnapsql -x Filename [-b] [-f] reattach
```

Creating the refreshed snapshot set

Snapshot all the volumes simultaneously using the `vxsnapsql create` command described in the section “[Creating the snapshot set](#)” on page 61. If you choose to reuse the name for the metadata file, then include the `-o` option.

- ◆ Type the command:

```
vxsnapsql -x billing_DB.xml -o create instance=billing  

component=billing_DB source=L:/driveletter=S  

source=M:/driveletter=T
```

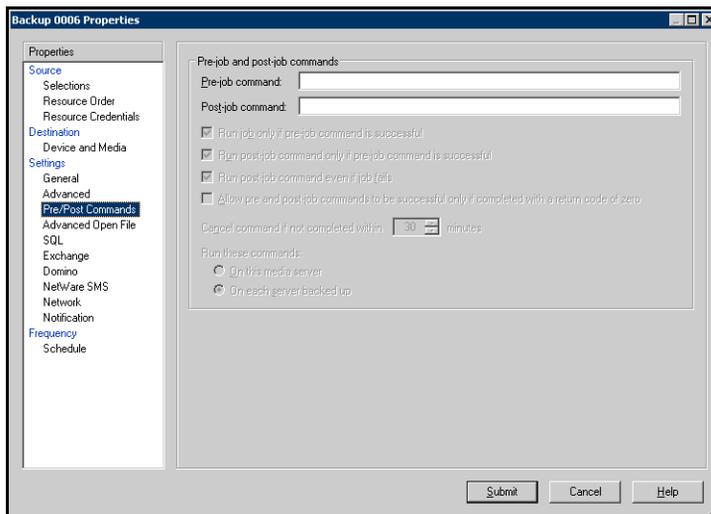
Automatically refreshing the snapshot set

The `vxsnapsql reattach` and `vxsnapsql create` commands can be called from either the `bpend_notify.bat` file in Veritas NetBackup or from a batch file in a pre/post command scheduled to run at the completion of a Veritas Backup Exec for Windows Servers backup job. Additionally, a script can be called by the Windows Task Scheduler to enable automatic updates of the snapshot set on a regular basis.

The example below describes how to add a post command to call the script from Veritas Backup Exec.

To call a refresh script from Veritas Backup Exec

- 1 Write a refresh script that calls the `vxassist snapback` and `vxsnapsql create` commands described in “[Manually refreshing the snapshot set](#)” on page 62 used to refresh the snapshot set.
- 2 Launch Veritas Backup Exec.
- 3 Click the **Job Setup** tab.
- 4 Right-click the job that is run to back up your SQL Server database and select **Properties**.
- 5 From the **Properties** list, expand **Settings** and click **Pre/Post Commands**.



- 6 Enter the path to your refresh script in the **Post-job command** text box.
- 7 Click **Submit**.

Refer to the *Veritas Backup Exec for Windows Servers Administrator's Guide* for more information.

Recovering a SQL Server 2000 database

Use the snapshot volumes in a snapshot set to restore a corrupt database. You can restore a database to the following points:

- A specified point in time
- The point of failure
- The point in time that the snapshot set was created (or last refreshed)

The complete syntax of the `vxsnapsql restore` command is:

```
vxsnapsql -x filename [-b] [-f] [-s]restore  
  {RestoreType=RECOVERY|NO_RECOVERY} [noLogs]  
  [logFiles=log1,log2,...]
```

To restore to a specified point in time

Caution: Before using this command, use your preferred method to backup the transaction logs within SQL Server. You must use the “overwrite existing media” option to create uniquely-named backup files.

- 1 Backup the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL Enterprise Manager GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 3 Restore the snapshot set and leave the database in a loading state. Type the command:

```
> vxsnapsql -x billing_DB.xml restore RestoreType=NO_RECOVERY
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnapsql create` command. The database and log snapshot volumes are restored and the SQL Server database is left in a loading state.
- 4 Use your preferred method to manually restore the backup transaction logs to the desired point-in-time.
- 5 To refresh the snapshot set use the `vxsnapsql create` command to create a new snapshot of all the volumes in the database.

To perform an automatic roll-forward recovery to the point of failure

Caution: Before using this command, use your preferred method to backup the transaction logs within SQL Server. You must use the “overwrite existing media” option to create uniquely-named backup files.

- 1 Backup the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL Enterprise Manager GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 3 Type the command:

```
vxsnapsql -x billing_DB.xml restore RestoreType=RECOVERY  
logFiles=c:\backup\tLog1.bak, c:\tLog2.bak
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnapsql create` command and `c:\backup\tLog1.bak`, `c:\tLog2.bak` are the paths to the transaction log backup files. After the most recent backup log is replayed, the SQL Server database is closed and left in an operational state.
- 4 To refresh the snapshot set use the `vxsnapsql create` command to create a new snapshot of all the volumes in the database.

To recover a database to the time of the snapshot set

- 1 Close the SQL Enterprise Manager GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Type the command:

```
vxsnapsql -x billing_DB.xml restore RestoreType=RECOVERY  
noLogs
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnapsql create` command. The volumes in the snapshot set are restored and the database is left in an operational state. The database is restored to the time the snapshot set was created or last refreshed.

- 3 To refresh the snapshot set use the `vxsnapsql create` command to create a new snapshot of all the volumes in the database.

Note: Refer to the troubleshooting chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for troubleshooting information.

Quick Recovery overview for SQL 2005

This chapter covers the following topics:

- [About Quick Recovery snapshot solutions](#)
- [About snapshot-assisted backups](#)
- [Advantages of Quick Recovery snapshots](#)
- [Quick Recovery process](#)
- [Methods of implementing Quick Recovery snapshots for SQL 2005](#)
- [About the components used in Quick Recovery](#)
- [VCS, Microsoft clustering, and VVR considerations](#)

About Quick Recovery snapshot solutions

Veritas Storage Foundation for Windows (SFW) provides the capability to create a point-in-time image of all the volumes associated with a SQL Server 2005 instance. This image, called a snapshot set, is a complete copy of the SQL Server 2005 instance at the specific point in time the snapshots are taken.

SFW uses Veritas FlashSnap™ technology along with the Microsoft Volume Shadow Copy Service (VSS) framework to quiesce the database and ensure a persistent snapshot of the production data.

Quick Recovery is the term for creating and maintaining the SFW snapshot sets on-host for use in quickly recovering databases in the event of corruption or an accidental or malicious update. The Quick Recovery solution provides fast recovery from logical errors and eliminates the time-consuming process of restoring data from tape. Databases can be recovered to the point in time when

the snapshot was taken or, by using current logs, rolled forward to the point of failure.

Using the SFW Quick Recovery Configuration wizard, you can create multiple snapshot sets for each SQL Server 2005 instance and set up schedules for creating and refreshing the snapshot sets. The snapshot sets can be maintained on-host as a Quick Recovery solution.

If you are using Veritas Volume Replicator (VVR) for replication, you can also synchronize snapshot sets on the secondary site. See *Veritas Volume Replicator, Administrator's Guide*.

SFW snapshots use a split-mirror snapshot method. The snapshot is a separate persistent volume that contains an exact duplicate of all the data on the original volume at the time the snapshot is taken. This type of persistent physical snapshot is also known as a Clone (HP) or a BCV (EMC). In contrast, copy-on-write snapshots, also known as metadata snapshots, only copy changed blocks to the snapshot and do not create a separate physical volume.

Because a snapshot set contains a split-mirror snapshot copy of each of the volumes in the database, the snapshot set requires the same amount of space as the original volumes.

Veritas FlashSnap technology is also integrated into the Veritas NetBackup 6.0 Advanced Client Option and Symantec Backup Exec 10d Advanced Disk-based Backup Option. These products are the preferred solution for on and off host snapshot-assisted backup.

About snapshot-assisted backups

Since a snapshot set created using SFW is an exact copy of the production data, it can be used with a backup application in place of the production volumes to produce regular backups to tape or other media. A snapshot set can be used for backup either on-host or moved to a secondary server.

A snapshot-assisted backup provides the following advantages:

- Reduces backup I/O load on the production volumes in an on-host environment, allowing backup to be completed with less impact to the users
- Eliminates backup I/O load on the production server in an off-host environment

Advantages of Quick Recovery snapshots

A Quick Recovery solution serves as a first line of defense to recover databases that have been subject to accidental or malicious updates. Quick Recovery is designed to augment your traditional backup methodology. Maintaining a

snapshot set requires just the few seconds it takes to detach a split-mirror snapshot from its original volume. On-host snapshot recovery is faster than restoring a full backup from tape or other media; on-host snapshot recovery reduces downtime and helps meet service-level agreements for application availability.

In addition to the primary benefit of recovery from logical errors, snapshot sets can be moved over a SAN to another server and used for other purposes including:

- Application tuning and testing—data can be updated and modified in a realistic environment without impacting users.
- Business reporting and decision analysis—up-to-date data is available with minimal impact on the production environment.

Quick Recovery process

The Quick Recovery process can be broken down into the following phases:

- Creating an initial snapshot set
This has two stages:
 - Preparing the mirror for the snapshot set
This stage takes a while and should be scheduled for a time of low activity.
 - Creating the initial snapshot set by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy
- Periodically refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again, as needed or according to a pre-set schedule
This stage is automated by setting up snapshot schedules using the Quick Recovery wizard or VSS Snapshot Scheduler wizard.
- Using a snapshot set to recover a corrupted SQL Server 2005 instance or single database

Methods of implementing Quick Recovery snapshots for SQL 2005

Veritas Storage Foundation for Windows provides more than one method of implementing Quick Recovery snapshots for SQL 2005.

- Quick Recovery Configuration Wizard
- VSS Snapshot Scheduler
- VSS Snapshot and Snapback wizards and the vxsnap utility

[Table 5-1](#) summarizes the methods and when you would use them.

Table 5-1 Methods of implementing Quick Recovery snapshots

Tasks	Method	For more information
<ul style="list-style-type: none"> ■ Set up and schedule multiple snapshot sets for multiple databases <p>or</p> <ul style="list-style-type: none"> ■ Perform major updates to an existing snapshot set schedule 	<p>From the Solutions Configuration Center: Quick Recovery Configuration Wizard</p>	<p>“About the Quick Recovery Configuration Wizard” on page 72</p>
<ul style="list-style-type: none"> ■ Add a new snapshot set schedule for one database 	<p>From the VEA: VSS Snapshot Scheduler Wizard</p> <p>You can also use the Quick Recovery Configuration Wizard.</p>	<p>“About the VSS Snapshot Scheduler Wizard” on page 73</p>
<ul style="list-style-type: none"> ■ Create a one-time snapshot as needed 	<p>From the VEA: VSS SQL Snapshot Wizard and VSS Snapback Wizard</p> <p>You can also use the Vxsnap utility from the CLI.</p>	<p>“About the VSS Snapshot and Snapback wizards and the vxsnap utility” on page 74</p>

About the Quick Recovery Configuration Wizard

The Quick Recovery Configuration Wizard provides the most complete method of configuring snapshot sets and is therefore recommended for initial implementation. The Quick Recovery Configuration Wizard enables you to schedule all phases of snapshot creation and maintenance:

- Preparing the mirror for the snapshot set.

- Creating the initial snapshot set by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy.
- Periodically refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again.

The wizard enables you to set up and schedule multiple snapshot sets for one or more databases in a SQL Server 2005 instance. You can set up one or more schedules for each snapshot set. You can schedule when to prepare the snapshot mirrors, when to create the initial snapshot sets, and when to refresh the snapshot sets, enabling you to establish a schedule that best suits your site. For example, you can schedule mirror preparation, the most time-consuming step, for a time of low activity. The scheduling capability automates the process of refreshing the snapshot sets. At the scheduled times, the snapshot volumes are automatically reattached, resynchronized, and then split again. Once configured and applied, snapshot schedules are maintained by the Veritas Scheduler Service, which runs in the background.

The snapshot creation process integrates with VSS to quiesce the database and then simultaneously snapshot the volumes in the database. The snapshot is done while the database is online and without disrupting the database operations.

The wizard also enables you to save all configuration settings to an XML file that can be imported for use on other systems or instances.

See [Chapter 7, “Implementing Quick Recovery for SQL 2005 with the configuration wizard”](#) on page 89.

About the VSS Snapshot Scheduler Wizard

You can use the VSS Snapshot Scheduler Wizard to add a snapshot schedule for a selected database. This wizard uses the same scheduling process as the Quick Recovery wizard. However, you can only schedule one database at a time.

Unlike the Quick Recovery Configuration Wizard, which enables you to perform all aspects of snapshot configuration, the VSS Snapshot Scheduler Wizard does not include preparing the snapshot volume. Therefore, you must use the Prepare command before running the VSS Snapshot Scheduler Wizard to schedule a new snapshot set. Also, you can only schedule one database at a time.

You might use this wizard to add a schedule to the initial configuration set up with the Quick Recovery Configuration Wizard. For example, you configured a daily snapshot for Quick Recovery use and now want to add a weekly snapshot on a different disk for backup use.

Note: Adding a snapshot schedule using the VSS Snapshot Scheduler will not update the template settings created with the Quick Recovery Configuration Wizard. If you want to keep the template settings up to date, you should instead run the Quick Recovery Configuration Wizard to modify the schedule.

See [Chapter 8, “Scheduling or creating an individual snapshot set for SQL 2005”](#) on page 109.

About the VSS Snapshot and Snapback wizards and the vxsnap utility

If you need to create an individual snapshot set, you can do so either from the command line, using the vxsnap command line utility, or from the VEA console using the VSS Snapshot and Snapback wizards.

Unlike using the Quick Recovery Configuration Wizard, which enables you to perform all aspects of snapshot configuration, you must prepare the mirror for the snapshot manually, using the Prepare command (or vxsnap utility). In addition, these methods do not include the capability of scheduling periodic refreshing of the snapshot set. Instead you would need to manually reattach the mirror, allowing it to resynchronize, and then create a snapshot set again from the resynchronized mirror.

Therefore, these methods are best suited for a one-time special-purpose snapshot. If you need to keep the snapshot set up to date, you should instead set up a snapshot schedule using the Quick Recovery Configuration Wizard or the VSS Snapshot Scheduler Wizard.

For more information about the VSS Snapshot and snapback wizards, see [Chapter 8, “Scheduling or creating an individual snapshot set for SQL 2005”](#) on page 109.

For more information about the vxsnap command line utility, see [Chapter 10, “Vxsnap utility command line reference for SQL 2005”](#) on page 149.

Note: The vxsnap commands must be invoked on a local system. On Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode.

Table 5-2 shows the vxsnap commands that perform the same actions as the VSS wizards.

Table 5-2 Actions of VSS wizard and vxsnap command

Action	VSS wizard	vxsnap command
Creates a snapshot set consisting of snapshots of all the volumes in the specified SQL Server instance.	VSS Snapshot	<code>create</code>
Reattaches and resynchronizes a snapshot set to the production database volumes.	VSS Snapback	<code>reattach</code>

About the components used in Quick Recovery

SFW Quick Recovery snapshots use Veritas FlashSnap and FastResync technology along with the Microsoft Volume Shadow Copy Service framework.

FlashSnap and FastResync

Veritas FlashSnap provides the ability to create and maintain the on-host point in time copies of volumes that are integral to the snapshot solutions. Both the original and snapshot volume may consist of multiple physical devices, as in the case of RAID 0+1 (Mirrored Striped) volumes. FlashSnap cannot be used with software RAID-5 volumes.

FastResync is a FlashSnap feature that optimizes the resynchronization of a snapshot volume and its original volume. FlashSnap uses FastResync technology to track the changed blocks in an original volume after a snapshot is detached. A Disk Change Object (DCO) volume is automatically created to store a record of these changes. When the snapshot volume is resynchronized with the original volume, only the changed data blocks are written to the snapshot volume. This greatly reduces the time and performance impact of resynchronization which means that a snapshot set can be refreshed with minimal impact to production.

Integration with Microsoft Volume Shadow Copy Service

SFW integrates with the Windows Volume Shadow Copy Service (VSS) as both a VSS Requestor and a VSS Provider. This integration is provided by FlashSnap.

The Volume Shadow Copy Service (VSS) process allows the databases of a SQL Server 2005 instance to be frozen before the snapshot operation occurs and then

thawed immediately after it. This quiescing allows for Microsoft supported and guaranteed persistent snapshots of your data.

FlashSnap integrates with VSS to create a snapshot set containing snapshot volumes of all the volumes associated with a SQL Server 2005 instance without taking the databases offline.

VSS framework

There are four components to the VSS framework: Requestor, Writer, Provider, and the Volume Shadow Copy Service itself.

Table 5-3 VSS framework components

Component	Action
Volume Shadow Copy Service	Talks to and coordinates the Requestor, Provider, and Writer.
Requestor	As a Requestor, the vxsnap component of FlashSnap notifies the VSS coordinator to initiate the VSS request to prepare the SQL Server 2005 instance for quiescing and later requests that the snap shot process begin.
Writer	As Writers, VSS-enabled applications such as SQL Server 2005 respond to requests to prepare and participate in the generation of snapshots, provide success/failure status, and provide information about the application including what is to be backed up and restored, and restore strategy.
Provider	As a Provider, FlashSnap creates the persistent snapshot.

VSS process

FlashSnap integrates with Volume Shadow Copy Service as both a VSS Requestor and a VSS Provider.

The following steps occur, in the order presented, during the snapshot process:

- Acting as a VSS Requestor, FlashSnap notifies the VSS coordinator service to prepare for a split-mirror snapshot of a SQL Server 2005 instance.
- The VSS coordinator service calls the SQL Server VSS Writer to find out which volumes contain the databases and transaction logs.
- The VSS coordinator service notifies the FlashSnap VSS Provider to prepare for the snapshot.
- Acting as a VSS Requestor, FlashSnap requests that the VSS coordinator service begin the snapshot call.

- The VSS coordinator service notifies the SQL Server VSS Writer to quiesce and freeze the databases in preparation for the snapshot. When this is accomplished, the SQL Server Writer informs the VSS coordinator service to proceed.
- The VSS coordinator service calls the FlashSnap Provider to create the split-mirror snapshot by detaching the snapshot volume from the original volume. The snapshot process takes a maximum of 10 seconds. After the snapshot volume is detached, the FlashSnap Provider informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the SQL Server VSS Writer to thaw (release the freeze) and resume normal I/O.

The following steps occur, in the order presented, during the restore process:

- Acting as a VSS Requestor, FlashSnap notifies the VSS coordinator service to prepare for a restore operation.
- The VSS coordinator service calls the SQL Server VSS Writer, which prepares for the restore operation.
- The FlashSnap utility restores the snapshot volumes. After the snapshot operation completes, the FlashSnap utility informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the SQL Server VSS Writer to process the post-restore operations.

VCS, Microsoft clustering, and VVR considerations

Certain requirements apply when implementing a snapshot solution with a VCS, Microsoft clustering, and VVR environment.

In a VCS or Microsoft clustering environment, observe the following precautions:

- The XML metadata file for each snapshot set along with files that store snapshot schedule information are created in a folder on the local drive by default. In a cluster environment, store these files on shared storage so that the files are available from all nodes in the cluster. The snapshot XML files should be stored separately from the volumes that are included in snapshots.
 - If you use the Quick Recovery Configuration Wizard to create the snapshot set, you can use the wizard to specify the file path to the appropriate volume.
 - If you use a VSS wizard to create the snapshot set, you can store the XML files in a location of your choice using the following method: Use a

text editor to create a text file named "**redirect.txt**." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, G:\BackupSets. Save the **redirect.txt** file in the default VSS XML file directory C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML on each node of the cluster.

- When using vxsnap utility commands that require the filename attribute, specify the full path to the location of the XML metadata file.
- If you plan to use the Quick Recovery or VSS Snapshot Scheduler wizard to specify scripts (commands) to be run before or after a snapshot, store the scripts on shared storage so that they are available to all nodes.
- If you set up a snapshot schedule with the Quick Recovery wizard and later add a node to the cluster, you can run the wizard again to synchronize schedules on the existing nodes with the new node.

In a VVR environment, observe the following precautions:

- Store the XML metadata file and other snapshot related files on a volume that is included in the replicated data set so that the metadata file is available at the secondary site. Additionally, if VCS is used, store the metadata file on a volume in the cluster disk group associated with the database so that the metadata file is available from all nodes in the cluster.
- During a point-in-time recovery, the volumes on the secondary site lose write-order fidelity. DCM automatically becomes active to ensure data consistency between the primary and secondary sites. While DCM is active, the volumes cannot be expanded by either manual or AutoGrow operations. You must perform a manual resynchronization of the secondary to deactivate DCM.

Preparing to implement Quick Recovery for SQL 2005

This chapter covers the following topics:

- [Tasks for preparing to implement Quick Recovery for SQL Server 2005](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring SQL Server storage with Veritas Storage Foundation for Windows](#)

Tasks for preparing to implement Quick Recovery for SQL Server 2005

[Table 6-1](#) outlines the high-level objectives and the tasks to complete each objective for preparing to implement SQL Server 2005 Quick Recovery.

Table 6-1 Tasks for preparing to implement SQL Server 2005 Quick Recovery

Objective	Tasks
“Reviewing the prerequisites” on page 80	■ Verifying hardware and software prerequisites and storage configuration best practices
“Reviewing the configuration” on page 83	■ Reviewing the sample configuration

Table 6-1 Tasks for preparing to implement SQL Server 2005 Quick Recovery

Objective	Tasks
“Configuring SQL Server storage with Veritas Storage Foundation for Windows” on page 84	<ul style="list-style-type: none">■ Creating a dynamic disk group■ Creating dynamic volumes■ Creating a SQL Server database and pointing the database and log paths to the SFW volumes

Reviewing the prerequisites

A Veritas Storage Foundation for Windows (SFW) Quick Recovery solution can be implemented on either a standalone system or a clustered system. Quick Recovery snapshots are supported in both Veritas Cluster Server (VCS) and Microsoft clusters.

This solution assumes that the required software is already installed and configured.

Refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for installation and configuration information for Veritas Storage Foundation for Windows or Veritas Storage Foundation HA for Windows.

Supported software

Quick Recovery snapshots require the following software:

- Veritas Storage Foundation 5.1 for Windows (SFW) with the FlashSnap option.
or
- Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the FlashSnap option.
- Any of the following Microsoft SQL servers and their operating systems:

- | | |
|---|--|
| <p>Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required) ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| <p>Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| <p>Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| <p>Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2008 for 64-bit Itanium (IA64) ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |

Storage requirements and best practices

Review the following best practices and requirements for snapshot storage:

- The hardware for Quick Recovery should include sufficient storage to be able to create dynamic volumes on separate disks or LUNs for the following purposes:
 - Databases
 - Transaction logs
 - Split-mirror snapshots of the database and transaction log volumes
- For the split-mirror snapshots, make sure to select disks or LUNs that are not used for production data.
- You can create a snapshot of more than one volume on the same disk or LUN as long as there is sufficient space available.

Configuration requirements and best practices

Review the following configuration requirements and best practices:

- The system and boot volumes must reside on a separate disk (Harddisk0) from the dynamic volumes used for the SQL user-defined databases and split-mirror snapshots.
- Disk groups must be of a Storage Foundation for Windows 4.0 or later version. Upgrade any disk groups created using an earlier version of Volume Manager for Windows before creating Quick Recovery snapshots.
- Quick Recovery snapshots are supported only on volumes belonging to an SFW dynamic disk group. They are not supported on volumes belonging to a Microsoft Disk Management Disk Group. For more information on Microsoft Disk Management Disk Groups, see *Veritas Storage Foundation Administrator's Guide*.
- Database and transaction logs must be stored on disks within a single dynamic disk group.
- Database and transaction logs should be on separate disks so that disk failure does not affect both the database and transaction logs.
- User-defined database and transaction logs may not be stored in the same volume as the SQL Server program files or system data files.
- In order to perform a roll-forward recovery to the point of failure, database and transaction logs must be in separate volumes.
- The snapshot XML files should be stored separately from the volumes that are included in snapshots.

- Locate snapshot volumes on separate disks from any database or log volumes so that the snapshot process will not interfere with database operations.
- Locate the snapshot volumes for each database on separate disks from snapshots of other databases. This is recommended so that the process of creating the snapshot of one database doesn't interfere with any operations on another database.
- Transaction logs should always be configured in a redundant layout. The preferred software layout is RAID 0+1 (mirrored striped) volumes as this provides better read and write performance than RAID 1 (mirrored) alone. The transaction log will generate the most I/O and thus should use the highest performance disks available.
- The preferred layout for the database is hardware RAID 5, software RAID 1 (mirrored with logging enabled) or software RAID 0+1 (mirrored striped).

Note: FlashSnap is not supported for software RAID 5 volumes.

Reviewing the configuration

A sample setup containing the following object names is used to illustrate how to create, refresh, and restore a snapshot set for a specified SQL Server database.

Table 6-2 Object names used in SQL Server Quick Recovery tasks

Name	Drive Letter	Object
billing		SQL Server instance name
billing_DB		component (SQL Server database name)
harddisk3		Disk where the mirrors will be created
billing_DG		disk group name for the Billing disk group
billing_data	L:	volume for Billing database user data
billing_log	M:	volume for Billing database log
billing_datasnap	S:	volume for Billing database user data snapshot
billing_logsnap	T:	volume for Billing database log snapshot

Table 6-2 Object names used in SQL Server Quick Recovery tasks (Continued)

Name	Drive Letter	Object
billing_DB.xml		name for the metadata file generated when the snapshot is created
tlog1.bak, tlog2.bak		log files saved at c:\backup prior to restoring

Configuring SQL Server storage with Veritas Storage Foundation for Windows

To use the SFW Quick Recovery snapshot functionality with SQL database components, you must store the database components on Veritas Storage Foundation for Windows (SFW) dynamic volumes. Configuring your SQL Server storage with SFW includes the following tasks:

- Creating one or more dynamic disk groups
- Creating volumes for the databases and transaction logs
- Creating a new database and set the appropriate paths to point to the new SFW volumes.

If your SQL Server environment is already configured with SFW, skip this section.

Creating dynamic disk groups

Create one or more dynamic disk groups.

Note: Disk groups must be of a Storage Foundation for Windows 4.0 or later version. You must upgrade any disk groups created using an earlier version of Volume Manager for Windows before implementing SFW snapshot solutions. Quick Recovery snapshots are supported only on volumes belonging to an SFW dynamic disk group. They are not supported on volumes belonging to a Microsoft Disk Management Disk Group.

To create a dynamic disk group from the VEA console

- 1 Click **Start > All Programs > Symantec > Veritas Enterprise Administrator** and if prompted to select a profile, select a profile (or Default).

- 2 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
For the local system you can specify **localhost**.
- 3 If prompted to do so, specify the user name, password, and domain for the system.
- 4 In the tree expand the system name and expand the storage agent.
- 5 Right-click **Disk Groups**, and click **New Dynamic Disk Group**.
- 6 On the **Welcome** screen of the New Dynamic Disk Group Wizard, click **Next**.
- 7 Enter a name for the disk group (for example, `billing_DG`).
- 8 For an off-host or cluster environment, choose from the following:
 - For a cluster environment, check the **Create cluster group** check box.
 - For an off-host environment, check the **Add private group protection** check box.
- 9 Select the appropriate disks in the Available disks list and click the **Add** button to move them to the Selected disks list.
- 10 Click **Next**.
- 11 Click **Next** to upgrade the selected disks.
- 12 Click **Finish** to create the new disk group.

To create a dynamic disk group from the command line

Type the following command:

```
> vxdbg -gbilling_DG init Harddisk1 Harddisk2 Harddisk3
```

where `-gbilling_DG` is the name of the dynamic disk group you want to create and `Harddisk1`, `Harddisk2`, and `Harddisk3` are the disks included in the dynamic disk group.

For the complete syntax for the `vxdbg init` command, see *Veritas Storage Foundation Administrator's Guide*.

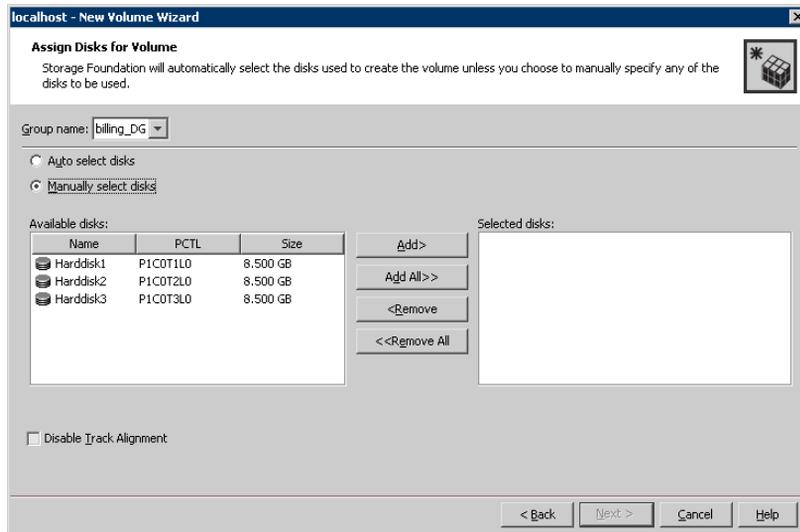
Creating dynamic volumes

Create volumes for the database and transaction log.

To create a dynamic volume from the VEA console

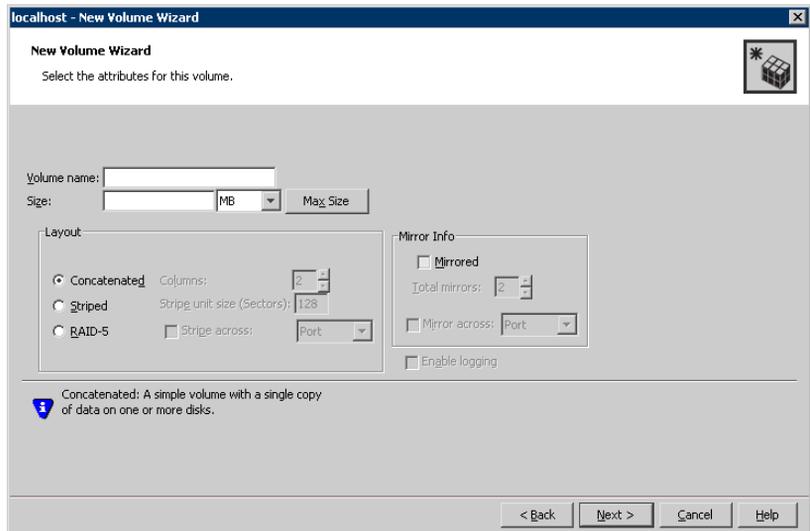
- 1 Start the VEA and connect to the appropriate host.
- 2 In the tree, expand the system name, expand the storage agent, and then expand **Disk Groups**.

- 3 Right-click on the disk group in which to create the volumes (for example, billing_DG), and click **New Volume**.
- 4 In the Welcome panel of the New Volume Wizard, click **Next**.
- 5 Select the disks for the volume:



- Confirm that the **Group name** is the correct disk group (for example, billing_DG). If necessary, select the correct disk group name from the drop-down menu.
- Specify automatic or manual disk selection. Symantec recommends using the **Manually select disks** option.
- Select the appropriate disks in the Available disks list, and click the **Add** button to move them to the Selected disks list.
- You may also check Disable Track Alignment to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

6 Specify the parameters of the volume:



- Enter the volume name (for example, `billing_data`).
 - Enter the size.
 - Select the layout.
 - Select the appropriate mirror options.
 - Click **Next**.
- 7 Assign a drive letter to the volume (for example, L: for the `billing_data` volume) and click **Next**.
- 8 Create an NTFS file system:
- Accept the default **Format this volume**.
 - Click **NTFS**.
 - Select an allocation size.
 - Accept the default file system label, which is the same as the volume name you entered previously or enter a file system label.
 - If desired, select **Perform a quick format**.
 - Click **Next**.
- 9 Review the volume specifications, then click **Finish** to create the new volume.
- 10 Repeat the previous steps as necessary to create volumes for the transaction log (for example, `billing_log`) and any other database in your configuration.

To create a volume from the command line

- 1 Type the CLI commands:

```
> vxassist [-b] -gbilling_DG make billing_data L:  
> vxassist [-b] -gbilling_DG make billing_log M:
```

This will create volumes in the billing_DG disk group named billing_data and billing_log on drive letters L and M respectively.

- 2 Modify and repeat this command as necessary to create additional volumes. For the complete syntax of the vxassist make command, see *Veritas Storage Foundation Administrator's Guide*.

Pointing the databases and log paths to the SFW volumes

Create a new database and point the database and log paths to the SFW volumes.

Note: To use Quick Recovery on an existing user database that is not on SFW volumes, you need to move the database and transaction log to the SFW volumes. Follow the procedures described in Microsoft Knowledge Base Article 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach (<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>)

To create a new SQL Server 2005 database

- 1 Open the SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 Expand the icon associated with your server.
- 3 Right-click on **Databases** and select **New Database**.
- 4 In the New Database page, enter a name for the new database.
- 5 Click the browse button (...) in the **Path** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 6 Select and edit other file properties as desired.
- 7 Click the browse button (...) in the **Path** column for the **Transaction Log** row and browse to the location of the volume you want to create for the transaction log, and click **OK**.
- 8 To add more data files if required, select Add, edit the properties in the new data file rows as required, and click **OK**.

Implementing Quick Recovery for SQL 2005 with the configuration wizard

This chapter covers the following topics:

- [About the Quick Recovery Configuration Wizard](#)
- [Tasks for implementing snapshot sets with the configuration wizard](#)
- [Reviewing the prerequisites](#)
- [Scheduling and creating snapshot sets](#)
- [Viewing the status of scheduled snapshots](#)
- [Troubleshooting scheduled snapshots](#)
- [Deleting or modifying schedules](#)
- [Synchronizing schedules after adding a cluster node](#)

About the Quick Recovery Configuration Wizard

The Quick Recovery Configuration Wizard provides the most complete method of configuring snapshot sets and is therefore recommended for initial implementation. The Quick Recovery Configuration Wizard enables you to schedule all phases of snapshot creation and maintenance:

- Preparing the mirror for the snapshot set.

- Creating the initial snapshot set by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy.
- Periodically refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again.

The wizard enables you to set up and schedule multiple snapshot sets for one or more databases in a SQL Server 2005 instance. You can set up one or more schedules for each snapshot set. You can schedule when to prepare the snapshot mirrors, when to create the initial snapshot sets, and when to refresh the snapshot sets, enabling you to establish a schedule that best suits your site. For example, you can schedule mirror preparation, the most time-consuming step, for a time of low activity. The scheduling capability automates the process of refreshing the snapshot sets. At the time scheduled for the snapshot, the snapshot volumes are automatically reattached, resynchronized, and then split again.

The snapshot creation process integrates with VSS to quiesce the database and then simultaneously snapshot the volumes in the database. The snapshot is done while the database is online and without disrupting the database operations.

Once configured and applied, snapshot schedules are maintained by the Veritas Scheduler Service, which runs in the background.

The wizard includes the following settings:

- Which database to snapshot
- Number of snapshot sets for each database
- Volume layout and disk location for each snapshot set
- When to prepare the selected volumes for the snapshots
- When and how often snapshots occur
You can schedule snapshot run days using days of the week, days of the month, and include or exclude dates. You can specify that a schedule recurs daily or uses some other recurrence interval. You can also schedule a daily time window in which the snapshots occur, and the intervals for refreshing snapshot sets within the time window.
- Backup type (Copy or Full)

Optionally, you can also specify scripts to run before and after snapshots.

After you configure the settings, you can do one or both of the following:

- Save the settings in a template file for reuse on other systems or instances. Even if you do not save the settings to a template file, they are still stored for

the databases to which they apply. You can access them later by selecting the same instance and database in the wizard.

- Implement the settings immediately.

Backup types for snapshot sets

When creating a snapshot set, you can specify that the snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

About snapshot templates

A snapshot template contains all the settings that you implemented with the Quick Recovery Configuration Wizard for selected databases of a particular instance. After configuring settings with the wizard, you can save the template settings to a template (.tpl) file (an XML file) for reuse on another database or instance. This enables you to re-use the same schedule or use it with minor modifications. If you apply the settings without saving, they are stored in the registry so that you can view or modify them in the wizard, but they are not available to import for another instance or database.

Before you can use the template settings on another system, you must first copy the template (.tpl) file to that system. You can then import the file while using the Quick Recovery Configuration Wizard.

If settings already exist for a database, you cannot select that database and import a template to overwrite those settings. You must instead modify the settings. However, you can import a template for another database for which settings do not exist.

Templates store the following information:

- The selected database
- The number of snapshot sets for the database
- When to prepare the snapshot mirrors
- The snapshot schedule and rules related to the schedule, including backup type (but not the names of optional scripts)
- The current date and time when the template is created

Templates and multiple components

When you apply a template that has multiple components, the wizard first attempts to match the names of components (databases) in the template to the databases you selected in the wizard.

If it matches a name, it applies the information stored under that name in the template to that database. For example, if you select the Billing database in the wizard and the template has settings for Billing, it applies those settings.

A database selected in the wizard may not match the names of any databases in the template. In that case, the wizard applies the information for the first unapplied component in the template to the first selected database that does not match any database name in the template. It continues in that sequence.

If you selected more databases in the wizard than the number of components in the template, the wizard prompts you to fill in any required information for the remaining databases.

Templates and schedule start dates

Templates contain settings for one or more schedule start dates, including a mirror preparation date, a "Start on" date, and a "Schedule takes effect on" date.

If you import a template after a start date has elapsed, the wizard tries to maintain the same delta between the template import date and a start date as the delta between the template creation date and a start date. It shifts the dates forward to maintain this delta.

Therefore, after importing the template, use the wizard to review the settings and adjust dates as necessary to meet your needs.

If you import a template that uses the current date as the "takes effect on" date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, a schedule includes two daily snapshots, one in the morning and one at night, and you import the schedule in the afternoon. In this case neither of the snapshots will occur on the current date. Instead the effective date is shifted ahead one day.

See "[Scheduling and creating snapshot sets](#)" on page 95.

Tasks for implementing snapshot sets with the configuration wizard

Table 7-1 outlines the high-level objectives and the tasks to complete each objective.

Table 7-1 Tasks for implementing snapshot sets with the configuration wizard

Objective	Tasks
“Reviewing the prerequisites” on page 94	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Scheduling and creating snapshot sets” on page 95	<ul style="list-style-type: none"> ■ Running the Quick Recovery Configuration Wizard to configure and schedule snapshot sets ■ Optionally, saving the configuration and schedules to a template file ■ Applying the configuration
“Viewing the status of scheduled snapshots” on page 103	<ul style="list-style-type: none"> ■ Viewing the status of scheduled snapshots to determine whether they were successful
“Deleting or modifying schedules” on page 106	<ul style="list-style-type: none"> ■ Deleting schedules that are no longer needed or modifying schedules

Reviewing the prerequisites

Before running the Quick Recovery Configuration Wizard, you should do the following:

- Ensure that your system hardware and software meets the requirements and review the recommendations and best practices.
See [“Reviewing the prerequisites”](#) on page 80.
- Set up SQL Server for use with Storage Foundation for Windows (SFW).
See [“Configuring SQL Server storage with Veritas Storage Foundation for Windows”](#) on page 84.
- Ensure that you have disks with enough space to store the snapshot volumes. Each snapshot set requires the same amount of space as the original volumes.
- Ensure that you are logged in as a domain administrator or as a member of the Domain Admins group.
- If a firewall exists between the wizard and any systems it needs access to, set the firewall to allow both ingoing and outgoing TCP requests on port 7419.
- In a clustered server environment, ensure that the scheduler service is configured with domain administrator privileges.
- To use a previously created template, copy it to the server that you are configuring.
- Ensure that the Microsoft Software Shadow Copy Provider service is running. For a clustered server environment, ensure that the service is running on all systems in the cluster.
- Although not required to run the wizard, the SQL Server VSS Writer service must be started for snapshot mirror preparation and scheduled snapshots to occur.

Scheduling and creating snapshot sets

You schedule and create snapshot sets using the Quick Recovery Configuration Wizard. You can also use the wizard to modify or delete existing schedules. The wizard also creates the snapshot mirrors required for the snapshots.

See “[About the Quick Recovery Configuration Wizard](#)” on page 89.

You should ensure that you meet the prerequisites before you begin using the wizard.

See “[Reviewing the prerequisites](#)” on page 94.

In addition, if you plan to reuse settings, you should be familiar with the information about snapshot templates.

See “[About snapshot templates](#)” on page 91.

To schedule and create snapshot sets

- 1 Start the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**), expand the application on the Solutions tab, expand **Quick Recovery Configuration**, and start the Quick Recovery Configuration Wizard.
- 2 In the Welcome panel, review the information to ensure you meet the requirements and click **Next**.
- 3 In the System Selection panel, specify the fully qualified host name or IP address of the system that is running the application (or specify localhost for the local system), and click **Next**.
 In a clustered server (high availability) environment, Symantec recommends that you specify the virtual server name or virtual IP address. Ensure that the disk groups for the application instance are imported to the selected system.
- 4 In the Instance Selection panel, specify the following options and click **Next**:

- | | |
|--------------------------------|--|
| Set up Quick Recovery for | Select the SQL instance for which you want to configure the snapshot sets.

Once you select an instance, its details are displayed in the panel. |
| Select objects or view details | Select the SQL databases for which you want to configure the snapshot sets.

Databases with names longer than 50 characters are not supported by the wizard. |

The wizard checks for existing schedules for the selected database. The wizard cleans up any obsolete schedule files and entries. In a clustered server environment, the wizard will synchronize schedules between cluster nodes to ensure that all existing schedules are available on the selected node.

- 5 The wizard validates that the volumes containing the schedule configuration files are mounted. If a volume is not mounted, the Mount Details panel displays the information about the missing drive. Mount the missing drive and click **Refresh**, then click **Next**. Otherwise, the schedule is deleted.
- 6 For existing schedules, the Synchronizing Schedules panel displays the status of schedule synchronization and cleanup. If schedule synchronization fails in a clustered environment, restart any cluster nodes that are down and then restart the wizard. Otherwise, click **Next**.
- 7 In the Template Selection panel, select from the following options and click **Next**:

Create or modify a template Create or modify the Quick Recovery settings for the selected components.

If at least one of the selected components has existing settings, this option is the only choice. In this case, the existing settings are displayed for you to view or modify as you continue through the wizard panels.

If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.

Import a template Import Quick Recovery settings from a template (.tpl) file created by the Quick Recovery wizard earlier. Browse to the file location and select the file. The file must be located on the system for which Quick Recovery is being configured. The file also must be of the same application type as the application you are configuring.

This option is available only if no settings exist for a selected component.

- 8 In the Number of Snapshot Sets panel, select how many sets of snapshots to create for each database. Remember that each snapshot set uses the same amount of space as the original volumes.

9 In the Snapshot Volume Assignment panel, expand the snapshot set and make the following selections:

Snapshot set	<p>In the box under the Snapshot Set column heading, optionally edit the snapshot set name. If you edit the name, ensure that the name is unique among all snapshot sets for all databases.</p> <p>If you are modifying an existing schedule, the snapshot set name cannot be changed.</p>
XML Metadata File Name	<p>Specify a name for the XML file that will be associated with the snapshot set. This file is used for recovery, so you may want to assign a name to easily identify it for that purpose. Ensure that the XML file for each snapshot set has a unique name.</p> <p>If you are modifying an existing schedule, the XML file name cannot be changed.</p>
Mirror Preparation Time	<p>Click Edit and in the Mirror Preparation dialog box, specify the date and time for mirror preparation and click OK.</p> <p>Preparing the snapshot mirrors slows down performance, so you may want to choose a time of low usage.</p>
Snapshot Disks	<p>Assign one or more disks to each snapshot volume. Click the icon next to the disk column to select from available disks and to select a concatenated or striped volume layout.</p>
Mount	<p>Optionally, click Set to set a drive letter or mount path.</p> <p>The drive letter specified may not be available when the snapshot operation is performed. When this occurs, the snapshot operation is performed, but no drive letter is assigned.</p>

File path for snapshot XML files

File path for the snapshot XML files. The path specified applies to all the snapshot sets that you are currently creating, for all currently selected databases.

If the field is not editable, you are modifying an existing snapshot set schedule, so the value cannot be changed.

For a non-clustered environment, the default location on Windows Server 2003 is:

```
C:\Documents and Settings\All Users\  
Application Data\VERITAS\VSSXML\  
application name
```

and on Windows Server 2008:

```
SystemDrive:\ProgramData\VERITAS\VSSXML\  
application name
```

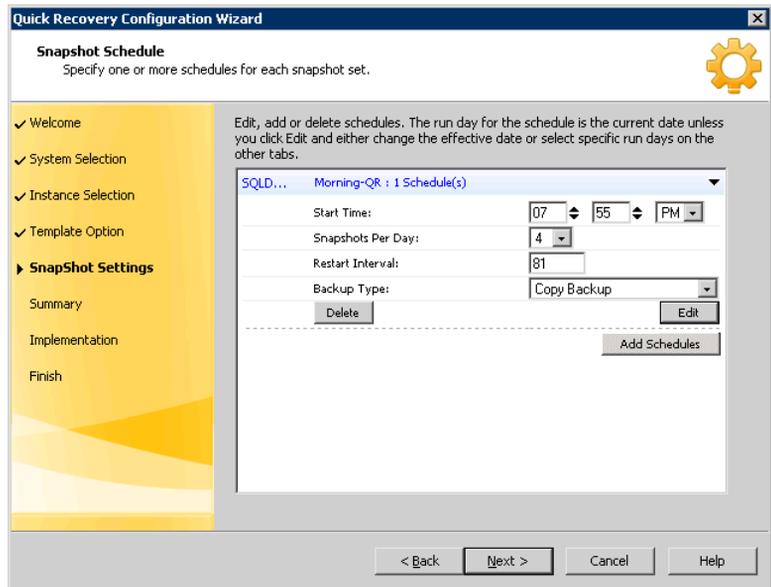
If a `redirect.txt` file has been created, the path in that file is shown in the field.

For new schedules for a clustered server environment, if there is no `redirect.txt` file, the field is empty. Enter the full path to a location on shared storage, for example: `G:\SnapshotSets`.

Symantec recommends storing the XML files on a separate volume from the volumes that are included in snapshots. Ensure that you use a consistent location. That way, you can easily find the XML metadata file, which is required for recovery.

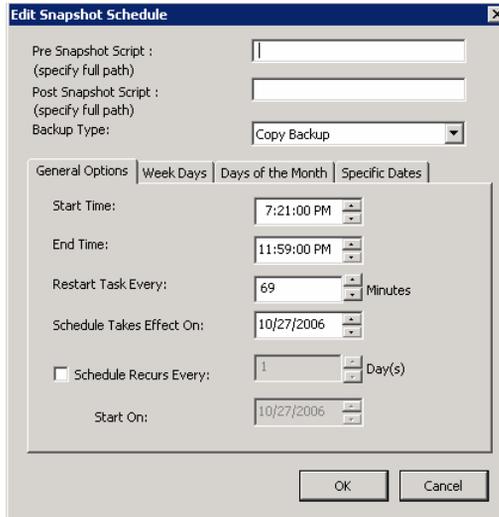
- 10 Repeat [step 9](#) for any additional snapshot sets and when you finish, click **Next**.

11 In the Snapshot Schedule panel, choose one of the following:



- To specify a simple schedule, edit the default schedule settings shown on this panel: the time to start the snapshot (at least 1 hour later than the scheduled mirror preparation time), the number of snapshots to take, the interval in minutes between snapshots, and the type of snapshot. By default, the simple schedule takes effect on the current date and does not recur. Continue with [step 16](#).
- To specify a recurring schedule, a time window, a different run date, or other schedule details, click **Edit**.
Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.

- 12 If you clicked **Edit**, in the Edit Snapshot Schedule dialog box, choose from the following settings and then make additional selections on the dialog box tabs:



Pre Snapshot Script

Optionally, specify the full path of a script to run before the scheduled snapshot occurs.

In a cluster environment, script files should be located on shared storage.

For security purposes, ensure that script files are not accessible to any user with fewer privileges than the user account under whose context the scheduler service is running. Otherwise, a user with fewer privileges might change or replace the file that is launched by the scheduler service.

Post Snapshot Script

Optionally, specify the full path of a script to run after the snapshot is complete.

Backup Type Select a backup type:

You can specify that snapshots be created as either a Full backup or Copy backup type. Any snapshots taken between the time a full backup is taken and the time a log backup is taken can be used as a base to restore and replay the logs. If the replay of logs is not necessary, then either type of snapshot can be used to restore the database to the time of the snapshot set. A Copy backup does not affect the sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

13 In the Edit Snapshot Schedule dialog box, on the General Options tab, you can specify the following:

Start Time	The time of the day to begin taking snapshots
End Time	The time of day to end taking snapshots. If a snapshot is in progress it is completed but a new one is not started after the end time.
Restart task every	The interval between snapshots, in minutes. For example, if the interval is 360 minutes and you schedule a snapshot start time of 12 P.M. and an end time of 7 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.
Schedule takes effect on	The date on which the specified schedule takes effect. If you specify no other run day information on the other tabs, this date is used as the only run date. If you specify run days on the other tabs, the schedule starts running on the earliest day specified. Therefore, if you want the current date included, specify it as an Include date on the Specific Dates tab.
Schedule recurs every	Enable this option to have the snapshot schedule continue to recur. Otherwise the schedule applies only for one day. Specify the number of days before restarting the snapshot schedule. For example, 1 day would mean the schedule takes effect daily, 2 days would mean every other day.
Start On	If you specify the Every option, specify the starting date.

- 14 In the Edit Snapshot Schedule dialog box, optionally make selections on the Week Days, Days of Month, and Specific Dates tabs as follows:

Week Days	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking 1st schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also specify the last day of the month.
Specific Dates	Select one or more specific dates to include in or to exclude from the schedule. Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules for the same snapshot set overlap for the same snapshot, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

- 15 When you are done making selections on the Edit Snapshot Schedule dialog box for this schedule, click **OK**.
- 16 In the Snapshot Schedule panel, choose from the following and when scheduling is complete for all snapshot sets, click **Next**:
- Edit schedules for any remaining snapshot sets
 - Click **Add Schedules** if you want to add a new schedule
 - Click **Delete** if you want to remove a schedule
- 17 In the Summary panel, choose one or both of the following:
- If you want to save the settings to a template file for reuse on other SQL databases, click **Save** and save the template to a file location of your choice.
If you do not save the settings to a file, you can still view or modify them by launching the wizard and selecting the same instance and database.
 - If you are ready to implement the template with its current settings, click **Apply**.

If you click **Apply** without saving the template, you are prompted to confirm.

If you have saved the settings to a template file and want to exit the wizard without applying the template, click **Cancel**.

- 18 In the Template Implementation panel, wait until the wizard shows that Quick Recovery configuration is complete and click **Next**.
- 19 Click **Finish**.

Viewing the status of scheduled snapshots

If a scheduled snapshot fails for some reason, the scheduler process will attempt to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 3 Expand the component for which you scheduled the snapshot.
- 4 Click **Scheduled Tasks**. The scheduled snapshots are listed in the pane on the right.
- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click **Scheduled Tasks** and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the Job History dialog box, view the schedule information. You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Troubleshooting scheduled snapshots

When scheduling snapshots using the Quick Recovery Configuration Wizard or the VSS Snapshot Scheduler Wizard, you can use the information in the following table to help avoid problems or troubleshoot situations.

Situation	Resolution
Snapshots do not occur on a scheduled date.	The date may have been excluded under “Specific Dates” on the schedule. Excluding a date takes precedence over other days scheduled.
A snapshot is not taken on the date that you create the schedule, although that date is specified as the “Schedule takes effect on date”.	The date shown in the Schedule takes effect on field is used as a run day only if no other run days are specified. If you specify other run days, the schedule starts running on the earliest day specified. If you want the current date included in a schedule, specify it as an Include date on the Specific Dates tab.
A scheduled snapshot does not occur for an imported template schedule.	If you import a template that uses the current date as the “takes effect on” date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, if the schedule includes two daily snapshots, one in the morning and one at night, and you import the schedule in the afternoon, neither of the snapshots occur on the current date. Both occur the following day.
While running the Quick Recovery Wizard to modify template settings, the existing settings are not displayed.	If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.
A schedule is unavailable to be edited in the Quick Recovery Wizard.	Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.
You want to use the VSS Scheduler wizard but the VSS Writers node is not shown in the VEA.	You may need to refresh the VEA display to see the node. For SQL, you must manually start the SQL Server VSS Writer Service. Then right-click the Storage Agent node in the VEA and click Refresh.
You are unable to locate a snapshot set XML file.	The VSS Snapshot Scheduler Wizard assigns a prefix of “VM_” to the name you assign.

Situation

Resolution

Drive letters assignments for snapshot volumes do not occur.

If time elapses between when you use the wizard to assign drive letters and when the snapshot operation occurs, a drive letter you assigned may become unavailable. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

You can assign or change drive letters or mount paths in Veritas Enterprise Administrator.

A scheduled snapshot fails after you have done a manual snapback (reattach) of one or more of the snapshot volumes in the snapshot set.

When a snapback is done manually for a volume rather than by using the VSS Snapback wizard or allowing the scheduler to handle it, the XML metadata file is not deleted. To take the next scheduled snapshot, the scheduler attempts to reattach all the volumes specified in the XML metadata file and fails because one or more are already reattached. Therefore, the snapshot also fails.

To ensure that the next scheduled snapshot works correctly, use the VSS Snapback wizard (or the vxsnap utility) to reattach using the XML file.

Two databases are on the same volume. You run the Quick Recovery wizard to schedule the first database. You then run it again to schedule a snapshot for the second database on the same disk. But the mirror preparation for the first database is not yet complete.

You can choose from the following ways to avoid these problems:

- Include both databases in the snapshot when running the wizard the first time.
- Select a different disk for the second snapshot mirror when running the wizard the second time.
- Wait for the mirror preparation for the first snapshot to finish before running the wizard to schedule the second snapshot.

The following problems can occur:

- When running the wizard, the second time, it does not update the available disk space shown for the disk where the first snapshot is scheduled.
- Mirror preparation for the second database fails.

In addition, when troubleshooting, you may want to review the following logs. If a schedule fails, check the scheduler service logs in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager  
5.1\logs\SchedService.log
```

If a snapshot fails, check the VxSnap.log file in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager 5.1\logs
```

Quick Recovery Configuration wizard log files are located in the following paths.

For Windows Server 2003:

```
C:\Documents and Settings\All Users\Application  
Data\VERITAS\winsolutions\log
```

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```

Deleting or modifying schedules

You can delete or modify a schedule that you created with the Quick Recovery Configuration Wizard by running the wizard again and deleting or editing the schedule on the Snapshot Schedule panel.

See “[Scheduling and creating snapshot sets](#)” on page 95.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the VEA console.

Note: The VEA can delete snapshot schedules only; it does not delete mirror preparation scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA does not update template settings created with the Quick Recovery Configuration Wizard.

To delete a schedule from the VEA

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 3 Expand the database component for which you scheduled the snapshot.
- 4 Click **Scheduled Tasks**. The scheduled snapshots are listed in the pane on the right.
- 5 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Synchronizing schedules after adding a cluster node

In a cluster environment, you may add a cluster node after you set up snapshot schedules with the Quick Recovery Configuration Wizard.

In such a case, you can ensure that the schedules are available on the new node by running the Quick Recovery Configuration Wizard again.

To synchronize schedules after adding a node

- 1 Start the Quick Recovery Configuration Wizard from the Solutions Configuration Center.
- 2 Continue through the wizard until the Synchronizing Schedules panel shows that synchronization between cluster nodes is complete.
- 3 Click **Cancel** to exit the wizard.

Scheduling or creating an individual snapshot set for SQL 2005

This chapter covers the following topics:

- [About scheduling or creating an individual snapshot set](#)
- [Tasks to schedule a new snapshot set](#)
- [Tasks to create a one-time snapshot set](#)
- [Reviewing the prerequisites](#)
- [Preparing the snapshot mirrors](#)
- [Scheduling a new snapshot set](#)
- [Creating a one-time snapshot set](#)
- [Refreshing a snapshot set manually](#)

About scheduling or creating an individual snapshot set

Typically you set up your initial snapshot schedules with the Quick Recovery Configuration Wizard.

See [Chapter 7, “Implementing Quick Recovery for SQL 2005 with the configuration wizard”](#) on page 89.

However, you can use the VSS Snapshot Scheduler Wizard instead of the Quick Recovery Configuration Wizard to add a snapshot schedule for all volumes of a selected database. Like the Quick Recovery Configuration Wizard, the scheduler

wizard enables you to automate the refreshing of snapshots according to the schedule that you define.

However, unlike the Quick Recovery Configuration Wizard, the VSS Snapshot Scheduler Wizard does not prepare the snapshot mirrors. You must use the Prepare command to prepare the snapshot mirrors before running the VSS Snapshot Scheduler Wizard for that database. In addition, you can only use the scheduler wizard to schedule one snapshot set for one database at a time.

Note: Adding a snapshot set using the VSS Snapshot Scheduler will not update the template settings created with the Quick Recovery Configuration Wizard. If you want to keep the template settings up to date, you should instead run the Quick Recovery Configuration Wizard to modify the schedule.

At times you may want to create a one-time snapshot of a specific volume or volumes. You can do so using either the vxsnap command line utility or from the Veritas Enterprise Administrator (VEA) console using the VSS Snapshot and Snapback wizards.

If you want to snapshot only a single volume rather than multiple volumes in the database, you can use the VEA Snapshot Volume Wizard rather than the VSS SQL Snapshot Wizard. See the *Storage Foundation for Windows Administrator's Guide*.

Tasks to schedule a new snapshot set

[Table 8-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 8-1 Tasks for scheduling a new snapshot set

Objective	Tasks
“Reviewing the prerequisites” on page 112	<ul style="list-style-type: none"> Verifying hardware and software prerequisites
“Preparing the snapshot mirrors” on page 112	<ul style="list-style-type: none"> Creating snapshot mirrors using the using the VEA Prepare command or the vxsnap utility
“Scheduling a new snapshot set” on page 114	<ul style="list-style-type: none"> Using the VSS Snapshot Scheduler Wizard to create the initial snapshot set and set up the schedule for keeping it refreshed.

Tasks to create a one-time snapshot set

[Table 8-2](#) outlines the high-level objectives and the tasks to complete each objective.

Table 8-2 Tasks for creating a one-time snapshot set

Objective	Tasks
“Reviewing the prerequisites” on page 112	■ Verifying hardware and software prerequisites
“Preparing the snapshot mirrors” on page 112	■ Creating snapshot mirrors using the VEA Prepare command or the vxsnap utility
“Creating a one-time snapshot set” on page 123	■ Creating the one-time snapshot set using the VEA or the vxsnap utility.

Reviewing the prerequisites

Review the following prerequisites:

- Ensure that your system hardware and software meets the requirements. See [“Reviewing the prerequisites”](#) on page 80 in [Chapter 6, “Preparing to implement Quick Recovery for SQL 2005”](#).
- Set up SQL Server for use with Storage Foundation for Windows (SFW). See [“Configuring SQL Server storage with Veritas Storage Foundation for Windows”](#) on page 84 in [Chapter 6, “Preparing to implement Quick Recovery for SQL 2005”](#).
- Ensure that the SQL Server VSS Writer service is started.
- Ensure that you have disks with enough space to store the snapshot volumes. Each snapshot set requires the same amount of space as the original volumes.
- The `vxsnap` commands must be invoked on a local system. On Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode.

Preparing the snapshot mirrors

To prepare the snapshot mirrors in order to create a snapshot from the VEA or from the `vxsnap` command line, you can use either of the following methods:

- The Prepare command from the VEA
You repeat the VEA console Prepare operation for each database volume.
- The `vxsnap prepare` command from the CLI
Use the `vxsnap prepare` command to prepare a mirror for each of the volumes associated with the database.

The snapshot mirrors remain attached to the original volumes and continue to be updated until you use the VSS SQL Snapshot Wizard, the `vxsnap create` command, or the VSS Snapshot Scheduler Wizard to create the snapshot set.

For the snapshot volumes, make sure to select disks or LUNs that are not used for production data. You can create more than one snapshot volume on the same disk as long as there is sufficient space available and as long as the snapshots are of different production volumes.

Also ensure that the SQL Server VSS Writer service is started.

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap** and then **Prepare**.

- 2 In the wizard, click **Next**.
- 3 Choose one of the following.

If the volume is not mirrored Choose **Manually select disks**, use the **Add** and **Remove** buttons to move the desired disks to the Selected disks box, and click **OK**.

If the volume is mirrored and no additional disks are available to create a new mirror Click on an existing plex and click **OK**.

If the volume is mirrored and there are additional disks available on your system Choose either to use an existing mirror for the snapshot or to create a new mirror.

- To create a new mirror, click **Select Disk**, use the **Add** and **Remove** buttons to move the desired disks to the **Selected disks** box, and click **OK**.
- To use an existing mirror, click **Select existing mirror for snap**, select the desired mirror, and click **OK**.

- 4 Repeat [step 1](#) and [step 3](#) to create a snapshot mirror for each volume associated with the database.

Note: Verify that the lower pane of the VEA console indicates that the resynchronization process is complete before creating the snapshot set.

To create the snapshot mirrors using the `vxsnap prepare` command

- ◆ Type the command, as in the following example:

```
> vxsnap prepare component=billing_DB/writer=SQLServerWriter
   source=L:/harddisk=harddisk3
   source=M:/harddisk=harddisk3
```

In this example, snapshot mirrors for the database and log volumes of the `billing_DB` database are created. The mirror for the volume mounted at `L:` is created on disk 3 and the mirror for the volume mounted at `M:` is also created on disk 3.

The complete syntax of the `vxsnap prepare` command is:

```
vxsnap prepare component=ComponentName/writer=WriterName
[server=ServerName][instance=InstanceName][-b]
[source=volume/harddisk=harddisk...]
```

Scheduling a new snapshot set

Before you run the VSS Snapshot Scheduler Wizard to schedule a snapshot set for a database, you must prepare a snapshot mirror for each of the volumes in the database.

See “[Preparing the snapshot mirrors](#)” on page 112.

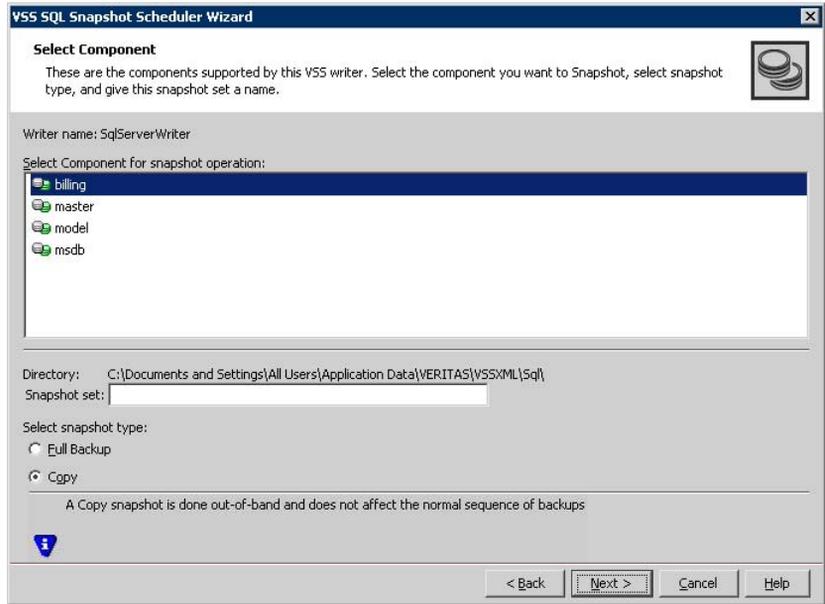
You can then use the VSS Snapshot Scheduler Wizard to schedule the initial snapshot set and to set up the schedule for keeping it refreshed.

When the scheduled snapshots occur, the snapshot mirrors are detached from the original volumes, creating separate on-host snapshot volumes as well as an XML file to store the SQL and snapshot volume metadata. The scheduled process integrates with VSS to quiesce the databases and then simultaneously snapshot the volumes in the database. This snapshot is done while the databases are online and without disrupting the database operations.

To schedule a snapshot for a selected component

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **VSS Writers** node. If the VSS Writers node is not shown, start the SQL Server VSS Writer service and then select the Storage Agent node and refresh the VEA display (**Actions > Refresh**).
- 3 Expand the **SQLServerWriter** node.
- 4 Right-click the instance, and click **Schedule VSS SQL Snapshot**.
- 5 In the Welcome panel, review the information and click **Next**.

6 Specify the snapshot set parameters as follows and then click **Next**:



Select component for snapshot operation

Snapshot set

Select the database for the snapshot set.

Enter a name for the snapshot set, for example, `billing`. The wizard creates the snapshot set metadata XML file with this name, with the prefix “VM_”.

The XML file is stored by default in the directory shown on the screen.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster.

To change the XML file location, use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the full path to the location of the XML file, for example, `G:\BackupSets`. Save the `redirect.txt` file in the default directory `C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML`.

Scheduling a new snapshot set

Select snapshot type

Select the snapshot type.

You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

- 7 In the Change Attributes panel, optionally change the attributes for the snapshot volumes and click **Next**:

Snapshot Volume Label

Displays the read-only label for the snapshot volume.

Drive Letter

Optionally, click a drive letter and select a new choice from the drop-down menu.

The drive letters specified may not be available when the snapshot is taken. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

Plex

Optionally, click a plex and select a new choice from the drop-down menu.

- 8 Optionally, in the Synchronized Snapshot panel (VVR only), select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

See *Veritas Volume Replicator Administrator's Guide*.

- 9 In the Schedule Information panel, on the General Options tab, specify the following:

The screenshot shows the 'VSS SQL Snapshot Scheduler Wizard' window. The title bar reads 'VSS SQL Snapshot Scheduler Wizard'. Below the title bar is the 'Schedule information' section with a sub-header 'Specify Schedule information for snapshots. Times entered on the general tab are times for the host.' There are four tabs: 'General Options' (selected), 'Days of Week', 'Days of Month', and 'Specific Dates'. The 'General Options' tab contains the following fields and options:

- Name of this schedule:
- Description of this schedule:
- Start time: End time:
- Schedule takes effect on:
- Restart task every minutes on each run day.
- Every days. Start on:
- Pre Command: Full path of command to run before snapshot.
- Post Command: Full path of command to run after snapshot.

At the bottom of the window are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

Name of this schedule	Enter a unique name for the snapshot set schedule. This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of the VSS writer name, the component name and a numbered suffix that increments with each schedule.
Description of this schedule	Optionally, enter a description to help you identify the schedule when you view information about the snapshot status.
Start Time	The time of the day to begin taking snapshots.
End Time	The time of day to end taking snapshots. If a snapshot is in progress it is completed but a new one is not started after the end time.
Schedule takes effect on	The date on which the specified schedule takes effect. The default is the current date.

Scheduling a new snapshot set

Restart task every	The interval between snapshots, in minutes. For example, if the interval is 360 minutes and you schedule a snapshot start time of 12 P.M. and an end time of 7 P.M, the snapshot occurs twice. If no interval is specified the snapshot occurs once.
Every	Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day. Specify the number of days before restarting the snapshot schedule. For example, 1 day would mean the schedule takes effect daily, 2 days would mean every other day.
Start On	If you enable the Every option, specify the starting date.
Pre Command	Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.
Post Command	Optionally, specify the full path of a command script to run after the snapshot is complete.

10 To specify run days for the schedule, make selections on the following tabs:

Days of Week	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also check the Last Day checkbox to schedule the snapshot for the last day of each month.
Specific Dates	Select one or more specific dates to include in or to exclude from the schedule. Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

- 11 Click **Next**.
- 12 Review the snapshot set and schedule details and click **Finish**.

Viewing the status of scheduled snapshots

If a scheduled snapshot fails for some reason, the scheduler process will attempt to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 3 Expand the component for which you scheduled the snapshot.
- 4 Click **Scheduled Tasks**. The scheduled snapshots are listed in the pane on the right.
- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click **Scheduled Tasks** and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the Job History dialog box, view the schedule information. You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Troubleshooting scheduled snapshots

When scheduling snapshots using the Quick Recovery Configuration Wizard or the VSS Snapshot Scheduler Wizard, you can use the information in the following table to help avoid problems or troubleshoot situations.

Situation	Resolution
Snapshots do not occur on a scheduled date.	The date may have been excluded under “Specific Dates” on the schedule. Excluding a date takes precedence over other days scheduled.
A snapshot is not taken on the date that you create the schedule, although that date is specified as the “Schedule takes effect on date”.	The date shown in the Schedule takes effect on field is used as a run day only if no other run days are specified. If you specify other run days, the schedule starts running on the earliest day specified. If you want the current date included in a schedule, specify it as an Include date on the Specific Dates tab.
A scheduled snapshot does not occur for an imported template schedule.	If you import a template that uses the current date as the “takes effect on” date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, if the schedule includes two daily snapshots, one in the morning and one at night, and you import the schedule in the afternoon, neither of the snapshots occur on the current date. Both occur the following day.
While running the Quick Recovery Wizard to modify template settings, the existing settings are not displayed.	If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.
A schedule is unavailable to be edited in the Quick Recovery Wizard.	Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.
You want to use the VSS Scheduler wizard but the VSS Writers node is not shown in the VEA.	You may need to refresh the VEA display to see the node. For SQL, you must manually start the SQL Server VSS Writer Service. Then right-click the Storage Agent node in the VEA and click Refresh.
You are unable to locate a snapshot set XML file.	The VSS Snapshot Scheduler Wizard assigns a prefix of “VM_” to the name you assign.

Situation

Drive letters assignments for snapshot volumes do not occur.

A scheduled snapshot fails after you have done a manual snapback (reattach) of one or more of the snapshot volumes in the snapshot set.

Two databases are on the same volume. You run the Quick Recovery wizard to schedule the first database. You then run it again to schedule a snapshot for the second database on the same disk. But the mirror preparation for the first database is not yet complete.

The following problems can occur:

- When running the wizard, the second time, it does not update the available disk space shown for the disk where the first snapshot is scheduled.
- Mirror preparation for the second database fails.

Resolution

If time elapses between when you use the wizard to assign drive letters and when the snapshot operation occurs, a drive letter you assigned may become unavailable. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

You can assign or change drive letters or mount paths in Veritas Enterprise Administrator.

When a snapback is done manually for a volume rather than by using the VSS Snapback wizard or allowing the scheduler to handle it, the XML metadata file is not deleted. To take the next scheduled snapshot, the scheduler attempts to reattach all the volumes specified in the XML metadata file and fails because one or more are already reattached. Therefore, the snapshot also fails.

To ensure that the next scheduled snapshot works correctly, use the VSS Snapback wizard (or the vxsnap utility) to reattach using the XML file.

You can choose from the following ways to avoid these problems:

- Include both databases in the snapshot when running the wizard the first time.
- Select a different disk for the second snapshot mirror when running the wizard the second time.
- Wait for the mirror preparation for the first snapshot to finish before running the wizard to schedule the second snapshot.

In addition, when troubleshooting, you may want to review the following logs.

If a schedule fails, check the scheduler service logs in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager  
5.1\logs\SchedService.log
```

If a snapshot fails, check the VxSnap.log file in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager 5.1\logs
```

Quick Recovery Configuration wizard log files are located in the following paths.

For Windows Server 2003:

```
C:\Documents and Settings\All Users\Application  
Data\VERITAS\winsolutions\log
```

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```

Deleting or modifying schedules

You can delete a schedule from the VEA console. To modify a schedule, run the wizard again and select the same instance and component.

Note: Deleting a snapshot schedule using the VEA does not update template settings created with the Quick Recovery Configuration Wizard.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the VEA console.

Note: The VEA can delete snapshot schedules only; it does not delete mirror preparation scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA does not update template settings created with the Quick Recovery Configuration Wizard.

To delete a schedule from the VEA

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 3 Expand the database component for which you scheduled the snapshot.
- 4 Click **Scheduled Tasks**. The scheduled snapshots are listed in the pane on the right.

- 5 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Creating a one-time snapshot set

Creating a one-time snapshot or snapshot set is a two-step process:

- The first step is to prepare the snapshot mirrors for the database volume or volumes. If you are creating a snapshot set after a snapback to refresh existing snapshot mirrors, you can skip this step. See “[Preparing the snapshot mirrors](#)” on page 112.
- The second step uses either the VSS SQL Snapshot Wizard or the `vxsnap create` command to create the snapshot set by detaching the snapshot mirrors from the original volumes. This step creates separate on-host snapshot volumes as well as a snapshot set XML file to store the application and snapshot volume metadata.

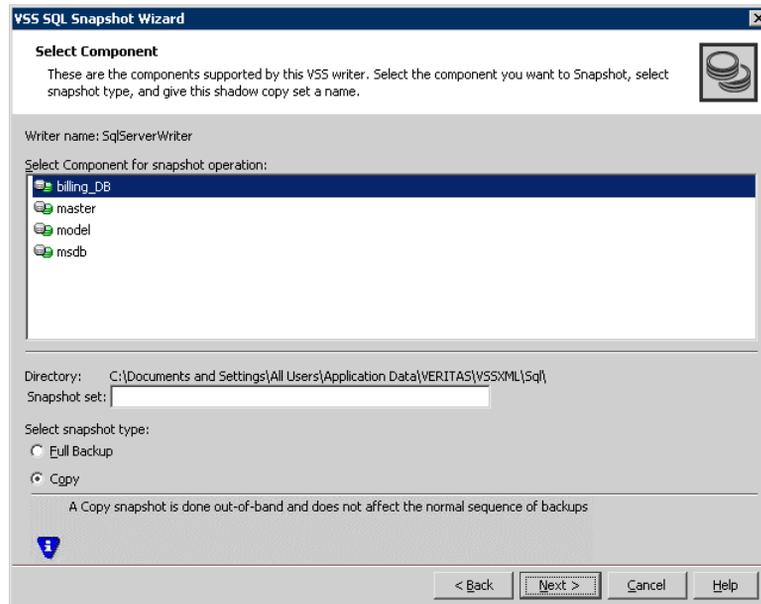
The VSS SQL Snapshot Wizard and `vxsnap create` command integrate with VSS to quiesce the database and then simultaneously snapshot the volumes in the database. This snapshot is done while the database is online and without disrupting the database operations. The resulting snapshot set provides a complete picture of the database at the point in time the command is issued.

The VSS SQL Snapshot Wizard can be run from either a local system or a remote node. The `vxsnap` utility must be run from the local system.

To create the snapshot set from the VEA console

- 1 From the VEA console, navigate to the system where the production volumes and snapshots mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **VSS Writers** node. If the VSS Writers node is not shown, start the SQL Server VSS Writer service and then select the Storage Agent node and refresh the VEA display (**Actions > Refresh**).
- 3 Expand the **SQLServerWriter** node.
- 4 Right-click the instance and click **VSS SQL Snapshot**.
- 5 In the wizard, review the Welcome page and click **Next**.

6 Specify the snapshot set parameters as follows and then click **Next**:



Select Component for snapshot operation

Snapshot set

Select the database for the snapshot set.

Enter a name for the snapshot set, for example, `billing_DB`. The wizard creates the snapshot set metadata XML file with this name. The XML metadata file is stored by default in the directory shown on the screen.

To change the XML file location, use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the full path to the location of the XML file, for example, `G:\BackupSets`. Save the `redirect.txt` file in the default directory `C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML`.

Select snapshot type	Select the snapshot type. You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.
----------------------	--

- 7 In the Change Attributes panel, optionally change the attributes for the snapshot volumes and click **Next**:

Snapshot Volume Label	Displays the read-only label for the snapshot volume.
Drive Letter	Optionally, click a drive letter and select a new choice from the drop-down menu.
Plex	Optionally, click a plex and select a new choice from the drop-down menu.

- 8 On the Synchronized Snapshot panel (VVR only), select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

See *Veritas Volume Replicator Administrator's Guide*.

- 9 Review the specifications of the snapshot set and click **Finish**.

To create the snapshot set from the command line

- ◆ Type the command, as in the following example:

```
> vxsnap -x billing_DB.xml create
  source=L:/Newvol=billing_data
  source=M:/Newvol=billing_log
  writer=SQLServerWriter component=billing_DB
  backuptype=full
```

where `billing_DB.xml` is a name you assign the metadata file that is used to reattach the snapshot set or to recover using that snapshot set.

The complete syntax of the `vxsnap create` command is:

```
vxsnap -x filename create [-gDynamicDiskGroupName]
source=Volume [/DriveLetter=DriveLetter]
[/DrivePath=DrivePath] [/Newvol=NewVolName][/Plex=PlexName]...
writer=WriterName component=ComponentName
[server=ServerName] [instance=InstanceName]
[backuptype=FULL|COPY][-O]
```

The *WriterName* and *ComponentName* are required. The component name is the name of the SQL database. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. You can use either a Full backup or Copy backup for restoring from a snapshot (recovery). However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Refreshing a snapshot set manually

Once a snapshot set has been created, it can be refreshed quickly since the time-consuming step of preparing the mirrors is not required.

Normally, if you want to periodically refresh a snapshot set, you set up the snapshot schedule using the VSS Snapshot Scheduler Wizard or the Quick Recovery Configuration Wizard.

However, if you should need to manually refresh a snapshot set, you can do so. To refresh the snapshot set requires the following tasks:

- [“Reattaching the split-mirror snapshots”](#) on page 127
- [“Creating the refreshed snapshot set”](#) on page 128

Note: The VSS Refresh option available in the VEA console from the VSS Writer object refreshes the display of the VSS Writer and components. It does not refresh the snapshot set.

Reattaching the split-mirror snapshots

The VSS Snapback wizard reattaches and resynchronizes an existing snapshot set so that it matches the current state of its original database. The wizard is available in the context menu of the VSS Writer object.

To snapback a snapshot set

- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the snapshot set.
- 2 From the VEA console URL bar, select the *<host name>* which is the system where the production volumes and snapshot mirrors are located, as the active host.
- 3 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 4 Right-click the writer node of the application and click **VSS Snapback**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set you want to snapback and click **Next**.
The XML metadata file contains all required information needed to snapback the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**. This file is deleted after the snapback operation has completed successfully.
- 7 If a message appears that indicates some volumes have open handles, confirm that all open handles are closed and then click Yes to proceed.
- 8 Verify that the snapback specifications are correct and click **Finish**.

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close the database application GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Type the command, as in the following example:

```
>vxsnap -x billing_DB.xml reattach  
writer=SQLServerWriter
```

where *billing_DB.xml* is the name of the XML metadata file created at the time the snapshot set was created.

The complete syntax for the `vxsnap reattach` command is:

```
vxsnap -x Filename [-f] [-b] reattach writer=WriterName
```

Creating the refreshed snapshot set

After you have reattached and resynchronized the snapshot set mirrors with the original volumes, using the VSS Snapback wizard or from the command line, create a new snapshot set of the database using either the VSS SQL Snapshot Wizard or the `vxsnap create` command.

See “[Creating a one-time snapshot set](#)” on page 123.

Recovering a SQL 2005 database

This chapter covers the following topics:

- [About recovering a SQL 2005 database](#)
- [Tasks for recovering a SQL 2005 database](#)
- [Prerequisites for recovering a SQL 2005 database](#)
- [Types of recovery](#)
- [Recovering using snapshots without log replay](#)
- [Recovering using snapshots and log replay](#)
- [Restoring snapshots and manually applying logs](#)
- [Recovering missing volumes](#)
- [Post-recovery steps](#)
- [Vxsnap restore command reference](#)

About recovering a SQL 2005 database

You can use the on-host Quick Recovery snapshot set to quickly recover a database after logical corruption.

You can also use a snapshot set to recover a database after production volumes are lost due to hardware failure. This recovery assumes that the failure does not affect the disk or disks where the snapshot set volumes are located.

You can use either the VSS SQL Restore Wizard from the Veritas Enterprise Administrator (VEA) console or the `vxsnap restore` command. Both these methods integrate with VSS to notify the SQL VSS Writer to prepare for the

restore before the snapback operation and then to complete post-restore processes afterwards.

After completing the recovery, you refresh the snapshot set.

Note: Refer to the Troubleshooting chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for troubleshooting information.

Tasks for recovering a SQL 2005 database

Table 9-1 outlines the high-level objectives and the tasks to complete each objective. The particular task required depends on the type of recovery you are using.

Table 9-1 Tasks for recovery using Quick Recovery snapshots

Objective	Tasks
“Prerequisites for recovering a SQL 2005 database” on page 131	■ Verifying the prerequisites for using VSS recovery
“Types of recovery” on page 131	■ Understanding the types of recovery
“Recovering using snapshots without log replay” on page 132	■ Recovering to a specified point in time using the VSS SQL Restore wizard or the vxsnap utility
“Recovering using snapshots and log replay” on page 134	■ Recovering to a point of failure using the SQL Restore wizard or the vxsnap utility
“Restoring snapshots and manually applying logs” on page 137	■ Recovering to the time of the snapshot set using the SQL Restore wizard or the vxsnap utility
“Recovering missing volumes” on page 139	■ Recovering missing production volumes from snapshot volumes, using the VSS SQL Restore wizard or the vxsnap utility
“Post-recovery steps” on page 145	■ Refreshing the snapshot set ■ Performing additional tasks in a VVR environment

Prerequisites for recovering a SQL 2005 database

You can use the VSS SQL Restore Wizard or vxsnap command line utility to recover a database from a snapshot set. Both the snapshot set and the snapshot set XML metadata file must be available.

If you are planning to recover missing volumes after hardware failure, additional prerequisites are described in the following topic:

[“Preparing for the recovery”](#) on page 140.

Types of recovery

[Table 9-2](#) gives an overview of the options you can select for recovery. The options are available from either the VSS SQL Restore Wizard or the vxsnap restore command.

Table 9-2 Recovery options

Selected option	Database state after recovery	Description
Recovery	online	Restore to the time of the snapshot set. Database and transaction log volumes are restored from the specified snapshot set database and log volumes. No additional transaction logs are applied.
Recovery, along with the option to restore missing volumes	offline	Restore to the time of the snapshot set. The specified snapshot set volumes are converted from read-only snapshot volumes to read/write volumes. No additional transaction logs are applied. You must then manually assign the drive letter/mount path of the missing production volumes to the previous snapshot volumes and bring the database online.

Table 9-2 Recovery options (Continued)

Selected option	Database state after recovery	Description
Recovery + log replay	online	<p>Restore the snapshot set and automatically roll-forward to the point of failure.</p> <p>Restores the database and transaction log volumes from the specified snapshot set and then automatically replays the specified logs to roll forward the recovery to the point of failure.</p> <p>Before using this option, you must back up your transaction logs within SQL Server. This log operation requires that at least one Full backup was created earlier.</p>
No Recovery	loading	<p>Restore the snapshot set and then manually apply logs in SQL.</p> <p>Restores the database and transaction log volumes from the specified snapshot set and leaves the database in a loading state. To bring the database back to an operational state, you must manually apply your backup transaction logs within SQL Server to the desired point in time.</p> <p>Before using this option, you must back up your transaction logs within SQL Server. This log operation requires that at least one Full backup was created earlier.</p>

Recovering using snapshots without log replay

The following procedure uses the Recovery option without the option to restore missing volumes. It restores the database from the snapshot set volumes to the time of the snapshot set. The database and transaction log volumes are restored but no additional transaction logs are applied.

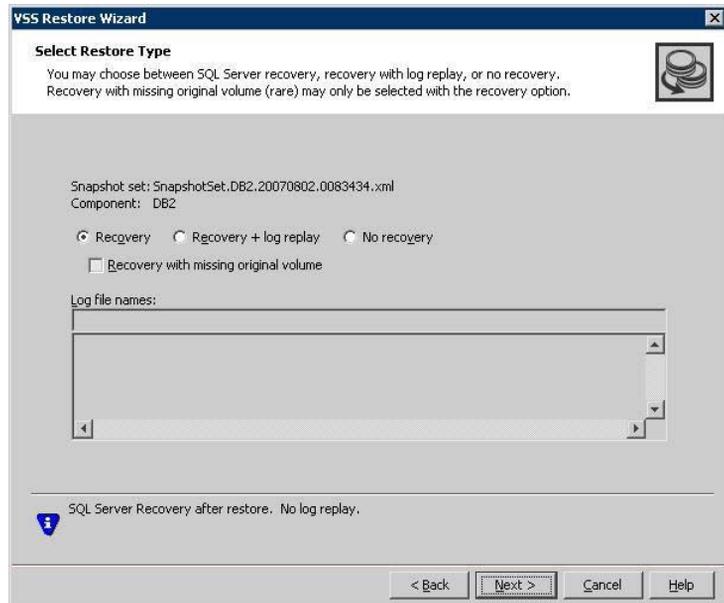
The Recovery option can also be used to recover a database for a missing volume. For the procedure to recover a missing volume, see the following topic: [“Recovering missing volumes”](#) on page 139

To recover a database without log replay using the VEA

- 1 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA), or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 From the VEA console, navigate to the system where the database volumes are located.
- 3 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 4 Right-click **SQLServerWriter** and click **VSS SQL Restore**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all required information needed to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 On the Select Restore Type panel, click **Recovery**.



- 8 You may receive a message “Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed.” Click **No**, close any open handles and retry the command.

- 9 Verify the restore specifications and click **Finish**.
The database is restored to the time the snapshot set was created or last refreshed. If you took it offline earlier, bring it back online.
- 10 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To recover without log replay using the `vxsnap restore` command

- 1 Close the SQL Enterprise Manager GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 Type the command as in the following example:

```
vxsnap -x billing_DB.xml restore RestoreType=RECOVERY noLogs
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command. The volumes in the snapshot set are restored and the database is left in an operational state. The database is restored to the time the snapshot set was created or last refreshed. If you took it offline earlier, bring it back online.
- 3 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Recovering using snapshots and log replay

The following procedure restores the database from the snapshot set volumes and applies the backed up transaction log files that you specify to automatically roll forward the recovery to the point of failure.

You can choose the procedure for the VSS SQL Restore Wizard or for the `vxsnap restore` command.

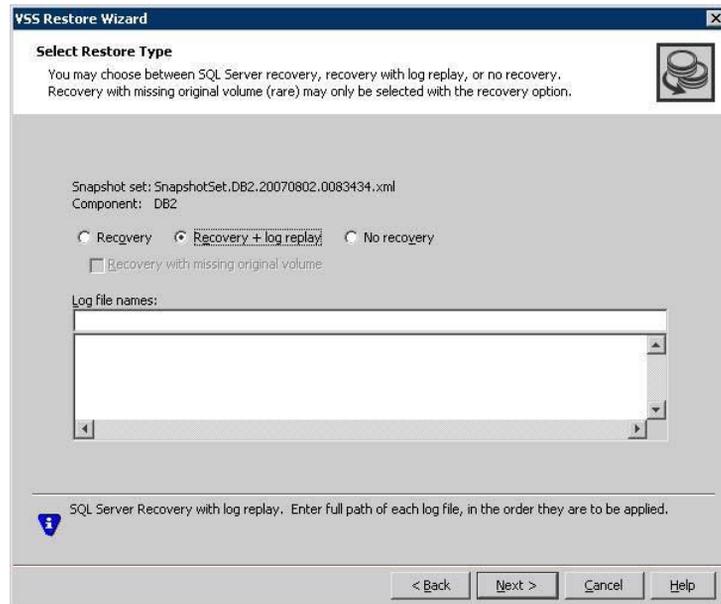
Warning: Before you begin, use your preferred method to backup the transaction logs within SQL Server. You must use the “overwrite existing media” option to create uniquely-named backup files.

To use log replay for an automatic roll-forward recovery to a point of failure using the VEA

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA), or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 3 From the VEA console, navigate to the system where the database volumes are located.
- 4 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 5 Right-click **SQLServerWriter** and click **VSS SQL Restore**.
- 6 Review the Welcome page and click **Next**.
- 7 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all required information needed to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.
- 8 On the Select Restore Type panel, do the following and click Next:
 - Click **Recovery + Log replay**.

- Enter the full path of each log file, in the order they are to be applied.



- 9 You may receive a message “Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed.” Click **No**, close any open handles and retry the command.
- 10 Verify the restore specifications and click **Finish**.
After the most recent backup log is replayed, the SQL Server database is closed and left in an operational state. If you took it offline earlier, bring it back online.
- 11 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To use log replay for an automatic roll-forward recovery to a point of failure using the `vxsnap restore` command

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.

- 3 Type the command as in the following example:

```
vxsnap -x billing_DB.xml restore RestoreType=RECOVERY
      logFiles=c:\backup\tLog1.bak, c:\tLog2.bak
```

 where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command and `c:\backup\tLog1.bak`, `c:\tLog2.bak` are the paths to the transaction log backup files. After the most recent backup log is replayed, the SQL Server database is closed and left in an operational state. If you took it offline earlier, bring it back online.
- 4 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Restoring snapshots and manually applying logs

The following procedure uses the No Recovery option to restore the database from the database and log snapshot volumes. Selecting this option leaves the database in a loading state. You can then manually apply backed up transaction logs to recover the database to the desired point in time.

You can use either the VSS SQL Restore Wizard from the VEA or the `vxsnap restore` command.

Warning: Before you begin, use your preferred method to backup the transaction logs within SQL Server. You must use the “overwrite existing media” option to create uniquely-named backup files.

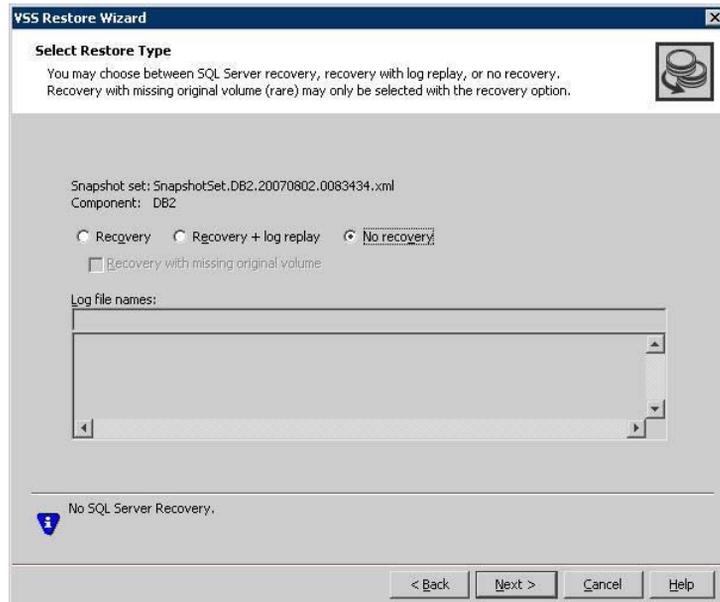
To restore using the No Recovery option in the VEA

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing the database volumes. It is also recommended to bring the database offline.
- 3 From the VEA console, navigate to the system where the database volumes are located.
- 4 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 5 Right-click **SQLServerWriter** and click **VSS SQL Restore**.

- 6 Review the Welcome page and click **Next**.
- 7 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all required information needed to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 8 On the Select Restore Type panel, click **No Recovery** and click **Next**.



- 9 You may receive a message "Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed." Click **No**, close any open handles and retry the command.
- 10 Verify the restore specifications and click **Finish**.
The database and log snapshot volumes are restored and the SQL Server database is left in a loading state.
- 11 Use your preferred method to manually restore the backup transaction logs to the desired point in time and then bring the database back online.
- 12 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To restore using the No Recovery option with the `vxsnap restore` command

- 1 Ensure that you have backed up the transaction logs within SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles or third-party system management tools that may be accessing the database volumes. It is also recommended to bring the database offline.
- 3 Type the command as in the following example:

```
> vxsnap -x billing_DB.xml restore RestoreType=NO_RECOVERY
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command.
The database and log snapshot volumes are restored and the SQL Server database is left in a loading state.
- 4 Use your preferred method to manually restore the backup transaction logs to the desired point-in-time and bring the database online.
- 5 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Recovering missing volumes

Typically missing volumes can result from a hardware failure. If the failure does not affect the disk or disks where the snapshot set resides, you can recover from the failure by restoring the missing volumes from the snapshot set.

Before performing the recovery, you complete several steps to prepare for it.

See “[Preparing for the recovery](#)” on page 140.

You can then recover the missing volumes using the VSS Restore Wizard or the `vxsnap` utility.

See “[Performing the recovery](#)” on page 143.

If the original volume is not present, the snapshot volume is changed from a read-only volume to a read-write volume. No additional transaction logs are applied.

Preparing for the recovery

Before using the VSS Restore Wizard or the vxsnap utility to recover missing production volumes from a snapshot set, you must prepare for the recovery. To prepare, you complete the following steps:

- Take any affected databases offline.
- Reassign the drive letters or mount points of the missing volumes to the snapshot volumes.
See [“Reassigning the drive letter or mount points of the missing volumes to the snapshot volumes”](#) on page 140.
- Replace the failed hardware and add the new disks to the dynamic disk group.
See [“Replacing hardware and adding disks to the dynamic disk group”](#) on page 142.

Reassigning the drive letter or mount points of the missing volumes to the snapshot volumes

If the production volume was missing, change the drive letter or mount point of the snapshot volume to the drive letter or mount point that was assigned to the missing production volume. If the production volume is healthy and available, do not make any changes.

There are two stages to the procedure:

- Note down the assigned drive letter or mount path of the production volume so that you can reassign the same drive letter or mount path to the snapshot volume. Then remove the existing drive letter or mount path from the production volume.
- Change the drive letter or mount path of the snapshot volume to the production volume drive letter or mount path.

The steps for changing a drive letter vary slightly from the steps for changing a mount point. Follow the procedure that best fits your environment.

To remove the drive letter or mount path from the production volume

- 1 Right-click the production volume, click **File System** and click **Change Drive Letter and Path**.
- 2 On the Drive Letter and Paths screen, click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.

To change a snapshot volume drive letter to the production volume drive letter

- 1 Right-click the snapshot volume, click **File System > Change Drive Letter and Path**.
- 2 On the Drive Letter and Paths screen, select **Modify**.
- 3 From the assign drive letter list, select the drive letter originally assigned to the production volume.
- 4 Click **OK**.

To change a snapshot volume mount point to a production volume drive letter

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 On the Drive Letter and Paths screen, click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new drive letter. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Select the drive letter originally assigned to the associated production volume.
- 8 Click **OK**.

To change a snapshot volume mount point to a production volume mount point

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 On the Drive Letter and Paths screen, click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new mount point. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Click **Mount as an empty NTFS folder** and click **Browse**.
- 8 Navigate to the folder where the production volume was mounted.
- 9 Click **OK**.

- 10 Click **OK** to assign the mount point.

Replacing hardware and adding disks to the dynamic disk group

Replace any defective hardware and add new disks to the dynamic disk group, as necessary. The number assigned to a new disk, for example `harddisk5`, may not be the same as the disk number of the failed disk. Note the new disk number(s). You will need the information to add the disks to the dynamic disk group.

To replace the hardware and add the new disks to the dynamic disk group

- 1 Replace the defective hardware.
- 2 In the Actions menu, click **Rescan**.
- 3 If the disk was previously used in another system and has a disk signature, proceed to [step 7](#).
or
If the new disk has never been used before, it is unsigned and needs a disk signature. In this case, the disk appears in the left pane of the VEA console and is marked with (No Signature), for example, `harddisk5 (No signature)`. Proceed to the next step.
- 4 Right-click on a new, unsigned disk and click **Write Signature**.
- 5 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 6 Click **OK**.
After a signature appears on a disk, the disk will display as a basic disk.
- 7 Add the disk to the dynamic disk group of the volumes associated with the production Exchange storage group. Right-click the new disk and click **Add Disk to Dynamic Disk Group**.
- 8 In the Welcome panel, click **Next**.
- 9 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 10 Click **Next**.
- 11 Review the confirmation information and click **Next**.
- 12 Click **Finish** to upgrade the selected disks from basic to dynamic and add them to the dynamic disk group.

Performing the recovery

Ensure that you have prepared for the recovery first.

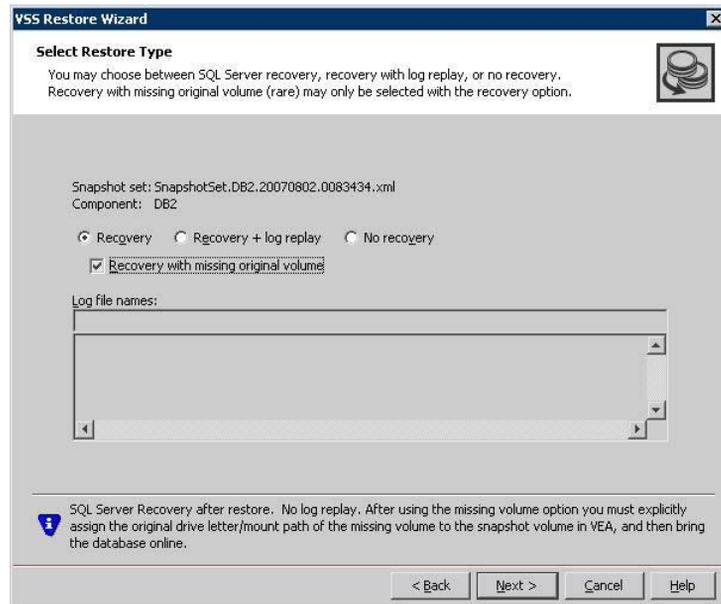
See “[Preparing for the recovery](#)”.

To recover a database with one or more missing volumes using the VEA

- 1 Close the SQL GUI and all Explorer windows, applications, consoles (except the VEA), or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 From the VEA console, navigate to the system where the database volumes are located.
- 3 Expand the system node, the Storage Agent node, and the **VSS Writers** node.
- 4 Right-click **SQLServerWriter** and click **VSS SQL Restore**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all required information needed to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 On the Select Restore Type panel, click **Recovery** and select **Recovery with missing original volume**.



- 8 You may receive a message “Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed.” Click **No**, close any open handles and retry the command.
- 9 Verify the restore specifications and click **Finish**.
The snapshot of the missing volume is changed from a read-only volume to a read-write volume.
- 10 If you have not already done so, in the VEA, ensure that the drive letter or mount path of the missing production volume is assigned to the snapshot volume.
- 11 Bring the database online.
If the production volume was missing, the snapshot volume is now changed to the production volume. The database is restored to the time the snapshot set was created or last refreshed.
- 12 To ensure that another split-mirror snapshot set is immediately available, use the VSS SQL Snapshot Wizard to create a new snapshot of all the volumes in the database.

To recover a database with one or more missing volumes using the `vxsnap restore` command

- 1 Close the SQL Enterprise Manager GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 2 Type the command as in the following example:

```
vxsnap -x billing_DB.xml -r restore RestoreType=RECOVERY noLogs
```

where `billing_DB.xml` is the name of the metadata file generated by the `vxsnap create` command.
The snapshot of the missing volume is changed from a read-only volume to a read-write volume.
- 3 If you have not already done so, in the VEA, reassign the drive letter or mount path of the missing volume to the snapshot volume.
- 4 Bring the database online.
If the production volume was missing, the snapshot volume is now changed to the production volume. The database is restored to the time the snapshot set was created or last refreshed.
- 5 To ensure that another split-mirror snapshot set is immediately available, use the `vxsnap create` command to create a new snapshot of all the volumes in the database.

Post-recovery steps

After you have performed any of the recovery methods, whether point-in-time or roll forward, create a new snapshot set.

See “[Creating a one-time snapshot set](#)” on page 123 in [Chapter 8, “Scheduling or creating an individual snapshot set for SQL 2005”](#)

In a VVR environment, there is an additional post-recovery step. During a point in time recovery in a VVR environment, the volumes on the secondary site lose write-order fidelity. DCM automatically becomes active to ensure data consistency between the primary and secondary sites. While DCM is active, the volumes cannot be expanded by either manual or AutoGrow operations. You must perform a manual resynchronization of the secondary to deactivate DCM.

To resynchronize the node at the secondary site with the node at the primary site:

- 1 Right-click on the primary RVG and click **Resynchronize Secondaries**.
- 2 Click **Yes** to resynchronize the nodes.

Vxsnap restore command reference

The following is a summary of the vxsnap restore command syntax and attribute descriptions.

The command syntax is as follows:

```
vxsnap -x Filename [-b] [-f] [-r] restore  
{RestoreType=[RECOVERY|NO_RECOVERY]}  
[noLogs|logFiles=tlog1,tlog2,...] writer=WriterName
```

The following attributes apply:

-x <i>Filename</i>	The metadata file created by the vxsnap create command. Each snapshot set must have a unique name for the metadata file.
-b	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
-f	Forces the operation. Make sure the volume is not in use before using this option.
-r	Recover even if original volume is not present. If this option is selected and the original volume is not present, the snapshot volume of the missing volume is changed from a read-only volume to a read-write volume. Use this option only with <code>Recovery noLogs</code> . After using this option you must explicitly assign the original drive letter/mount path of the missing volume to the snapshot volume in the VEA and then bring the database online.
RestoreType= [RECOVERY NO_RECOVERY]	Specifies the type of database recovery, either recovery or no recovery: <code>RECOVERY</code> can be used with either the <code>noLogs</code> or <code>logFiles=<i>tlog1,tlog2,...</i></code> attributes. <code>NO_RECOVERY</code> restores from the specified snapshot set to the time of the snapshot. No logs are applied and the database is left in an loading state so that you can manually replay backup logs to a specific point in time.

<i>noLogs</i>	Database and transaction log files are restored from the snapshot set. No transaction backup logs are applied. The database is left in an operational state.
<i>logFiles=tlog1,tlog2,...</i>	Transaction log backup files to be applied with the RECOVERY option to achieve a point of failure recovery. Each transaction log must have a unique name and be created using the “overwrite existing media” option within SQL Server. The database is left in an online state.
<i>writer=WriterName</i>	The name for the SQL Server VSS Writer; used to located the default directory to search for the XML metadata file. Specify SQLServerWriter .

Vxsnap utility command line reference for SQL 2005

This chapter covers the following topics:

- [About the Vxsnap utility](#)
- [Vxsnap keywords](#)

About the Vxsnap utility

The command line utilities are available in the Veritas Storage Foundation for Windows installation directory.

The vxsnap utility integrates with Windows Volume Shadow Copy Service (VSS) to create split-mirror snapshots of all the volumes in the specified component.

Note the following requirements:

- Disk groups must be of a Storage Foundation for Windows 4.0 or later version. You must upgrade any disk groups created using an earlier version of Volume Manager for Windows before using the vxsnap utility
- The CLI commands run only on the server. They will not run on the Veritas Storage Foundation for Windows client.
- The vxsnap commands must be invoked on a local system.
- For Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode.

- vxsnap is used for SQL 2005; SQL 2000 uses a different version called vxsnapsql.

Vxsnap keywords

The vxsnap utility has the following keywords:

prepare	Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes. Note: Either the <code>prepare</code> or <code>start</code> keyword may be used in the CLI; however <code>prepare</code> is recommended.
create	Creates simultaneous snapshots of all volumes in the specified component, providing a point-in-time snapshot set.
reattach	Reattaches and resynchronizes an existing snapshot set to the original database volumes.
restore	Restores a SQL database from a snapshot set.

You can view an online description of the command syntax by typing the following:

```
vxsnap keyword -?
```

vxsnap prepare

Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes.

Syntax

```
vxsnap prepare component=ComponentName/writer=WriterName  
[server=ServerName][instance=InstanceName][-b]  
source=volume/harddisk=harddisk...
```

Attributes

The following attributes apply:

<code>component=ComponentName</code>	Name of the database. The command prepares mirrors for both the database and log volumes of the SQL database.
<code>writer=WriterName</code>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer. Required if you specify the component.
<code>server=ServerName</code>	SQL Server server name.
<code>instance=InstanceName</code>	SQL Server instance name.
<code>-b</code>	Resynchronizes the volume in the background. A snapshot cannot be made until the resynchronization is complete.
<code>source=Volume</code>	Indicates the source volume for the snapshot mirror specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\".
<code>harddisk=Harddisk</code>	Name of the disk where the snapshot mirror is prepared, for example, harddisk2.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Example

```
vxprepare component=billing_DB/writer=SQLServerWriter  
source=L:/harddisk=harddisk3  
source=M:/harddisk=harddisk3
```

This command will create snapshot mirrors of all the volumes contained in the `billing_DB` database. The snapshot mirror of the volume mounted on L: will be created on harddisk 3 and the snapshot mirror of the volume mounted on M: will also be created on harddisk 3.

vxsnap create

Creates snapshot(s) of the specified volume(s) or SQL database (component).
Allows volumes to be snapshotted simultaneously.

Separate source volumes and attributes with forward slashes, not spaces. Source and snapshot volume attributes are paired. You must specify the source volume if you choose to specify the snapshot volume plex, drive letter, drive path, label, or volume name.

Syntax

```
vxsnap -x Filename create [-gDynamicDiskGroupName]
source=Volume [/DriveLetter=DriveLetter]
[/DrivePath=DrivePath] [/Newvol=NewVolName]
[/Plex=PlexName]...
[writer=WriterName] [component=ComponentName]
[server=ServerName] [instance=InstanceName]
[backuptype=FULL|COPY] [-o]
```

Attributes

The following attributes apply:

<code>-x <i>Filename</i></code>	<p>Indicates the name to be assigned to the XML metadata file that will be created with the command. The file name must include the ".xml" extension.</p> <p>The default path to the file is under the SFW program files directory (normally C:\Documents and Settings\All Users\Application Data\Veritas\VSSXML\SQL). If you wish to place the file in another directory, specify a full path before the file name, for example J:\XML\Image1.xml.</p>
<code>-g<i>DynamicDiskGroupName</i></code>	<p>Indicates an SFW dynamic disk group name. If any volume is specified using the SFW volume name, this option is required.</p>
<code>source=<i>Volume</i></code>	<p>Indicates the source volume for the split-mirror snapshot specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\". Repeat this parameter for each volume associated with the specified component.</p>

<code>[/plex=<i>PlexName</i>]</code>	Specifies the name of the mirror or plex that is to be detached. Use this parameter if there are multiple snap plexes available to be snapshotted.
<code>[/DriveLetter=<i>DriveLetter</i>]</code>	The drive letter to be assigned to the new snapshot volume.
<code>[/DrivePath=<i>DrivePath</i>]</code>	The drive path to be assigned to the new snapshot volume. The drive path must reference an empty local NTFS folder, which was created beforehand. The path must include the drive letter and folder to be mounted, for example, C:\DB1VOL.
<code>[/Newvol=<i>NewVolName</i>]</code>	Specifies the name of the new snapshot volume that is created. If the name is not specified using this option, the form "SnapVolume01" is created. The full device path becomes: \Device\HarddiskDmVolumes\ DiskGroupName\ <i>NewVolName</i>
<code>writer=<i>WriterName</i></code>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer. Required if you specify the component.
<code>component=<i>ComponentName</i></code>	Name of the database. The command prepares mirrors for both the database and log volumes of the SQL database.
<code>server=<i>ServerName</i></code>	SQL Server server name.
<code>instance=<i>InstanceName</i></code>	SQL Server instance name.
<code>backuptype=<i>FULL</i> <i>COPY</i></code>	Specifies the type of backup, either a Full or Copy. If no option is specified then Copy is the default. To back up logs in SQL Server so that you can restore the database using SQL log replay, at least one Full backup must have been created earlier.

- o Allows an existing XML file of the same name to be overwritten. If -O is not specified the vxsnap create command does not overwrite an existing XML file of the same name and the operation fails.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Examples

```
vxsnap -x billing_DB.xml create
source=G:/Newvol=billing_data
source=H:/Newvol=billing_log
writer=SQLServerWriter component=billing_DB
backuptype=full
```

This example creates a snapshot set based on the component billing_DB which contains volume G, the database volume, and volume H, the log volume. The snapshot volumes are named billing_data and billing_log, respectively. The XML file, billing_DB.xml, is used to store the VSS metadata that identifies the snapshot set. This file is used in the command to reattach the snapshot set to resynchronize it and in the command to recover the database using the snapshot set.

vxsnap reattach

Reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes.

Syntax

```
vxsnap -x Filename [-f] [-b] reattach writer=WriterName
```

Attributes

The following attributes apply:

- x *Filename* The file created by the vxsnap create command. Each snapshot set must have a unique name for the metadata file.

Note: This file is deleted after the reattach operation has completed successfully.

<code>-f</code>	Forces the reattach. Make sure the volume is not in use by another application before using this command. Use this option with care.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>writer=WriterName</code>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer.

Note: Make sure that the snapshot volumes are not in use before using this command.

Example

```
vxsnap -x billing_DB.xml reattach  
writer=SQLServerWriter
```

This command uses the information in the `snapdata.xml` file to reattach and resynchronize all the volumes in the snapshot set. This `xml` file is deleted after the reattach operation has completed successfully. The snapshot volumes remain synchronized with the original volumes until the `vxsnap create` command is issued.

Note: Any text string that contains spaces must be enclosed in quotation marks.

vxsnap restore

Uses the snapshot volumes in a snapshot set created by the `vxsnap create` command to recover a corrupted or missing SQL Server database.

Exclusive access to the SQL Server database is required for this operation.

Before using this command verify that the source volumes and the snapshot volumes are not in use.

Syntax

```
vxsnap -x Filename [-b] [-f] [-r] restore  
{RestoreType=[RECOVERY|NO_RECOVERY]}  
{noLogs|logFiles=tlog1,tlog2,...} writer=WriterName
```

Attributes

The following attributes apply:

<code>-x <i>Filename</i></code>	The metadata file created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>-f</code>	Forces the operation. Make sure the volume is not in use before using this option.
<code>-r</code>	Recover even if original volume is not present. If this option is selected and the original volume is not present, the snapshot volume of the missing volume is changed from a read-only volume to a read-write volume. Use this option only with <code>Recovery noLogs</code> . After using this option you must explicitly assign the original drive letter/mount path of the missing volume to the snapshot volume in the VEA and then bring the database online.
<code>RestoreType=</code> <code>[RECOVERY NO_RECOVERY]</code>	Specifies the type of database recovery, either recovery or no recovery: <code>RECOVERY</code> can be used with either the <code>noLogs</code> or <code>logFiles=<i>tlog1,tlog2,...</i></code> attributes. <code>RECOVERY</code> leaves the database in an online state. To back up logs so that you can restore the database using SQL log replay, at least one Full backup must have been created earlier. <code>NO_RECOVERY</code> restores from the specified snapshot set to the time of the snapshot. No logs are applied and the database is left in an loading state so that you can manually replay backup logs to a specific point in time.

<i>noLogs</i>	Database and transaction log files are restored from the snapshot set. No transaction backup logs are applied. The database is left in an operational state.
<i>logFiles=tlog1,tlog2,...</i>	Transaction log backup files to be applied with the RECOVERY option to achieve a point of failure recovery and leave the database in an online state. Each transaction log must have a unique name and be created using the “overwrite existing media” option within SQL Server.
<i>writer=WriterName</i>	The name for the SQL Server VSS Writer; used to located the default directory to search for the XML metadata file. Specify SQLServerWriter .

Examples

Following are examples of the main types of restore operation:

- Recovering using snapshots without log replay

```
vxsnapl -x TestDB.xml restore RestoreType=RECOVERY  
noLogs
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and brings the database online. The database is restored to the time the snapshot set was created or last refreshed.

You can use the **-r** option with the **RECOVERY noLogs** restore type if a production volume is missing due to hardware failure:

```
vxsnap -x TestDB.xml -r restore RestoreType=RECOVERY  
noLogs
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set. Any missing volume is changed from a read-only volume to a read-write volume. After using the **-r** option you must explicitly assign the original drive letter/mount path of the missing production volume to the snapshot volume in the VEA. You then bring the database online.

- Recovering using snapshots and log replay

```
vxsnap -x TestDB.xml restore RestoreType=RECOVERY  
logFiles=c:\backup\tLog1.bak, c:\tLog2.bak
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and then applies the specified transaction log

backups (c:\backup\tLog1.bak and c:\tLog2.bak) and brings the database online.

- Restoring snapshots and manually applying logs

```
vxsnap -x TestDB.xml restore RestoreType=NO_RECOVERY
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and leaves the database in a loading state so that backup logs can be manually restored to a specific point in time.

Microsoft clustering solutions

This section includes the following:

- [Deploying SFW with MSCS: New SQL 2000 installation](#)
- [Deploying SFW with MSCS: New SQL 2005 installation](#)
- [Deploying SFW with Microsoft failover clustering: New SQL 2005 installation](#)
- [Deploying SFW with MSCS and SQL Server in a campus cluster](#)
- [Deploying SFW with Microsoft failover clustering and SQL Server in a campus cluster](#)
- [Deploying SFW and VVR with MSCS: New SQL 2000 installation](#)
- [Deploying SFW and VVR with MSCS: New SQL 2005 installation](#)
- [Deploying SFW and VVR with Microsoft failover clustering: New SQL 2005 installation](#)

About Microsoft clustering solutions

Microsoft clustering may be used with Veritas Storage Foundation for Windows to provide high availability for SQL Server. Microsoft clustering may be used with Veritas Storage Foundation for Windows and Veritas Volume Replicator to provide replication support for SQL Server. Using VVR with Microsoft Clustering provides a replicated backup of your SQL Server data, which can be used for recovery after an outage or disaster. However, this solution does not provide the automated failover capability for disaster recovery that can be achieved using VVR with VCS.

Microsoft clustering solutions are covered in separate chapters according to operating system:

- Microsoft Cluster Server (MSCS) on Windows Server 2003 (SQL 2000 and SQL 2005)
- Microsoft failover clustering on Windows Server 2008 (SQL 2005 only)

Deploying SFW with MSCS: New SQL 2000 installation

This chapter covers the following topics:

- [Tasks for a new SQL Server 2000 installation with SFW and MSCS \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing an MSCS cluster](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Creating the SQL virtual server group](#)
- [Creating the MSDTC resource](#)
- [Installing SQL Server 2000](#)
- [Implementing a dynamic quorum resource](#)
- [Testing the cluster](#)

Tasks for a new SQL Server 2000 installation with SFW and MSCS (Windows Server 2003)

You can install and configure Storage Foundation for Windows with MSCS and SQL Server 2000 for a new installation on Windows Server 2003. This environment involves an active/passive configuration with one to one failover capability for high availability.

If you will use Veritas Volume Replicator and replication, see “[Deploying SFW and VVR with MSCS: New SQL 2000 installation](#)” on page 373.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 11-1 Tasks for deploying SFW with MSCS for SQL Server 2000

Objective	Tasks
“ Reviewing the requirements ” on page 164	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 166	<ul style="list-style-type: none"> ■ Understanding a typical active/passive SQL configuration in a two-node cluster ■ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 169	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“ Establishing an MSCS cluster ” on page 171	<ul style="list-style-type: none"> ■ Reviewing general guidelines to establish an MSCS cluster
“ Installing SFW with MSCS/Failover Cluster option ” on page 172	<ul style="list-style-type: none"> ■ Modifying the driver signing options for Windows 2003 remote systems ■ Installing Veritas Storage Foundation for Windows (automatic installation) ■ Restoring the driver signing options for Windows 2003 remote systems
“ Configuring SFW disk groups and volumes ” on page 181	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes
“ Managing disk group and volumes ” on page 189	<ul style="list-style-type: none"> ■ Importing and deporting from cluster nodes
“ Creating the SQL virtual server group ” on page 191	<ul style="list-style-type: none"> ■ Creating a SQL Server cluster group ■ Creating the virtual server IP address ■ Creating the disk group resource
“ Creating the MSDTC resource ” on page 193	<ul style="list-style-type: none"> ■ Creating the MSDTC resource for SQL Server
“ Installing SQL Server 2000 ” on page 194	<ul style="list-style-type: none"> ■ Installing SQL Server and any required patches ■ Verifying SQL installation

Table 11-1 Tasks for deploying SFW with MSCS for SQL Server 2000

Objective	Tasks
<p>“Implementing a dynamic quorum resource” on page 196</p>	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the cluster group ■ Changing the quorum resource to a dynamic mirrored quorum resource.
<p>“Testing the cluster” on page 199</p>	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation.

Requirements for Veritas Storage Foundation for Windows

Before you install SFW, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List located at the following website to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>.

Supported software for MSCS and SFW

- Veritas Storage Foundation 5.1 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
- Microsoft SQL 2000 servers and their operating systems:

Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)

- Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
- Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

Microsoft SQL Server 2000 (64-bit) Enterprise Edition

- Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)

Microsoft SQL Server 2000 (64-bit) Standard Edition or Enterprise Edition (SP4 required)

- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions)
- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions)

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 11-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

- One CD-ROM drive accessible to the system on which you are installing SFW.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- MSCS requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft SQL Server documentation for instructions on creating a reverse lookup zone.
- MSCS requires two disks for SQL: one for SQL database files and one for SQL log files.
- Each system requires 1 GB of RAM for SFW.
- SFW requires administrator privileges to install the software.

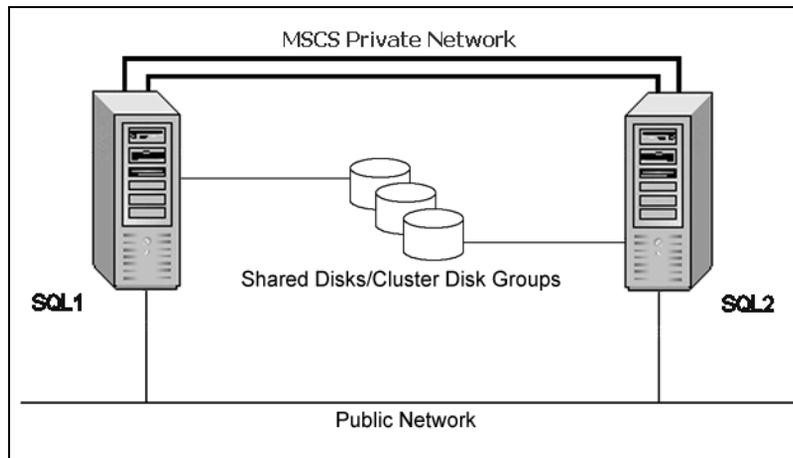
Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Reviewing the configuration

You can create a virtual server in an active/passive SQL Server configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node.

The following figure illustrates a typical active/passive configuration. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured on the active node (SYSTEM1). If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 11-1 Active/passive configuration



Some key points about the configuration:

- An MSCS cluster must be running before you can install SFW.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- SFW enables you to create a dynamic mirrored quorum. The quorum resource maintains the cluster database and critical recovery information in a recovery log.
 In an MSCS cluster without SFW, the quorum disk is a point of failure because MSCS only supports a basic physical disk and does not enable you to mirror the quorum resource.
 The main advantage of SFW is that it provides a dynamic mirrored quorum resource for MSCS. If the quorum resource fails, the mirror takes over for the resource. In this configuration, Symantec recommends creating a three-way mirror for the quorum to provide additional fault tolerance. If possible, do not use the disks assigned to the quorum for any other purpose. You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that SQL is working in the cluster before adding the dynamic quorum volume.
- SFW enables you to add fault-tolerance to data volumes. Symantec recommends mirroring log volumes and a mirrored striped RAID layout for data volumes. SFW offers multiple disk groups, mirrors, capacity management and automatic volume growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, dynamic multi-pathing, and enhanced snapshot capabilities with FlashSnap. Refer to the *Veritas Storage Foundation Administrator's Guide* for details on these features.

Sample configuration

A sample setup is used through this guide to illustrate the installation and configuration tasks.

During the configuration process you will create virtual IP addresses for the following:

- Cluster IP address: used by MSCS cluster
- SQL virtual server: the IP address should be the same on all nodes

You should have these IP addresses available before you start deploying your environment.

The following names describe the objects created and used during the installation and configuration:

Name	Object
SYSTEM1 & SYSTEM2	server names

Reviewing the configuration

SQL_GROUP	Microsoft SQL Server resource group
SQLCLUST	Microsoft SQL Server virtual cluster (underscores not supported)
SQLVS	Microsoft SQL Server virtual server
INST1	Microsoft SQL Server instance name
INST1_DG	disk group for Microsoft SQL volumes
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
SQLVS_QRM	volume for storing the MSCS cluster quorum
QUORUM_DG	quorum volume disk group

Configuring the storage hardware and network

Use the following procedures to configure the storage hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent MSCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 5 In the Public Status dialog box, on the General tab, click **Properties**.
- 6 In the Public Properties dialog box, on the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.

- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Establishing an MSCS cluster

Before installing SFW, you must establish an MSCS cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To establish an MSCS cluster (general guidelines)

- 1 Verify that the quorum disk has been created before installing MSCS and configuring a cluster. (For IA64 systems, the quorum must be created using MBR instead of GPT or it will not be visible.)
- 2 Configure the shared storage and create a partition with drive letter “Q” for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster (SYSTEM1) using MSCS Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Verify that the node can access the shared storage.
- 4 Connect the second node to the shared storage.
- 5 Add the second node (SYSTEM2) using Cluster Administrator on that system.
- 6 Test the cluster by using the Move Group command to move the cluster resources to the second node.
SYSTEM2 becomes the active cluster node.

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running an MSCS cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the MSCS cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 172.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 174.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 180.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
- Moving the online groups

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The table below describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 11-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a MSCS configuration.

To install the product

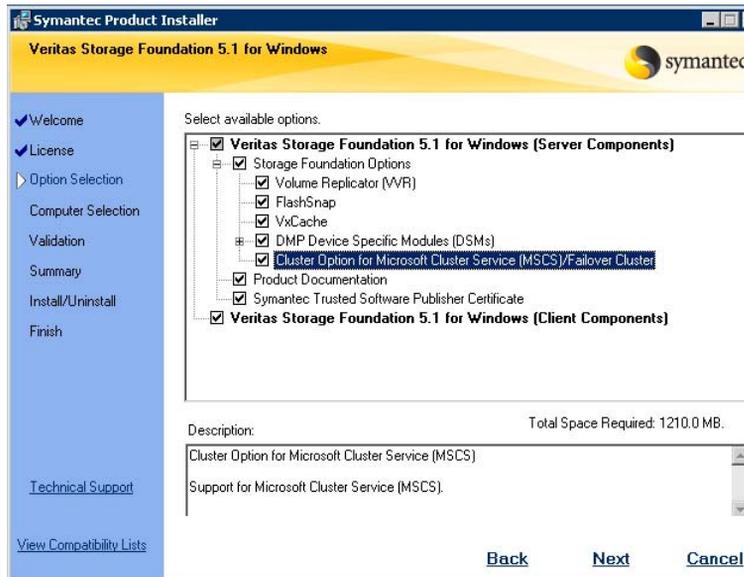
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

- 3 Click **Storage Foundation 5.1 for Windows**.



- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for "**I accept the terms of the license agreement**," and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
- 8 Repeat for additional license keys.
To remove a license key, click the key to select it and click **Remove**.
To see the license key's details, click the key.
- 9 Click **Next**.

- 10 Specify the product options by selecting the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.

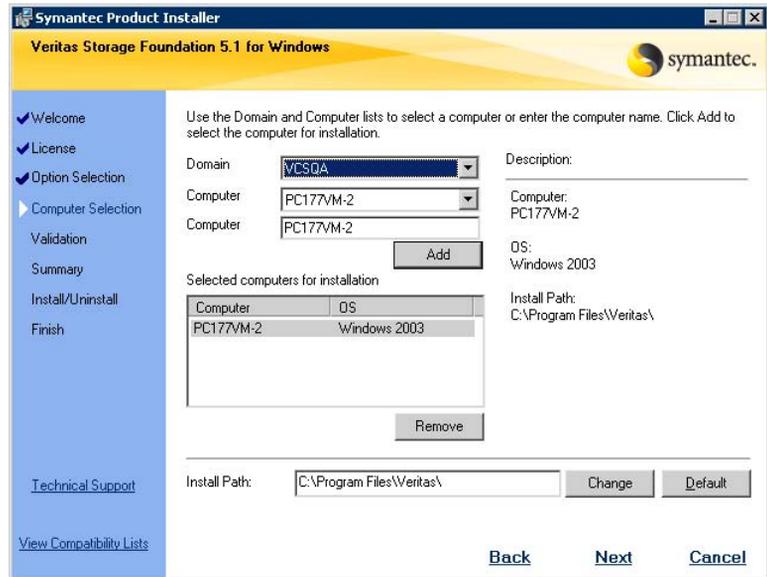


Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.

Note that under Veritas Dynamic Multi-pathing, you can select DMP Device Specific Modules (DSMs).

- 11 Click **Next**.
- 12 Verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is checked, to install the client component and click **Next**.

13 Select the domain and the PC computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 14 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 15 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 16 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that Microsoft clustering allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

Dynamic Multi-pathing

Additionally, if you selected the Dynamic Multi-pathing option, a warning appears:

- For DMP installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.

- For DMP DSM installations—the time required to install the Veritas Dynamic Multi-pathing DSM feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 17 Review the information and click **Install**. Click **Back** to make changes.
 - 18 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.
If the installation is successful on all systems, the installation report screen appears.
If a security alert asks you to accept the Veritas driver software, click **Yes**. This alert appears if your local computer has its driver signing options set to Warn. If your local computer has its driver signing options set to Block, installation fails.
 - 19 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 20](#) if you are installing SFW on the local node only.
 - Proceed to [step 22](#) if you are installing SFW on local and remote systems.
 - 20 To complete the installation, click **Finish**.
 - 21 Click **Yes** to reboot the system and complete the installation.
 - 22 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
 - 23 Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
 - 24 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available.
 - 25 Click **Next**.
 - 26 Click **Finish**.
 - 27 Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW Installation
- Resetting the driver signing options

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.
- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 172.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for SQL. A dynamic disk group is a collection of one or more disks that behaves as a single storage repository. Within each disk group, you can have dynamic volumes with different layouts.

Configuring disk groups and volumes involves the following tasks:

- “[Planning disk groups and volumes](#)” on page 181
- “[Creating dynamic cluster disk groups](#)” on page 183
- “[Creating dynamic volumes](#)” on page 185
- “[Managing disk group and volumes](#)” on page 189

Planning disk groups and volumes

Before installing SQL, you create disk groups and volumes using the VEA console installed with SFW.

A dynamic cluster disk group is a collection of one or more disks that behave as a single storage repository and which can potentially be accessed by different computers. Within each disk group, you can have dynamic volumes with different layouts. Before creating a disk group, consider:

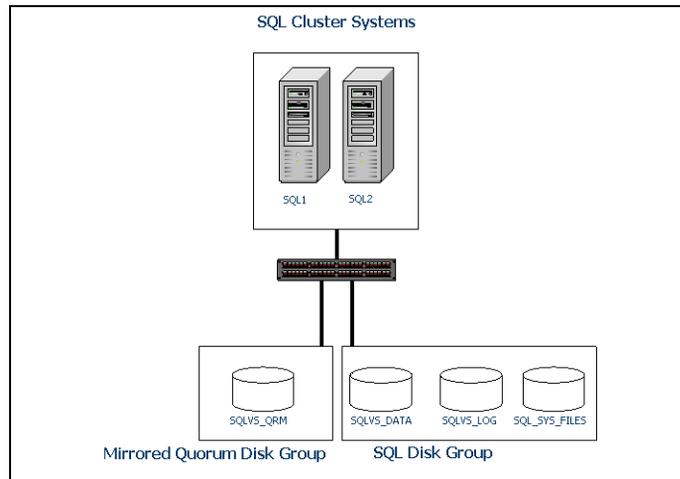
- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.
- The number of disk groups and volumes that are needed for SQL Server
The number of disk groups for SQL depends on the planned organization of the data. The application program files need to be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage in a cluster disk group. You create at least one disk group for the system data files. You may want to create additional disk groups for user databases. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- The disk groups and volumes for the mirrored quorum resource
You will need a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. You can create the quorum disk group at the same time you create application disk groups, although it is not required for installing the application. You can wait until the end of setting up the environment to

convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that SQL is working in the cluster before adding the dynamic quorum volume.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Figure 11-2 shows a detailed view of the disk groups and volumes for SQL.

Figure 11-2 SFW disk groups and volumes for SQL virtual server SQLVS in MSCS setup



SQL disk group INST1_DG contains three volumes:

- INST1_DB1_VOL: Contains the SQL database.
- INST1_DB1_LOG: Contains the transaction log.
- INST1_DATA_FILES: Contains the system data files.

The mirrored quorum disk group and mirrored quorum volume will be created in "[Implementing a dynamic quorum resource](#)" on page 196.

Use the following procedures to create the appropriate disk groups and volumes. This section assumes you are using one database.

Creating dynamic cluster disk groups

When the tasks described in this section are completed, you will have a dynamic cluster disk group with volumes on shared storage. The dynamic cluster disk groups will be ready to be shared between nodes in the cluster.

Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it will be required later during the SQL installation process.

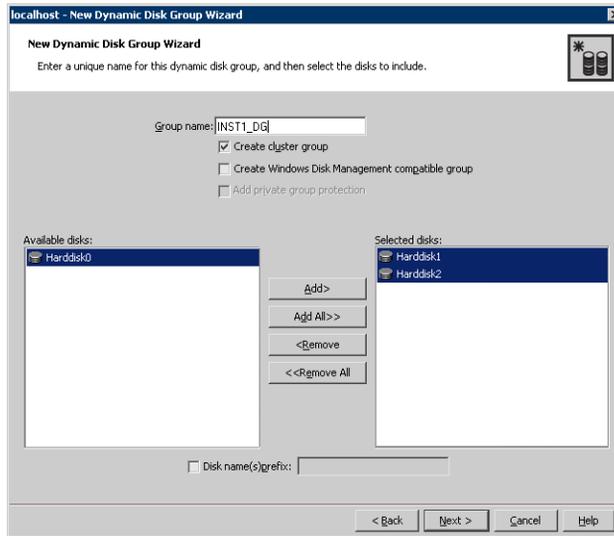
To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on a client system and can be used to manage all the servers remotely.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

Creating dynamic volumes

This section will guide you through the process of creating a volume on a dynamic disk group.

When creating a disk group to support a SQL Server 2000 solution, it is best to separate SQL data files from SQL log files and place them on separate volumes.

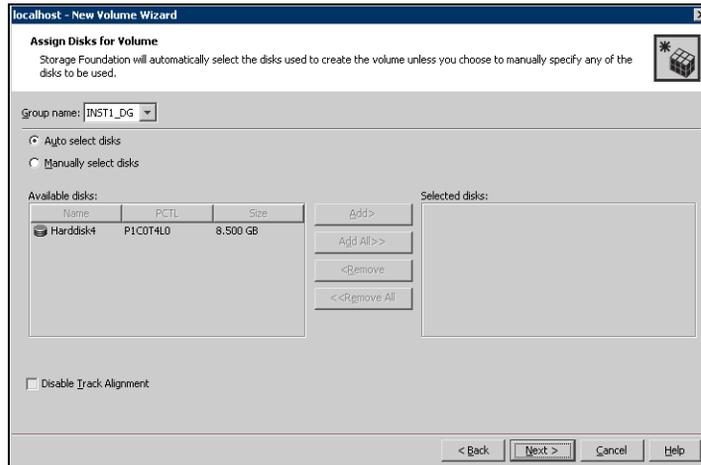
Repeat the procedure below to create the following volumes on the first node of the cluster:

- INST1_DATA_FILES: For storing the SQL system databases.
- INST1_DB1_VOL: For storing the user database.
- INST1_DB1_LOG: For storing the user database log.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

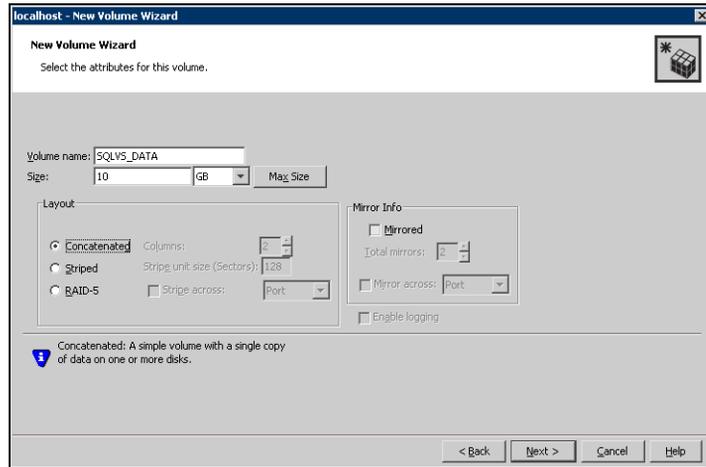


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

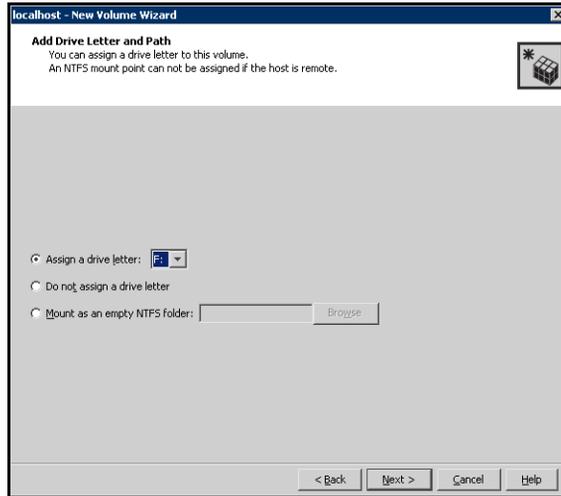
- 8 Click **Next**.

9 Specify the parameters of the volume.



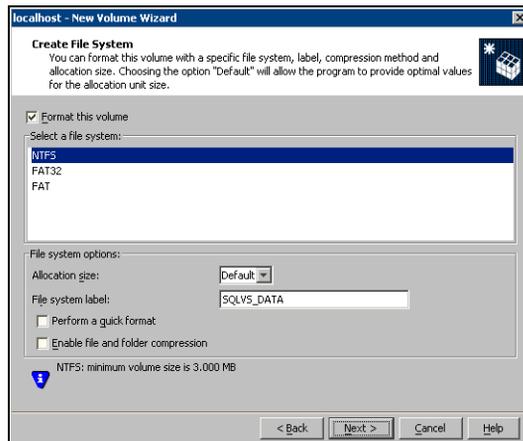
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
 - 14 Repeat these steps to create additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk group and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter

- Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Creating the SQL virtual server group

Before installing SQL you must create the SQL Server cluster group and add the appropriate resources.

Note: Before creating the resources, start the cluster service on all the nodes in the cluster.

To create an SQL Server cluster group

- 1 Launch the Cluster Administrator by selecting **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator**.
Make sure you are connected to the required cluster.
- 2 Create a new group by selecting the **Groups** node from the tree that is displayed in the left hand pane. Right-click to display the **Groups** menu. Select **New > Group** option from the menu. The **New Group** window appears.
- 3 Specify a name for the group in the **Name** field.
 - In the New Group Wizard specify a name `SQL_GROUP` for the SQL cluster group.
 - If required specify a description for this resource in the field provided. Click **Next**.
- 4 The Preferred Owners page appears. Make sure that all the preferred owners are added to the **Preferred Owners** list.
- 5 Click **Finish** to create the group.

You can now start adding resources to it.

Creating an IP address resource

A separate valid IP address for the SQL virtual server is necessary to install SQL Server on more than one node.

To create an IP Address resource

- 1 Right click on the SQL cluster group (`SQL_GROUP`) and select **New > Resource**.
- 2 In the Resource creation wizard, configure the IP address. Specify a name for the **IP Address** resource.
Add a **Description** if required.

- 3 Select the **IP address** from the **Resource Type** field drop down list. Click **Next**.
- 4 In the **Possible Owners** page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 5 In the **Dependencies Page**, make sure the **Resource Dependencies** pane is empty, and click **Next**.
- 6 On the **TCP/IP Address Parameters** page, set the TCP/IP parameters.
 - Enter the IP address and the corresponding subnet mask.
 - Make sure the Network is set to **Public** and click **Finish** to create the **IP Address** resource.
- 7 Bring the resource online.

Creating the SQL disk group resource

SQL virtual server installation requires a separate volume, INST1_DATA_FILES on which the system database files will be placed. You must create a Volume Manager Disk Group resource for the disk group that contains this volume. Creating this resource will enable SQL to monitor the system database files.

To create the disk group resource

- 1 If your cluster administrator is already open then proceed to the Step 2. To launch the Cluster Administrator select from **Start > Setting > Control Panel > Administrative Tools > Cluster Administrator**. You can create a short cut for the cluster administrator on the desktop to avoid accessing it every time from this path.
- 2 In the left pane of the cluster administrator select the SQL_GROUP Group and right-click. Select **New > Resource** from the menu that appears. The New Resource wizard appears.
- 3 Specify a name for the disk group resource, for example, SQL_DG_RES in the **Name** field.
If required, you can add a description about the resource in the **Description** field.
- 4 Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource Type** field drop down list.
Click **Next**.
- 5 By default, in the Possible Owners page, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 6 On the dependencies page, click **Next**. You do not need to set any dependency for a Disk Group resource.

- 7 On the **Volume Manager Disk Group Parameters** page, select the created disk group. Click **Finish**.

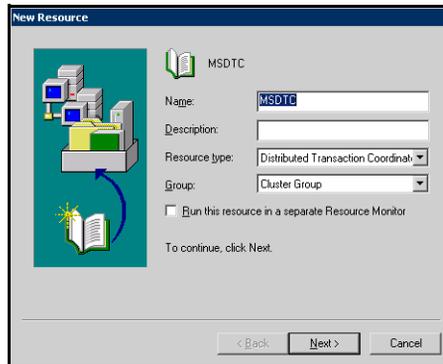
The specified disk group resource, SQL_DG_RES resource is created under the SQL_GROUP group.

Creating the MSDTC resource

Prior to installing SQL Server, create the MSDTC resource. This procedure is required for multiple instances of SQL.

To create the MSDTC resource

- 1 From Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**), right-click **Cluster Group**, click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify a name for the MSDTC resource. If necessary, add a description about the resource.



- 3 Select **Distributed Transaction Coordinator** from the **Resource type** list and click **Next**.
- 4 In the Possible Owners dialog box, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 5 In the Dependencies dialog box, select the cluster IP address, cluster name, and physical disk resources from the **Available Resources** list and add them to the **Resource dependencies** list. The volume with the SQL Server system data files must be included. Click **Finish**.
- 6 Click **OK**.
- 7 Bring the MSDTC resource online. In the left pane, expand the Groups icon.

- 8 Click **Cluster Group**.
- 9 Right-click **Bring Online**. The state changes to online.

Installing SQL Server 2000

This section provides some useful tips on how to install SQL Server 2000 on the primary and secondary sites. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

Before you proceed with installing SQL note the following points:

- The SQL Setup program detects that the system is a cluster, and prompts you for information appropriately at the beginning of the install process.
- The Setup program automatically installs a new, separate instance of SQL Server binaries on the local disk of each server in the cluster. The binaries are installed in exactly the same path on each cluster node, so it is important to ensure that each node has a local drive letter in common with all the other nodes in the cluster.
- The Setup program also installs the system databases on the specified cluster (shared) disk. System database files must be on a clustered disk so that they can be shared between the nodes (and failed over when necessary), because these databases contain specific user login and database object information that must be the same for each node. The virtual server name will allow users access to the online node.

Caution: Installation of a named instance of SQL Server 2000 virtual server on a Windows 2003-based cluster will fail. See: <http://support.microsoft.com/kb/815431>

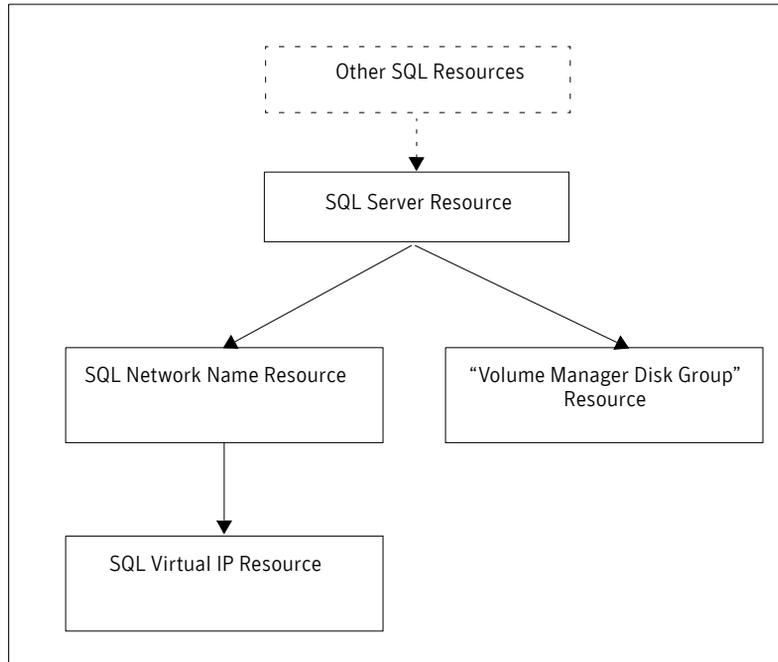
To install SQL Server

- 1 Verify the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters) See “[Managing disk group and volumes](#)” on page 189.
- 2 Launch the Microsoft SQL Server Installation Wizard.

- 3 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 and higher patches after installing SQL Server.
- 4 In the **Computer Name** wizard page, select the **Virtual Server** option, and specify a name for the virtual SQL server name; for example, `SQLVS`. Note this name as you will need to use the same name when installing on the secondary cluster nodes. Click **Next** to continue.
- 5 In the Failover Clustering wizard page, specify the following information:
 - The IP address that is intended for the SQL virtual server in the **IP Address** field.
 - The appropriate subnet for the IP in the **Subnet** field.
 - The appropriate public network that you have configured in the **Network to use** field. By default, the configured public network will be selected. However, if there are more than one network cards configured for public network then you can select the appropriate one from the list. Click **Next**.
- 6 The Cluster Disk Selection wizard page appears. This screen allows you to specify the logical disk from the shared disk array that will be used for the SQL Server 2000 system database files. Select the drive letter of the volume, `INST1_DATA_FILES`. Click **Next**.
- 7 In the Cluster Management screen, specify the nodes in the cluster on which you want SQL to fail over. Make sure these are a part of the Configured Nodes box on this wizard page. Once you are sure that all the required nodes are in the Configured nodes box, click **Next**.
- 8 The Remote Information wizard page appears. In this page, specify the administrative user name and password that is valid on all the nodes. It is recommended that the domain account identification be used so that it is acceptable on all nodes. Click **Next**.
- 9 In the Instance Name dialog box, the **Default** option is selected. Since this is the first instance of SQL being installed, leave the Default option selected and click **Next**.
- 10 Follow the wizard page instructions to complete the SQL installation on all the nodes of the cluster.

Once SQL is installed, the SQL Server Resource with dependencies on the SQL Network Name and the Volume Manager Disk Group resource is automatically created. The following dependency graph indicates the dependencies that are established.

Dependency graph after the SQL installation is completed.



Applying the SQL SP4 patch

After installing SQL 2000, you need to apply the SP4 or higher patch for each of the nodes.

Verifying SQL installation

Click **Start > Programs > Microsoft SQL Server**. Select **Enterprise Manager** from the menu that appears to start the SQL Server Enterprise Manager.

Implementing a dynamic quorum resource

One of the key advantages of using SFW with Microsoft clustering is the ability to create a mirrored quorum resource that adds fault tolerance to the quorum and protects the cluster. Complete the following tasks:

- Create a dynamic cluster disk group for the quorum resource with a mirrored volume
- Create the quorum resource for the Cluster Group

- Change the quorum resource to a dynamic mirrored quorum resource.

Creating a dynamic cluster disk group and a mirrored volume for the quorum resource

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using three (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

To create a three-way mirrored volume in the New Volume wizard, select the Concatenated layout, select the Mirrored check box, and specify three mirrors. See “[Configuring SFW disk groups and volumes](#)” on page 181 for details on a creating cluster disk groups and volumes.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by Microsoft clustering.

To create a three-way mirrored volume using the New Volume wizard

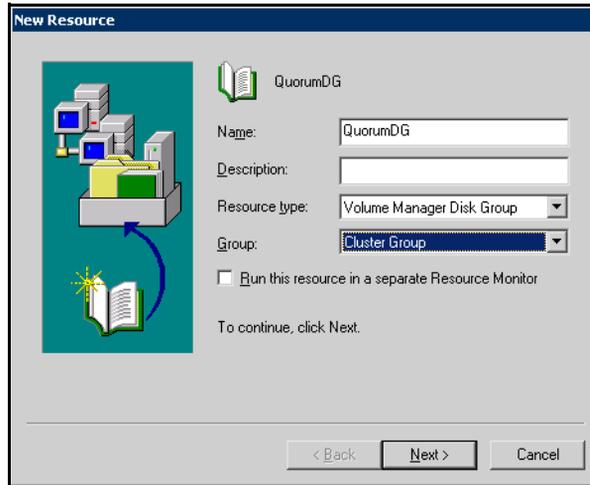
- 1 Create the cluster disk group with three small disks.
- 2 Create a volume with the three disks, in the sample, this is INST1_QUORUM.
- 3 Select the **Concatenated** layout, select the **Mirrored** check box, and specify three mirrors.
For full details on a creating cluster disk groups and volumes, see “[Creating dynamic cluster disk groups](#)” on page 183.

Creating the quorum resource for the cluster group

To create a quorum resource for the cluster group

- 1 Verify that the Cluster Group is online on the same node where you created the disk group.
- 2 To create the quorum resource, open Cluster Administrator (**Start > Control Panel > Administrative Tools > Cluster Administrator**).
- 3 In the left pane of the Cluster Administrator, select and right click the SQL_GROUP Group. Select **New**, then **Resource** from the menu that appears.

- 4 In the New Resource dialog box, Specify a name for the quorum resource, for example, QUORUM_DG.
If necessary, add a description about the resource.



- 5 Select **Volume Manager Disk Group** from the **Resource type** list and click **Next**.
- 6 In the Possible Owners dialog box, click **Next**.
- 7 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a quorum resource.
- 8 In the Volume Manager Disk Group Parameters dialog box, select the disk group and click **Finish**.
- 9 Click **OK**.
- 10 Bring the newly added resource online.

Changing the quorum resource to a dynamic mirrored quorum resource

To change the quorum to a dynamic mirrored quorum resource

- 1 From Cluster Administrator, right-click the cluster name in the configuration tree, and click **Properties**.
- 2 Select the Quorum tab of the Properties window.
- 3 Select the name of the dynamic quorum disk group resource that was added.
- 4 Click **OK**.

Testing the cluster

You can verify your installation by moving the cluster group between nodes to see if it fails over properly. The ultimate test of the cluster's failover capability involves shutting down the node that is currently online and bringing it back up after the cluster fails over to the other node.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.

- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Deploying SFW with MSCS: New SQL 2005 installation

This chapter covers the following topics:

- [Tasks for a new SQL Server 2005 installation with SFW and MSCS \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing an MSCS cluster](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Creating the SQL virtual server group](#)
- [Creating the MSDTC resource](#)
- [Installing SQL Server 2005](#)
- [Implementing a dynamic mirrored quorum resource](#)
- [Verifying the cluster configuration](#)

Tasks for a new SQL Server 2005 installation with SFW and MSCS (Windows Server 2003)

You can install and configure Storage Foundation for Windows with MSCS and SQL Server 2005 for a new installation on Windows Server 2003. This environment involves an active/passive configuration with one to one failover capability for high availability.

If you will use Veritas Volume Replicator and replication, see:

[“Deploying SFW and VVR with MSCS: New SQL 2005 installation”](#) on page 441.

If you are deploying SQL Server 2005 on Windows Server 2008, see:

[Chapter 13, “Deploying SFW with Microsoft failover clustering: New SQL 2005 installation”](#) on page 241

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 12-1 Tasks for deploying SFW with MSCS for SQL Server 2005

Objective	Tasks
“Reviewing the requirements” on page 204	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 206	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive SQL configuration in a two-node cluster ■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 209	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Establishing an MSCS cluster” on page 210	<ul style="list-style-type: none"> ■ Reviewing general guidelines to establish an MSCS cluster
“Installing SFW with MSCS/Failover Cluster option” on page 211	<ul style="list-style-type: none"> ■ Modifying the driver signing options for Windows 2003 remote systems ■ Installing Veritas Storage Foundation for Windows (automatic installation) ■ Restoring the driver signing options for Windows 2003 remote systems
“Configuring SFW disk groups and volumes” on page 220	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes
“Managing disk group and volumes” on page 228	<ul style="list-style-type: none"> ■ Exporting and importing to cluster nodes
“Creating the SQL virtual server group” on page 230	<ul style="list-style-type: none"> ■ Creating a SQL Server cluster group ■ Creating the virtual server IP address ■ Creating the disk group resource
“Creating the MSDTC resource” on page 232	<ul style="list-style-type: none"> ■ Creating the MSDTC resource for SQL Server

Table 12-1 Tasks for deploying SFW with MSCS for SQL Server 2005

Objective	Tasks
“Installing SQL Server 2005” on page 233	<ul style="list-style-type: none"> ■ Installing SQL Server ■ Verifying SQL Server installation
“Implementing a dynamic mirrored quorum resource” on page 237	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the cluster group ■ Changing the quorum resource to a dynamic mirrored quorum resource.
“Verifying the cluster configuration” on page 239	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation.

Supported software for MSCS and SFW

- Veritas Storage Foundation 5.0 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
- Microsoft SQL servers and their operating systems:

Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required)	■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required)
	■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required)
	■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)
	■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)
	■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)
Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required)	■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)
	■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)
	■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 12-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

- One CD-ROM drive accessible to the system on which you are installing SFW.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- MSCS requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends having two adapters for the private network and routing each private network adapter through a separate hub or switch to avoid single points of failure.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft SQL Server documentation for instructions on creating a reverse lookup zone.
- MSCS requires two disks for SQL: one for SQL database files and one for SQL log files.
- Each system requires 1 GB of RAM for SFW.
- Each system requires a minimum of 1 GB of RAM for SQL Server 2005; refer to your Microsoft documentation for more information.
- SFW requires administrator privileges to install the software.

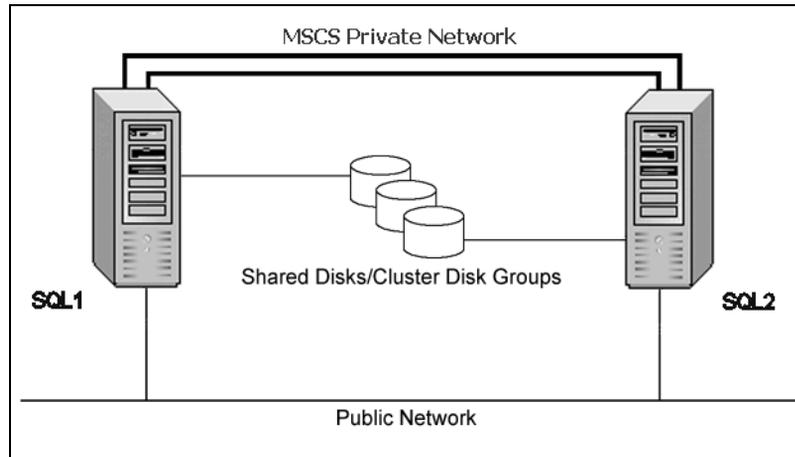
Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Reviewing the configuration

This chapter describes the tasks necessary to create a virtual server in an active/passive SQL configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node.

The following figure illustrates a typical active/passive configuration. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured on the active node (SYSTEM1). If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 12-1 Active/passive configuration



Some key points about the configuration:

- An MSCS cluster must be running before you can install SFW.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- SFW enables you to create a dynamic mirrored quorum. The quorum resource maintains the cluster database and critical recovery information in a recovery log.

In an MSCS cluster without SFW, the quorum disk is a point of failure because MSCS only supports a basic physical disk and does not enable you to mirror the quorum resource.

The main advantage of SFW is that it provides a dynamic mirrored quorum resource for MSCS. If the quorum resource fails, the mirror takes over for the resource. In this configuration, Symantec recommends creating a three-way mirror for the quorum to provide additional fault tolerance. If possible, do not use the disks assigned to the quorum for any other purpose. You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that SQL is working in the cluster before adding the dynamic quorum volume.

- SFW enables you to add fault-tolerance to data volumes. Symantec recommends mirroring log volumes and a mirrored striped RAID layout for data volumes. SFW offers multiple disk groups, mirrors, capacity management and automatic volume growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, dynamic multi-pathing, and enhanced snapshot capabilities with FlashSnap. Refer to the *Veritas Storage Foundation for Windows Administrator's Guide* for details on these features.

Sample configuration

A sample setup is used through this guide to illustrate the installation and configuration tasks.

During the configuration process you will create virtual IP addresses for the following:

- Cluster IP address: used by MSCS cluster
- SQL virtual server: the IP address should be the same on all nodes

You should have these IP addresses available before you start deploying your environment.

The following names describe the objects created and used during the installation and configuration.

Name	Object
SYSTEM1 & SYSTEM2	server names
SQL_GROUP	Microsoft SQL Server resource group

Reviewing the configuration

SQLCLUST	Microsoft SQL Server virtual cluster (underscores not supported)
SQLVS	Microsoft SQL Server virtual server
INST1	Microsoft SQL Server instance name
INST1_DG	disk group for Microsoft SQL volumes
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
SQLVS_QRM	volume for storing the MSCS cluster quorum
QUORUM_DG	quorum volume disk group

Configuring the storage hardware and network

Use the following procedures to configure the storage hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent MSCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 5 In the Public Status dialog box, on the General tab, click **Properties**.
- 6 In the Public Properties dialog box, on the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.

- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Establishing an MSCS cluster

Before installing SFW, you must establish an MSCS cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To establish an MSCS cluster (general guidelines)

- 1 Verify that the quorum disk has been created before installing MSCS and configuring a cluster. (For IA64 systems, the quorum must be created using MBR instead of GPT or it will not be visible.)
- 2 Configure the shared storage and create a partition with drive letter "Q" for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster (SYSTEM1) using MSCS Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Verify that the node can access the shared storage.
- 4 Connect the second node to the shared storage.
- 5 Add the second node (SYSTEM2) using Cluster Administrator on that system.
- 6 Test the cluster by using the Move Group command to move the cluster resources to the second node.
SYSTEM2 becomes the active cluster node.

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running an MSCS cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the MSCS cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 211.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 213.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 219.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
- Moving the online groups

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The table below describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 12-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
 If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a MSCS configuration.

To install the product

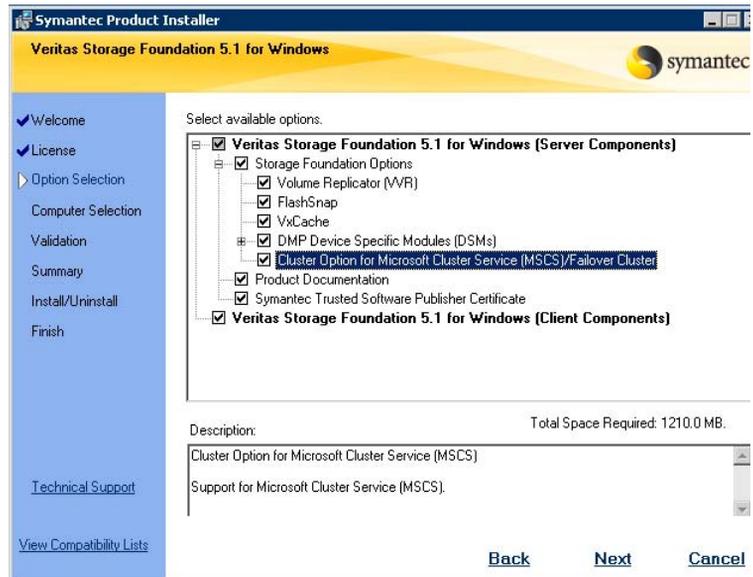
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

- 3 Click **Storage Foundation 5.1 for Windows**.



- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for "**I accept the terms of the license agreement,**" and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
- 8 Repeat for additional license keys.
To remove a license key, click the key to select it and click **Remove**.
To see the license key's details, click the key.
- 9 Click **Next**.

- 10 Specify the product options by selecting the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.

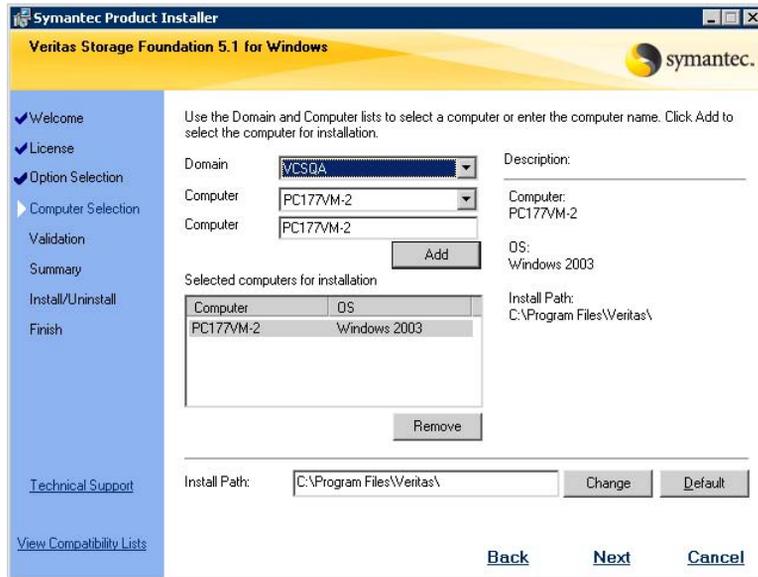


Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.

Note that under Veritas Dynamic Multi-pathing, you can select DMP Device Specific Modules (DSMs).

- 11 Click **Next**.
- 12 Verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is checked, to install the client component and click **Next**.

13 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 14 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 15 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
 If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 16 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that Microsoft clustering allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

Dynamic Multi-pathing

Additionally, if you selected the Dynamic Multi-pathing option, a warning appears:

- For DMP installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.

- For DMP DSM installations—the time required to install the Veritas Dynamic Multi-pathing DSM feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 17 Review the information and click **Install**. Click **Back** to make changes.
 - 18 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.
If the installation is successful on all systems, the installation report screen appears.
If a security alert asks you to accept the Veritas driver software, click **Yes**. This alert appears if your local computer has its driver signing options set to Warn. If your local computer has its driver signing options set to Block, installation fails.
 - 19 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 20](#) if you are installing SFW on the local node only.
 - Proceed to [step 22](#) if you are installing SFW on local and remote systems.
 - 20 To complete the installation, click **Finish**.
 - 21 Click **Yes** to reboot the system and complete the installation.
 - 22 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
 - 23 Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
 - 24 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available.
 - 25 Click **Next**.
 - 26 Click **Finish**.
 - 27 Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW Installation
- Resetting the driver signing options

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.
- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 211.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for SQL. A dynamic disk group is a collection of one or more disks that behaves as a single storage repository. Within each disk group, you can have dynamic volumes with different layouts.

Configuring disk groups and volumes involves the following tasks:

- [Planning disk groups and volumes](#)
- [Creating dynamic cluster disk groups](#)
- [Creating dynamic volumes](#)
- [Managing disk group and volumes](#)

Planning disk groups and volumes

Before installing SQL, you must create disk groups and volumes using the VEA console installed with SFW.

A dynamic cluster disk group is a collection of one or more disks that behave as a single storage repository and which can potentially be accessed by different computers.

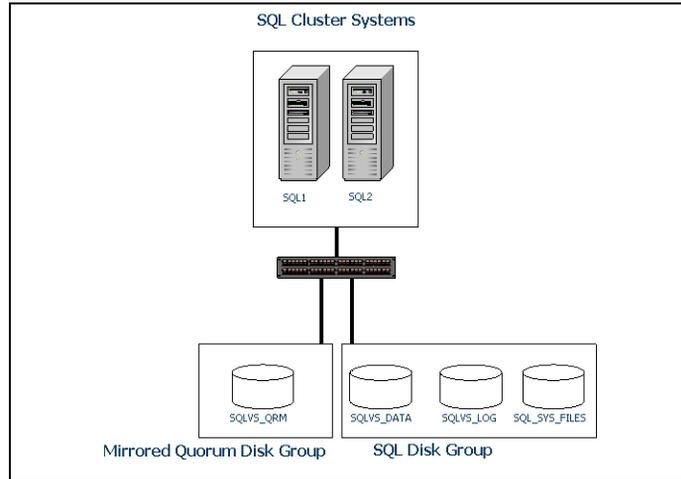
Before creating a disk group, consider:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.
- Use of a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. As noted earlier, you can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that SQL is working in the cluster before adding the dynamic quorum volume.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Below is a detailed view of the disk groups and volumes for SQL:

Figure 12-2 SFW disk groups and volumes for SQL virtual server SQLVS in MSCS setup



SQL disk group SQLVS contains three volumes:

- INST1_DB1_VOL: Contains the SQL database. Each database in an SQL storage group typically resides on a separate volume.
- SQLVS_LOG: Contains the transaction log.
- INST1_DATA_FILES: Contains volume for Microsoft SQL Server system data files.

The mirrored quorum disk group and mirrored quorum volume will be created in [“Implementing a dynamic mirrored quorum resource”](#) on page 237.

Use the following procedures to create the appropriate disk groups and volumes. This section assumes you are using one database.

Creating dynamic cluster disk groups

You create a dynamic cluster disk group with volumes on shared storage so that they can be shared between nodes in the cluster.

Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it will be required later during the SQL the installation process.

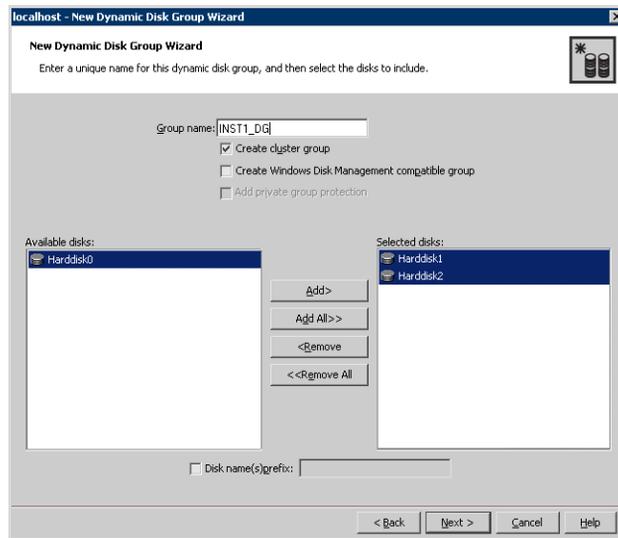
To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on client system and can be used to manage all the servers remotely.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

- Click **Next**.

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the new disk group.

Creating dynamic volumes

This section will guide you through the process of creating a volume on a dynamic disk group.

When creating a disk group to support a SQL Server 2005 solution, it is best to separate SQL data files from SQL log files and place them on separate volumes.

Repeat the procedure below to create the following volumes on the first node of the cluster:

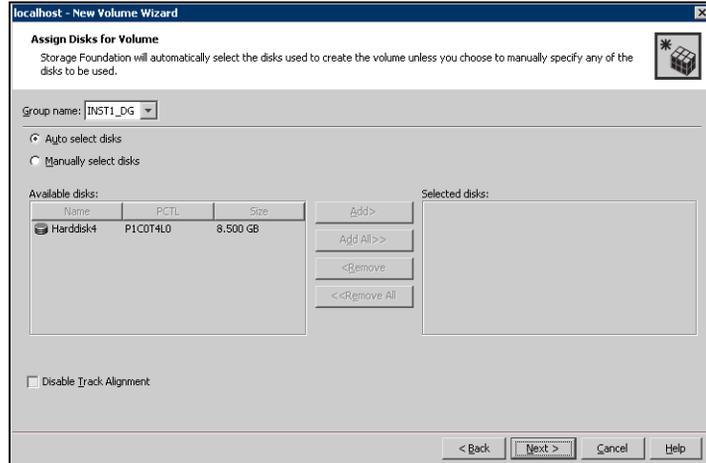
- **INST1_DATA_FILES**: For storing the SQL system databases.
- **INST1_DB1_VOL**: For storing the user database.
- **SQLVS_LOG**: For storing the user database log.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example **INST1_DG**.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

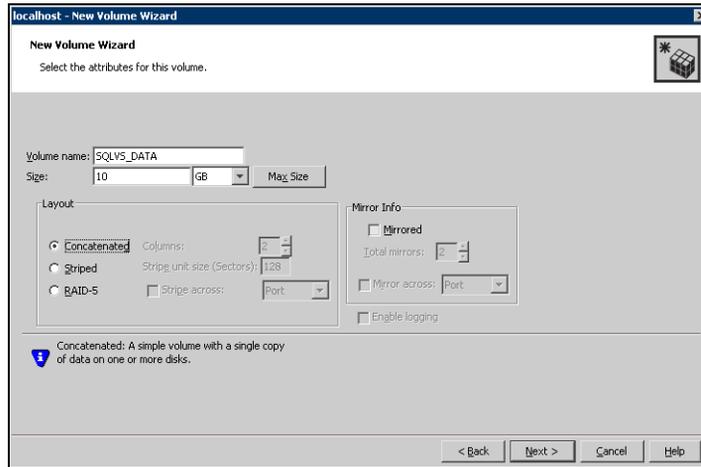


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

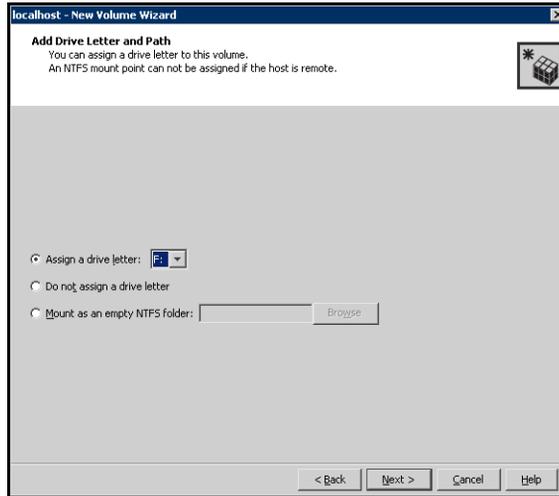
- 8 Click **Next**.

9 Specify the parameters of the volume.



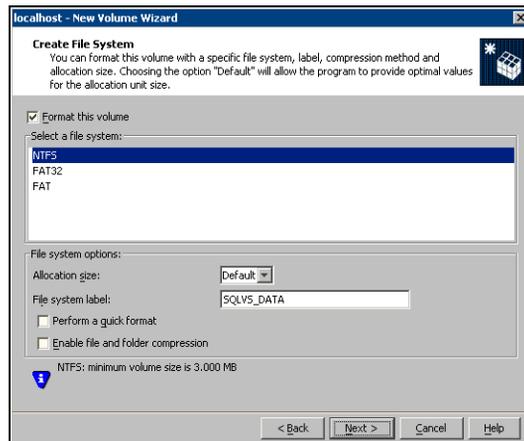
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

13 Click **Finish** to create the new volume.

14 Repeat these steps to create additional volumes.

Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk group and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.

- To assign a drive letter
Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Creating the SQL virtual server group

Before installing SQL you must create the SQL server cluster group and add the appropriate resources.

Note: Before creating the resources, start the cluster service on all the nodes in the cluster.

To create an SQL Server cluster group

- 1 Launch the Cluster Administrator by selecting **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator**.
Make sure you are connected to the required cluster.
- 2 Create a new group by selecting the **Groups** node from the tree that is displayed in the left hand pane. Right-click to display the **Groups** menu. Select **New > Group** option from the menu. The **New Group** window appears.
- 3 Specify a name for the group in the **Name** field.
 - In the New Group Wizard specify a name `SQL_GROUP` for the SQL cluster group.
 - If required specify a description for this resource in the field provided. Click **Next**.
- 4 The Preferred Owners page appears. Make sure that all the preferred owners are added to the **Preferred Owners** list.
- 5 Click **Finish** to create the group.

You can now start adding resources to it.

Creating an IP address resource

A separate valid IP address for the SQL virtual server is necessary to install SQL Server on more than one node.

To create an IP Address resource

- 1 Right click on the SQL cluster group (SQL_GROUP) and select **New > Resource**.
- 2 In the Resource creation wizard, configure the IP address:
 - Specify a name for the **IP Address** resource and add a **Description** if required.
 - Select the **IP address** from the **Resource Type** field drop down list. Click **Next**.
- 3 In the **Possible Owners** page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 4 In the **Dependencies Page**, make sure the **Resource Dependencies** pane is empty, and click **Next**.
- 5 On the **TCP/IP Address Parameters** page, set the TCP/IP parameters by entering the IP address and the corresponding subnet mask.
- 6 Set the Network to **Public** and click **Finish** to create the **IP Address** resource.
- 7 Bring the resource online.

Creating the SQL disk group resource

SQL virtual server installation requires a separate volume, INST1_DATA_FILES on which the system database files will be placed. You must create a Volume Manager Disk Group resource for the disk group that contains this volume. Creating this resource will enable SQL to monitor the system database files.

To create the disk group resource

- 1 If the Cluster Administrator is already open, then proceed to Step 2. To launch the Cluster Administrator, select it from **Start > Setting > Control Panel > Administrative Tools > Cluster Administrator**. You can create a short cut for the Cluster Administrator on the desktop to avoid accessing it every time from this path.
- 2 In the left pane of the Cluster Administrator, select and right-click the SQL_GROUP Group. Select **New > Resource** from the menu. The New Resource wizard appears.

- 3 Specify a name for the disk group resource, for example, `SQL_DG_RES` in the **Name** field.
If required, you can add a description about the resource in the **Description** field.
Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource Type** field drop down list.
Click **Next**.
- 4 By default, in the Possible Owners page, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 5 On the Dependencies page, click **Next**. You do not need to set any dependency for a Disk Group resource.
- 6 On the **Volume Manager Disk Group Parameters** page, select the created disk group. Click **Finish**.

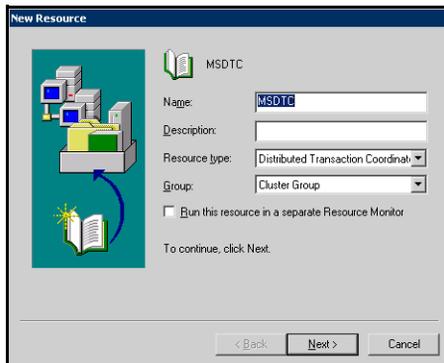
The specified disk group resource, `SQL_DG_RES` resource is created under the `SQL_GROUP` group.

Creating the MSDTC resource

Prior to installing SQL Server, create the MSDTC resource. This procedure is required for multiple instances of SQL.

To create the MSDTC resource

- 1 From Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**), right-click **Cluster Group**, click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify a name for the MSDTC resource. If necessary, add a description about the resource.



- 3 Select **Distributed Transaction Coordinator** from the **Resource type** list and click **Next**.
- 4 In the Possible Owners dialog box, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 5 In the Dependencies dialog box, select the cluster IP address, cluster name, and physical disk resources from the **Available Resources** list and add them to the **Resource dependencies** list. The volume with the SQL Server system data files must be included. Click **Finish**.
- 6 Click **OK**.
- 7 Bring the MSDTC resource online. In the left pane, expand the Groups icon.
- 8 Click **Cluster Group**.
- 9 Right-click **Bring Online**. The state changes to online.

Installing SQL Server 2005

This section provides some useful tips on how to install SQL Server 2005. As you progress through the installation, use these guidelines to create an installation that will function properly in your environment.

Note: Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

Before you proceed with installing SQL note the following points:

- Verify the cluster disk group is imported to the first node and the volumes are mounted. See “[Managing disk group and volumes](#)” on page 228.
- The Setup program automatically installs a new, separate instance of SQL Server binaries on the local disk of each server in the cluster. The binaries are installed in exactly the same path on each cluster node, so it is important to ensure that each node has a local drive letter in common with all the other nodes in the cluster.
- The Setup program also installs the system databases on the specified cluster (shared) disk. System databases must be on a clustered disk so that they can be shared between the nodes (and failed over when necessary), because these databases contain specific user login and database object information that must be the same for each node. The virtual server name will allow users access to the online node.

To install SQL Server 2005

- 1 Begin the SQL Server 2005 installation, following the instructions from Microsoft. To begin the installation, navigate to the installation directory and launch **splash.hta**.
- 2 Review the hardware and software requirements for SQL 2005.
- 3 Under the **Install** section, select **Server components, tools, Books Online, and samples**.
- 4 Continue with the installation, following the instructions from Microsoft. Complete the SQL Server Component Update, System Configuration Check, and Registration Information pages.
- 5 In the Components to Install dialog box, select the **SQL Server Database Services**.
To cluster SQL Server on MSCS, you must select the **Create a SQL Server failover cluster** option.
Select the optional components:
 - Analysis Service. If this option is selected, the option **Create an Analysis Server failover cluster** must also be selected.
 - Notification Services
 - Integration Services
 - Workstation ComponentsClick the **Advanced** option.
- 6 In the **Feature Selection** dialog box, specify the path for SQL Server data files and other services.
- 7 Expand **SQL Server Database Services** and select **Data Files**.
- 8 Select **Browse** to reset the installation path.
- 9 Set the installation path in the Change Folders dialog box to the drive letter and location of the volume created for the SQL Server system data files (INST1_DATA_FILES). Allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.
This must be the same as the path on all nodes.
- 10 If you selected the **Analysis Services** option in [step 5](#) on page 234, expand **Analysis Services**, select **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files.
This must be the same as the path on all nodes.
- 11 Click **Next**.
- 12 In the **Instance Name** dialog box, enter an instance name or accept the default. Click **Next**.

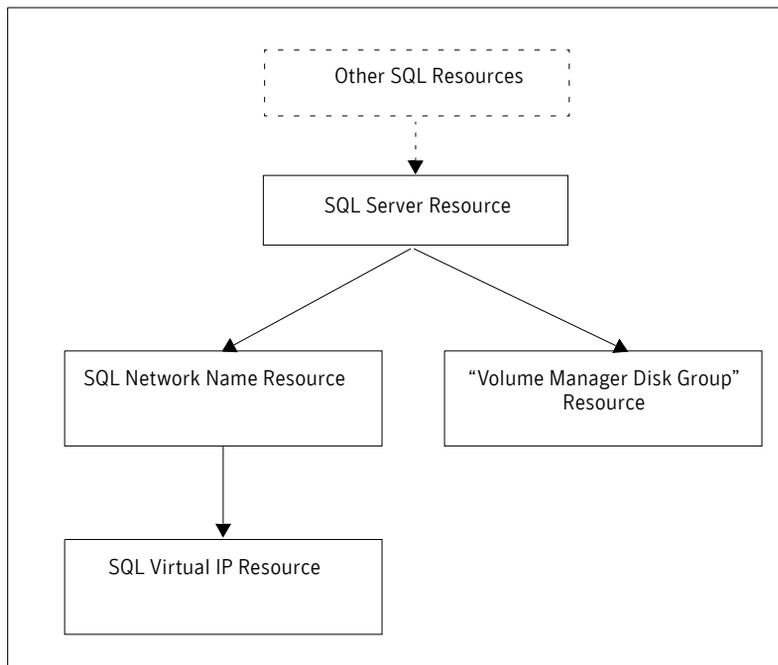
Only one default instance is allowed per cluster.

Use the same instance name when installing SQL Server 2005 on the first node and on all failover nodes.

- 13 In the **Virtual Server Name** page, specify a unique name for the virtual SQL server, for example, *SQLVS*. Make a note of this name as you will need to use the same name when installing on the secondary cluster nodes. Click **Next**.
- 14 In the **Virtual Server Configuration** page, configure the virtual server. Select the appropriate public network that you have configured in the **Network to use** field. By default, the configured public network will be selected. However, if there are more than one network cards configured for public network then you can select the appropriate one from the list.
- 15 Specify the IP address that is intended for the SQL virtual server in the **IP Address** field.
- 16 The appropriate **Network address** and **Network subnet** will be displayed. Click **Next**.
- 17 In the **Cluster Group Selection** page, specify the cluster group with the logical disk from the shared disk array that will be used for the SQL Server 2005 system database files. You can also specify a custom path in the **Data Files** field. Click **Next**.
- 18 In the **Cluster Node Configuration** page, specify the nodes in the cluster for SQL failover by selecting them from **Available nodes**, and clicking **Add** to add them to the **Selected Nodes** box. Once all the required nodes are in the **Selected Nodes** box, click **Next**.
- 19 In the **Remote Account Information** page, specify the administrative password that is valid on all the nodes. Click **Next**.
 - In the **Service Account** page, specify the type of account and information.
 - Select the type of service account. Select **Use the built-in System account** if you do not want to replicate data. Otherwise, select **Use a domain user account**.
 - If you chose to run the service in the context of a domain user, specify the information for the user.
 - Click **Next**.
- 20 In the **Domain Groups for Clustered Services** page, use the browse button at the right to select a **DomainName** and **GroupName** for each of the selected SQL Server options. Click **Next**.

- 21 In the **Authentication Mode** page, select **Mixed Mode** (recommended option), and specify your password.
- 22 Follow the wizard page instructions to complete the SQL installation on all the nodes of the cluster.
Once SQL is installed, the SQL Server Resource with dependencies on the SQL Network Name and the Volume Manager Disk Group resource is created. The following dependency graph indicates the dependencies that are established.

Dependency graph after the SQL installation is completed



Verifying SQL installation

Click **Start > Programs > Microsoft SQL Server**. Select **Enterprise Manager** from the menu that appears to start the SQL Server Enterprise Manager.

Implementing a dynamic mirrored quorum resource

One of the key advantages of using SFW with MSCS is the ability to create a mirrored quorum resource that adds fault tolerance to the quorum and protects the cluster. Complete the following tasks:

- Create a dynamic cluster disk group for the quorum resource with a mirrored volume
- Create the quorum resource for the Cluster Group
- Change the quorum resource to a dynamic mirrored quorum resource.

Creating a dynamic cluster disk group and a mirrored volume for the quorum resource

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using three (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

For full details on a creating cluster disk groups and volumes, see “[Creating dynamic cluster disk groups](#)” on page 222.

To create a three-way mirrored volume using the New Volume wizard

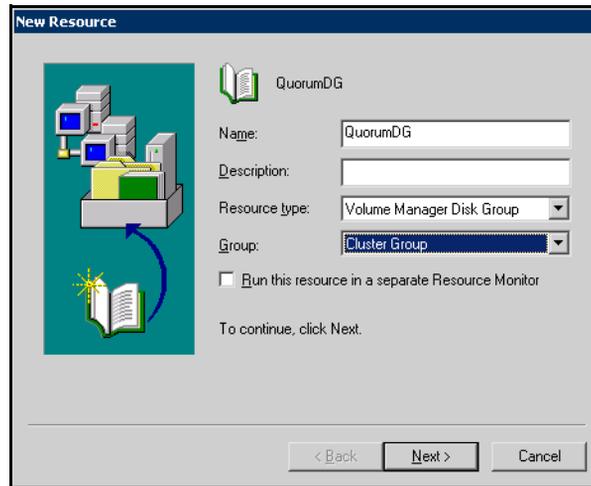
- 1 Create the cluster disk group with three small disks.
- 2 Create a volume with the three disks, in the sample this is SQLVS_QRM.
- 3 Select the **Concatenated** layout, select the **Mirrored** check box, and specify three mirrors.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by MSCS.

Creating the quorum resource for the cluster group

To create a quorum resource for the cluster group

- 1 Verify that the Cluster Group is online on the same node where you created the disk group.
- 2 Create the quorum resource by opening Cluster Administrator (**Start > Control Panel > Administrative Tools > Cluster Administrator**).
- 3 In the left pane of the Cluster Administrator, select the `SQL_GROUP` Group and right-click. Select **New**, then **Resource** from the menu that appears.
- 4 In the New Resource dialog box, specify a name for the quorum resource, for example, `QUORUM_DG`.
If necessary, add a description about the resource.



- 5 Select **Volume Manager Disk Group** from the **Resource type** list and click **Next**.
- 6 In the Possible Owners dialog box, click **Next**.
- 7 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a quorum resource.
- 8 In the Volume Manager Disk Group Parameters dialog box, select the disk group and click **Finish**.
- 9 Click **OK**.
- 10 Bring the newly added resource online.

Changing the quorum resource to a dynamic mirrored quorum resource

To change the quorum to a dynamic mirrored quorum resource

- 1 From Cluster Administrator, right-click the cluster name in the configuration tree, and click **Properties**.
- 2 Select the Quorum tab of the Properties window.
- 3 Select the name of the dynamic quorum disk group resource that was added.
- 4 Click **OK**.

Verifying the cluster configuration

You can verify your installation by moving the cluster group between nodes to see if it fails over properly. The ultimate test of the cluster's failover capability involves shutting down the node that is currently online and bringing it back up after the cluster fails over to the other node.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.

- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Deploying SFW with Microsoft failover clustering: New SQL 2005 installation

This chapter covers the following topics:

- [Tasks for a new SQL Server 2005 installation with SFW and Microsoft failover clustering \(Windows Server 2008\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing a Microsoft failover cluster](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Creating the SQL Server virtual server group](#)
- [Installing SQL Server 2005](#)
- [Implementing a dynamic mirrored quorum resource](#)
- [Verifying the cluster configuration](#)

Tasks for a new SQL Server 2005 installation with SFW and Microsoft failover clustering (Windows Server 2008)

You can install and configure Storage Foundation for Windows with Microsoft failover clustering and SQL Server 2005 on Windows Server 2008. This environment involves an active/passive configuration with one to one failover capability for high availability.

If you will use Veritas Volume Replicator and replication, see [“Deploying SFW and VVR with Microsoft failover clustering: New SQL 2005 installation”](#) on page 511.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 13-1 Tasks for deploying SFW with failover clustering for SQL Server 2005

Objective	Tasks
“Reviewing the requirements” on page 244	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 245	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive SQL configuration in a two-node cluster ■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 249	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Establishing a Microsoft failover cluster” on page 250	<ul style="list-style-type: none"> ■ Reviewing general guidelines to establish a Microsoft failover cluster on Windows Server 2008
“Installing SFW with MSCS/Failover Cluster option” on page 253	<ul style="list-style-type: none"> ■ Installing Veritas Storage Foundation for Windows (automatic installation)
“Configuring SFW disk groups and volumes” on page 261	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes
“Managing disk group and volumes” on page 269	<ul style="list-style-type: none"> ■ Setting up a SFW environment with SQL

Table 13-1 Tasks for deploying SFW with failover clustering for SQL Server 2005 (Continued)

Objective	Tasks
“Creating the SQL Server virtual server group” on page 271	<ul style="list-style-type: none"> ■ Creating a SQL Server cluster group ■ Creating the disk group resource
“Installing SQL Server 2005” on page 272	<ul style="list-style-type: none"> ■ Installing SQL ■ Verifying SQL Server installation
“Implementing a dynamic mirrored quorum resource” on page 276	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the Cluster Group ■ Changing the quorum resource to a dynamic mirrored quorum resource.
“Verifying the cluster configuration” on page 278	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation.

Supported software for Microsoft failover clusters and SFW

The following software is supported:

- Veritas Storage Foundation 5.1 for Windows (SFW)
Include the following option along with any others applicable to your environment:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
- For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SQL Server 2005 SP2 required) ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition

Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SQL Server 2005 SP2 required) ■ Windows Server 2008 for 64-bit Itanium (IA64) ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 13-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

- One CD-ROM drive accessible to the system on which you are installing SFW.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- Microsoft failover clustering requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends having two adapters for the private network and routing each private network adapter through a separate hub or switch to avoid single points of failure.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft SQL Server documentation for instructions on creating a reverse lookup zone.
- Microsoft failover clustering requires two disks for SQL: one for SQL database files and one for SQL log files.
- Each system requires 1 GB of RAM.
- SFW requires administrator privileges to install the software.

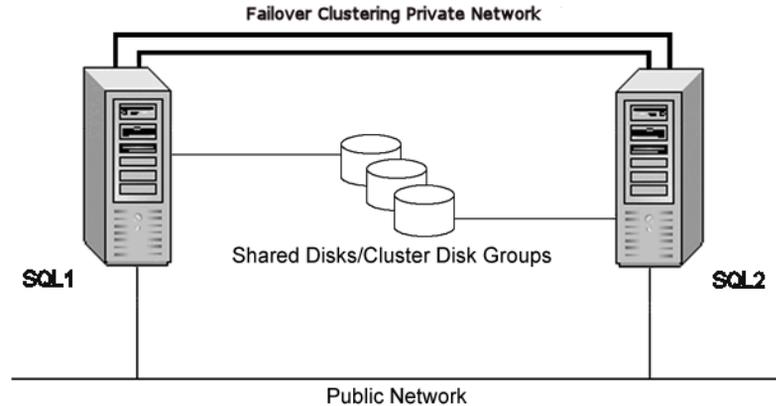
Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Reviewing the configuration

You create a virtual server in an active/passive SQL configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node.

The following figure illustrates a typical active/passive configuration. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured on the active node (SYSTEM1). If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 13-1 Active/passive configuration



Some key points about the configuration:

- A Microsoft failover cluster must be running before you can install SFW.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- SFW enables you to create a dynamic mirrored quorum. The quorum resource maintains the cluster database and critical recovery information in a recovery log.

In an failover cluster without SFW, the quorum disk is a point of failure because Microsoft failover clustering only supports a basic physical disk and does not enable you to mirror the quorum resource.

The main advantage of SFW is that it provides a dynamic mirrored quorum resource for Microsoft failover clustering. If the quorum resource fails, the mirror takes over for the resource. In this configuration, Symantec recommends creating a three-way mirror for the quorum to provide additional fault tolerance. If possible, do not use the disks assigned to the quorum for any other purpose.

You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables

you to verify that SQL is working in the cluster before adding the dynamic quorum volume.

- SFW enables you to add fault-tolerance to data volumes. Symantec recommends mirroring log volumes and a mirrored striped RAID layout for data volumes. SFW offers multiple disk groups, mirrors, capacity management and automatic volume growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, dynamic multi-pathing, and enhanced snapshot capabilities with FlashSnap. Refer to the *Veritas Storage Foundation for Windows Administrator's Guide* for details on these features.

Sample configuration

A sample setup is used through this guide to illustrate the installation and configuration tasks.

During the configuration process you will create virtual IP addresses for the following:

- Cluster IP address: used by Microsoft cluster
- SQL virtual server: the IP address should be the same on all nodes

You should have these IP addresses available before you start deploying your environment.

The following names describe the objects created and used during the installation and configuration.

Name	Object
SYSTEM1 & SYSTEM2	server names
SQL_GROUP	Microsoft SQL Server resource group
SQLCLUST	Microsoft SQL Server virtual cluster (underscores not supported)
SQLVS	Microsoft SQL Server virtual server
INST1	Microsoft SQL Server instance name
INST1_DG	disk group for Microsoft SQL volumes
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database

INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
SQLVS_QRM	volume for storing the Microsoft Failover Cluster quorum
QUORUM_DG	quorum volume disk group

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent the Microsoft cluster from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.
- 13 Close the window.

Establishing a Microsoft failover cluster

Before installing SFW, you must first verify that Microsoft failover clustering is enabled (if a new installation of Windows Server 2008), and then establish a Microsoft failover cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To enable Microsoft failover clustering

- 1 In Server Manager, select **Features** in the left pane (tree view) and then click **Add Features** (link on the right side of the screen).
- 2 In the Add Features Wizard, check the **Failover Clustering** option, and click **Next**.
- 3 Click **Install**.
- 4 When the installation is complete, click **Close**.

To establish a Microsoft failover cluster

- 1 Ensure that you have met the hardware prerequisites for a failover cluster. You can run the Microsoft wizard to validate the configuration. See the Microsoft documentation for details.
- 2 Configure the shared storage and create a volume with drive letter “Q” for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends a minimum of 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster using Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 4 In the action pane, click **Create a Cluster**. The Create Cluster Wizard will start.
If this is the first time this wizard has been run, the Before You Begin page will appear. Review the information that is displayed and then click **Next**. You can hide this page on subsequent uses of the wizard; if this has been done, the first page might be the Select Servers page.
- 5 In the Select Servers panel, type the name of the first node in the Enter server name field and click **Add**. You can also use the Browse button to browse the Active Directory for the computers you want to add.
Repeat this step for the second node.
- 6 After both nodes have been added to the list of Selected Servers, click **Next**.
- 7 Based on the information on the validation warning screen, assess your hardware configuration, and select one of the options. Symantec recommends that you select **Yes**, which starts the Validate a Configuration wizard. Follow the wizard instructions.
- 8 In the Access Point for Administering the Cluster screen, in the Cluster Name field, type the name for the failover cluster. This is the name that you use to connect to and administer the cluster.
- 9 In the Address field of the network area, type the appropriate IP address and then click **Next**.
- 10 In the Confirmation screen, verify that the cluster configuration is correct, and then click **Next** to create the cluster.
- 11 Review the Summary page and then click **Finish** to close the wizard.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use these disks when you create your SFW cluster disk groups, you must first remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs. After creating the SFW cluster disk groups, you will add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running a Microsoft failover cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. Our example uses a two node configuration, so the inactive system is the second node. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft failover cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 253.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 254.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 260.

Pre-installation tasks

Perform the following pre-installation tasks:

- Moving the online groups

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.

If there is more than one resource group, you must repeat this step until all the resource groups are moved.

- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a Microsoft failover cluster configuration.

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

Note: Before you install Storage Foundation for Windows, make sure that the node is inactive.

To install the product

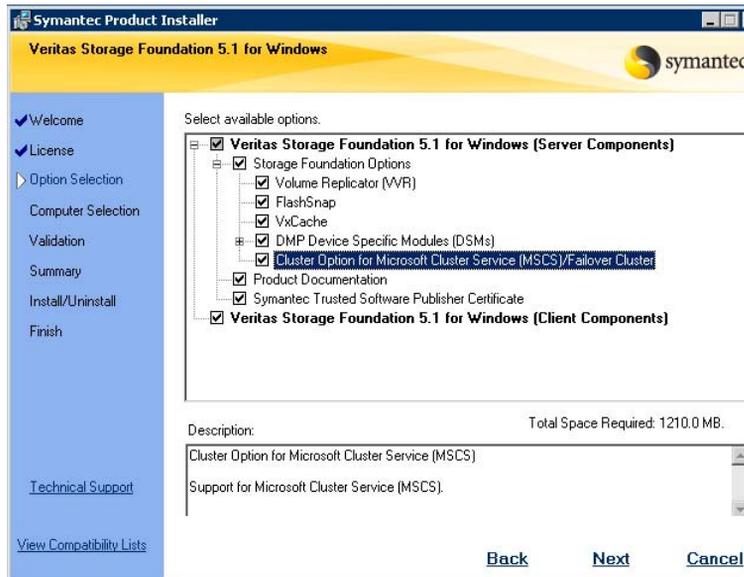
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

3 Click **Storage Foundation 5.1 for Windows**.



- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for "**I accept the terms of the license agreement,**" and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
- 8 Repeat for additional license keys.
To remove a license key, click the key to select it and click **Remove**.
To see the license key's details, click the key.
- 9 Click **Next**.

- 10 Specify the product options by selecting the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.

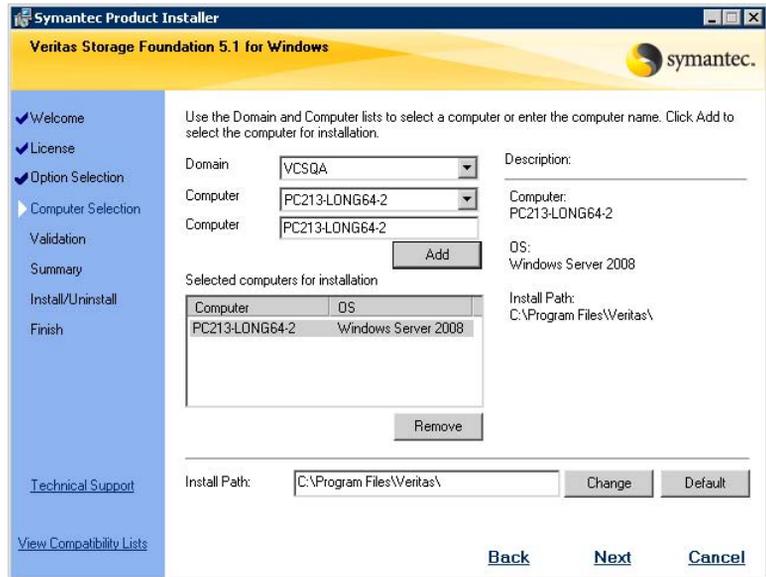


Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.

Note that under Veritas Dynamic Multi-pathing, you can select DMP Device Specific Modules (DSMs).

- 11 Verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is checked, to install the client component and click **Next**.

12 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path	Optionally, change the installation path. <ul style="list-style-type: none">■ To change the path, select a computer in the Selected computers for installation field, type the new path, and click Change.■ To restore the default path, select a computer and click Default. The default path is: C:\Program Files\Veritas For 64-bit installations, the default path is: C:\Program Files (x86)\Veritas
--------------	--

- 13 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 14 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 15 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that Microsoft clustering allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

Dynamic Multi-pathing

Additionally, if you selected the Dynamic Multi-pathing option, a warning appears:

- For DMP DSM installations—the time required to install the Veritas Dynamic Multi-pathing DSM feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during

installation. After installation, reconnect additional physical paths before rebooting the system.

- 16 Review the information and click **Install**. Click **Back** to make changes.
- 17 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.
If the installation is successful on all systems, the installation report screen appears.
If a security alert asks you to accept the Veritas driver software, click **Yes**.
- 18 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 19](#) if you are installing SFW on the local node only.
 - Proceed to [step 21](#) if you are installing SFW on local and remote systems.
- 19 To complete the installation, click **Finish**.
- 20 Click **Yes** to reboot the system and complete the installation.
- 21 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
- 22 Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
- 23 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available.
- 24 Click **Next**.
- 25 Click **Finish**.
- 26 Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW installation

Moving the online groups

You can move the resource groups from the current system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open the Failover Cluster Management tool. (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click the resource group and then click **Move this service or application to another node > Move to node [name of original node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved back to the original node.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that all the resource groups have moved back to the original system.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the Microsoft failover cluster.

See “[SFW installation tasks](#)” on page 253.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for SQL. A dynamic disk group is a collection of one or more disks that behaves as a single storage repository. Within each disk group, you can have dynamic volumes with different layouts.

Configuring disk groups and volumes involves the following tasks:

- [“Planning disk groups and volumes”](#) on page 261
- [“Creating dynamic cluster disk groups”](#) on page 263
- [“Creating dynamic volumes”](#) on page 265
- [“Managing disk group and volumes”](#) on page 269

Planning disk groups and volumes

Before installing SQL Server, you must create disk groups and volumes using the VEA console installed with SFW.

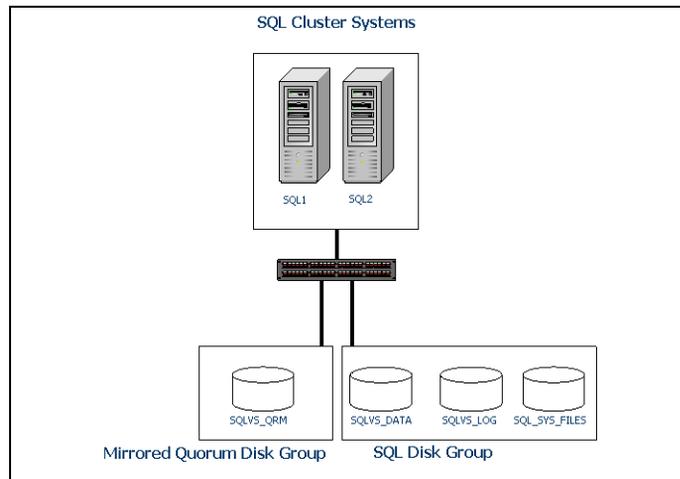
Before creating a disk group, consider:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.
- The number of disk groups and volumes that are needed for SQL Server
The number of disk groups for SQL depends on the planned organization of the data. The application program files need to be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage in a cluster disk group. You create at least one disk group for the system data files. You may want to create additional disk groups for user databases. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- The disk groups and volumes for the mirrored quorum resource
You will need a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. You can create the quorum disk group at the same time you create application disk groups, although it is not required for installing the application.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Typically, a SFW disk group corresponds to a SQL virtual server group. Below is a detailed view of the disk groups and volumes for SQL:

Figure 13-2 SFW disk groups and volumes for SQL virtual server SQLVS in Microsoft failover clustering setup



An example SQL disk group INST1_DG contains three volumes:

- INST1_DB1_VOL: Contains a SQL user database.
- SQLVS_LOG: Contains the transaction log.
- INST1_DATA_FILES: Contains the Microsoft SQL Server system data files.

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

You can create the quorum disk group and mirrored volumes now although they are not required for installing the application. For more details on the quorum resource, see [“Implementing a dynamic mirrored quorum resource”](#) on page 276.

Creating dynamic cluster disk groups

When the tasks described in this section are completed, you will have a dynamic cluster disk group with volumes on shared storage. The dynamic cluster disk groups will be ready to be shared between nodes in the cluster.

A dynamic cluster disk group is a collection of one or more disks that behave as a single storage repository and which can potentially be accessed by different computers. Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it will be required later during the SQL installation process.

To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on the client system and can be used to manage all the servers remotely.

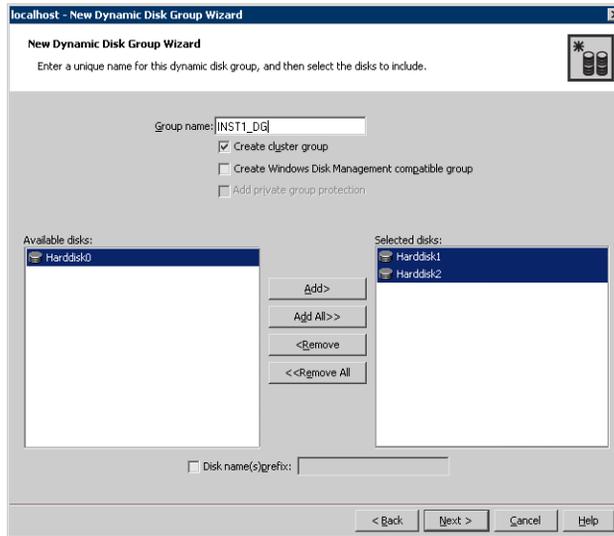
Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

Creating dynamic volumes

This section will guide you through the process of creating a volume on a dynamic disk group.

When creating a disk group to support a SQL Server 2005 solution, it is best to separate SQL data files from SQL log files and place them on separate volumes.

Repeat the procedure below to create the following volumes on the first node of the cluster:

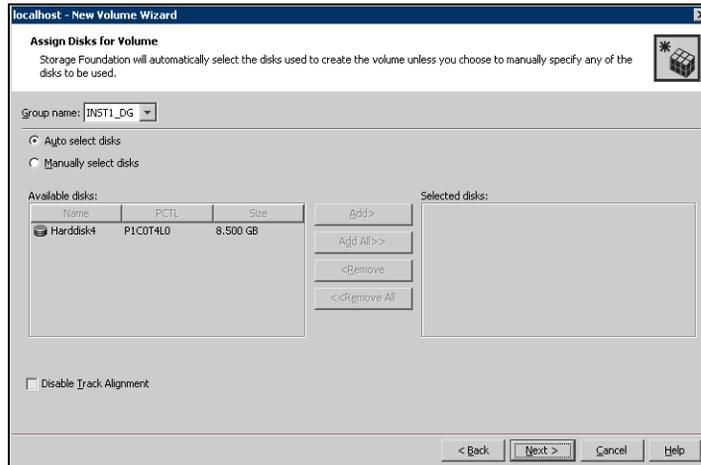
- INST1_DATA_FILES: For storing the SQL system databases.
- INST1_DB1_VOL: For storing the user database.
- SQLVS_LOG: For storing the user database log.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

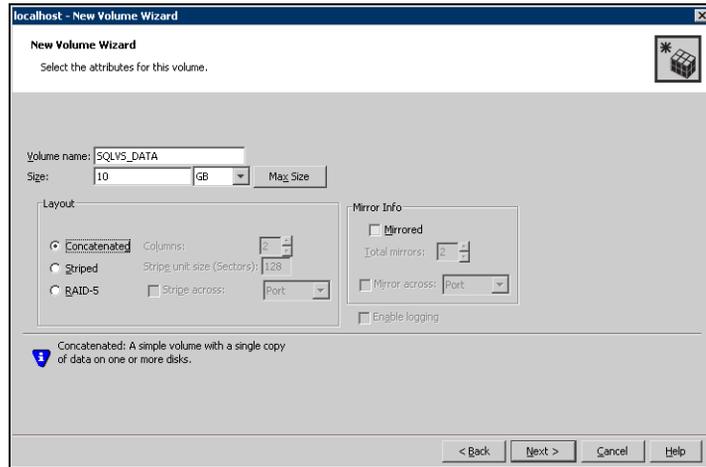


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

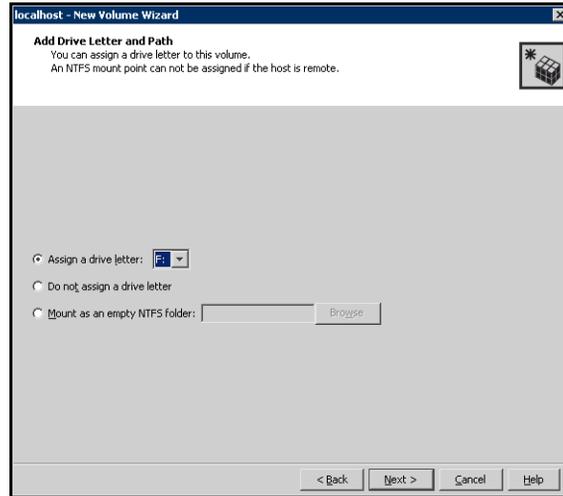
- 8 Click **Next**.

9 Specify the parameters of the volume.



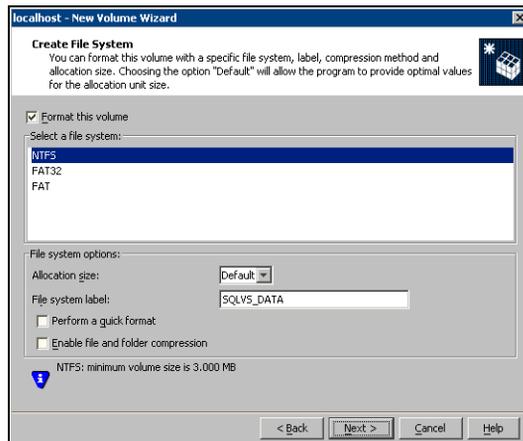
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
 - 14 Repeat these steps to create additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk group and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.

- To assign a drive letter
Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Creating the SQL Server virtual server group

Before installing SQL Server, you must create the SQL server cluster group.

In addition, SQL Server installation requires a separate volume, INST1_DATA_FILES on which the system database files will be placed. Before installation, you must add a Volume Manager Disk Group resource for the disk group that contains this volume. You add the Volume Manager Disk Group resource to the SQL Server cluster group after it is created.

Note: Before creating the resources, start the cluster service on all the nodes in the cluster.

To create an SQL Server cluster group

- 1 Launch Failover Cluster Management by selecting **Start > Administrative Tools > Failover Cluster Management**. Make sure you are connected to the required cluster.
- 2 Create a new group by selecting the **Services and Applications** node from the tree that is displayed in the left hand pane. Right-click and select **More Actions > Create Empty Service or Application**. An empty group named New service or application is created.
- 3 Specify a name for the group by right-clicking it and selecting **Rename** from the drop down menu.
- 4 Type the name of the new group (for example, SQL_GROUP) in the **Name** field.
SQL installation requires that you add the You can now add the Volume Manager Disk Group resource to the application group.

Creating a Volume Manager Disk Group resource for the application

You add a Volume Manager Disk Group resource for each SFW disk group that you created for the SQL Server application. When configuring the resource, you must type the exact name of the existing SFW disk group.

You create the Volume Manager Disk Group resource before SQL installation. SQL installation will automatically create the appropriate dependency for the SQL Server resource on the disk group resource.

To create a Volume Manager Disk Group resource for the application

- 1 If Failover Cluster Management is already open, then proceed to Step 2.

To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.

- 2 In the left pane of Failover Cluster Management, right-click the SQL Server cluster group (for example, SQL_GROUP) and select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 3 In the center panel under Disk Drives, double-click **New Volume Manager Disk Group** to open its Properties dialog box.
- 4 On the General tab of the Properties dialog box, type a name for the resource.
For example, type SQL_DG_RES.
- 5 On the Properties tab, in the Disk Group Name field, type the name of the disk group you previously created for the application (for example, INST1_DG), and click **OK** to close the dialog box.
- 6 Right-click the newly named resource and select **Bring this resource online**.
- 7 If you created more than one disk group for the application, repeat this procedure to add another Volume Manager Disk Group resource for another disk group.

Installing SQL Server 2005

This section provides some useful tips on how to install SQL Server 2005. As you progress through the installation, use these guidelines to create an installation that will function properly in your environment.

Note: Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

Before you proceed with installing SQL Server, note the following points:

- Verify the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters) See “[Managing disk group and volumes](#)” on page 269.
- The Setup program automatically installs a new, separate instance of SQL Server binaries on the local disk of each server in the cluster. The binaries are installed in exactly the same path on each cluster node, so it is important to ensure that each node has a local drive letter in common with all the other nodes in the cluster.
- The Setup program also installs the system databases on the specified cluster (shared) disk. System databases must be on a clustered disk so that they can be shared between the nodes (and failed over when necessary),

because these databases contain specific user login and database object information that must be the same for each node. The virtual server name will allow users access to the online node.

To install SQL Server 2005

- 1 Begin the SQL Server 2005 installation, following the instructions from Microsoft. To begin the installation, navigate to the installation directory and launch **splash.hta**.
- 2 Review the hardware and software requirements for SQL Server 2005.
- 3 Under the **Install** section, select **Server components, tools, Books Online, and samples**.
- 4 Continue with the installation, following the instructions from Microsoft. Complete the SQL Server Component Update, System Configuration Check, and Registration Information pages.

- 5 In the Components to Install dialog box, select the **SQL Server Database Services**.

To cluster SQL Server using Microsoft failover clustering, you must select the **Create a SQL Server failover cluster** option.

Select the optional components:

- Analysis Service. If this option is selected, the option **Create an Analysis Server failover cluster** must also be selected.
- Notification Services
- Integration Services
- Workstation Components

Click the **Advanced** option.

- 6 In the **Feature Selection** dialog box, specify the path for SQL Server data files and other services.
- 7 Expand **SQL Server Database Services** and select **Data Files**.
- 8 Select **Browse** to reset the installation path.
- 9 Set the installation path in the Change Folders dialog box to the drive letter and location of the volume created for the SQL Server system data files (INST1_DATA_FILES). Allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. This must be the same as the path on all nodes.
- 10 If you selected the **Analysis Services** option above, expand **Analysis Services**, select **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. This must be the same as the path on all nodes.

- 11 Click **Next**.
- 12 In the **Instance Name** dialog box, enter an instance name or accept the default. Click **Next**.

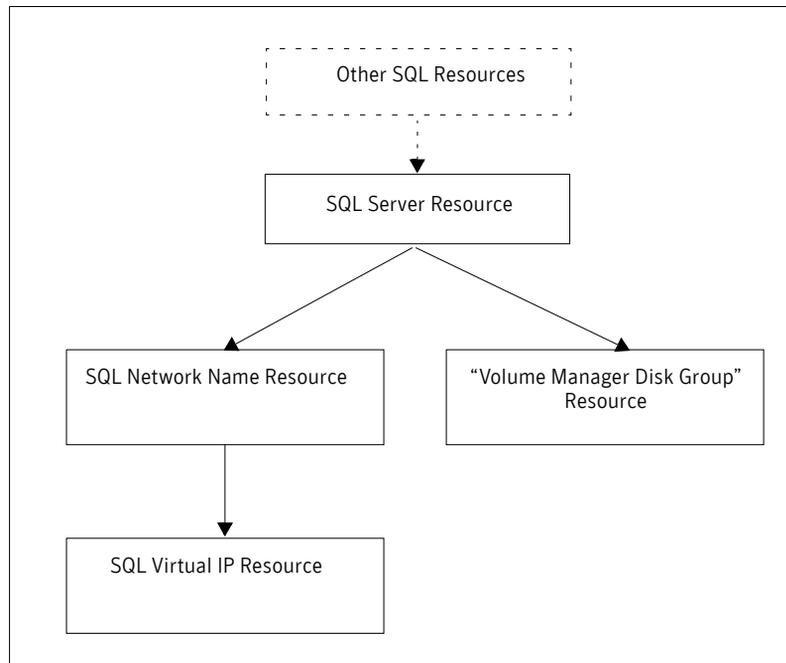
Only one default instance is allowed per cluster.

Use the same instance name when installing SQL Server 2005 on the first node and on all failover nodes.
- 13 In the **Virtual Server Name** page, specify a unique name for the virtual SQL server, for example, *SQLVS*. Make a note of this name as you will need to use the same name when installing on the secondary cluster nodes. Click **Next**.
- 14 In the **Virtual Server Configuration** page, configure the virtual server. Select the appropriate public network that you have configured in the **Network to use** field. By default, the configured public network will be selected. However, if there are more than one network cards configured for public network then you can select the appropriate one from the list.
- 15 Specify the IP address that is intended for the SQL virtual server in the **IP Address** field.
- 16 The appropriate **Network address** and **Network subnet** will be displayed. Click **Next**.
- 17 In the **Cluster Group Selection** page, specify the cluster group with the logical disk from the shared disk array that will be used for the SQL Server 2005 system database files. You can also specify a custom path in the **Data Files** field.

Click **Next**.
- 18 In the **Cluster Node Configuration** page, specify the nodes in the cluster for SQL failover by selecting them from **Available nodes**, and clicking **Add** to add them to the **Selected Nodes** box. Once all the required nodes are in the **Selected Nodes** box, click **Next**.
- 19 In the **Remote Account Information** page, specify the administrative password that is valid on all the nodes. Click **Next**.
 - In the **Service Account** page, specify the type of account and information.
 - Select the type of service account. Select **Use the built-in System account** if you do not want to replicate data. Otherwise, select **Use a domain user account**.
 - If you chose to run the service in the context of a domain user, specify the information for the user.
 - Click **Next**.

- 20 In the **Domain Groups for Clustered Services** page, use the browse button at the right to select a **DomainName** and **GroupName** for each of the selected SQL Server options. Click **Next**.
- 21 In the **Authentication Mode** page, select **Mixed Mode** (recommended option), and specify your password.
- 22 Follow the wizard page instructions to complete the SQL installation on all the nodes of the cluster.
 Once SQL is installed, the SQL Server Resource with dependencies on the SQL Network Name and the Volume Manager Disk Group resource is created. The following dependency graph indicates the dependencies that are established.

Dependency graph after the SQL installation is completed



Verifying SQL Server 2005 installation

Click **Start** > **Programs** > **Microsoft SQL Server 2005**. Select **SQL Server Management Studio** from the menu that appears. Once the **SQL Server Management Studio** window opens, you have verified that SQL Server 2005 is installed.

Implementing a dynamic mirrored quorum resource

One of the key advantages of using SFW with Microsoft clustering is the ability to create a mirrored quorum resource that adds fault tolerance to the quorum and protects the cluster. Complete the following tasks:

- Create a dynamic cluster disk group for the quorum resource with a mirrored volume
- Create the quorum resource for the Cluster Group
- Change the quorum resource to a dynamic mirrored quorum resource.

Creating a dynamic cluster disk group and a mirrored volume for the quorum resource

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using three small disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

To create a three-way mirrored volume using the New Volume wizard

- 1 Create the cluster disk group with three small disks.
- 2 Create a volume with the three disks, in the sample this is SQLVS_QRM.
- 3 Select the **Concatenated** layout, click the **Mirrored** check box, and specify three mirrors.

For full details on a creating cluster disk groups and volumes, see “[Creating dynamic cluster disk groups](#)” on page 263.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by Microsoft clustering.

Creating the quorum resource for the cluster group

You must add the Volume Manager Disk Group resource for the quorum.

To add the Volume Manager Disk Group resource for the quorum

- 1 If Failover Cluster Management is already open, then proceed to Step 2. To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.
- 2 Verify that the cluster is online on the same node where you created the disk group.
- 3 In the left pane of Failover Cluster Management, right-click **Services and Applications** and select **More Actions > Create Empty Service or Application**.
- 4 Right-click the new group and rename it, for example `QUORUM`.
- 5 Right-click `QUORUM` and select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 6 Right-click **New Volume Manager Disk Group** in the center pane and click **Properties**.
- 7 In the General tab of the Properties dialog box, type a name for the resource in the Resource Name field, for example, `QUORUM_DG_RES`.
- 8 On the Properties tab, in the Disk Group Name field, type the name of the disk group that you previously created for the quorum, and click **OK** to close the dialog box.
- 9 Right-click the Quorum disk group resource (for example, `QUORUM_DG_RES`) in the left pane and select **Bring this resource online**.
The specified disk group resource, `QUORUM_DG_RES` resource, is created under the Quorum group (for example, `QUORUM`).

Changing the quorum resource to a dynamic mirrored quorum resource

Use the following procedure to configure the cluster quorum settings and change the quorum resource to a dynamic mirrored quorum resource.

To change the quorum to a dynamic mirrored quorum resource

- 1 In Failover Cluster Management, right-click the cluster node in the configuration tree, and select **More Actions > Configure Cluster Quorum Settings**.

The Configure Cluster Quorum Wizard opens.

- 2 Review the screen and click **Next**.
- 3 Select either the **Node and Disk Majority** or **No Majority: Disk Only** radio button, and click **Next**.
- 4 Select the storage resource that you want to assign as the disk witness for the quorum and click **Next**.
This is the Volume Manager Disk Group resource that you previously created for the quorum disk group, for example, `QUORUM_DG_RES`.
- 5 Review the information in the Confirmation screen and click **Next**.
- 6 Click **Finish** to close the wizard.

Verifying the cluster configuration

You can verify your installation by moving the cluster group between nodes to see if it fails over properly. The ultimate test of the cluster's failover capability involves shutting down the node that is currently online and bringing it back up after the cluster fails over to the other node.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Failover Cluster Management to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Failover Cluster Management tool (**Start > All Programs > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open Failover Cluster Management. Click **Start > All Programs > Administrative Tools > Failover Cluster Management** from any node in the cluster.
- 3 In Failover Cluster Management, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move the resource groups back to the original node, restart the node you shut down in [step 1](#), select the resource group, and use **Move this service or application to another node > Move to node [name of node]** to move the resource group.

Deploying SFW with MSCS and SQL Server in a campus cluster

This chapter covers the following topics:

- [Tasks for a new SQL Server installation with SFW and MSCS in a campus cluster \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the network and storage](#)
- [Establishing an MSCS cluster](#)
- [Creating the MSDTC resource](#)
- [Installing SFW](#)
- [Creating disk groups and volumes](#)
- [Implementing a dynamic quorum resource](#)
- [Setting up a group for SQL Server in MSCS](#)
- [Installing the application on the cluster nodes](#)
- [Verifying the cluster configuration](#)

Tasks for a new SQL Server installation with SFW and MSCS in a campus cluster (Windows Server 2003)

On Windows Server 2003, you can install and configure Veritas Storage Foundation for Windows with MSCS and Microsoft SQL Server 2000 or 2005. This chapter presents a campus clustering example using a two-node cluster. For information on deploying SFW with Microsoft failover clustering, which runs under Windows Server 2008, see:

[Chapter 15, “Deploying SFW with Microsoft failover clustering and SQL Server in a campus cluster” on page 329.](#)

[Table 14-1](#) lists the high-level objectives for deploying SFW with MSCS in a campus cluster, as well as the tasks within each objective:

Table 14-1 Task list for deploying SQL Server with SFW and MSCS in a campus cluster

Objectives	Tasks
“Reviewing the requirements” on page 284	<ul style="list-style-type: none">■ Verify hardware and software prerequisites.
“Reviewing the configuration” on page 288	<ul style="list-style-type: none">■ Review the configuration requirements.■ Overview of MSCS campus cluster, and recovery scenarios
“Configuring the network and storage” on page 296	<ul style="list-style-type: none">■ Install the hardware for Site A. The server and storage array are connected to the SAN. Leave the cables for the NICs unconnected, and do not yet connect the switch to site B.■ Install the hardware in the same manner for Site B.

Table 14-1 Task list for deploying SQL Server with SFW and MSCS in a campus cluster (Continued)

Objectives	Tasks
<p>“Establishing an MSCS cluster” on page 298</p>	<ul style="list-style-type: none"> ■ Install and configure the operating system and MSCS on Server A. ■ Configure the storage and create a partition for the cluster quorum disk on Site A. ■ Create the first node of the cluster on Server A. ■ Install and configure the operating system and MSCS on Server B. ■ Connect the two nodes. ■ Create the second node of the cluster on Server B. ■ Test the cluster by moving the resources to Server B. Server B becomes the active node. Do not move them back to Server A at this point.
<p>“Creating the MSDTC resource” on page 301</p>	<ul style="list-style-type: none"> ■ Create the MSDTC resource.
<p>“Installing SFW” on page 303</p>	<ul style="list-style-type: none"> ■ Install SFW on Node A (Node B active). ■ Install SFW on Node B (Node A active).
<p>“Creating disk groups and volumes” on page 312</p>	<ul style="list-style-type: none"> ■ In SFW on Node A, create two or more dynamic cluster disk groups on the storage, one or more for the application data files and one for the mirrored quorum.
<p>“Implementing a dynamic quorum resource” on page 321</p>	<ul style="list-style-type: none"> ■ If not done earlier, create a dynamic disk group for the quorum with a mirrored volume. ■ Make that disk group into a Volume Manager Disk Group type resource in the default Cluster Group. ■ Change the quorum resource to the dynamic mirrored quorum resource.
<p>“Setting up a group for SQL Server in MSCS” on page 324</p>	<ul style="list-style-type: none"> ■ Create a group within MSCS for the SQL Server application. ■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.

Table 14-1 Task list for deploying SQL Server with SFW and MSCS in a campus cluster (Continued)

Objectives	Tasks
“Installing the application on the cluster nodes” on page 325	<ul style="list-style-type: none">■ Install the application program files on the local drive of the first node.■ Install files relating to the data and logs on the shared storage.■ Move the cluster resources to the second node.■ Make sure that the volumes on the second node have the same drive letters or mount points as they had on the first node.■ Install the application on the second node.
“Verifying the cluster configuration” on page 326	<ul style="list-style-type: none">■ Verify the cluster configuration by switching service groups or shutting down an active cluster node

Reviewing the requirements

Reviewing the prerequisites and the configuration allows you to gain an overall understanding of the configuration and its requirements.

See the following topics:

- [Supported software](#)
- [System requirements](#)
- [Disk space requirements](#)

Supported software

You can check the Late Breaking News information on the Support web site for any recent updates to this list of supported software.

The following software is supported:

- Veritas Storage Foundation 5.1 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster.

For a Microsoft SQL Server 2000 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|---|
| Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required) | <ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) |
| Microsoft SQL Server 2000 (64-bit) Enterprise Edition | <ul style="list-style-type: none"> ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| Microsoft SQL Server 2000 (64-bit) Standard Edition or Enterprise Edition (SP4 required) | <ul style="list-style-type: none"> ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions) |

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|---|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
|---|--|

- | | | |
|--|---|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | ■ | Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SP1 or higher required) | ■ | Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| | ■ | Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) |
| | ■ | Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | ■ | Windows Server 2008 for 64-bit Itanium (IA64) |
| | ■ | Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |

System requirements

To deploy SFW with MSCS in a campus cluster, your system must meet the following requirements:

- One CD-ROM drive accessible to each system on which you are installing MSCS.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access the storage.
- MSCS requires at least two network adapters per system (one network adapter to connect each system to the public network and one network adapter for the private network on each system). Symantec recommends having two adapters for the private network and routing each private network adapter through a separate hub or switch to avoid single points of failure.
- Refer to application documentation to determine disk space requirements for your application.
- Each system requires 1 GB of RAM.
- The configuration requires two sites with a storage array for each site, with an equal number of disks at each site for the mirrored volumes.

- Interconnects between the clusters are required for the storage and the network.
- Systems to be clustered must be configured as part of a Windows Server 2003 domain. Each system in an MSCS cluster must be in the same domain and must be using the same operating system version.
- To install and work with the SFW and MSCS software, you must have an account with Administrator privileges. You must also have a license key to install SFW.
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported. Six network interface cards, three for each server (two each for the private network and one for the public network). You also need a static IP address for the cluster itself.

Note: To determine the approved hardware for SFW, see the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp>.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 14-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

Note: Plan for an equal number of disks on the two sites, because each disk group should contain the same number of disks on each site.

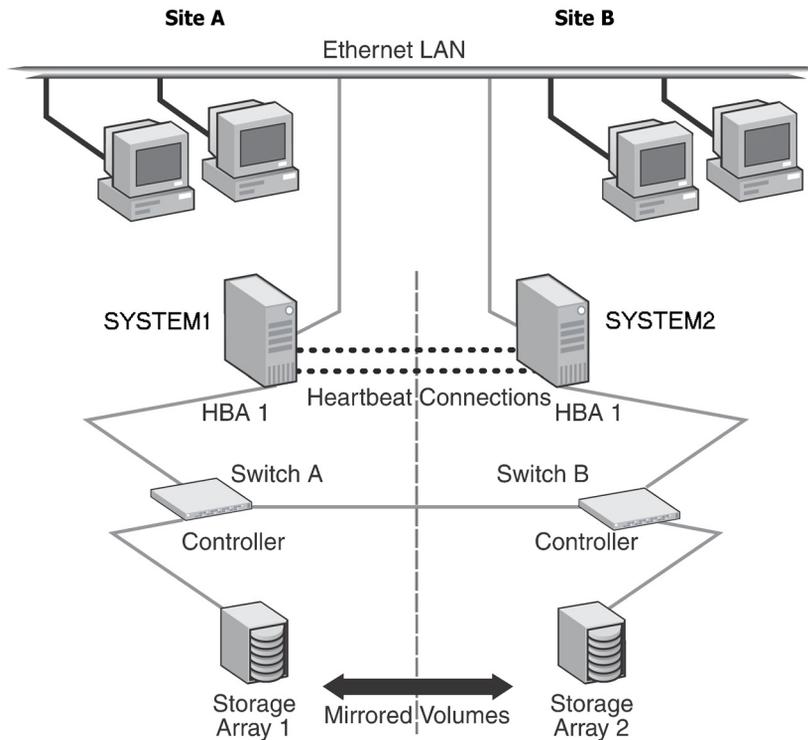
Reviewing the configuration

This configuration example describes a two-node campus cluster with each node at a separate site.

For an overview of campus clusters with MSCS or for recovery scenarios, see

- [“Overview of campus clustering with MSCS”](#) on page 289
- [“MSCS campus cluster failure scenarios”](#) on page 290

Figure 14-1 MSCS campus clustering configuration example



The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array and contains mirrored data of the storage on the other array. Each disk group should contain the same number of disks on each site for the mirrored volumes.

MSCS uses the quorum architecture, where the cluster database resides in the quorum resource. If you use MSCS for clustering, adding SFW to the configuration protects the quorum disk from being a single point of failure in

the cluster because SFW provides dynamic volumes and software mirroring of the quorum device. To avoid a single point of failure, set up the quorum as a dynamic mirrored device. This example includes the dynamic mirrored quorum and requires setting up two or more dynamic cluster disk groups in SFW—one or more cluster disk groups for the application and data and one for the dynamic mirrored quorum.

The example configuration does not include Dynamic Multi-pathing (DMP). For instructions on how to add DMP to a clustering configuration, see *Veritas Storage Foundation and High Availability Solutions, Solutions Guide*.

When you are installing SFW and MSCS together, remember the following:

- To install SFW, an MSCS cluster must be running.
Before you install SFW, you must set up the hardware and install the operating system and MSCS on all systems and establish the MSCS cluster. Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Therefore, use a “rolling install” procedure to install SFW first on the inactive cluster node. Then move the cluster resources to the other node and install on the now inactive node.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

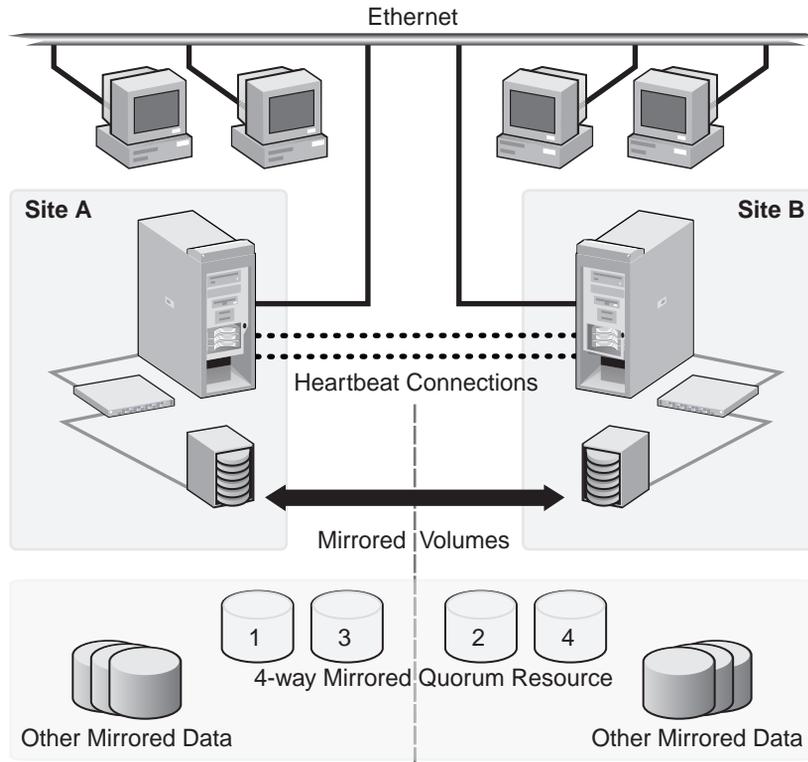
- Using SFW also offers other advantages over using MSCS alone. SFW lets you add fault tolerance to your data volumes. Mirroring of log volumes is recommended, and a mirrored striped RAID layout is recommended for your data volumes. SFW also offers multiple disk groups, multiple mirrors, capacity management and Automatic Volume Growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, Dynamic Multi-pathing, and enhanced snapshot capabilities with FlashSnap.

Overview of campus clustering with MSCS

Figure 14-2 on page 290 shows an MSCS campus cluster configuration with mirrored storage across clusters and a mirrored quorum resource. The 4-way mirrored quorum has an extra set of mirrors for added redundancy. Although a campus cluster setup with MSCS can work without Storage Foundation for Windows, SFW provides key advantages over using MSCS alone. Through a dynamic mirrored volume that functions on multiple disks across multiple sites,

SFW protects the quorum resource in the cluster from being the single point of failure in the cluster.

Figure 14-2 Typical MSCS campus clustering configuration



Most customers use hardware RAID to protect the quorum disk, but that does not work when a natural disaster takes down the primary node and its attached storage. If the quorum resource is lost to the cluster, the cluster fails, because none of the cluster servers can gain control of the quorum resource and ultimately the cluster. MSCS alone cannot provide fault tolerance to the quorum disk.

MSCS campus cluster failure scenarios

This section focuses on the failure and recovery scenarios with an MSCS campus cluster and SFW installed.

For information about the quorum resource and arbitration in MSCS, see

“[MSCS quorum and quorum arbitration](#)” on page 294.

Table 14-3 lists failure situations and the outcomes that occur:

Table 14-3 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
Application fault May mean the services stopped for an application, a NIC failed, or a database table went offline.	Failover	If the services stop for an application failure, the application automatically fails over to the other site.
Server failure (Site A) May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.	Failover	Assuming a two-node cluster pair, failing a single node results in a cluster failover. There will be a temporary service interruption for cluster resources that are moved from the failed node to the remaining live node.
Server failure (Site B) May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.	No interruption of service.	Failure of the passive site (Site B) does not interrupt service to the active site (Site A).
Partial SAN network failure May mean that SAN fiber channel cables were disconnected to Site A or Site B Storage.	No interruption of service.	Assuming that each of the cluster nodes has some type of Dynamic Multi-pathing (DMP) solution, removing one SAN fiber cable from a single cluster node should not effect any cluster resources running on that node, because the underlying DMP solution should seamlessly handle the SAN fiber path failover.
Private IP Heartbeat Network Failure May mean that the private NICs or the connecting network cables failed.	No interruption of service.	With the standard two-NIC configuration for a cluster node, one NIC for the public cluster network and one NIC for the private heartbeat network, disabling the NIC for the private heartbeat network should not effect the cluster software and the cluster resources, because the cluster software will simply route the heartbeat packets through the public network.

Table 14-3 List of failure situations and possible outcomes (Continued)

Failure Situation	Outcome	Comments
<p>Public IP Network Failure May mean that the public NIC or LAN network has failed.</p>	<p>Failover. Mirroring continues.</p>	<p>When the public NIC on the active node, or public LAN fails, clients cannot access the active node, and failover occurs.</p>
<p>Public and Private IP or Network Failure May mean that the LAN network, including both private and public NIC connections, has failed.</p>	<p>No interruption of service. No Public LAN access. Mirroring continues.</p>	<p>The site that owned the quorum resource right before the “network partition” remains as owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource.</p>
<p>Lose Network Connection (SAN & LAN), failing both heartbeat and connection to storage May mean that all network and SAN connections are severed, for example if a single pipe is used between buildings for the Ethernet and storage.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>The node/site that owned the quorum resource right before the “network partition” remains as owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource. By default MSCS clussvc service will try to auto-start every minute, so after LAN/SAN communication has been re-established, MSCS clussvc will auto-start and will be able to re-join the existing cluster.</p>
<p>Storage Array failure on Site A, or on Site B May mean that a power cord was unplugged, or a storage array failure caused the array to stop responding.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>The campus cluster is divided equally between two sites with one array at each site. Completely failing one storage array should not effect on the cluster or any cluster resources that are currently online. However, you will not be able to move any cluster resources between nodes after this storage failure, because neither node will be able to obtain a majority of disks within the cluster disk group.</p>
<p>Site A failure (power) Means that all access to site A, including server and storage, is lost.</p>	<p>Manual failover.</p>	<p>If the failed site contains the cluster node that owned the quorum resource, then the overall cluster would be offline and cannot be online on the remaining live site without manual intervention.</p>

Table 14-3 List of failure situations and possible outcomes (Continued)

Failure Situation	Outcome	Comments
Site B failure (power) Means that all access to site B, including server and storage, is lost.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	If the failed site did not contain the cluster node that owned the quorum resource, then the cluster would still be alive with whatever cluster resources that were online on that node right before the site failure.

Dealing with a failover situation

In summary, the site scenarios that can occur when there is a cluster server failure include the following:

- If the site not owning the quorum volume and the cluster goes offline, the quorum and data volumes stay online at the other site and other cluster resources stay online or move to that site. Storage Foundation for Windows lets the owning cluster node remain online with 50% ownership of the disks in the quorum group.
- If the site owning the quorum volume goes offline, the remaining site cannot gain control of the quorum volume because it cannot reserve a majority of disks in the quorum group. This is a safeguard to prevent multiple nodes from onlining members of a cluster disk group to which they have access.

Caution: Manual failover of a cluster between two sites should be performed only after coordination between the two sites to ensure that the primary server has in fact failed. If the primary server is still active and you manually import a cluster disk group containing the MSCS quorum to the secondary (failover) server, a split-brain situation occurs. There may be data loss if the split-brain situation occurs because each plex of the mirrored volume may be updated independently when the same disk group is imported on both nodes.

MSCS quorum and quorum arbitration

This section explains the quorum and quorum arbitration in MSCS.

Quorum

The quorum resource maintains the cluster database, as well as critical recovery information, in a recovery log. The quorum resource must be available to all nodes through a SCSI or Fibre Channel bus. With MSCS alone, the quorum disk must be located on a single physical disk. However, with SFW, the quorum disk can be a mirrored volume that spans multiple disks and cluster nodes.

The quorum resource also determines ownership of the cluster. When a node that is controlling the cluster goes offline, other nodes use a challenge/defense protocol to determine which node can have control of the quorum resource and the cluster.

Cluster ownership of the quorum resource

The MSCS challenge/defense protocol uses a low-level bus reset of the SCSI buses between the machines to attempt to gain control of the quorum resource.

After a SCSI bus reset, the reservation that each server had been holding on the quorum disk is lost. Each server has about 10 seconds to re-establish that reservation, which would in turn let the other servers know that it is still functioning, even though the other servers would not necessarily be able to communicate with it.

If the active cluster server does not re-establish the SCSI reservation on the quorum resource within the time limit, the applications that were on the server transfer to the server that establishes the SCSI reservation first. The new server servicing the application may now be a bit slower, but clients still get their applications serviced. The IP (Internet Protocol) address and network names move, applications are reconstituted according to the defined dependencies, and clients are still serviced, without any question as to the state of the cluster.

The challenge/defense protocol is more complex when the quorum device is a volume in a Storage Foundation for Windows disk group. For a server to take ownership of the disk group containing the cluster quorum device, SFW on that server must successfully import the disk group, obtaining SCSI reservations on more than half of its disks. Because a campus cluster configuration has an even number of disks on each site, failover cannot occur automatically. After a site failure, you must use the manual CLI command `vxclus enable` to bring the cluster disk groups online on the secondary node.

The vxclus utility

Storage Foundation for Windows provides the `vxclus` command line utility to allow forcing a failover to the secondary site. The command `vxclus enable` creates an entry in the Registry that enables the cluster disk group to be brought online on a node with a minority of the disks. After you run `vxclus enable`, you can bring the disk group resource online in MSCS Cluster Administrator. After the cluster disk group is brought online, the `vxclus` functionality is disabled.

To bring a cluster online on a node with a minority of the cluster disks

- 1 Use the following `vxclus` command for each disk group on your cluster node:

```
vxclus enable -g<DynamicDiskGroupName>
```

You are asked to confirm the use of this command.

Caution: When you bring a cluster disk group online with a minority of cluster disks, make sure that a majority of the disk group disks are NOT online on any other cluster node before (and after) onlining the disk group. If a majority of disk group disks are online on another node, data can be corrupted.

- 2 If the cluster service has stopped because of a dynamic quorum resource failure, start the cluster service (`clusvc`).
- 3 Use MSCS Cluster Administrator to bring the cluster disk groups online.

For more information on the `vxclus` utility, see the “Command Line Interface” chapter of the *Storage Foundation Administrator’s Guide*. The `vxclus` utility also provides support for booting from a SAN, but you must have a hardware storage array that supports the capability.

Configuring the network and storage

Use the following procedures to configure the storage hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent MSCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 5 In the Public Status dialog box, on the General tab, click **Properties**.
- 6 In the Public Properties dialog box, on the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.

- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Establishing an MSCS cluster

Before you install SFW, you must install the operating system along with MSCS and then establish an MSCS cluster. After setting up the cluster under MSCS, then you can install SFW and add SFW support with SFW disk groups and volumes.

Note: The steps outlined in this section are general and do not contain specific details. Refer to Microsoft documentation for more complete information.

The tasks for installing the cluster are:

- [“Installing and configuring the operating system and MSCS on Server A”](#) on page 298
- [“Configuring the shared storage and creating a partition for the Cluster quorum disk”](#) on page 299
- [“Creating the first node of the cluster on Server A”](#) on page 299
- [“Installing and configuring the operating system and MSCS on Server B”](#) on page 299
- [“Connecting the two nodes”](#) on page 300
- [“Creating the second node of the cluster on Server B”](#) on page 300
- [“Verifying the cluster configuration”](#) on page 300

Further descriptions of these tasks follow.

Installing and configuring the operating system and MSCS on Server A

This topic summarizes the steps for installing the operating system and configuring the network settings for Server A.

To install and configure the operating system and MSCS on Server A

- 1 Install the Windows Server 2003 operating system on Server A. MSCS is installed as part of the operating system.
- 2 Use the Internet Protocol (TCP/IP) window to identify the static Server A network addresses for the public and private networks in the cluster.
- 3 Make sure a domain is set up that can be used by the cluster nodes, which must be members of the same domain.

- 4 Select **Administrative Tools > Active Directory > Users and Computers** and set up a cluster account for the cluster. Microsoft recommends having a separate user account under which the cluster can run.

Configuring the shared storage and creating a partition for the Cluster quorum disk

Configuring the shared storage and creating a partition for the cluster quorum disk, consists of the following tasks:

- Configure the disks for the storage array attached to Server A.
- Use **Disk Management** to create a partition for the cluster quorum disk on a basic disk that will be used as the quorum disk when the first node of the cluster is created.

Microsoft recommends 500 MB as the partition size and includes the entire disk as a cluster resource.

Creating the first node of the cluster on Server A

Create the first node of the cluster on Server A. Refer to the Microsoft documentation for details.

After you establish the cluster on Server A, make sure that you can see the storage array's disks from Server A.

Installing and configuring the operating system and MSCS on Server B

Repeat the same installation steps for Server B as you used for Server A.

See [“Installing and configuring the operating system and MSCS on Server A”](#) on page 298.

Connecting the two nodes

Make the necessary connections between the two sites. The cluster is already active on Server A, so now MSCS controls the cluster storage on Server A, and the operating system cannot access both nodes of the storage at the same time.

To connect the two nodes

- 1 Connect the corresponding cables between the three network cards on the two sites.
- 2 Connect the two switches at the two sites through the storage interconnect.
- 3 Test the connectivity between the two sites, by doing the following:
 - Test the IP addresses of all the network adapter cards in the cluster.
 - Bring up the command window and type `ping ipaddress`, where the *ipaddress* is the corresponding network adapter in the other node.

Creating the second node of the cluster on Server B

Create the second node of the cluster on Server B. Refer to the Microsoft documentation for details.

Verifying the cluster configuration

After the configuration is complete, use the following procedure to verify failover.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.

- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

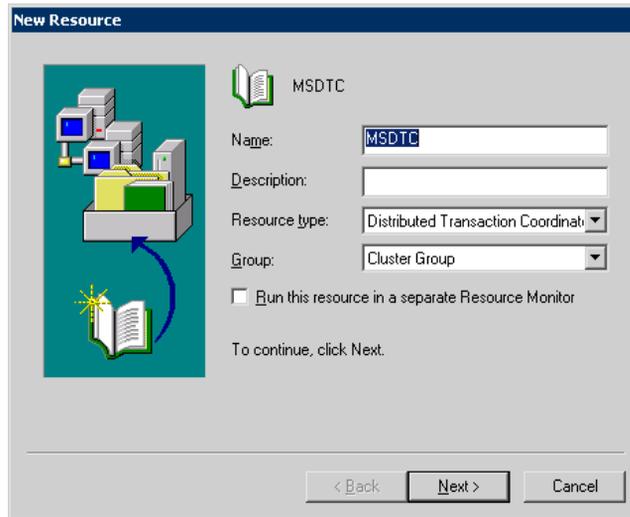
Creating the MSDTC resource

Before you install SQL Server, create the MSDTC resource. You can create this resource now or just before installing SQL.

To create the MSDTC resource

- 1 In Cluster Administrator (**Start > Administrative Tools > Cluster Administrator**), right-click **Cluster Group**, click **New**, and click **Resource**.

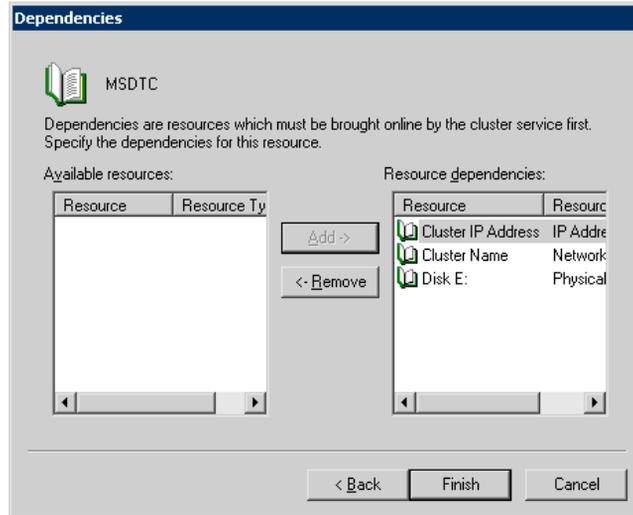
- 2 In the New Resource dialog box, specify the following options and then click **Next**.



Name	Type a name for the MSDTC resource.
Description	If necessary, type a description of the resource
Resource type	Click Distributed Transaction Coordinator .

- 3 In the Possible Owners dialog box, all the nodes in the cluster are listed as possible owners. Click **Next**.

- 4 In the Dependencies dialog box, select the cluster IP address, cluster name, and physical disk resources from the Available Resources list, add them to the Resource dependencies list, and click **Finish**.



- 5 Click **OK**.
- 6 In the left pane, expand the Groups icon.
- 7 Click **Cluster Group**.
- 8 Right-click **Bring Online**.
The state changes to online.

Installing SFW

This section assumes you are running an MSCS cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the MSCS cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 304.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 306.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 311.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
- Moving the online groups

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The table below describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 14-4 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed

Table 14-4 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.

- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a MSCS configuration.

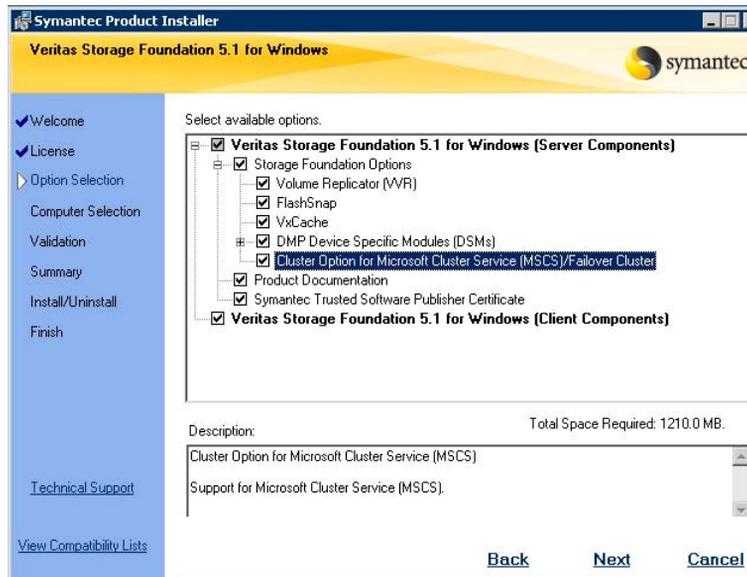
To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.
- 3 Click **Storage Foundation 5.1 for Windows**.



- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.

- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **“I accept the terms of the license agreement,”** and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
- 8 Repeat for additional license keys.
 To remove a license key, click the key to select it and click **Remove**.
 To see the license key’s details, click the key.
- 9 Click **Next**.
- 10 Specify the product options by selecting the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.

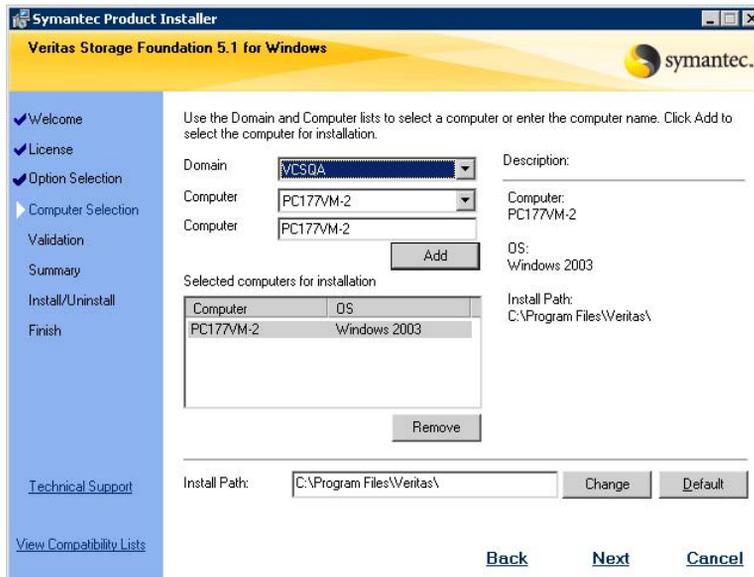


Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.

Note that under Veritas Dynamic Multi-pathing, you can select DMP Device Specific Modules (DSMs).

- 11 Click **Next**.
- 12 Verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is checked, to install the client component and click **Next**.

13 Select the domain and the computers for the installation and click **Next**.



- Domain** Select a domain from the list.
Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.
- Computer** To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.
To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.
Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 14 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 15 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 16 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that Microsoft clustering allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

Dynamic Multi-pathing

Additionally, if you selected the Dynamic Multi-pathing option, a warning appears:

- For DMP installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.

- For DMP DSM installations—the time required to install the Veritas Dynamic Multi-pathing DSM feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 17 Review the information and click **Install**. Click **Back** to make changes.
 - 18 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.
If the installation is successful on all systems, the installation report screen appears.
If a security alert asks you to accept the Veritas driver software, click **Yes**. This alert appears if your local computer has its driver signing options set to Warn. If your local computer has its driver signing options set to Block, installation fails.
 - 19 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 20](#) if you are installing SFW on the local node only.
 - Proceed to [step 22](#) if you are installing SFW on local and remote systems.
 - 20 To complete the installation, click **Finish**.
 - 21 Click **Yes** to reboot the system and complete the installation.
 - 22 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
 - 23 Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
 - 24 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available.
 - 25 Click **Next**.
 - 26 Click **Finish**.
 - 27 Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW Installation
- Resetting the driver signing options

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.
- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 304.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Creating disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and dynamic volumes on the cluster storage, which in a campus cluster consists of at least two storage arrays.

Before you create disk groups and volumes, consider the following items:

- The type of volume configurations that are required
- The number of LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs which depend on the traffic load
- The disk groups and number of disks on each site

Note: For campus clusters, each disk group must contain an equal number of disks on each site.

- Types of volumes required and location of the plex of each volume in the storage array

Note: Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

Create two or more dynamic cluster disk groups on the storage—one or more for the application data files and one for the mirrored quorum.

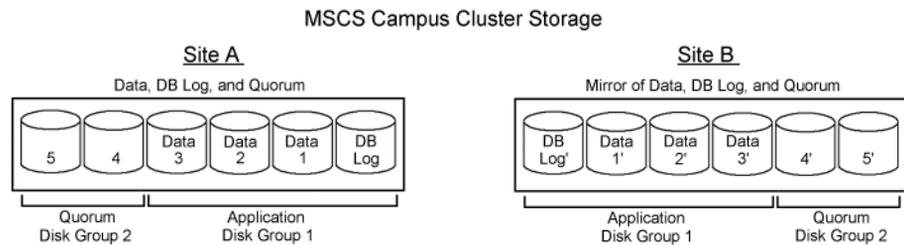
For the Microsoft SQL Server application data files, you could create a separate disk group for each database. It is best to separate data files from log files and place them in separate volumes. For example, you might create a SQL disk group, INST1_DG, containing three volumes:

- INST1_DB1_VOL: Contains the SQL database.
- INST1_DB1_LOG: Contains the transaction log.
- INST1_DATA_FILES: Contains Microsoft SQL Server system data files.

Figure 14-3 shows a typical MSCS campus cluster setup of disks. This example has only one application disk group that spans the storage arrays at both sites. The data and database log on Site A are mirrored to Site B. Each mirrored volume can have more than two disks, but must have an even number, such as four. All the application data could be in one large mirrored volume with multiple disks, but the same number of disks are required on both sites for the mirroring. It is recommended that the log volumes be on separate disks from the data.

In [Figure 14-3](#), a four-way mirror for the quorum volume provides additional redundancy. The minimum configuration is a two-way mirror. If possible, use small disks for the quorum volume. Microsoft recommends 500 MB for the quorum volume.

Figure 14-3 MSCS campus cluster disks and disk groups example



Configuring the disks and volumes

Ensure that each disk group contains an equal number of disks on each site, and that each volume is a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

See the following sections:

- [“Creating a dynamic \(cluster\) disk group”](#) on page 314
- [“Creating a volume”](#) on page 316

Considerations when creating new volumes

- For campus clusters, when you create a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.
- Symantec recommends using either simple mirrored (concatenated) or striped mirrored for the new volumes. Striped mirrored gives you better performance compared to concatenated.

When selecting striped mirrored, select two columns to stripe one enclosure that is mirrored to the second enclosure.

- You cannot select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.
- Replicating the system databases is not required or recommended. Therefore, you can place the system databases on volumes that are a part of the disk group but are not being replicated.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and, if prompted, select a profile.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. If prompted, provide the user name, password, and domain.
- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.
The internal names for the disks which the current system can access for available storage are displayed, with names **Harddisk1**, **Harddisk2**, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a dynamic (cluster) disk group

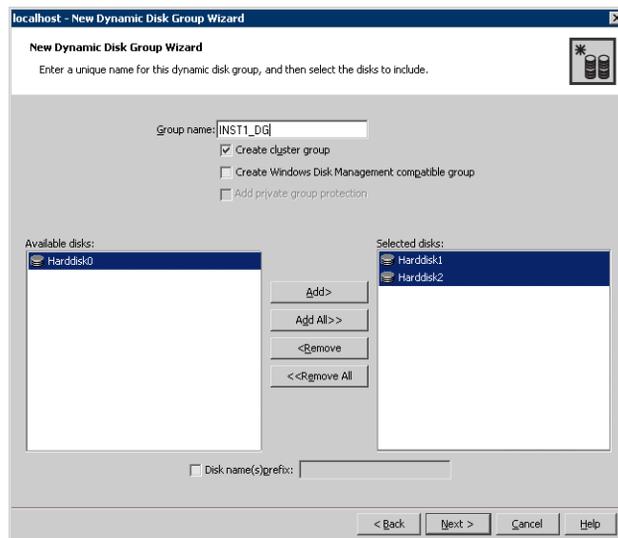
Use the following procedure to create a dynamic cluster disk group.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.

- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
 Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

■ Click **Next**.

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the new disk group.

Proceed to create the appropriate volumes on each disk.

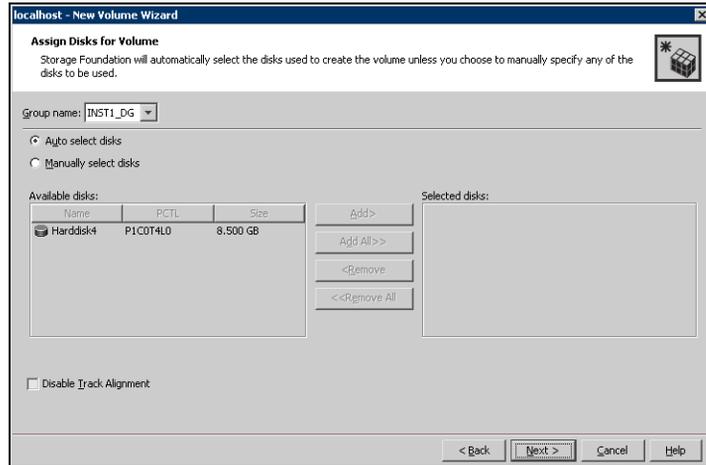
Creating a volume

Use the following procedure to create dynamic volumes.

To create dynamic volumes

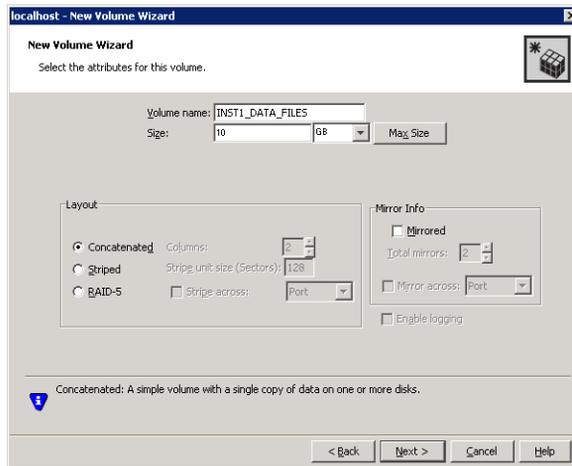
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume; make sure the appropriate disk group name appears in the Group name drop-down list.



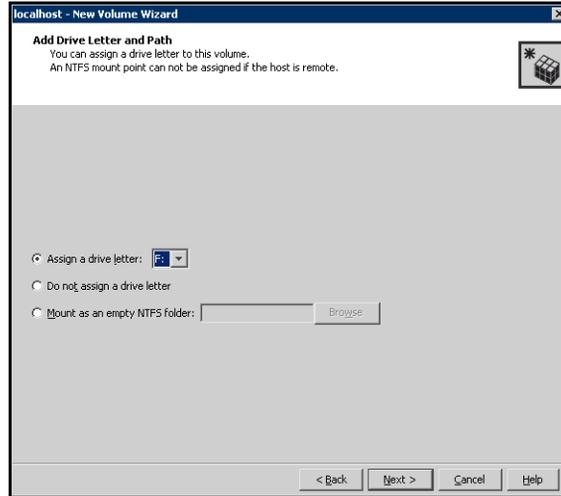
- 7 Select auto or manual disk selection and enable or disable track alignment.
 - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
 - To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.

9 Specify the parameters of the volume.

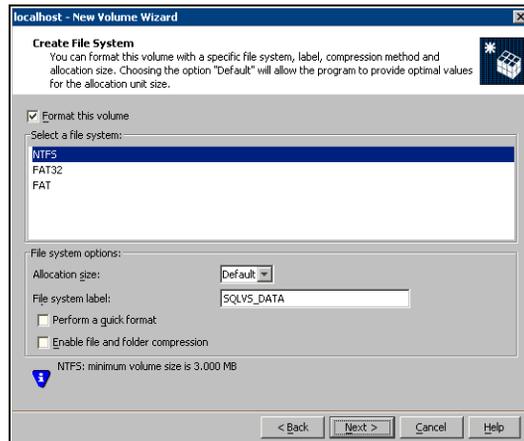


- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - Select a volume layout type. For campus clusters, select either **Concatenated** or **Striped**. Since campus clusters are mirrored volumes, you also must select the **Mirrored** checkbox.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
 - In the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
 - Verify that **Enable logging** is not selected and click **Next**.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.
- 12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and select **NTFS**.
- Select an allocation size or accept the Default.

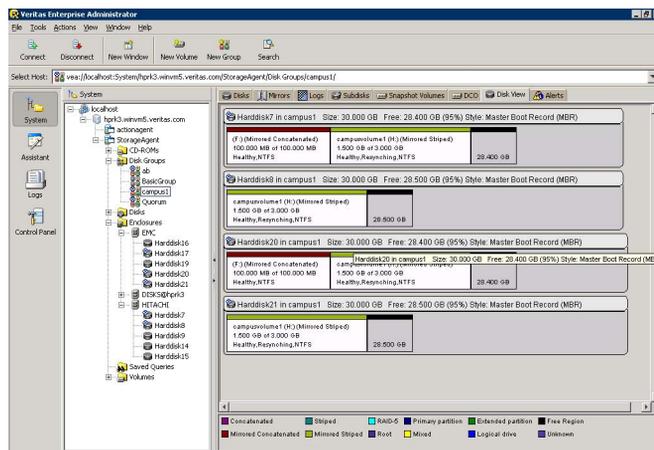
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance. Click **Next**.

13 Click **Finish** to create the new volume.

14 Repeat these steps to create additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Figure 14-4 View of disks with volumes in VEA Console



Implementing a dynamic quorum resource

One of the key advantages of using SFW with MSCS is that you can create a mirrored quorum resource that adds fault tolerance to the quorum, thus protecting the cluster from failure if the disk that the quorum is on fails. In the following procedure, you transfer the cluster's quorum resource from a physical disk resource to a mirrored dynamic quorum resource. The tasks for creating a mirrored quorum resource are:

- [“Creating a dynamic cluster disk group for the quorum, mirrored”](#) on page 321
- [“Making the quorum cluster disk group an MSCS resource”](#) on page 322
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 323

Creating a dynamic cluster disk group for the quorum, mirrored

If you have not already done so, use SFW to create a dynamic disk group for the quorum. The minimum number of disks for the mirrored quorum is two disks. Symantec recommends using four small disks for the mirrored quorum for additional redundancy.

If possible, use small disks, because the disk group will only be used for the quorum volume, which Microsoft recommends to be 500 MB. To create a four-way mirrored volume in the New Volume wizard, select the **Concatenated** layout, click the **Mirrored** checkbox, and specify four mirrors. For full details on creating cluster disk groups and volumes, see:

[“Creating disk groups and volumes”](#) on page 312.

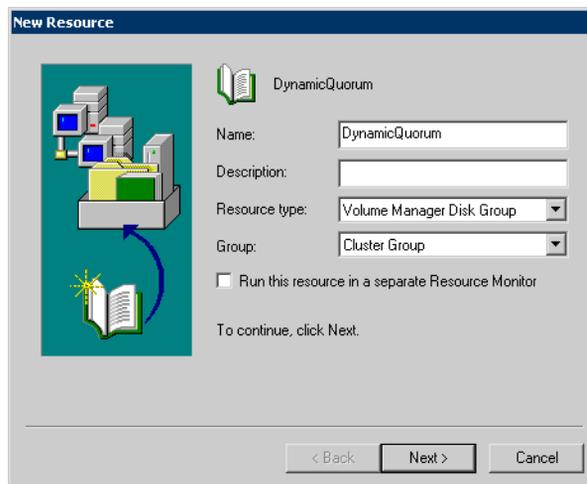
Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by MSCS.

Making the quorum cluster disk group an MSCS resource

The dynamic cluster disk group that you prepared for the quorum needs to be added as a resource to the default Cluster Group in MSCS. Complete this step now if you have not done it earlier.

To make the quorum disk group an MSCS resource

- 1 Verify that the Cluster Group is online on the same node where you created the cluster disk group for the quorum.
- 2 Right-click on that disk group and select **New > Resource**. The New Resource window appears.

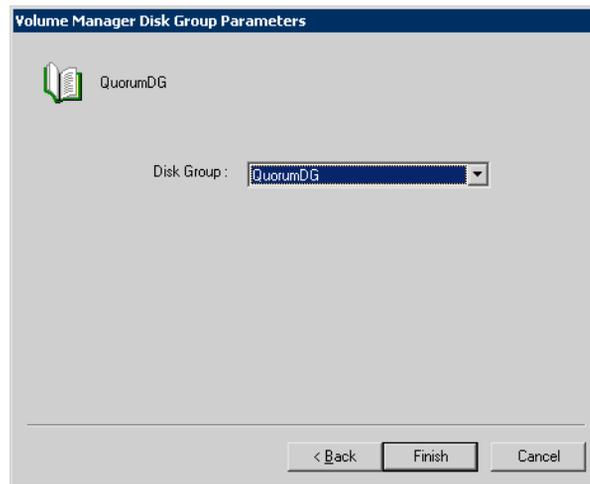


- 3 On the New Resource, window, do the following:
 - Specify a name for the disk group resource in the **Name** field, such as "QuorumDG."
 - If necessary, you can add a description of the resource in the **Description** field.
 - Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource type** field drop-down list.

Note: The resource name has not been changed to Storage Foundation Disk Group.

- Generally, make sure that **Run this resource in a separate Resource Monitor** is not checked.
- Click **Next**.

- 4 On the Possible Owners screen, by default all the nodes in the cluster are listed as possible owners. Click **Next**.
- 5 On the Dependencies screen, click **Next**. (You do not need to set any dependencies for a disk group resource.)
- 6 Make sure the appropriate SFW quorum cluster dynamic disk group is selected from the drop-down list for the resource, and click **Finish** to complete the operation.



Changing the quorum resource to the dynamic mirrored quorum resource

Use Cluster Administrator to change the quorum resource from a physical disk resource to a dynamic disk quorum resource.

To change the quorum resource to the dynamic mirrored quorum resource

- 1 From Cluster Administrator, right-click the cluster name in the tree view to bring up its context menu.
- 2 Select **Properties**, which displays the Properties window.
- 3 Click the **Quorum** tab of the Properties window.
- 4 Select the name of the dynamic quorum disk group as the resource to be used for the quorum resource.
- 5 Click **OK**.

Setting up a group for SQL Server in MSCS

Using MSCS, you set up a group for the application that contains the SFW disk group or groups that were created for the application.

Add the appropriate SFW disk groups as resources to the application group. You must add the SFW disk groups as the following resource type:

Volume Manager Disk Group

For SQL Server 2000 only, add the IP address resource.

After you install the application, the application resource and its dependencies are automatically configured.

Installing the application on the cluster nodes

You must install the application program files on the same local drive of all the cluster nodes. You install the application data and log files (or other files related to the application data) on the shared storage.

For any specific requirements for the application in an MSCS environment, see the Microsoft documentation.

Checklist for installing the application on the first node

- Applications may have built-in procedures for running on a cluster. Consult the application documentation to determine whether these procedures are available.
- Some applications, such as Microsoft Exchange Server and Microsoft SQL Server, install on both nodes at once.
- All nodes of the clustered application must share the same virtual name and IP address.
- When you install the application, remember not to accept the default location for the application data and log files. Instead, click to browse to the dynamic volumes that were prepared previously.

Checklist for installing the application on the second node

- To install the application on the second node, move the cluster resources to the second node.
- Make sure that the shared volumes, when accessed on the second node, have the corresponding drive letters or mount points that they had when accessed from the first node. To change a drive letter or mount point, see [“To add or change a drive letter or mount point”](#) on page 325.
- If you are installing a database, you may need to stop the database service on the first node while the shared disks are being manipulated by the installation on the second node. Then restart the service after the application is installed.

To add or change a drive letter or mount point

- 1 In VEA, right-click on the volume for which the drive letter will be added or changed.
- 2 Select **File System** and click **Change Drive Letter and Path**.
- 3 In the Drive Letter and Paths window, add or change a drive letter, or add or change a mount point.

- To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter.
- To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Change the drive letter.
- To add a mount point, click the **Add** radio button, click the **Mount as an empty NTFS folder** radio button, browse to select an empty folder or click the **New Folder** button to create a new folder.

Note: A mount point is also referred to as a “drive path.”

- To change a mount point, you must remove it and add it again. (See the bullet above). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.
- Click **OK**.

Verifying the cluster configuration

After you complete the configuration, verify that failover occurs as desired.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.

- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Deploying SFW with Microsoft failover clustering and SQL Server in a campus cluster

This chapter covers the following topics:

- [Tasks for deploying SFW with Microsoft failover clustering in a campus cluster \(Windows Server 2008\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the network and storage](#)
- [Establishing a Microsoft failover cluster](#)
- [Installing SFW](#)
- [Creating disk groups and volumes](#)
- [Implementing a dynamic quorum resource](#)
- [Setting up a group for SQL Server in the failover cluster](#)
- [Installing the application on the cluster nodes](#)
- [Verifying the cluster configuration](#)

Tasks for deploying SFW with Microsoft failover clustering in a campus cluster (Windows Server 2008)

On Windows Server 2008, you can install and configure Veritas Storage Foundation for Windows with Microsoft failover clustering and Microsoft SQL Server 2005. This chapter presents a campus clustering example using a two-node cluster.

The table below outlines the high-level objectives and the tasks for each objective:

Table 15-1 Task list for deploying SFW with Microsoft failover clustering in a campus cluster

Objectives	Tasks
“ Reviewing the requirements ” on page 331	<ul style="list-style-type: none"> ■ Verify hardware and software prerequisites.
“ Reviewing the configuration ” on page 334	<ul style="list-style-type: none"> ■ Review the configuration requirements. ■ Overview of a campus cluster using Microsoft clustering and recovery scenarios.
“ Configuring the network and storage ” on page 343	<ul style="list-style-type: none"> ■ Install and configure the hardware for each node in the cluster. ■ Verify the DNS settings and binding order for all systems.
“ Establishing a Microsoft failover cluster ” on page 345	<ul style="list-style-type: none"> ■ Enable the Microsoft failover clustering feature. ■ Ensure that you have met the hardware requirements for a failover cluster. ■ Run the Microsoft wizard to validate the configuration. ■ Use Failover Cluster Management to create the first node of the cluster. ■ Create the second node of the cluster. ■ Connect the two nodes.
“ Installing SFW ” on page 347	<ul style="list-style-type: none"> ■ Install SFW on Node A (Node B active). ■ Install SFW on Node B (Node A active).

Table 15-1 Task list for deploying SFW with Microsoft failover clustering in a campus cluster (Continued)

Objectives	Tasks
“Creating disk groups and volumes” on page 355	<ul style="list-style-type: none"> ■ In SFW on Node A, create two or more dynamic cluster disk groups on the storage, one or more for the application data files and one for the mirrored quorum.
“Implementing a dynamic quorum resource” on page 365	<ul style="list-style-type: none"> ■ If not done earlier, create a dynamic disk group for the quorum with a mirrored volume. ■ Add the volume manager disk group for the quorum. ■ Change the quorum resource to the dynamic mirrored quorum resource.
“Setting up a group for SQL Server in the failover cluster” on page 368	<ul style="list-style-type: none"> ■ Create a group within failover clustering for the SQL Server application. ■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.
“Installing the application on the cluster nodes” on page 369	<ul style="list-style-type: none"> ■ Install the application program files on the local drive of the first node. ■ Install files relating to the data and logs on the shared storage. ■ Move the cluster resources to the second node. ■ Make sure that the volumes on the second node have the same drive letters or mount points as they had on the first node. ■ Install the application on the second node.
“Verifying the cluster configuration” on page 371	<ul style="list-style-type: none"> ■ Verify the cluster configuration by either moving all the resource groups from one node to another or by simulating a failover by shutting down the active cluster node.

Reviewing the requirements

Reviewing the requirements and the configuration allows you to gain an overall understanding of the configuration and its requirements.

See the following topics:

- [Supported software](#)
- [System requirements](#)
- [Disk space requirements](#)

Supported software

You can check the Late Breaking News information on the Support web site for any recent updates to this list of supported software.

The following software is supported:

- Veritas Storage Foundation 5.1 for Windows (SFW)
Include the following option along with any others applicable to your environment:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
- For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SQL Server 2005 SP2 required) ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition

Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SQL Server 2005 SP2 required) ■ Windows Server 2008 for 64-bit Itanium (IA64) ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition

System requirements

- One CD-ROM drive accessible to each system on which you are installing Microsoft clustering.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access the storage.
- Microsoft clustering requires at least two network adapters per system (one network adapter to connect each system to the public network and one network adapter for the private network on each system). Symantec recommends having two adapters for the private network and routing each

private network adapter through a separate hub or switch to avoid single points of failure.

- Refer to application documentation to determine disk space requirements for your application.
- Each system requires 1 GB of RAM.
- The configuration requires two sites with a storage array for each site, with an equal number of disks at each site for the mirrored volumes.
- Interconnects between the clusters are required for the storage and the network.
- Systems to be clustered must be configured as part of a Windows Server 2008 domain. Each system in a cluster with Microsoft failover clustering must be in the same domain and must be using the same operating system version.
- To install and work with the SFW and Microsoft clustering software, you must have an account with Administrator privileges. You must also have a license key to install SFW.
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported with Windows Server 2003, but DHCP is supported with Windows 2008. Six network interface cards, three for each server (two each for the private network and one for the public network). You also need a static IP address for the cluster itself.

Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 15-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB

Table 15-2 Disk space requirements

Installation options	Install directory/drive
Client components	354 MB

Note: Plan for an equal number of disks on the two sites, because each disk group should contain the same number of disks on each site.

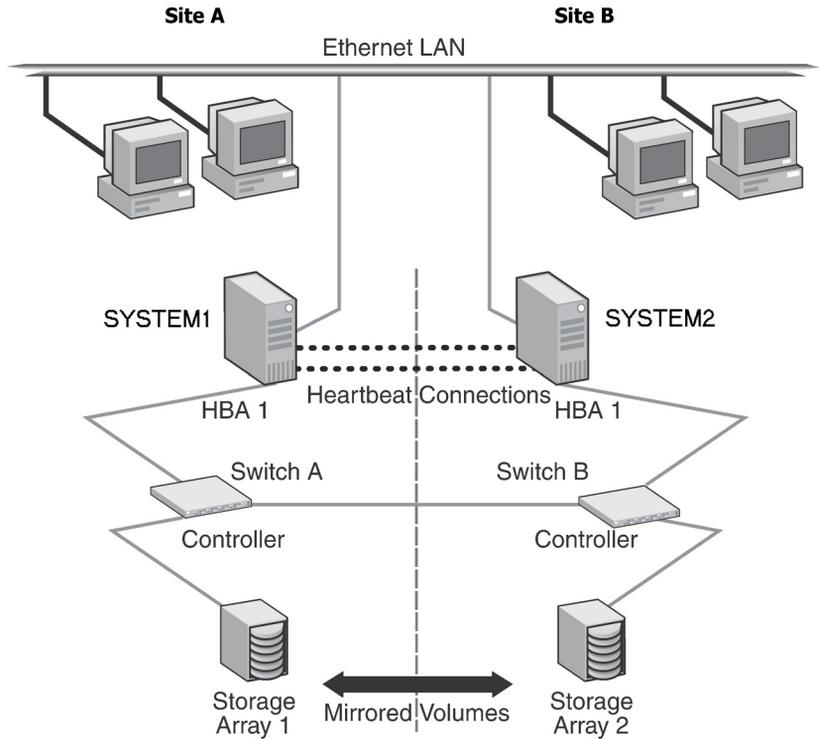
Reviewing the configuration

This configuration example describes a two-node campus cluster with each node at a separate site.

For an overview of campus clusters with Microsoft clustering or for recovery scenarios, see the following:

- [“Overview of campus clustering with Microsoft clustering”](#) on page 336
- [“Campus cluster failure with Microsoft clustering scenarios”](#) on page 337

Figure 15-1 Campus clustering with Microsoft clustering configuration example



The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array and contains mirrored data of the storage on the other array. Each disk group should contain the same number of disks on each site for the mirrored volumes.

Microsoft clustering uses the quorum architecture, where the cluster database resides in the quorum resource. If you are using Microsoft clustering, adding SFW to the configuration protects the quorum disk from being a single point of failure in the cluster because SFW provides dynamic volumes and software mirroring of the quorum device. To avoid a single point of failure, set up the quorum as a dynamic mirrored device. This example includes the dynamic mirrored quorum and requires setting up two or more dynamic cluster disk groups in SFW— one or more cluster disk groups for the application and data and one for the dynamic mirrored quorum.

The example configuration does not include DMP. For instructions on how to add DMP to a clustering configuration, see *Veritas Storage Foundation and High Availability Solutions, Solutions Guide*.

When you are installing SFW and Microsoft clustering together, remember the following:

- A cluster using Microsoft clustering must be running to install SFW. You need to set up the hardware and install the operating system and Microsoft clustering on all systems and establish the failover cluster before installing SFW. Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Use a “rolling install” procedure to install SFW first on the inactive cluster node. Then move the cluster resources to the other node and install on the now inactive node.

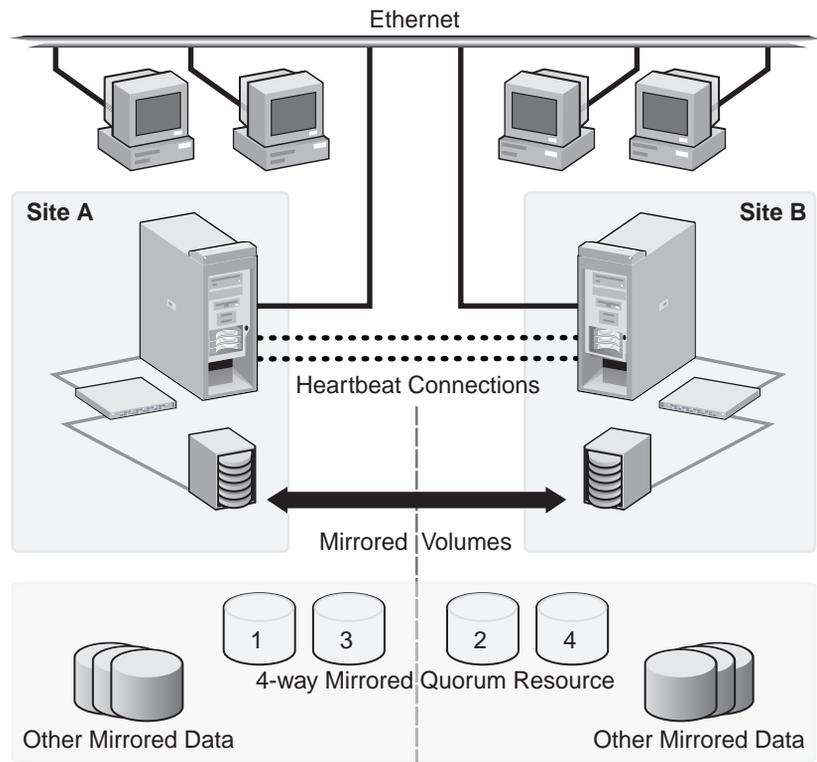
Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- After SFW is installed, create one or more cluster disk groups with SFW and set up the volumes for your application. At the same time, you can create the mirrored volume for the dynamic quorum resource.
- SFW allows you to add fault tolerance to your data volumes. Mirroring of log volumes is recommended, and a mirrored striped RAID layout is recommended for your data volumes. SFW also offers multiple disk groups, multiple mirrors, capacity management and Automatic Volume Growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, Dynamic Multi-pathing, and enhanced snapshot capabilities with FlashSnap.

Overview of campus clustering with Microsoft clustering

[Figure 15-2](#) on page 337 shows a campus cluster configuration with Microsoft clustering. It features mirrored storage across clusters and a mirrored quorum resource. The figure shows a four-way mirrored quorum that has an extra set of mirrors for added redundancy. Although a campus cluster setup with Microsoft clustering can work without Storage Foundation for Windows, SFW provides key advantages over using Microsoft clustering alone. Through a dynamic mirrored volume that functions on multiple disks across multiple sites, SFW protects the quorum resource in the cluster from being the single point of failure in the cluster.

Figure 15-2 Typical campus clustering configuration with Microsoft clustering



Most customers use hardware RAID to protect the quorum disk, but that will not work when a natural disaster takes down the primary node and its attached storage. If the quorum resource is lost to the cluster, the cluster fails, because none of the cluster servers can gain control of the quorum resource and ultimately the cluster. Microsoft clustering alone cannot provide fault tolerance to the quorum disk.

Campus cluster failure with Microsoft clustering scenarios

This section focuses on the failure and recovery scenarios with a campus cluster with Microsoft clustering and SFW installed.

For information about the quorum resource and arbitration in Microsoft clustering, see

[“Microsoft clustering quorum and quorum arbitration”](#) on page 341.

Table 15-3 lists failure situations and the outcomes that occur:

Table 15-3 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
<p>Application fault</p> <p>May mean the services stopped for an application, a NIC failed, or a database table went offline.</p>	Failover	If the services stop for an application failure, the application automatically fails over to the other site.
<p>Server failure (Site A)</p> <p>May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.</p>	Failover	Assuming a two-node cluster pair, failing a single node results in a cluster failover. Service is temporarily interrupted for cluster resources that are moved from the failed node to the remaining live node.
<p>Server failure (Site B)</p> <p>May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.</p>	No interruption of service.	Failure of the passive site (Site B) does not interrupt service to the active site (Site A).
<p>Partial SAN network failure</p> <p>May mean that SAN fiber channel cables were disconnected to Site A or Site B Storage.</p>	No interruption of service.	Assuming that each of the cluster nodes has some type of Dynamic Multi-pathing (DMP) solution, removing one SAN fiber cable from a single cluster node should not effect any cluster resources running on that node, because the underlying DMP solution should seamlessly handle the SAN fiber path failover.
<p>Private IP Heartbeat Network Failure</p> <p>May mean that the private NICs or the connecting network cables failed.</p>	No interruption of service.	With the standard two-NIC configuration for a cluster node, one NIC for the public cluster network and one NIC for the private heartbeat network, disabling the NIC for the private heartbeat network should not effect the cluster software and the cluster resources, because the cluster software simply routes the heartbeat packets through the public network.

Table 15-3 List of failure situations and possible outcomes (Continued)

Failure Situation	Outcome	Comments
<p>Public IP Network Failure</p> <p>May mean that the public NIC or LAN network has failed.</p>	<p>Failover.</p> <p>Mirroring continues.</p>	<p>When the public NIC on the active node, or public LAN fails, clients cannot access the active node, and failover occurs.</p>
<p>Public and Private IP or Network Failure</p> <p>May mean that the LAN network, including both private and public NIC connections, has failed.</p>	<p>No interruption of service.</p> <p>No Public LAN access.</p> <p>Mirroring continues.</p>	<p>The site that owned the quorum resource right before the “network partition” remains the owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource.</p>
<p>Lose Network Connection (SAN & LAN), failing both heartbeat and connection to storage</p> <p>May mean that all network and SAN connections are severed; for example, if a single pipe is used between buildings for the Ethernet and storage.</p>	<p>No interruption of service.</p> <p>Disks on the same node are functioning.</p> <p>Mirroring is not working.</p>	<p>The node/site that owned the quorum resource right before the “network partition” remains the owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource. By default, the Microsoft clustering clussvc service tries to auto-start every minute, so after LAN/SAN communication has been re-established, the Microsoft clustering clussvc auto-starts and will be able to re-join the existing cluster.</p>
<p>Storage Array failure on Site A, or on Site B</p> <p>May mean that a power cord was unplugged, or a storage array failure caused the array to stop responding.</p>	<p>No interruption of service.</p> <p>Disks on the same node are functioning.</p> <p>Mirroring is not working.</p>	<p>The campus cluster is divided equally between two sites with one array at each site. Completely failing one storage array should have no effect on the cluster or any cluster resources that are online. However, you cannot move any cluster resources between nodes after this storage failure, because neither node will be able to obtain a majority of disks within the cluster disk group.</p>

Table 15-3 List of failure situations and possible outcomes (Continued)

Failure Situation	Outcome	Comments
Site A failure (power) Means that all access to site A, including server and storage, is lost.	Manual failover.	If the failed site contains the cluster node that owned the quorum resource, then the overall cluster is offline and cannot be onlined on the remaining live site without manual intervention.
Site B failure (power) Means that all access to site B, including server and storage, is lost.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	If the failed site did not contain the cluster node that owned the quorum resource, then the cluster is still alive with whatever cluster resources that were online on that node right before the site failure.

Dealing with a failover situation

In summary, the site scenarios that can occur when there is a cluster server failure include the following possibilities:

- If the site not owning the quorum volume and the cluster goes offline, the quorum and data volumes stay online at the other site and other cluster resources stay online or move to that site. Storage Foundation for Windows allows the owning cluster node to remain online with 50% ownership of the disks in the quorum group.
- If the site owning the quorum volume goes offline, the remaining site cannot gain control of the quorum volume because it cannot reserve a majority of disks in the quorum group. This is a safeguard to prevent multiple nodes from onlining members of a cluster disk group to which they have access.

Caution: Manual failover of a cluster between two sites should be performed only after coordination between the two sites to ensure that the primary server has actually failed. If you manually import a cluster disk group containing the Microsoft clustering quorum to the secondary (failover) server when the primary server is still active, this causes a split-brain situation. If the split-brain situation occurs, you may lose data because each plex of the mirrored volume may be updated independently when the same disk group is imported on both nodes.

Microsoft clustering quorum and quorum arbitration

This section provides an explanation of the quorum and quorum arbitration in Microsoft clustering.

Quorum

The quorum resource maintains the cluster database, as well as critical recovery information, in a recovery log. The quorum resource has to be available to all nodes through a SCSI or Fibre Channel bus. With Microsoft clustering alone, the quorum disk must be located on a single physical disk. However, with SFW, the quorum disk can be a mirrored volume that spans multiple disks and cluster nodes.

The quorum resource also determines ownership of the cluster. When a node that is controlling the cluster goes offline, other nodes use a challenge/defense protocol to determine which node can have control of the quorum resource and the cluster.

Cluster ownership of the quorum resource

The Microsoft clustering challenge/defense protocol uses a low-level bus reset of the SCSI buses between the machines to attempt to gain control of the quorum resource.

After a SCSI bus reset, the reservation that each server had been holding on the quorum disk is lost. Each server then has roughly 10 seconds to re-establish that reservation, which would in turn let the other servers know that it is still functioning, even though the other servers would not necessarily be able to communicate with it.

If the active cluster server does not re-establish the SCSI reservation on the quorum resource within the time limit, all applications that were on the server will then transfer to the server that establishes the SCSI reservation first. The new server servicing the application may now be a bit slower, but clients will still get their applications serviced. The IP (Internet Protocol) address and network names will move, applications will be reconstituted according to the defined dependencies, and clients will still be serviced, without any question as to the state of the cluster.

The challenge/defense protocol is more complex when the quorum device is a volume in a Storage Foundation for Windows disk group. For a server to take ownership of the disk group containing the cluster quorum device, SFW on that server must successfully import the disk group, obtaining SCSI reservations on more than half of its disks. Because a campus cluster configuration has an even number of disks on each site, failover cannot occur automatically. The manual

CLI command, `vxclus enable` must be used to bring the cluster disk groups online on the secondary node after a site failure.

The vxclus utility

Storage Foundation for Windows provides the `vxclus` command line utility to allow forcing a failover to the secondary site. The command `vxclus enable` creates an entry in the Registry that enables the cluster disk group to be brought online on a node with a minority of the disks. Once `vxclus enable` is executed, you can bring the disk group resource online in Failover Cluster Management. After the cluster disk group is brought online, the `vxclus` functionality is disabled.

To bring a cluster online on a node with a minority of the cluster disks

- 1 Use the following `vxclus` command for each disk group on your cluster node:

```
vxclus enable -g<DynamicDiskGroupName>
```

You will be asked to confirm the use of this command.

Caution: When bringing a cluster disk group online with a minority of cluster disks, make sure that a majority of the disk group disks are NOT online on any other cluster node before (and after) onlining the disk group. If a majority of disk group disks are online on another node, data corruption can occur.

- 2 If the cluster service has stopped because of a dynamic quorum resource failure, start the cluster service (`clusvc`).
- 3 Then, using Failover Cluster Management, bring the cluster disk groups online.

For more information on the `vxclus` utility, see the “Command Line Interface” chapter of the *Storage Foundation Administrator’s Guide*. The `vxclus` utility also provides support for booting from a SAN, but you must have a hardware storage array that supports the capability.

Configuring the network and storage

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent the Microsoft cluster from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.
- 13 Close the window.

Establishing a Microsoft failover cluster

Before installing SFW, you must first verify that Microsoft failover clustering is enabled (if a new installation of Windows Server 2008), and then establish a Microsoft failover cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To enable Microsoft failover clustering

- 1 In Server Manager, select **Features** in the left pane (tree view) and then click **Add Features** (link on the right side of the screen).
- 2 In the Add Features Wizard, check the **Failover Clustering** option, and click **Next**.
- 3 Click **Install**.
- 4 When the installation is complete, click **Close**.

To establish a Microsoft failover cluster

- 1 Ensure that you have met the hardware prerequisites for a failover cluster. You can run the Microsoft wizard to validate the configuration. See the Microsoft documentation for details.
- 2 Configure the shared storage and create a volume with drive letter “Q” for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends a minimum of 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster using Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 4 In the action pane, click **Create a Cluster**. The Create Cluster Wizard will start.
If this is the first time this wizard has been run, the Before You Begin page will appear. Review the information that is displayed and then click **Next**. You can hide this page on subsequent uses of the wizard; if this has been done, the first page might be the Select Servers page.
- 5 In the Select Servers panel, type the name of the first node in the Enter server name field and click **Add**. You can also use the Browse button to browse the Active Directory for the computers you want to add. Repeat this step for the second node.
- 6 After both nodes have been added to the list of Selected Servers, click **Next**.

- 7 Based on the information on the validation warning screen, assess your hardware configuration, and select one of the options. Symantec recommends that you select **Yes**, which starts the Validate a Configuration wizard. Follow the wizard instructions.
- 8 In the Access Point for Administering the Cluster screen, in the Cluster Name field, type the name for the failover cluster. This is the name that you use to connect to and administer the cluster.
- 9 In the Address field of the network area, type the appropriate IP address and then click **Next**.
- 10 In the Confirmation screen, verify that the cluster configuration is correct, and then click **Next** to create the cluster.
- 11 Review the Summary page and then click **Finish** to close the wizard.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use these disks when you create your SFW cluster disk groups, you must first remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs. After creating the SFW cluster disk groups, you will add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

Connecting the two nodes

Make the necessary connections between the two sites. The cluster is already active on Server A, so Microsoft clustering is now in control of the cluster storage on Server A, and both nodes of the storage cannot be accessed at the same time by the operating system.

To connect the two nodes

- 1 Connect corresponding cables between the three network cards on the two sites.
- 2 Connect the two switches at the two sites through the storage interconnect.
- 3 Test the connectivity between the two sites. Test the IP addresses of all the network adapter cards in the cluster. Bring up the command window and type `ping ipaddress`, where the *ipaddress* is the corresponding network adapter in the other node.

Installing SFW

This section assumes you are running a Microsoft failover cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. Our example uses a two node configuration, so the inactive system is the second node. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft failover cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 347.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 348.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 353.

Pre-installation tasks

Perform the following pre-installation tasks:

- Moving the online groups

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.

If there is more than one resource group, you must repeat this step until all the resource groups are moved.

- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a Microsoft failover cluster configuration.

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

Note: Before you install Storage Foundation for Windows, make sure that the node is inactive.

To install the product

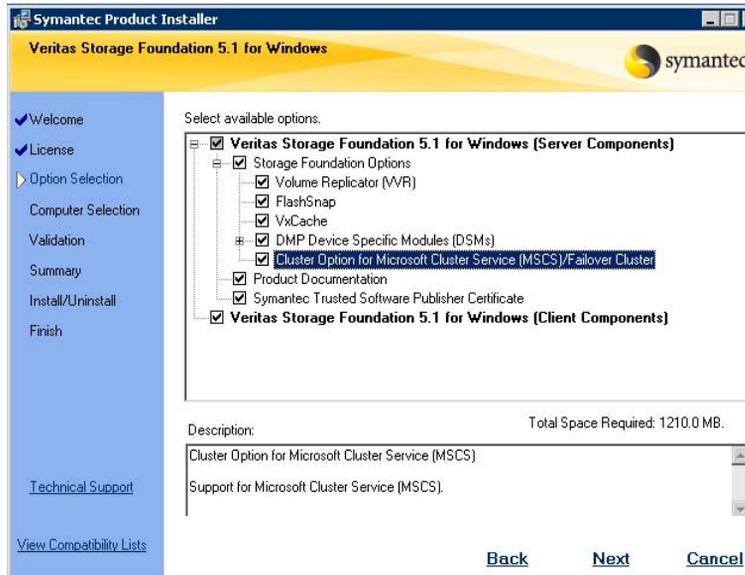
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

3 Click **Storage Foundation 5.1 for Windows**.



- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for "**I accept the terms of the license agreement,**" and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
- 8 Repeat for additional license keys.
To remove a license key, click the key to select it and click **Remove**.
To see the license key's details, click the key.
- 9 Click **Next**.

- 10 Specify the product options by selecting the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.

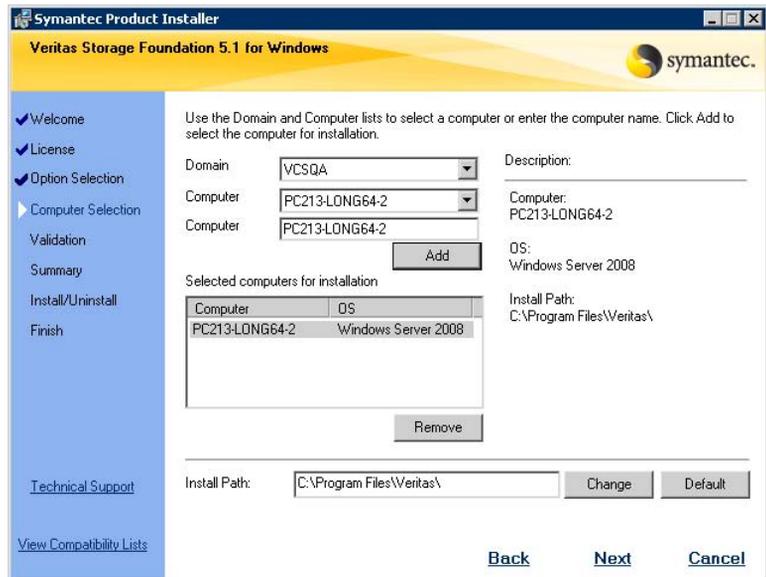


Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.

Note that under Veritas Dynamic Multi-pathing, you can select DMP Device Specific Modules (DSMs).

- 11 Verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is checked, to install the client component and click **Next**.

12 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path	Optionally, change the installation path. <ul style="list-style-type: none">■ To change the path, select a computer in the Selected computers for installation field, type the new path, and click Change.■ To restore the default path, select a computer and click Default. The default path is: C:\Program Files\Veritas For 64-bit installations, the default path is: C:\Program Files (x86)\Veritas
--------------	--

- 13 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 14 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 15 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that Microsoft clustering allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

Dynamic Multi-pathing

Additionally, if you selected the Dynamic Multi-pathing option, a warning appears:

- For DMP DSM installations—the time required to install the Veritas Dynamic Multi-pathing DSM feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during

installation. After installation, reconnect additional physical paths before rebooting the system.

- 16 Review the information and click **Install**. Click **Back** to make changes.
- 17 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.
If the installation is successful on all systems, the installation report screen appears.
If a security alert asks you to accept the Veritas driver software, click **Yes**.
- 18 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 19](#) if you are installing SFW on the local node only.
 - Proceed to [step 21](#) if you are installing SFW on local and remote systems.
- 19 To complete the installation, click **Finish**.
- 20 Click **Yes** to reboot the system and complete the installation.
- 21 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
- 22 Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
- 23 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available.
- 24 Click **Next**.
- 25 Click **Finish**.
- 26 Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW installation

Moving the online groups

You can move the resource groups from the current system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open the Failover Cluster Management tool. (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click the resource group and then click **Move this service or application to another node > Move to node [name of original node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved back to the original node.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that all the resource groups have moved back to the original system.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the Microsoft failover cluster.

See “[SFW installation tasks](#)” on page 347.

Creating disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and dynamic volumes on the cluster storage, which in a campus cluster consists of at least two storage arrays.

Before you create disk groups and volumes, consider the following items:

- The type of volume configurations that are required
- The number of LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs which depend on the traffic load
- The disk groups and number of disks on each site

Note: For campus clusters, each disk group *must* contain an equal number of disks on each site.

- Types of volumes required and location of the plex of each volume in the storage array

Note: Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

Create two or more dynamic cluster disk groups on the storage—one or more for the application data files and one for the mirrored quorum.

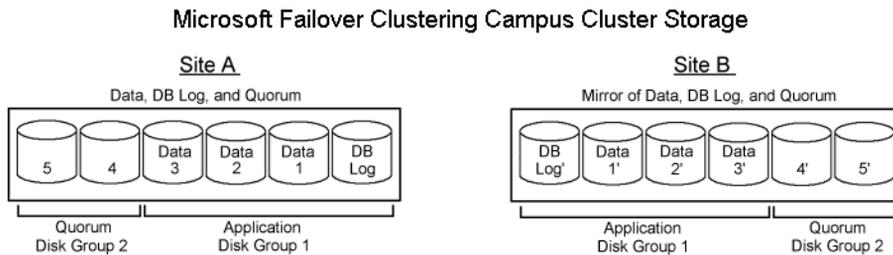
For the Microsoft SQL Server application data files, you could create a separate disk group for each database. It is best to separate data files from log files and place them in separate volumes. For example, you might create a SQL disk group, INST1_DG, containing three volumes:

- INST1_DB1_VOL: Contains the SQL database.
- INST1_DB1_LOG: Contains the transaction log.
- INST1_DATA_FILES: Contains Microsoft SQL Server system data files.

The illustration that follows shows a typical Microsoft failover cluster with a campus cluster setup of disks. This example has only one application disk group that spans the storage arrays at both sites. The data and database log on Site A are mirrored to Site B. Each mirrored volume can have more than two disks, but must have an even number, such as four. All the application data could be in one large mirrored volume with multiple disks, but the same number of disks are required on both sites for the mirroring. It is recommended that the log volumes be on separate disks from the data.

In the example, a four-way mirror for the quorum volume provides additional redundancy. The minimum configuration would be a two-way mirror. If possible, use small disks for the quorum volume. Microsoft recommends 500 MB for the quorum volume.

Figure 15-3 Microsoft failover cluster with campus cluster disks and disk groups example



Configuring the disks and volumes

Ensure that each disk group contains an equal number of disks on each site, and that each volume is a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

See the following sections:

- [“Considerations when creating new volumes”](#) on page 356
- [“Creating a dynamic \(cluster\) disk group”](#) on page 357
- [“Creating a volume”](#) on page 359

Considerations when creating new volumes

- For campus clusters, when you create a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.

- Symantec recommends using either simple mirrored (concatenated) or striped mirrored for the new volumes. Striped mirrored gives you better performance compared to concatenated.
 When selecting striped mirrored, select two columns in order to stripe one enclosure that is mirrored to the second enclosure.
- You cannot select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.
 The internal names for the disks which the current system can access for available storage are displayed, with names **Harddisk1**, **Harddisk2**, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a dynamic (cluster) disk group

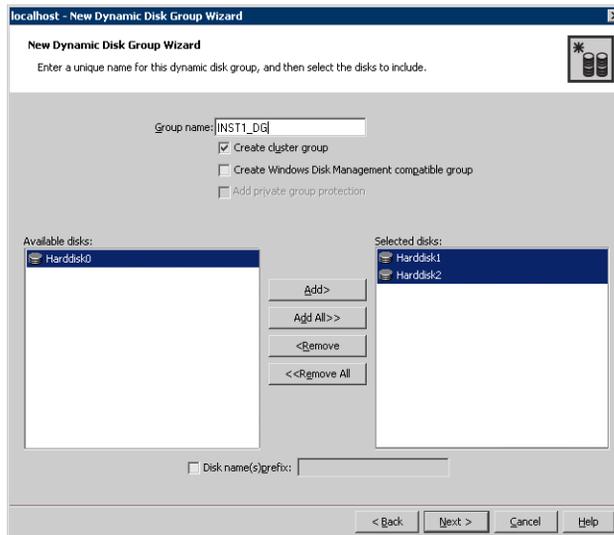
Use the following procedure to create a dynamic cluster disk group.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that

contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

■ Click **Next**.

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the new disk group.

Proceed to create the appropriate volumes on each disk.

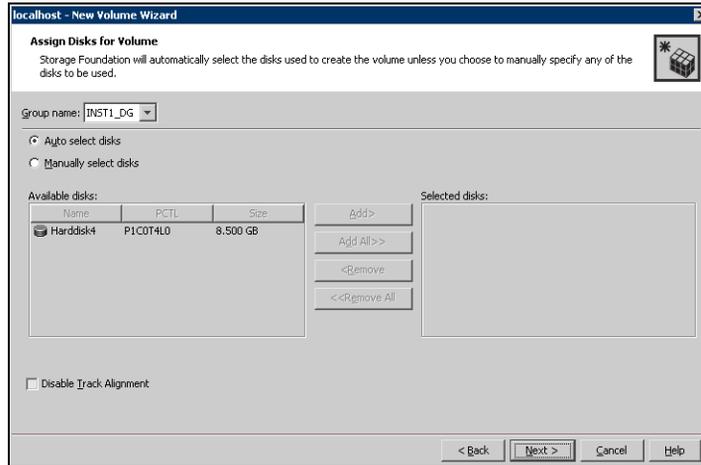
Creating a volume

Use the following procedure to create dynamic volumes.

To create dynamic volumes

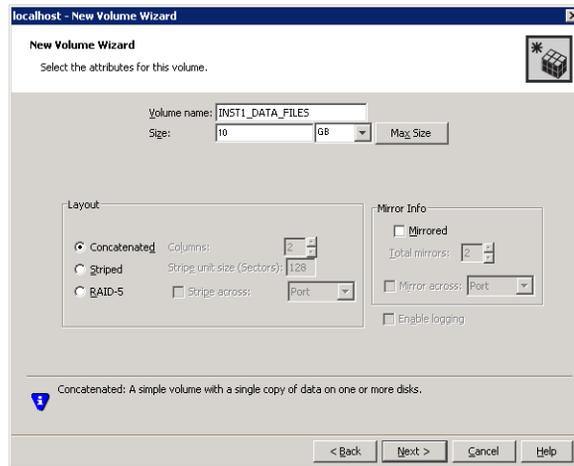
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume; make sure the appropriate disk group name appears in the Group name drop-down list.



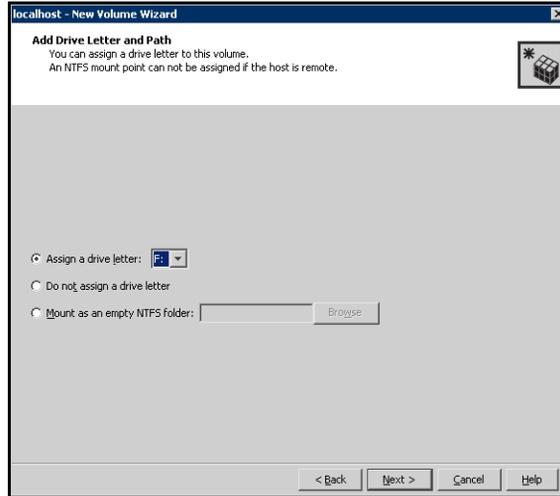
- 7 Select auto or manual disk selection and enable or disable track alignment.
 - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3COT2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
 - To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.

9 Specify the parameters of the volume.

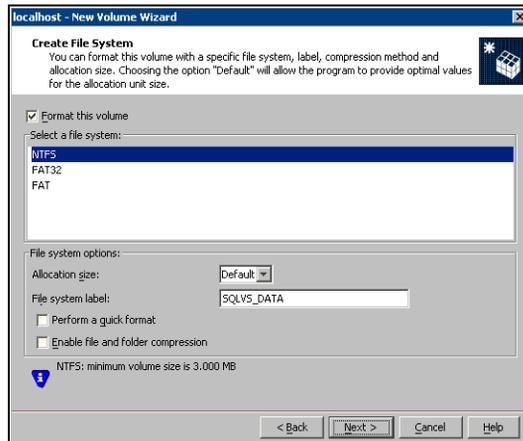


- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - Select a volume layout type. For campus clusters, select either **Concatenated** or **Striped**. Since campus clusters are mirrored volumes, you also must select the **Mirrored** checkbox.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
 - In the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
 - Verify that **Enable logging** is not selected and click **Next**.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.
- 12 Create an NTFS file system.

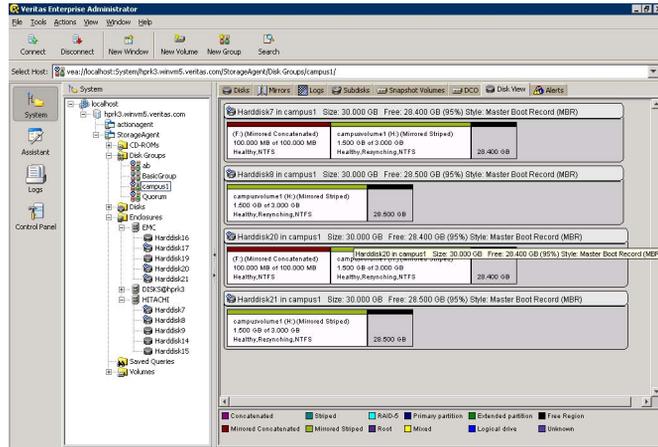


- Make sure the **Format this volume** checkbox is checked and select **NTFS**.
- Select an allocation size or accept the Default.

- The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance. Click **Next**.
- 13 Click **Finish** to create the new volume.
- 14 Repeat these steps to create additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Figure 15-4 View of disks with volumes in VEA Console



Implementing a dynamic quorum resource

One of the key advantages of using SFW with Microsoft clustering is that you can create a mirrored quorum resource that adds fault tolerance to the quorum. The tasks for creating a mirrored quorum resource are:

- [“Creating a dynamic cluster disk group and a mirrored volume for the quorum resource”](#) on page 365
- [“Adding the volume manager disk group for the quorum”](#) on page 365
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 366

Creating a dynamic cluster disk group and a mirrored volume for the quorum resource

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using four (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

To create a four-way mirrored volume using the New Volume wizard

- 1 Create the cluster disk group with four small disks.
- 2 Create a volume with the four disks.
- 3 Select the **Concatenated** layout, click the **Mirrored** check box, and specify four mirrors.

For full details on creating cluster disk groups and volumes, see [“Creating disk groups and volumes”](#) on page 355.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by Microsoft clustering.

Adding the volume manager disk group for the quorum

You must add the Volume Manager Disk Group resource for the quorum.

To add the Volume Manager Disk Group resource for the quorum

- 1 If Failover Cluster Management is already open, then proceed to Step 2.

To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.

- 2 Verify that the cluster is online on the same node where you created the disk group.
- 3 In the left pane of Failover Cluster Management, right-click **Services and Applications** and select **More Actions > Create Empty Service or Application**.
- 4 Right-click the new group and rename it, for example `QUORUM`.
- 5 Right-click `QUORUM` and select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 6 Right-click **New Volume Manager Disk Group** in the center pane and click **Properties**.
- 7 In the General tab of the Properties dialog box, type a name for the resource in the Resource Name field, for example, `QUORUM_DG_RES`.
- 8 On the Properties tab, in the Disk Group Name field, type the name of the disk group that you previously created for the quorum, and click **OK** to close the dialog box.
- 9 Right-click the Quorum disk group resource (for example, `QUORUM_DG_RES`) in the left pane and select **Bring this resource online**.
The specified disk group resource, `QUORUM_DG_RES` resource, is created under the Quorum group (for example, `QUORUM`).

Changing the quorum resource to the dynamic mirrored quorum resource

Use the following procedure to configure the cluster quorum settings and change the quorum resource to a dynamic mirrored quorum resource.

To change the quorum to a dynamic mirrored quorum resource

- 1 In Failover Cluster Management, right-click the cluster node in the configuration tree, and select **More Actions > Configure Cluster Quorum Settings**.
The Configure Cluster Quorum Wizard opens.
- 2 Review the screen and click **Next**.
- 3 Select either the **Node and Disk Majority** or **No Majority: Disk Only** radio button, and click **Next**.
- 4 Select the storage resource that you want to assign as the disk witness for the quorum and click **Next**.

This is the Volume Manager Disk Group resource that you previously created for the quorum disk group, for example, `QUORUM_DG_RES`.

- 5 Review the information in the Confirmation screen and click **Next**.
- 6 Click **Finish** to close the wizard.

Setting up a group for SQL Server in the failover cluster

Using Failover Cluster Management, set up a cluster group for the SQL Server application.

Add the appropriate SFW disk groups as resources to the application group. You must add the SFW disk groups as the following resource type:

Volume Manager Disk Group

Once you install the application, the application resource and its dependencies are automatically configured.

To create a Volume Manager Disk Group resource for the application

- 1 If Failover Cluster Management is already open, then proceed to Step 2. To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.
- 2 In the left pane of Failover Cluster Management, right-click the SQL Server cluster group (for example, SQL_GROUP) and select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 3 In the center panel under Disk Drives, double-click **New Volume Manager Disk Group** to open its Properties dialog box.
- 4 On the General tab of the Properties dialog box, type a name for the resource.
For example, type SQL_DG_RES.
- 5 On the Properties tab, in the Disk Group Name field, type the name of the disk group you previously created for the application (for example, INST1_DG), and click **OK** to close the dialog box.
- 6 Right-click the newly named resource and select **Bring this resource online**.
- 7 If you created more than one disk group for the application, repeat this procedure to add another Volume Manager Disk Group resource for another disk group.

Installing the application on the cluster nodes

The application program files must be installed on the same local drive of all the cluster nodes. The application data and log files or other files related to the application data are installed on the shared storage.

Refer to the Microsoft documentation for any specific requirements for the application in a failover cluster environment.

Pointers for installing the application on the first node

- Applications may have built-in procedures for running on a cluster. Consult the application documentation to determine whether these procedures are available.
- Some applications install on both nodes at once.
- All nodes of the clustered application need to share the same virtual name and IP address.
- Remember not to accept the default location for the application data and log files. Instead, browse to the dynamic volumes that were prepared previously.

Pointers for installing the application on the second node

- In Failover Cluster Management, move the cluster resources to the second node.
- Make sure that the shared volumes, when accessed on the second node, have the corresponding drive letters or mount points that they had when accessed from the first node.
- If you are installing a database, you may need to stop the database service on the first node while the shared disks are being manipulated by the installation on the second node. Then restart the service after the application is installed.

To add or change a drive letter or mount point

- 1 In VEA, right-click on the volume for which the drive letter will be added or changed.
- 2 Select **File System** and click **Change Drive Letter and Path**. The Drive Letter and Paths window appears. You can choose from the following:
 - To add a drive letter, click **Add**. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.
 - To change a drive letter, click **Modify**. The **Assign a drive letter** drop-down list becomes available. Change the drive letter and click **OK**.

- To add a mount point, click **Add**, click the **Mount as an empty NTFS folder** radio button, browse to select an empty folder or click the **New Folder** button to create a new folder, and click **OK** to mount the volume.
- To change a mount point, you must remove it and then select the Add option to add it back. To remove it, select it in the Drive Letter and Paths window and click **Remove**.

Verifying the cluster configuration

After completing the configuration, verify that failover occurs as desired.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Failover Cluster Management to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node. Do not simulate failover in a production environment.

To move online groups

- 1 Open the Failover Cluster Management tool (**Start > All Programs > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open Failover Cluster Management. Click **Start > All Programs > Administrative Tools > Failover Cluster Management** from any node in the cluster.
- 3 In Failover Cluster Management, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move the resource groups back to the original node, restart the node you shut down in [step 1](#), select the resource group, and use **Move this service or application to another node > Move to node [name of node]** to move the resource group.

Deploying SFW and VVR with MSCS: New SQL 2000 installation

This chapter covers the following topics:

- [Tasks for a new SQL Server 2000 installation with SFW, VVR, and MSCS \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing an MSCS cluster](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Creating the SQL virtual server group](#)
- [Creating the MSDTC resource](#)
- [Installing SQL Server 2000](#)
- [Implementing a dynamic mirrored quorum resource](#)
- [Verifying the cluster configuration](#)
- [Creating a parallel environment on the secondary site](#)
- [VVR components overview](#)
- [Creating resources for VVR](#)
- [Configuring VVR: Setting up an RDS](#)

- [Creating the RVG resource \(primary and secondary sites\)](#)
- [Setting the SQL server resource dependency on the RVG resource](#)
- [Working with a solution: Normal operations and recovery procedures](#)

Tasks for a new SQL Server 2000 installation with SFW, VVR, and MSCS (Windows Server 2003)

This chapter describes how to install and configure Storage Foundation for Windows and Veritas Volume Replicator (VVR) with MSCS and SQL Server 2000 in a new installation. You can install and configure high availability, replication, and SQL components. This environment involves an active/passive configuration with one to one failover capability. After setting up a SFW environment with high availability for SQL on a primary site, you can create a secondary or “failover” site for replication. Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details on VVR.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 16-1 Tasks for deploying SFW and VVR with MSCS for SQL Server 2000

Objective	Tasks
“Reviewing the requirements” on page 375	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 377	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive SQL configuration in a two-node cluster
“Configuring the storage hardware and network” on page 380	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Establishing an MSCS cluster” on page 381	<ul style="list-style-type: none"> ■ Reviewing general guidelines to establish an MSCS cluster
“Installing SFW with MSCS/Failover Cluster option” on page 383	<ul style="list-style-type: none"> ■ Modifying the driver signing options for Windows 2003 remote systems ■ Installing SFW ■ Configuring VxSAS ■ Restoring the driver signing options for Windows 2003 remote systems

Table 16-1 Tasks for deploying SFW and VVR with MSCS for SQL Server 2000

Objective	Tasks
“Configuring SFW disk groups and volumes” on page 395	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes
“Creating the SQL virtual server group” on page 407	<ul style="list-style-type: none"> ■ Creating a SQL Server cluster group ■ Creating the virtual server IP address ■ Creating the disk group resource
“Installing SFW with MSCS/Failover Cluster option” on page 383	<ul style="list-style-type: none"> ■ Creating the MSDTC resource for SQL Server
“Installing SQL Server 2000” on page 410	<ul style="list-style-type: none"> ■ Installing SQL and any required patches ■ Verifying SQL installation
“Implementing a dynamic mirrored quorum resource” on page 414	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the cluster group ■ Changing the quorum resource to a dynamic mirrored quorum resource.
“Verifying the cluster configuration” on page 417	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node
“Creating a parallel environment on the secondary site” on page 418	<ul style="list-style-type: none"> ■ Using the previous objectives for the secondary site from reviewing the prerequisites through testing the cluster.
“Creating resources for VVR” on page 420	<ul style="list-style-type: none"> ■ Creating an IP address for the Replicated Volume Group (RVG). ■ Creating a network name resource for the Replicated Volume Group (RVG) ■ Creating the Replicator Log volumes for VVR.
“Configuring VVR: Setting up an RDS” on page 422	<ul style="list-style-type: none"> ■ Configuring VVR by setting up an RDS

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

Supported software for MSCS with VVR

- Veritas Storage Foundation 5.1 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster and the Veritas Volume Replicator Option
- Microsoft SQL 2000 servers and their operating systems:

<p>Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)</p>	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
<p>Microsoft SQL Server 2000 (64-bit) Enterprise Edition</p>	<ul style="list-style-type: none"> ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)
<p>Microsoft SQL Server 2000 (64-bit) Standard Edition or Enterprise Edition (SP4 required)</p>	<ul style="list-style-type: none"> ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions)

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 16-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

- One CD-ROM drive accessible to the system on which you are installing SFW.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- MSCS requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft SQL Server documentation for instructions on creating a reverse lookup zone.
- MSCS requires two disks for SQL: one for SQL database files and one for SQL log files.
- Each system requires 1 GB of RAM for SFW.
- SFW requires administrator privileges to install the software.
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.

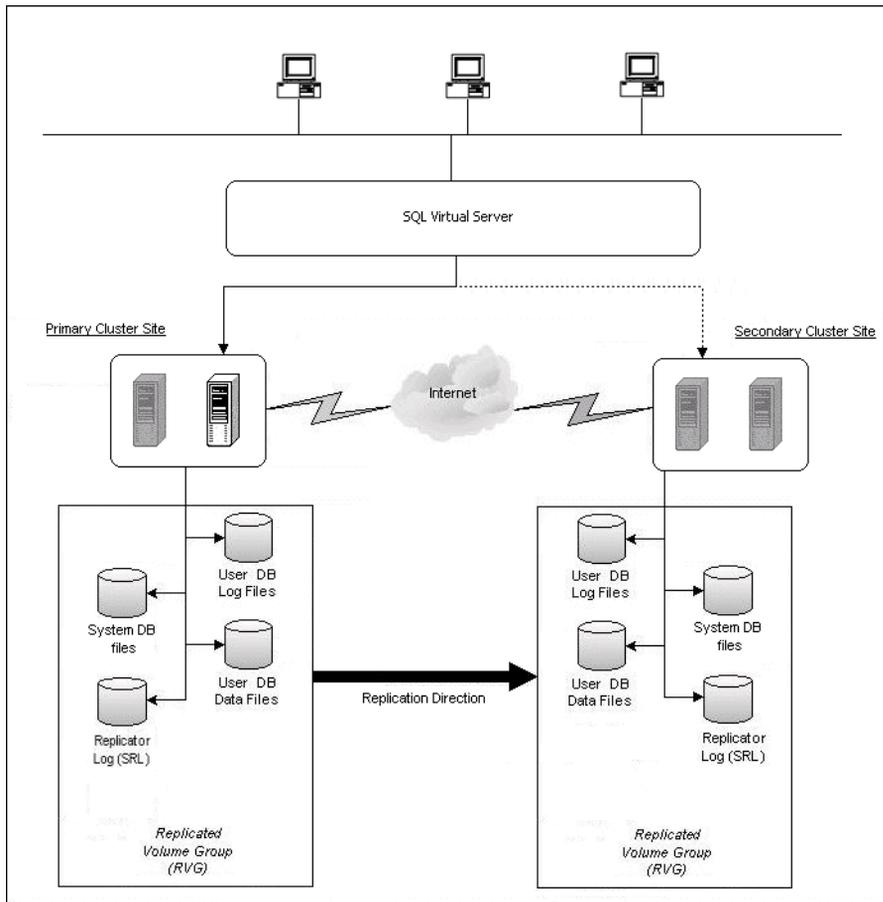
Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Reviewing the configuration

The following figure illustrates a typical clustered VVR configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2000 application data is stored on the volumes that are under the control of the RVG. A separate disk group is created for the quorum volume which is not replicated.

[Figure 16-1](#) shows a typical VVR configuration.

Figure 16-1 Typical VVR configuration



If the Microsoft SQL Server 2000 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the replication solution is activated. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over. The data that was replicated to the secondary site is used to restore the SQL services to clients.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site. The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary Site

SYSTEM1 & SYSTEM2	server names
SQL_GROUP	Microsoft SQL Server resource group
SQLCLUST	Microsoft SQL Server virtual cluster (underscores not supported)
SQLVS	Microsoft SQL Server virtual server
INST1	Microsoft SQL Server instance name
INST1_DG	disk group for Microsoft SQL volumes
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
SQLVS_QRM	volume for storing the MSCS cluster quorum
QUORUM_DG	quorum volume disk group

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
----------------------	--

All the other parameters are the same as on the primary site.

DR Components

INST1_RDS	VVR Replicated Data Set (RDS) name
INST1_RVG	VVR Replicated Volume Group (RVG) name

INST1_REPLOG	VVR Replicator log volume
INST1_RVG_RES	MSCS Replicated Volume Group Resource name

Configuring the storage hardware and network

Use the following procedures to configure the storage hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent MSCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 5 In the Public Status dialog box, on the General tab, click **Properties**.

- 6 In the Public Properties dialog box, on the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Establishing an MSCS cluster

Before installing SFW, you must establish an MSCS cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To establish an MSCS cluster (general guidelines)

- 1 Verify that the quorum disk has been created before installing MSCS and configuring a cluster. (For IA64 systems, the quorum must be created using MBR instead of GPT or it will not be visible.)
- 2 Configure the shared storage and create a partition with drive letter "Q" for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster (SYSTEM1) using MSCS Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Verify that the node can access the shared storage.
- 4 Connect the second node to the shared storage.
- 5 Add the second node (SYSTEM2) using Cluster Administrator on that system.
- 6 Test the cluster by using the Move Group command to move the cluster resources to the second node.
SYSTEM2 becomes the active cluster node.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running a Microsoft cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft cluster simultaneously.

SFW installation tasks

The product installer enables you to install the software for Veritas Storage Foundation 5.1 for Windows. The installer automatically installs SFW. You must select the options to install VVR, and the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation. The steps in this section are based on a server installation.

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 383.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 385.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 390.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
- Moving the online groups

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 16-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are online on the system where you are installing SFW. You must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install Veritas Storage Foundation for Windows on a Microsoft cluster configuration.

To install the product

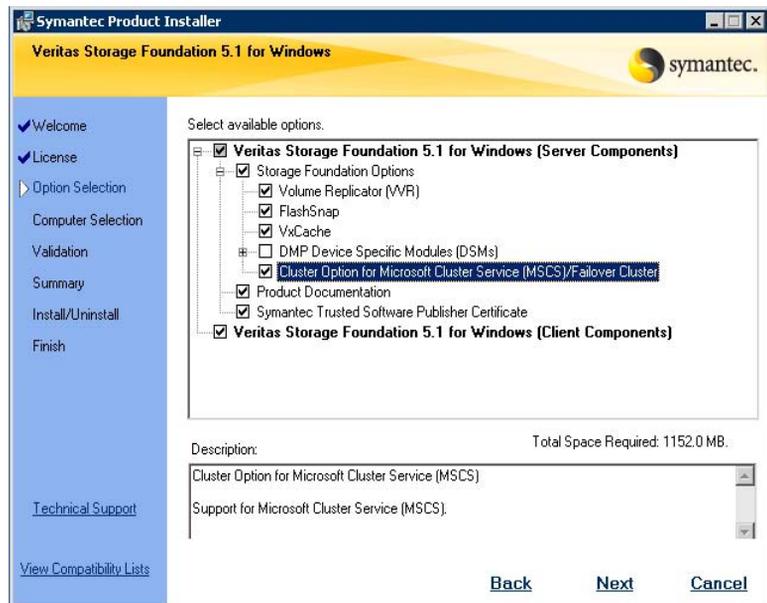
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

- 3 Click **Storage Foundation 5.1 for Windows**.



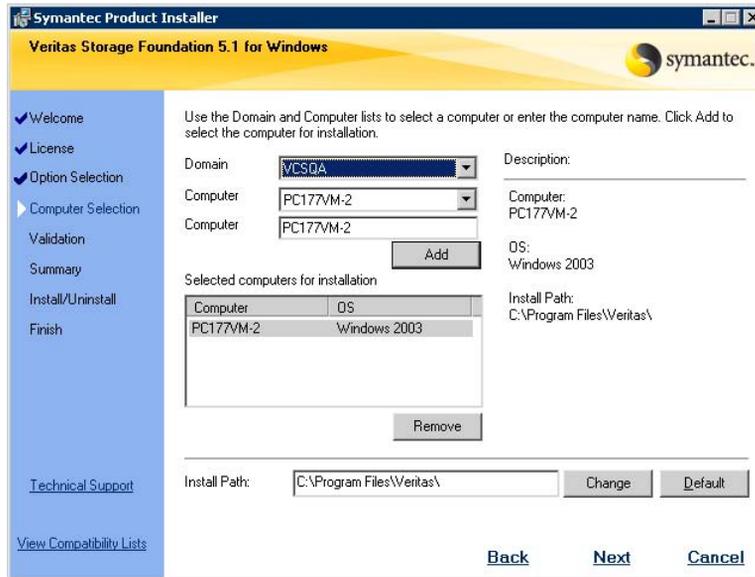
- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for "**I accept the terms of the license agreement,**" and click **Next**.
- 7 Enter the product license key before adding license keys for features.
- 8 Enter the license key in the top field and click **Add**.
- 9 Repeat for additional license keys.
To remove a license key, click the key to select it and click **Remove**.
To see the license key's details, click the key.
- 10 Click **Next**.

11 Specify the product options:



- 12 Select the **Volume Replicator (VVR)** option.
- 13 Select the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.
- 14 Do not select the Dynamic Multi-pathing option.
 Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.
 Click **Next**.
- 15 To install the client components on all the computers in your installation, verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is selected, to install the client component.
 Click **Next**.

16 Select the domain and the computers for the installation and click **Next**.



- Domain** Select a domain from the list.
Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.
- Computer** To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.
To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.
Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 17 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 18 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
 If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 19 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that MSCS allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

- 20 Review the information and click **Install**. Click **Back** to make changes.
- 21 The Installation Status screen displays status messages and the progress of the installation.
 If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.
 If the installation is successful on all systems, the installation report screen appears.

If your local computer has its driver signing options set to Block or Warn then installation fails.

- 22 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 23](#) if you are installing SFW on the local node only.
 - Proceed to [step 25](#) if you are installing SFW on local and remote systems.
- 23 To complete the installation, click **Finish**.
- 24 Click **Yes** to reboot the system and complete the installation.
- 25 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
 - Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
 - When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
 - Click **Finish**.
 - Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW Installation for the other systems in the MSCS cluster
- Configure the VxSAS service
- Resetting the driver signing options

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.
- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 383.

Configuring the VxSAS service

Complete the following procedure to configure this required service for VVR. This procedure should not be done until you have installed SFW on all cluster systems. Otherwise, you will get an error message from the VxSAS wizard if you try to select a system without SFW installed.

You can run the VxSAS wizard from the secondary site once SFW is installed on all cluster systems; at that time, you can run the wizard for both the primary and secondary site systems. The MSCS groups can be either online or offline.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. Accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.

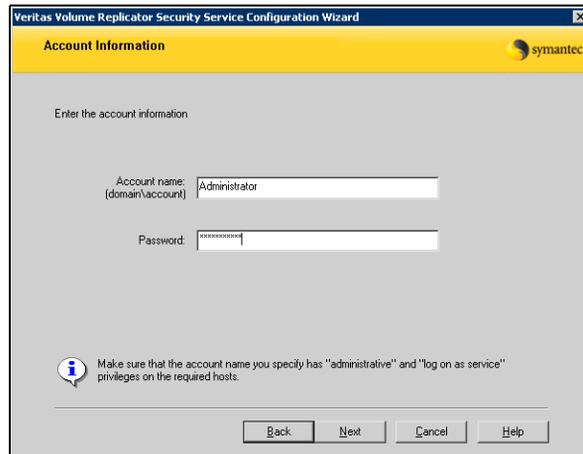
The welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

2 Complete the Account Information wizard page as follows:

Account name (domain\account) Enter the administrative account name in the Account name field.

Password Specify a password in the **Password** field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same user name and password when configuring the VxSAS service on the other hosts.



3 After providing the required information, click **Next**.

4 Select the required domain to which the hosts that you want to configure belong, from the Domain Selection wizard page.

Selecting Domains The Available Domains pane lists all the domains that are present in the Windows network neighborhood.

Select the required domain by moving the appropriate name from the Available Domains pane to the Selected Domains pane, either by double-clicking it or using the arrow button.

Adding a Domain If the domain name that you require is not displayed, then add it by using the **Add Domain** option. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected Domains list.

After specifying the domain, click **Next**.

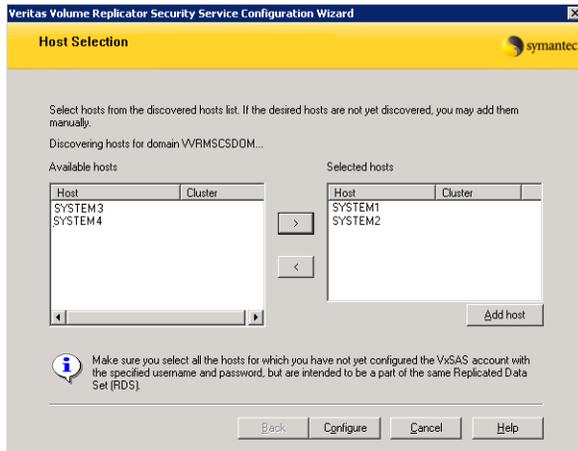
5 Select the required hosts from the Host Selection page.

Selecting Hosts The Available Hosts pane lists the hosts that are present in the specified domain.

Select the required host by moving the appropriate name from the Available Hosts list to the Selected Hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a Host If the host name you require is not displayed, then add it using **Add Host** option. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected Hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.



6 After the configuration completes, the Configuration Results page is displayed. If the operation is successful, then the Status column displays the appropriate message to indicate that the operation was successful.

If the operation was not successful, then the page displays the status as failed and the corresponding details on why the account update failed. It also displays the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

- 7 Click **Finish** to exit the wizard.

Resetting the driver signing options

You must reset the driver signing options to its previous state. This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for SQL. A dynamic cluster disk group is a collection of one or more disks that behave as a single storage repository and which can potentially be accessed by different computers. Within each disk group, you can have dynamic volumes with different layouts.

Configuring disk groups and volumes involves the following tasks:

- [Planning disk groups and volumes](#)
- [Creating dynamic disk groups](#)
- [Creating dynamic volumes](#)
- [Managing disk groups and volumes](#)

Planning disk groups and volumes

Before installing SQL, you must create disk groups and volumes using the VEA console installed with SFW.

A dynamic cluster disk group is a collection of one or more disks that behave as a single storage repository and which can potentially be accessed by different computers. Within each disk group, you can have dynamic volumes with different layouts.

Before creating a disk group, consider:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.
- The number of disk groups and volumes that are needed for SQL Server
- The number of disk groups for SQL depends on the planned organization of the data. The application program files need to be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage in a cluster disk group. You create at least one disk group for the system data files. You may want to create additional disk groups for user databases. Symantec recommends that you place SQL Server user database files and log files on separate volumes. Replicating the system databases is not required or recommended. Make sure that the system databases are not placed on volumes that will be replicated.
- The disk groups and volumes for the mirrored quorum resource

You will need a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. You can create the quorum disk group at the same time you create application disk groups, although it is not required for installing the application. You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that SQL is working in the cluster before adding the dynamic quorum volume. Make sure that the disk group for the mirrored quorum resource is not replicated.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

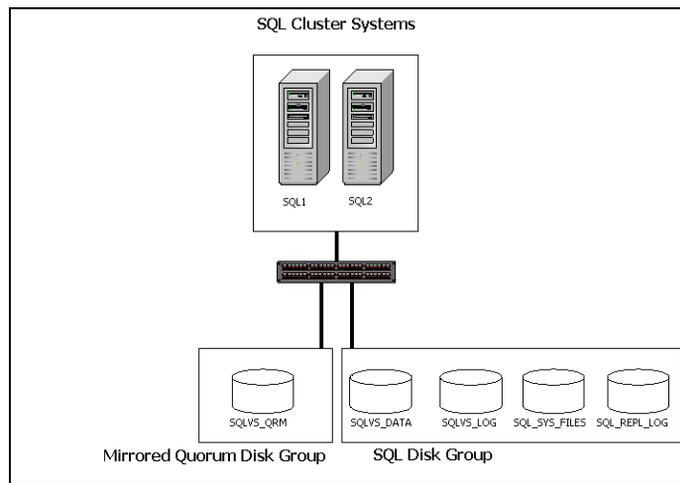
VVR does *not* support these types of volumes:

- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with the Dirty Region Log (DRL)
- Volumes with a comma in their names
- For the Replicator Log volume, in addition to the above types also make sure that the volume does not have a DCM.

Caution: Do not use volume types that are not supported by VVR.

Below is a detailed view of the disk groups and volumes for SQL:

Figure 16-2 SFW disk groups and volumes for SQL virtual server INST1-VS in MSCS setup



SQL disk group INST1 contains four volumes:

- INST1_DB1_VOL: Contains the SQL database. Each database in an SQL storage group typically resides on a separate volume.
- INST1_DB1_LOG: Contains the transaction log.
- INST1_DATA_FILES: Contains volume for Microsoft SQL Server system data files.
- INST1_REPLOG: Contains the replicator log for VVR.

Caution: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.

The mirrored quorum disk group and mirrored quorum volume will be created in “[Implementing a dynamic mirrored quorum resource](#)” on page 414.

Use the following procedures to create the appropriate disk groups and volumes. This section assumes you are using one database.

Creating dynamic disk groups

When the tasks described in this section are completed, you will have a dynamic cluster disk group with volumes on shared storage. The dynamic cluster disk groups will be ready to be shared between nodes in the cluster.

Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it will be required later during the SQL the installation process.

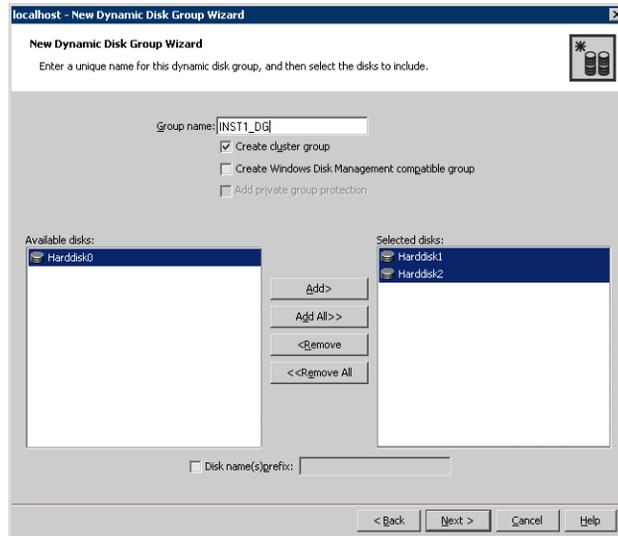
To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on client system and can be used to manage all the servers remotely.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

- Click **Next**.

- 7** Click **Next** to accept the confirmation screen with the selected disks.
- 8** Click **Finish** to create the new disk group.

Creating dynamic volumes

This section will guide you through the process of creating a volume on a dynamic disk group. When creating a disk group to support a SQL Server 2000 solution, it is best to separate SQL data files from SQL log files and place them on separate volumes. Repeat the procedure below to create the following volumes on the first node of the cluster:

- **INST1_DATA_FILES**: For storing the SQL system databases.
- **INST1_DB1_VOL**: For storing the user database.
- **INST1_DB1_LOG**: For storing the user database log.
- **INST1_REPLOG**: Contains the replicator log for VVR.

Caution: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

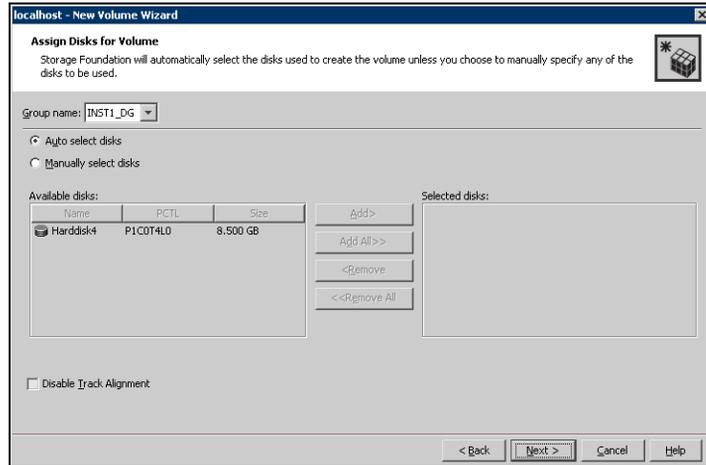
You can create the **INST1_REPLOG** volume at this time or during the process of “[Configuring VVR: Setting up an RDS](#)” on page 422.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

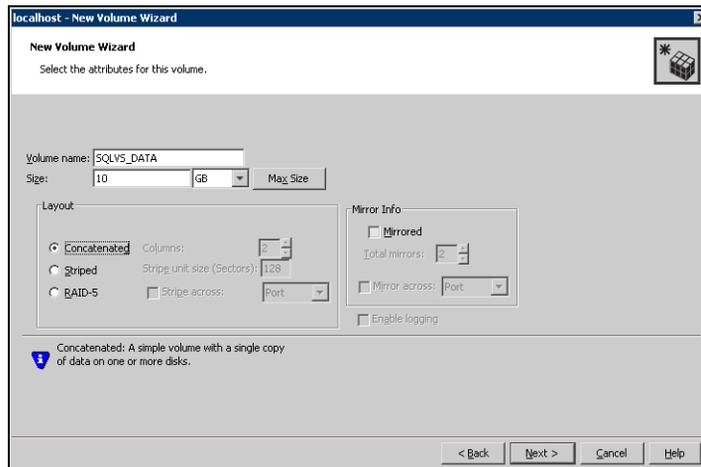
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example **INST1_DG**.

- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.



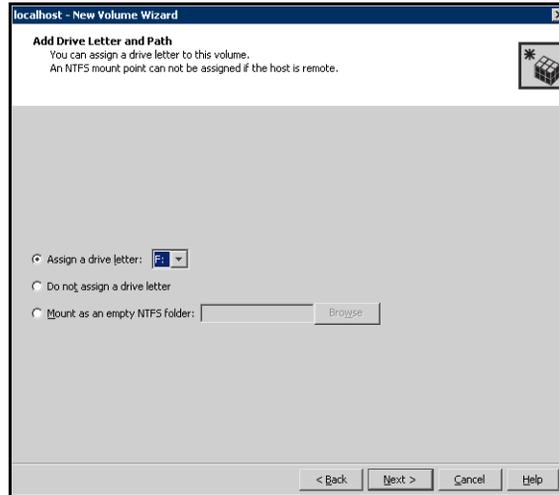
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
 You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the parameters of the volume.



- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the **Mirror Info** area, select the appropriate mirroring options.
 - Verify that **Enable Logging** is not selected.
 - Click **Next**.
- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on

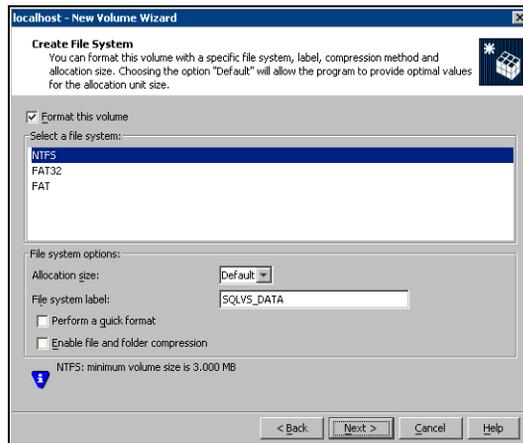
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
- For the Replicator Log volume only: Clear the Format this volume check box.
- Click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

11 Click **Finish** to create the new volume.

12 Repeat these steps to create additional volumes.

Managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter
Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Creating the SQL virtual server group

Before installing SQL you must create the SQL server cluster group and add the appropriate resources.

Note: Before creating the resources, start the cluster service on all the nodes in the cluster.

To create an SQL Server cluster group

- 1 Launch the Cluster Administrator by selecting **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator**.
Make sure you are connected to the required cluster.
- 2 Create a new group by selecting the **Groups** node from the tree that is displayed in the left hand pane. Right-click to display the **Groups** menu. Select **New > Group** option from the menu. The **New Group** window appears.
- 3 Specify a name for the group in the **Name** field.
 - In the New Group Wizard specify a name `SQL_GROUP` for the SQL cluster group.
 - If required specify a description for this resource in the field provided. Click **Next**.
- 4 The Preferred Owners page appears. Make sure that all the preferred owners are added to the **Preferred Owners** list.
- 5 Click **Finish** to create the group.

You can now start adding resources to it.

Creating an IP address resource

A separate valid IP address for the SQL virtual server is necessary to install SQL Server on more than one node.

To create an IP address resource

- 1 Right click on the Volume Manager disk group, **INST1_DG** from the example, and select **New > Resource**.
- 2 In the Resource creation wizard, configure the IP address.
 - Specify a name for the **IP Address** resource.
 - Add a **Description** if required.
 - Select the **IP address** from the **Resource Type** field drop down list.
 - Click **Next**.
- 3 In the **Possible Owners** page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 4 In the **Dependencies Page**, make sure the **Resource Dependencies** pane is empty, and click **Next**.
- 5 On the **TCP/IP Address Parameters** page, set the TCP/IP parameters.
 - Enter the IP address.
 - Enter the corresponding subnet mask.
 - Make sure the Network is set to **Public**.
 - Click **Finish** to create the **IP Address** resource.
- 6 Bring the resource online.

Creating the disk group resource

SQL virtual server installation requires a separate volume, **INST1_DATA_FILES** on which the system database files will be placed. You must create a Volume Manager Disk Group resource for the disk group that contains this volume. Creating this resource will enable SQL to monitor the system database files.

To create the disk group resource

- 1 If your cluster administrator is already open then proceed to the Step 2. To launch the Cluster Administrator select from **Start > Setting > Control Panel > Administrative Tools > Cluster Administrator**. You can create a short cut for the cluster administrator on the desktop to avoid accessing it every time from this path.

- 2 In the left pane of the cluster administrator select the SQL_GROUP Group and right-click. Select **New > Resource** from the menu that appears. The New Resource wizard appears.
- 3 Specify a name for the disk group resource, for example, SQL_DG_RES in the **Name** field.
If required, you can add a description about the resource in the **Description** field.
Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource Type** field drop down list.
Click **Next**. The Possible Owners page appears.
- 4 By default, all the nodes in the cluster are listed as possible owners. Click **Next**. The **Dependencies** page appears.
- 5 On the dependencies page, click **Next**. You do not need to set any dependency for a Disk Group resource.
- 6 On the **Volume Manager Disk Group Parameters** page select the created disk group. Click **Finish**.

The specified disk group resource, SQL_DG_RES resource is created under the SQL_GROUP group.

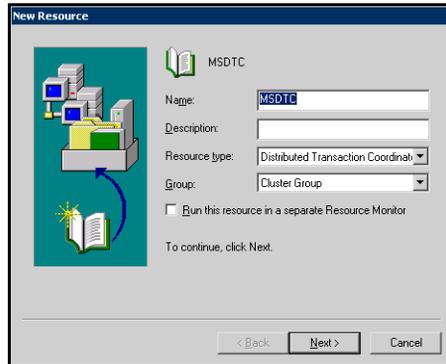
Creating the MSDTC resource

Prior to installing SQL Server, create the MSDTC resource. This procedure is required for multiple instances of SQL.

To create the MSDTC resource

- 1 From Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**), right-click **Cluster Group**, click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify a name for the MSDTC resource.

If necessary, add a description about the resource.



- 3 Select **Distributed Transaction Coordinator** from the **Resource type** list and click **Next**.
- 4 In the Possible Owners dialog box, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 5 In the Dependencies dialog box, select the cluster IP address, cluster name, and physical disk resources from the **Available Resources** list and add them to the **Resource dependencies** list. The volume with the SQL Server system data files must be included. Click **Finish**.
- 6 Click **OK**.
- 7 Bring the MSDTC resource online. In the left pane, expand the Groups icon.
- 8 Click **Cluster Group**.
- 9 Right-click **Bring Online**. The state changes to online.

Installing SQL Server 2000

This section provides some useful tips on how to install SQL Server 2000 on the primary and secondary sites. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

Before you proceed with installing SQL note the following points:

- The SQL Setup program detects that the system is a cluster, and prompts you for information appropriately at the beginning of the install process.
- The Setup program automatically installs a new, separate instance of SQL Server binaries on the local disk of each server in the cluster. The binaries are installed in exactly the same path on each cluster node, so it is important to ensure that each node has a local drive letter in common with all the other nodes in the cluster.
- The Setup program also installs the system databases on the specified cluster (shared) disk. System database files must be on a clustered disk so that they can be shared between the nodes (and failed over when necessary), because these databases contain specific user login and database object information that must be the same for each node. The virtual server name will allow users access to the online node.

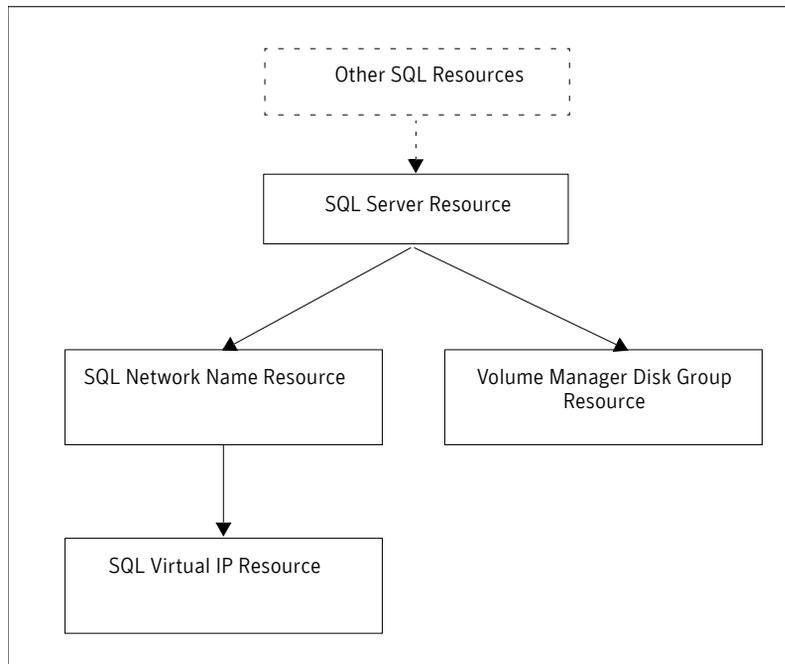
Caution: Installation of a named instance of SQL Server 2000 virtual server on a Windows 2003-based cluster will fail. See: <http://support.microsoft.com/kb/815431>

To install SQL Server

- 1 Verify the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters) See “[Managing disk groups and volumes](#)” on page 405.
- 2 Launch the Microsoft SQL Server Installation Wizard.
- 3 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 4 In the **Computer Name** wizard page select the **Virtual Server** option. This will require you to specify a name for the virtual SQL server name, for example, SQLVS. Make a note of this name as you will need to use the same name when installing on the secondary cluster nodes. Click **Next** to continue.
- 5 In the Failover Clustering wizard page you will need to specify the following information:
 - Specify the IP address that is intended for the SQL virtual server in the **IP Address** field.
 - Specify the appropriate subnet for the IP in the **Subnet** field.

- Select the appropriate public network that you have configured in the **Network to use** field. By default, the configured public network will be selected. However, if there are more than one network cards configured for public network then you can select the appropriate one from the list. Click **Next**.
- 6 The Cluster Disk Selection wizard page appears. This screen allows you to specify the logical disk from the shared disk array that will be used for the SQL Server 2000 system database files. Select the drive letter of the volume, `INST1_DATA_FILES`.
Click **Next**.
 - 7 In the Cluster Management screen specify the nodes in the cluster on which you want SQL to fail over. Make sure these are a part of the Configured Nodes box on this wizard page. Once you are sure that all the required nodes are in the Configured nodes box, click **Next**.
 - 8 The Remote Information wizard page appears. In this page specify the administrative user name and password that is valid on all the nodes. It is recommended that the domain account identification be used so that it is acceptable on all nodes. Click **Next**.
 - 9 In the Instance Name dialog box the **Default** option is selected. Since this is the first instance of SQL being installed leave the Default option selected and click **Next**.
 - 10 Follow the wizard page instructions to complete the SQL installation on all the nodes of the cluster.
Once SQL is installed, the SQL Server Resource with dependencies on the SQL Network Name and the Volume Manager Disk Group resource is automatically created. The following dependency graph indicates the dependencies that are established.

Dependency graph after the SQL installation is completed



Applying the SQL SP4

After installing SQL 2000 it is necessary to apply the SP4 patch for each of the nodes.

Verifying SQL installation

Click **Start > Programs > Microsoft SQL Server**. Select **Enterprise Manager** from the menu that appears to start the SQL Server Enterprise Manager.

Implementing a dynamic mirrored quorum resource

One of the key advantages of using SFW with MSCS is the ability to create a mirrored quorum resource that adds fault tolerance to the quorum and protects the cluster. Complete the following tasks:

- Create a dynamic cluster disk group for the quorum resource with a mirrored volume
- Create the quorum resource for the cluster group
- Change the quorum resource to a dynamic mirrored quorum resource.

Creating a dynamic cluster disk group and a mirrored volume for the quorum resource

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using three (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

For information on creating a cluster disk group and volumes, see “[Configuring SFW disk groups and volumes](#)” on page 395.

To create a three-way mirrored volume using the New Volume wizard

- 1 Create the cluster disk group with three small disks.
- 2 Create a volume with the three disks, in the sample this is INST1_QUORUM.
- 3 Select the **Concatenated** layout, select the **Mirrored** check box, and specify three mirrors.

For full details on a creating cluster disk groups and volumes, see “[Creating dynamic disk groups](#)” on page 398.

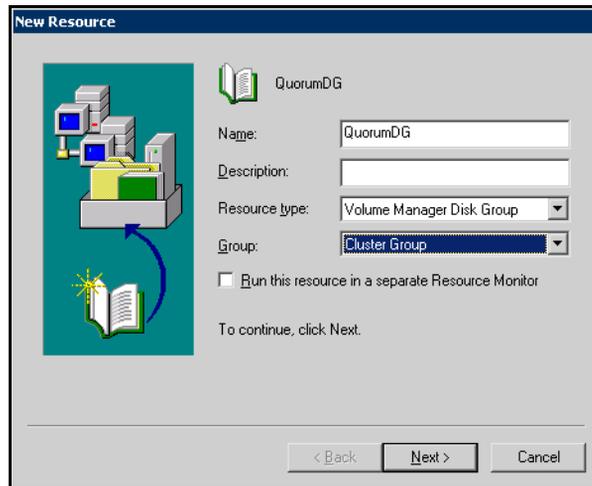
Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by MSCS.

Creating the quorum resource for the cluster group

To create a quorum resource for the cluster group

- 1 Verify that the Cluster Group is online on the same node where you created the disk group.
- 2 Create the quorum resource. Open Cluster Administrator (**Start > Control Panel > Administrative Tools > Cluster Administrator**).
- 3 In the left pane of the cluster administrator select the `SQL_GROUP` Group and right-click. Select **New**, then **Resource** from the menu that appears.
- 4 In the New Resource dialog box, specify a name for the quorum resource, for example, `QUORUM_DG`.

If necessary, add a description about the resource.



- 5 Select **Volume Manager Disk Group** from the **Resource type** list and click **Next**.
- 6 In the Possible Owners dialog box, click **Next**.
- 7 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a quorum resource.
- 8 In the Volume Manager Disk Group Parameters dialog box, select the disk group and click **Finish**.
- 9 Click **OK**.
- 10 Bring the newly added resource online.

Changing the quorum resource to a dynamic mirrored quorum resource

To change the quorum to a dynamic mirrored quorum resource

- 1 From Cluster Administrator, right-click the cluster name in the configuration tree, and click **Properties**.
- 2 Select the Quorum tab of the Properties window.
- 3 Select the name of the dynamic quorum disk group resource that was added.
- 4 Click **OK**.

Verifying the cluster configuration

You can verify your installation by moving the cluster group between nodes to see if it fails over properly. The ultimate test of the cluster's failover capability involves shutting down the node that is currently online and bringing it back up after the cluster fails over to the other node.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node. Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Creating a parallel environment on the secondary site

After setting up a SFW environment with MSCS on the primary site, complete the same tasks on the secondary site prior to the SQL installation. See “[Tasks for a new SQL Server 2000 installation with SFW, VVR, and MSCS \(Windows Server 2003\)](#)” on page 374.

During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:

- Cluster disk group name
- Volume names and sizes
- Drive letters

Before starting the installation make sure you take the SQL IP resource offline on the primary site. This will also offline the dependent resources.

Specify the same name for the SQL virtual server. Make sure the name is the same as that on the primary site.

After completing the tasks listed on page 374, you will have a clustered secondary site with:

- SFW installed
- MSCS option configured
- SQL installed on all the nodes

The next step is to set up replication between the two sites.

VVR components overview

You configure the following Veritas Volume Replicator components:

Replicated Volume Group (RVG)	<p>An RVG is made up of one or more volumes in a SFW disk group. The updates made on the RVG on the primary host are sent to a configured secondary host. Thus, on the secondary host there is a corresponding RVG with a disk group of the same name and volumes with the same names. The data volumes should be the same size. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG.</p> <p>An RVG within a disk group is the container for replication, so if you have multiple disk groups, you will need to create a separate RVG for each disk group. It is possible to have more than one RVG in a disk group; however, the RVG cannot span across disk groups.</p>
Replicated Data Set (RDS)	<p>An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).</p>
Replicator Log volume	<p>Each RVG must have a Replicator Log associated with it. The Replicator Log volume at the primary site holds a copy of any RVG updates that are sent to the secondary site. The Replicator Log on the secondary site is held in reserve so that it can be used if the primary site becomes nonfunctional and the secondary site needs to become the new primary site. The log volumes at the two sites must have the same name. Symantec recommends having Replicator Log volumes of the same size at the primary site and the secondary site.</p>

Creating resources for VVR

VVR resources must be created on the primary and secondary sites:

- IP address for replication
- Network name resource

Creating an IP address resource

A separate valid IP address for VVR replication is necessary because on the secondary cluster before a disaster, the SQL IP must be offline whereas the VVR IP must to be online.

To create an IP address resource

- 1 Right click on the Volume Manager disk group and select **New > Resource**.
- 2 In the Resource creation wizard, configure the IP address.
 - Specify a name for the **IP Address** resource.
 - Add a **Description** if required.
 - Select the **IP address** from the **Resource Type** field drop down list and click **Next**.
- 3 In the **Possible Owners** page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 4 In the **Dependencies Page**, make sure the **Resource Dependencies** pane is empty, and click **Next**.
- 5 On the **TCP/IP Address Parameters** page, set the TCP/IP parameters.
 - Enter the IP address.
 - Enter the corresponding subnet mask.
 - Set the Network to **Public** and click **Finish** to create the **IP Address** resource.
- 6 Bring the resource online.

Creating a network name resource

To create a network name resource

- 1 Right-click on the `SQL_GROUP` group and select **New > Resource**.
- 2 In the Resource creation wizard, create a Network Name resource.
 - Specify the **Network Name**.
Add a **Description** if required.
 - Specify the resource type by selecting **Network Name** from the Resource Type field drop down list and click **Next**.
- 3 In the **Possible Owners** page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 4 On the **Dependencies** page, select the IP Address resource you just created for the RVG from the **Available Resources** pane.
- 5 Add it to the **Resource Dependencies** pane and click **Next**.
- 6 In the **Name** field on the **Network Name Parameters** page, specify any name except the node and SQL Virtual Server names. Click **Finish**.

Note: The network name for the RVG must be different for the primary and secondary cluster.

- 7 Repeat the same procedure to create the IP and the Network Name resource for the secondary site.
- 8 Bring the resources online.

Configuring VVR: Setting up an RDS

For each disk group you created for the application, you set up a Replicated Data Set (RDS) on the primary and secondary hosts. The Setup Replicated Data Set Wizard enables you to configure an RDS for both sites.

Before running the wizard, verify the following:

- Verify that the disk groups and volumes for the SQL user database files and log files have been created. The Replicator Log volume can be created while running the wizard if not created earlier.
- Verify that VxSAS has been configured.
- Verify that the SQL virtual server IP resource is offline on the secondary site. This would also offline all the dependent SQL resources.

VVR does not support these types of volumes:

- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with the Dirty Region Log (DRL)
- Volumes with a comma in their names
- For the Replicator Log volume, in addition to the above types also make sure that the volume does not have a DCM.

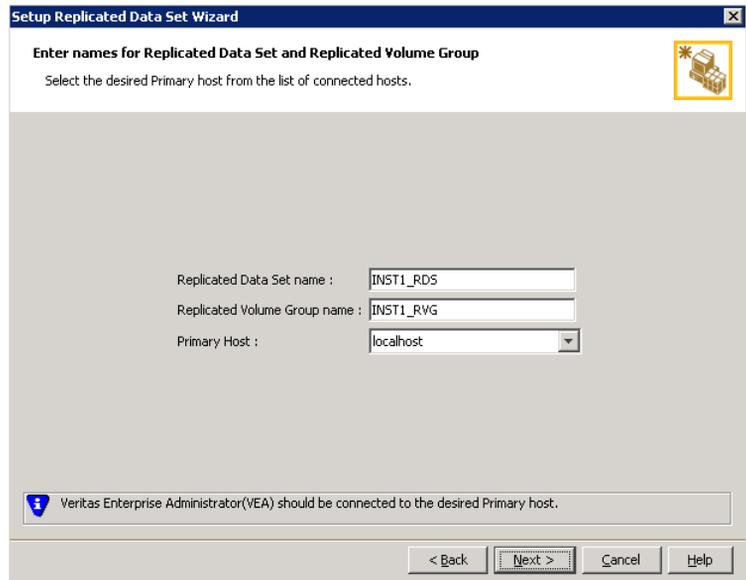
Caution: Do not use volume types that are not supported by VVR.

The following procedure enables you to set up an RDS on the primary and secondary sites and to start replication.

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.
- 3 Read the Welcome page and click **Next**.

- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

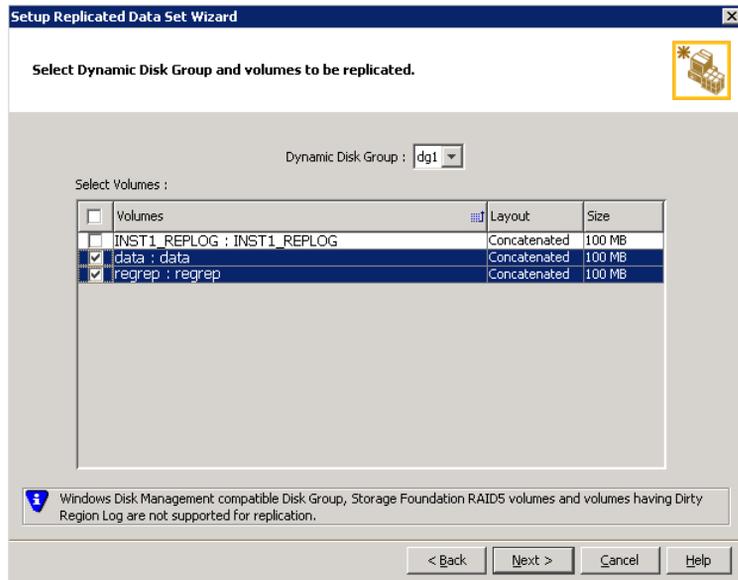


By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.

- 5 Click **Next**.

- 6 Select from the table the dynamic disk group and data volumes that will undergo replication.

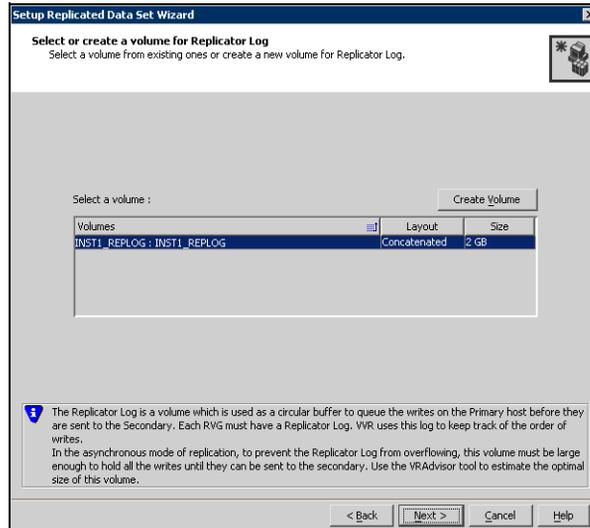


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

- 7 Click **Next**.

8 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name Enter the name for the volume in the **Name** field.

Size Enter a size for the volume in the **Size** field.

Layout Select the desired volume layout.

- Disk Selection**
- Choose **Select disks automatically** if you want VVR to select the disks for the Replicator Log.
 - Choose **Select disks manually** to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select **Add** to move the disks into the selected disks pane.

- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 9 Review the information on the summary page and click **Create Primary RVG**.
 - 10 After the Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?
Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.
Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.
 - 11 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field. Click **Next**. If the Secondary host is not connected to VEA, the wizard tries to connect it when you click **Next**. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.
 - 12 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

 - the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the Primary

Otherwise, the RDS setup wizard enables you to create the required volumes manually.

 - Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
 - 13 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.

This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.

 - If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.

- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.

Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.

When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.

14 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

Setup Replicated Data Set Wizard

Edit replication settings
 Edit replication settings or click next.

Primary side IP: 10.217.53.214

Secondary side IP: 10.217.53.215

Replication Mode: Synchronous Override

Replicator Log Protection: AutoDCM

Primary RLINK Name: Pri_RLINK

Secondary RLINK Name: Sec_RLINK

Advanced

DHCP addresses are not supported by VVR.

< Back Next > Cancel Help

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not

wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP	Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Secondary side IP	Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Replication Mode	Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous . The default is synchronous override. Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously. If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as MISSING.
Replicator Log Protection	The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **DCM** option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.

The **Off** option disables Replicator Log Overflow protection.

In the case of the Bunker node. Replicator Log protection is set to **Off**, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

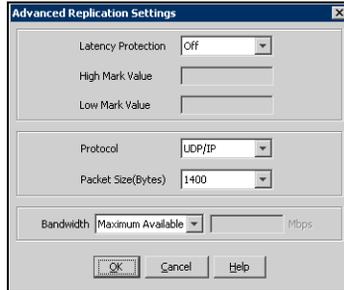
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK Name	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- 15 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



- Latency protection** Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.
- **Off** is the default option and disables latency protection.
 - **Fail** enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.
 - **Override** enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Caution: Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

- High Mark Value** Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Click **OK** to close the dialog box.

16 Click **Next**.

17 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from
Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.
- 18** Review the information.
Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

Creating the RVG resource (primary and secondary sites)

To enable a disaster recovery setup, once VVR is configured you will need to create an RVG resource with dependency on the VVR IP resource and the SQL group IP address resource. You will then need to modify the dependency in the SQL group IP address resource to remove the dependency on the resource and add dependency on the RVG resource.

To create a Replicated Volume Group (RVG) resource

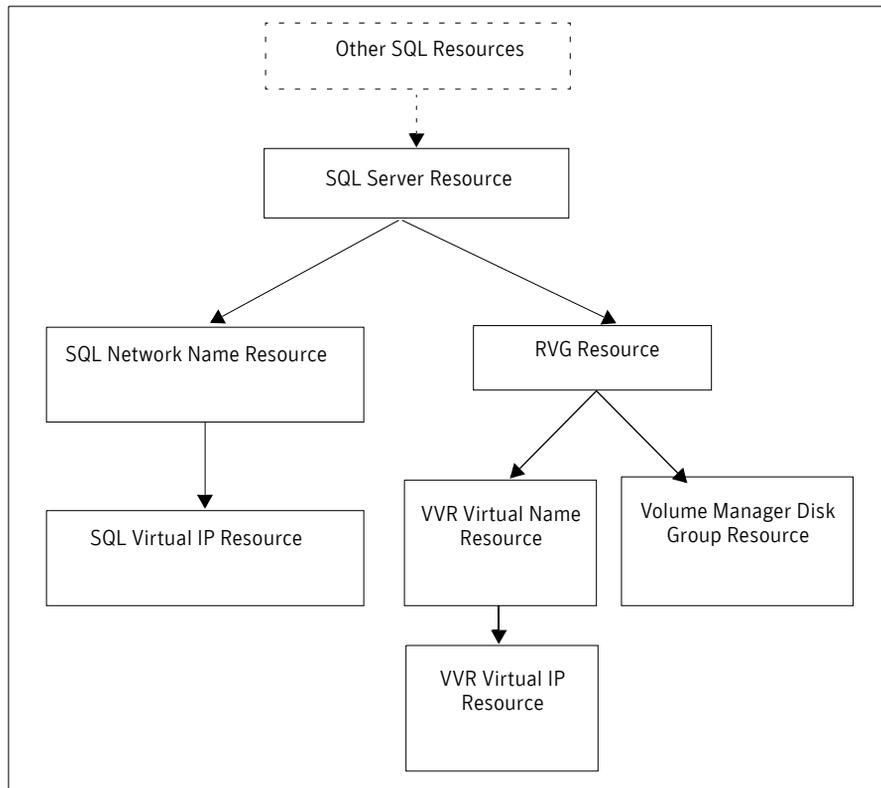
- 1 Right click on the `SQL_GROUP` group that you have created and select **New > Resource**. The New Resource wizard appears.
- 2 Specify a name for the **Replicated Volume Group** resource in the **Name** field. If required, you can add a description about the resource in the **Description** field.
Specify the resource type by selecting **Replicated Volume Group** from the **Resource Type** field drop down list. Click **Next**. The **Possible Owners** page appears.
Configure a separate resource monitor process for the RVG resource by selecting the **Run this resource in a separate Resource Monitor** checkbox provided in the **New Resource** wizard.
- 3 By default all the nodes in the cluster are listed as possible owners. Click **Next**. The **Dependencies** page appears.
- 4 On the dependencies page, select the VVR IP resource, and the Disk Group resource, from the **Available Resources** and add it to **Resource Dependencies**. Click **Next**.
- 5 On the **Replicated Volume Group Parameters** page select the created RVG. Click **Finish**.
- 6 Repeat the steps to create the RVG resource on the secondary site.

Setting the SQL server resource dependency on the RVG resource

To set the SQL server resource dependency on the RVG resource

- 1 Make sure the **SQL Server** resource is offline before attempting to modify the dependencies.
- 2 Right-click on the **SQL Server** resource and select **Properties > Dependencies** tab. This will display the Dependencies page.
- 3 Click **Modify**. Select the **Replicated Volume Group** resource from the **Available Resources** and add it to **Resource Dependencies**. Remove the **Disk Group** resource from **Resource Dependencies**. Click **OK**.

The cluster configuration is now complete. Online the entire `SQL_GROUP` group on the primary cluster. The following dependency graph indicates the dependencies that have been established.



Working with a solution: Normal operations and recovery procedures

This section gives considerations for normal VVR operations and also describes the recovery process.

Monitoring the status of the replication

Under normal operating conditions you can monitor the status of the replication using:

- VEA GUI
- Command Line Interface (CLI)
- Performance Monitor (perfmon)
- Alerts

For details, refer to the “Monitoring Replication” Chapter in the *Veritas Volume Replicator Administrator’s Guide*.

Performing planned migration

For maintenance purposes, or for testing the readiness of the secondary host you may want to migrate the application to the secondary host. These are a generic set of tasks that you may need to perform.

To detach the user database

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Expand the Databases node. Right-click on the required user database and select **All Tasks > Detach**.

Note that the `master`, `model`, and `tempdb`, databases cannot be detached.

To take the RVG resource offline

Take the RVG resource offline on both the clusters.

To transfer the Primary role

Transfer the Primary role to the Secondary using the **Migrate** option.

- 1 From the **VEA** screen, right-click the Primary RVG and select **Migrate**.
- 2 Select the Secondary host and click **OK**. The replication role is migrated to the Secondary host.

To assign drive letters to the volumes

Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as that of the original Primary.

To bring the RVG resource online

- 1 Bring the RVG resource online on both the clusters.
- 2 Bring the `SQL_GROUP` group online on the new Primary.

To attach the databases

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click on the required Database node and select **All Tasks > Attach Database**.
- 4 In the Attach Database dialog specify the name of the Master Data File (MDF) file that corresponds to the database, which you want to attach. Use the browse (...) button to search. For more information refer to the Microsoft documentation.

You can now verify that the SQL runs fine on the new Primary with the replicated data. After verifying you can revert back the roles to its original state using the same set of tasks described above.

Note: Any changes that you make to the data on the new Primary will get replicated to the original Primary, which is now the Secondary.

Replication recovery procedures

This section provides information on bringing up an SQL server on the Secondary host, in the event of a disaster. It also explains how to migrate the Primary role back to the original primary host once it is in a good state after a disaster.

Bringing up SQL on the secondary host

To recover the SQL data

- 1 From the left-pane in the VEA GUI console on the Secondary host, right-click on the desired secondary RVG node inside the replication network. Select the **Take Over** option. The **Take Over** dialog box is displayed.
 - By default, the **Enable Fast-Failback Logging** option is selected if the data volumes have DCM logs associated with them. You can use this option to perform takeover with fast-failback logging.
The DCM is activated for fast-failback logging and the new incoming writes are marked on the DCM of the New Primary.
If the replication status of Secondary RVG was `Inactive` when the Primary failed, then the **Enable Fast-Failback Logging** option is unavailable for selection. In this case you can perform **Take Over** without using fast-failback logging.
 - Select the **Synchronize Automatically** option if you want the new Primary and the original Primary to get synchronized automatically, after the original Primary recovers.
If you have not selected this option, the original Primary, after it recovers will be in the `Acting as Secondary` state. To synchronize this original Primary with the new Primary use the **Resynchronize Secondaries** option from new Primary RVG's right-click menu. When the resynchronization starts, the original Primary which was in the `Acting as Secondary` state is converted to a Secondary of the new Primary. The new Primary now starts replaying the DCM to update the Secondary with the writes that were written to the DCM.
- 2 If you do not want to use the **Enable Fast-Failback Logging** option, clear the checkbox, and click **OK** to perform Take Over without the fast-failback logging.
After takeover is complete, to add the Secondary hosts of the original Primary as Secondary hosts of the new Primary, delete the existing RVGs of the original Secondary hosts and then add them as a part of the new Primary.

- 3 If you have chosen to perform the Take Over operation without using fast-failback logging and the original Primary becomes available again, convert it to a Secondary using the **Make Secondary** option. Then resynchronize the original Primary with the new Primary using the **Synchronize Automatically** option. Depending on the size of the data volume this may take quite a while.
Only after the synchronization is complete can you migrate the Primary role back to the original Primary.
After takeover, the existing Secondary becomes the new Primary.
- 4 Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as that of the original Primary.
- 5 Bring the SQL_GROUP group online.

To attach the databases

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click on the required Database node and select **All Tasks > Attach Database**.
- 4 In the Attach Database dialog specify the name of the Master Data File (MDF) file that corresponds to the database, which you want to attach. Use the browse (...) button to search. For more information refer to the Microsoft documentation.

Now you can start using SQL on the new Primary.

Restoring the primary host

After a disaster, if the original Primary becomes available again you may want to revert the role of the Primary back to this host.

To detach the user database on the new Primary

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Expand the Databases node. Right-click on the required user database and select **All Tasks > Detach**.
Note that the `master`, `model`, and `tempdb`, databases cannot be detached.

To migrate the Primary role back to the original Primary

- 1 Take the RVG resource offline on both the clusters.
- 2 Depending on whether you performed Takeover with or without fast-failback option do one of the following.

Takeover with the Fast-failback option

The original Primary, after it has recovered, will be in the **Acting as Secondary** state. If the original Primary is not in the **Acting as Secondary** state, verify whether your network connection has been restored.

To synchronize this original Primary and the new Primary, use the **Resynchronize Secondaries** option from new Primary's right-click menu.

Takeover without the Fast-failback option

After performing a takeover without fast-failback, you must convert the original Primary to a Secondary using the **Make Secondary** option.

Note: Before performing the **Make Secondary** operation, the original Primary's RVG and the new Primary's RVG will be shown in separate RDSs. However, after this operation they will be merged under a single RDS.

After the **Make Secondary** operation, the original Primary will be converted to a secondary. Right-click on this secondary RVG and select **Start Replication with Synchronize Automatically** option

- 3 After the synchronization is complete, perform a migrate operation to transfer the Primary role back to the original Primary. To do this, right-click on the Primary RVG and select **Migrate** option from the menu that appears.
- 4 Ensure that the volumes have retained the same drive letters that existed before the disaster.
- 5 Bring the RVG resource online on the Secondary.
- 6 Bring the SQL_GROUP group online on the original Primary.

To attach the databases on the original Primary

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click on the required Database node and select **All Tasks > Attach Database**.

- 4 In the Attach Database dialog, specify the name of the Master Data File (MDF) file that corresponds to the database, which you want to attach. Use the browse (...) button to search. For more information refer to the Microsoft documentation.

Deploying SFW and VVR with MSCS: New SQL 2005 installation

This chapter covers the following topics:

- [Tasks for a new SQL Server 2005 installation with SFW, VVR, and MSCS \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing an MSCS cluster](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Creating the SQL virtual server group](#)
- [Creating the MSDTC resource](#)
- [Installing SQL Server 2005](#)
- [Implementing a dynamic quorum resource](#)
- [Verifying the cluster configuration](#)
- [Creating a parallel environment on the secondary site](#)
- [VVR components overview](#)
- [Creating resources for VVR](#)
- [Configuring VVR: Setting up an RDS](#)

- [Creating the RVG resource \(primary and secondary sites\)](#)
- [Setting the SQL server resource dependency on the RVG resource](#)
- [Working with the solution: Normal operations and recovery procedures](#)

Tasks for a new SQL Server 2005 installation with SFW, VVR, and MSCS (Windows Server 2003)

You can install and configure Storage Foundation for Windows and Veritas Volume Replicator (VVR) with MSCS and SQL Server 2005 on Windows Server 2003. You can install and configure high availability, replication, and SQL components. This environment involves an active/passive configuration with one to one failover capability. After setting up a SFW environment with high availability for SQL on a primary site, you can create a secondary or “failover” site for replication. Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details on VVR.

If you are deploying SQL Server 2005 on Windows Server 2008, see:

[Chapter 18, “Deploying SFW and VVR with Microsoft failover clustering: New SQL 2005 installation” on page 511.](#)

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 17-1 Tasks for deploying SFW and VVR with MSCS for SQL Server 2005

Objective	Tasks
“Reviewing the requirements” on page 444	<ul style="list-style-type: none">■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 446	<ul style="list-style-type: none">■ Understanding a typical Active/Passive SQL configuration in a two-node cluster
“Configuring the storage hardware and network” on page 449	<ul style="list-style-type: none">■ Setting up the storage hardware for a cluster environment■ Verifying the DNS entries for the systems on which SQL will be installed
“Establishing an MSCS cluster” on page 450	<ul style="list-style-type: none">■ Reviewing general guidelines to establish an MSCS cluster

Table 17-1 Tasks for deploying SFW and VVR with MSCS for SQL Server 2005

Objective	Tasks
“Installing SFW with MSCS/Failover Cluster option” on page 452	<ul style="list-style-type: none"> ■ Modifying the driver signing options for Windows 2003 remote systems ■ Installing SFW ■ Configuring VxSAS ■ Restoring the driver signing options for Windows 2003 remote systems
“Configuring SFW disk groups and volumes” on page 464	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes
“Creating the SQL virtual server group” on page 476	<ul style="list-style-type: none"> ■ Creating a SQL Server cluster group ■ Creating a virtual server IP address ■ Creating the disk group resource
“Creating the MSDTC resource” on page 478	<ul style="list-style-type: none"> ■ Creating the MSDTC resource for SQL Server
“Installing SQL Server 2005” on page 479	<ul style="list-style-type: none"> ■ Installing SQL ■ Verifying SQL installation
“Implementing a dynamic quorum resource” on page 484	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the Cluster Group ■ Changing the quorum resource to a dynamic mirrored quorum resource.
“Verifying the cluster configuration” on page 486	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node
“Creating a parallel environment on the secondary site” on page 488	<ul style="list-style-type: none"> ■ Using the previous objectives for the secondary site from reviewing the prerequisites through testing the cluster.
“Creating resources for VVR” on page 490	<ul style="list-style-type: none"> ■ Creating an IP address for the Replicated Volume Group (RVG). ■ Creating a network name resource for the Replicated Volume Group (RVG)
“Configuring VVR: Setting up an RDS” on page 492	<ul style="list-style-type: none"> ■ Create an RDS using the VVR wizard.

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

Supported software for MSCS with VVR

- Veritas Storage Foundation 5.1 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster and the Veritas Volume Replicator Option
- Microsoft SQL servers and their operating systems:

Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required)

- Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required)
- Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required)
- Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)
- Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)
- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)

Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required)

- Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)
- Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)
- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 17-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

- One CD-ROM drive accessible to the system on which you are installing SFW.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- MSCS requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft SQL Server documentation for instructions on creating a reverse lookup zone.
- MSCS requires two disks for SQL: one for SQL database files and one for SQL log files.
- Each system requires 1 GB of RAM for SFW.
- SFW requires administrator privileges to install the software.
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.

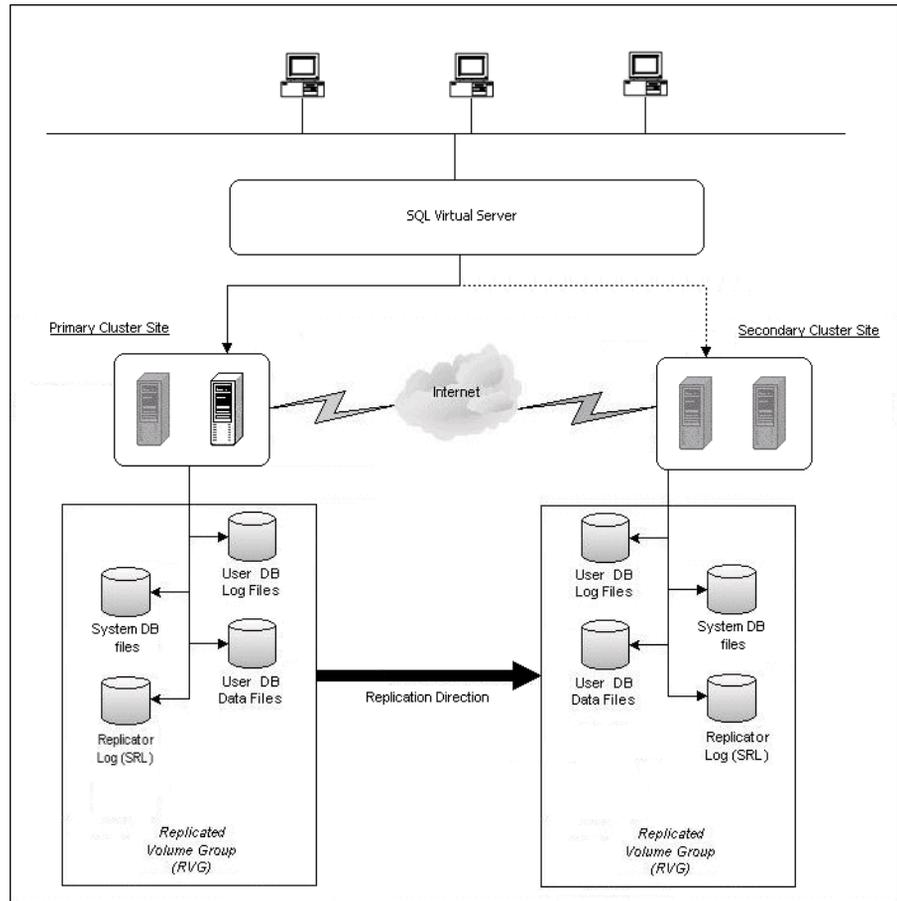
Note: Refer to the Hardware Compliance List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Reviewing the configuration

The following figure illustrates a typical clustered VVR configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2005 application data is stored on the volumes that are under the control of the RVG. A separate disk group is created for the quorum volume which is not replicated.

[Figure 17-1](#) shows a typical VVR configuration.

Figure 17-1 Typical VVR configuration



If the Microsoft SQL Server 2005 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the replication solution is activated. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over. The data that was replicated to the secondary site is used to restore the SQL services to clients.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site. The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary Site

SYSTEM1 & SYSTEM2	server names
SQL_GROUP	Microsoft SQL Server virtual server group
SQLCLUST	Microsoft SQL Server virtual cluster (underscores not supported)
SQLVS	Microsoft SQL Server virtual server
SQL_IP	Microsoft SQL virtual server IP address resource
INST1	Microsoft SQL Server instance name
INST1_DG	disk group for Microsoft SQL volumes
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
SQLVS_QRM	volume for storing the MSCS cluster quorum
QUORUM_DG	quorum volume disk group for mirroring the quorum

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
----------------------	--

All the other parameters are the same as on the primary site.

DR Components

INST1_RDS	VVR Replicated Data Set (RDS) name
INST1_RVG	VVR Replicated Volume Group (RVG) name
INST1_REPLOG	VVR Replicator log volume

INST1_RVG_RES	MSCS Replicated Volume Group Resource name
VVR_IP	Microsoft SQL RVG IP address resource

Configuring the storage hardware and network

Use the following procedures to configure the storage hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent MSCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 5 In the Public Status dialog box, on the General tab, click **Properties**.

- 6 In the Public Properties dialog box, on the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Establishing an MSCS cluster

Before installing SFW, you must establish an MSCS cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To establish an MSCS cluster (general guidelines)

- 1 Verify that the quorum disk has been created before installing MSCS and configuring a cluster. (For IA64 systems, the quorum must be created using MBR instead of GPT or it will not be visible.)
- 2 Configure the shared storage and create a partition with drive letter "Q" for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster (SYSTEM1) using MSCS Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Verify that the node can access the shared storage.
- 4 Connect the second node to the shared storage.
- 5 Add the second node (SYSTEM2) using Cluster Administrator on that system.
- 6 Test the cluster by using the Move Group command to move the cluster resources to the second node.
SYSTEM2 becomes the active cluster node.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running a Microsoft cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft cluster simultaneously.

SFW installation tasks

The product installer enables you to install the software for Veritas Storage Foundation 5.1 for Windows. The installer automatically installs SFW. You must select the options to install VVR, and the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation. The steps in this section are based on a server installation.

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 452.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 454.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 459.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
- Moving the online groups

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 17-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are online on the system where you are installing SFW. You must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install Veritas Storage Foundation for Windows on a Microsoft cluster configuration.

To install the product

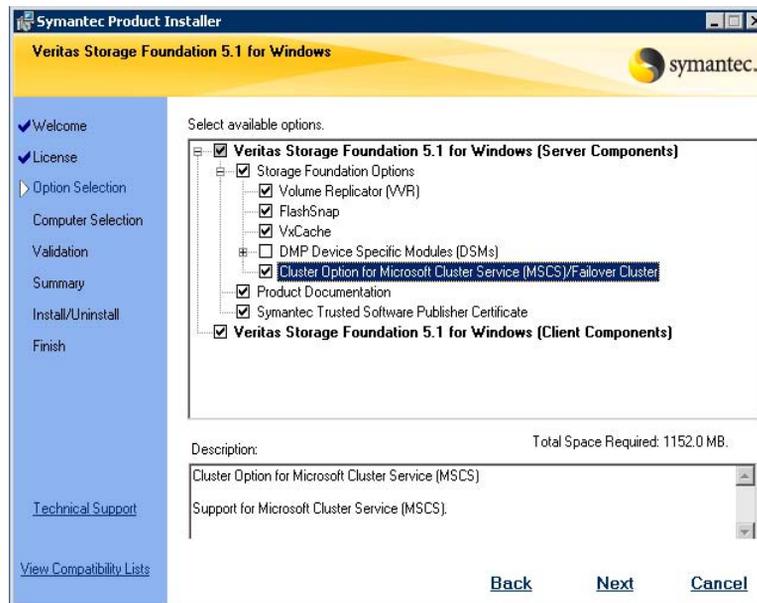
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

3 Click **Storage Foundation 5.1 for Windows**.



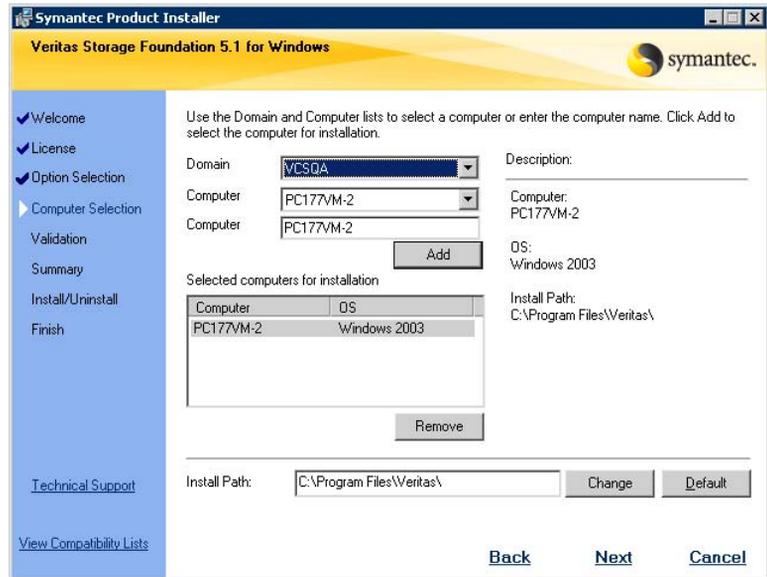
- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for "**I accept the terms of the license agreement,**" and click **Next**.
- 7 Enter the product license key before adding license keys for features.
- 8 Enter the license key in the top field and click **Add**.
- 9 Repeat for additional license keys.
To remove a license key, click the key to select it and click **Remove**.
To see the license key's details, click the key.
- 10 Click **Next**.

11 Specify the product options:



- 12 Select the **Volume Replicator (VVR)** option.
- 13 Select the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.
- 14 Do not select the Dynamic Multi-pathing option.
Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.
Click **Next**.
- 15 To install the client components on all the computers in your installation, verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is selected, to install the client component.
Click **Next**.

16 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

- Install Path
- Optionally, change the installation path.
- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
 - To restore the default path, select a computer and click **Default**.
The default path is:
C:\Program Files\Veritas
For 64-bit installations, the default path is:
C:\Program Files (x86)\Veritas

17 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.

18 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.

19 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that MSCS allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

20 Review the information and click **Install**. Click **Back** to make changes.

21 The Installation Status screen displays status messages and the progress of the installation.

If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.

If the installation is successful on all systems, the installation report screen appears.

If your local computer has its driver signing options set to Block or Warn then installation fails.

- 22 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 23](#) if you are installing SFW on the local node only.
 - Proceed to [step 25](#) if you are installing SFW on local and remote systems.
- 23 To complete the installation, click **Finish**.
- 24 Click **Yes** to reboot the system and complete the installation.
- 25 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
 - Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
 - When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
 - Click **Finish**.
 - Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW Installation for the other systems in the MSCS cluster
- Configure the VxSAS service
- Resetting the driver signing options

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.
- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 452.

Configuring the VxSAS service

Complete the following procedure to configure this required service for VVR. This procedure should not be done until you have installed SFW on all cluster systems. Otherwise, you will get an error message from the VxSAS wizard if you try to select a system without SFW installed.

You can run the VxSAS wizard from the secondary site once SFW is installed on all cluster systems; at that time, you can run the wizard for both the primary and secondary site systems. The MSCS groups can be either online or offline.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. Accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

To configure the VxSAS service

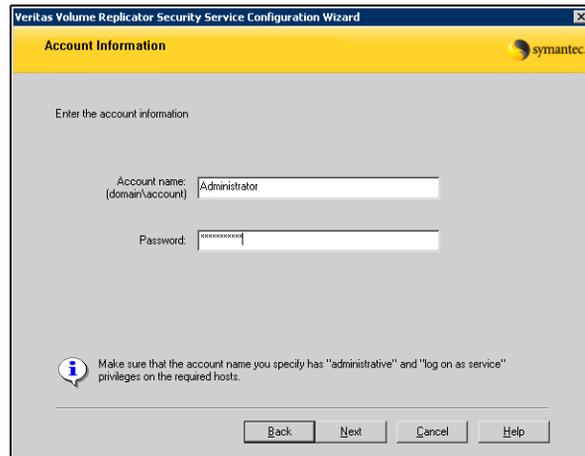
- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.

The welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

2 Complete the Account Information wizard page as follows:

- | | |
|----------------------------------|--|
| Account name
(domain\account) | Enter the administrative account name in the Account name field. |
| Password | Specify a password in the Password field. |

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same user name and password when configuring the VxSAS service on the other hosts.



3 After providing the required information, click **Next**.

4 Select the required domain to which the hosts that you want to configure belong, from the Domain Selection wizard page.

Selecting Domains The Available Domains pane lists all the domains that are present in the Windows network neighborhood.
 Select the required domain by moving the appropriate name from the Available Domains pane to the Selected Domains pane, either by double-clicking it or using the arrow button.

Adding a Domain If the domain name that you require is not displayed, then add it by using the **Add Domain** option. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected Domains list.

After specifying the domain, click **Next**.

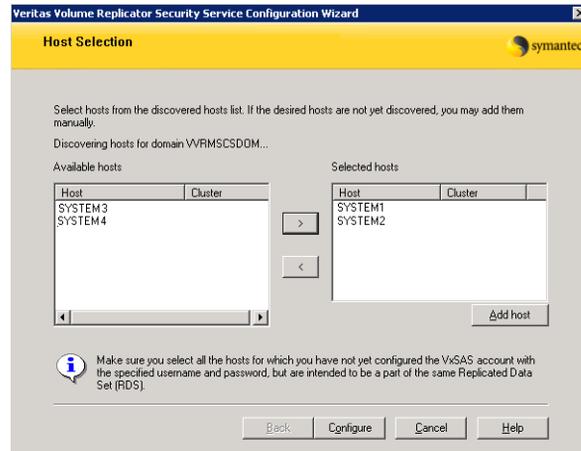
- 5 Select the required hosts from the Host Selection page.

Selecting Hosts The Available Hosts pane lists the hosts that are present in the specified domain.

Select the required host by moving the appropriate name from the Available Hosts list to the Selected Hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a Host If the host name you require is not displayed, then add it using **Add Host** option. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected Hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.



- 6 After the configuration completes, the Configuration Results page is displayed. If the operation is successful, then the Status column displays the appropriate message to indicate that the operation was successful.

If the operation was not successful, then the page displays the status as failed and the corresponding details on why the account update failed. It also displays the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

- 7 Click **Finish** to exit the wizard.

Resetting the driver signing options

You must reset the driver signing options to its previous state. This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for SQL. A dynamic disk group is a collection of one or more disks that behaves as a single storage repository. Within each disk group, you can have dynamic volumes with different layouts.

Configuring disk groups and volumes involves the following tasks:

- [“Planning disk groups and volumes”](#) on page 464
- [“Creating dynamic cluster disk groups”](#) on page 466
- [“Creating dynamic volumes”](#) on page 469
- [“Managing disk groups and volumes”](#) on page 474

Planning disk groups and volumes

Before installing SQL, you must create disk groups and volumes using the VEA console installed with SFW.

Before creating a disk group, consider:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.
- The number of disk groups and volumes that are needed for SQL Server
The number of disk groups for SQL depends on the planned organization of the data. The application program files need to be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage in a cluster disk group. You create at least one disk group for the system data files. You may want to create additional disk groups for user databases. Symantec recommends that you place SQL Server user database files and log files on separate volumes. Replicating the system databases is not required or recommended. Make sure that the system databases are not placed on volumes that will be replicated.
- The disk groups and volumes for the mirrored quorum resource
You will need a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. You can create the quorum disk group at the same time you create application disk groups, although it is not required for installing the application. You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume;

this enables you to verify that SQL is working in the cluster before adding the dynamic quorum volume.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

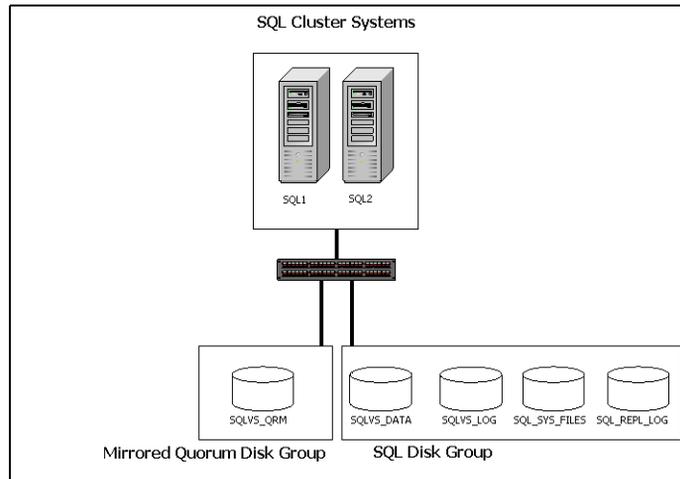
VVR does *not* support these types of volumes:

- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with the Dirty Region Log (DRL)
- Volumes with a comma in their names
- For the Replicator Log volume, in addition to the above types also make sure that the volume does not have a DCM.

Caution: Do not use volume types that are not supported by VVR.

Below is a detailed view of the disk groups and volumes for SQL:

Figure 17-2 SFW disk groups and volumes for SQL virtual server INST1-VS in MSCS setup



SQL disk group INST1_DG contains four volumes:

- **INST1_DATA_FILES:** Contains the SQL database. Each database in an SQL storage group typically resides on a separate volume.

- INST1_DB1_LOG: Contains the transaction log.
- INST1_DB1_VOL: Contains volume for Microsoft SQL Server system data files.
- INST1_REPLOG: Contains the replicator log for VVR.

The mirrored quorum disk group and mirrored quorum volume will be created in “[Implementing a dynamic quorum resource](#)” on page 484.

Use the following procedures to create the appropriate disk groups and volumes. This section assumes you are using one database.

Creating dynamic cluster disk groups

When the tasks described in this section are completed, you will have a dynamic cluster disk group with volumes on shared storage. The dynamic cluster disk groups will be ready to be shared between nodes in the cluster.

A dynamic cluster disk group is a collection of one or more disks that behave as a single storage repository and which can potentially be accessed by different computers. Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it will be required later during the SQL the installation process.

To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on client system and can be used to manage all the servers remotely.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

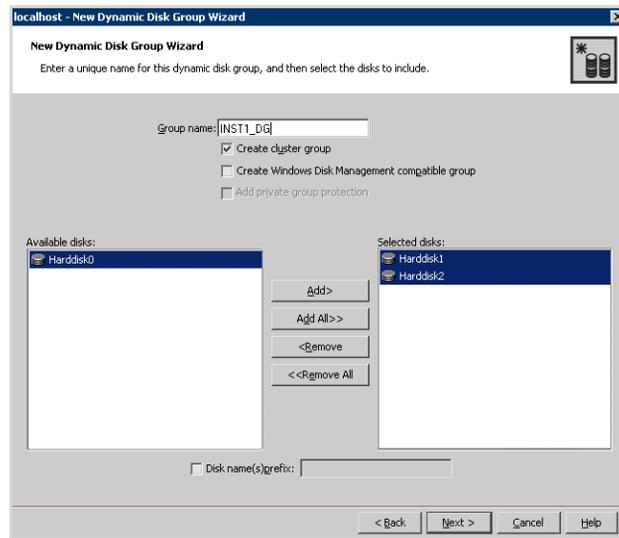
To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.

To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.
For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Creating dynamic volumes

This section will guide you through the process of creating a volume on a dynamic disk group. When creating a disk group to support a SQL Server 2005 solution, it is best to separate SQL data files from SQL log files and place them on separate volumes. Repeat the procedure below to create the following volumes on the first node of the cluster:

- INST1_DATA_FILES: For storing the SQL system databases.
- INST1_DB1_VOL: For storing the user database.
- INST1_DB1_LOG: For storing the user database log.
- INST1_REPLOG: For storing the replicator log required by VVR.

Caution: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

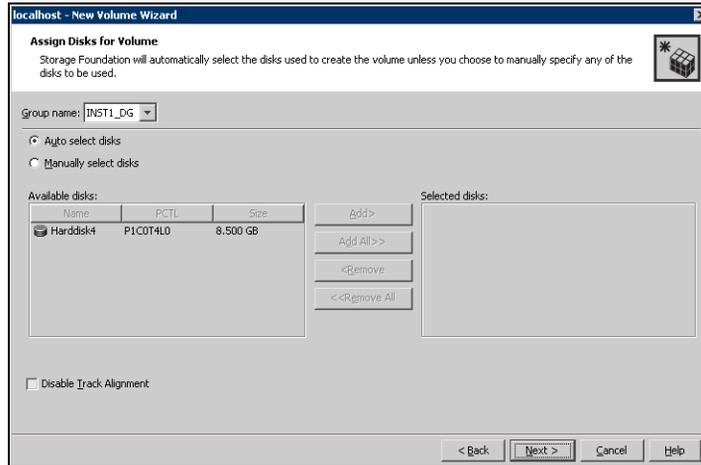
You can create the INST1_REPLOG volume at this time or during the process of [“Configuring VVR: Setting up an RDS”](#) on page 492.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

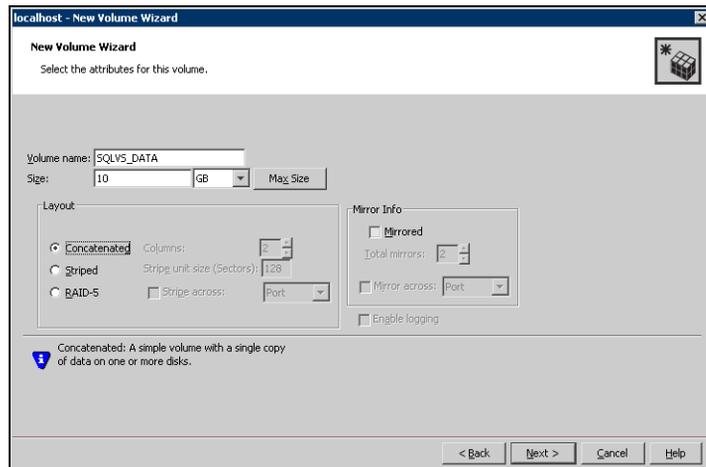
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.

- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.



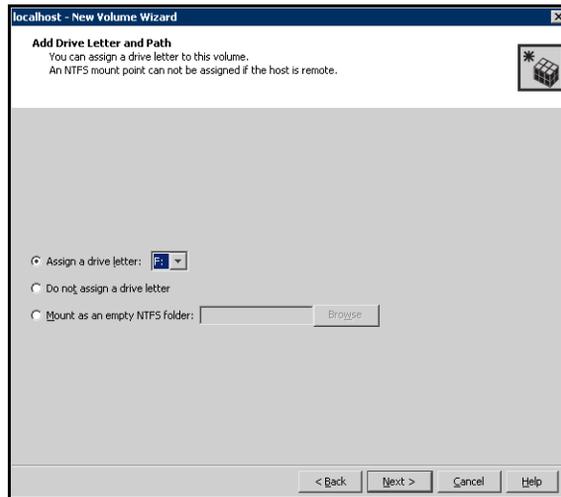
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the parameters of the volume.



- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the **Mirror Info** area, select the appropriate mirroring options.
 - Verify that **Enable Logging** is not selected.
 - Click **Next**.
- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on

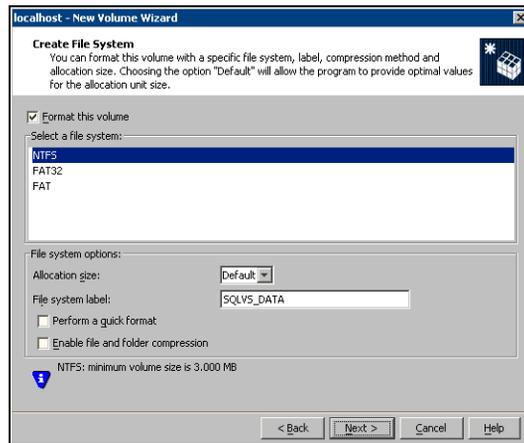
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
- For the Replicator Log volume only: Clear the Format this volume check box.
- Click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

11 Click **Finish** to create the new volume.

12 Repeat these steps to create additional volumes.

Managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter
Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Creating the SQL virtual server group

Before installing SQL Server 2005 you must create the SQL Server cluster group and add the appropriate resources.

Note: Before creating the resources, start the cluster service on all the nodes in the cluster.

To create an SQL Server cluster group

- 1 Launch the Cluster Administrator by selecting **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator**.
Make sure you are connected to the required cluster.
- 2 Create a new group by selecting the **Groups** node from the tree that is displayed in the left hand pane. Right-click to display the **Groups** menu. Select **New > Group** option from the menu. The **New Group** window appears.
- 3 Specify a name for the group in the **Name** field.
 - In the New Group Wizard specify a name `SQL_GROUP` for the SQL cluster group.
 - If required specify a description for this resource in the field provided. Click **Next**.
- 4 The Preferred Owners page appears. Make sure that all the preferred owners are added to the **Preferred Owners** list.
- 5 Click **Finish** to create the group.

You can now start adding resources to it.

Creating an IP address resource

A separate valid IP address for the SQL virtual server is necessary to install SQL Server on more than one node.

To create an IP address resource

- 1 Right click on the Volume Manager disk group, **INST1_DG** from the example, and select **New > Resource**.
- 2 In the Resource creation wizard, configure the IP address.
 - Specify a name for the **IP Address** resource and add a **Description** if required.
 - Select the **IP address** from the **Resource Type** field drop down list and click **Next**.
- 3 In the **Possible Owners** page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 4 In the **Dependencies Page**, make sure the **Resource Dependencies** pane is empty, and click **Next**.
- 5 On the **TCP/IP Address Parameters** page, set the TCP/IP parameters.
 - Enter the IP address.
 - Enter the corresponding subnet mask.
 - Make sure the Network is set to **Public** and click **Finish** to create the **IP Address** resource.
- 6 Bring the resource online.

Creating the disk group resource

SQL virtual server installation requires a separate volume, **SQL_DATA_FILES** on which the system database files will be placed. You must create a Volume Manager Disk Group resource for the disk group that contains this volume. Creating this resource will enable SQL to monitor the system database files.

To create the disk group resource

- 1 If your cluster administrator is already open then proceed to the Step 2. To launch the Cluster Administrator select from **Start > Setting > Control Panel > Administrative Tools > Cluster Administrator**. You can create a short cut for the cluster administrator on the desktop to avoid accessing it every time from this path.

- 2 In the left pane of the cluster administrator select the `SQL_GROUP` Group and right-click. Select **New > Resource** from the menu that appears. The New Resource wizard appears.
- 3 Specify a name for the disk group resource, for example, `SQL_DG_RES` in the **Name** field.
If required, you can add a description about the resource in the **Description** field.
Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource Type** field drop down list.
Click **Next**. The Possible Owners page appears.
- 4 By default, all the nodes in the cluster are listed as possible owners. Click **Next**. The **Dependencies** page appears.
- 5 On the dependencies page, click **Next**. You do not need to set any dependency for a Disk Group resource.
- 6 On the **Volume Manager Disk Group Parameters** page select the created disk group. Click **Finish**.

The specified disk group resource, `SQL_DG_RES` resource is created under the `SQL_GROUP` group.

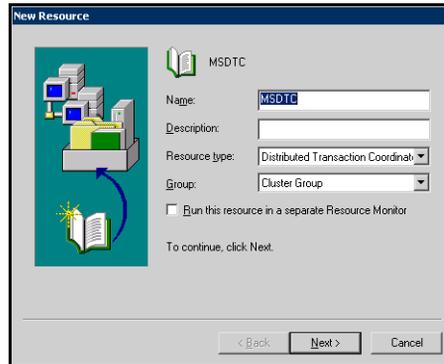
Creating the MSDTC resource

Prior to installing SQL Server, create the MSDTC resource. This procedure is required for multiple instances of SQL.

To create the MSDTC resource

- 1 From Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**), right-click **Cluster Group**, click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify a name for the MSDTC resource.

If necessary, add a description about the resource.



- 3 Select **Distributed Transaction Coordinator** from the **Resource type** list and click **Next**.
- 4 In the Possible Owners dialog box, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 5 In the Dependencies dialog box, select the cluster IP address, cluster name, and physical disk resources from the **Available Resources** list and add them to the **Resource dependencies** list. The volume with the SQL Server system data files must be included. Click **Finish**.
- 6 Click **OK**.
- 7 Bring the MSDTC resource online. In the left pane, expand the Groups icon.
- 8 Click **Cluster Group**.
- 9 Right-click **Bring Online**. The state changes to online.

Installing SQL Server 2005

This section provides some useful tips on how to install SQL Server 2005 on the primary and secondary sites. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

Before you proceed with installing SQL note the following points:

- Verify the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters) See “[Managing disk groups and volumes](#)” on page 474.
- The Setup program automatically installs a new, separate instance of SQL Server binaries on the local disk of each server in the cluster. The binaries are installed in exactly the same path on each cluster node, so it is important to ensure that each node has a local drive letter in common with all the other nodes in the cluster.
- The Setup program also installs the system databases on the specified cluster (shared) disk. System database files must be on a clustered disk so that they can be shared between the nodes (and failed over when necessary), because these databases contain specific user login and database object information that must be the same for each node. The virtual server name will allow users access to the online node.

To install SQL Server 2005

- 1 Begin the SQL Server 2005 installation, following the instructions from Microsoft. To begin the installation, navigate to the installation directory and launch **splash.hta**.
- 2 Review the hardware and software requirements for SQL 2005.
- 3 Under the **Install** section, select **Server components, tools, Books Online, and samples**.
- 4 Continue with the installation, following the instructions from Microsoft. Complete the SQL Server Component Update, System Configuration Check, and Registration Information pages.
- 5 In the Components to Install dialog box, select the **SQL Server Database Services**.
- 6 To cluster SQL Server on MSCS, select the **Create a SQL Server failover cluster** option.
- 7 Select the optional components:
 - Analysis Service. If this option is selected, the option **Create an Analysis Server failover cluster** must also be selected.
 - Notification Services
 - Integration Services
 - Workstation Components
- 8 Click the **Advanced** option.

- 9 In the **Feature Selection** dialog box, specify the path for SQL Server data files and other services.
 - Expand **SQL Server Database Services** and select **Data Files**.
 - Select **Browse** to reset the installation path.
 - Set the installation path in the Change Folders dialog box to the drive letter and location of the volume created for the SQL Server system data files (INST1_DATA_FILES). Allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

Note: This must be the same as the path on all nodes.

- If you selected the **Analysis Services** option in [step 7](#), expand **Analysis Services**, select **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **Next**.

Note: This must be the same as the path on all nodes.

- 10 In the **Instance Name** dialog box, enter an instance name or accept the default. Click **Next**.

Only one default instance is allowed per cluster.

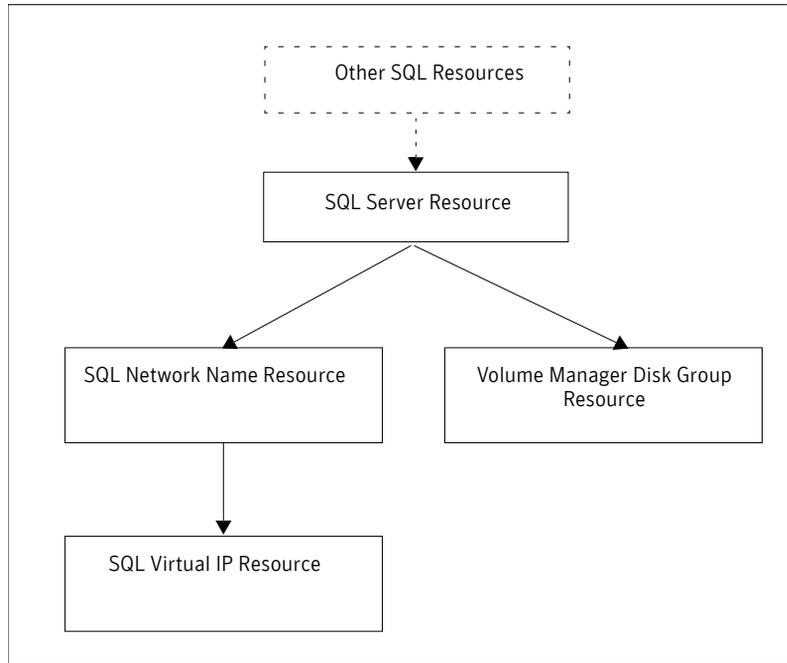
Note: Use the same instance name when installing SQL Server 2005 on the first node and on all failover nodes.

- 11 In the **Virtual Server Name** page, specify a unique name for the virtual SQL server, for example, *SQLVS*. Make a note of this name as you will need to use the same name when installing on the Secondary cluster nodes. Click **Next**.
- 12 In the **Virtual Server Configuration** page, configure the virtual server.
 - Select the appropriate public network that you have configured in the **Network to use** field. By default, the configured public network will be selected. However, if there are more than one network cards configured for public network then you can select the appropriate one from the list.
 - Specify the IP address that is intended for the SQL virtual server in the **IP Address** field.
 - The appropriate **Network address** and **Network subnet** are displayed. Click **Next**.
- 13 In the **Cluster Group Selection** page, specify the cluster group with the logical disk from the shared disk array that will be used for the SQL Server 2005 system database files. You can also specify a custom path in the **Data Files** field. Click **Next**.

- 14 In the **Cluster Node Configuration** page, specify the nodes in the cluster for SQL failover by selecting them from **Available nodes**, and clicking **Add** to add them to the **Selected Nodes** box. Once all the required nodes are in the **Selected Nodes** box, click **Next**.
- 15 In the **Remote Account Information** page, specify the administrative password that is valid on all the nodes. Click **Next**.
- 16 In the **Service Account** page, specify the type of account and information.
 - Select the type of service account. Select **Use the built-in System account** if you do not want to replicate data. Otherwise, select **Use a domain user account**.
 - If you chose to run the service in the context of a domain user, specify the information for the user. Click **Next**.
- 17 In the **Domain Groups for Clustered Services** page, use the browse button at the right to select a **DomainName** and **GroupName** for each of the selected SQL Server options. Click **Next**.
- 18 In the **Authentication Mode** page, select **Mixed Mode** (recommended option), and specify your password.
- 19 Follow the wizard page instructions to complete the SQL installation on all the nodes of the cluster.

Once SQL is installed, the SQL Server Resource with dependencies on the SQL Network Name and the Volume Manager Disk Group resource is created. The following dependency graph indicates the dependencies that are established.

Dependency graph after the SQL installation is completed.



Verifying SQL installation

Click **Start > Programs > Microsoft SQL Server**. Select **Enterprise Manager** from the menu that appears to start the SQL Server Enterprise Manager.

Implementing a dynamic quorum resource

One of the key advantages of using SFW with MSCS is the ability to create a mirrored quorum resource that adds fault tolerance to the quorum and protects the cluster. Complete the following tasks:

- Create a dynamic cluster disk group for the quorum resource with a mirrored volume
- Create the quorum resource for the Cluster Group
- Change the quorum resource to a dynamic mirrored quorum resource.

Creating a dynamic cluster disk group and a mirrored volume

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using three (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

To create a three-way mirrored volume using the New Volume wizard

- 1 Create the cluster disk group with three small disks.
- 2 Create a volume with the three disks, in the sample this is SQLVS_QRM.
- 3 Select the **Concatenated** layout, select the **Mirrored** check box, and specify three mirrors.

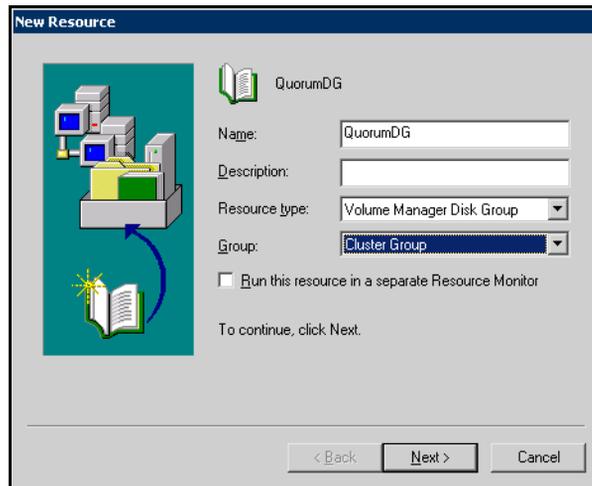
For full details on a creating cluster disk groups and volumes, see [“Creating dynamic cluster disk groups”](#) on page 466.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by MSCS.

Creating the quorum resource for the cluster group

To create a quorum resource for the cluster group

- 1 Verify that the Cluster Group is online on the same node where you created the disk group.
- 2 Create the quorum resource. Open Cluster Administrator (**Start > Control Panel > Administrative Tools > Cluster Administrator**).
- 3 In the left pane of the Cluster Administrator, select the SQL_GROUP Group and right-click. Select **New**, then select **Resource** from the menu that appears.
- 4 In the New Resource dialog box, specify a name for the quorum resource, for example, QUORUM_DG.
 If necessary, add a description about the resource.



- 5 Select **Volume Manager Disk Group** from the **Resource type** list. Click **Next**.
- 6 In the Possible Owners dialog box, click **Next**.
- 7 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a quorum resource.

- 8 In the Volume Manager Disk Group Parameters dialog box, select the disk group. Click **Finish**.
- 9 Click **OK**.
- 10 Bring the newly added resource online.

Changing the quorum resource to a dynamic mirrored quorum resource

To change the quorum to a dynamic mirrored quorum resource

- 1 From Cluster Administrator, right-click the cluster name in the configuration tree, and click **Properties**.
- 2 Select the Quorum tab of the Properties window.
- 3 Select the name of the dynamic quorum disk group resource that was added.
- 4 Click **OK**.

Verifying the cluster configuration

You can verify your installation by moving the cluster group between nodes to see if it fails over properly. The ultimate test of the cluster's failover capability involves shutting down the node that is currently online and bringing it back up after the cluster fails over to the other node.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.

- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Creating a parallel environment on the secondary site

After setting up a SFW environment with MSCS on the primary site, complete the same tasks on the secondary site prior to the SQL installation. See [“Tasks for a new SQL Server 2005 installation with SFW, VVR, and MSCS \(Windows Server 2003\)”](#) on page 442.

During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:

- Cluster disk group name
- Volume names and sizes
- Drive letters

Before starting the installation make sure you take the SQL IP resource offline on the primary site. This will also offline the dependent resources.

Specify the same name for the SQL virtual server. Make sure the name is the same as that on the primary site.

After completing the tasks listed on page 442, you will have a clustered secondary site with:

- SFW installed
- MSCS option configured
- SQL installed on all the nodes

The next step is to set up replication between the two sites.

VVR components overview

You configure the following Veritas Volume Replicator components:

Replicated Volume Group (RVG)	<p>An RVG is made up of one or more volumes in a SFW disk group. The updates made on the RVG on the primary host are sent to a configured secondary host. Thus, on the secondary host there is a corresponding RVG with a disk group of the same name and volumes with the same names. The data volumes should be the same size. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG.</p> <p>An RVG within a disk group is the container for replication, so if you have multiple disk groups, you will need to create a separate RVG for each disk group. It is possible to have more than one RVG in a disk group; however, the RVG cannot span across disk groups.</p>
Replicated Data Set (RDS)	<p>An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).</p>
Replicator Log volume	<p>Each RVG must have a Replicator Log associated with it. The Replicator Log volume at the primary site holds a copy of any RVG updates that are sent to the secondary site. The Replicator Log on the secondary site is held in reserve so that it can be used if the primary site becomes nonfunctional and the secondary site needs to become the new primary site. The log volumes at the two sites must have the same name. Symantec recommends having Replicator Log volumes of the same size at the primary site and the secondary site.</p>

Creating resources for VVR

VVR resources must be created on the primary and secondary site cluster:

- IP address for replication
- Network name resource

Creating an IP address resource

A separate valid IP address for VVR replication is necessary because on the secondary cluster before a disaster, the SQL virtual server IP address must be offline whereas the VVR IP address must to be online.

To create an IP address resource

- 1 Right click on the Volume Manager disk group and select **New > Resource**.
- 2 In the Resource creation wizard, configure the IP address. Specify a name for the **IP Address** resource.
Add a **Description** if required.
- 3 Select the **IP address** from the **Resource Type** field drop down list. Click **Next**.
- 4 In the Possible Owners page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 5 In the Dependencies Page, make sure the **Resource Dependencies** pane is empty, and click **Next**.
- 6 On the TCP/IP Address Parameters page, set the TCP/IP parameters. Enter the IP address and the corresponding subnet mask.
- 7 Make sure the Network is set to **Public** and click **Finish** to create the IP Address resource.
- 8 Bring the resource online.

Creating a network name resource

To create a network name resource

- 1 Right-click on the SQL_GROUP group and select **New > Resource**.
- 2 In the Resource creation wizard, create a Network Name resource. Specify the **Network Name**.
Add a **Description** if required.

- 3 Specify the resource type by selecting **Network Name** from the **Resource Type** field drop down list. Click **Next**.
- 4 In the Possible Owners page, click **Next**. All the nodes in the cluster are listed as possible owners by default.
- 5 On the Dependencies page, select the IP Address resource you just created for the RVG from the Available Resources pane. Add it to the Resource Dependencies pane and click **Next**.
- 6 In the **Name** field on the Network Name Parameters page, specify any name except the node and SQL Virtual Server names. Click **Finish**.

Note: The network name for the RVG must be different for the primary and secondary cluster.

- 7 Repeat the same procedure to create the IP and the Network Name resource for the secondary site.
- 8 Bring the resources online.

Configuring VVR: Setting up an RDS

For each disk group you created for the application, you set up a Replicated Data Set (RDS) on the primary and secondary hosts. The Setup Replicated Data Set Wizard enables you to configure an RDS for both sites.

Before running the wizard, verify the following:

- Verify that the disk groups and volumes for the SQL user database files and log files have been created. The Replicator Log volume can be created while running the wizard if not created earlier.
- Verify that VxSAS has been configured.
- Verify that the SQL IP virtual server resource is offline on the secondary site. This would also offline all the dependent SQL resources.

VVR does not support these types of volumes:

- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with the Dirty Region Log (DRL)
- Volumes with a comma in their names
- For the Replicator Log volume, in addition to the above types also make sure that the volume does not have a DCM.

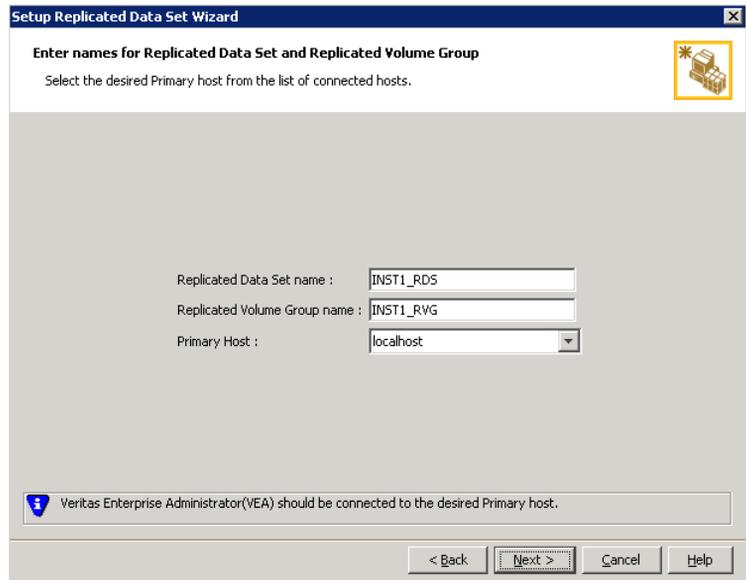
Caution: Do not use volume types that are not supported by VVR.

The following procedure enables you to set up an RDS on the primary and secondary sites and to start replication.

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.
- 3 Read the Welcome page and click **Next**.

- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

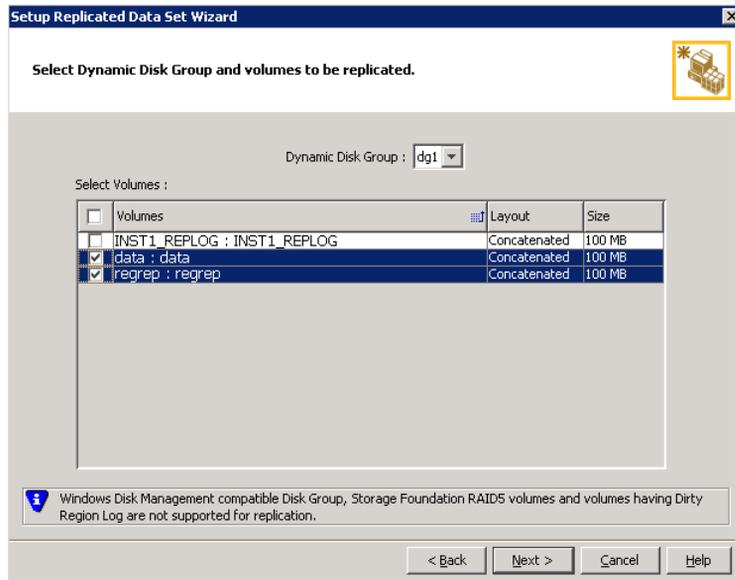


By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.

- 5 Click **Next**.

- 6 Select from the table the dynamic disk group and data volumes that will undergo replication.

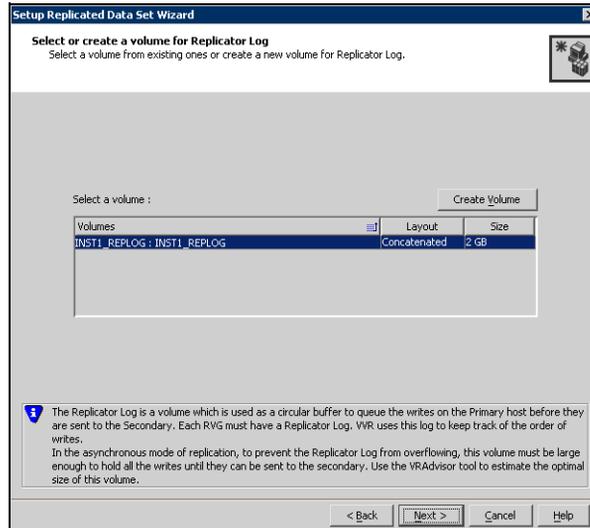


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

- 7 Click **Next**.

8 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name Enter the name for the volume in the **Name** field.

Size Enter a size for the volume in the **Size** field.

Layout Select the desired volume layout.

- Disk Selection**
- Choose **Select disks automatically** if you want VVR to select the disks for the Replicator Log.
 - Choose **Select disks manually** to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select **Add** to move the disks into the selected disks pane.

- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 9 Review the information on the summary page and click **Create Primary RVG**.
 - 10 After the Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?
Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.
Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.
 - 11 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field. Click **Next**. If the Secondary host is not connected to VEA, the wizard tries to connect it when you click **Next**. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.
 - 12 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

 - the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the Primary

Otherwise, the RDS setup wizard enables you to create the required volumes manually.

 - Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
 - 13 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.

This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.

 - If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.

- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.

Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.

When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.

14 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

The screenshot shows the 'Setup Replicated Data Set Wizard' dialog box, specifically the 'Edit replication settings' page. The title bar reads 'Setup Replicated Data Set Wizard'. Below the title bar, the text 'Edit replication settings' is displayed, followed by the instruction 'Edit replication settings or click next.' There is a small icon of a server rack with a star in the top right corner. The main area contains several configuration fields:

- Primary side IP: 10.217.53.214
- Secondary side IP: 10.217.53.215
- Replication Mode: Synchronous Override
- Replicator Log Protection: AutoDCM
- Primary RLINK Name: Pri_RLINK
- Secondary RLINK Name: Sec_RLINK

Below these fields is an 'Advanced' button. At the bottom of the dialog, there is a warning message: 'DHCP addresses are not supported by VVR.' The bottom of the dialog features navigation buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not

wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP	Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Secondary side IP	Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Replication Mode	Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous . The default is synchronous override. Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously. If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as MISSING.
Replicator Log Protection	The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **DCM** option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.

The **Off** option disables Replicator Log Overflow protection.

In the case of the Bunker node. Replicator Log protection is set to **Off**, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

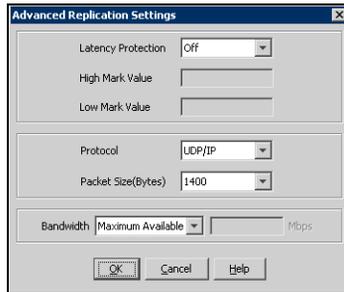
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK Name	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- 15 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

- **Off** is the default option and disables latency protection.
- **Fail** enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.
- **Override** enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Caution: Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Click **OK** to close the dialog box.

16 Click **Next**.

17 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from Checkpoint If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.
- 18 Review the information.
Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

Creating the RVG resource (primary and secondary sites)

To enable a disaster recovery setup, once VVR is configured, you will need to create an RVG resource with dependency on the VVR IP resource and the SQL group IP address resource. You will then need to modify the dependency in the SQL group IP address resource to remove the dependency on the resource and add dependency on the RVG resource.

To create a Replicated Volume Group (RVG) resource

- 1 Right click on the `SQL_GROUP` group that you have created and select **New > Resource**. The New Resource wizard appears.
- 2 Specify a name for the Replicated Volume Group resource in the Name field. If required, you can add a description about the resource in the Description field.
- 3 Specify the resource type by selecting **Replicated Volume Group** from the Resource Type field drop down list. Click **Next**.
- 4 In the Possible Owners page, configure a separate resource monitor process for the RVG resource. Select the **Run this resource in a separate Resource Monitor** checkbox provided in the New Resource wizard.
- 5 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 6 On the Dependencies page, select the VVR IP resource and the Disk Group resource from the **Available Resources** and add it to **Resource Dependencies**. Click **Next**.
- 7 On the Replicated Volume Group Parameters page, select the created RVG. Click **Finish**.
- 8 Repeat the steps to create the RVG resource on the Secondary site.

Setting the SQL server resource dependency on the RVG resource

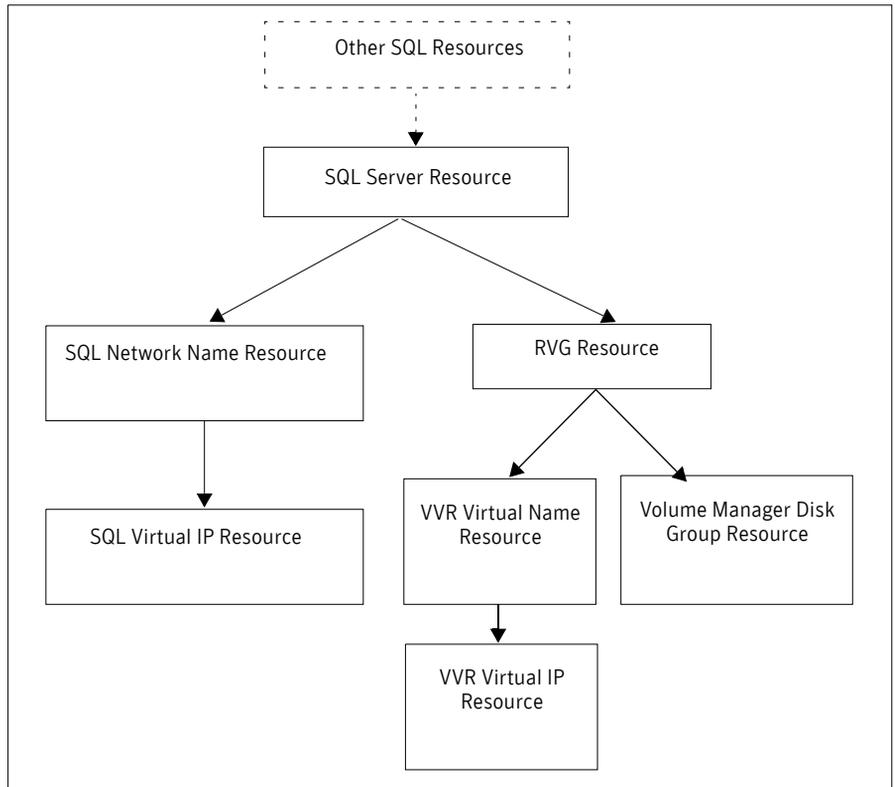
To set the SQL server resource dependency on the RVG resource

- 1 Make sure the SQL Server resource is offline before attempting to modify the dependencies.
- 2 Right-click the **SQL Server** resource and select **Properties > Dependencies** tab. This displays the Dependencies page.
- 3 Click **Modify**.
- 4 Select the **Replicated Volume Group** resource from the Available Resources and add it to **Resource Dependencies**.
- 5 Remove the **Disk Group** resource from **Resource Dependencies**.
- 6 Click **OK**.

The cluster configuration is now complete. Online the entire `SQL_GROUP` group on the primary cluster.

The following dependency graph indicates the dependencies that have been established.

Figure 17-3 Dependency graph



Working with the solution: Normal operations and recovery procedures

This section gives considerations for normal VVR operations and also describes the recovery process.

Normal operations

Monitoring the status of the replication

Under normal operating conditions you can monitor the status of the replication using:

- VEA GUI
- Command Line Interface (CLI)
- Performance Monitor (perfmon)
- Alerts

For details, refer to the “Monitoring Replication” Chapter in the *Veritas Volume Replicator Administrator’s Guide*.

Performing planned migration

For maintenance purposes, or for testing the readiness of the Secondary host you may want to migrate the application to the Secondary host. These are a generic set of tasks that you may need to perform.

To detach the user database

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Expand the Databases node. Right-click on the required user database and select **All Tasks > Detach**.

Note that the `master`, `model`, and `tempdb`, databases cannot be detached.

To take the RVG resource offline

Take the RVG resource offline on both the clusters.

To transfer the Primary role

Transfer the Primary role to the Secondary using the **Migrate** option.

- 1 From the VEA screen, right-click the Primary RVG and select **Migrate**.
- 2 Select the Secondary host and click **OK**. The replication role is migrated to the Secondary host.

To assign drive letters to the volumes

Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as that of the original Primary.

To bring the RVG resource online

- 1 Bring the RVG resource online on both the clusters.
- 2 Bring the `SQL_GROUP` group online on the new Primary.

To attach the databases

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click on the required Database node and select **All Tasks > Attach Database**.
- 4 In the Attach Database dialog, specify the name of the Master Data File (MDF) file that corresponds to the database, which you want to attach. Use the browse (...) button to search. For more information, refer to the Microsoft documentation.

You can now verify that the SQL runs fine on the new Primary with the replicated data. After verifying, you can revert back the roles to its original state using the same set of tasks described above.

Note: Any changes that you make to the data on the new Primary will get replicated to the original Primary, which is now the Secondary.

Replication recovery procedures

This section provides information on bringing up an SQL server on the Secondary host, in the event of a disaster. It also explains how to migrate the Primary role back to the original Primary host once it is in a good state after a disaster.

Bringing up SQL on the secondary host

To recover the SQL data

- 1 From the left-pane in the VEA GUI console on the Secondary host, right-click on the desired secondary RVG node inside the replication network.
- 2 Select **Takeover** and follow the instructions in the wizard to perform the takeover operation. You can choose to perform takeover with the following options:
 - Perform the **Takeover with fast-failback** option to restore the original Primary easily once it becomes available again. When performing Takeover with fast-failback, make sure that you do not select the **Synchronize Automatically** option.
 - Perform the **Takeover without fast-failback** option. In this case, you need to perform a complete synchronization of the original Primary with the new Primary. This may take quite a while depending on the size of the data volume. Only after the synchronization is complete can you migrate the Primary role back to the original Primary.
After takeover, the existing Secondary becomes the new Primary.
- 3 Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as that of the original Primary.
- 4 Bring the SQL_GROUP group online.

To attach the databases

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click on the required Database node and select **All Tasks > Attach Database**.
- 4 In the Attach Database dialog, specify the name of the Master Data File (MDF) file that corresponds to the database that you want to attach. Use the

browse (...) button to search. For more information, refer to the Microsoft documentation.

Now you can start using SQL on the new Primary.

Restoring the primary host

After a disaster, if the original Primary becomes available again you may want to revert the role of the Primary back to this host.

To detach the user database on the new Primary

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Expand the Databases node. Right-click on the required user database and select **All Tasks > Detach**.

Note that the `master`, `model`, and `tempdb`, databases cannot be detached.

To migrate the Primary role back to the original Primary

- 1 Take the RVG resource offline on both the clusters.
- 2 Depending on whether you performed **Takeover** with or without fast-failback option, do one of the following:
 - For Takeover with the Fast-failback option, the original Primary, after it has recovered, will be in the `Acting as Secondary` state. If the original Primary is not in the `Acting as Secondary` state, verify whether your network connection has been restored.
To synchronize this original Primary and the new Primary, use the **Resynchronize Secondaries** option from the right-click menu of the new Primary.
 - For Takeover without the Fast-failback option, after you have performed this operation, you must convert the original Primary to a Secondary using the **Make Secondary** option.

Note: Before performing the **Make Secondary** operation, the original Primary's RVG and the new Primary's RVG will be shown in separate RDSs. However, after this operation they will be merged under a single RDS.

After the **Make Secondary** operation, the original Primary will be converted to a secondary. Right-click this secondary RVG and select **Start Replication** with **Synchronize Automatically** option.

- 3 After the synchronization is complete, perform a migrate operation to transfer the Primary role back to the original Primary. To do this, right-click the Primary RVG and select **Migrate** from the menu.
- 4 Ensure that the volumes have retained the same drive letters that existed before the disaster.
- 5 Bring the RVG resource online on the Secondary.
- 6 Bring the `SQL_GROUP` group online on the original Primary.

To attach the databases on the original Primary

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click the required Database node and select **All Tasks > Attach Database**.

In the Attach Database dialog, specify the name of the Master Data File (MDF) file that corresponds to the database that you want to attach. Use the browse (...) button to search. For more information, refer to the Microsoft documentation.

Deploying SFW and VVR with Microsoft failover clustering: New SQL 2005 installation

This chapter covers the following topics:

- [Tasks for a new SQL 2005 installation with SFW, VVR, and Microsoft failover clustering \(Windows Server 2008\)](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing an Microsoft failover cluster](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Completing the primary site configuration](#)
- [Creating a parallel environment on the secondary site](#)
- [VVR components overview](#)
- [Creating resources for VVR](#)
- [Setting up the replicated data sets \(RDS\) for VVR](#)
- [Creating the RVG resource \(primary and secondary sites\)](#)
- [Setting the SQL Server resource dependency on the RVG resource](#)

- Working with the solution: Normal operations and recovery procedures

Tasks for a new SQL 2005 installation with SFW, VVR, and Microsoft failover clustering (Windows Server 2008)

You can install and configure Storage Foundation for Windows (SFW) and Veritas Volume Replicator (VVR) with Microsoft failover clustering and SQL Server 2005. This chapter describes the deployment on Windows Server 2008. This environment involves an active/passive configuration with one-to-one failover capability. After setting up a SFW environment with high availability for SQL Server 2005 on a primary site, you can create a secondary or “failover” site for replication. Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details on VVR.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 18-1 Tasks for deploying SFW and VVR with failover clustering for SQL 2005

Objective	Tasks
“ Reviewing the prerequisites ” on page 514	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 515	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive SQL configuration in VVR cluster
Part 1	Configuring the primary site
“ Configuring the storage hardware and network ” on page 518	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“ Establishing a Microsoft failover cluster ” on page 250	<ul style="list-style-type: none"> ■ Reviewing general guidelines to establish a Microsoft failover cluster on Windows Server 2008
“ Installing SFW with MSCS/Failover Cluster option ” on page 522	<ul style="list-style-type: none"> ■ Installing SFW (automatic installation) ■ Configuring VxSAS
“ Configuring SFW disk groups and volumes ” on page 533	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes

Table 18-1 Tasks for deploying SFW and VVR with failover clustering for SQL 2005 (Continued)

Objective	Tasks
“Creating the SQL Server virtual server group” on page 271	<ul style="list-style-type: none"> ■ Creating a SQL Server cluster group ■ Creating the disk group resource
“Installing SQL Server 2005” on page 272	<ul style="list-style-type: none"> ■ Installing SQL ■ Verifying SQL installation
“Implementing a dynamic mirrored quorum resource” on page 276	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the Cluster Group ■ Changing the quorum resource to a dynamic mirrored quorum resource.
“Verifying the cluster configuration” on page 278	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node
Part 2	Configuring the secondary site
“Creating a parallel environment on the secondary site” on page 545	<ul style="list-style-type: none"> ■ Reviewing the special requirements for creating the parallel environment on the secondary site ■ Following the procedures in the tasks for configuring the primary site
Part 3	Adding the VVR components and resources
“Creating resources for VVR” on page 547	<ul style="list-style-type: none"> ■ Creating an IP address for the Replicated Volume Group (RVG). ■ Creating a network name resource for the Replicated Volume Group (RVG)
“Setting up the replicated data sets (RDS) for VVR” on page 548	<ul style="list-style-type: none"> ■ Setting up an RDS for each application disk group
“Creating the RVG resource (primary and secondary sites)” on page 559	<ul style="list-style-type: none"> ■ Creating an RVG resource
“Setting the SQL Server resource dependency on the RVG resource” on page 560	<ul style="list-style-type: none"> ■ Setting up SQL Server resource dependencies

Reviewing the prerequisites

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

Supported software for Microsoft failover clusters with SFW

- Veritas Storage Foundation 5.1 for Windows (SFW)
Include the following option along with any others applicable to your environment:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
- For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SQL Server 2005 SP2 required) ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition

Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SQL Server 2005 SP2 required) ■ Windows Server 2008 for 64-bit Itanium (IA64) ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 18-2 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

- One CD-ROM drive accessible to the system on which you are installing SFW.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- Microsoft clustering requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft SQL Server documentation for instructions on creating a reverse lookup zone.
- Microsoft clustering requires two disks for SQL: one for SQL database files and one for SQL log files.
- Each system requires a minimum of 1 GB of RAM.
- SFW requires administrator privileges to install the software.
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node. Using static IP addresses for the public network and private network cards is highly recommended and is required for a VVR configuration.

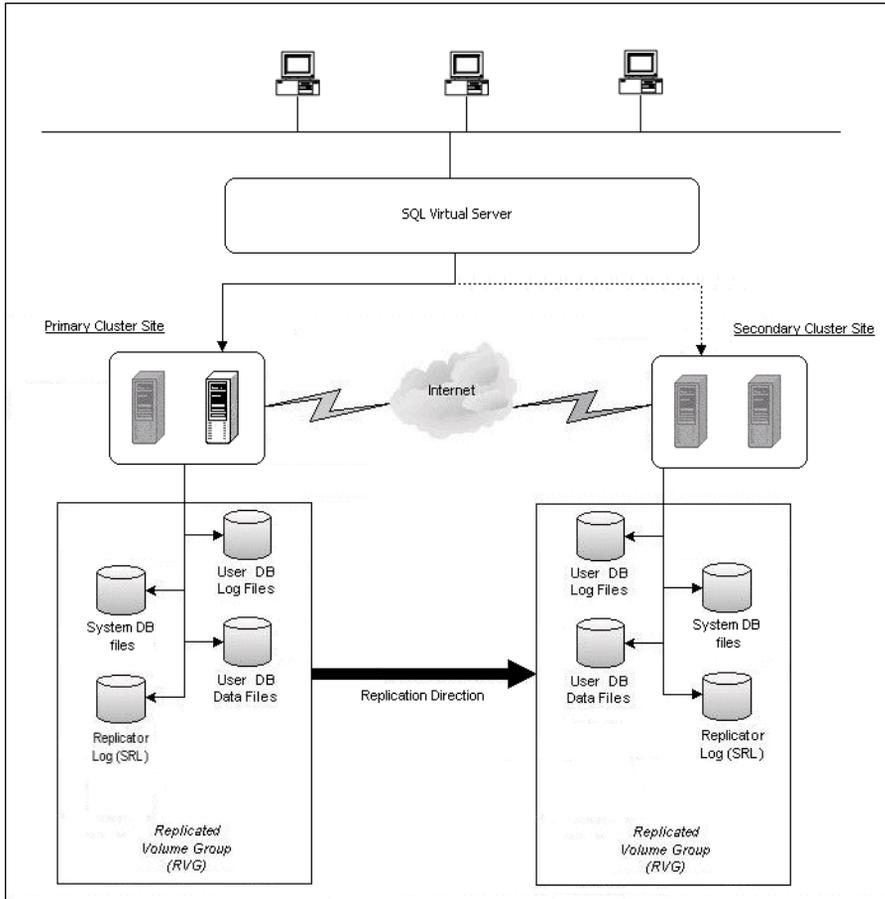
Note: Refer to the Hardware Compliance List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Reviewing the configuration

The following figure illustrates a typical clustered VVR configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the Replicated Volume Group (RVG). The Microsoft SQL Server 2005 application data is stored on the volumes that are under the control of the RVG. A separate disk group is created for the quorum volume, which is not replicated.

Figure 18-1 shows a typical VVR configuration.

Figure 18-1 Typical VVR configuration



If the Microsoft SQL Server 2005 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the replication solution is activated. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over. The data that was replicated to the secondary site is used to restore the SQL services to clients.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site. The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary Site

SYSTEM1 & SYSTEM2	server names
SQL_GROUP	Microsoft SQL Server virtual server group
SQLCLUST	Microsoft SQL Server virtual cluster (underscores not supported)
SQLVS	Microsoft SQL Server virtual server
SQL_IP	Microsoft SQL virtual server IP address resource
INST1	Microsoft SQL Server instance name
INST1_DG	disk group for Microsoft SQL volumes
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
SQLVS_QRM	volume for storing the Microsoft failover cluster quorum
QUORUM_DG	quorum volume disk group for mirroring the quorum

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
----------------------	--

All the other parameters are the same as on the primary site.

DR Components

INST1_RDS	VVR Replicated Data Set (RDS) name
INST1_RVG	VVR Replicated Volume Group (RVG) name
INST1_REPLOG	VVR Replicator log volume

INST1_RVG_RES	Replicated Volume Group Resource name
VVR_IP	Microsoft SQL RVG IP address resource

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent the Microsoft cluster from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.
- 13 Close the window.

Establishing an Microsoft failover cluster

Before installing SFW, you must first verify that Microsoft failover clustering is enabled (if a new installation of Windows Server 2008), and then establish a Microsoft failover cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To enable Microsoft failover clustering

- 1 In Server Manager, select **Features** in the left pane (tree view) and then click **Add Features** (link on the right side of the screen).
- 2 In the Add Features Wizard, check the **Failover Clustering** option, and click **Next**.
- 3 Click **Install**.
- 4 When the installation is complete, click **Close**.

To establish a Microsoft failover cluster

- 1 Ensure that you have met the hardware prerequisites for a failover cluster. You can run the Microsoft wizard to validate the configuration. See the Microsoft documentation for details.
- 2 Configure the shared storage and create a volume with drive letter “Q” for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends a minimum of 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster using Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 4 In the action pane, click **Create a Cluster**. The Create Cluster Wizard will start.
If this is the first time this wizard has been run, the Before You Begin page will appear. Review the information that is displayed and then click **Next**. You can hide this page on subsequent uses of the wizard; if this has been done, the first page might be the Select Servers page.
- 5 In the Select Servers panel, type the name of the first node in the Enter server name field and click **Add**. You can also use the Browse button to browse the Active Directory for the computers you want to add. Repeat this step for the second node.
- 6 After both nodes have been added to the list of Selected Servers, click **Next**.

- 7 Based on the information on the validation warning screen, assess your hardware configuration, and select one of the options. Symantec recommends that you select **Yes**, which starts the Validate a Configuration wizard. Follow the wizard instructions.
- 8 In the Access Point for Administering the Cluster screen, in the Cluster Name field, type the name for the failover cluster. This is the name that you use to connect to and administer the cluster.
- 9 In the Address field of the network area, type the appropriate IP address and then click **Next**.
- 10 In the Confirmation screen, verify that the cluster configuration is correct, and then click **Next** to create the cluster.
- 11 Review the Summary page and then click **Finish** to close the wizard.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use these disks when you create your SFW cluster disk groups, you must first remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs. After creating the SFW cluster disk groups, you will add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

Installing SFW with MSCS/Failover Cluster option

This section assumes that you are running a Microsoft failover cluster and that you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft failover cluster simultaneously.

SFW installation tasks

The product installer enables you to install the software for Veritas Storage Foundation 5.1 for Windows. The installer automatically installs SFW. You must select the options to install VVR, and the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation. The steps in this section are based on a server installation.

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 522.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 523.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 528.

Pre-installation tasks

Perform the following pre-installation task:

- Moving the online groups

Moving the online groups

If your resource groups are online on the system where you are installing SFW. You must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).

- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install Veritas Storage Foundation for Windows on a MSCS configuration.

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

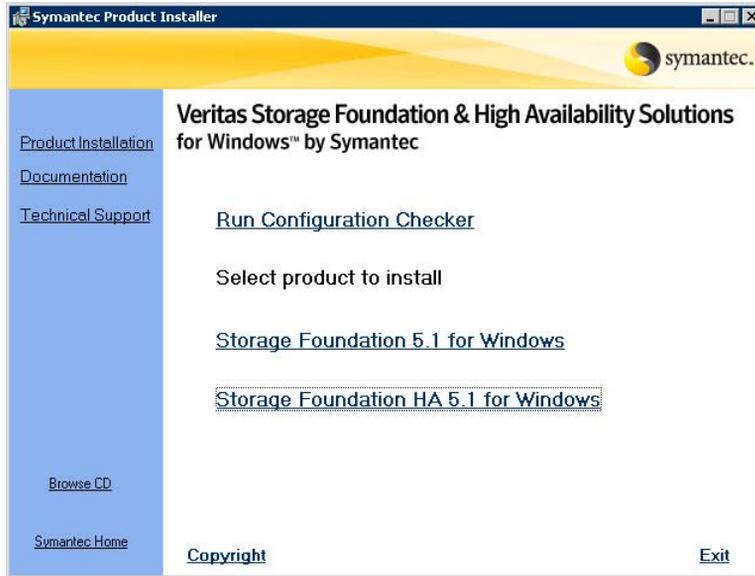
Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

To install the product

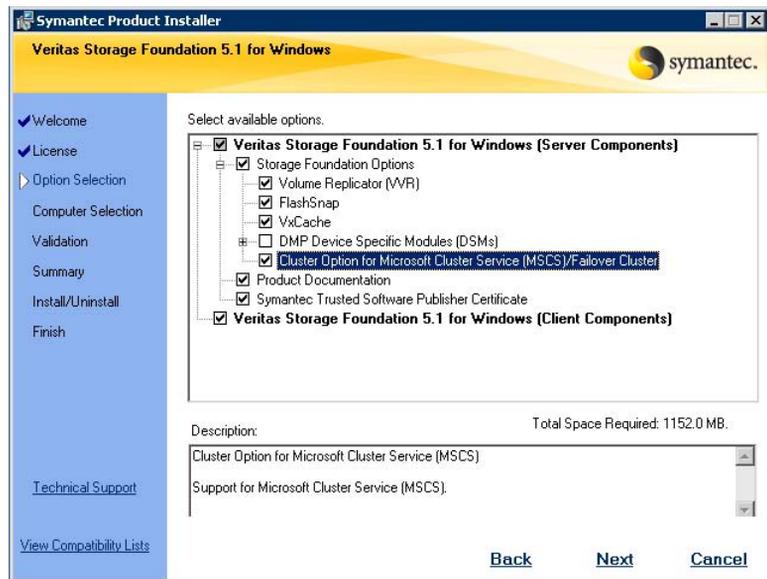
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

3 Click **Storage Foundation 5.1 for Windows**.



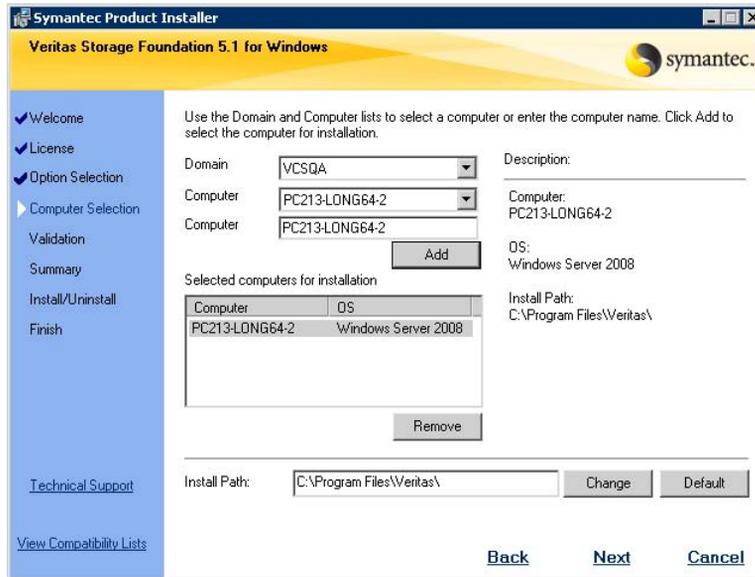
- 4 Click **Complete/Custom** to begin installation. The **Administrative Console** link allows you to install only the Client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for "**I accept the terms of the license agreement**," and click **Next**.
- 7 Enter the product license key before adding license keys for features.
- 8 Enter the license key in the top field and click **Add**.
- 9 Repeat for additional license keys.
To remove a license key, click the key to select it and click **Remove**.
To see the license key's details, click the key.
- 10 Click **Next**.

11 Specify the product options:



- 12 Select the **Volume Replicator (VVR)** option.
- 13 Select the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** and any additional options applicable to your environment.
- 14 Do not select the Dynamic Multi-pathing option.
Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.
Click **Next**.
- 15 To install the client components on all the computers in your installation, verify that the **Veritas Storage Foundation 5.1 for Windows (Client Components)** check box is selected, to install the client component.
Click **Next**.

16 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path	<p>Optionally, change the installation path.</p> <ul style="list-style-type: none"> ■ To change the path, select a computer in the Selected computers for installation field, type the new path, and click Change. ■ To restore the default path, select a computer and click Default. <p>The default path is: C:\Program Files\Veritas For 64-bit installations, the default path is: C:\Program Files (x86)\Veritas</p>
--------------	--

- 17 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 18 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the computers list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 19 Read the information in the warning box that appears after validation and click **OK**.

Quorum Arbitration

The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume. The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that Microsoft clustering allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the MSCS Support chapter of the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

- 20 Review the information and click **Install**. Click **Back** to make changes.
- 21 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, the status screen shows a failed installation. Click **Next** to review the report, address the reason for failure, and retry this step on that computer.

If the installation is successful on all systems, the installation report screen appears.

If your local computer has its driver signing options set to Block or Warn then installation fails.

- 22 Review or print the report and review log files. Click **Next**.
 - Proceed to [step 23](#) if you are installing SFW on the local node only.
 - Proceed to [step 25](#) if you are installing SFW on local and remote systems.
- 23 To complete the installation, click **Finish**.
- 24 Click **Yes** to reboot the system and complete the installation.
- 25 Reboot the remote nodes. You can neither select computers where the installation has failed nor the local computer.
 - Click the check box next to the remote nodes that you want to reboot and click **Reboot**.
 - When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
 - Click **Finish**.
 - Click **Yes** to reboot the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
- Completing the SFW Installation for the other systems in the failover cluster
- Configuring the VxSAS service

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open the Failover Cluster Management tool. (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click the resource group and then click **Move this service or application to another node > Move to node [name of original node]**.

If there is more than one resource group, you must repeat this step until all the resource groups are moved back to the original node.

- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that all the resource groups have moved back to the original system.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 522.

Configuring the VxSAS service

Complete the following procedure to configure this required service for VVR. This procedure should not be done until you have installed SFW on all cluster systems. Otherwise, you will get an error message from the VxSAS wizard if you try to select a system without SFW installed.

You can run the VxSAS wizard from any site once SFW is installed on all cluster systems; at that time, you can run the wizard for both the primary and secondary site systems. The MSCS groups can be either online or offline.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator’s Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.

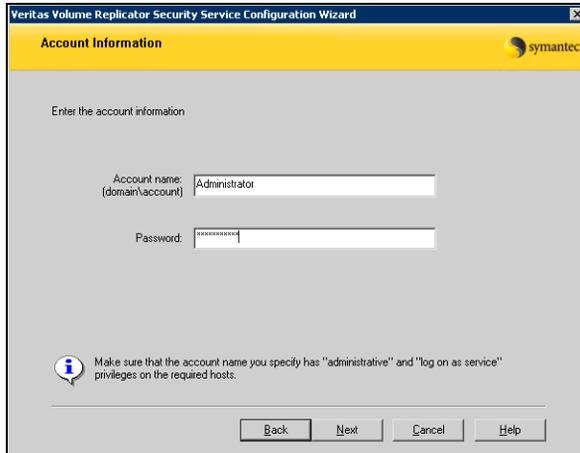
The welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account) Enter the administrative account name in the Account name field.

Password Specify a password in the **Password** field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same user name and password when configuring the VxSAS service on the other hosts.



The screenshot shows a window titled "Veritas Volume Replicator Security Service Configuration Wizard" with a yellow header bar containing the Symantec logo. The main area is titled "Account Information" and contains the instruction "Enter the account information". There are two input fields: "Account name: (domain\account)" with the text "Administrator" entered, and "Password:" with a masked password "*****". At the bottom, there is an information icon and a note: "Make sure that the account name you specify has 'administrative' and 'log on as service' privileges on the required hosts." Below the note are four buttons: "Back", "Next", "Cancel", and "Help".

- 3 After providing the required information, click **Next**.

- 4 Select the required domain to which the hosts that you want to configure belong, from the Domain Selection wizard page.

Selecting Domains The Available Domains pane lists all the domains that are present in the Windows network neighborhood.
Select the required domain by moving the appropriate name from the Available Domains pane to the Selected Domains pane, either by double-clicking it or using the arrow button.

Adding a Domain If the domain name that you require is not displayed, then add it by using the **Add Domain** option. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected Domains list.

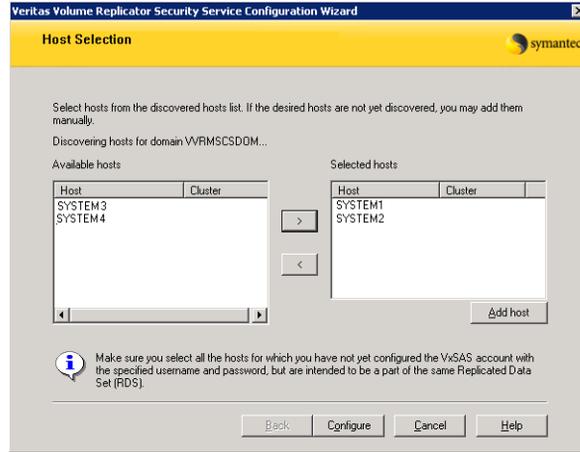
After specifying the domain, click **Next**.

- 5 Select the required hosts from the Host Selection page.

Selecting Hosts The Available Hosts pane lists the hosts that are present in the specified domain.
Select the required host by moving the appropriate name from the Available Hosts list to the Selected Hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a Host If the host name you require is not displayed, then add it using **Add Host** option. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected Hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.



- 6 After the configuration completes, the Configuration Results page is displayed. If the operation is successful, then the Status column displays the appropriate message to indicate that the operation was successful.

If the operation was not successful, then the page displays the status as failed and the corresponding details on why the account update failed. It also displays the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

- 7 Click **Finish** to exit the wizard.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for SQL Server. A dynamic disk group is a collection of one or more disks that behaves as a single storage repository. Within each disk group, you can have dynamic volumes with different layouts.

Configuring disk groups and volumes involves the following tasks:

- [“Planning disk groups and volumes”](#) on page 533
- [“Creating dynamic cluster disk groups”](#) on page 535
- [“Creating dynamic volumes”](#) on page 537
- [“Managing disk groups and volumes”](#) on page 542

Planning disk groups and volumes

Before installing SQL, you must create disk groups and volumes using the VEA console installed with SFW.

Before creating a disk group, consider:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.
- The number of disk groups and volumes that are needed for SQL Server
The number of disk groups for SQL Server depends on the planned organization of the data. The application program files need to be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage in a cluster disk group. You create at least one disk group for the system data files. You may want to create additional disk groups for user databases. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- The disk group and volume for the mirrored quorum resource
You will need a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends a minimum of 500 MB for the quorum disk. You can create the quorum disk group at the same time you create application disk groups, although it is not required for installing the application.

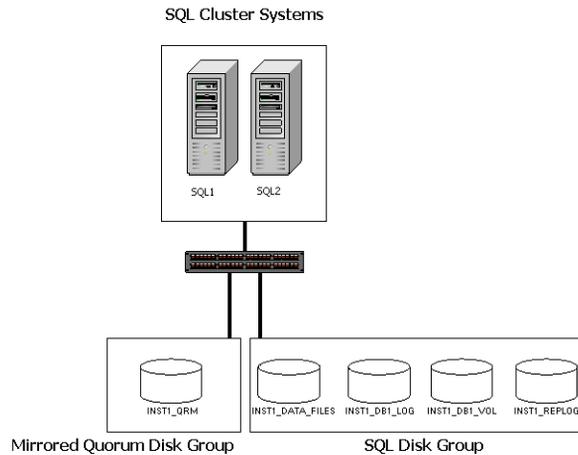
Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

VVR does not support these types of volumes:

- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with the Dirty Region Log (DRL)
- Volumes with a comma in their names
- For the Replicator Log volume, in addition to the above types also make sure that the volume does not have a DCM.

Typically, a SFW disk group corresponds to a SQL virtual server group. Below is a detailed view of the disk groups and volumes for SQL:

Figure 18-2 SFW disk groups and volumes for SQL virtual server INST1-VS in failover clustering setup



SQL disk group INST1_DG contains four volumes:

- INST1_DATA_FILES: Contains the SQL database. Each database in an SQL storage group typically resides on a separate volume.
- INST1_DB1_LOG: Contains the transaction log.
- INST1_DB1_VOL: Contains volume for Microsoft SQL Server system data files.
- INST1_REPLOG: Contains the replicator log for VVR.

Use the following procedures to create the appropriate disk groups and volumes. This section assumes you are using one database.

Creating dynamic cluster disk groups

When the tasks described in this section are completed, you will have a dynamic cluster disk group with volumes on shared storage. The dynamic cluster disk groups will be ready to be shared between nodes in the cluster.

A dynamic cluster disk group is a collection of one or more disks that behave as a single storage repository and which can potentially be accessed by different computers. Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it will be required later during the SQL installation process.

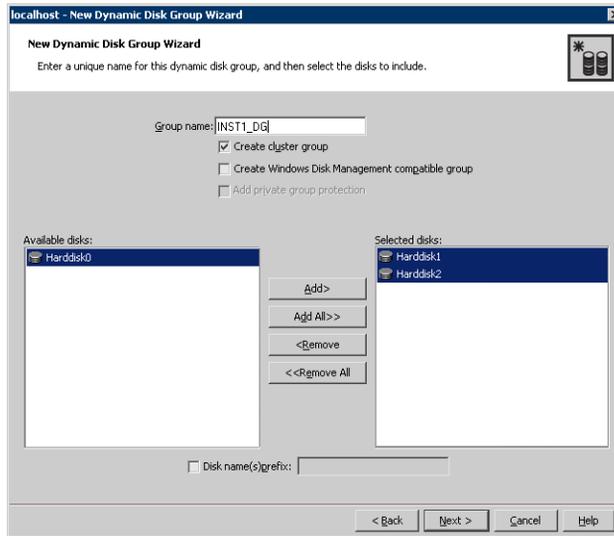
To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on client system and can be used to manage all the servers remotely.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

Creating dynamic volumes

This section will guide you through the process of creating a volume on a dynamic disk group. When creating a disk group to support a SQL Server 2005 solution, it is best to separate SQL data files from SQL log files and place them on separate volumes. Repeat the procedure below to create the following volumes on the first node of the cluster:

- **INST1_DATA_FILES**: For storing the SQL system databases.
- **INST1_DB1_VOL**: For storing the user database.
- **INST1_DB1_LOG**: For storing the user database log.
- **INST1_REPLOG**: For storing the replicator log required by VVR.

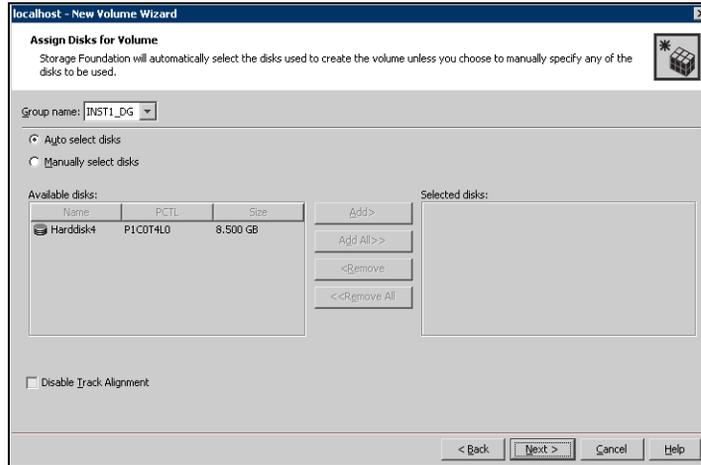
Caution: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

You can create the **INST1_REPLOG** volume at this time or during the process of configuring VVR. See “[Setting up the replicated data sets \(RDS\) for VVR](#)” on page 548.

To create dynamic volumes

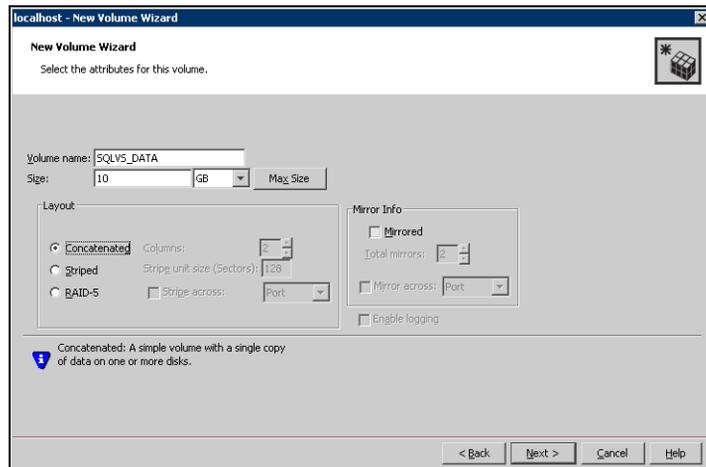
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example **INST1_DG**.
- 5 At the New Volume wizard opening screen, click **Next**.

6 Select the disks for the volume.



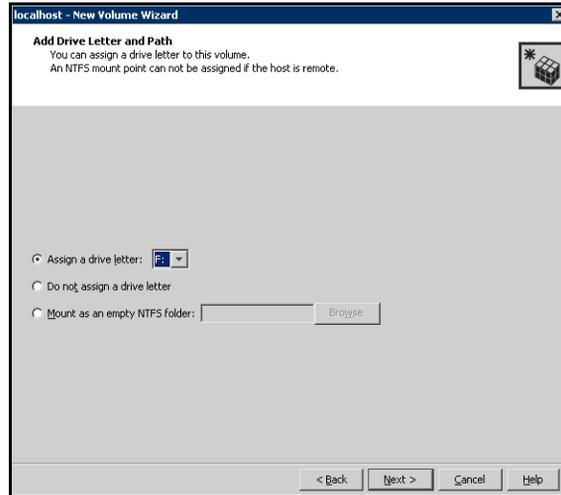
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the parameters of the volume.



- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the **Mirror Info** area, select the appropriate mirroring options.
 - Verify that **Enable Logging** is not selected.
 - Click **Next**.
- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on

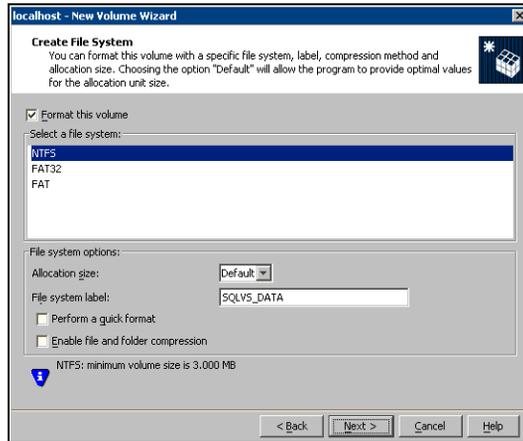
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
- For the Replicator Log volume only: Clear the Format this volume check box.
- Click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

11 Click **Finish** to create the new volume.

12 Repeat these steps to create additional volumes.

Managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter
Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Completing the primary site configuration

The remainder of the tasks for the primary site configuration are identical to the tasks described in the chapter for configuring SQL Server 2005 for high availability:

[Chapter 18, “Deploying SFW and VVR with Microsoft failover clustering: New SQL 2005 installation”](#).

See the following topics in that chapter to complete configuring the primary site:

- [“Creating the SQL Server virtual server group”](#) on page 271
- [“Installing SQL Server 2005”](#) on page 272
- [“Implementing a dynamic mirrored quorum resource”](#) on page 276
- [“Verifying the cluster configuration”](#) on page 278

When done, proceed with the guidelines in this chapter for configuring the secondary site:

- [“Creating a parallel environment on the secondary site”](#)

Creating a parallel environment on the secondary site

After configuring Microsoft failover clustering, Storage Foundation for Windows, and SQL Server 2005 on the primary site, complete the same tasks on the secondary site. Follow the task list table.

See [“Tasks for a new SQL 2005 installation with SFW, VVR, and Microsoft failover clustering \(Windows Server 2008\)”](#) on page 512.

In addition, note the following special requirements for configuring the secondary site:

- During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - Cluster disk group name
 - Volume names and sizes
 - Drive letters
- Before starting the installation make sure you take the SQL IP resource offline on the primary site. This will also offline the dependent resources.
- When installing SQL Server on the secondary site, specify the same name for the SQL virtual server as was used on the primary site.

After completing the tasks listed on page 512, you will have a clustered secondary site with:

- SFW installed
- SQL Server group virtual server group configured
- SQL Server installed on all the nodes

You can now continue with the tasks to set up replication between the two sites.

See the following topics:

- [“VVR components overview”](#) on page 546
- [“Creating resources for VVR”](#) on page 547
- [“Setting up the replicated data sets \(RDS\) for VVR”](#) on page 548
- [“Creating the RVG resource \(primary and secondary sites\)”](#) on page 559
- [“Setting the SQL Server resource dependency on the RVG resource”](#) on page 560

VVR components overview

You configure the following Veritas Volume Replicator components:

Replicated Volume Group (RVG)	<p>An RVG is made up of one or more volumes in a SFW disk group. The updates made on the RVG on the primary host are sent to a configured secondary host. Thus, on the secondary host there is a corresponding RVG with a disk group of the same name and volumes with the same names. The data volumes should be the same size. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG.</p> <p>An RVG within a disk group is the container for replication, so if you have multiple disk groups, you will need to create a separate RVG for each disk group. It is possible to have more than one RVG in a disk group; however, the RVG cannot span across disk groups.</p>
Replicated Data Set (RDS)	<p>An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).</p>
Replicator Log volume	<p>Each RVG must have a Replicator Log associated with it. The Replicator Log volume at the primary site holds a copy of any RVG updates that are sent to the secondary site. The Replicator Log on the secondary site is held in reserve so that it can be used if the primary site becomes nonfunctional and the secondary site needs to become the new primary site. The log volumes at the two sites must have the same name. Symantec recommends having Replicator Log volumes of the same size at the primary site and the secondary site.</p>

Creating resources for VVR

Create the resources for VVR replication at the primary and secondary sites using the Failover Cluster Management tool. You create a network name resource and IP address resource to be used for VVR replication.

A separate valid IP address is necessary for VVR replication, because on the secondary cluster before a disaster, the application IP must be offline whereas the VVR IP must be online.

You create the resources for the primary site and then repeat the procedure to create the resources on the secondary site.

To create a Network Name resource and IP address resource for VVR replication

- 1 Right-click on the application group and select **Add a Resource > Client Access Point**.
- 2 In the Client Access Point panel of the New Resource Wizard, specify the following:
 - In the **Name** field, specify a name for the Network Name resource. The default is the name of the group you selected. Specify any name except the node and the virtual server name. The network name you assign when creating the resource for the secondary site must be different from the network name for the primary site.
 - Select the network and specify the IP address.Click **Next**.
- 3 In the Confirmation panel, review the information and click **Next**.
- 4 When configuration is complete, click **Finish**.
- 5 Repeat the same procedure to create the IP and the Network Name resource at the secondary site.
- 6 Bring the resources online.

Setting up the replicated data sets (RDS) for VVR

For each disk group you created for the application, you set up a Replicated Data Set (RDS) on the primary and secondary hosts. The Setup Replicated Data Set Wizard enables you to configure an RDS for both sites.

Before you begin, verify the prerequisites:

- Verify that the disk groups and volumes for the SQL user database files and log files have been created. The Replicator Log volume can be created while running the wizard if not created earlier.
- Verify that VxSAS has been configured.
- Verify that the SQL IP virtual server resource is offline on the secondary site. This would also offline all the dependent SQL resources.

VVR does not support these types of volumes:

- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with the Dirty Region Log (DRL)
- Volumes with a comma in their names
- For the Replicator Log volume, in addition to the above types also make sure that the volume does not have a DCM.

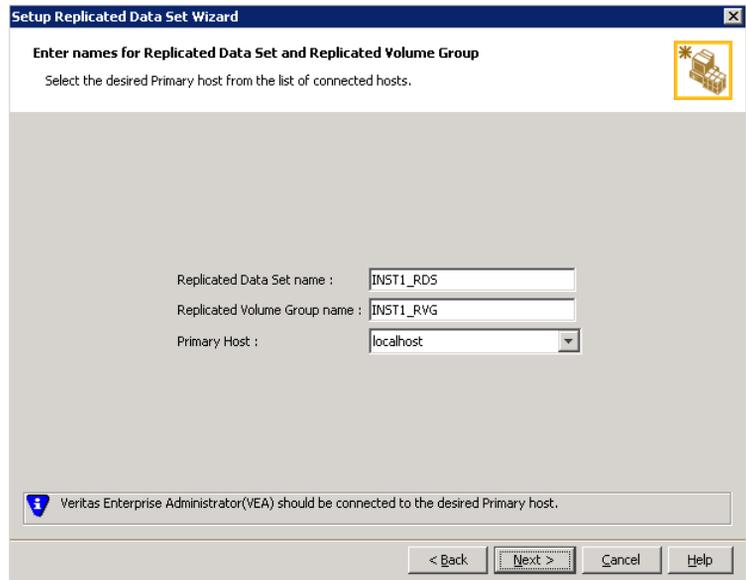
Caution: Do not use volume types that are not supported by VVR.

The following procedure enables you to set up an RDS on the primary and secondary sites and start replication.

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.
- 3 Read the Welcome page and click **Next**.

- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

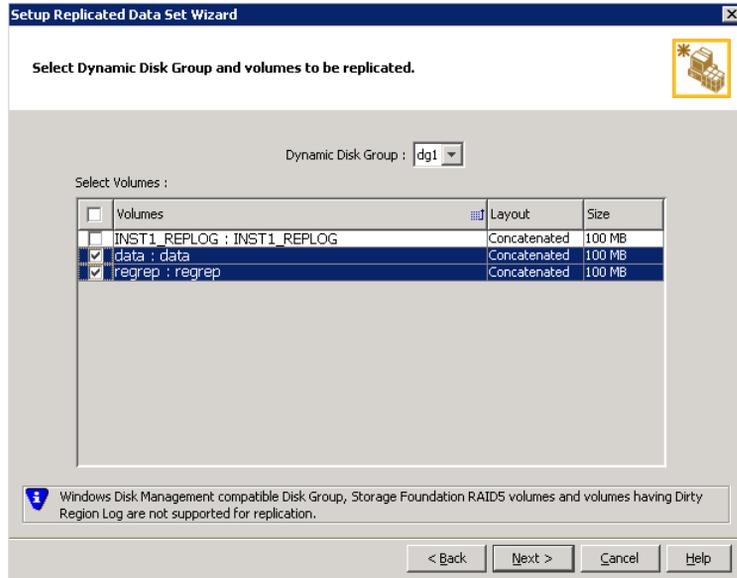


By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.

- 5 Click **Next**.

- 6 Select from the table the dynamic disk group and data volumes that will undergo replication.

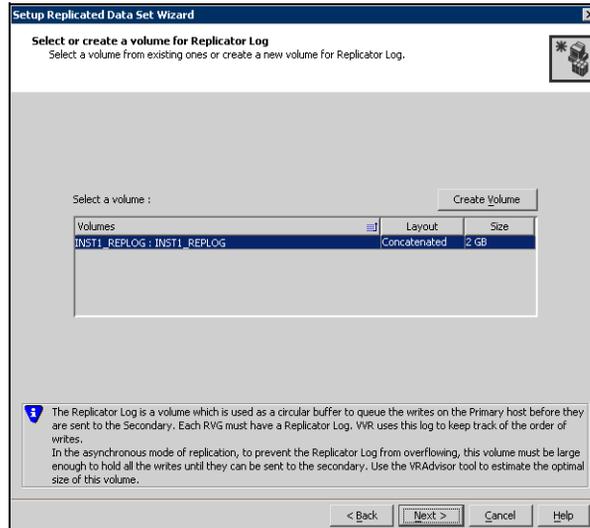


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

- 7 Click **Next**.

8 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name Enter the name for the volume in the **Name** field.

Size Enter a size for the volume in the **Size** field.

Layout Select the desired volume layout.

- Disk Selection**
- Choose **Select disks automatically** if you want VVR to select the disks for the Replicator Log.
 - Choose **Select disks manually** to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select **Add** to move the disks into the selected disks pane.

- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 9 Review the information on the summary page and click **Create Primary RVG**.
 - 10 After the Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?
Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.
Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.
 - 11 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field. Click **Next**. If the Secondary host is not connected to VEA, the wizard tries to connect it when you click **Next**. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.
 - 12 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

 - the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the Primary

Otherwise, the RDS setup wizard enables you to create the required volumes manually.

 - Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
 - 13 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.

This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.

 - If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.

- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.

Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.

When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.

14 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

The screenshot shows the 'Setup Replicated Data Set Wizard' dialog box, specifically the 'Edit replication settings' page. The title bar reads 'Setup Replicated Data Set Wizard'. Below the title bar, the text 'Edit replication settings' is displayed, followed by the instruction 'Edit replication settings or click next.' There is a small icon of a server rack with a star in the top right corner. The main area contains several configuration fields:

- Primary side IP: 10.217.53.214
- Secondary side IP: 10.217.53.215
- Replication Mode: Synchronous Override
- Replicator Log Protection: AutoDCM
- Primary RLINK Name: Pri_RLINK
- Secondary RLINK Name: Sec_RLINK

Below these fields is an 'Advanced' button. At the bottom of the dialog, there is a warning message: 'DHCP addresses are not supported by VVR.' and navigation buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not

wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP	Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Secondary side IP	Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Replication Mode	Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous . The default is synchronous override. Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously. If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as MISSING.
Replicator Log Protection	The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **DCM** option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.

The **Off** option disables Replicator Log Overflow protection.

In the case of the Bunker node. Replicator Log protection is set to **Off**, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

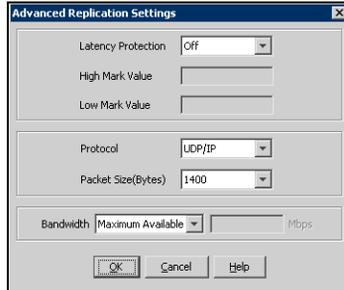
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK Name	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- 15 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



- Latency protection** Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.
- **Off** is the default option and disables latency protection.
 - **Fail** enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.
 - **Override** enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Caution: Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

- High Mark Value** Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Click **OK** to close the dialog box.

16 Click **Next**.

17 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from
Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.
- 18** Review the information.
Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

Creating the RVG resource (primary and secondary sites)

To enable a disaster recovery setup, once VVR is configured, you will need to create an RVG resource with dependency on the VVR IP resource and the SQL group IP address resource. You will then need to modify the dependency in the SQL group IP address resource to remove the dependency on the resource and add dependency on the RVG resource.

To create a Replicated Volume Group (RVG) resource

- 1 In Failover Cluster Management, expand Services and Applications, right-click the SQL Server virtual server group that you have created and select **Add a resource > More resources > Add Replicated Volume Group**. The New Replicated Volume Group appears in the center panel under Disk Drives.
- 2 Right-click **New Replicated Volume Group** and click **Properties**.
- 3 On the General tab of the Properties dialog box, in the Resource Name field, type a name for the RVG resource.
- 4 On the Dependencies tab, add the dependencies for the RVG resource:
 - Click the box **Click here to add a dependency**
 - From the Resource drop-down list, select the network name you created for the RVG. Click **Insert**.
 - Click the box **Click here to add a dependency**
 - From the Resource drop-down list, select the Volume Manager Disk Group resource created for the application disk group. Click **Insert**.
- 5 On the Properties tab, specify the following:
 - In the rvgName field, type the same name that you assigned the RVG on the General tab.
 - In the dgName field, type the name assigned in the VEA to the application disk group.
- 6 Click **OK** to close the Properties dialog box.
- 7 Right-click the RVG resource and click **Bring this resource online**.
- 8 Repeat the same steps to create the RVG resource at the secondary site.

Setting the SQL Server resource dependency on the RVG resource

The SQL Server resource was earlier set to depend on a Volume Manager Disk Group resource that corresponded to the disk group created for the application. After you add the RVG resource for that disk group, you must change the dependency. You set the SQL Server resource to depend on the RVG resource instead.

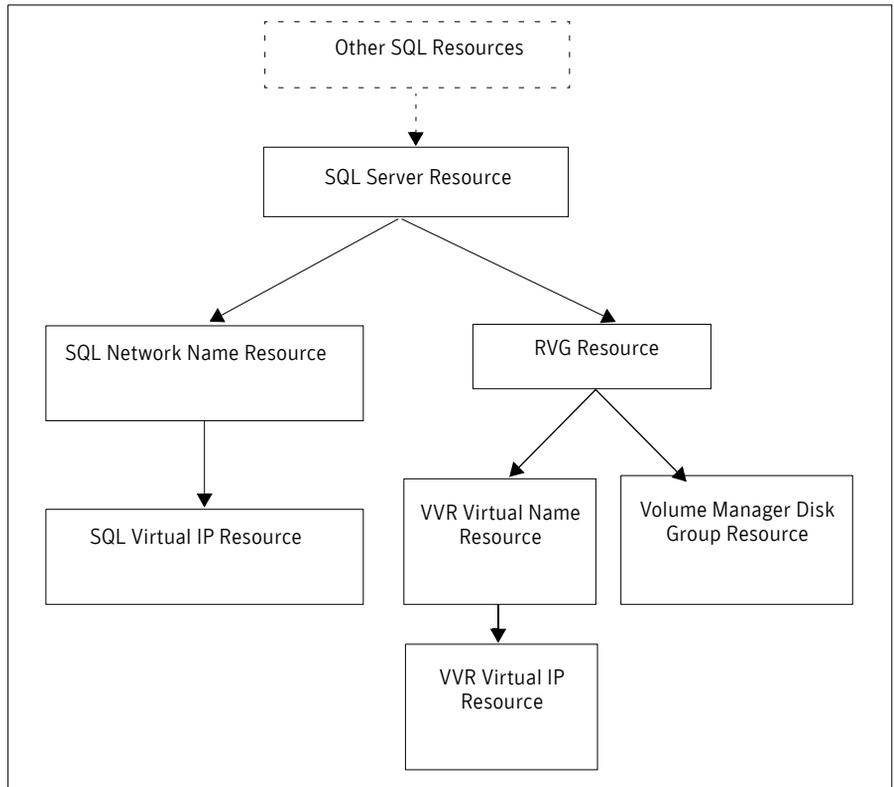
To set the SQL Server application resource dependency on the RVG resource

- 1 Make sure the SQL Server resource is offline before attempting to modify the dependencies. Right-click the resource and click **Take this resource offline**.
- 2 Right-click the SQL Server resource and click **Properties**.
- 3 In the Dependencies tab of the Properties dialog box:
 - Click the box **Click here to add a dependency**.
 - Select the Replicated Volume Group resource from the dropdown list of available resources.
 - Select the Volume Manager Disk Group (VMDG) resource from the dependencies list and click **Delete**.
- 4 Click **OK** to close the Properties dialog box.

The cluster configuration is now complete. Bring online the entire application group on the primary cluster.

The following dependency graph indicates the dependencies that have been established.

Figure 18-3 Dependency graph



Working with the solution: Normal operations and recovery procedures

This section gives considerations for normal VVR operations and also describes the recovery process.

Normal operations

Monitoring the status of the replication

Under normal operating conditions you can monitor the status of the replication using:

- VEA GUI
- Command Line Interface (CLI)
- Performance Monitor (perfmon)
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Volume Replicator Administrator’s Guide*.

Performing planned migration

For maintenance purposes, or for testing the readiness of the Secondary host you may want to migrate the application to the Secondary host. These are a generic set of tasks that you may need to perform.

To detach the user database

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Expand the Databases node. Right-click on the required user database and select **All Tasks > Detach**.

Note that the `master`, `model`, and `tempdb`, databases cannot be detached.

To take the RVG resource offline

Take the RVG resource offline on both the clusters.

To transfer the Primary role

Transfer the Primary role to the Secondary using the **Migrate** option.

- 1 From the VEA screen, right-click the Primary RVG and select **Migrate**.

- 2 Select the Secondary host and click **OK**. The replication role is migrated to the Secondary host.

To assign drive letters to the volumes

Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as that of the original Primary.

To bring the RVG resource online

- 1 Bring the RVG resource online on both the clusters.
- 2 Bring the `SQL_GROUP` group online on the new Primary.

To attach the databases

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click on the required Database node and select **All Tasks > Attach Database**.
- 4 In the Attach Database dialog, specify the name of the Master Data File (MDF) file that corresponds to the database, which you want to attach. Use the browse (...) button to search. For more information, refer to the Microsoft documentation.

You can now verify that the SQL runs fine on the new Primary with the replicated data. After verifying, you can revert back the roles to its original state using the same set of tasks described above.

Note: Any changes that you make to the data on the new Primary will get replicated to the original Primary, which is now the Secondary.

Replication recovery procedures

This section provides information on bringing up an SQL server on the Secondary host, in the event of a disaster. It also explains how to migrate the Primary role back to the original Primary host once it is in a good state after a disaster.

Bringing up SQL on the secondary host

To recover the SQL data

- 1 From the left-pane in the VEA GUI console on the Secondary host, right-click on the desired secondary RVG node inside the replication network. Select the **Take Over** option. The **Take Over** dialog box is displayed.
 - By default, the **Enable Fast-Failback Logging** option is selected if the data volumes have DCM logs associated with them. You can use this option to perform takeover with fast-failback logging. The DCM is activated for fast-failback logging and the new incoming writes are marked on the DCM of the New Primary. If the replication status of Secondary RVG was *Inactive* when the Primary failed, then the **Enable Fast-Failback Logging** option is unavailable for selection. In this case you can perform **Take Over** without using fast-failback logging.
 - Select the **Synchronize Automatically** option if you want the new Primary and the original Primary to get synchronized automatically, after the original Primary recovers. If you have not selected this option, the original Primary, after it recovers will be in the *Acting as Secondary* state. To synchronize this original Primary with the new Primary use the **Resynchronize Secondaries** option from new Primary RVG's right-click menu. When the resynchronization starts, the original Primary which was in the *Acting as Secondary* state is converted to a Secondary of the new Primary. The new Primary now starts replaying the DCM to update the Secondary with the writes that were written to the DCM.
- 2 If you do not want to use the **Enable Fast-Failback Logging** option, clear the checkbox, and click **OK** to perform Take Over without the fast-failback logging. After takeover is complete, to add the Secondary hosts of the original Primary as Secondary hosts of the new Primary, delete the existing RVGs of the original Secondary hosts and then add them as a part of the new Primary.

- 3 If you have chosen to perform the Take Over operation without using fast-failback logging and the original Primary becomes available again, convert it to a Secondary using the **Make Secondary** option. Then resynchronize the original Primary with the new Primary using the **Synchronize Automatically** option. Depending on the size of the data volume this may take quite a while.
Only after the synchronization is complete can you migrate the Primary role back to the original Primary.
After takeover, the existing Secondary becomes the new Primary.
- 4 Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as that of the original Primary.
- 5 Bring the SQL_GROUP group online.

To attach the databases

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click on the required Database node and select **All Tasks > Attach Database**.
- 4 In the Attach Database dialog, specify the name of the Master Data File (MDF) file that corresponds to the database that you want to attach. Use the browse (...) button to search. For more information, refer to the Microsoft documentation.

Now you can start using SQL on the new Primary.

Restoring the primary host

After a disaster, if the original Primary becomes available again you may want to revert the role of the Primary back to this host.

To detach the user database on the new Primary

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Expand the Databases node. Right-click on the required user database and select **All Tasks > Detach**.
Note that the `master`, `model`, and `tempdb`, databases cannot be detached.

To migrate the Primary role back to the original Primary

- 1 Take the RVG resource offline on both the clusters.
- 2 Depending on whether you performed **Takeover** with or without fast-failback option, do one of the following:
 - For Takeover with the Fast-failback option, the original Primary, after it has recovered, will be in the *Acting as Secondary* state. If the original Primary is not in the *Acting as Secondary* state, verify whether your network connection has been restored.
To synchronize this original Primary and the new Primary, use the **Resynchronize Secondaries** option from the right-click menu of the new Primary.
 - For Takeover without the Fast-failback option, after you have performed this operation, you must convert the original Primary to a Secondary using the **Make Secondary** option.

Note: Before performing the **Make Secondary** operation, the original Primary's RVG and the new Primary's RVG will be shown in separate RDSs. However, after this operation they will be merged under a single RDS.

After the **Make Secondary** operation, the original Primary will be converted to a secondary. Right-click this secondary RVG and select **Start Replication** with **Synchronize Automatically** option.

- 3 After the synchronization is complete, perform a migrate operation to transfer the Primary role back to the original Primary. To do this, right-click the Primary RVG and select **Migrate** from the menu.
- 4 Ensure that the volumes have retained the same drive letters that existed before the disaster.
- 5 Bring the RVG resource online on the Secondary.
- 6 Bring the `SQL_GROUP` group online on the original Primary.

To attach the databases on the original Primary

- 1 Open the Enterprise Manager. Select **Start > Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 Expand the SQL Server Group and the corresponding server under it.
- 3 Right-click the required Database node and select **All Tasks > Attach Database**.

In the Attach Database dialog, specify the name of the Master Data File (MDF) file that corresponds to the database that you want to attach. Use the browse (...) button to search. For more information, refer to the Microsoft documentation.

Index

A

Administrative Console 386, 455, 524

B

backup types for snapshot sets 91

backups, snapshot-assisted 70

C

campus cluster

changing quorum resource to dynamic quorum resource 371

configuration 288, 334

creating SFW cluster disk groups and volumes 355

installing and configuring hardware 296, 343

installing application on cluster nodes 369

making quorum cluster disk group a cluster resource 322, 365

Microsoft cluster configuration 284, 288, 331, 334

Microsoft cluster failure scenarios 293, 340

prerequisites 284, 331

setting up a group for the application 321, 365

verifying Microsoft cluster configuration 371
Vxclus 295

clustering concepts

ownership of quorum 294, 341

quorum 294, 341

commands for snapshots

vxsnap 149

vxsnap create 151

vxsnap prepare 150

vxsnap reattach 154

vxsnap restore 155

vxsnapsql create 61

vxsnapsql reattach 62

vxsnapsql restore 65

vxsnapsql start 61

Complete/Custom 386, 455, 524

configuring

disk groups (Microsoft cluster DR) 397, 465, 534

disk groups (Microsoft cluster HA) 182, 221, 262

SQL Server storage with Veritas Storage Foundation for Windows 56, 84

volumes (Microsoft cluster DR) 397, 465, 534

volumes (Microsoft cluster HA) 182, 221, 262

D

database

preferred hardware layout 42, 83

dependencies for SQL Server 2000 196

dependencies for SQL Server 2005 236, 275

disaster recovery (DR)

configuring storage on the primary (Microsoft cluster) 220, 261, 395, 464, 533

creating an RVG resource (Microsoft cluster) 433, 503

Disk Change Object (DCO) 40

disk groups

creating 183, 222, 263

planning (Microsoft cluster) 182, 220, 261, 355, 395, 465, 534

disk space requirements 287, 333

drive letters 140

driver signing options

resetting 394, 463

dynamic disk groups

creating with the command line 57, 85

creating with VEA 56, 84

dynamic volumes

creating with the command line 60, 88

creating with VEA 58, 85

F

failover cluster *see* Microsoft cluster

FastResync 40, 75

FlashSnap

integration with Microsoft Volume Shadow

- Copy Service 76
- integration with SQL Server Virtual Device Interface (VDI) 40
- overview 39, 75

H

- hardware recovery
 - adding disks to the dynamic disk group 142
 - replacing hardware 142

I

- installing and configuring hardware 296, 343
- IP address resource for SQL 191, 231, 408
- IP address resource for VVR 420, 477, 490

M

- Microsoft cluster
 - creating dynamic quorum resource
 - Windows Server 2003 196
 - Windows Server 2008 276
 - creating IP resource for SQL 2000 191
 - creating IP resource for SQL 2005 231
 - creating MSDTC resource 193, 231
 - creating SQL 2000 group 191
 - creating SQL 2005 group
 - Windows Server 2003 230
 - Windows Server 2008 271
 - creating SQL disk group resource
 - Windows Server 2003 192, 231
 - Windows Server 2008 271
 - dependencies for SQL Server 2000 196
 - dependencies for SQL Server 2005 236, 275
 - establishing the cluster
 - Windows Server 2003 171, 210, 298, 381, 450
 - Windows Server 2008 321, 345, 520
 - installing SQL Server 2000 194
 - installing SQL Server 2005 233, 272
 - setting up a group for application 365
 - snapshot considerations 42
 - verifying configuration
 - Windows Server 2003 199, 239
 - Windows Server 2008 278
- mirrors for Quick Recovery snapshots (SQL 2000) 61
- mirrors for Quick Recovery snapshots (SQL 2005) 95, 112, 113

- mount points
 - changing 140
- MSCS *see* Microsoft cluster
- MSDTC resource 193, 231

N

- Network Name resource, RVG 421, 490

P

- pre and post snapshot script file locations 78

Q

- Quick Recovery
 - file locations 78
 - Microsoft clustering considerations 77
 - VVR considerations 77
- Quick Recovery Configuration Wizard
 - overview 72, 89
 - prerequisites 94
 - running 95
- Quick Recovery for SQL 2000
 - best practices recommendations 41
 - components 39
 - creating initial snapshot set 61
 - methods 38
 - Microsoft cluster considerations 42
 - overview 51
 - prerequisites 53
 - recovery 65
 - storage configuration best practices 41
 - storage configuration requirements 40
 - VCS considerations 42
 - VVR considerations 42
- Quick Recovery for SQL 2005
 - components 75
 - hardware requirements 82
 - methods 72
 - overview 69
 - prerequisites 80
 - software requirements 80
 - storage configuration recommendations 82
- quorum resource
 - implementing on Windows Server 2003 237
 - implementing on Windows Server 2008 276

R

- recovery

- SQL Server 2000 65
- SQL Server 2005 65, 129
- replication
 - creating a Replicated Volume Group (RVG) 433, 503, 559
 - creating an IP resource for VVR 420, 490, 547
- requirements
 - disk space 287, 333
- resetting
 - driver signing options 394, 463
- resource
 - IP address for SQL 191, 231, 408
 - IP address for VVR 420, 477, 490
 - MSDTC 193
 - quorum
 - Windows Server 2003 237
 - Windows Server 2008 276
 - Volume Manager Disk Group
 - Windows Server 2003 192, 231
 - Windows Server 2008 271
- resynchronizing
 - secondary site 42
 - snapshot (SQL 2000) 40
 - snapshot (SQL 2005) 126
- RVG Network Name resource 421, 490

S

- sample configurations
 - disk group (Microsoft cluster) 182, 221, 262, 397, 465, 534
 - Quick Recovery 55, 83
 - SQL Server (Microsoft cluster DR) 379, 448, 517
 - SQL Server (Microsoft cluster HA) 167, 207, 247
- schedules for snapshot sets
 - creating with Quick Recovery Configuration Wizard 95
 - creating with VSS Snapshot Scheduler Wizard 114
 - deleting 106, 107, 122
 - modifying 106, 122
 - synchronizing after adding a cluster node 107
 - troubleshooting 104
- SFW
 - installing (Microsoft cluster DR) 383, 452, 522
 - installing (Microsoft cluster HA) 172, 211, 253
- snapshot sets *See* snapshots
- snapshot volumes

- changing drive letter 140
- changing mount points 140
- snapshots
 - Microsoft cluster considerations 77
 - VVR considerations 77
- snapshots for SQL 2000
 - automatically refreshing 64
 - creating 61
 - creating mirrors 61
 - reattaching 62
 - refreshing 62
 - using for recovery 65
- snapshots for SQL 2005
 - backup types 91
 - creating one-time snapshot set 123
 - creating with Quick Recovery Configuration Wizard 95
 - creating with VSS Snapshot Scheduler Wizard 114
 - deleting 122
 - deleting schedules 106
 - modifying 122
 - modifying schedules 106
 - reattaching manually 126
 - reattaching split-mirror 127
 - refreshing 126
 - scheduling with Quick Recovery Configuration Wizard 95
 - scheduling with VSS Snapshot Scheduler Wizard 114
 - templates 91
 - troubleshooting 104
 - using for recovery 129
 - viewing status 103, 119
- Solutions Configuration Center
 - context sensitivity 22
 - overview 21
 - running wizards remotely 29
 - starting 22
 - wizard descriptions 29
- SQL Server storage
 - configuration overview 56, 84
 - creating dynamic disk groups 56, 84
 - creating dynamic volumes 58, 85
 - requirements and best practices 41, 82
- SQL Server Virtual Device Interface (VDI) 40
- synchronized snapshots 116
- synchronizing schedules in a cluster 107

T

- templates for snapshot sets
 - description 91
 - multiple components 92
 - schedule start dates 92
- transaction logs, best practice recommendations 41

V

VCS

- snapshot considerations 42

volume

- quorum resource 197

Volume Manager Disk Group resource

- creating on Windows Server 2003 231
- creating on Windows Server 2008 271

Volume Shadow Copy Service 75

volumes

- creating 185, 224, 265, 313, 316, 359
- creating (campus cluster) 356
- creating on primary 316, 359
- planning (Microsoft cluster) 181, 220, 261, 397, 464, 533
- sample configuration (Microsoft cluster) 182, 221, 262, 397, 465, 534

VSS defined 75

VSS Snapback Wizard

- overview 74
- using 127

VSS Snapshot Scheduler Wizard

- overview 73
- using 114

VSS Snapshot Wizard

- overview 74
- using 123

VVR

- IP address resource 420, 477, 490
- resynchronize secondary node 42
- RVG Network Name resource 421, 490
- snapshot considerations 42

Vxclus utility 295, 342

vxsnap command reference 149

vxsnap commands

- create 151
- reattach 154
- restore 155
- start 150

vxsnapsql

- command syntax 43

- integration with SQL Server VDI 40

vxsnapsql commands

- create 61
- reattach 62
- restore 65
- start 61