

# Veritas Storage Foundation<sup>™</sup> and High Availability Solutions, Solutions Guide

Windows 2000, Windows Server 2003

5.0

# Symantec Storage Foundation and HA Solutions SFW Solutions Guide

Copyright © 2007 Symantec Corporation. All rights reserved.

Storage Foundation 5.0 for Windows

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](http://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://entsupport.symantec.com> 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

About the solutions guides .....	19
About Quick Recovery .....	20
About high availability .....	20
About campus clusters .....	20
About disaster recovery .....	20
About MSCS solutions .....	21
How this guide is organized .....	21

### Chapter 2 Using the Solutions Configuration Center

About the Solutions Configuration Center .....	23
Starting the Configuration Center .....	24
Available options from the Configuration Center .....	24
About running the Configuration Center wizards .....	28
Following the workflow in the Configuration Center .....	29

### Chapter 3 SFW best practices for storage

Best practices for storage availability .....	31
Best practices configuring SFW disk groups for availability .....	33
Best practices for storage performance .....	34
Best practices for I/O performance tuning .....	35
Best practices for storage capacity management .....	38

## Section 2 Quick Recovery

### Chapter 4 Quick Recovery overview

About the Quick Recovery solution .....	44
Need for implementing the SFW Quick Recovery solution .....	44
Understanding the underlying components of	
SFW's Quick Recovery process .....	46
FlashSnap .....	46
FastResync (FR) .....	47
Microsoft Volume Shadow Copy Service (VSS) .....	47
Overview of the Quick Recovery process .....	48
Creating initial snapshots .....	48
Refreshing a snapshot .....	48
Recovering a database .....	49
Other applications for point-in-time snapshots .....	50
Off-host backups .....	50
Reporting and analysis .....	51
Application testing and training .....	52

### Chapter 5 Quick Recovery example

Example of Quick Recovery of an Oracle database .....	54
Create split-mirror snapshot of database .....	54
Recover database using split-mirror snapshot and database logs .....	55
More on FlashSnap: Tips and references .....	58

## Section 3 High Availability

### Chapter 6 High availability: Overview

About high availability .....	61
About clusters .....	61

### Chapter 7 Deploying SFW HA for high availability: New installation

About the high availability solution .....	63
Tasks for a new high availability (HA) installation—additional applications .....	64

Before you begin .....	65
Disk space requirements .....	66
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA) .....	66
Reviewing the configuration .....	69
Configuring the storage hardware and network .....	70
Installing SFW HA .....	71
Setting Windows driver signing options .....	71
Installing Storage Foundation HA for Windows .....	72
Resetting the driver signing options .....	77
Configuring disk groups and volumes .....	77
Planning disk groups and volumes .....	77
Creating dynamic cluster disk groups .....	79
Creating dynamic volumes .....	81
Managing disk groups and volumes .....	84
Importing a disk group and mounting a volume .....	84
Unmounting a volume and deporting a disk group .....	85
Configuring the cluster .....	85
Configuring Web console .....	97
Configuring notification .....	98
Installing and configuring the application or server role .....	101
Configuring a File Share server role .....	101
Configuring a Print Share server role .....	101
Installing and configuring the IIS application .....	102
Installing and configuring Microsoft Virtual Server .....	103
Installing additional applications .....	103
Configuring the service group .....	104
Configuring the File Share service group .....	104
Configuring the PrintShare service group .....	112
Configuring the IIS service group .....	120
Configuring the MSVirtual Machine service group .....	126
Configuring the service group for any additional applications .....	129
Configuring Application Dependencies .....	148
Verifying the cluster configuration .....	151
Possible tasks after completing the configuration .....	152
Configuring the Cluster Management Console connection .....	152
Modifying the existing cluster configuration .....	157
Modifying the application service groups .....	164
Chapter 8	Adding DMP to a clustering configuration
About dynamic multi-pathing .....	170
Overview of configuration tasks for adding DMP ASLs or DMP DSMs .....	171
Reviewing prerequisites .....	171

Reviewing the configuration .....	173
Steps for a new cluster configuration .....	174
DMP ASLs .....	174
DMP DSMs .....	175
Steps for an existing cluster configuration .....	176
DMP ASLs .....	176
DMP DSMs .....	177

## Section 4 Campus Clustering

### Chapter 9 Introduction to campus clustering

Sample campus cluster configuration .....	181
Differences between campus clusters and local clusters .....	183

### Chapter 10 Deploying SFW HA for campus cluster

Reviewing the requirements .....	187
Disk space requirements .....	187
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA) .....	187
Reviewing the configuration .....	191
Overview of campus clustering with VCS .....	192
Reinstating faulted hardware .....	193
Setting the ForceImport attribute .....	195
Installing and configuring the hardware .....	196
Installing Windows and configuring network settings .....	197
Installing Veritas Storage Foundation HA for Windows .....	198
Setting Windows driver signing options .....	198
Installing Storage Foundation HA for Windows .....	199
Resetting the driver signing options .....	204
Configuring the cluster .....	204
Configuring Web console .....	215
Configuring notification .....	216
Creating disk groups and volumes .....	220
Configuring the disks and volumes .....	221
Creating a dynamic (cluster) disk group .....	222
Creating a volume .....	223
Installing the application on cluster nodes .....	229
Creating VCS service groups .....	232
Verifying the cluster configuration .....	233

## Section 5 Disaster Recovery

### Chapter 11 Disaster recovery: Overview

About a disaster recovery solution .....	238
Need for implementing a disaster recovery solution .....	240
Overview of the recovery process .....	241
Components of VVR that enable disaster recovery .....	242
Understanding replication .....	242
Modes of replication .....	242
Features of VVR that help in disaster recovery .....	243

### Chapter 12 Deploying disaster recovery: New application installation

Tasks for a new disaster recovery installation— additional applications .....	246
Before you begin .....	249
Disk space requirements .....	249
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA) .....	249
Reviewing the configuration .....	252
Supported disaster recovery configurations for service group dependencies .....	253
Configuring the storage hardware and network .....	254
Managing disk groups and volumes .....	255
Importing a disk group and mounting a volume .....	256
Unmounting a volume and deporting a disk group .....	256
Setting up the secondary site: Configuring SFW HA and setting up a cluster .....	257
Installing SFW HA .....	257
Setting Windows driver signing options .....	257
Installing Storage Foundation HA for Windows .....	258
Resetting the driver signing options .....	263
Configuring the cluster .....	263
Configuring Web console .....	274
Configuring notification .....	275
Verifying that your application or server role is configured for HA at the primary site .....	278
Configuring the VVR security service .....	278
Configuring disaster recovery .....	281
Assigning user privileges (secure clusters only) .....	282
Cloning the storage on the secondary site using the DR wizard .....	283

Installing and configuring the application or server role .....	287
Installing the FileShare application .....	287
Installing the PrintShare application .....	287
Installing the IIS application .....	287
Installing the Microsoft Virtual Machine application .....	288
Installing additional applications .....	288
Cloning the service group configuration from the primary to the secondary site .....	289
Configuring replication and global clustering .....	292
Verifying the disaster recovery configuration .....	298
Establishing secure communication within the global cluster (optional) .....	300
Possible task after creating the DR environment:	
Adding a new failover node .....	302
Preparing the new node .....	302
Preparing the existing DR environment .....	302
Modifying the replication and application service groups .....	303
Reversing replication direction .....	303
Maintaining: Normal operations and recovery procedures .....	305
Normal operations: Monitoring the status of the replication .....	305
Performing planned migration .....	305
Disaster recovery procedures .....	306
Recovery procedures for service group dependencies .....	307

## Chapter 13    Testing fault readiness by running a fire drill

About disaster recovery fire drills .....	311
About the Fire Drill Wizard .....	311
Tasks for configuring and running fire drills .....	313
Prerequisites for a fire drill .....	313
Fire Drill Wizard actions .....	314
Preparing the fire drill configuration .....	315
Running a fire drill .....	318
Restoring the fire drill system to a prepared state .....	319
Deleting the fire drill configuration .....	320

## Section 6    MSCS Solutions

### Chapter 14    MSCS solutions overview

About high availability .....	325
About campus clustering .....	326
About the SFW-MSCS-VVR configuration .....	327

Chapter 15	Deploying SFW with MSCS	
	Reviewing the requirements .....	332
	Supported software .....	332
	Disk space requirements .....	332
	System requirements .....	333
	Reviewing the configuration .....	334
	Configuring the network and storage .....	335
	Establishing an MSCS cluster .....	337
	Installing SFW .....	337
	SFW installation tasks .....	337
	Pre-installation tasks .....	338
	Installing Veritas Storage Foundation for Windows .....	339
	Post-installation tasks .....	344
	Creating SFW disk groups and volumes .....	346
	Planning disk groups and volumes .....	346
	Creating dynamic cluster disk groups .....	348
	Creating dynamic volumes .....	350
	Setting up a group for the application in MSCS .....	356
	Installing the application on cluster nodes .....	358
	Completing the setup of the application group in MSCS .....	360
	Implementing a dynamic quorum resource .....	361
	Creating a dynamic cluster disk group for the Quorum Resource with mirrored volume .....	362
	Creating the quorum resource for the cluster group .....	362
	Changing the quorum resource to a dynamic mirrored quorum resource .....	363
	Verifying the cluster configuration .....	364

## Chapter 16 Deploying SFW with MSCS in a campus cluster

Reviewing the prerequisites .....	370
Supported software .....	370
System requirements .....	370
Disk space requirements .....	371
Reviewing the configuration .....	372
Overview of campus clustering with MSCS .....	374
MSCS campus cluster failure scenarios .....	375
MSCS quorum and quorum arbitration .....	378
Installing and configuring the hardware .....	381
Establishing the cluster under MSCS .....	382
Installing and configuring the operating system and MSCS on server A .....	382
Configuring the shared storage and creating a partition for the Cluster quorum disk .....	383
Creating the first node of the cluster on server A .....	383
Installing and configuring the operating system and MSCS on server B .....	383
Connecting the two nodes .....	383
Creating the second node of the cluster on server B .....	384
Verifying the cluster configuration .....	384
Installing SFW .....	385
SFW installation tasks .....	385
Pre-installation tasks .....	385
Installing Veritas Storage Foundation for Windows .....	387
Post-installation tasks .....	391
Creating disk groups and volumes .....	393
Configuring the disks and volumes .....	394
Creating a dynamic (cluster) disk group .....	395
Creating a volume .....	397
Setting up a group for the application in MSCS .....	403
Installing the application on the cluster nodes .....	406
Completing the setup of the application group in MSCS .....	408
Changing the quorum resource to a dynamic quorum resource .....	410
Creating a dynamic cluster disk group for the quorum, mirrored .....	410
Making the quorum cluster disk group an MSCS resource .....	411
Changing the quorum resource to the dynamic mirrored quorum resource .....	412
Verifying the cluster configuration .....	413

Chapter 17	Deploying SFW and VVR with MSCS	
	Part 1: Setting up the cluster on the primary site	419
	Reviewing the prerequisites and the configuration	419
	Installing and configuring the hardware	423
	Installing Windows and configuring network settings	423
	Establishing the cluster under MSCS (Primary site)	423
	Installing SFW (Primary site)	424
	Installing Veritas Volume Replicator Security Services (VxSAS)	424
	Creating SFW disk groups and volumes	427
	Setting up a group for the application in MSCS	428
	Installing the application (Primary site)	430
	Completing the setup of the application group in MSCS	430
	Changing the quorum resource to a dynamic quorum resource	431
	Testing of the cluster on the primary site	434
	Part 2: Setting up the cluster on the secondary site	436
	Repeating cluster configuration steps for the secondary site	436
	Part 3: Adding the VVR components for replication	439
	VVR components overview	439
	Configuring the Replicator Log volumes for VVR	440
	Setting up the Replicated Data Sets (RDS) for VVR	442
	Creating an RVG resource and setting the dependencies	450
	Part 4: Maintaining normal operations and recovery procedures	453
	Normal operations: Monitoring the status of the replication	453
	Performing planned migration	453
	Disaster recovery procedures	454
Section 7	Server Consolidation	
Chapter 18	Server consolidation overview	
	Server consolidation definition	459
	Need for implementing server consolidation	459
	Advantages of using SFW with server consolidation	460
	Overview of the server consolidation process	462

<b>Chapter 19</b>	<b>Server consolidation configurations</b>	
	Typical server consolidation configuration .....	464
	Proof of concept .....	464
	Server consolidation configuration 1 – many to one .....	465
	About this configuration .....	465
	Reviewing the configuration requirements .....	468
	Preparing to consolidate .....	470
	Migrating the data to the large server .....	471
	Adding the storage array .....	472
	Completing the consolidation process .....	472
	Server consolidation configuration 2 – many	
	to two: Adding clustering and DMP .....	473
	About this configuration .....	473
	Reviewing the configuration requirements .....	476
	Adding the new hardware .....	477
	Establishing the MSCS cluster .....	478
	Adding SFW support to the cluster .....	478
	Setting up MSCS cluster groups for the applications .....	479
	Installing applications on the second computer .....	479
	Completing the setup of the application group in MSCS .....	480
	Changing the quorum resource to the dynamic quorum resource ....	480
	Verifying the cluster configuration .....	480
	Enabling DMP .....	480
	SFW features that support server consolidation .....	481
	Automatic volume growth .....	481
	Features that support storage in a SAN .....	482
	Performance monitoring .....	482
	Server consolidation customer success story .....	482

## Section 8 Appendix

<b>Appendix A</b>	<b>Deploying Disaster Recovery: Manual implementation</b>	
	Part 1: Setting up the cluster on the primary site .....	492
	Reviewing the requirements .....	492
	Disk space requirements .....	492
	Requirements for Veritas Storage Foundation High Availability for	
	Windows (SFW HA) .....	493
	Reviewing the configuration .....	496
	Installing and configuring the hardware .....	497
	Installing Windows and configuring network settings .....	497

Installing SFW HA (Primary site) .....	499
Setting Windows driver signing options .....	499
Installing Storage Foundation HA for Windows .....	500
Configuring VVR security service .....	504
Resetting the driver signing options .....	506
Configuring the cluster (Primary site) .....	506
Configuring disk groups and volumes (Primary site) .....	523
Planning disk groups and volumes .....	523
Creating dynamic cluster disk groups .....	525
Creating dynamic volumes .....	527
Installing the application on cluster nodes (Primary site) .....	530
Creating VCS service groups (primary site) .....	532
About VCS service groups .....	532
Service group example with a generic database application .....	533
Verifying the cluster configuration .....	542
Part 2: Setting up the cluster on the secondary site .....	544
Creating a parallel environment on the secondary site .....	544
Configuring disk groups and volumes (Secondary site) .....	545
Installing the application (Secondary site) .....	545
Configuring the Service group for VCS (Secondary site) .....	545
Verifying the cluster configuration (Secondary site) .....	545
Part 3: Adding the VVR components for replication .....	546
VVR components overview .....	546
Configuring the Replicator Log volumes for VVR .....	547
Setting up the replicated data sets (RDS) for VVR .....	549
Creating the VVR RVG Service group .....	557
Part 4: Adding GCO components for wide-area recovery .....	561
Prerequisites for a global cluster environment .....	561
Linking clusters by adding a remote cluster .....	562
Converting a local Service group to a global group .....	564
Additional global cluster administration tasks .....	566
Part 5: Maintaining: Normal Operations and recovery procedures .....	568
Normal operations: Monitoring the status of the replication .....	568
Performing planned migration .....	568
Disaster recovery procedures .....	569
Index.....	571



# Introduction

- [Introducing Veritas Storage Foundation and High Availability Solutions](#)
- [Using the Solutions Configuration Center](#)
- [SFW best practices for storage](#)



# Introducing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [About the solutions guides](#)
- [About Quick Recovery](#)
- [About high availability](#)
- [About campus clusters](#)
- [About disaster recovery](#)
- [About MSCS solutions](#)
- [How this guide is organized](#)

## About the solutions guides

The *Veritas Storage Foundation and High Availability Solutions, Solutions Guide* contains solutions for the following:

- Quick Recovery
- High availability (HA)
- Campus clusters
- Disaster recovery (DR)
- MSCS Solutions

Solutions for Quick Recovery and MSCS Solutions are in *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft Exchange*.

Separate guides are available for Microsoft Exchange and Microsoft SQL solutions.

## About Quick Recovery

Quick Recovery is the process of creating and maintaining on-host point-in-time images of dynamic volumes that can be used to quickly recover from logical errors in data files.

Quick Recovery is designed to augment your traditional backup methodology.

## About high availability

The term high availability (HA) refers to a state where data and applications are highly available because software or hardware is in place to maintain the continued functioning in the event of computer failure. High availability can refer to any software or hardware that provides fault tolerance, but generally the term has become associated with clustering. Local clustering provides high availability through database and application failover. Veritas Storage Foundation HA for Windows (SFW HA) includes Veritas Storage Foundation and Veritas Cluster Server and provides the capability for local clustering.

## About campus clusters

A campus cluster is a single cluster that stretches over two sites using fiber channel connectivity, with SAN connections for data mirroring and network connections for cluster communication. Although two sites are the most common, more than two can be used for additional redundancy.

Campus clusters provide disaster protection when an entire site goes down by locating the clustered servers in different buildings or areas. This solution provides a level of high availability that is above mirroring or clustering at a single site and is an alternative to using replication software.

## About disaster recovery

Wide area disaster recovery (DR) provides the ultimate protection for data and applications in the event of a disaster. If a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or

thousands of miles away. Veritas Storage Foundation HA for Windows (SFW HA) provides the capability for implementing disaster recovery.

## About MSCS solutions

Microsoft Cluster Server (MSCS) may be used with Veritas Storage Foundation for Windows to provide high availability for any application or server role.

MSCS may be used with Veritas Storage Foundation for Windows and Veritas Volume Replicator to provide high availability and replication support.

## How this guide is organized

Where applicable, the *Veritas Storage Foundation and High Availability Solutions, Solutions Guide* is organized to follow the workflow in the Solutions Configuration Center.

See [Chapter 2, “Using the Solutions Configuration Center”](#).

When setting up a site for disaster recovery using the Configuration Center, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration. Likewise, in this guide, you first follow the instructions in the high availability section and then continue with the appropriate chapter in the disaster recovery section.

The Solutions Configuration Center includes a number of wizards that were not available in earlier versions of the product, including a Disaster Recovery wizard. The earlier methods of setting up disaster recovery manually, without the wizard, are available in an appendix section.



# 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](#)

## About the Solutions Configuration Center

The Storage Foundation and High Availability Solutions Configuration Center guides you through setting up your SFW HA environment. The Configuration Center provides solutions for Microsoft Exchange, Microsoft SQL Server 2005, and for additional applications.

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

- High availability at a single site for a new installation
- Wide area disaster recovery involving two or more sites
- Quick Recovery for on-host recovery from logical errors in application data (available for Microsoft Exchange and for Microsoft SQL Server 2005)
- Fire drill to test the fault readiness of your disaster recovery environment

# Starting the Configuration Center

You can start the Configuration Center by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.

## Available options from the Configuration Center

The Solutions Center is context sensitive to the application. For example, the Solution Guides listed in the right pane match the selected application.

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

**Figure 2-1** Solutions Configuration Center for Microsoft Exchange

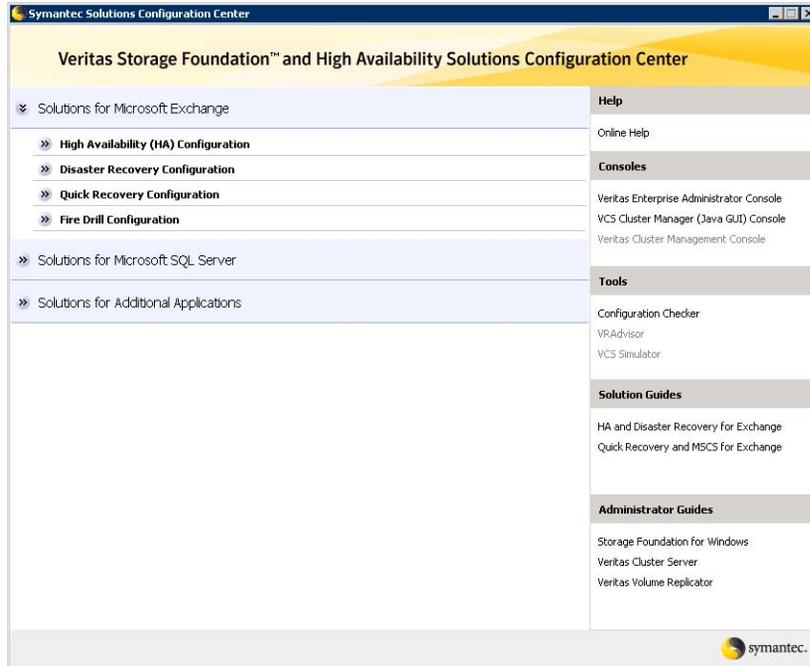


Figure 2-2 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 2-2 Solutions Configuration Center for Microsoft SQL Server

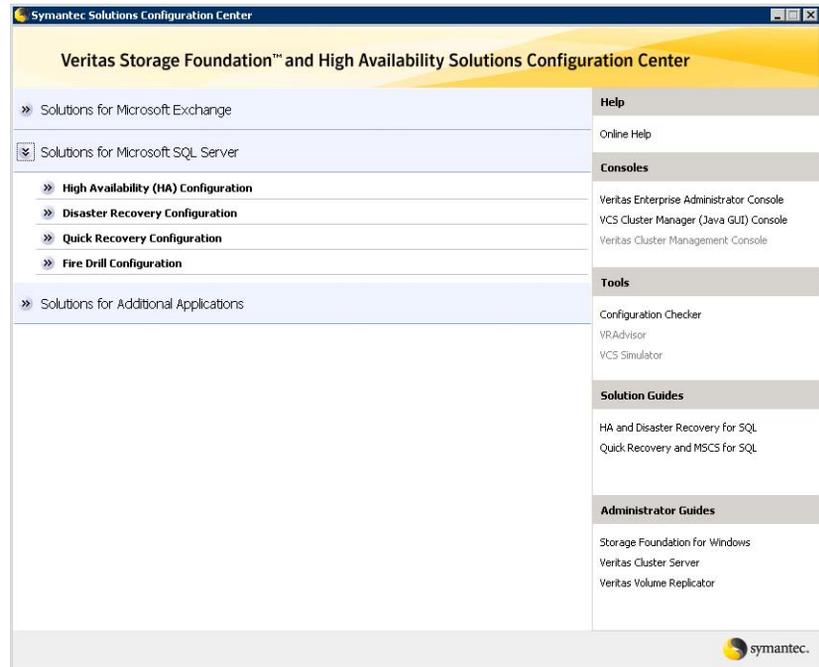
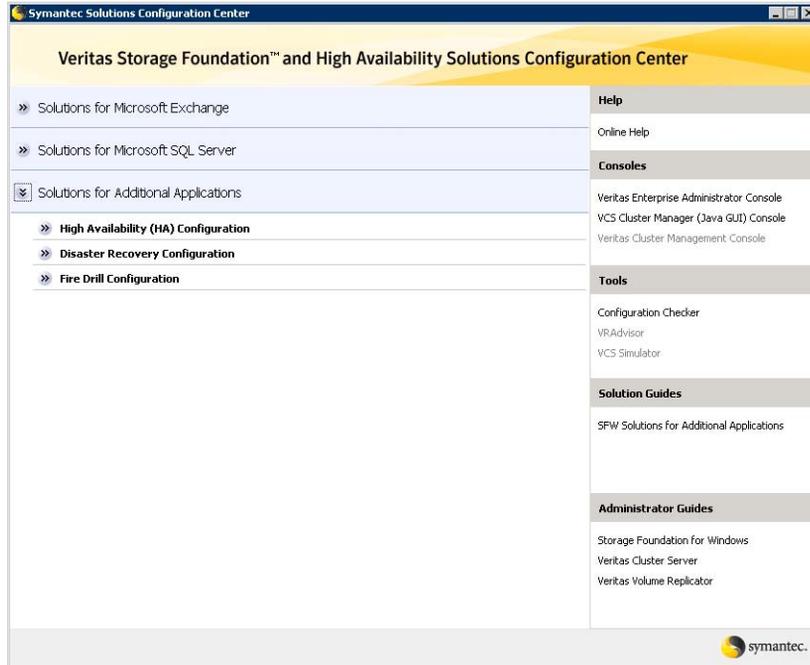


Figure 2-3 shows the choices available when you click Solutions for Additional Applications.

Figure 2-3 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 2-4 shows one of the steps for implementing high availability for Exchange.

Figure 2-4 Context-sensitive step for Exchange



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

Figure 2-5 Context-sensitive step for SQL Server



Figure 2-6 shows one of the steps for implementing high availability for additional applications.

Figure 2-6 Context-sensitive step for additional applications



## About running the Configuration Center wizards

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

Veritas Cluster Wizard	Sets up the VCS cluster
Disaster Recovery Wizard	Configures wide area disaster recovery involving two sites Requires first configuring high availability on the first site
Quick Recovery 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
Application Agent for Exchange Setup Wizard	Installs and configures Exchange for the high availability environment If Exchange is already installed, refer to the documentation for further instructions.
Application Agent for Exchange Configuration Wizard	Configures the service group for Exchange high availability
Database Agent for SQL 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.

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:

- FileShare Configuration Wizard    Configures FileShare for high availability.
- PrintShare 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.
- 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 2-7 shows the high-level overview of the workflow steps for configuring high availability for Exchange from the Solutions Configuration Center.

**Figure 2-7**      Workflow for configuring Exchange high availability



Figure 2-8 shows the high-level overview of the workflow steps for configuring high availability for SQL Server from the Solutions Configuration Center.

**Figure 2-8** Workflow for configuring SQL Server high availability

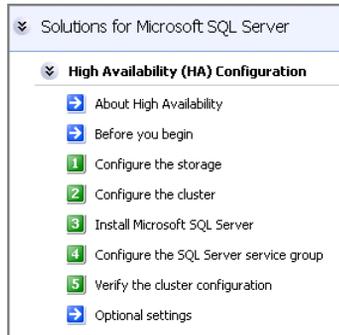
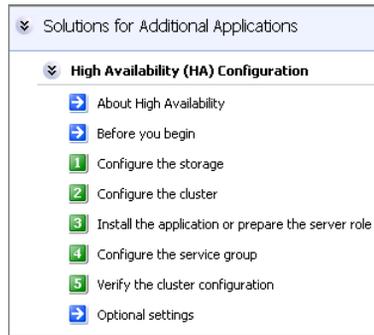


Figure 2-9 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

**Figure 2-9** Workflow for configuring high availability for additional applications



# SFW best practices for storage

This chapter includes the following topics

- [Best practices for storage availability](#)
- [Best practices configuring SFW disk groups for availability](#)
- [Best practices for storage performance](#)
- [Best practices for I/O performance tuning](#)
- [Best practices for storage capacity management](#)

## Best practices for storage availability

For high-availability of storage, use the following best practices to ensure continued access to data:

- [Adding software mirrors for critical data](#)
- [Locating data objects for optimum recovery](#)
- [Managing three-way software mirrors for reliability](#)
- [Software striping and mirroring on top of hardware RAID for high availability](#)

### Adding software mirrors for critical data

For data that is absolutely critical to enterprise operation, use three-way mirrored volumes.

Using host-based volume management to construct the mirrors takes very little CPU (server) bandwidth, because the data and log writes are concurrent. Additionally, connect the disks composing a three-way mirrored volume to the

hosts using independent paths (cables, host bus adapters, connectors) to protect against path failure as well as disk failure.

### **Locating data objects for optimum recovery**

When laying out volumes on disks with SFW, locate data objects that depend on each other on separate volumes, on separate disks. This ensures that a single disk failure does not destroy both data and its recovery mechanism.

Enterprise server environments often have interdependent sets of data. For example, the datasets in a database, its archive logs, and its redo log all depend on each other. If a volume holding database data fails, causing data loss, the typical practice would be the following steps:

- Repair the cause of the failure (for example, replace one or more disks).
- Restore the database to some baseline from a backup copy.
- Play the archive and redo logs against the restored copy to bring the database state as close to current as possible.

If the database logs reside on the same volume as the data, however, both data and logs are inaccessible, and database recovery is impossible.

### **Managing three-way software mirrors for reliability**

The storage on a broken-off mirror should be restored to the original volume during periods of low application I/O load, if possible, because resynchronization and regeneration are I/O-intensive activities that can adversely affect application performance.

If a mirror is regularly broken off from a three-way mirrored volume, note that susceptibility to failure is greater during the interval between the third-mirror breakoff and the completion of resynchronization after the third-mirror storage is returned to the mirrored volume. By performing the restoration during periods of low I/O load, the susceptibility-to-failure window is minimized.

### **Software striping and mirroring on top of hardware RAID for high availability**

For data protection in addition to the performance of striping, apply mirroring on top of striping. Perform the striping first and then mirror the striped volumes.

Striped volumes store large amounts of low-value or easily reproduced data where rapid access is required. Because striping alone will not maintain the availability of the data in a disk failure, consider mirroring also.

# Best practices configuring SFW disk groups for availability

Storage Foundation for Windows supports multiple disk groups. Disk groups provide a way of organizing physical disks in a system into logical entities and simplify storage management for systems with large numbers of disks.

Disk groups are useful for managing storage in clusters, as well as convenient for organizing and managing disk storage resources for applications. SFW allows moving disks between host systems, providing an easy method of transferring storage from one system to another.

For high-availability with SFW disk groups, use the following best practices to ensure continued access to data:

- [Configuring disk groups for separate storage capacity pools or common pools](#)
- [Allocating disk groups for availability in clusters](#)

## Configuring disk groups for separate storage capacity pools or common pools

Creating multiple disk groups creates separate storage capacity pools.

Effective use of subdisks is key to efficient disk group structure. The subdisks composing any given volume must be allocated from disks within a single disk group. Raw physical storage in one of these pools is available exclusively for use within the pool and cannot be used in other disk groups, unless an administrator specifically moves a disk from one disk group to another.

System administrators must decide on the basis of projected application and administrative needs whether to use disk groups to create disjointed storage pools or to manage all storage as a common pool.

Effective configuration of disk groups depends on an organization's application needs.

- If a critical application requires frequent volume expansion, allocating its storage in a private disk group helps guarantee that capacity is available when required. When storage capacity is added to the system, it is not absorbed by other applications.
- If a critical application unexpectedly requires additional storage and none is available in the disk group from which its volumes are allocated, the application will fail, even if the required amount of storage is available in other disk groups.

In general, multiple pools give the administrator greater flexibility, whereas a common pool may be more convenient for applications.

## Allocating disk groups for availability in clusters

In a cluster, each application that fails over independently of other applications should have its data stored on volumes in disk groups exclusive to that application.

In a clustered environment, the disk group is the unit in which storage fails over from one computer to another. Only entire disk groups fail over. Thus, volumes that hold data for applications that are required to fail over should belong to disk groups that hold data for that application only. This allows an application's storage to fail over with the application and have no adverse effects on other applications or their associated storage. The disk groups should also be part of the application's resource group, so that failover can occur.

# Best practices for storage performance

For optimal performance of storage, use the following best practices to ensure fast access to data:

- [Best practices for storage performance](#)
- [Host-based mirroring for increased read performance and failure tolerance](#)
- [Using software RAID 5 for read-mostly data](#)

## Host-based mirroring for increased read performance and failure tolerance

Use host-based mirroring of virtual disks to increase overall system read performance and failure tolerance.

SFW provides host-based volume management to RAID subsystems, increasing overall data availability and I/O performance. With host-based volume management, software RAID can be applied across RAID subsystems from the same or different vendors, thus aggregating all the desirable properties of RAID subsystems.

In a mirrored configuration, read requests are handled in a round-robin fashion. The round-robin algorithm distributes read requests across all members, or "plexes," of a mirrored volume. Mirroring can increase read performance significantly.

Additionally, by configuring the hardware RAID subsystem-based virtual disks exported by different controllers as members of a host-based mirrored volume, the host-based mirrored volume provides protection against I/O bus, host bus adapter, enclosure power and cooling, RAID controller, and disk failures.

## Software striping across hardware for increased performance

Use SFW to combine multiple hardware arrays connected to the host via multiple buses in a single large striped volume, for higher transfer rates with some applications.

Host-based volume management can be used to aggregate the performance of multiple hardware subsystems by striping data across two or more virtual disks, each managed by a different RAID controller. Construct stripes across similar devices for the best use of storage. Because certain high-bandwidth applications, such as audio-visual streaming, have data transfer requirements that surpass the capability of a hardware array controller attached to the host by a single connection, the ability to aggregate the bandwidth of multiple data buses is needed.

## Using software RAID 5 for read-mostly data

Host-based RAID-5 volumes are recommended for read-mostly data, because noticeable performance degradation may occur due to the overhead that writes generate.

Host-based RAID-5 volumes should be avoided in applications in which the rate of updates is high (more than about 10% of the aggregate I/O request-handling capacity of the disks constituting the volume), unless sufficient host CPU cycles are available. Disk controller RAID-5 volumes equipped with nonvolatile write-back cache may be used for more write-intensive applications (up to about 40% of the aggregate I/O request capacity of the disks composing the volume).

# Best practices for I/O performance tuning

SFW enables administrators to “tune” any type of striped volume, including RAID-5 and mirrored striped volumes, by adjusting the stripe unit size. This feature is particularly useful for optimizing the I/O performance of these volume types.

Most I/O-bound applications can be characterized as one of the following:

- I/O-request intensive, making I/O requests faster than the hardware to which they are made can satisfy them  
With rare exceptions, transaction-oriented applications (for example, credit verification, point of sale, order taking) are I/O-request intensive. See “[Striping for I/O-request-intensive applications](#)” on page 36.
- Data-transfer intensive, moving large single streams of data between memory and storage

Scientific, engineering, audio, video, and imaging applications are typically data-transfer intensive.

See “[Striping for data-transfer-intensive applications](#)” on page 37.

If a striped volume will be used predominantly for one or the other of these I/O load types, the stripe unit size can be set at volume creation to optimize I/O performance.

## Striping for I/O-request-intensive applications

A good compromise stripe unit size for I/O-request-intensive applications is one that results in a 3% to 5% probability of splitting in a uniform distribution of requests. For example, a 2 KB (four-block) database page size would have an ideal stripe unit size of 100 blocks. This would typically be rounded up to the nearest power of two (128 blocks, or 65,536 bytes) for simplicity.

I/O-request-intensive applications are typically characterized by small (for example, 2 to 16 KB) data transfers for each request. These applications are I/O bound because they make so many I/O requests, not because they transfer large amounts of data.

For example, an application that makes 1,000 I/O requests per second with an average request size of 2 KB uses at most 2 MB per second of data transfer bandwidth. Because each I/O request occupies a disk completely for the duration of its execution, the way to maximize I/O throughput for I/O-request-intensive applications is to maximize the number of disks that can be executing requests concurrently. Clearly, the largest number of concurrent I/O requests that can be executed on a volume is the number of disks that contribute to the volume’s storage. Each application I/O request that “splits” across two stripe units occupies two disks for the duration of its execution, reducing the number of requests that can be executed concurrently and thus the efficiency of I/O response.

Therefore, try to minimize the probability that I/O requests “split” across stripe units in I/O-request-intensive applications.

The following factors influence whether an I/O request with a random starting address will split across two stripe units:

- The request starting address relative to the starting address of the storage allocation unit (the file extent)
- The size of the request relative to the stripe unit size

Most database management systems will allocate pages in alignment with the blocks in a file, so that requests for the first page will almost never split across stripe units. However, database requests for two or more consecutive pages may split across stripe units. In this case, larger stripe unit sizes reduce the probability of split I/O requests. However, the primary objective of striping data

across a volume is to cause I/O requests to be spread across the volume's disks. Too large a stripe unit size is likely to reduce this spreading effect.

## Striping for data-transfer-intensive applications

The ideal stripe unit size for data-transfer-intensive applications that use a striped volume is the typical I/O request size of the application, divided by the number of data disks in the stripe. For example, if an application typically makes requests for 256 KB, an ideal stripe size for a four-disk striped volume would be 64 KB (256 KB/4).

Data-transfer-intensive applications typically request a large amount of data with every request, between 64 KB and 1 MB, or more. When a large amount of data is requested, the data-transfer phase of the request represents the majority of the request execution time. Thus, reducing data-transfer time improves I/O performance.

A single disk can transfer data only as fast as the data passes under the disk's read-write head. For example, a disk that rotates at 10,000 RPM and has 200 blocks on a certain track cannot transfer data to or from that track any faster than 17.06 MB per second (200 blocks x 512 bytes per block/0.006 seconds per revolution). An application request for 500 KB would require five platter revolutions, or 30 milliseconds, to execute. If the request were addressed to a volume of five identical disks created with SFW, each disk would ideally deliver one-fifth of the data, and the request would complete in a shorter time.

In general, if a striped volume is optimized for data-transfer-intensive applications, each application I/O request will split evenly across all of the volume's disks (or all but the disk containing parity data in the case of a RAID-5 volume).

## Best practices for storage capacity management

Maintaining a percentage of unallocated storage capacity in a disk group is a useful means of managing online storage to avoid application failures. When an application requires more storage, its volumes can be extended quickly and easily by a system administrator while it's online, using the unallocated capacity. If volume expansion causes unallocated capacity to drop below a safety threshold, the event can be displayed in the GUI provided with SFW. Additional storage should then be installed and added to the disk group to maintain an adequate cushion for anticipated application requirements.

For storage capacity management, use the following best practices to ensure the best allocation of data:

- [Managing storage allocation for flexibility](#)
- [Aggregating hardware RAID for very large volumes](#)
- [Managing unallocated space for free space savings](#)
- [Reserving spares for failure-tolerant volume recovery](#)

### Managing storage allocation for flexibility

One way to maximize the flexibility of storage allocation is to manage the disks in a disk group in units of a single capacity or of a small number of discrete capacities. This maximizes SFW's flexibility to allocate storage when new subdisks are required for new volumes, for volume extension, or for moving a subdisk from one disk to another.

To ensure that the amount of unallocated storage in each disk group is adequate, an appropriate level of unallocated storage be maintained. The distribution of unallocated storage across disks must allow for management operations such as failure-tolerant volume expansion to be carried out without violating volume failure tolerance and performance restrictions.

For example, if an additional mirror must be added to a mirrored striped volume, each subdisk of the added mirror must be located either on the same disk as the subdisk it extends, or on a disk separate from any of the volume's existing subdisks. (A subdisk is defined as a number of consecutively addressed blocks on a disk.) Subdisks are created by SFW as building blocks from which volumes are created. When an administrator makes a request to extend a volume, SFW checks the unallocated space in the disk group containing the volume to make sure that extension is possible. An administrator must maintain a distribution of unallocated capacity that allows such operations.

## Aggregating hardware RAID for very large volumes

Combine LUNs from multiple RAID controllers with SFW to construct a very large volume capable of holding a very large database or file system, spanning multiple LUNs across controllers. This combined capability can give users better access to their data than if the file system or database is split across multiple LUNs.

For any type of hardware RAID used in an array, the size of a database or file system is limited to the maximum size of a logical unit number (LUN) in the particular hardware array. However, this limitation is removed with advanced host-based volume management.

## Managing unallocated space for free space savings

Any policy for maintaining a minimum percentage of a disk group's capacity as unallocated space should include a cap to avoid maintaining wastefully large amounts of free space.

How much unallocated capacity to maintain depends strongly on application characteristics. In most cases, however, there are lower and upper bounds beyond which less or more unallocated storage would be of little use.

For example, an installation may observe a policy of maintaining a level of 8% to 10% of a disk group's total capacity as unallocated space. As the capacity of the disk group grows, however, the amount of unallocated space maintained by this policy can grow beyond any reasonable expectation of exploiting it effectively. If unallocated space is typically used in quantities of around 1 to 10 GB to relocate subdisks or to accommodate data processing peaks, then growing the disk group to 1 TB total capacity would mean that 100 GB are reserved for this purpose. If the typical number of subdisk moves or volume adds is one or two, a significant amount of storage capacity would never be used.

## Reserving spares for failure-tolerant volume recovery

Reserve one or more spare disks for every 10 disks that are part of a failure-tolerant volume, with a minimum of one spare disk for any disk group that contains failure-tolerant volumes.

Storage capacity is managed in subdisk units, but entire disks usually fail. Because entire disks fail, spare capacity reserved for recovering from disk failures should be entire disks whose capacity is at least as large as that of the largest disk in a failure-tolerant volume in the disk group.

When a disk fails, all non-failure-tolerant volumes having subdisks on it fail, and all failure-tolerant volumes become degraded.



# Quick Recovery

This section presents the SFW quick recovery solution, a process that uses split-mirror snapshots of production data and a transaction log to recover a database that has been corrupted or that has missing data.

This section has the following chapters:

- [Chapter 4, “Quick Recovery overview” on page 43](#)
- [Chapter 5, “Quick Recovery example” on page 53](#)



# Quick Recovery overview

Veritas FlashSnap is an option to SFW that is a highly efficient procedure involving multiple commands that allows detaching of a mirrored volume. Once the volume is detached, it can be used for a variety of purposes. This chapter focuses on the SFW Quick Recovery solution, which uses split-mirror snapshots to recover from logical errors in data files. It also gives a summary of the other uses for split-mirror snapshots.

The chapter's topics are:

- [“About the Quick Recovery solution”](#) on page 44
- [“Need for implementing the SFW Quick Recovery solution”](#) on page 44
- [“Understanding the underlying components of SFW’s Quick Recovery process”](#) on page 46
- [“Overview of the Quick Recovery process”](#) on page 48
- [“Other applications for point-in-time snapshots”](#) on page 50

[“Example of Quick Recovery of an Oracle database”](#) on page 54 provides an example of a Quick Recovery solution with an Oracle database.

---

**Note:** This chapter gives a general overview of SFW’s Quick Recovery solution. For detailed information about implementing a Quick Recovery solution with SFW and Microsoft Exchange Server, see the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft Exchange*. For detailed information about implementing a Quick Recovery solution with SFW and Microsoft SQL Server, see the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft SQL*.

---

## About the Quick Recovery solution

Quick Recovery is the process of using on-host point-in-time copies of production data and a transaction log to recover a database that has been corrupted or that has missing data. If a database becomes corrupted, for example, you could reload the original data from the most recent snapshot, and then use the transaction log to bring the database current to the point before the corruption.

The SFW Quick Recovery solution uses on-host, disk-based snapshots to provide fast recovery from logical errors and eliminates the time-consuming process of restoring data from tape.

If you are using Microsoft Exchange 2003, SFW has recovery procedures for Microsoft Exchange storage groups or individual databases within an Exchange storage group. Additionally, Quick Recovery of Microsoft SQL 2005 databases is supported. Those procedures are provided through SFW's `vxsnap restore` command and the VSS Snapshot wizards.

SFW also provides the `vxsnapsql` utility for Quick Recovery which 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. The `vxsnapsql` utility is supported for SQL 2000 and SQL 2005.

## Need for implementing the SFW Quick Recovery solution

Advantages of using SFW's Quick Recovery solution:

- **Faster than Restoring from Tape or Other Media**  
On-host snapshot recovery is faster than restoring a full backup from tape or other media; this reduces downtime and helps meet service-level agreements for application availability. A Quick Recovery solution serves as a first line of defense to recover a corrupted database or missing data. The impact on system performance of maintaining a Quick Recovery image is limited to the brief time of detaching a split-mirror snapshot from its original volume.
- **A Less Costly and More Flexible Solution than Array-based Snapshots**  
SFW's split-mirror snapshots are based on the FlashSnap technology. FlashSnap puts the snapshot logic on the host system itself, so you can use any storage you have or might acquire to create snapshots. This is in

contrast to a split-mirror image created through a hardware storage array, where you are limited to only the storage provided by the array vendor. FlashSnap provides several benefits over a hardware-based approach:

- You can use virtually any storage hardware to create snapshots, including expensive arrays and simple JBOD storage.
- The volumes that are copied can span physical devices.
- The original and snapshot volumes can use different vendors' storage arrays.
- **Integration with Microsoft Server 2003 Volume Shadow Copy Service**  
SFW integrates with the Windows Server 2003 Volume Shadow Copy Service (VSS) as both a VSS requestor and a VSS provider. This integration is provided by FlashSnap and SFW's `vxsnap` command line utility and VSS Snapshot wizards. The VSS process enables a VSS-aware application, such as Exchange 2003, to be quiesced before the snapshot operation occurs and then resumed immediately after it. This pause of the application can produce Microsoft supported and guaranteed snapshots of your data. It protects the integrity of your data.
- **Allows Multiple Snapshots at One Time**  
SFW offers the option to create simultaneous, multiple split-mirror snapshots. These snapshots can be done either through the GUI **Snap Shot** command or through the `vxsnap` CLI command.

## Understanding the underlying components of SFW's Quick Recovery process

SFW's Quick Recovery solution uses Veritas FlashSnap and FastResync technology to leverage the Microsoft Volume Shadow Copy Service (VSS) capability to pause and resume a VSS-aware application.

### FlashSnap

FlashSnap provides the ability to create and maintain the on-host point-in-time copies that are integral to the Quick Recovery solution. FlashSnap is the multi-step process used to create and maintain split-mirror snapshots that are copies of the original volumes they mirror. Both the original and snapshot volumes 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. The **Prepare** procedure may take considerable time because it involves creating a mirror, but it has to be done only the first time you perform the snap commands sequence.

---

**Note:** The **Prepare** command replaces the **Snap Start** command in the GUI. Both `prepare` and `snapstart` keywords are available in the CLI, however `prepare` is the recommended keyword.

---

#### **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 mirror 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, the snapshot volume can be used to quickly restore a consistent, point-in-time image to the original volume.

#### **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.

The FlashSnap commands listed above are implemented through the SFW GUI. There are also command line equivalents, using the `vxassist` or `vxsnap` command.

## FastResync (FR)

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. When the snapshot volume is resynchronized with the original volume by 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 Quick Recovery image can be refreshed with minimal impact on production.

FR is automatically enabled for a volume when the prepare operation is performed on the volume through the GUI **Prepare** command or the command line interface `vxassist snapstart` command.

## Microsoft Volume Shadow Copy Service (VSS)

Microsoft Volume Shadow Copy Service (VSS) is a Windows Server 2003 service that provides the capability of creating snapshots or volume shadow copies. A volume shadow copy is a volume that represents a duplicate of the state of the original volume at the time the copy began. SFW integrates VSS into its snapshot function through the `vxsnap` command. Because SFW is a VSS requestor, it can initiate VSS snapshots at any time.

In the Windows Server 2003 version, `vxsnap` makes use of both FlashSnap and VSS technology to create high-quality snapshots that can be done when application files are open. VSS can quiesce the application for the moment when the snapshot is created and then resume the application immediately after the snapshot; but a VSS-aware application must be used, such as Microsoft Exchange Server 2003.

The Windows 2000 version of `vxsnap` does not make use of VSS. It relies solely on SFW's FlashSnap technology. VSS is not supported in Windows 2000.

Both versions of `vxsnap` allow you to name the snapshot volume. They also both require you to use the Prepare command first to create a mirror ready to be used for a snapshot. For more information on how VSS and SFW work together, see Chapter 8 of the *Veritas Storage Foundation Administrator's Guide*.

## Overview of the Quick Recovery process

The Quick Recovery process can be broken down into three phases: creating, refreshing, and recovering.

### Creating initial snapshots

Split-mirror snapshots should be created on a regular schedule, following the backup of the database from tape. You can snapshot a database volume by itself or you can use the SFW GUI **Snap Shot** command or the `vxsnap` utility to snapshot one or more database volumes and any database log volumes simultaneously. If you are using Windows Server 2003 and have an application that is VSS-aware, such as Microsoft Exchange Server 2003 or Microsoft SQL Server 2005, you have the advantage of creating VSS snapshots. By taking VSS-enabled snapshots, you can create snapshot images without needing to take the database offline. Additionally, SFW offers the `vxsnapsql` utility for Microsoft SQL 2000 or Microsoft SQL 2005.

Creating a snapshot is a two-step process. The first step, **Prepare**, creates the snapshot mirror attached to the original volume. The second step, **Snap Shot**, detaches the snapshot mirror from the original volume and creates a separate on-host split-mirror snapshot volume.

Once a snapshot has been created, it can be refreshed quickly without repeating the time-consuming **Prepare** step.

### Refreshing a snapshot

Periodically refresh or update your snapshot or set of snapshots so they contain a current copy of the original volumes. Refreshing a snapshot is a two-step process. During the first step, the **Snap Back** operation reattaches a snapshot volume to its original volume and uses Fast Resync to automatically update the snapshot mirror and synchronize it with the original volume, applying only the changes tracked in the Disk Change Object (DCO) volume. This process takes less time than the traditional method of copying the entire original volume to the returning mirror. In the second step, the **Snap Shot** operation is performed to detach the snapshot mirrors again, creating a new point-in-time copy of the database. If you are creating multiple snapshots, the SFW GUI **Snap Shot** command or the `vxsnap` CLI command must be used to snapshot all the database and log volumes simultaneously. This step is done without taking the database offline.

The **Snap Back** and `vxsnap` commands can be called from either the **bpend\_notify.bat** file in Veritas NetBackup or from a batch file in a pre/post command to run at the completion of a Veritas Backup Exec for Windows Servers backup job. Additionally, a script could be written and used with the

Windows Task Scheduler to automatically update the snapshot or set of snapshots on a regular basis.

## Recovering a database

In the event a database needs to be recovered, you can use the snapshot or set of snapshots to restore the data.

---

**Caution:** Data corruption can occur if the FlashSnap utility does not have exclusive access to the volumes accessed in the **Snap Back** command. Before running the **Snap Back** command when using the snapshot data as the source, close any Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.

---

Storage Foundation 5.0 for Windows provides recovery support for Microsoft Exchange storage groups or individual databases within an Exchange storage group. Through SFW's `vxsnap restore` command or the VSS Restore wizard, the VSS hot snapshots can be used for a point-in-time recovery of the storage group or a roll-forward recovery to the point of failure of either the storage group or an individual database within it.

Refer to the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft Exchange* for detailed procedures on using FlashSnap with Microsoft Exchange Server 2003 to perform hot snapshots and to implement recovery procedures.

For Microsoft SQL, you can use the snapshot volumes in a snapshot set to restore a corrupt database. You can restore a database to a specified point in time, the point of failure, or the point in time that the snapshot set was created (or last refreshed).

Refer to the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft SQL* for detailed procedures on Quick Recovery in a Microsoft SQL environment.

## Other applications for point-in-time snapshots

This section describes several of the possible applications for using FlashSnap’s snapshots for off-host processing. Topics include:

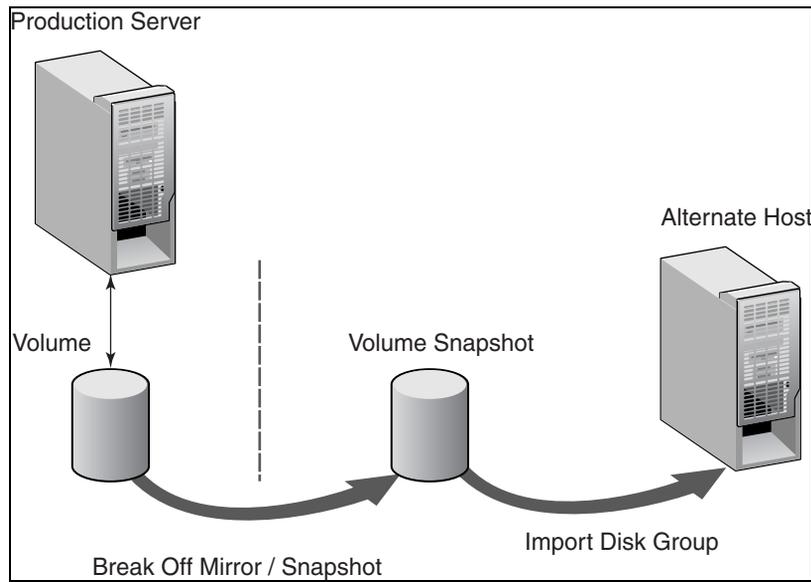
- “Off-host backups” on page 50
- “Reporting and analysis” on page 51
- “Application testing and training” on page 52

### Off-host backups

The more frequent your backups, the less data lost or, in the case of a database with a transaction log, the faster your recovery. Incremental backups reduce the backup time but increase recovery time. For organizations with little or no backup window, off-host backups offer a good solution, particularly as the amount of data to be managed grows.

Because backups take place on another host, the backup window is of less concern, and you can make full backups each day. This speeds recovery time in the event a problem does occur.

**Figure 4-1** Mirror break-off and import of the snapshot to the alternate host



FlashSnap simplifies the process of making snapshot volumes available for off-host processing with the Disk Group Split and Join feature. Using this feature,

administrators can split one or more volume snapshots into another disk group, then “deport” the disk group. The alternate host, running Storage Foundation for Windows, can then import that disk group and its volumes for off-host processing.

When the off-host processing is complete, you can rejoin the snapshot volume and its disk group in a similar manner, deporting it from the secondary host, importing it to the primary host, and rejoining the original disk group.

FlashSnap snapshots can be backed up with Veritas NetBackup or Backup Exec.

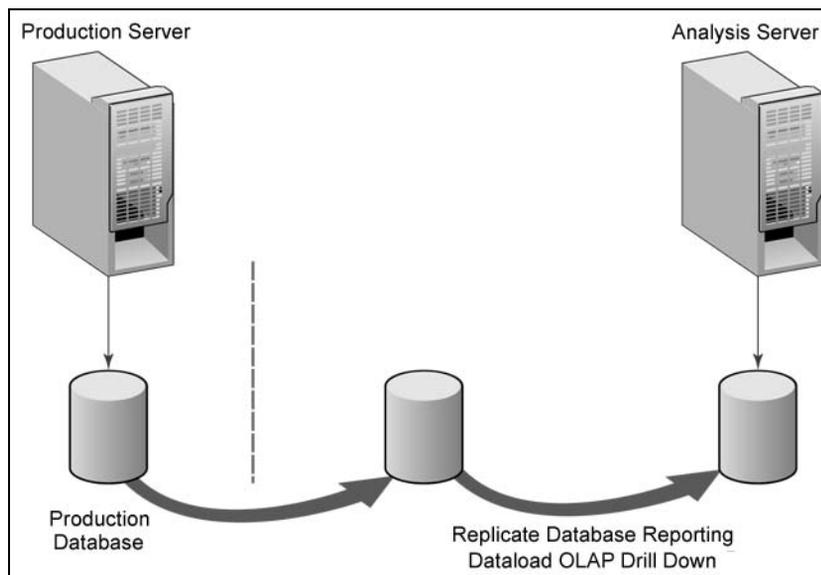
## Reporting and analysis

Decision support and business intelligence are data-intensive activities that are critical to many organizations. Analysts and others frequently need to access up-to-date or even real-time data in their analysis. Retail organizations, for example, want to spot sales trends as they occur, and typically require at least daily updates. Financial institutions likewise must keep a close eye on current transactions to spot trends or potential problems quickly.

Unfortunately, reporting and analysis data needs typically conflict with the performance requirements of transactional database applications. Reporting and analysis activities generally implement a few selected statements that scan a large number of records, and may include complex processing. This will have an impact on the many simpler write and update activities characteristic of a transactional system. For this reason, among others, many companies load data from operational systems into data warehouses specifically designed and tuned for analytic queries. But even the process of creating the data loads can have a performance impact on your operational systems, causing most organizations to schedule these Extract, Transform, and Load (ETL) processes during off-hours, such as in the middle of the night.

You can solve this problem by creating point-in-time snapshots of the production systems to be used for reporting and analysis purposes. You can either run reports directly against the snapshot volumes or use the snapshots to extract data for a data load to the warehouse.

Figure 4-2 Extract, transform, and load (ETL) process



Because taking the snapshot itself has a very brief, limited impact on the production system, you can generate fresh data for analysis on a regular basis. You can even create a replica of the production database on a secondary system to be accessible for “drill-down” analysis from OLAP applications. Again, in the off-host scenario, the analysis has no impact on the production system.

## Application testing and training

Software testing and training are other valuable applications for FlashSnap point-in-time copies. These are needs that cannot be addressed by simple data replication, because you need to be able to update and modify the copy of the data used for testing. FlashSnap addresses these needs easily.

By taking a snapshot and loading it on a host used for testing or development, you can provide developers and QA staff with the most realistic test data possible. By actually using a point-in-time copy of the production data, you can anticipate the behavior of the application in the production setting. You also save the time of creating and maintaining test data sets. This data can also be used for training purposes.

# Quick Recovery example

This chapter provides a Quick Recovery example using an Oracle database.

The chapter's topics are:

- [“Example of Quick Recovery of an Oracle database”](#) on page 54
- [“More on FlashSnap: Tips and references”](#) on page 58

## Example of Quick Recovery of an Oracle database

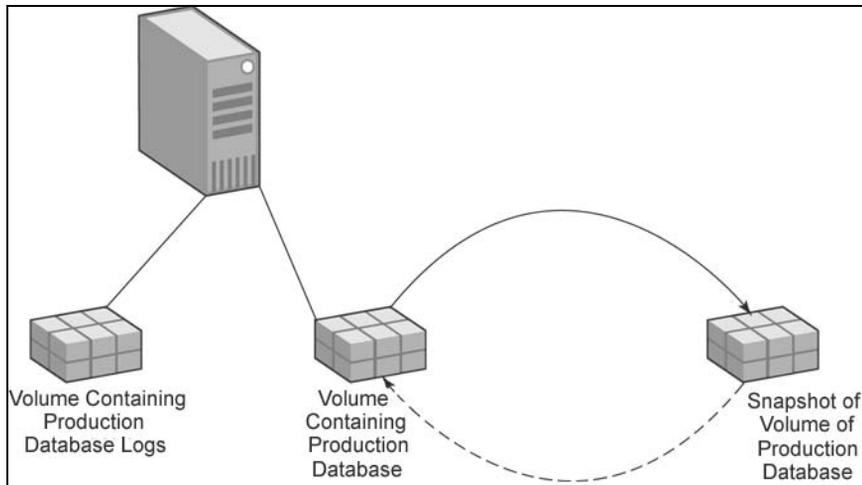
This example demonstrates how SFW's split-mirror snapshot can be used to recover an Oracle database after its data has become corrupted. The advantage of using the snapshot process is that it is much faster than recovering the database from tape backup. The process assumes that these split-mirror snapshots would take place on a regular schedule following the regular backup of the database.

This example does not require Microsoft Volume Shadow Copy Service (VSS), so it can be used with Windows 2000 or Windows Server 2003.

### Create split-mirror snapshot of database

The illustration below shows the snapshot step. The arrow pointing back to the original volume indicates that the snapshot volume can be rejoined to the original volume, updated, and ready to create a refreshed snapshot.

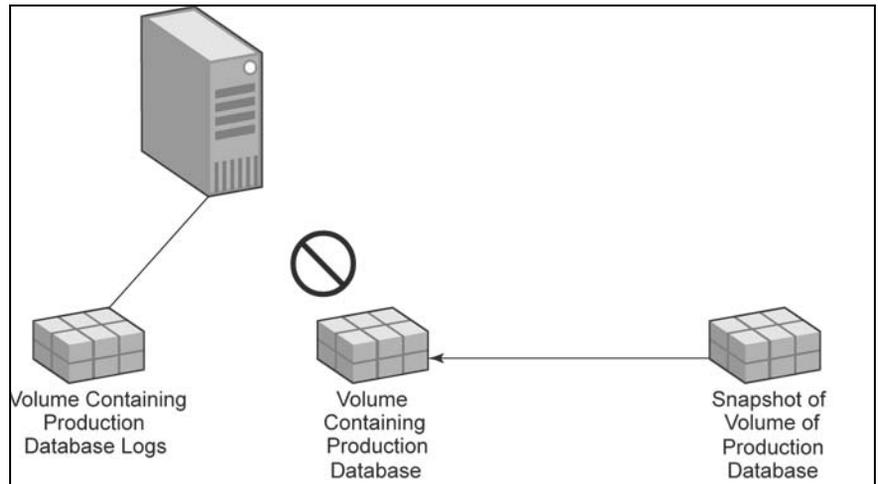
**Figure 5-1** Creating a backup of the database with a snapshot



## Recover database using split-mirror snapshot and database logs

The following illustration shows the situation where there has been a database failure. The snapshot volume is located on the right. The arrow pointing back to the production database volume represents the recovery of the database using the snapshot and applying the logs to bring the database to the level just before the failure occurred.

**Figure 5-2** Database recovery from a snapshot



### Overview of tasks

The main tasks for this example are:

- Snapshot the Oracle datafile volume.
- Resume normal processing with the Oracle datafile.
- Simulate Oracle datafile corruption.
- Recover the Oracle datafile.

### Specific steps

#### Prerequisites

This example assumes:

- Experience in Oracle database backup and recovery.
- Experience with the SFW FlashSnap procedures.

## Setup

- The Oracle database must be running on a single server.
- The Oracle database must be in ARCHIVELOG mode.
- The volume that contains the datafile for the Oracle database must meet the following requirements:
  - The volume must not be the system/boot volume.
  - The volume must be a SFW volume.

## To snapshot the Oracle datafile volume

- 1 Open the Oracle database and verify that the tablespace you want to work with is running normally.
- 2 In SFW, prepare the volume that contains the datafile of the tablespace.  
The CLI command is:

```
vxassist -g<DynamicDiskGroupName> snapstart  
<DriveLetter>
```
- 3 In Oracle, ALTER the tablespace with the BEGIN BACKUP option to prepare the database logs for backup creation mode.
- 4 In SFW, use **Snap Shot** to make a snapshot of the Oracle datafile volume.  
The CLI command is:

```
vxassist -g<DynamicDiskGroupName> snapshot  
<DriveLetter>
```
- 5 In Oracle, ALTER the tablespace with the END BACKUP option to set the database logs to normal mode.
- 6 In Oracle, archive the current database log to keep it at the same level as the snapshot.

This completes the process of implementing the snapshot of the database and saving it for later use. This process offers the most benefit if it is done on a periodic basis.

## Resuming normal processing with the Oracle datafile

- ◆ In Oracle, update the tables in the tablespace to create database log activity.

## To simulate Oracle datafile corruption

- 1 In Oracle Enterprise Manager, offline the tablespace.
- 2 Use Windows Explorer to locate and open the datafile volume.
- 3 Delete the datafile.

**To recover the Oracle datafile**

- 1 In Oracle, take the datafile offline.
- 2 Use **Snap Back** in the VEA GUI to reattach the snapshot volume. Use the **Resynchronize using the snapshot** option.
- 3 In Oracle, use RECOVER TABLESPACE to apply the database logs to bring the replica to the level just before the datafile corruption occurred.
- 4 In Oracle, bring the datafile online.
- 5 Verify that the tablespace in the datafile has been recovered.

---

**Note:** This example uses a single snapshot of the datafile of the tablespace of an Oracle database. It could also be done by using multiple, simultaneous snapshots that include both the data file and the log.

---

## More on FlashSnap: Tips and references

The following FlashSnap tips may be helpful:

- Use related disk group names.  
For example, when doing off-host processing from the GUI, use “database” for the original disk group name and “database\_snap” for the snapshot disk group name.  
A disk group name can be a maximum of 18 characters long.
- When using FlashSnap with a database application, store all database files and related transaction logs on disks contained within a single dynamic disk group.
- For easy identification, the volumes within a disk group should begin with the name of the disk group. For example:  
    DiskGroup1\_VolumeName1  
    DiskGroup1\_VolumeName2  
    DiskGroup1\_VolumeName3
- For more information on FlashSnap, see the *Veritas Storage Foundation Administrator’s Guide*.
- For more information on Quick Recovery, see the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft Exchange* and the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft SQL*.

# High Availability

This section focuses on SFW and SFW HA's local clustering support features. It includes step-by-step configuration examples with VCS.

This section has the following chapters:

- [Chapter 6, “High availability: Overview” on page 61](#)
- [Chapter 7, “Deploying SFW HA for high availability: New installation” on page 63](#)
- [Chapter 8, “Adding DMP to a clustering configuration” on page 169](#)



# High availability: Overview

This overview of high availability has the following topics:

- [“About high availability”](#) on page 61
- [“About clusters”](#) on page 61

## About high availability

“High availability” maintains continued functioning of applications in the event of computer failure, where data and applications are continuously available using redundant software and hardware. “High availability” can refer to any software or hardware that provides fault tolerance, but generally it has become associated with clustering. This section will focus on local clustering configurations that use Veritas Cluster Server (VCS) with Veritas Storage Foundation for Windows.

## About clusters

A cluster is a group of independent computers working together as a single system to ensure that mission-critical applications and resources are highly available. The cluster is managed as a single system, shares a common namespace, and is specifically designed to tolerate component failures and to support the addition or removal of components in a way that is transparent to users.

Keeping data and applications functioning 24 hours day and seven days a week is the necessary for critical applications today. Clustered systems have several advantages, including fault tolerance, high availability, scalability, simplified management, and support for rolling upgrades.



# Deploying SFW HA for high availability: New installation

This chapter covers the following topics:

- [About the high availability solution](#)
- [Tasks for a new high availability \(HA\) installation—additional applications](#)
- [Before you begin](#)
- [Installing SFW HA](#)
- [Configuring disk groups and volumes](#)
- [Configuring the cluster](#)
- [Installing and configuring the application or server role](#)
- [Configuring the service group](#)
- [Verifying the cluster configuration](#)
- [Possible tasks after completing the configuration](#)

## About the high availability solution

This chapter provides the steps for setting up a High Availability (HA) solution, using SFW HA in a new installation. The chapter describes the process for any generic application or server role and specifically for File Share, PrintShare, IIS and MSVirtual Machines.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA solutions for the example applications or server roles.

See [“Using the Solutions Configuration Center”](#) on page 23.

For examples of the SFW HA solution with Microsoft Exchange or Microsoft SQL Server, see the other Solutions Guides included with this release: *Veritas Storage Foundation and High Availability Solutions High Availability and Disaster Recovery Solutions Guide for Microsoft Exchange* and *Veritas Storage Foundation and High Availability Solutions High Availability and Disaster Recovery Solutions Guide for Microsoft SQL*.

## Tasks for a new high availability (HA) installation—additional applications

This chapter provides information on how to install and configure the high availability and application components.

- active/passive      One application instance per node with one to one failover capabilities. 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.
- active/active      Multiple application instances per cluster node.  
For example, in a two-node cluster with two application instances, a different instance is online on each of the two servers. If a failure occurs, the instance on the failing node is brought online on the other server, resulting in two instances online on one server.

[Table 7-1](#) outlines the high-level objectives for implementing the configuration and the tasks for each objective:

**Table 7-1**      Task List

Objectives	Tasks
<a href="#">“Before you begin”</a> on page 65	<ul style="list-style-type: none"><li>■ Verify hardware and software prerequisites</li><li>■ Review the requirements</li><li>■ Review the configuration</li><li>■ Configure the storage hardware and network</li></ul>

**Table 7-1** Task List (Continued)

Objectives	Tasks
“Installing SFW HA” on page 71	<ul style="list-style-type: none"> <li>■ Verify the driver signing options for Windows 2003 systems</li> <li>■ Install SFW HA</li> <li>■ Restore driver signing options for Windows 2003 systems</li> </ul>
“Configuring disk groups and volumes” on page 77	<ul style="list-style-type: none"> <li>■ Planning your storage layout</li> <li>■ Create disk groups</li> <li>■ Create volumes</li> <li>■ Managing disk groups and volumes</li> </ul>
“Configuring the cluster” on page 85	<ul style="list-style-type: none"> <li>■ Use the VCS Configuration wizard to set up the cluster</li> </ul>
“Installing and configuring the application or server role” on page 101	<ul style="list-style-type: none"> <li>■ As necessary, Install the application program files on the local drive of the first node</li> <li>■ Install files relating to the data and logs on the shared storage</li> <li>■ Deport the disk groups on the first node and import them on the second node</li> <li>■ Install the application on the second node</li> </ul>
“Configuring the service group” on page 104	<ul style="list-style-type: none"> <li>■ Use the applicable wizard to create and configure the VCS service group or groups</li> <li>■ Bring the service group online</li> </ul>
“Verifying the cluster configuration” on page 151	<ul style="list-style-type: none"> <li>■ Switch the service group to the second node</li> <li>■ Shut down an active cluster node</li> </ul>
“Possible tasks after completing the configuration” on page 152	<ul style="list-style-type: none"> <li>■ Configure the Cluster Management Console connection</li> <li>■ Modify the cluster configuration</li> <li>■ Modify the application or server role service group</li> </ul>

## Before you begin

Review these product installation requirements for your systems before installation. Minimum requirements and Veritas recommended requirements may vary.

## Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Installation on a non-system drive requires space on both the system drive and the non-system drive.

[Table 7-2](#) estimates disk space requirements for SFW HA.

**Table 7-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
SFW HA + all options + client components	1675 MB	Non-system space: 1675 MB System space: 345 MB
SFW HA + all options	1230 MB	Non-system space: 1230 MB System space: 285 MB
Client components	630 MB	Non-system space: 630 MB System space: 115 MB

## Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.0 Hardware Compatibility List to confirm supported hardware at: <http://entsupport.symantec.com>

For a Disaster Recovery configuration select the Global Clustering Option and optionally select Veritas Volume Replicator.

### Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA)
- Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
  - or
  - Windows Server 2003 Web Edition (limited to file share support for SFW HA), Windows Server 2003 Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 for all editions)
  - or
  - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
  - or

Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP1 required for all editions)

or

Windows Server 2003 for Intel Xeon (EM64T) or AMD Opteron: Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition

or

Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

## System requirements

Systems must meet the following requirements:

- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See “[Best practices](#)” on page 69.
- 1 GB of RAM for each system.
- All servers must run the same operating system, service pack level, and system architecture.

## Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows 2000 or Windows Server 2003 domain.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.
- Static IP addresses for the following purposes:
  - One static IP address available per site for each application virtual server
  - One IP address for each physical node in the cluster
  - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.

- For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.  
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.  
See the *Veritas Cluster Server Bundled Agents Reference Guide*.

## Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

## Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.

## Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).  
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

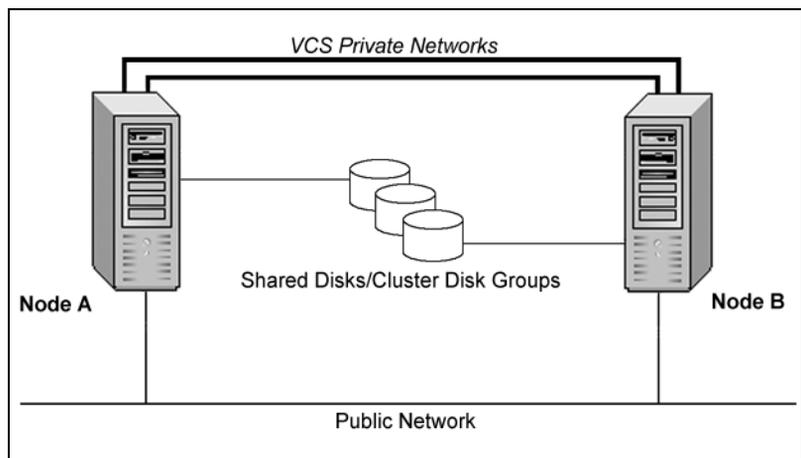
## Reviewing the configuration

This example configuration is one of the most common configurations for a cluster. It is a new installation with two servers and one storage array, in an active/passive configuration where the active node of the cluster hosts the virtual server and 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 example describes a generic database application.

The example configuration does not include dynamic multi-pathing.

See “[Adding DMP to a clustering configuration](#)” on page 169.

**Figure 7-1** SFW HA Active/Passive configuration with two servers



## 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 VCS 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.
  - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk.  
Use Windows Disk Management (**Start > Control Panel > Administrative Tools > Computer Management**) or Veritas Enterprise Administrator (VEA) on each system to verify 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 Right-click on 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 and Dial-up 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 VCS private network.
- 5 From the status window, click **Properties**.

- 6 In 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**.

## Installing SFW HA

The product installer enables you to install the software for Veritas Storage Foundation HA 5.0 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. For a disaster recovery configuration, select the option to install GCO. If you plan to use VVR for replication, you must also select the option to install VVR.

## Setting Windows driver signing options

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

[Table 7-3](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

**Table 7-3** 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.

**Table 7-3** Installation behavior with unsigned drivers (Continued)

<b>Driver Signing Setting</b>	<b>Installation behavior on the local system</b>	<b>Installation behavior on remote systems</b>
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.

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

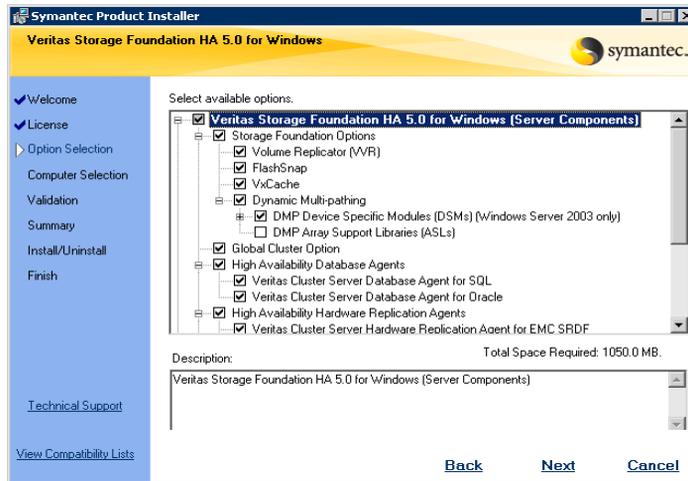
**To install the product**

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

3 Click **Storage Foundation HA 5.0 for Windows**.

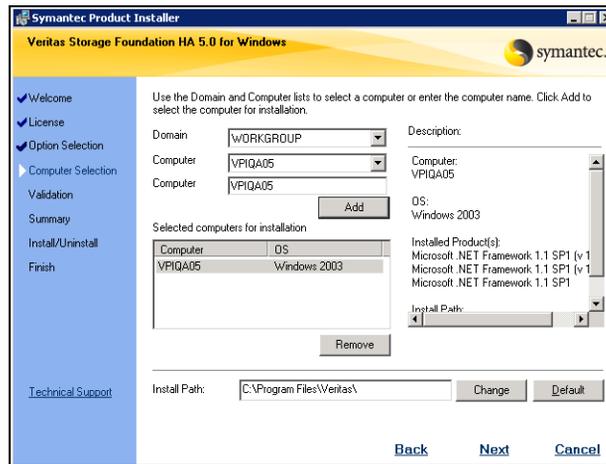


- 4 Do one of the following:
  - Click **Complete/Custom** to begin installation.
  - Click the **Administrative Console** link to install only the client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you 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**.  
 If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 8 Repeat for additional license keys. Click **Next**
  - 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 Select the appropriate SFW product options and click **Next**.



- |                           |   |
|---------------------------|---|
| Client                    | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console.<br>Required to install the Solutions Configuration Center which provides information and wizards to assist configuration |
| Global Cluster Option     | Required for a disaster recovery configuration only.  |
| Veritas Volume Replicator | If you plan to use VVR for replication, you must also select the option to install VVR.   |

**10** 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

**11** 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.

- 12 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 list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
  - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
  - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs 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.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.  
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

## Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

### 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 disk groups and volumes

Use Veritas Storage Foundation for Windows to create cluster disk groups and dynamic volumes for the application on the shared storage. A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different RAID layouts. Configuring disk groups and volumes involves the following tasks:

- [“Planning disk groups and volumes”](#) on page 77
- [“Creating dynamic cluster disk groups”](#) on page 79
- [“Creating dynamic volumes”](#) on page 81

## Planning disk groups and volumes

The requirements for disk groups and volumes depend on the type of application or server role. Review the requirements and best practices for your application or server role:

- [Planning your File Share storage](#)
- [Planning your Print Share storage](#)
- [Planning your IIS storage](#)
- [Planning your Microsoft Virtual Machine storage](#)
- [Planning your storage for additional applications](#)

### Planning your File Share storage

Considerations for planning the File Share storage:

- Make sure that the disk group and volumes for the file server shared directory are configured on shared storage.
- When configuring a new set up, first create the disk groups and volumes on the shared storage and then create the directory structure for the file shares on the shared storage.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage using the practices recommended by Microsoft.

## Planning your Print Share storage

Considerations for planning the Print Share storage:

- Make sure you create separate volumes on shared storage for the print spooler and for the RegRep volume that contains the list of registry keys that must be replicated between the cluster systems.

---

**Note:** Symantec recommends that you place the RegRep and spooler data in separate directories on separate volumes on the shared storage.

---

## Planning your IIS storage

Considerations for planning the IIS storage:

- Make sure that the disk groups and volumes which will host the directory and files for the web sites are on the shared storage.
- For a new IIS installation, make sure that the directory for the web sites is created on volumes on the shared storage.
- For existing web sites, stop the sites and then move the website content to volumes on the shared storage. You must also reconfigure the home directory location for the web site in IIS and then restart the web site again.

## Planning your Microsoft Virtual Machine storage

Make sure the volumes that contain the shared virtual disk files for the virtual machines are located on the shared storage.

## Planning your storage for additional applications

The information provided in this section is generic to any application. Make sure you create the appropriate disk groups and volumes to hold the application data. If your application requires replication of registry keys between the cluster systems, then Symantec recommends that you create a dedicated RegRep

volume so that its MountV dependency is not linked with any other application-specific resources in the group.

Decide how you want to organize the disk groups and the number and type of volumes you want to create. Some considerations are:

- The number of disk groups that are needed  
The number of disk groups depends on your application and the planned organization of the data. VCS requires that the application program files be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage. Typically, a main organizational unit in your application would be contained in a single disk group.
- The type of volumes you want to create
  - Mirrored and RAID-5 volumes provide fault tolerance for critical data.
  - Striped volumes add performance capability.
  - Volumes that are both mirrored and striped offer both performance and fault tolerance.

---

**Note:** If you plan to use replication software, such as VVR, do not use software RAID-5 volumes. This does not apply to hardware RAID-5.

---

Recommendations:

- Use mirrored volumes for logs.
- Use striped or mirrored striped volumes for data.
- The implications of backup and restore operations for the disk group setup.
- The sizes of databases and logs, which depend on the traffic load.

## Creating dynamic cluster disk groups

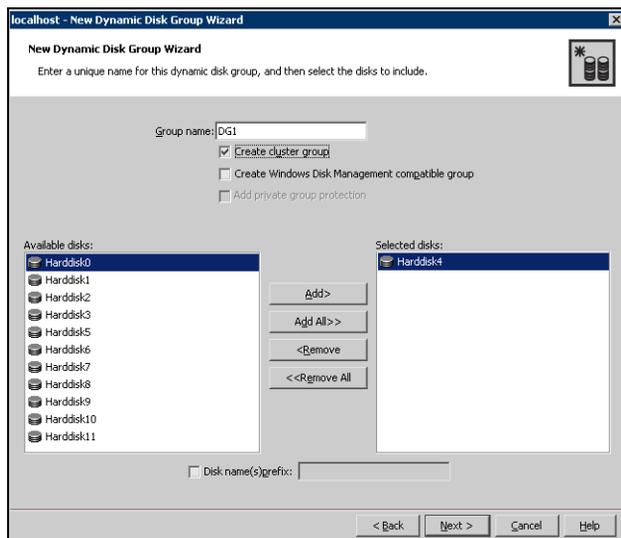
Follow the steps in this section to create one or more disk groups for your application.

### To create a dynamic (cluster) disk group

- 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:



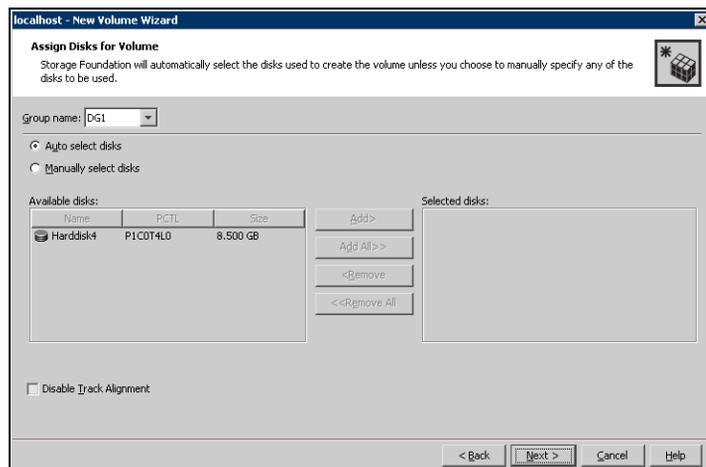
- Enter the disk group name (for example, DG1).
  - Click the checkbox for **Create cluster group**.
  - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter 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.
  - 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

Once the disk groups are created, make the disks within them usable by creating the dynamic volumes that will store data.

### 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.
- 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.

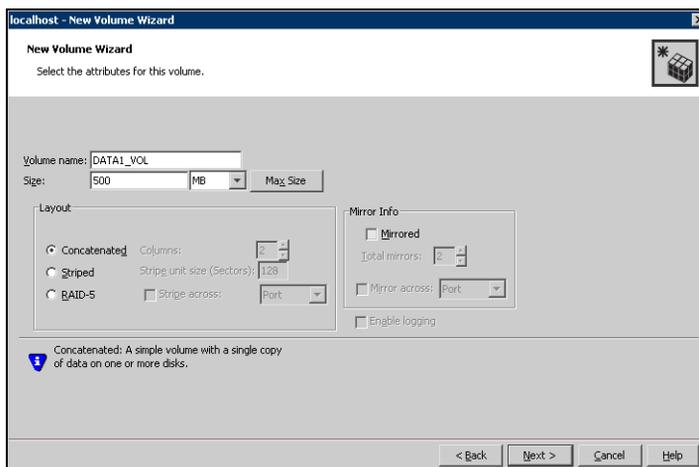


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

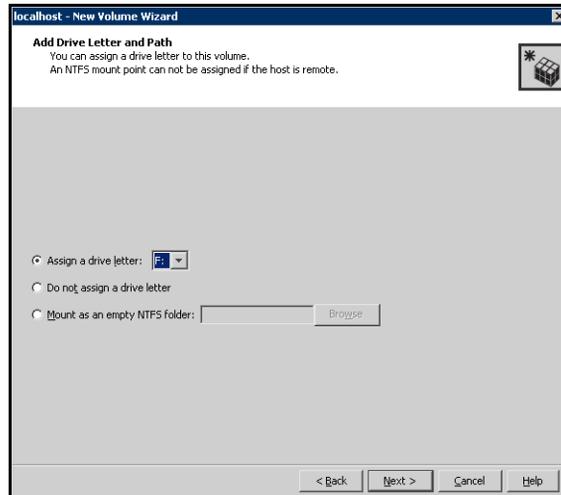
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling **Track Alignment** means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

- 8 Click **Next**.
- 9 Specify the volume attributes.



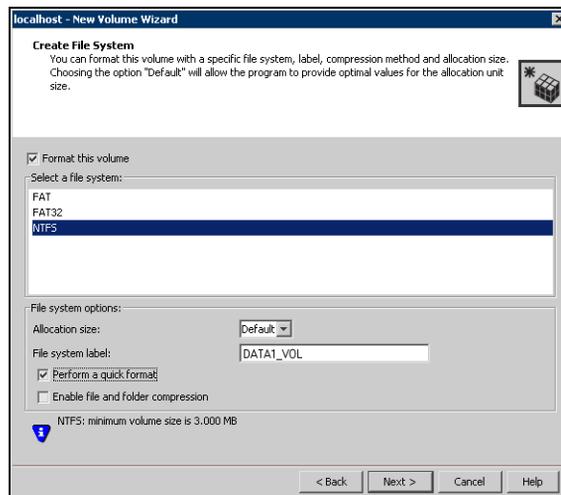
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
  - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
  - If you are creating a striped volume, the **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 groups and volumes

This section includes the following procedures:

- Importing a disk group and mounting a shared volume
- Unmounting a volume and deporting a disk group

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**.

## Configuring the cluster

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

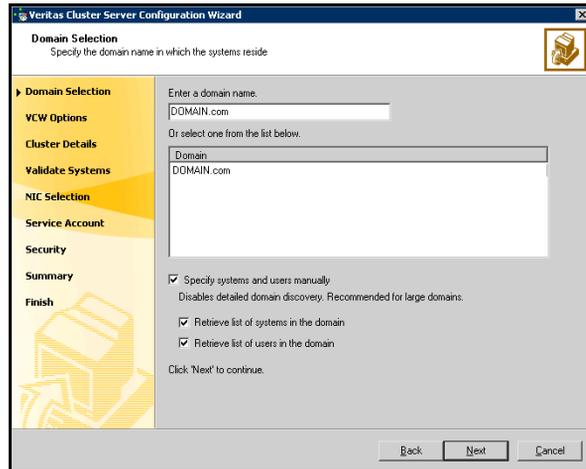
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
  - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
  - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
  - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

#### To configure a VCS cluster

- 1 Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

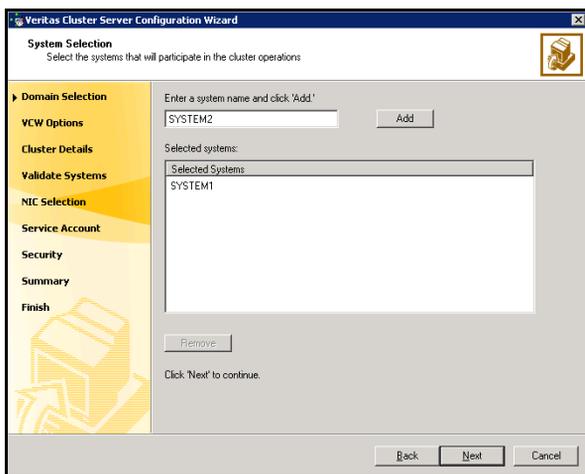
Proceed to [step 7](#) on page 89.

To specify systems and user names manually (recommended for large domains):

- Check the **Specify systems and users manually** check box.  
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

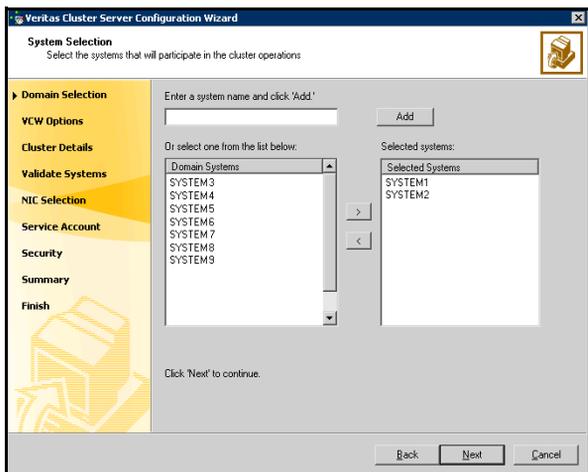
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 88. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



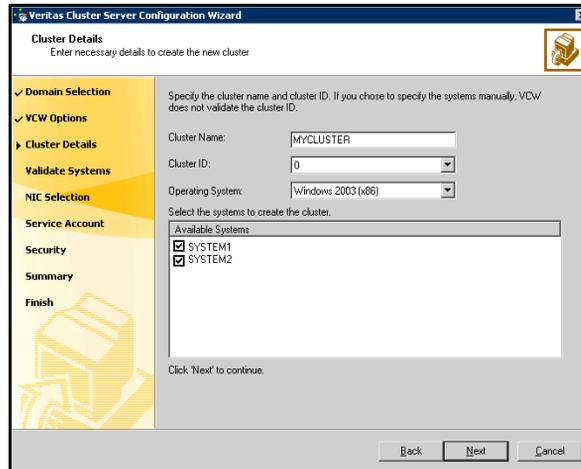
Proceed to [step 7](#) on page 89.

- 6 On the System Selection panel, specify the systems to form a cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.

- 7 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 8 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



**Cluster Name** Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

**Cluster ID** Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

---

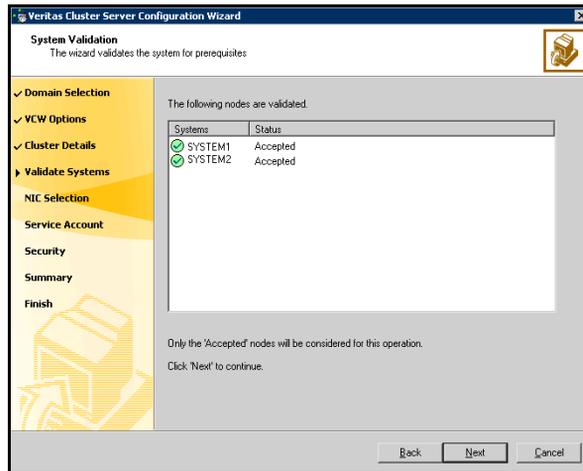
**Warning:** If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

---

**Operating System** From the drop-down list, select the operating system that the systems are running.

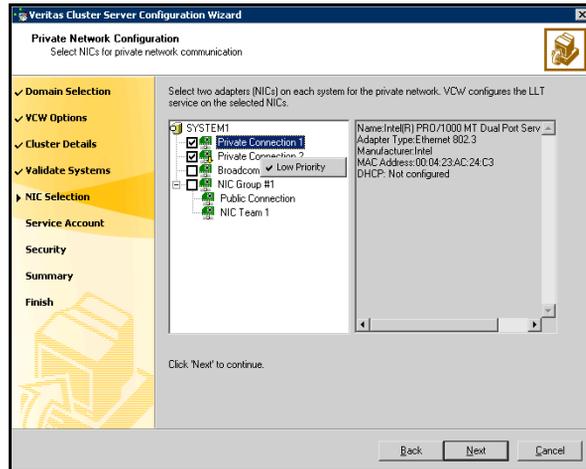
**Available Systems** Select the systems that will be part of the cluster. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.  
If you chose to configure a private link heartbeat in [step 8](#) on page 89, proceed to the next step. Otherwise, proceed to [step 11](#) on page 91.

- 10 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

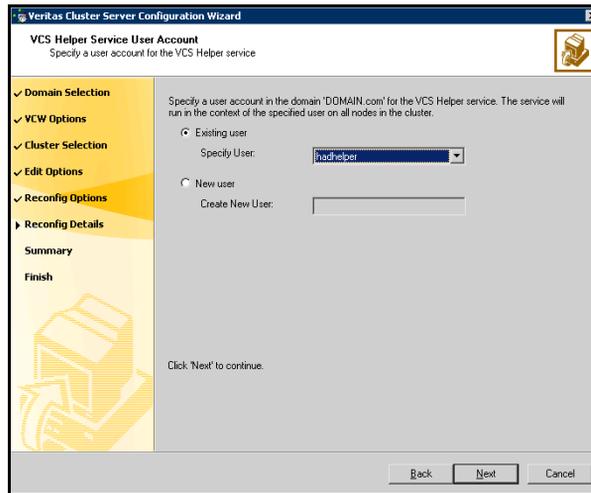


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of at least one NIC for that system. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.

If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

- 11 On the VCS Helper Service User Account panel, specify the name of a domain user context for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper

Service user context to access the network. Do not use the Administrator account for the VCS Helper service.



- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list,
  - If you chose not to retrieve the list of users in [step 4](#) on page 87, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

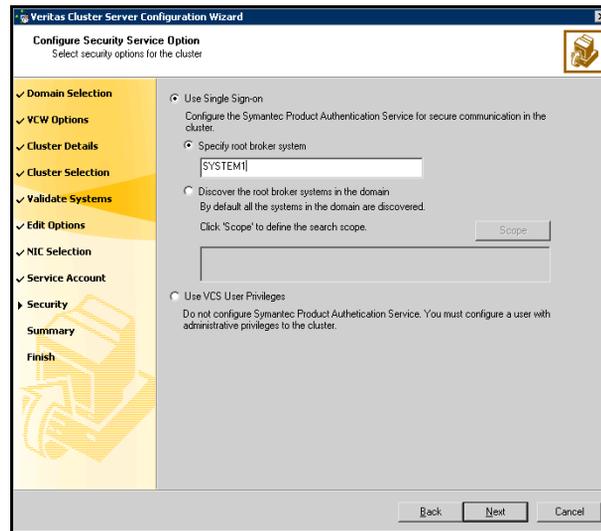
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 12 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

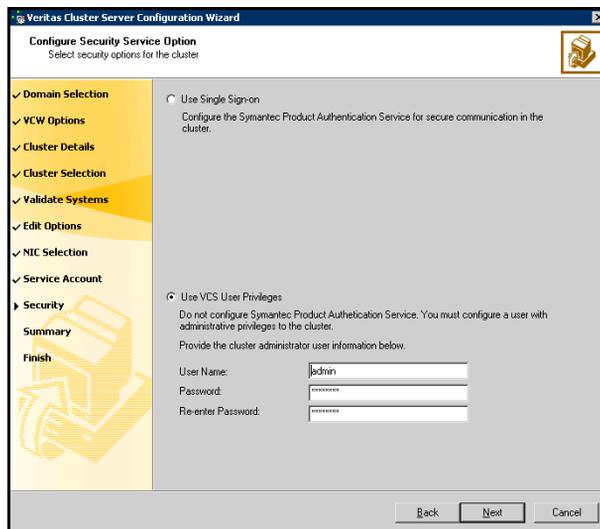
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition. For example, to search for systems managed by Administrator, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.
- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use VCS user privilege:



- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

The default user name for the VCS administrator is admin and the default password is password. Both are case-sensitive. Use this account

to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

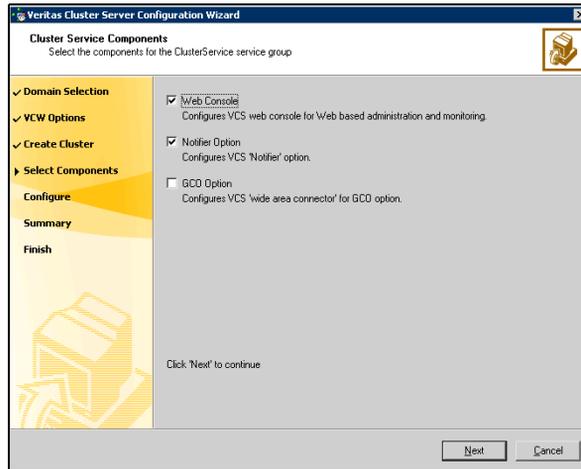
- Click **Next**.

- 13 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 14 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters. To configure the ClusterService group later, click **Finish**. At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

- 15 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.



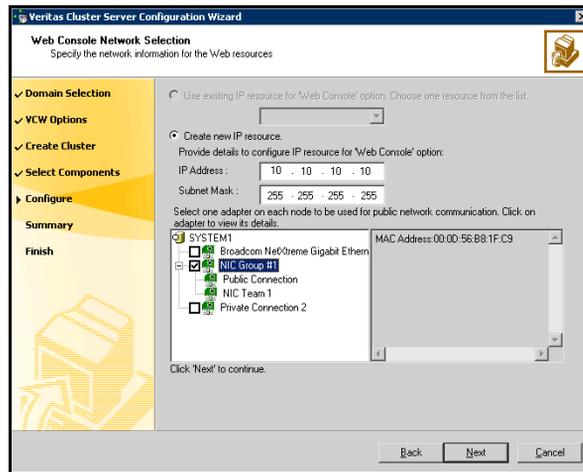
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 97.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 98.

## Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

### To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



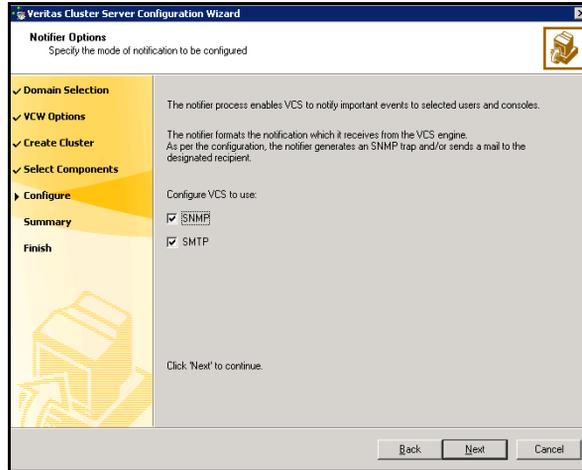
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
  - If you choose to configure a new IP address, type the IP address and associated subnet mask.
  - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
  - 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 98. Otherwise, click **Finish** to exit the wizard.

## Configuring notification

This section describes steps to configure notification.

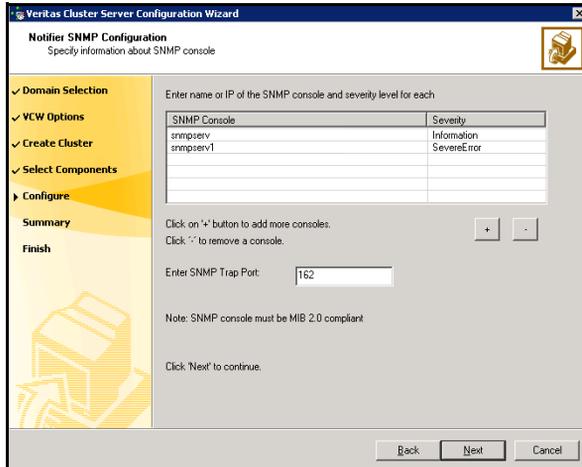
### To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.

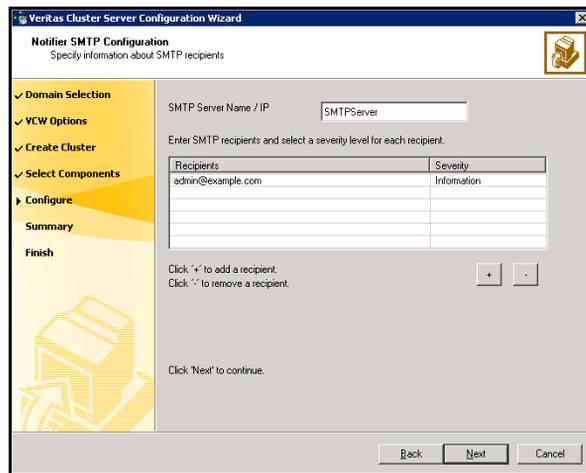


You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

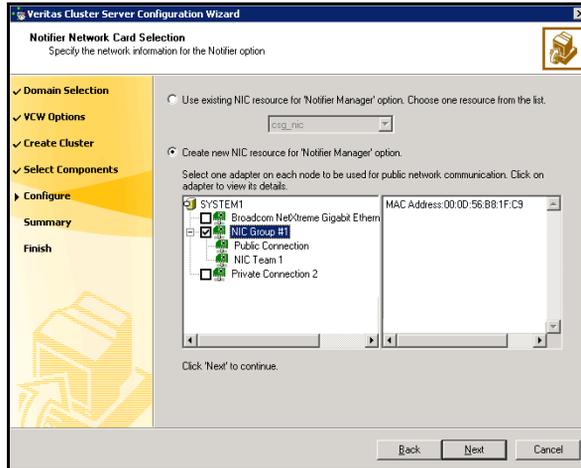


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
  - Click the corresponding field in the Severity column and select a severity level for the console.
  - Click '+' to add a field; click '-' to remove a field.
  - Enter an SNMP trap port. The default value is "162".
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
  - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
  - 6 Click **Configure**.
  - 7 Click **Finish** to exit the wizard.

# Installing and configuring the application or server role

This section provides considerations for installing and configuring your application or server role. See the following topics:

- [Configuring a File Share server role](#)
- [Configuring a Print Share server role](#)
- [Installing and configuring the IIS application](#)
- [Installing and configuring Microsoft Virtual Server](#)
- [Installing additional applications](#)

## Configuring a File Share server role

Points to note when configuring a File Share:

- Make sure that the disk group and volumes for the file server shared directory are configured on shared storage.
- When configuring a new set up, first create the disk groups and volumes on the shared storage and then create the directory structure for the file shares on the shared storage.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage using the practices recommended by Microsoft.
- The FileShare agent is installed automatically with SFW HA.

## Configuring a Print Share server role

Points to note when configuring a Print Share:

- Make sure the printer is connected to the network and is configured with an IP address.
- Install software drivers for the network printer on all systems in the cluster.

### To add a print driver

- 1 Open the **Printers** Control Panel.
- 2 Click **File > Server Properties**.
- 3 In the Print Server Properties dialog box, click the **Drivers** tab.
- 4 Click **Add**. This launches the Add Printer Driver wizard.

- 5 Follow the wizard instructions to add the printer driver on the system. You must add the driver on each system that will be part of the service group.

## Installing and configuring the IIS application

Points to note when installing IIS:

- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify that the sites to be monitored are on shared storage.
- Import the cluster disk groups and mount the volumes that contain the website data, on the first node.
- For a new IIS installation, while creating new web sites, create the site folder on the shared storage and place the site content in that folder.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- For existing web sites, stop the sites and then move the website content to volumes on the shared storage. You must also reconfigure the home directory location for the website in IIS and then restart the website again.
- Verify the port numbers assigned to IIS sites are not used by other applications or sites.
- Synchronize the IIS configuration on all nodes hosting the service group.

### To synchronize the IIS configuration on Windows 2003 systems

- 1 Synchronize the IIS configuration on all nodes that will host the IIS service group. Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system.

For example, the following command copies the IIS metabase to *target\_system*. You must enter a valid user name and password for the target system.

```
%systemroot%\System32> iiscnfg /copy /ts target_system /tu  
user_name /tp password
```

- 2 Stop and restart IIS Admin Service on all nodes.

### To synchronize the IIS configuration on Windows 2000 systems

Synchronize the IIS configuration on all nodes that will host the IIS service group by running the `iissync` utility, located at `%windir%\system32\inetsrv\iissync`.

For more information about the `iissync` utility, see the IIS documentation.

## Installing and configuring Microsoft Virtual Server

Points to note when installing MS Virtual Server:

- Verify Microsoft Virtual Server is installed and configured identically on all nodes hosting the service group.
- Install the operating system and the applications that you want to make highly available on the virtual machine.
- Install and configure Virtual Machine Additions *on each virtual machine* if you plan to enable detailed monitoring for the virtual machine resources.
- Verify the Microsoft Virtual Server configuration files reside locally on each node.
- Make sure the name of the virtual machine is unique in the cluster.

## Installing additional applications

Following are some very generic points for installing any application:

- Make sure that the disk groups and volumes are mounted on the node before installing the application.
- VCS requires the application program files to be installed on the same local drive on all nodes. For example, if you install the application program files on drive C of one node, installation of these same files on all other nodes must be on drive C.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- The data files and any associated files, such as log files, should be installed on the shared storage.

## Configuring the service group

The Solutions Configuration Center provides wizards to configure the service groups for the additional SFW HA applications or server roles. It also supports the Application Configuration Wizard which can be used to configure any other application for which application specific wizards have not been provided. Depending on the application that you have installed, complete the appropriate procedure to configure the service group:

- [Configuring the File Share service group](#)
- [Configuring the PrintShare service group](#)
- [Configuring the IIS service group](#)
- [Configuring the MSVirtual Machine service group](#)
- [Configuring the service group for any additional applications](#)

### Configuring the File Share service group

Configuring the File Share service group involves creating a FileShare service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

#### Prerequisites

- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Verify that the directories to be shared reside on shared drives.
- Mount the drives containing the shared directories from the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- Verify that Veritas Command Server is running on all systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
  - A unique virtual computer name to be assigned to the file share server. This is the name by which clients will access the server. The virtual name must not exceed 15 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
  - A unique virtual IP address to be assigned to the file share server.

- The list of directories to be shared.

The wizard enables you to add existing shares to the VCS configuration. However, you cannot add special shares (shares created by the operating system for administrative and system use). For example, you cannot add the shares ADMIN\$, print\$, IPC\$, and *DriveLetter\$* to the VCS configuration.

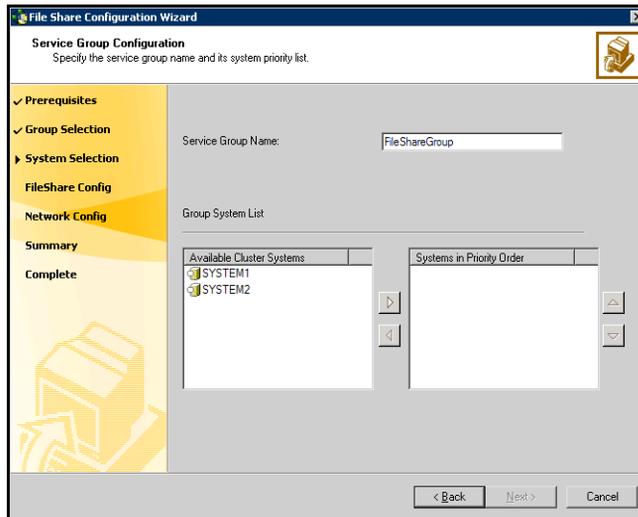
## Creating the FileShare service group

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide*. for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

### To configure a FileShare

- 1 Start the File Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.  
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > File Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.

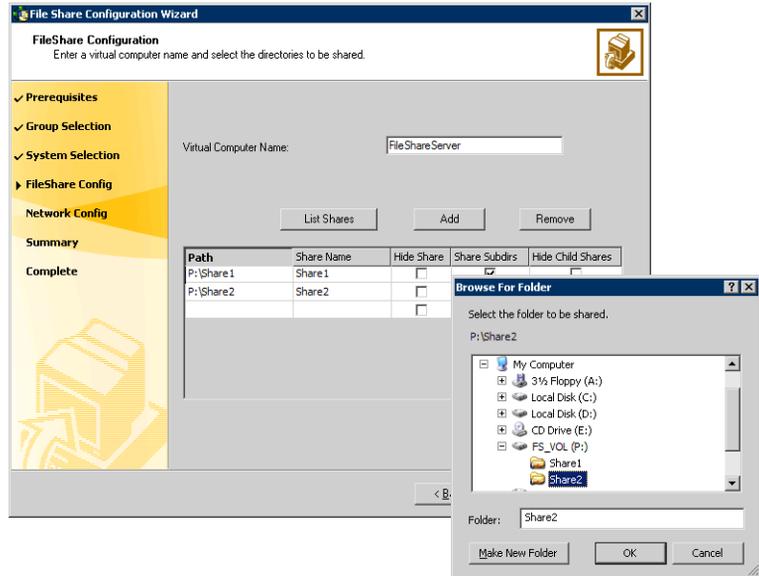
- 4 On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name	Type a name for the File Share service group.
Available Cluster Systems	Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list. To remove a system from the service group's system list, click the system from the Systems in Priority Order box and click the left arrow. To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the File Share Configuration panel, specify the configuration information for the FileShare resources to be created and then click **Next**.

The wizard begins validating your configuration. Various messages indicate the validation status.

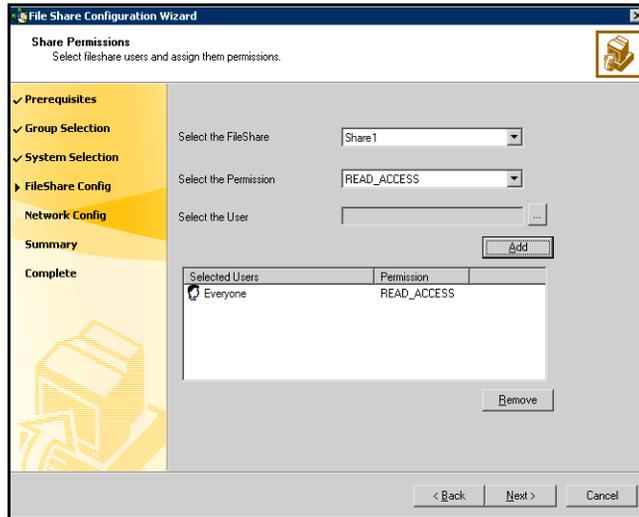


- Virtual Computer Name      Type a unique virtual computer name by which the server will be known to clients. The virtual name must not exceed 15 characters.
  
- List Shares                      Click **List Shares** to view the existing shares on the shared storage, then select a share and click **Add**.
  
- Add                                  Click **Add** to add a file share.
  
- Path                                  Type the path of the directories to be shared or click the field and then click the ellipse icon (...) to browse for folders. The selected directories must meet the following conditions:
  - The selected drive, the mount path, and the file path must not exist in the VCS configuration.
  - The directories to be shared must reside on shared, non-system drives.

The wizard validates the selected directory and displays an error message if the directory does not meet any of the conditions.

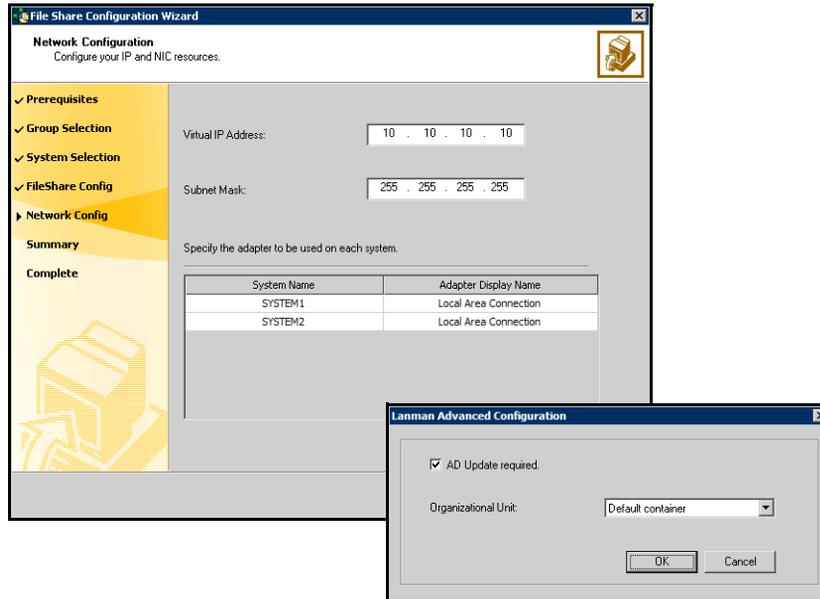
Share Name	If a selected directory is already shared, the Share Name column lists the names by which it is shared. You can select a listed share name to make an existing share highly available. You can also create a new share for the same directory by typing a new share name.
Hide Share	Check <b>Hide Share</b> check box to make the new share a hidden share.
Share Subdirs	Check the <b>Share Subdirs</b> check box to share the subdirectories.
Hide Child Shares	Check the <b>Hide Child Shares</b> check box to hide the shared subdirectories.
Remove	To remove a file share from the configuration, click to select the file share, and then click <b>Remove</b> .

- 6 On the Share Permissions panel, specify the users for the file shares, assign permissions to them and then click **Next**.



- |                       |  |
|-----------------------|--|
| Select the FileShare  | From the drop-down list, select the file share with which to associate user permissions, or select the default <b>All FileShares</b> to set the same permissions for all file shares.  |
| Select the Permission | From the drop-down list, select the permission to be associated with the user.   |
| Select the User       | Click ... (ellipsis button), select a user, and click <b>OK</b> .  |
| Add                   | Click <b>Add</b> to add the specified user to the Selected Users list. By default, all selected users are given READ_ACCESS permission.  |
| Selected Users        | Displays a list of selected users and the file share permissions. You can configure a maximum of 50 users for each file share. To configure more users, create a user group.<br>To change the file share permission associated with a user, click a user name in the Selected Users list and then select the desired permission from the Select the Permission drop-down list. |
| Remove                | To deny file share access to a user, click the user name in the Selected Users list and click <b>Remove</b> .  |

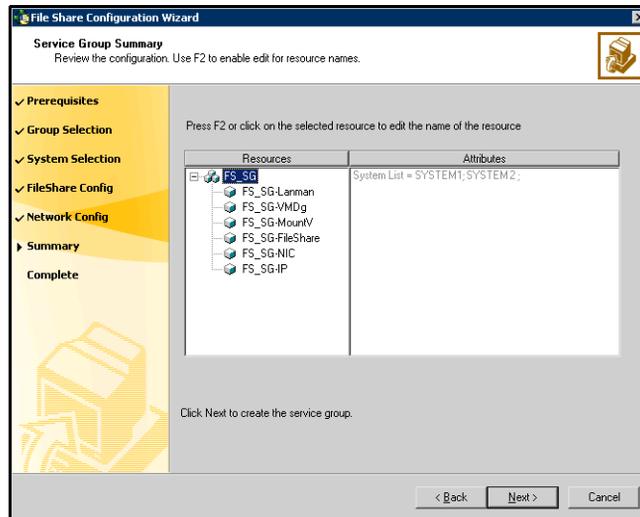
- 7 On the Network Configuration panel, specify information related to your network and then click **Next**.



- |                    |  |
|--------------------|--|
| Virtual IP Address | Type a unique virtual IP address for the virtual server.   |
| Subnet Mask        | Type the subnet to which the virtual server belongs.   |
| Advanced Settings  | Click <b>Advanced Settings...</b> to specify additional details for the Lanman resource.<br>On the Lanman Advanced Configuration dialog box, complete the following: <ol style="list-style-type: none"><li>1 Check <b>AD Update required</b> check box to enable the Lanman resource to update the Active Directory with the virtual name.</li><li>2 From the Organizational Unit drop-down list, select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."</li><li>3 Click <b>OK</b>.<br/>The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.</li></ol> |

**Adapter Display Name** Displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.

- On the Service Group Summary panel, review the service group configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



**Resources** Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required. To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

**Attributes** Displays the attributes and their configured values, for a resource selected in the Resources list.

- 9 In the completion dialog box, check **Bring the service group online** check box if you want to bring the service group online on the local system, and then click **Finish**.

## Configuring the PrintShare service group

Configuring the Print Share service group involves creating a PrintShare service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

### Prerequisites

- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Verify that VCS Command Server is running on all systems in the cluster.
- Verify that the network printer has an IP address assigned.
- Symantec recommends creating spooler and the replication directories on different disk partitions or volumes.
- Mount the drives with the spooler and the replication directories on the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- Verify that the software drivers for the network printers are installed on all systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
  - A unique virtual computer name to be assigned to the print share server.  
This is the name by which clients will access the server. The virtual name must not exceed 15 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
  - A unique virtual IP address to be assigned to the print share server.
  - The network printer's IP address.

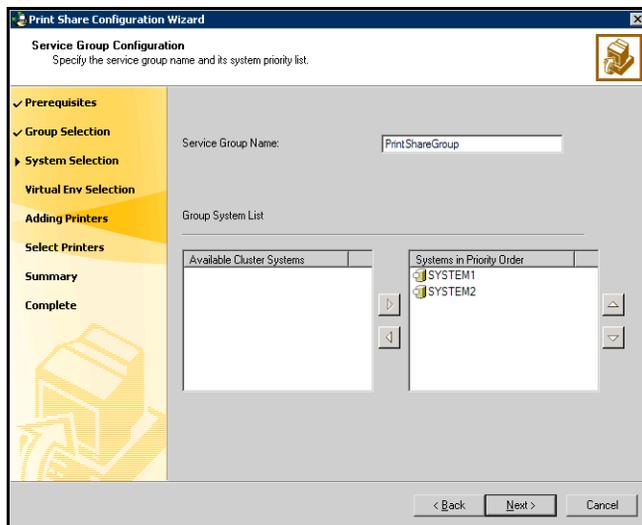
### Creating the PrintShare service group

To create a new Print Share service group perform the following tasks:

- Create a new service group with a PrintSpool resource and bring it online. This also involves configuring the Lanman resource on which the PrintSpool resource depends. See [To create a Print Share service group with a PrintSpool resource](#).
- Add a network printer to the virtual computer created by the Lanman resource. Create a new TCP/IP port for the printer. See [To add the network printer to the virtual computer](#).
- Configure a PrintShare resource in your service group and bring it online. See [To configure a PrintShare resource for the service group](#).

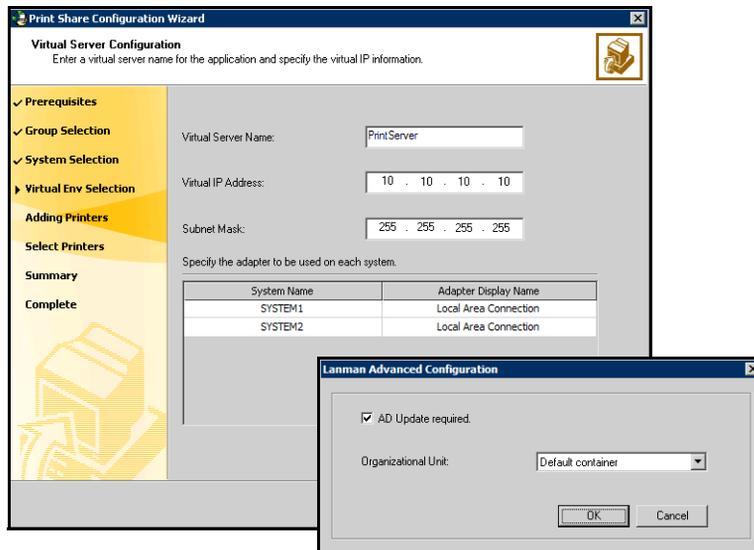
**To create a Print Share service group with a PrintSpool resource**

- 1 Start the Print Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Print Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group details and click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



- Service Group Name** Type a name for the Print Share service group.
- Available Cluster Systems** Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.  
To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.  
To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

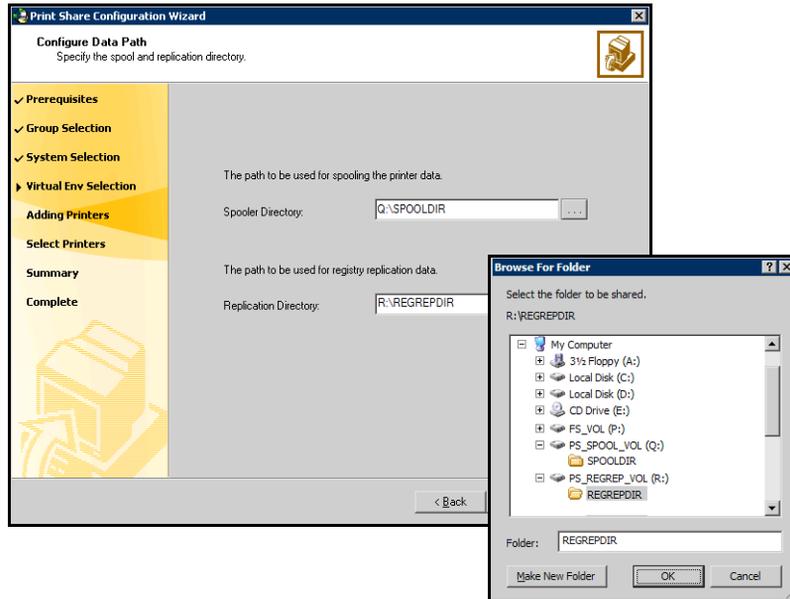
- 5 On the Virtual Server Configuration panel, specify information related to your network and then click **Next**.



- Virtual Server Name** Type a unique virtual computer name by which the server will be known to clients. Note that the virtual name must not exceed 15 characters.

Virtual IP Address	Type a unique virtual IP address for the virtual server.
Subnet Mask	Type the subnet to which the virtual server belongs.
Advanced Settings	<p>Click <b>Advanced Settings...</b> to specify additional details for the Lanman resource.</p> <p>On the Lanman Advanced Configuration dialog box, complete the following:</p> <ol style="list-style-type: none"><li>1 Check <b>AD Update required</b> check box to enable the Lanman resource to update the Active Directory with the virtual name.</li><li>2 From the Organizational Unit drop-down list, select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."</li><li>3 Click <b>OK</b>.</li></ol> <p>The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.</p>
Adapter Display Name	<p>Displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow.</p> <p>For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.</p>

- 6 On the Configure Data Path panel, specify the spool and registry replication directories and then click **Next**.



Spooler Directory      Type the path or click ... (ellipsis button) to browse for the directory. All print commands will be spooled at this location.

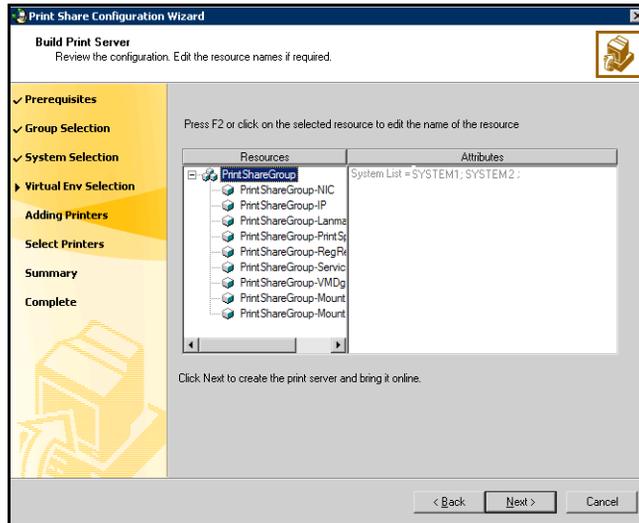
Replication Directory      Type the path or click ... (ellipsis button) to browse for the directory. All changes related to the printer registry keys will be logged at this location.

The selected directories must fulfill the following conditions:

- The selected drive, the mount path, and the file path must not exist in the VCS configuration.
- The directories to be shared must reside on shared, non-system drives.  
Symantec recommends creating the directories for replication and spooling on different mounts.

- 7 On the Build Print Server panel, review the configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**. The wizard starts running

commands to add the PrintSpool resource and the resources on which the PrintSpool resource depends, including the Lanman and ServiceMonitor resources.



- Resources                      Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.  
    To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.
- Attributes                      Displays the attributes and their configured values, for a resource selected in the Resources list.

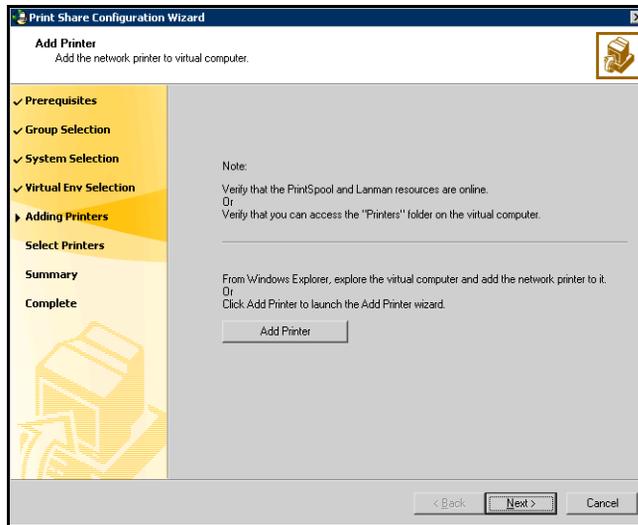
## 8 Bring the PrintSpool resource online.

Proceed to the next step to add the network printer to the virtual computer created by the Lanman resource and to create a new TCP/IP port for the printer.

### To add the network printer to the virtual computer

- 1 Launch the Add Printer wizard to add the network printer to the virtual computer. Before starting the Add Printer wizard, verify that the PrintSpool and Lanman resources are online in your configuration.

To launch the Add Printer wizard, return to the Print Share Configuration Wizard and click **Add Printer** on the Add Printer panel, or in Windows Explorer, search for the virtual computer, explore the virtual computer by double-clicking its name and on the virtual computer's Printers folder, double-click **Add Printer**.

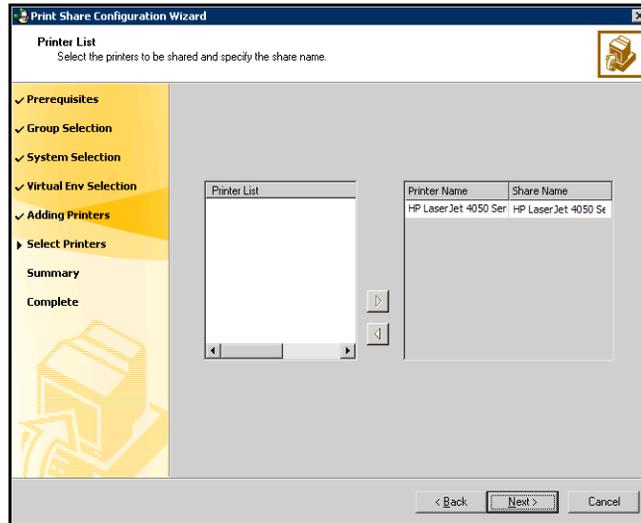


- 2 In the Add Printer wizard, review the information in the Welcome panel and click **Next**.
- 3 Follow the wizard instructions to add the network printer to the virtual computer.  
In the Printer Sharing dialog box, always choose the **Do not share this printer** option.  
Repeat these steps for each additional printer to be installed.
- 4 Return to the Print Share Configuration Wizard, and proceed to the next step to configure a PrintShare resource in your service group and bring it online.

#### To configure a PrintShare resource for the service group

- 1 On the Add Printer panel, click **Next**.

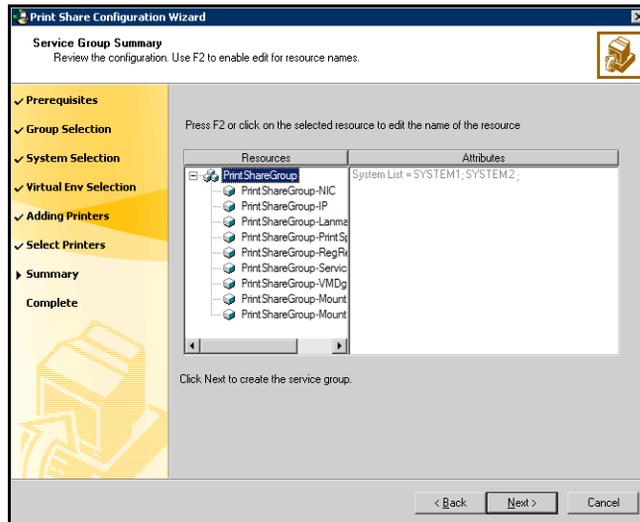
- 2 On the Printer List panel, specify the printers to be included in the Print Share service group and then click **Next**.



- Printer List** Click to select the printer, and then click the right arrow to include the selected printers in your service group. To remove a selected printer from your service group, click the printer from the Printer Name list and click the left arrow.
- Share Name** Type a unique share name for the printer by which it will be known to clients. If you previously chose to share the printer, VCS uses the printer's share name.

- 3 On the Service Group Summary panel, review the service group configuration and then click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration.

Click **Yes**. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



- Resources** Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.  
To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.
- Attributes** Displays the attributes and their configured values, for a resource selected in the Resources list.

In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

## Configuring the IIS service group

Configuring the IIS service group involves creating a IIS service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

## Prerequisites

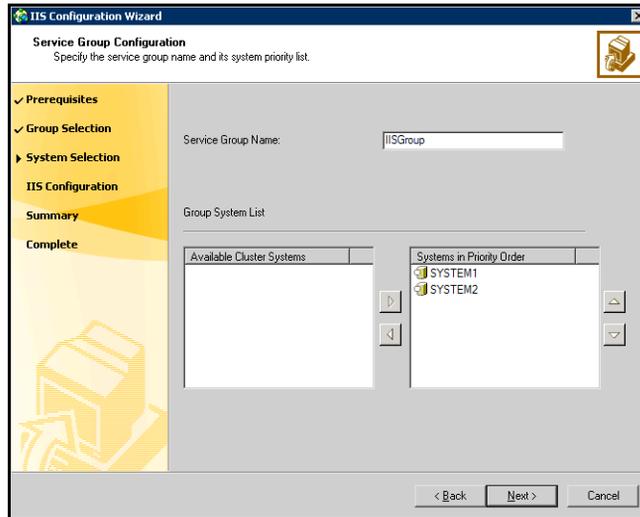
- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Mount the drives containing the shared directories from the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
  - IIS sites to be monitored.
  - Application pools associated with each site.
  - Port numbers associated with each site.
  - Virtual IP addresses and computer names associated with the sites. The virtual IP addresses and the virtual computer names must have forward and reverse entries in the DNS.

## Creating the IIS service group

### To create an IIS service group

- 1 Start the IIS Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.  
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > IIS Configuration Wizard**.
- 1 Start the IIS Configuration Wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > IIS Configuration Wizard**)
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.

- 4 On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the IIS service group.

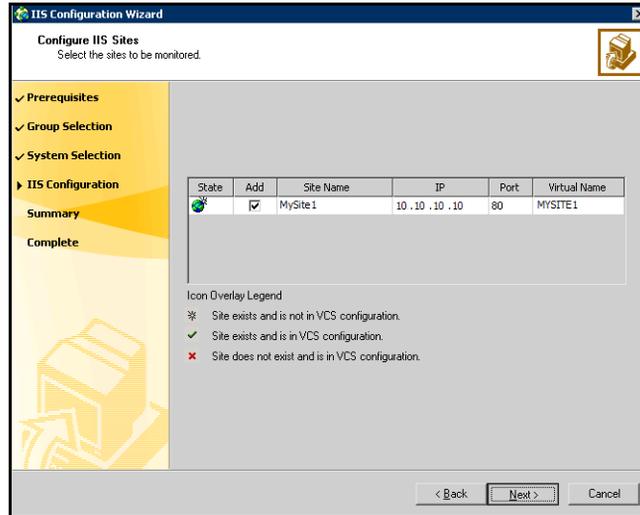
Available Cluster Systems

Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

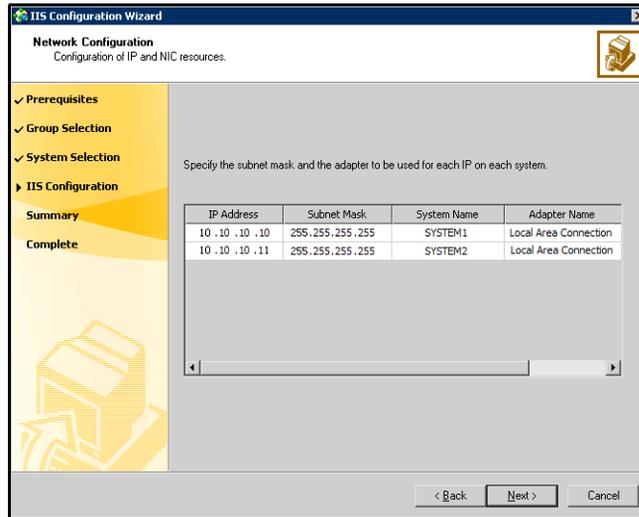
To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the Configure IIS Sites panel, add and remove sites from the service group, configure IP addresses, ports, and virtual computer names, and then click **Next**.



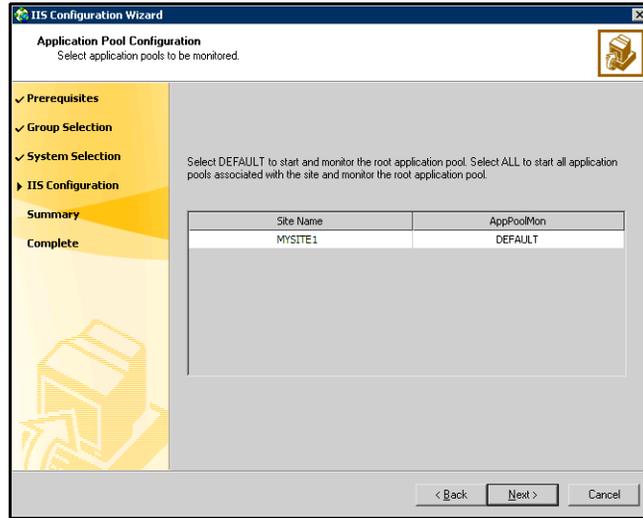
- Add Check the check box corresponding to the site to be configured in VCS.
- IP Type the virtual IP address for each site to be configured. Make sure that each virtual IP address is associated with only one virtual computer name and vice-versa.
- Port Type the port number for each site to be configured.

- 6 On the Network Configuration panel, specify information related to the virtual IP addresses and click **Next**.



- |              |   |
|--------------|---|
| IP Address   | Displays the virtual IP addresses. The wizard groups systems by the virtual IP addresses associated with the systems. |
| Subnet Mask  | Type the subnet mask associated with each virtual IP address.   |
| Adapter Name | Select the adapter associated with the virtual IP address on each system.   |

- 7 On the Application Pool Configuration panel, select the monitoring options for application pools associated with each site and click **Next**.



Site Name

Displays the site names.

AppPoolMon

For each site, select the monitoring options from the AppPoolMon list.

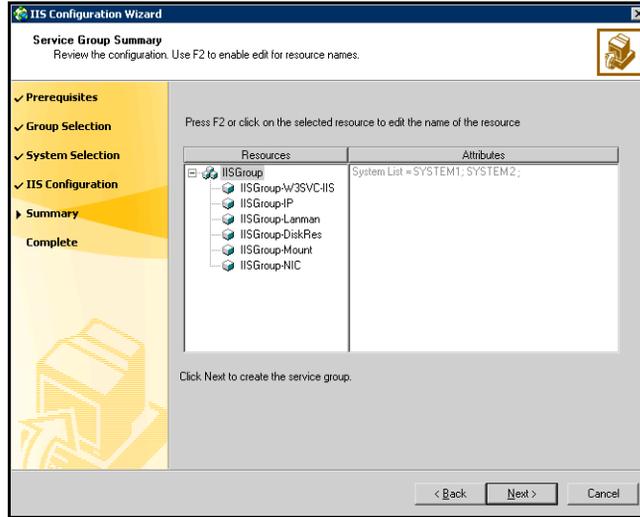
**NONE**—The agent will not monitor the application pool associated with the site.

**DEFAULT**—Starts and monitors the root application pool associated with the site.

**ALL**—Starts all application pools associated with the site and monitors root application pool.

- 8 On the Service Group Summary panel, review the service group configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click

**Yes.** The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



Resources

Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

Attributes

Displays the attributes and their configured values, for a resource selected in the Resources list.

- 9 In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

## Configuring the MSVirtual Machine service group

Configuring the MSVirtual Machine service group involves creating a MSVirtual Machine service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

## Prerequisites

- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the shared drives required by the applications are mounted.
- Mount the drives containing the shared directories of the virtual machine, on which the wizard will be run. Unmount the drives from other systems in the cluster.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Disable the firewall on each node that will host the service group.
- You must have the following information ready. The wizard will prompt you for this information:
  - The name of the virtual machine.
  - Destination on shared disks for the virtual hard disk files.
  - Network adapters on physical nodes to be associated with network adapters on the virtual machine.
  - Information about monitoring heartbeats (optional).

## Creating the MSVirtual Machine service group

### To create the MSVirtualMachine service group

- 1 Start the MSVirtual Machine Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.  
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > MSVirtual Machine Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, select the **Create service group** option and click **Next**.
- 4 Enter a name for the service group and specify the systems on which to configure the service group.
  - Enter a name for the service group.
  - In the **Available Cluster Systems** box, select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the **Systems in Priority Order** box and click the left arrow.

- To change a system's priority in the service group's system list, click the system from the **Systems in Priority Order** and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.

5 Specify details about the virtual machine.

- Select the virtual machine.
- For each virtual disk, specify a destination folder where the virtual hard disk files will be moved. Click the Browse icon to browse for folders.
- To enable detail monitoring for the virtual machine, select the **Monitor Heartbeats** check box and enter failed heartbeat threshold in the **No. of Monitor Cycles** field.

The threshold defines the number of consecutive monitor cycles the agent waits to detect heartbeats from the virtual machine before declaring the resource as faulted.

- Click **Next**.

6 Select an adapters corresponding to the virtual machine on each system.

- For each system in the cluster, enter or click a network adapter name to be associated with the network adapters on the virtual machine. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.

The fields for the virtual IP address and subnet mask are disabled by design.

- Click **Next**.

7 Review the service group configuration.

The **Resources** box lists the configured resources. Click on a resource to view its attributes and their configured values in the **Attributes** box.

- The wizard assigns unique names to resources. Change names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.

- Click **Next**.

- A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.  
The wizard starts running commands to create the service group. Various messages indicate the status of these commands.
- 8 In the completion dialog box, select the check box if you want to bring the service group online on the local system.
  - 9 Click **Finish**.

## Configuring the service group for any additional applications

Configuring the service group for any additional application involves creating an application service group and defining the attribute values for its resources. This can be done using the Application Configuration Wizard. After the service group is created, you must configure the shares to mount automatically at startup.

### Prerequisites

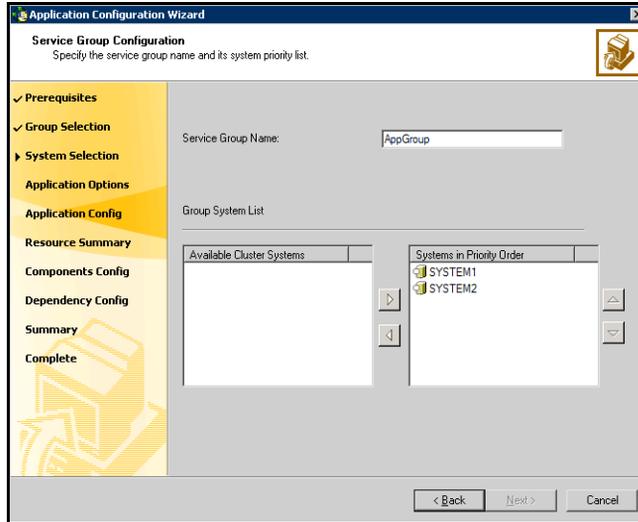
- Verify that the binaries of the application to be configured are present on the nodes on which the service group will be configured.
- Verify that the shared drives required by the applications are mounted.
- Before running the wizard, make sure you have the following information ready:
  - Type of applications for which resources are to be configured.
  - Shared storage used by the applications.
  - Registry replication information.
  - Network information.

### Creating the application service group

- 1 Start the Application Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.  
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Application Configuration Wizard**.  
or

Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Configuration Wizard**.

- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, click **Create service group** and click **Next**.
- 4 Specify the service group name and system list.

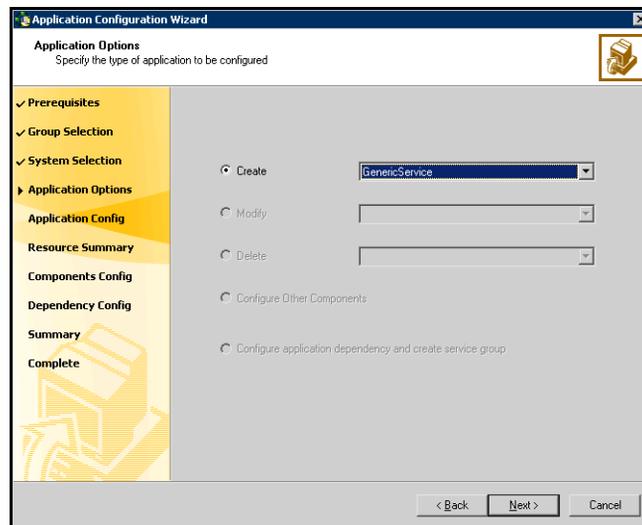


- Enter a name for the service group.
  - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list.  
To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.
  - To change a system's priority in the service group's system list, select the system in the Systems in Priority Order list and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
  - Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 The Application Options panel provides you the option to specify the type of application to be configured. The available options are:
    - **Generic Service**: Configures a service using the Generic Service agent. The agent brings services online, takes them offline, and monitors their status. See "[Configuring a GenericService resource](#)" on page 131.

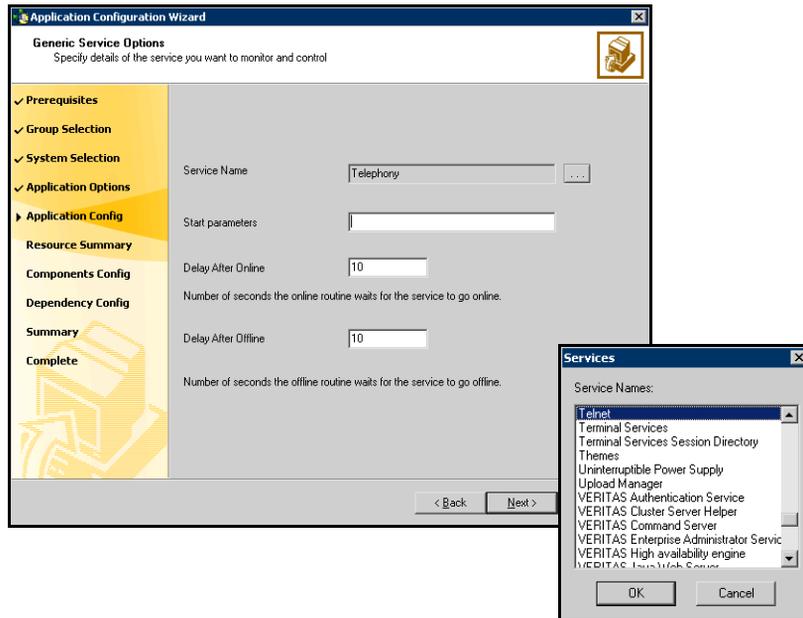
- **Process:** Configures a process using the Process agent. The agent brings processes online, takes them offline, and monitors their status. See [“Configuring processes”](#) on page 135.
- **Service Monitor:** Configures a service using the ServiceMonitor agent. The agent monitors a service or starts a user-defined script and interprets the exit code of the script. See [“Configuring a ServiceMonitor resource”](#) on page 139.

## Configuring a GenericService resource

- 1 In the Application Options panel, click **Create**, select **GenericService** from the corresponding drop-down list, and click **Next**.

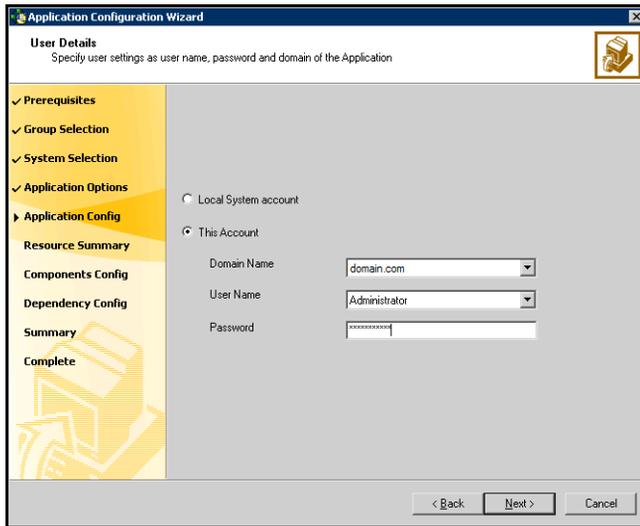


- 2 Select the service name for which you wish to configure a GenericService resource. Also specify the attributes for the resource.



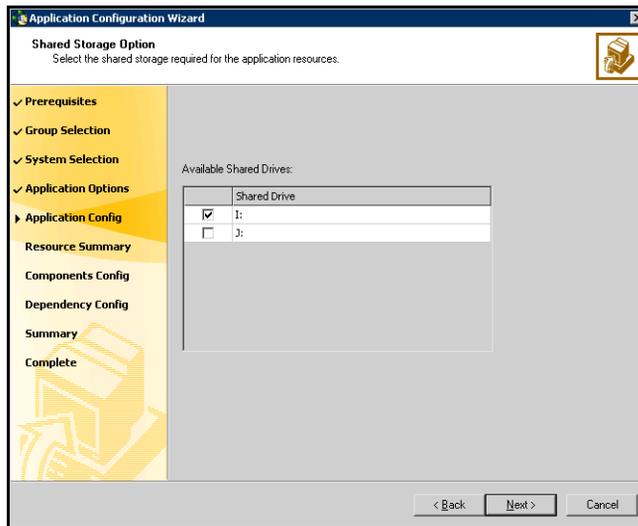
- Click the icon (...) adjacent to the Service Name text box.
- In the Services dialog box, select a service and click **OK**. The selected service appears in the Service Name text box.
- In the **Start Parameters** text box, provide the start parameters for the service, if any.
- In the **Delay After Online** text box, specify the number of seconds the agent waits after the service is brought online before starting the monitor routine.
- In the **Delay After Offline** text box, specify the number of seconds the agent waits after the service is taken offline before starting the monitor routine.
- Click **Next**.

3 Specify the information about the user in whose context the service will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the GenericService resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

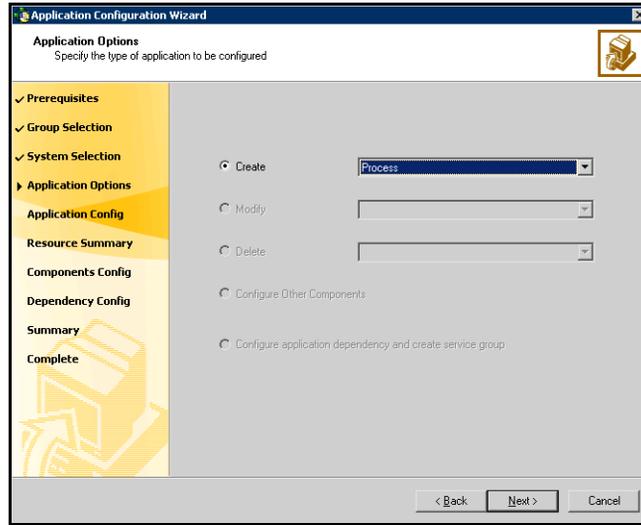


- In the Available Shared Drives box, select the check box adjacent to the shared drive.
  - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the GenericService resource. Click **Back** to make changes. Otherwise, click **Next**.
  - 6 The Application Options panel appears. Select one of the following options:
    - To configure another GenericService resource, repeat [step 1](#) through [step 5](#).
    - To configure a Process resource, proceed to “[Configuring processes](#)” on page 135 for instructions.
    - To configure a ServiceMonitor resource, proceed to “[Configuring a ServiceMonitor resource](#)” on page 139 for instructions.
    - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 142 for instructions.

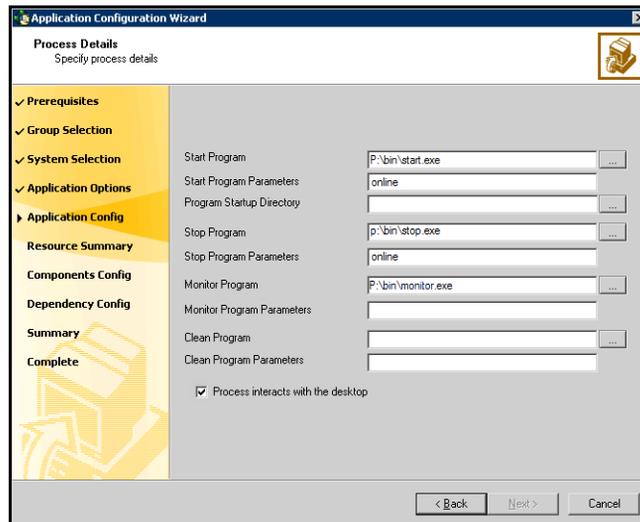
If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 148.

## Configuring processes

- 1 In the Application Options panel, click **Create**, select **Process** from the corresponding list, and click **Next**.



- 2 Specify the details for the process.



- In the **Start Program** text box, specify the complete path of the program that will start the process to be monitored by VCS. You can

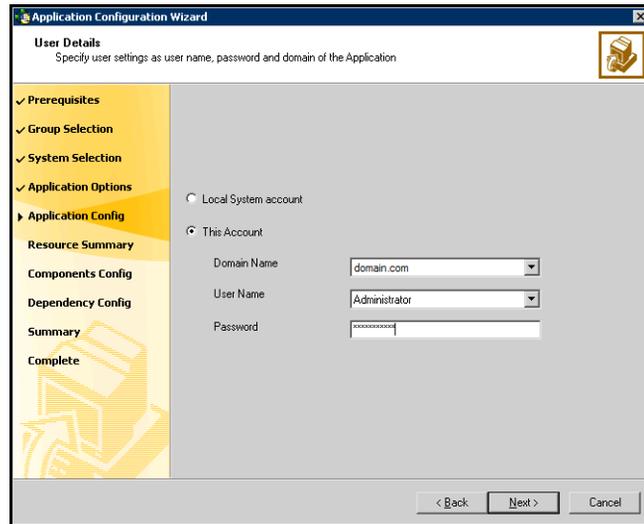
choose to either type in the location of the program or browse for it using the (...) icon.

- In the **Start Program Parameters** text box, specify the parameters used by the Process agent start program.
- In the **Program Startup Directory** text box, enter the complete path of the Process agent program or browse for it by clicking the (...) icon.
- In the **Stop Program** text box, enter the complete path of the program that will stop the process started by the Start Program or browse for it by clicking the (...) icon.
- In the **Stop Program Parameters** text box, specify the parameters used by the stop program.
- In the **Monitor Program** text box, enter the complete path of the program that monitors the Start Program or browse for it by clicking the (...) icon.

If you do not specify a value for this attribute, VCS monitors the Start Program. If the Start Program is a script to launch another program, you must specify a monitor program.

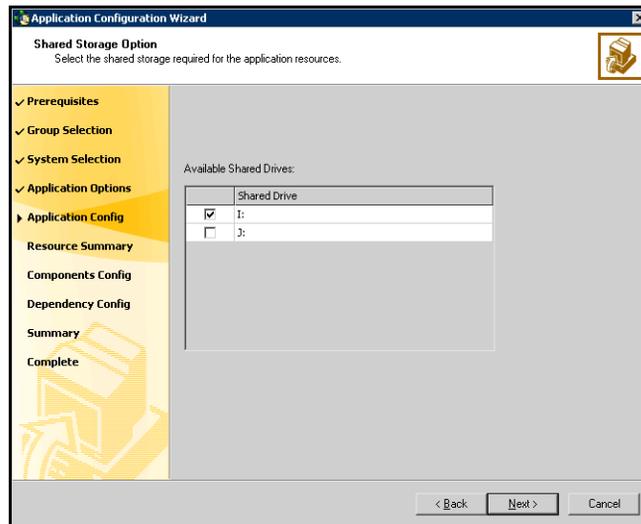
- In the **Monitor Program Parameters** text box, specify the parameters used by the monitor program.
- In the **Clean Program** text box, enter the complete path of the Clean process or browse for it by clicking the (..) icon.
- If no value is specified, the agent kills the process indicated by the Start Program.
- In the **Clean Program Parameters** text box, specify the parameters used by the Clean program.
- Select the **Process interacts with the desktop** check box if you want the process to interact with your Windows desktop. Setting this option enables user intervention for the process.
- Click **Next**.

3 Specify information about the user in whose context the process will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the Process resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

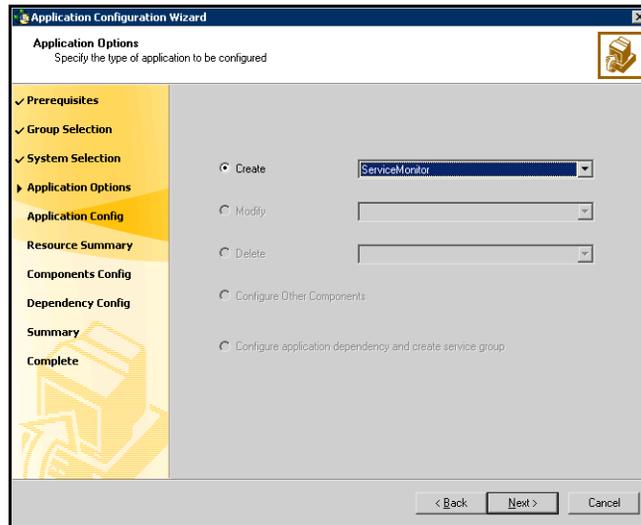


- From the Available Shared Drives box, select the check box adjacent to the shared drive.
  - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the Process resource. Click **Back** to make changes. Otherwise, click **Next**.
  - 6 The Application Options panel appears. Select one of the following options:
    - To configure another Process resource, repeat [step 1](#) through [step 5](#).
    - To configure a GenericService resource, see “[Configuring a GenericService resource](#)” on page 131 for instructions.
    - To configure a ServiceMonitor resource, proceed to “[Configuring a ServiceMonitor resource](#)” on page 139 for instructions.
    - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 142 for instructions.

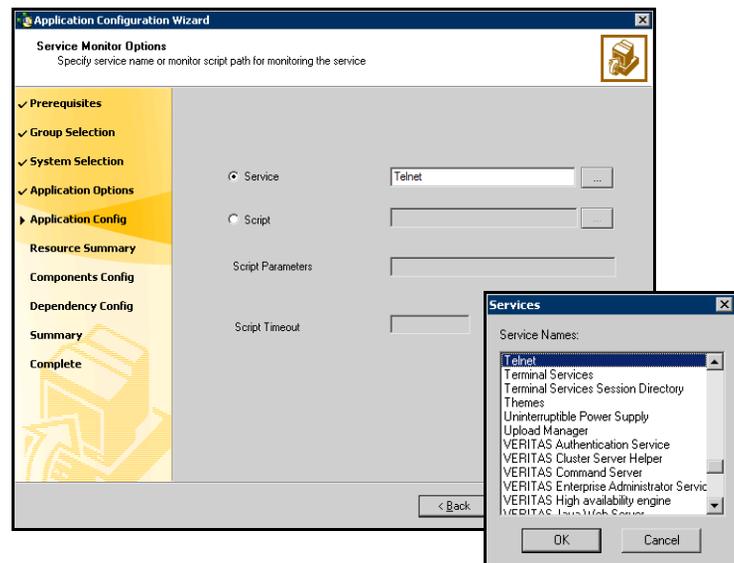
If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 148.

## Configuring a ServiceMonitor resource

- 1 In the Application Options panel, click **Create**, select **ServiceMonitor** from the corresponding drop-down list, and click **Next**.



- 2 Specify the service to be monitored or a user-defined script to monitor a service.



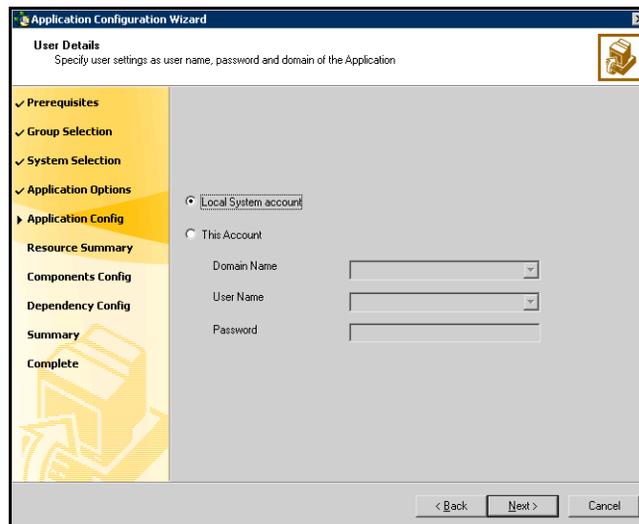
If you want VCS to monitor the service:

- Select the **Service** option and click the icon (...) adjacent to the **Service Name** text box.
- In the Service dialog box, select the service and click **OK**. The selected service name appears in the **Service Name** text box. Alternatively, You may also type in the service name to be monitored.
- Click **Next**.

If you want a script to monitor the service:

- Specify the complete path for the script using the Browse button (...).
- Specify the parameters for the script.
- Specify the time in seconds for the agent to receive a return value from the monitor script.
- Click **Next**.

3 Specify the user information in whose context the service will be monitored.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.

If the service selected in [step 2](#) on page 139 is running in the context of a local system account, the **This Account** option is disabled. Similarly, if

the service is running in the context of any other user account, the **Local System account** option is disabled.

■ Click **Next**.

ServiceMonitor resource belongs to the category of *persistence* resources. Such resources do not depend on other VCS resources, including shared storage. Hence, the Shared Storage Option panel does not appear if you select the ServiceMonitor option.

- 4 In the Application Resource Summary panel, review the summary of the ServiceMonitor resource. Click **Back** to make changes. Otherwise, click **Next**.
- 5 The Application Options panel appears. Select one of the following options:
  - To configure another ServiceMonitor resource, repeat [step 1](#) through [step 4](#).
  - To configure a GenericService resource, see “[Configuring a GenericService resource](#)” on page 131 for instructions.
  - To configure a Process resource, see “[Configuring processes](#)” on page 135 for instructions.
  - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 142 for instructions.

If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 148.

## Configuring VCS components

Applications configured using GenericService or Process resources may require File Share components, network components, or Registry Replication resources. You can configure these VCS components *only* for service groups created using the wizard.

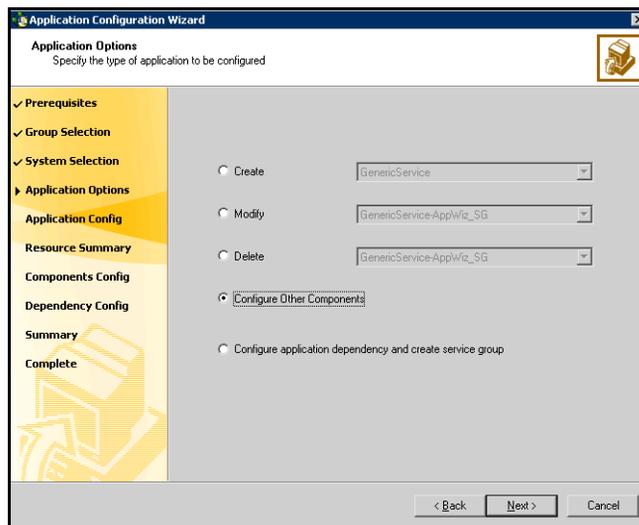
---

**Note:** Configure these components only after configuring all application resources. The wizard creates a service group after these components are configured. To add more application resources, you must rerun the wizard in the Modify mode.

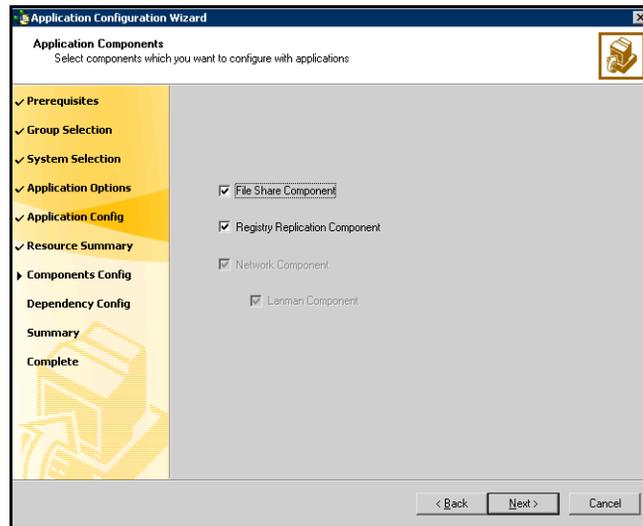
---

### To configure VCS components

- 1 In the Application Options panel, click **Configure Other Components**.



2 Select the VCS component to be configured for your applications.



The available options are:

- **File Share Component:** Select this to configure a FileShare resource for your application. To configure a FileShare resource, proceed to the next step.

If you select to configure the File Share component, the **Network Component** check box is checked by default.

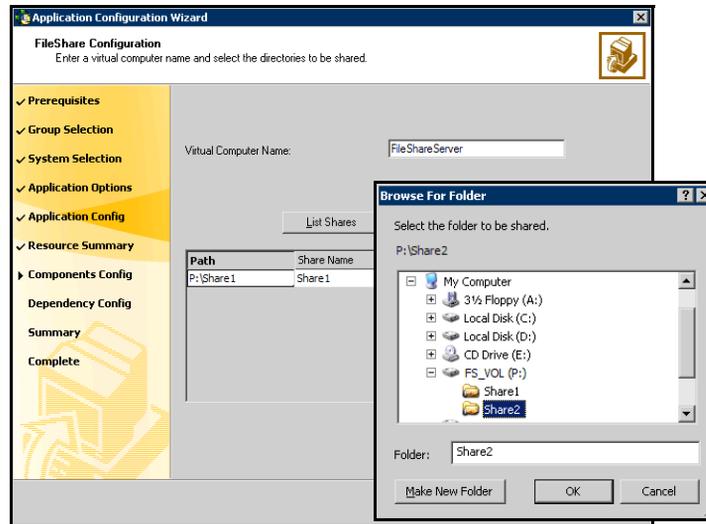
- **Registry Replication Component:** Select this option to configure registry replication for your application. To configure a Registry Replication resource, proceed to [step 5](#) on page 146.
- **Network Component:** Select this option to configure network components for your application. If you wish to configure a virtual computer name, check **Lanman component** also. To configure a network resource, proceed to [step 6](#) on page 147.

The wizard does not enable the **Lanman Component** check box unless the **Network Component** check box is checked.

### To configure a FileShare resource

The File Share Configuration panel appears only if you chose File Share component in the Application Component panel.

- 3 Specify the configuration information for the FileShare resource to be created.

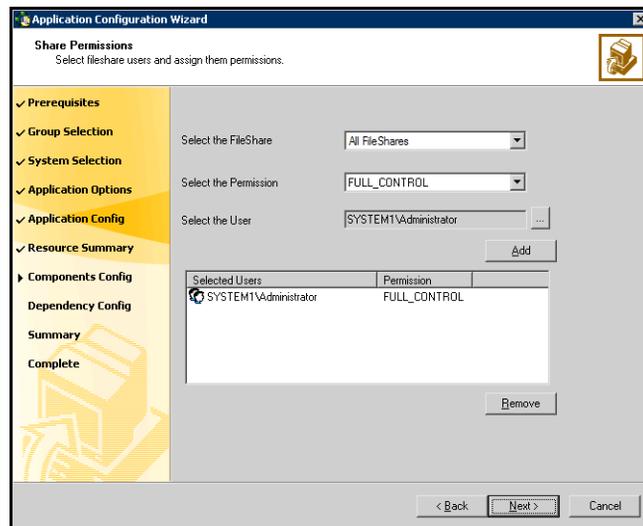


- Enter a unique virtual computer name by which the server will be known to clients.  
The virtual name must not exceed 15 characters.
- Click **List Shares** to view all the existing shares on the shared storage. Select a share and click **Add**.
- In the **Path** column, type or click the path of the directories to be shared. Click the Edit icon (...) to browse for folders. The selected directories must meet the following conditions:
  - The selected drive, the mount path, and the file path must not exist in the VCS configuration.
  - The directories to be shared must reside on shared, non-system drives.

The wizard validates the selected directory and displays an error message if the directory does not meet any of the conditions.

- If a selected directory is already shared, the Share Name column lists the names by which it is shared. You can selected a listed share name to make an existing share highly available. You can also create a new share for the same directory by typing a new share name.

- To make the new share a hidden share, check **Hide Share**. To share subdirectories, check **Share Subdirs**. To hide shared subdirectories, check **Hide Child Shares**.
  - Click **Add** to add a file share. Repeat [step n](#) through [step n](#) for each file share to be created. Click **Remove** to remove a File Share from the configuration.
  - Click **Next**. The wizard begins validating your configuration. Various messages indicate the validation status. After the validations are completed, the Share Permissions panel appears.
- 4 Specify the users for the file share and assign permissions to them.



- From the Select the FileShare drop-down list, select the file share with which to associate user permissions, or select the default **All FileShares** to set the same permissions for all file shares.
- From the Select the Permission drop-down list, select the permission to be associated with the user.
- Click the ... (ellipsis button) adjacent to the Select the User list, select a user to be assigned permissions for the selected file share, and click **OK**. Then click **Add** to add the specified user to the Selected Users list. The selected user will be assigned the permission selected in the Select the Permission list. By default, all selected users are given READ\_ACCESS permission.
- The Selected Users list displays a list of selected users and their file share permissions. You can configure a maximum of 50 users for each file share. To configure more users, create a user group.

- To change the file share permission associated with a user, click a user name in the Selected Users list and then select the desired permission from the Select the Permission drop-down list.
- To deny file share access to a user, click the user name in the Selected Users list and click **Remove**.

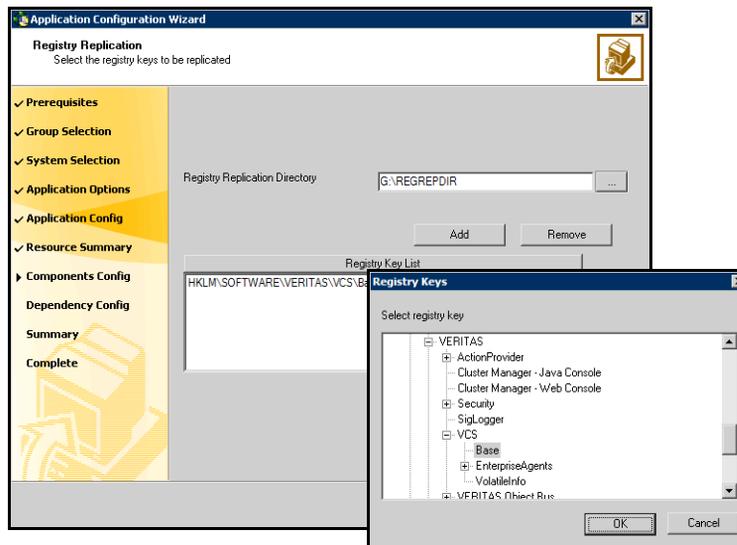
Repeat the process for each file share with which to associate users and permissions.

- Click **Next**.  
If you selected Registry Replication from the Application Component panel, proceed to the next step. Otherwise, proceed to [step 6](#) on page 147.

### To configure Registry Replication

The RegistryReplication panel appears only if you chose to configure the Registry Replication Component in the Application Component panel.

- 5 Specify the registry keys to be replicated.



- Specify the directory on the shared disk in which the registry changes are logged.
- Click **Add**.
- In the Registry Keys dialog box, select the registry key to be replicated.
- Click **OK**.
- The selected registry key is added to Registry KeyList box. Click **Next**.

If you chose Network Component from the Application Component panel, proceed to the next step. Otherwise, proceed to [step 7](#) on page 147.

### To configure network components

The Virtual Computer Configuration panel appears only if you chose to configure the Network Component in the Application Component panel.

- 6 Specify information related to your network.
  - Enter a unique virtual computer name by which the node will be visible to the other nodes. Note that the virtual name must not exceed 15 characters.

Note that the Virtual Computer Name text box is displayed only if you chose to configure the Lanman Component in Application Component panel. However, if you chose to configure a FileShare resource in the service group, the Virtual Computer Name text box is not displayed. In such a case, the Lanman resource uses the virtual computer name specified for the FileShare resource.
  - Enter a unique virtual IP address for the virtual server.
  - Enter the subnet to which the virtual server belongs.
  - Click **Advanced...** to specify additional details for the Lanman resource.
    - Check **AD Update required** to enable the Lanman resource to update the Active Directory with the virtual name.
    - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
  - Click **OK**.
  - For each system in the cluster, select the public network adapter name. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.

Note that the wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if applicable. Verify that you select the adapters assigned to the public network, not the private.
  - Click **Next**.
- 7 The Application Options panel is displayed. Select one of the following options:
  - To configure additional VCS components, repeat [step 1](#) on page 142 through [step 6](#) on page 147.

- To configure a GenericService resource, see “[Configuring a GenericService resource](#)” on page 131 for instructions.
- To configure a Process resource, see “[Configuring processes](#)” on page 135 for instructions.
- To configure a Service Monitor resource, see “[Configuring a ServiceMonitor resource](#)” on page 139 for instructions.

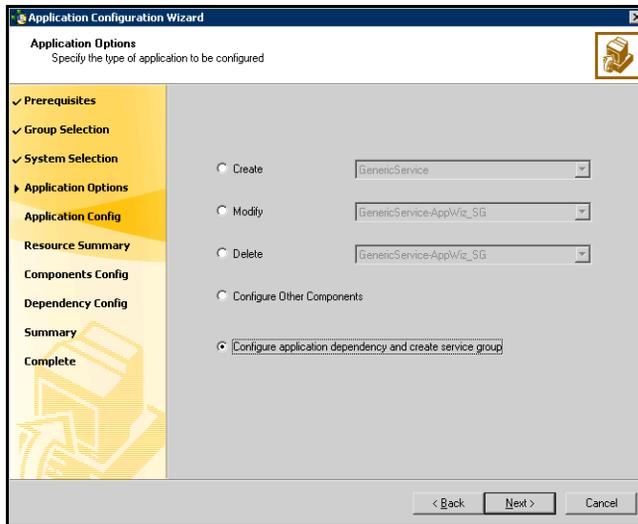
If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 148.

## Configuring Application Dependencies

The Application Configuration Wizard enables you to create service group for the application resources and other VCS components configured using the wizard. This section describes how to create the service group using the wizard.

### To create a service group

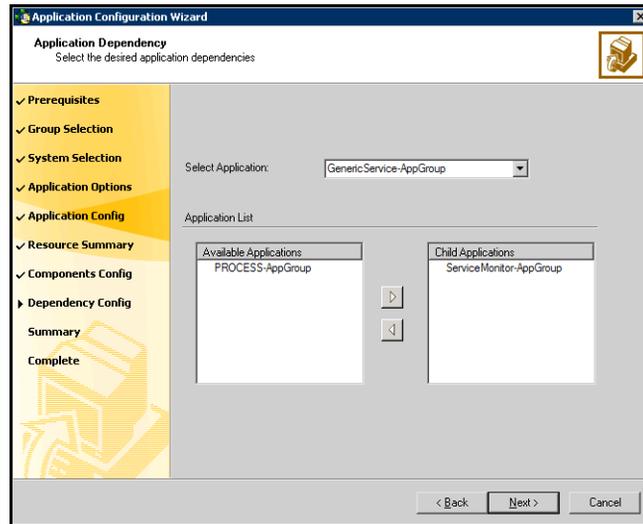
- 1 In the Application Options panel, click **Configure application dependency and create service group**.



The option is enabled only if:

- resources and VCS components are already configured using the wizard.
- you clicked **Modify Service Groups** in the Wizard Options panel.

2 Specify the dependency between the applications.

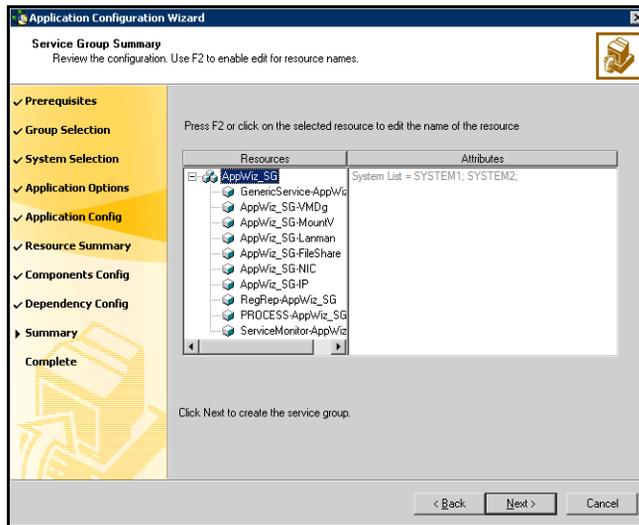


You must have at least two resources configured for Application Dependency panel to appear. Of the two resources, one should either be a GenericService or a Process resource.

- From the Select Application list, select the application that would depend on other applications. The selected application becomes the parent application.
- From the Available Applications list, select the application on which the parent application would depend and click the right-arrow icon to move the application to the Child Applications list.
- To remove an application from the Child Applications list, select the application in the list and click the left arrow.
- Repeat these steps for all such applications for which you want to create a dependency.
- Click **Next**.

The Application Dependency panel enables you to link resources configured using the wizard. If these resources are dependent on other services outside the VCS environment, you should first configure resources for such services and then create the appropriate dependency.

### 3 Review the service group configuration.



The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

- The wizard assigns unique names to resources. Change names of resource, if required.  
To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.
  - Click **Next**.
  - A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.  
The wizard starts running commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 4 In the completion panel, check **Bring the service group online** if you want to bring the service group online on the local system.
  - 5 Click **Finish** to create the service group and exit the Application Configuration Wizard.

# Verifying the cluster configuration

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Veritas Cluster Manager (Java Console) to switch all the service groups from one node to another.
- Simulate a local cluster failover by shutting down an active cluster node.

## To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
  - Click **Switch To**, and click the appropriate node from the menu.
  - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.  
If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

## To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node:
  - Restart the node you shut down in [step 1](#).
  - Click **Switch To**, and click the appropriate node from the menu.
  - In the dialog box, click **Yes**.  
The service group you selected is taken offline and brought online on the node that you selected.

## Possible tasks after completing the configuration

After completing the configuration, you may want to make some changes to the cluster configuration or modify the application service groups. Depending on your specific requirements perform one of the following operations:

- Configuring the Cluster Management Console to centrally administer multiple clusters.  
See “[Configuring the Cluster Management Console connection](#)” on page 152.
- Modifying the existing cluster configuration to add additional nodes, removing existing nodes, or configure the Web Console.  
See “[Modifying the existing cluster configuration](#)” on page 157.
- Modifying the replication service group configuration.  
See “[Modifying the application service groups](#)” on page 164.

## Configuring the Cluster Management Console connection

The Veritas Cluster Management Console (CMC) is a centralized management solution for high-availability application environments based on Veritas Cluster Server. CMC can be configured to locally manage a single cluster or to centrally manage multiple clusters.

CMC comprises of the following components:

- *Management Server*  
The management server accepts and processes the operational commands and the configuration inputs that users enter through CMC. The management server communicates with the VCS High Availability engine (HAD). Install the CMC Management Server only if you plan to centrally manage multiple clusters. You must install the management server on a standalone system that is outside any cluster but available on the local network.
- *Cluster Connector*  
The cluster connector is an agent that enables the management server to communicate with clusters through intervening firewalls. You must install the cluster connector on each cluster that is separated from the management server by a firewall. If there are no firewalls between the management server and the clusters, you can configure the clusters to use direct connection instead.  
In each cluster, the cluster connector runs on one node at a time, but is installed on all nodes and is configured for failover.

This section describes how to install the cluster connector on VCS clusters. For more information on CMC and its components, see the *Veritas Cluster Management Console Implementation Guide*.

## Prerequisites for installing the cluster connector

- You must stop all VCS Web consoles, VCS Java consoles, and agent wizards that are running on any cluster nodes before you install the cluster connector
- When you install the cluster connector, Symantec Product Authentication Service 4.3.x must be available on the system from which you run the installer. If you install from a standalone system, you must manually install the authentication service on that system before you install the cluster connector. If you install from a cluster node that is also a member of the target cluster, the installer provides the authentication service automatically.
- When installing the cluster connector on 64-bit Windows platforms from a 32-bit system, the default installation directory is C:\Program Files. Symantec recommends that you change the 64-bit installation directory to C:\Program Files (x86).
- Ensure that your network and DNS configuration provide proper name resolution. Otherwise, the cluster connector cannot resolve the management server host name when attempting to connect to the management server.
- The cluster connector requires the management server network address. For example, mgmtserver1.symantecexample.com.
- A CMC service account password. You must have set this account password while installing the management server.
- The root hash of the management server. Use the `vssat showbrokerhash` command and copy the root hash of the management server. Note that you must run this command from the C:\Program Files\Veritas\Security\Authentication\bin directory on the management server.

## Installing the cluster connector on Windows clusters

Perform this procedure to use the cluster connector for management server communications with a supported Windows cluster.

### To install the cluster connector on a Windows cluster

- 1 Insert the Veritas™ Cluster Management Console *for Windows* disc into the disc drive on the local system.
- 2 Locate the \installer\installer directory.
- 3 Double-click **setup.bat**.
- 4 In the Welcome to the Veritas Cluster Management Console Installation Manager dialog box, read the introduction and then click **Next**.
- 5 In the Installation and Configuration Options dialog box, click **Add Clusters or clustered systems to a management server**, and then click **Next**.
- 6 In the Cluster Connector Cluster Selection dialog box, follow the dialog box instructions exactly as specified, and then click **Next**.  
The installer performs a check for WMI on the specified nodes to ensure that they are ready for the cluster connector installation.
- 7 When prompted, enter user account information for each cluster. If a cluster is secure, you are prompted for a domain name in addition to a user name and password that is valid for the cluster.
- 8 In the Cluster Connector Directory Selection dialog box, do one of the following and then click **Next**:
  - Leave the default directories provided
  - Double-click on a directory, or click a directory and then press F2, and then specify another directory
  - Click **Reset all** to specify new directories on each node
- 9 In the Management Server Information dialog box, provide the IP address for the management server to which the cluster connector is intended to connect.  
You cannot change the port specification, 14145, but it is provided to help you to prevent port conflicts when configuring other software. The other ports used by the Cluster Management Console are 8181 (HTTP), 8443 (HTTPS), and 2994 (DBMS; this port can be shared with other Symantec products)
- 10 In the Services Account Password dialog box:
  - Enter a password for the user account that the cluster connector uses for management server communications
  - Enter the root hash of the authentication broker used by the authentication broker installed on the management serverThe password is the password that was entered for the cluster connector service account during management server installation.

To retrieve the root hash of the management server authentication broker, run the following command:

```
\program files\veritas\security\authentication\bin\vssat  
showbrokerhash
```

The output of this command looks similar to the following:

```
Root Hash:          9dfde3d9aaebee084f8e35819c1fed7e6b01d2ae
```

Enter or copy the alphanumeric string into the Root Hash text box (the string you receive is different from the one shown).

- 11 In the Summary dialog box, review the information you have specified and, if satisfactory, click **Next** to accept it and start the installation. The Installing Veritas Cluster Management Console dialog box displays a progress bar and a status message window for the installation.
- 12 After the installation is complete, click **Next**.
- 13 In the Completed the Symantec Veritas Cluster Management Console Installation Manager dialog box, click **Finish**.

The installer creates log files at C:\Documents and Settings\All Users\Application Data\Veritas\Cluster Management Console. The file names are Install\_GUI\_0.log and Install\_MSI\_0.log. The installer creates Install\_GUI\_0.log on the system from which you run the cluster connector installation. The installer creates Install\_MSI\_0.log on the target systems.

## Avoiding service group faults on Windows clusters configured in secure mode

If you install the cluster connector on a Windows cluster that is configured in secure mode, the cluster connector service account, CMC\_CC@CMC\_SERVICES, might fail to authenticate on the cluster nodes. The installer reports an error about the failed authentication.

If the service account authentication fails, the ClusterConnector resource faults on the cluster, causing the CMC service group to fault. If the CMC service group faults, the ClusterConnector.log file contains the error message:

```
Can not get Cache Credential for CMC_CC
```

You must rectify any clock skew that exists among the cluster or management server systems before attempting the following procedure.

### To avoid service group faults on Windows clusters configured in secure mode

- 1 On a cluster node, obtain a command prompt and change to the following directory:

```
Veritas\Security\Authentication\bin
```

This directory may be in one of the following paths:

```
C:\Program Files
```

or

```
C:\Program Files\Common Files
```

- 2 Set up a trust relationship between the authentication broker on the management server and the authentication broker on the local cluster node.

Type the following command:

```
vssat setuptrust --broker MS_IPAddress:[2821  
(optional)]--securitylevel high --hash Hash_From_MS
```

- 3 Authenticate the CMC\_CC@CMC\_SERVICES account on the local node. Type the following command:

```
"vssat authenticate --domain vx:CMC_SERVICES --prplname CMC_CC  
--password password_for_CMC_CC_user_created_during_MS_install  
--broker MS_IPAddress:2821
```

Usage for this command is

```
vssat authenticate --domain <type:name> [--prplname <prplname>  
[--password <password>]] [--broker <host:port>]
```

Repeat these steps on each node in the cluster.

## Uninstalling the cluster connector

You must run the cluster connector uninstallation on a cluster node. Use the setup program to remove the cluster connector from each cluster node.

### To uninstall the cluster connector from Windows clusters

- 1 Insert the Veritas™ Cluster Management Console *for Windows* disc into the disc drive on the local system.
- 2 Locate the \installer\installer directory for Cluster Management Console in the \windows folder.
- 3 Double-click the **setup.bat** file.
- 4 In the Welcome to the Veritas Cluster Management Console Installation Manager dialog box, read the introduction and then click **Next**.
- 5 In the Installation and Configuration Options dialog box, click **Uninstall cluster connectors** and then click **Next**.
- 6 Follow the prompts in the uninstallation wizard. When available, click **Finish** to close the wizard.

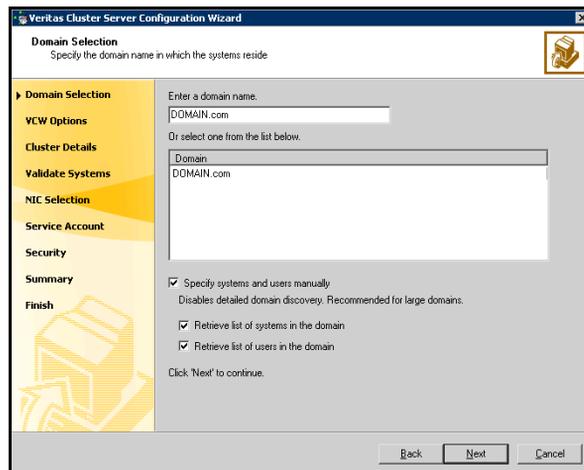
## Modifying the existing cluster configuration

### To add a node to a VCS cluster

- 1 Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)

Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.

- 2 Read the information on the Welcome screen and click **Next**.
- 3 On the Configuration Options panel, choose the **Cluster Operations** option and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



To discover information about all the systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

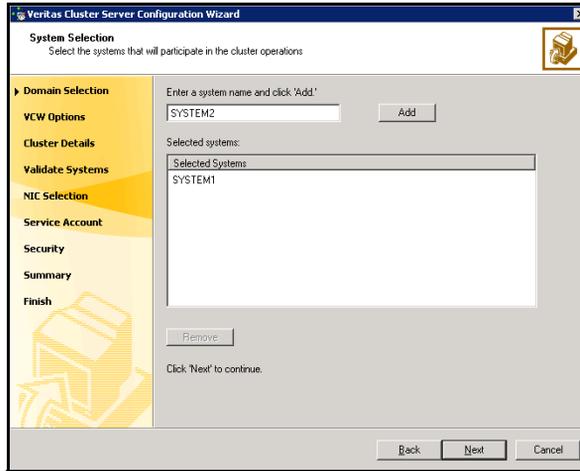
Proceed to [step 7](#) on page 160.

To specify systems and user names manually (recommended for large domains):

- Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

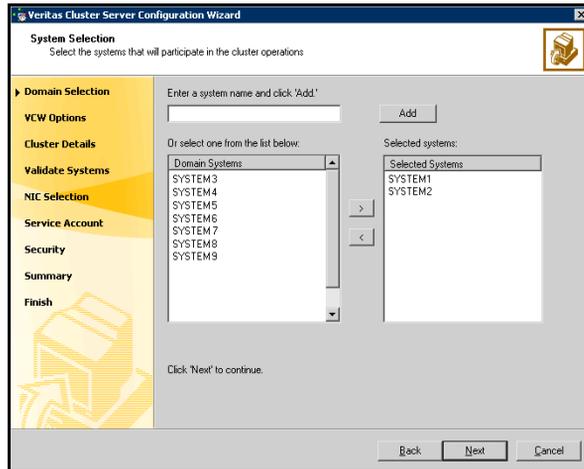
If you selected the **Retrieve system list from domain** check box, proceed to [step 6](#) on page 159. Otherwise proceed to the next step.

- 5 On the System Selection panel, complete the following and click **Next**.



- Type the name of a node in the cluster and click **Add**.
  - Type the name of the system to be added to the cluster and click **Add**.
- If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.
- Proceed to [step 7](#) on page 160.

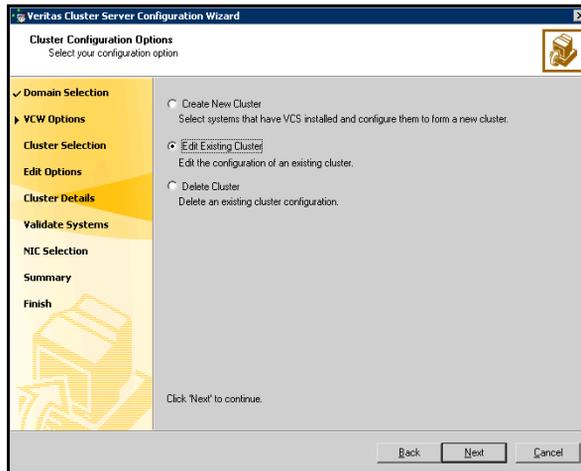
- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems.



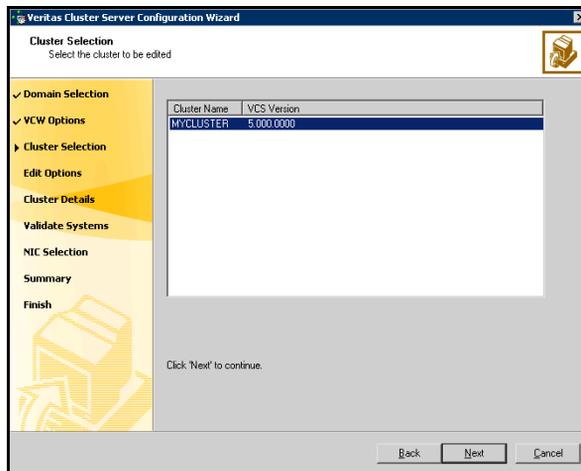
Enter the system name and click **Add** to add the system to the Selected Systems list. Alternatively, you can select the systems from the Domain Systems list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

- 7 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.

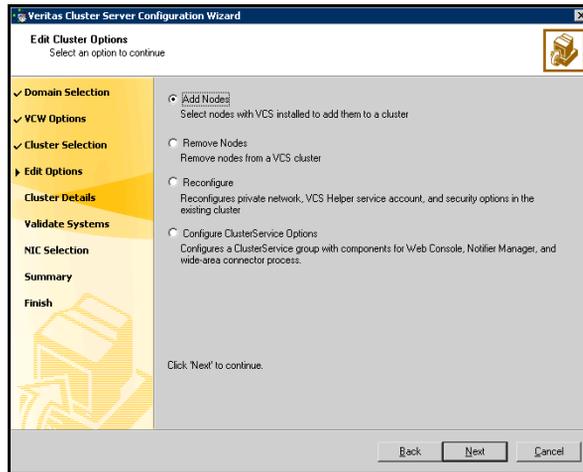


- 8 On the Cluster Selection panel, select the cluster to be edited and click **Next**.



If you chose to specify the systems manually in [step 4](#), only the clusters configured with the specified systems are displayed.

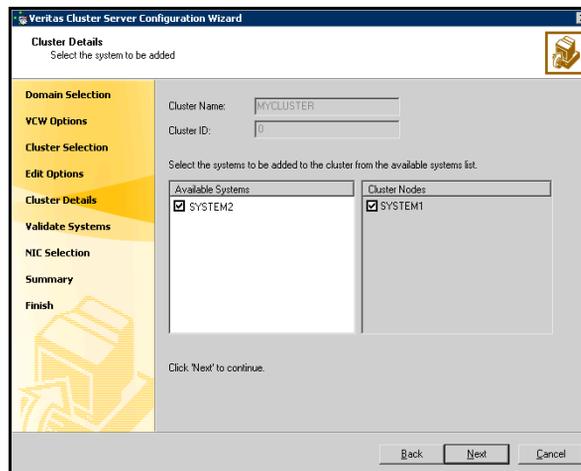
- 9 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**.



In the Cluster User Information dialog box, enter the user name and password for a user with administrative privileges to the cluster and click **OK**.

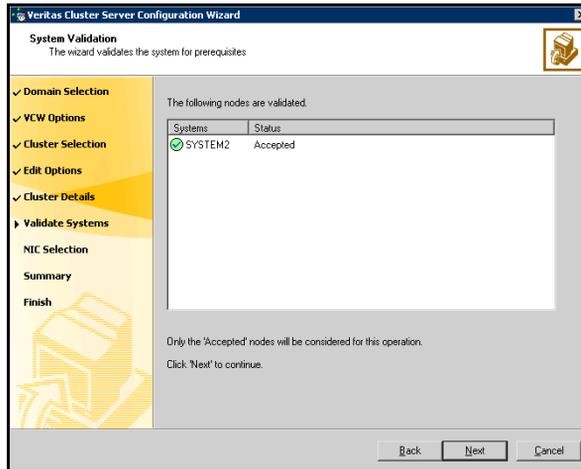
The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges (a cluster that is not a secure cluster).

- 10 On the Cluster Details panel, check the check boxes next to the systems to be added to the cluster and click **Next**.



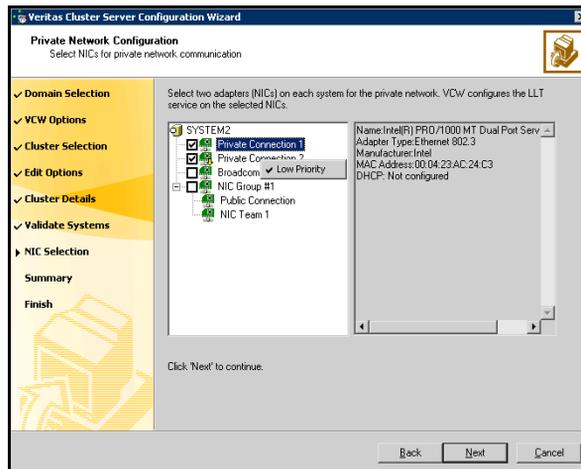
The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

- 11 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.



If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.

- 12 On the Private Network Configuration panel, select two NICs for VCS private network communication, for each system being added, and then click **Next**.



- Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.

- If you have only two NICs on a selected system, make sure you lower the priority of the NIC that is used for public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
  - If your configuration contains teamed NICs, the wizard groups them as NIC Group #N where N is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.
- 13 On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.
  - 14 Specify the credentials for the user in whose context the VCS Helper service runs.
  - 15 Review the summary information and click **Add**.
  - 16 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

## Modifying Values for ClusterService Group Attributes

Modify the following ClusterService group attributes on all the newly added nodes to include local values:

- MACAddress attributes of all the NIC resources
- MACAddress attributes of all the IP resources
- StartProgram, StopProgram, and MonitorProgram attributes of the wac resource
- InstallDir attribute of VCSWeb resource

You can modify these values from the VCS Java Console or Web Console.

## Modifying the application service groups

You may want to modify existing application service groups. Use one of the following options depending on your specific application environment:

- [Modifying the FileShare service group](#)
- [Modifying the PrintShare service group](#)
- [Modifying the IIS service group](#)
- [Modifying the MSVirtual Machine service group](#)
- [Modifying any other application service group](#)

### Modifying the FileShare service group

To modify a FileShare service group

- 1 Start the File Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > File Share Configuration Wizard**.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Wizard Options panel, choose the **Modify service group** option, select the service group to be modified, and click **Next**.
- 5 Follow the wizard instructions and make required modifications to the service group configuration.

### Modifying the PrintShare service group

The Print Share Configuration Wizard enables you to modify a Print Share service group. Following are some points to note before modifying the service group:

- If the Print Share service group is online, you must run the wizard from a system on which the service group is online. You can then add and remove resources to the configuration using the wizard; you cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.

- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

#### To modify the Print Share service group

- 1 Start the Print Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Print Share Configuration Wizard**.
- 3 Read the information on the Welcome panel and click **Next**.
- 4 In the Wizard Options panel, choose the **Modify service group** option, select the service group to be modified, and click **Next**.
- 5 Follow the wizard instructions and make required modifications to the service group configuration.  
If you are modifying the service group to remove a PrintShare resource, make sure you offline the resource before deleting it.

### Modifying the IIS service group

The IIS configuration wizard enables you to modify an IIS service group.

- If the IIS service group is online, you must run the wizard from a system on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

#### To modify the IIS service group

- 1 Start the IIS Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center**, expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > IIS Configuration Wizard**.

- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Wizard Options panel, choose the **Modify service group** option, select the service group to be modified, and click **Next**.
- 5 Follow the wizard instructions and make required modifications to the service group configuration.

## Modifying the MSVirtual Machine service group

### To modify the MSVirtual Machine service group

- 1 Start the MSVirtual Machine Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center**, expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > MSVirtual Machine Configuration Wizard**.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Wizard Options panel, choose the **Modify service group** option, select the service group to be modified, and click **Next**.
- 5 Follow the wizard instructions and make required modifications to the service group configuration.

## Modifying any other application service group

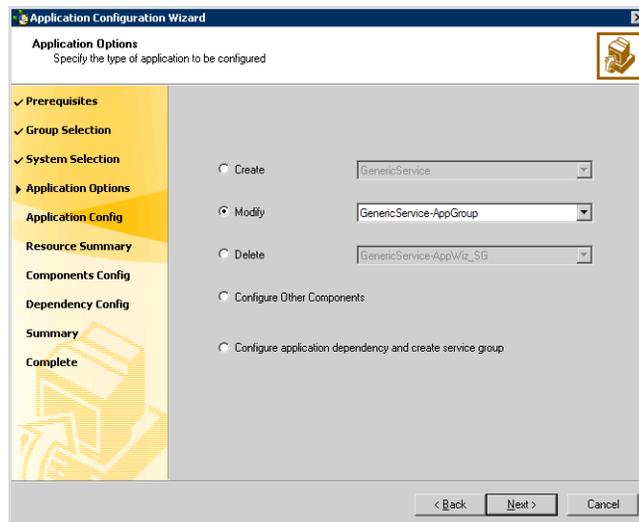
This section describes how to modify a service group using the Application Configuration Wizard. Following are some points to note before modifying the service group:

- If the service group to be modified is online, you must run the wizard from a system on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

**Note:** Symantec recommends that you do not use the wizard to modify service groups that were not created using the wizard.

### To modify a service group

- 1 Start the Application Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center**, expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Application Configuration Wizard**.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Wizard Options panel, click **Modify service group**. From the Service Groups list, select the service group containing the resource that you want to modify.
- 5 In the Service Group Configuration panel, click **Next**.
- 6 Click **Modify** and select the resource you want to modify and then click **Next**.



The **Modify** option is enabled only if:

- Service and Process resources are already configured using the wizard.
- You selected the **Modify Service Groups** option in the **Wizard Options** dialog box.

- 7 Depending on the resource you chose to modify from the Application Options page, you would either get the Generic Service Options, Process Details, or Service Monitor Options dialog box. Make required changes in the appropriate dialog box and click **Next**.
- 8 In the User Details dialog box, specify the user information and click **Next**.
- 9 In the Application Resource Summary dialog box, review the summary of the resource. Click **Back** to make changes. Otherwise, click **Next**.
- 10 The Application Options dialog box appears. Repeat [step 6](#) through [step 9](#) for each resource that you want to modify.
- 11 After modifying the required resources, you can use the wizard to:
  - Add additional resources to the service group.
  - Delete resources from the service group.
  - Add VCS components to the service group.

# Adding DMP to a clustering configuration

This chapter describes how to add Dynamic Multi-pathing (DMP) to a clustering configuration. Dynamic Multi-pathing is an optional software component in SFW that provides redundant path support for your storage. This support is provided by DMP ASLs (DMP Array Support Libraries) or DMP DSMs (DMP Device Specific Modules).

The steps for adding Dynamic Multi-pathing are given for both a new cluster configuration and an existing cluster configuration. The cluster configuration can be either VCS or MSCS. The steps for enabling Dynamic Multi-pathing are done after a cluster is up and running and tested. This chapter has the following topics:

- [“About dynamic multi-pathing”](#) on page 170
- [“Overview of configuration tasks for adding DMP ASLs or DMP DSMs”](#) on page 171
- [“Reviewing prerequisites”](#) on page 171
- [“Reviewing the configuration”](#) on page 173
- [“Steps for a new cluster configuration”](#) on page 174
- [“Steps for an existing cluster configuration”](#) on page 176

## About dynamic multi-pathing

A second host bus adapter in each computer allows redundant paths to the storage for fault tolerance purposes. The Dynamic Multi-pathing software controls the usage of the paths and allows only one path to the storage to operate at a time. However, if one path fails, the Dynamic Multi-pathing software will automatically transfer the storage to the second path.

---

**Note:** DMP ASLs cannot coexist with DMP DSMs. You cannot install both on the same cluster.

---

In installing DMP ASLs or DMP DSMs you must connect only one host bus adapter path while you are setting up the hardware and installing the DMP ASL or DMP DSM software and any other software. This is true for both a new installation and an existing installation. For DMP ASLs, at the end of the installation and configuration process you enable DMP ASLs control. If you connect the two paths before the DMP ASLs are enabled, data corruption can result. DMP DSM control is enabled automatically.

---

**Caution:** Do not connect the second path to the storage (that is, the second host bus adapter on each server) to the SAN until DMP ASLs are installed and the storage array is included under the DMP ASLs. If you allow two paths to the storage without DMP ASLs control, data can become corrupted.

---

After you install the DMP ASLs software, any arrays attached to the system come up as excluded from DMP ASLs control. Once the DMP ASLs software is running, you can include the arrays attached to the system, and only then attach a second host bus adapter to the storage arrays. For DMP DSMs, enabling DMP DSMs control is done automatically.

For the hardware setup step for each server, install the second host bus adapter in each computer, but do not connect it to the switch. After SFW and the cluster are set up and working and the DMP ASLs or DMP DSMs software is installed and running, then perform the additional steps to activate the DMP ASLs or DMP DSMs.

## Overview of configuration tasks for adding DMP ASLs or DMP DSMs

There are six tasks to add DMP ASLs or DMP DSMs on each server. The order in which the tasks are done in relation to the rest of the configuration depends on whether you are installing a new configuration or upgrading an existing cluster configuration. Detailed steps are presented in later sections. In summary, the tasks are the following:

- Install a second host bus adapter in each server. Do not connect the second path to each additional switch at this point. The path must be left unconnected.  
For a new install, this step can be done with the initial hardware configuration before the servers are running.  
For an existing cluster system, a rolling upgrade procedure is used to allow installation of the hardware and software on the inactive node on the cluster. When the installation of one node is complete, switch the active node and complete the hardware and software installation on the remaining node that is inactive.
- Install the DMP ASLs or DMP DSMs software on the inactive node. Installing DMP ASLs or DMP DSMs requires a reboot, and this avoids rebooting the active node of the cluster.
- With SFW on the server, bring up Dynamic Multi-pathing and include the disks on the storage array to bring the array under DMP ASLs control. This step is done automatically for DMP DSMs.
- Using appropriate cables, connect the second path on Server A to the second switch. Configure the switch, if necessary. Do the same for the second switch on Server B.
- Verify that both paths are under DMP ASLs or DMP DSMs control.
- Access the Array Settings dialog for each array and make sure that the array load balancing settings are set to active/passive.

## Reviewing prerequisites

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

## Supported software

Veritas Storage Foundation 5.0 for Windows (SFW) or Veritas Storage Foundation HA 5.0 for Windows (SFW HA) with the Dynamic Multi-pathing (DMP) option.

If you already have the SFW or SFW HA software installed and want to add Dynamic Multi-pathing, purchase a license key for the Dynamic Multi-pathing option and use **Add or Remove Programs** from the Windows Control Panel to add the option.

## Hardware requirements (Two-server cluster)

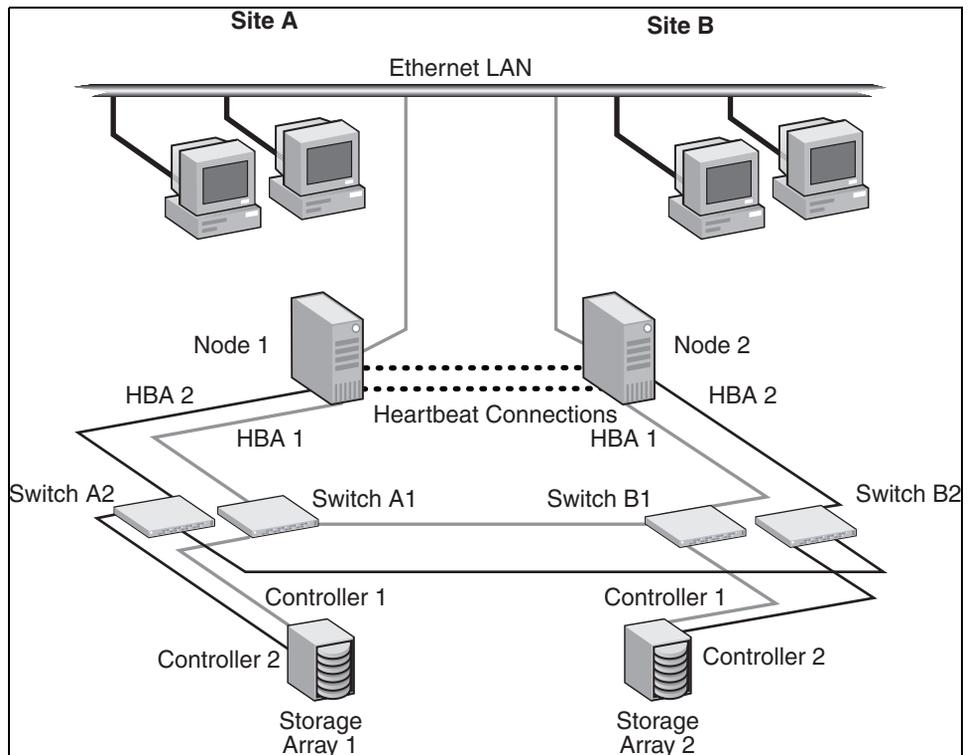
- 2 HBAs with appropriate cabling
- 2 fibre switches

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.0 Hardware Compatibility List to confirm supported hardware at: <http://entsupport.symantec.com>

## Reviewing the configuration

The purpose of the configuration is to add Dynamic Multi-pathing (DMP) to a clustering configuration. Dynamic Multi-pathing is an optional software component in SFW and SFW HA that provides redundant path support for your storage. You need to have two switches per server, one switch to accommodate each path. The diagram illustrates the four different host bus adapter paths.

**Figure 8-1** Adding fault tolerance with DMP



# Steps for a new cluster configuration

Select the procedure for DMP ASLs or DMP DSMs.

- “[DMP ASLs](#)” on page 174
- “[DMP DSMs](#)” on page 175

## DMP ASLs

Use the following procedure to set up DMP ASLs.

### To set up DMP ASLs

- 1 When the hardware is installed initially for each server, install the second host bus adapter in each computer. Leave the second path unconnected to the second switch.

---

**Caution:** Do not connect the second path (that is, the second host bus adapter on each server) to the storage connected to the SAN until DMP ASLs are installed and the storage array is included under DMP ASLs at the end of the configuration process. If you allow two paths to the storage without DMP ASLs control, data can become corrupted.

---

- 2 Follow the instructions for installing the cluster as a new install in [Chapter 7, “Deploying SFW HA for high availability: New installation”](#) for VCS and in [Chapter 15, “Deploying SFW with MSCS” on page 329](#) for MSCS, except in the sections on installing SFW HA and SFW, you will need to add the additional license key for Dynamic Multi-pathing in the License Information screen, and you will need to click the checkbox for **Dynamic Multi-pathing** in the Options screen.  
When you have tested the cluster and verified it, then you can enable DMP ASLs, as shown in the steps that follow.
- 3 With SFW on the first server, bring up Dynamic Multi-pathing and include the disks on the storage array. To include the storage array, open the VEA console to display the Array Settings screen for the storage array by doing the following:
  - a In the tree view under the Disks folder, select a disk from the storage array.
  - b In the right pane, click the **Paths** tab for the disk. Only one path should display in the **Paths** tab.
  - c Right-click the path and select **Array Settings** from the path context menu that comes up.

- d In the Array Settings dialog box, clear the **Exclude** checkbox.
- 4 Using appropriate cables, connect the second path on Server A to the second switch.
  - a Connect the remaining paths (as shown on the diagram).
  - b Complete any necessary configuration of the switch.
- 5 Go to **Actions** and select **Rescan** to verify that two paths are shown on the **Paths** tab in the GUI.
- 6 Access the Array Settings dialog for the storage array and make sure that the array load balancing settings are set to active/passive. A cluster disk requires the active/passive settings.
- 7 Complete [step 3](#) through [step 6](#) on Server B.

This completes the addition of DMP ASLs as part of a new clustering configuration.

## DMP DSMs

Use the following procedure to set up DMP DSMs.

### To set up DMP DSMs

---

**Note:** DMP DSMs cannot co-exist with DMP ASLs. Uninstall DMP ASLs before installing DMP DSMs.

---

- 1 Install additional hardware and its appropriate drivers.
- 2 Connect only one path from the array to the computer.
- 3 Install the appropriate DMP DSMs using **Add or Remove Programs** from the Windows Control Panel. For more information, see “Adding or Removing Features” in the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.
- 4 Reconnect the additional physical path.
- 5 Reboot the system.
- 6 Verify that the additional path is displayed.

## Steps for an existing cluster configuration

Select the procedure for DMP ASLs or DMP DSMs.

- “DMP ASLs” on page 176
- “DMP DSMs” on page 177

### DMP ASLs

Use the following procedure to set up DMP ASLs.

#### To set up DMP ASLs

- 1 Verify that your data is backed up before proceeding.
- 2 On the cluster node that is inactive, shut down the server and install the second host bus adapter. Leave the second path unconnected to the second switch.

---

**Caution:** Do not connect the second path (that is, the second host bus adapter on each server) to the storage connected to the SAN until DMP ASLs are installed and the storage array is included under DMP ASLs at the end of the configuration process. If you allow two paths to the storage without DMP ASLs control, data can become corrupted.

---

- 3 Reboot the server on the inactive node, access the installation media, and install the Dynamic Multi-pathing option using **Add or Remove Programs** from the Windows Control Panel. For more information, see “Adding or Removing Features” in the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.
  - a Enter the additional license key for Dynamic Multi-pathing in the License Information screen.
  - b Click the option for Dynamic Multi-pathing in the Options screen.
  - c Reboot again at the end of the installation.
- 4 On the server that has the active node of the cluster, use the **Move Group** command in MSCS to move the cluster resources to the other server. Then repeat [step 2](#) and [step 3](#) for the server that is now inactive—installing the second host bus adapter and the DMP software option, with the necessary reboots.
- 5 Move the cluster resources back to the server that had control of the cluster at the beginning of the procedure.

- 6 With SFW on the first server, bring up Dynamic Multi-pathing and include the disks on the storage array. To include the storage array, open the VEA console to display the Array Settings screen for the storage array by doing the following:
  - a In the tree view under the Disks folder, select a disk from the storage array.
  - b In the right pane, click the **Paths** tab for the disk. Only one path should display in the **Paths** tab, since the disk is not yet under DMP ASLs control.
  - c Right-click the path and select **Array Settings** from the path context menu that comes up.
  - d In the Array Settings dialog box, clear the **Exclude** checkbox.
- 7 Using appropriate cables, connect the second path on Server A to Switch A2.
  - a Connect the path through Server A HBA 2 and Controller 2 of the storage array.
  - b Complete any necessary configuration of the switch.
- 8 Go to **Actions** and select **Rescan** to verify that two paths are shown on the **Paths** tab in the GUI.
- 9 Access the Array Settings dialog for the storage array and make sure that the array load balancing settings are set to active/passive. A cluster disk requires the active/passive settings.
- 10 Complete [step 5](#) through [step 8](#) on Server B.

This completes the addition of DMP ASLs as an upgrade to an existing clustering configuration.

## DMP DSMs

Use the following procedure to set up DMP DSMs.

### To set up DMP DSMs

---

**Note:** Before installing DMP DSMs, you must remove any existing installation of DMP DSMs, regardless of version, and any device specific modules from other vendors.

---

- 1 Move resources to another node or take the resources offline.
- 2 Install additional hardware and its appropriate drivers.
- 3 Connect only one path from the array to the computer.

- 4 Install the appropriate DMP DSM using **Add or Remove Programs** from the Windows Control Panel. For more information, see “Adding or Removing Features” in the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.
- 5 Reconnect the additional physical path.
- 6 Reboot the system.
- 7 Verify that the additional path is displayed.

# Campus Clustering

This SFW solutions section describes solutions that are available within Veritas Storage Foundation™ for Windows itself. SFW solutions include mirrored striped volumes, VxCache and Automatic Volume Growth.

The section includes the following chapters:

- [Chapter 9, “Introduction to campus clustering” on page 181](#)
- [Chapter 10, “Deploying SFW HA for campus cluster” on page 185](#)



# Introduction to campus clustering

A campus cluster is a single cluster that stretches over two sites using fiber channel connectivity, with SAN connections for data mirroring and network connections for cluster communication. Although two sites are the most common, more than two can be used for additional redundancy.

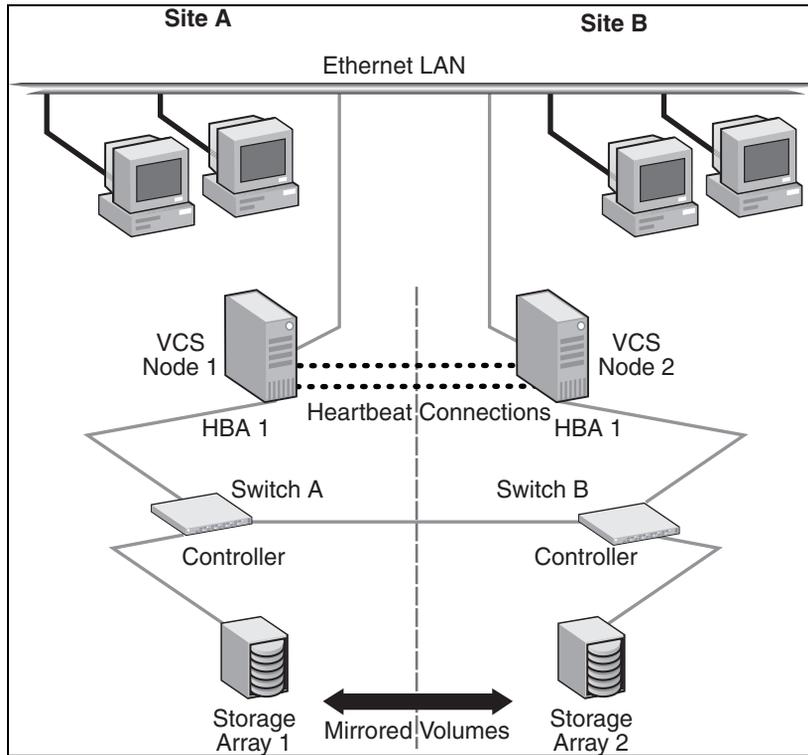
Clusters are usually located across a campus or a city but can range over much wider distances if their infrastructure supports it, using Fibre Channel SANs and long-wave optical technologies.

Campus clusters provide disaster protection when an entire site goes down by locating the clustered servers in different buildings or areas. This solution provides a level of high availability that is above mirroring or clustering at a single site and is an alternative to using replication software.

## Sample campus cluster configuration

The following sample configuration represents a campus cluster with two sites, Site A and Site B.

Figure 9-1 Typical campus clustering configuration



With SFW, a campus cluster can be set up using a Veritas Cluster Server (VCS) configuration. Both configurations involve setting up a single cluster with two nodes that are in separate buildings and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number of disks and contains mirrored data of the storage on the other array. SFW provides the mirrored storage and the disk groups that make it possible to fail over the storage by deporting the disk groups on one node and importing them on the other.

If a site failure occurs in a two-node campus cluster, the remaining cluster node will not be able to bring the cluster disk groups online because it cannot reserve a majority of disks in the disk groups. To allow for failover to the other site, a procedure forces the import to the other node, allowing a cluster disk group to be brought online on another node when that node has a minority of the cluster disks.

Implementing these force import procedures should be done with care. The primary site may appear to have failed but what really has happened is that both the storage interconnect between sites and the heartbeats have been lost. In that

case, cluster disk groups can still be online on the primary node. If a force import is done so that the data can be accessed on the secondary site, the cluster disks will be online on both sites, risking data corruption.

## Differences between campus clusters and local clusters

The procedures for setting up a campus cluster are nearly the same as those for local clusters, except that a campus cluster has the nodes located in separate buildings, so the hardware setup requires SAN interconnects that allows these connections. Also, in a campus cluster, each node has its own storage array rather than having a shared storage array between the two clusters. Both local clusters and campus clusters have SFW dynamic disk groups and volumes, but the volumes on each campus cluster node are mirrors of one another.



# Deploying SFW HA for campus cluster

This chapter presents a VCS campus clustering configuration example with SFW HA. For a campus clustering configuration with MSCS, see [Chapter 16, “Deploying SFW with MSCS in a campus cluster” on page 367](#).

The table below outlines the configuration’s high-level objectives and the tasks for each objective.

**Table 10-1** Task list

Objectives	Tasks
<a href="#">“Reviewing the requirements” on page 187</a>	<ul style="list-style-type: none"><li>■ Verify hardware and software prerequisites.</li></ul>
<a href="#">“Reviewing the configuration” on page 191</a>	<ul style="list-style-type: none"><li>■ Review the configuration requirements.</li><li>■ Overview of VCS campus cluster, and recovery scenarios</li></ul>
<a href="#">“Installing and configuring the hardware” on page 196</a>	<ul style="list-style-type: none"><li>■ Install the hardware for Site A.</li><li>■ Install the hardware in the same manner for Site B.</li><li>■ Make all the necessary hardware connections between the two cluster nodes.</li></ul>
<a href="#">“Installing Windows and configuring network settings” on page 197</a>	<ul style="list-style-type: none"><li>■ Install the operating system on both nodes.</li><li>■ Make necessary networking settings on both nodes.</li></ul>

**Table 10-1** Task list

Objectives	Tasks
<a href="#">“Installing Veritas Storage Foundation HA for Windows”</a> on page 198	<ul style="list-style-type: none"> <li>■ Set the Windows driver signing options to “Ignore” for remote nodes if you are using Windows Server 2003.</li> <li>■ Install SFW HA on both nodes with the push install. Doing so installs SFW, VCS, and the Veritas Cluster Server Enterprise Agent on both cluster nodes.</li> </ul>
<a href="#">“Configuring the cluster”</a> on page 204	<ul style="list-style-type: none"> <li>■ Use the VCS Configuration wizard to set up the cluster.</li> </ul>
<a href="#">“Creating disk groups and volumes”</a> on page 220	<ul style="list-style-type: none"> <li>■ Create dynamic cluster disk groups.</li> <li>■ Create dynamic volumes.</li> </ul>
<a href="#">“Installing the application on cluster nodes”</a> on page 229	<ul style="list-style-type: none"> <li>■ Install the application program files on the local drive of the first node.</li> <li>■ Install files relating to the data and logs on the shared storage.</li> <li>■ Deport the disk groups on the first node and import them on the second node.</li> <li>■ Install the application on the second node.</li> </ul>
<a href="#">“Creating VCS service groups”</a> on page 232	<ul style="list-style-type: none"> <li>■ Use an appropriate method to create and configure the VCS service group or groups.</li> <li>■ Bring the service group online.</li> </ul>
<a href="#">“Verifying the cluster configuration”</a> on page 233	<ul style="list-style-type: none"> <li>■ Switch the service group to the second node.</li> <li>■ Switch it back to the first node.</li> </ul>

## Reviewing the requirements

Review these product installation requirements for your systems before installation. Minimum requirements and Symantec recommended requirements may vary.

---

**Note:** Before you install the SFW or SFW HA software, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.0 Hardware compatibility List to confirm supported hardware:  
<http://entsupport.symantec.com>

---

### Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Installation on a non-system drive requires space on both the system drive and the non-system drive.

[Table 10-2](#) estimates disk space requirements for SFW HA.

**Table 10-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
SFW HA + all options + client components	1675 MB	Non-system space: 1675 MB System space: 345 MB
SFW HA + all options	1230 MB	Non-system space: 1230 MB System space: 285 MB
Client components	630 MB	Non-system space: 630 MB System space: 115 MB

### Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.0 Hardware Compatibility List to confirm supported hardware at: <http://entsupport.symantec.com>

For a Disaster Recovery configuration select the Global Clustering Option and optionally select Veritas Volume Replicator.

## Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA)
- Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)  
*or*  
Windows Server 2003 Web Edition (limited to file share support for SFW HA), Windows Server 2003 Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 for all editions)  
*or*  
Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition  
*or*  
Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP1 required for all editions)  
*or*  
Windows Server 2003 for Intel Xeon (EM64T) or AMD Opteron: Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition  
*or*  
Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

## System requirements

Systems must meet the following requirements:

- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs.  
See “[Best practices](#)” on page 190.
- 1 GB of RAM for each system.
- All servers must run the same operating system, service pack level, and system architecture.

## Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows 2000 or Windows Server 2003 domain.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.
- Static IP addresses for the following purposes:
  - One static IP address available per site for each application virtual server
  - One IP address for each physical node in the cluster
  - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
  - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

## Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

## Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.

## Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).  
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

## Campus cluster requirements

- Interconnects between the clusters are required for the storage and the network.
- The configuration requires a storage array for each site, with an equal number of disks at each site for the mirrored volumes.

---

**Note:** Plan for an equal number of disks on the two sites, because each disk group must contain the same number of disks on each site.

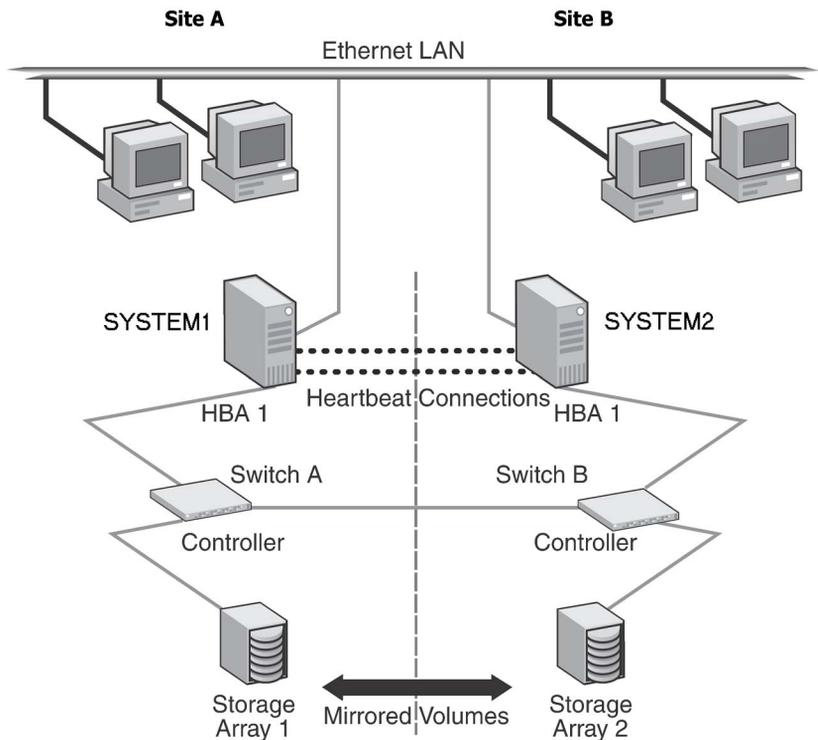
---

## Reviewing the configuration

This configuration example describes the most common configuration, a two-node campus cluster with each node at a separate site.

For an overview of campus clusters with VCS, see “[Overview of campus clustering with VCS](#)” on page 192.

**Figure 10-1** VCS 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 with an equal number of disks and contains mirrored data of the storage on the other array. The example describes a generic database application.

Plan for an equal number and size of disks on the two sites, because each disk group must contain the same number of disks on each site for the mirrored volumes.

The configuration does not include DMP. For instructions on how to add DMP to a clustering configuration, see the DMP chapter, [“Adding DMP to a clustering configuration”](#) on page 169.

## Overview of campus clustering with VCS

This overview focuses on the recovery with a VCS campus cluster. Automated recovery is handled differently in a VCS campus cluster than with a VCS local cluster.

The table below lists failure situations and the outcomes that occur with the two different settings for the ForceImport attribute of the VMDg resource. This attribute can be set to 1 (automatically forcing the import of the disk groups to the another node) or 0 (not forcing the import). Information on how to set the ForceImport attribute are given in [“Setting the ForceImport attribute”](#) on page 195.

**Table 10-3** Failure Situations

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
<b>1) Application fault</b> May mean the services stopped for an application, a NIC failed, or a database table went offline.	Application automatically moves to other site.	Service Group failover is automatic on the standby or preferred system or node.
<b>2) Server failure</b> May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.	Application automatically moves to other site. 100% of the disks are still available.	Service Group failover is automatic on the standby or preferred system or node. 100% of the mirrored disks are still available.
<b>3) Failure of disk array or all disks</b> Remaining disks in mirror are still accessible from the other site.	No interruption of service. Remaining disks in mirror are still accessible from other site.	The Service Group does not failover. 50% of the mirrored disk is still available at remaining site.
<b>4) Site failure</b> All access to the server and storage is lost.	Manual intervention required to move application. Can't import with only 50% of the disks available.	Application automatically moves to the other site.

**Table 10-3** Failure Situations

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
<b>5) Split-brain situation (loss of both heartbeats)</b> If the public network link is used as a low-priority heartbeat, it is assumed that link is also lost.	No interruption of service. Can't import disks because original site still has the SCSI reservation.	No interruption of service. Failover does not occur due to Service Group resources remaining online on the original nodes. Example: Online node has SCSI reservation to own disk.
<b>6) Storage interconnect lost</b> Fibre interconnect severed.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	No interruption of service. Service Group resources remain online, but 50% of the mirror disk becomes detached.
<b>7) Split-brain situation and storage interconnect lost</b> If a single pipe is used between buildings for the Ethernet and storage, this situation can occur.	No interruption of service. Can't import with only 50% of disks available. Disks on the same node are functioning. Mirroring is not working.	Automatically imports disks on secondary site. Now disks are online in both locations—data can be kept from only one.

## Reinstating faulted hardware

Once a failure occurs and an application is migrated to another node or site, it is important to know what will happen when the original hardware is reinstated.

For failure scenarios 3 through 7 above, the table below lists the behavior when various hardware components affecting the configuration (array or disks, site hardware, networking cards or cabling, storage interconnect, etc.) are reinstated after failure. Situations 1 and 2 have no effect when reinstated. Keep in mind

that the cluster has already responded to the initial failure as indicated in the table above.

**Table 10-4** Behavior exhibited when hardware is reinstated

<b>Failure Situation, before Reinstating the Configuration</b>	<b>ForceImport set to 0 (import not forced)</b>	<b>ForceImport set to 1 (automatic force import)</b>
<b>3) Failure of disk array or all disks</b> Remaining disks in mirror are still accessible from the other site.	No interruption of service. Resync the mirror from the remote site.	Same behavior.
<b>4) Site failure</b> All access to the server and storage is lost.	Inter-node heartbeat communication is restored and the original cluster node becomes aware that the application is online at the remote site. Resync the mirror from the remote site.	Same behavior.
<b>5) Split-brain situation (loss of both heartbeats)</b>	No interruption of service.	Same behavior.
<b>6) Storage interconnect lost</b> Fibre interconnect severed.	No interruption of service. Resync the mirror from the original site.	Same behavior.
<b>7) Split-brain situation and storage interconnect lost</b>	No interruption of service. Resync the mirror from the original site.	VCS alerts administrator that volumes are online at both sites. Resync the mirror from the copy with the latest data.

While the outcomes of using both settings of the ForceImport attribute for most scenarios are the same, the ForceImport option provides automatic failover in the event of site failure. This advantage comes at the cost of potential data loss if all storage and network communication paths between the sites are severed. Choose an option that is suitable given your cluster infrastructure, uptime requirements, and administrative capabilities.

## Setting the ForceImport attribute

After the VCS campus cluster is configured, set the ForceImport attribute. The command for implementing the force import setting in VCS is:

```
hares -modify <vmdg_resource_name> ForceImport 1|0
```

ForceImport is a flag that defines whether the agent forcibly imports the disk group when exactly half the disks are available. The value 1 indicates the agent imports the configured disk group when half the disks are available. The value 0 indicates it does not. Default is 0. This means that the disk group will be imported only when SFW acquires control over the majority of the disks.

---

**Caution:** Set this attribute to 1 only after verifying the integrity of your data. If due caution is not exercised before setting this attribute to 1, you risk potential data loss.

---

### Example

```
hares -modify vmdg_Dg1 ForceImport 1
```

Import is forced on vmdg\_Dg1.

## Installing and configuring the hardware

This topic gives the general steps for the hardware installation. For complete details on installing the hardware, refer to the hardware documentation.

### To set up the hardware

- 1 Install the hardware for Site A, using the manufacturers' instructions.
  - a Install three network interface cards.
    - One for the public network.
    - Two are recommended for the private network.Use independent hubs or switches for each VCS communication network (GAB and LLT). GAB supports hub-based or switch network paths, or two-system clusters with direct network links.

---

**Note:** To prevent lost heartbeats on the private networks and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Symantec also recommends removing TCP/IP from private NICs to lower system overhead. Contact the NIC manufacturer for details on this process.

---

- b Install the host adapter.
  - c Install the switch and the storage array.
  - d Verify that the system can access the storage devices.
- 2 Install the hardware in the same manner for Site B.
- 3 Make the necessary hardware connections to connect the two clusters together.
  - a Connect corresponding cables between the three networking cards on the two sites.
  - b Connect the two switches at the two sites through the storage interconnect.
  - c 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 Windows and configuring network settings

This section focuses on network configuration procedures that are requirements for a VCS setup. Refer to Microsoft documentation for complete information on installing the operating system and network configuration. The network configuration steps given here are for Windows Server 2003. If you are using Windows 2000, refer to the Windows 2000 documentation for the corresponding procedures.

## To install Windows and configure network settings

- 1 Install the Windows 2000 or Windows Server 2003 operating system on each node.
- 2 Configure the necessary settings for the network cards and the domain setup on each node.
- 3 Verify DNS settings for all systems on which the application will be installed.
  - a Open the Control Panel (**Start>Control Panel**).
  - b Open **Network and Dial-up Connections**.
  - c Make sure that the public network adapter is the first bound adapter:
    - From the **Advanced** menu, click **Advanced Settings**.
    - On the **Adapters and Bindings** tab, verify that 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 and click **OK**.
  - d In the Network and Dial-up Connections window, double-click the adapter for the public network.

---

**Note:** When enabling DNS name resolution, make sure that you use the public network adapters, not those configured for the VCS private network.

---

- e From the status window, click **Properties**.
- f On the **General** tab:
  - Select **Internet Protocol (TCP/IP)**.
  - Click **Properties**.
- g Select the **Use the following DNS server addresses** option.
- h Verify that the value for the IP address of the DNS server is correct.
- i Click **Advanced**.

- j On the **DNS** tab, make sure the **Register this connection's address in DNS** checkbox is selected.
  - k Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field, and click **OK**.
- 4 Use Windows Disk Management on each system to verify that the attached shared disks are visible.

## Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.0 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. For a disaster recovery configuration, select the option to install GCO. If you plan to use VVR for replication, you must also select the option to install VVR.

### Setting Windows driver signing options

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

[Table 10-5](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

**Table 10-5** 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.

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

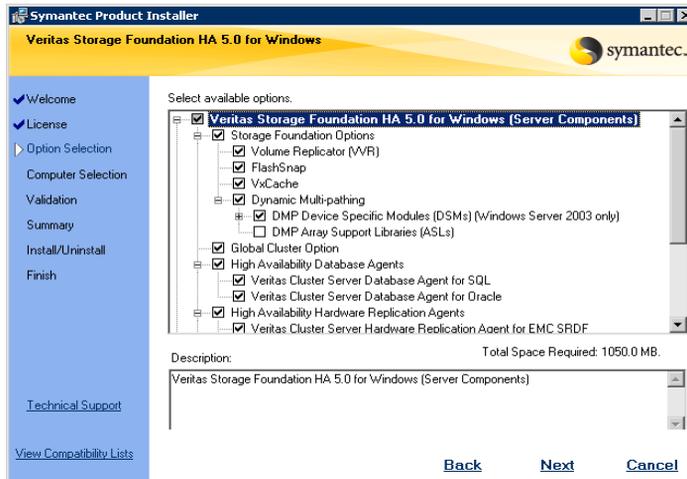
### To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.
- 3 Click **Storage Foundation HA 5.0 for Windows**.



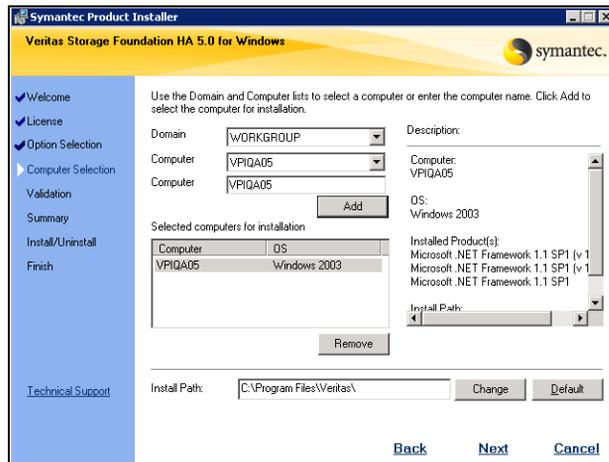
- 4 Do one of the following:

- Click **Complete/Custom** to begin installation.
  - Click the **Administrative Console** link 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**.  
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
  - 8 Repeat for additional license keys. Click **Next**
    - 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 Select the appropriate SFW product options and click **Next**.



Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console. Required to install the Solutions Configuration Center which provides information and wizards to assist configuration
Global Cluster Option	Required for a disaster recovery configuration only.
Veritas Volume Replicator	If you plan to use VVR for replication, you must also select the option to install VVR.

10 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

- 11 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.
- 12 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 list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
  - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
  - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs 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.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.  
 If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

## Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

### 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 the cluster

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

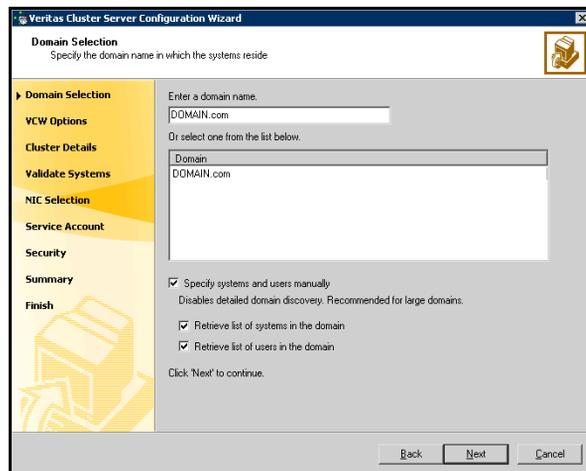
Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
  - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
  - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
  - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

### To configure a VCS cluster

- 1 Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

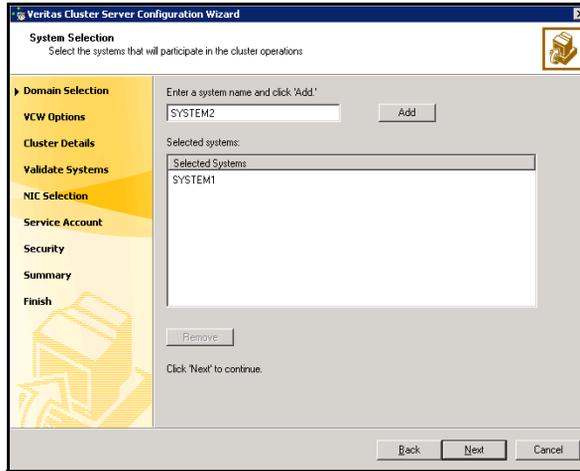
Proceed to [step 7](#) on page 207.

To specify systems and user names manually (recommended for large domains):

- Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

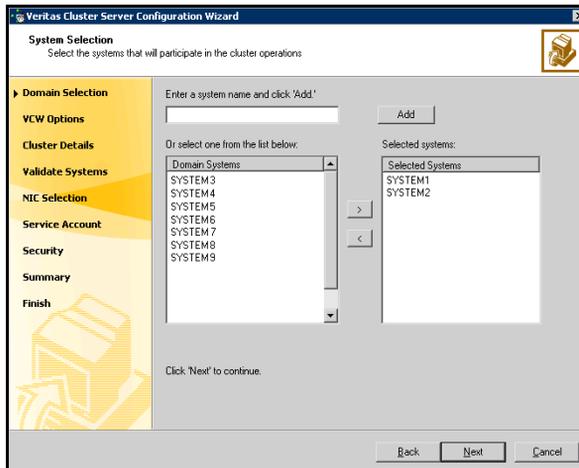
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 206. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



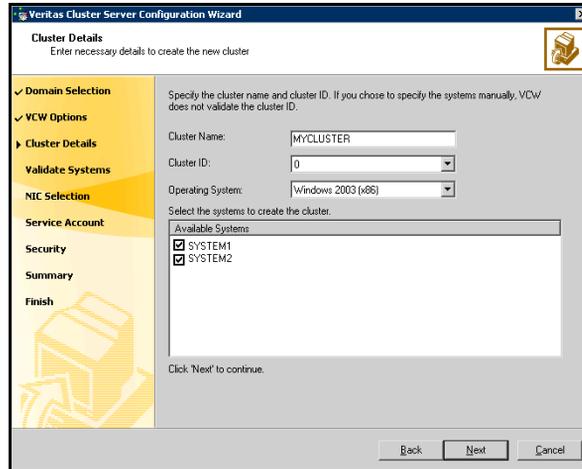
Proceed to [step 7](#) on page 207.

- 6 On the System Selection panel, specify the systems to form a cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.

- 7 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 8 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

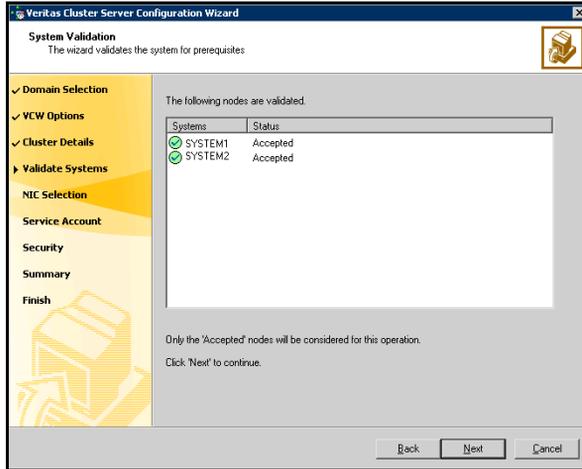
---

**Warning:** If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

---

Operating System	From the drop-down list, select the operating system that the systems are running.
Available Systems	Select the systems that will be part of the cluster. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click <b>Yes</b> to configure a private link heartbeat.

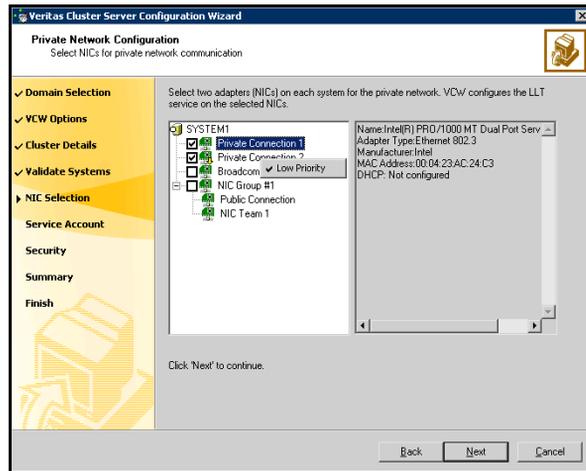
- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 8](#) on page 207, proceed to the next step. Otherwise, proceed to [step 11](#) on page 209.

- 10 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

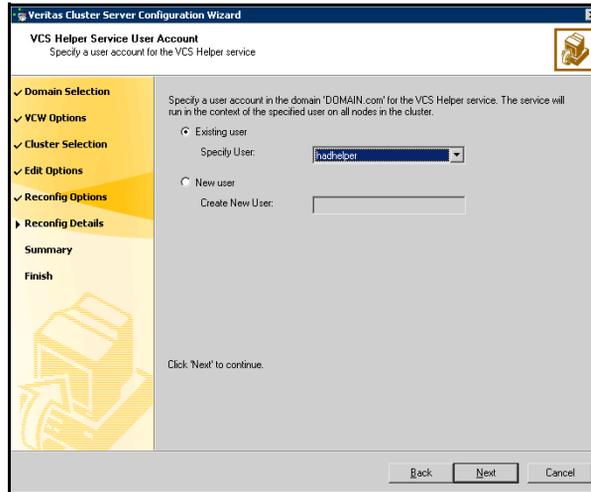


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of at least one NIC for that system. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.

If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

- 11 On the VCS Helper Service User Account panel, specify the name of a domain user context for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper

Service user context to access the network. Do not use the Administrator account for the VCS Helper service.



- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list,
  - If you chose not to retrieve the list of users in [step 4](#) on page 205, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

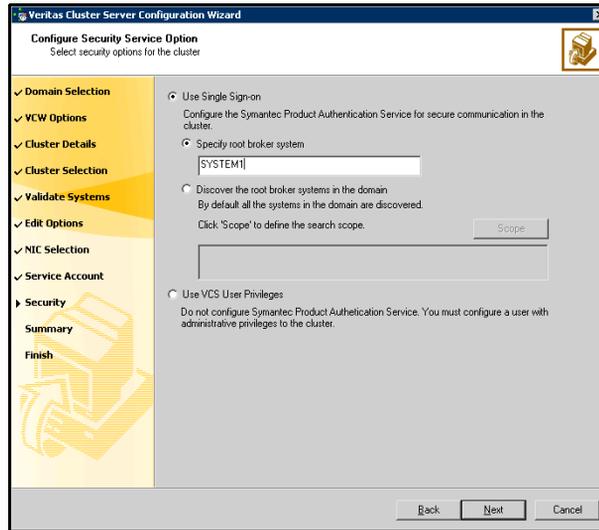
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

12 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

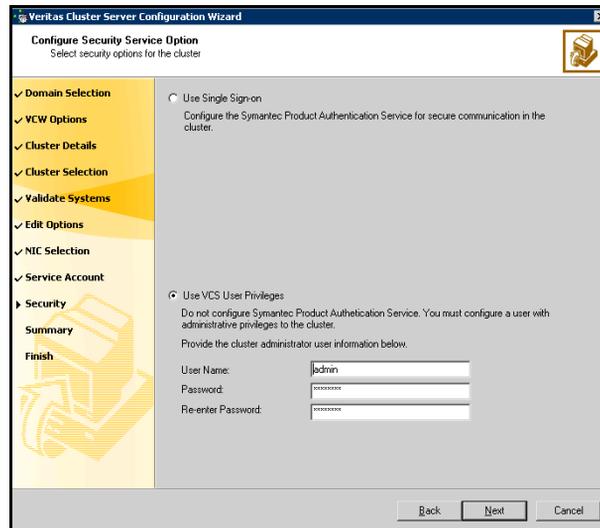
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition. For example, to search for systems managed by Administrator, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.
- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use VCS user privilege:



- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

The default user name for the VCS administrator is admin and the default password is password. Both are case-sensitive. Use this account

to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

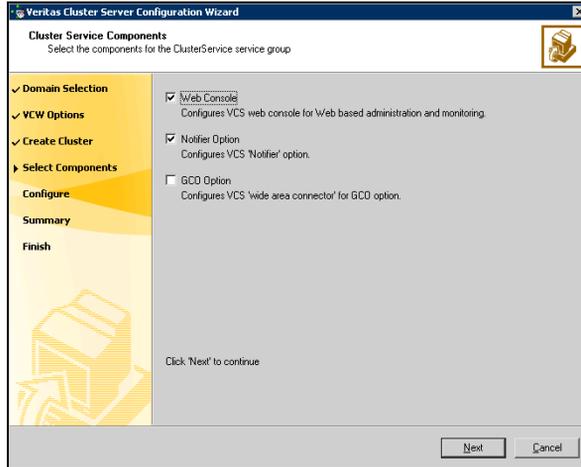
- Click **Next**.

- 13 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 14 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters. To configure the ClusterService group later, click **Finish**. At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

- 15 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.



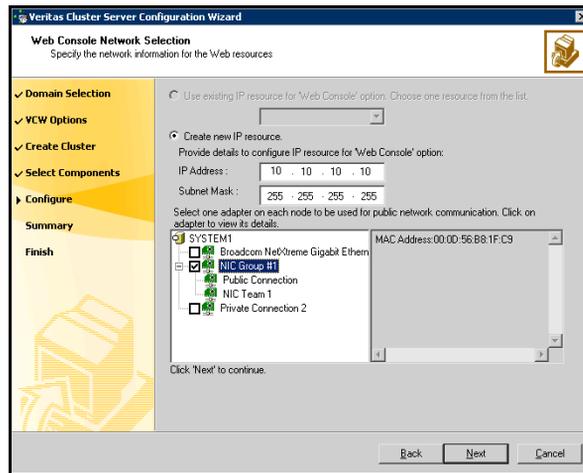
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 215.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 216.

## Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

### To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



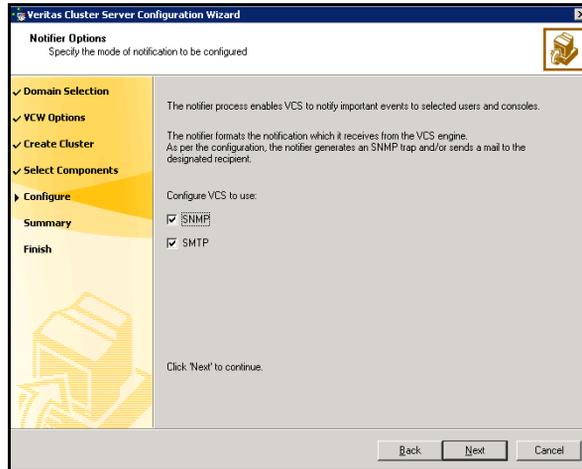
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
  - If you choose to configure a new IP address, type the IP address and associated subnet mask.
  - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
  - 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 216. Otherwise, click **Finish** to exit the wizard.

## Configuring notification

This section describes steps to configure notification.

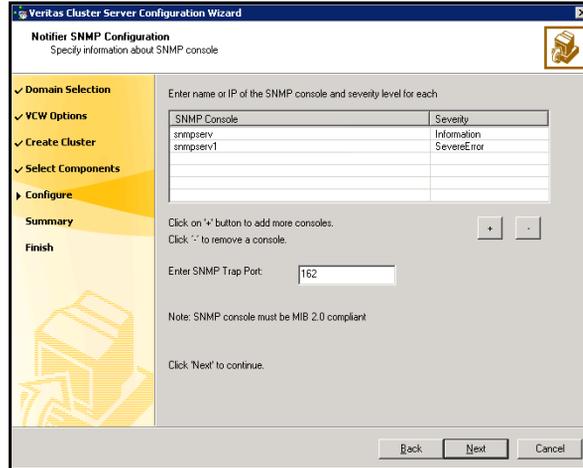
### To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.



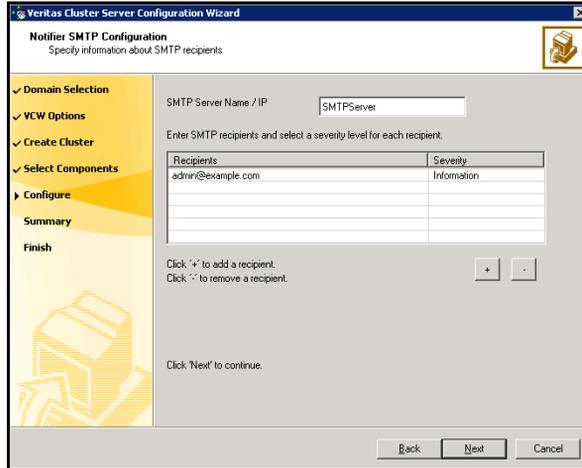
You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.



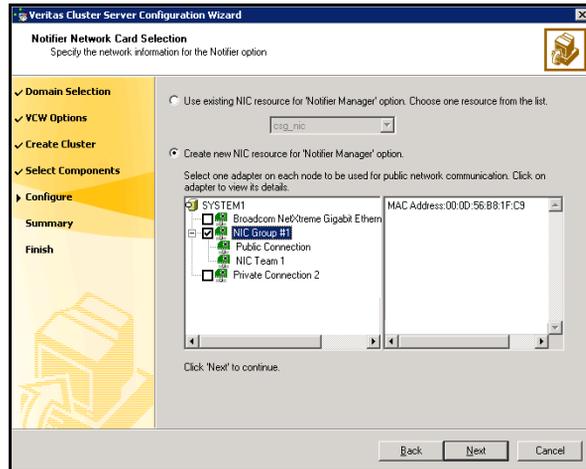
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- Enter an SNMP trap port. The default value is "162".

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
  - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
  - 6 Click **Configure**.
  - 7 Click **Finish** to exit the wizard.

## 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 two storage arrays. Create one or more dynamic cluster disk groups on the storage.

Before you create a disk group, 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 that depend on the traffic load
- The disk groups and number of disks on each site
- Types of volumes required and location of the plex of each volume in the storage array

---

**Note:** For campus clusters, each disk group must contain an equal number of disks on each site.

---

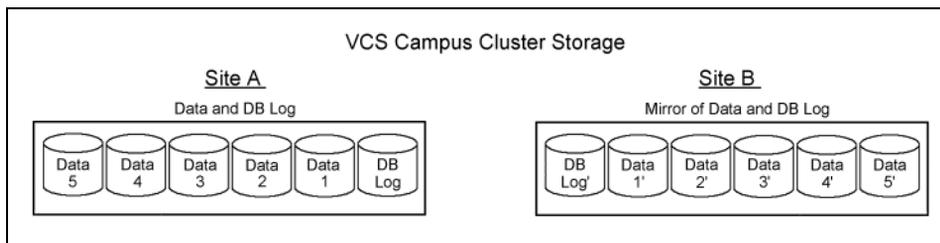
---

**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.

---

The illustration that follows shows a VCS campus cluster configuration of disks. This example has one 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 does not need to be limited to two disks, but can have four disks for greater resiliency. All the data on one site could be in one large mirrored volume with multiple disks, but this also requires the same number of disks on both sites for the mirroring. It is recommended that the log volumes be on separate disks from the data.

**Figure 10-2** VCS campus cluster disks example



## Configuring the disks and volumes

Ensure that each disk group contains an equal number of disks on each site. Each volume must 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.

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 222
- [“Creating a volume”](#) on page 223

### Considerations when creating new volumes

Consider the following when creating new volumes.

- For campus clusters, when creating 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.

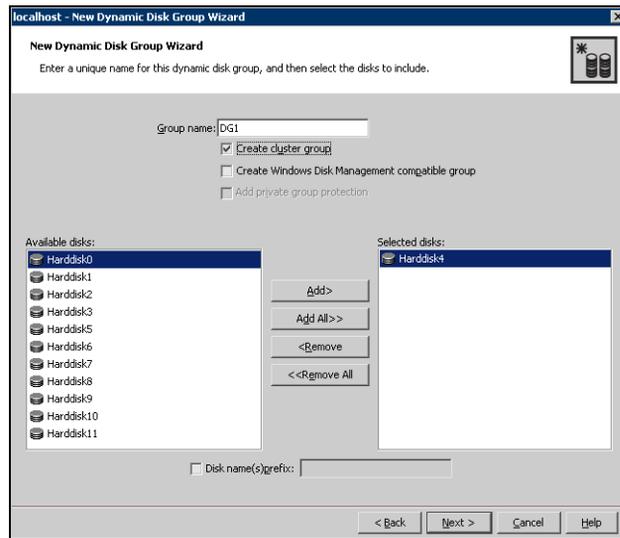
## Creating a dynamic (cluster) disk group

Use the following procedure to create a dynamic disk group.

### To create a dynamic (cluster) disk group

- 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:



- Enter the disk group name (for example, DG1).
  - 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.
  - 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

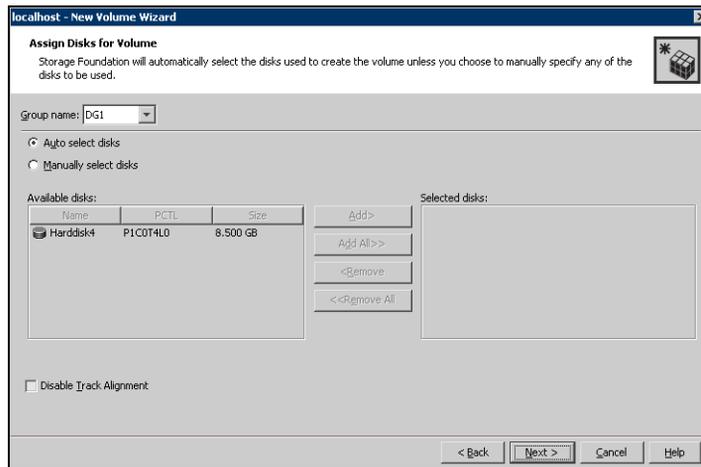
Use the following procedure to create dynamic volumes.

### 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.
- 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.

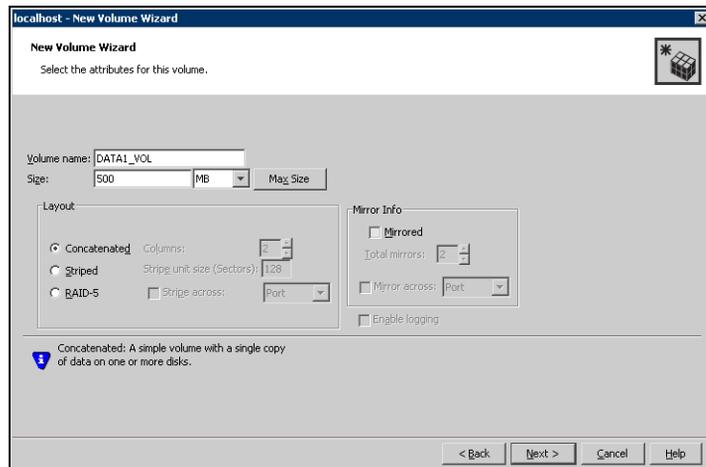


- 7 Select auto or manual disk selection and enable or disable track alignment.
  - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:

- 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.

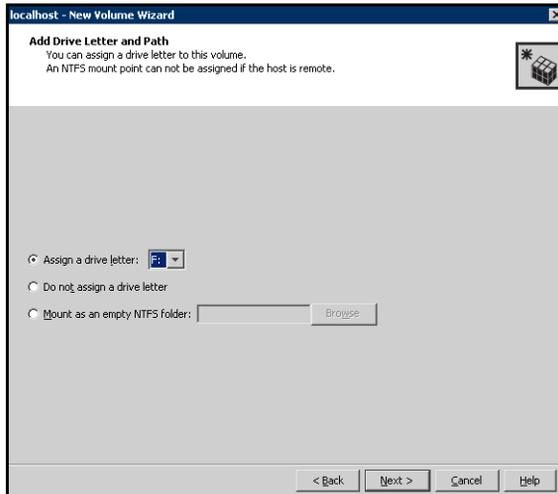
8 Click **Next**.

9 Specify the volume attributes.



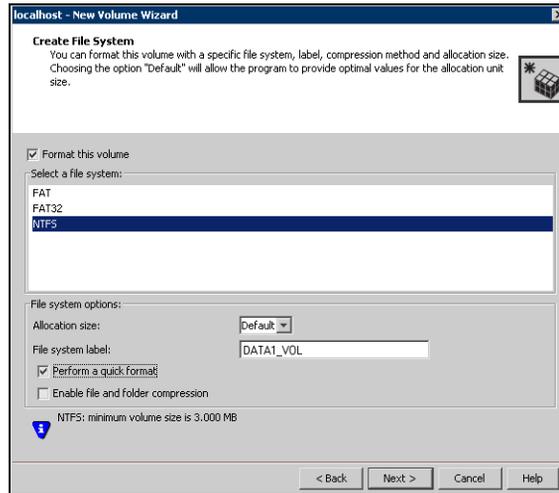
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume.

- 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.

---



# Installing the application on cluster nodes

VCS requires that the application program files be installed on the same local drive of all cluster nodes and that the application data and log files or other files related to the application data be installed on the shared storage.

## 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.
- Make sure that the disk groups and volumes are imported and thus mounted on the server before you install the application.
- If you have just created the disk groups and volumes, they will be mounted and accessible. When a disk group is created, it is automatically imported on that node. You can verify that the disk group and volumes are accessible if you can see the disk group and volume icons in VEA for the server.
- 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 when installing the application. Instead, click to browse to the dynamic volumes that were prepared previously.

## Pointers for installing the application on the second node

- To install the application on the second node, deport any disk groups from the first node and import them on the second node. Steps for deporting and importing disk groups are in the section [“Deporting and importing a disk group”](#) on page 230.
- 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 the instructions in the section [“To add or change a drive letter or mount point”](#) on page 230.
- 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.

## Deporting and importing a disk group

This section describes the steps for deporting and importing a disk group in order to install the application on the second node.

### To deport a disk group on the first node

- 1 If VEA is not already running, start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Enterprise Administrator**). If the Storage Foundation Assistant automatically opens, close it.
- 2 Navigate to **dynamic disk groups** on the node on which the dynamic disk group is currently imported.
- 3 Right-click the dynamic disk group to be deported and click **Deport**.

### To import the dynamic disk group on the second node

- 1 Start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Enterprise Administrator**). If the Storage Foundation Assistant automatically opens, close it.
- 2 Navigate to **dynamic disk groups** on the node to which you will import the dynamic disk group.
- 3 Right-click the dynamic disk group to be imported and click **Import**.  
No drive letter may be associated with an existing dynamic volume when it is imported to a computer for the first time. In such a case, use VEA to add or change drive letters. You need to make sure that drive letters or mount points for the volumes on the second node are the same as were used on the first node.

### 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.
- 3 To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.
- 4 To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Select the new drive letter and click **OK**.
- 5 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, and click **OK** to mount the volume.

---

**Note:** A mount point is also referred to as a “drive path.”

---

- 6 To change a mount point, you must remove it and recreate it ([step 5](#)). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.

## Creating VCS service groups

In order for VCS to be able to monitor and fail over an application in a cluster, the application must be included in a VCS service group.

A service group is a collection of resources working together to provide application services to clients. It can also relate to a print or a file share that does not contain a specific application. A service group's resources fail over as a group to another cluster node when there is an application failure or server failure on the active node.

VCS provides multiple methods for creating a service group. If you have Microsoft Exchange Server or SQL Server as the application, VCS provides a wizard for each of these, but you need to purchase the VCS enterprise agents for these programs. There are also separate wizards for file and print servers. In addition, there are several ways to create a service group through VCS Java Console, as well as a generic Application Configuration wizard. If you prefer to use the command line, that method can be used to create a service group as well.

Creating a VCS service group provides the following:

- Defining the cluster resources and their attributes.
- Setting their dependencies; for example, a NIC resource depends on an IP resource.
- Logically grouping the resources together.
- Providing capabilities for monitoring the service group and taking it online or offline.

For an example of installing a service group with the Application Configuration wizard, see the section “[Configuring the service group](#)” on page 104 in [Chapter 7, “Deploying SFW HA for high availability: New installation”](#).

For instructions on how to create a service group for Microsoft Exchange Server or Microsoft SQL Server, see the other Solutions Guides included with this release:

- *Veritas Storage Foundation and High Availability Solutions, Solutions Guide for Microsoft Exchange*
- *Veritas Storage Foundation and High Availability Solutions, Solutions Guide for Microsoft SQL*

To install a service group for a file or print server, or by using the Java Console, refer to the *Veritas Cluster Server Administrator's Guide*.

# 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 Veritas Cluster Manager (Java Console) to switch all the service groups from one node to another.
- Simulate a local cluster failover by shutting down an active cluster node.

## To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
  - Click **Switch To**, and click the appropriate node from the menu.
  - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.  
If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

## To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node:
  - Restart the node you shut down in [step 1](#).
  - Click **Switch To**, and click the appropriate node from the menu.
  - In the dialog box, click **Yes**.  
The service group you selected is taken offline and brought online on the node that you selected.



# Disaster Recovery

This section includes the following chapters:

- [Disaster recovery: Overview](#)
- [Deploying disaster recovery: New application installation](#)
- [Testing fault readiness by running a fire drill](#)



# Disaster recovery: Overview

Topics in this section include:

- [“About a disaster recovery solution”](#) on page 238
- [“Need for implementing a disaster recovery solution”](#) on page 240
- [“Overview of the recovery process”](#) on page 241
- [“Components of VVR that enable disaster recovery”](#) on page 242

## About a disaster recovery solution

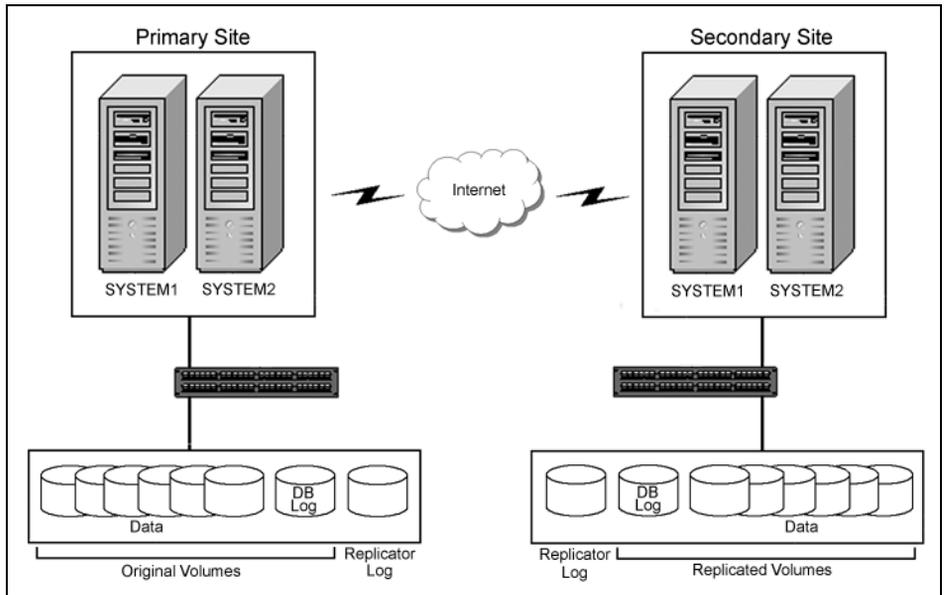
A disaster recovery (DR) solution is a series of procedures used to safely and efficiently restore application data and services in the event of a catastrophic failure. A typical disaster recovery solution requires that you have a source host on the *primary* site and a destination host on the *secondary* site. The application data is stored on the primary site and replicated to the secondary site by using a tool such as the Veritas Volume Replicator. The primary site provides data and services during normal operation. If a disaster occurs on the primary site and its data is destroyed, a secondary host can take over the role of the primary host to make the data accessible. The application can be restarted on that host.

This chapter is an overview of the Veritas Volume Replicator disaster recovery solution that can be used with SFW HA and VCS. For a SFW and VVR configuration with MSCS, see [Chapter 17, “Deploying SFW and VVR with MSCS” on page 415](#).

The example configuration in this section includes a SFW HA-VVR configuration with VCS configuration. Both configurations are described with a generic database application that includes both data and a database log.

The illustration below shows the SFW HA-VVR configuration with VCS. The example has one disk group on each site for the application. Note that a VVR Replicator Log is needed on each site. If there is more than one disk group, an additional Replicator Log is required for each disk group.

Figure 11-1 SFW HA-VVR configuration with VCS



## Need for implementing a disaster recovery solution

Two major trends affecting businesses today are reliance on data and geographic distribution. Continuous, consistent, fast, and reliable access to data is important. If a disaster occurs, quick availability of data becomes important. One of the ways of achieving this is by using replication.

A well-designed disaster recovery solution prepares a business for unexpected disasters and provides the following benefits in the event of a disaster:

- Minimizes economic loss due to the unavailability or loss of data.
- Ensures safe and efficient recovery of data and services.
- Minimizes decision making during the disaster recovery.
- Reduces reliance on key individuals.
- Minimizes data loss during recovery and ensures availability of the most recent data.

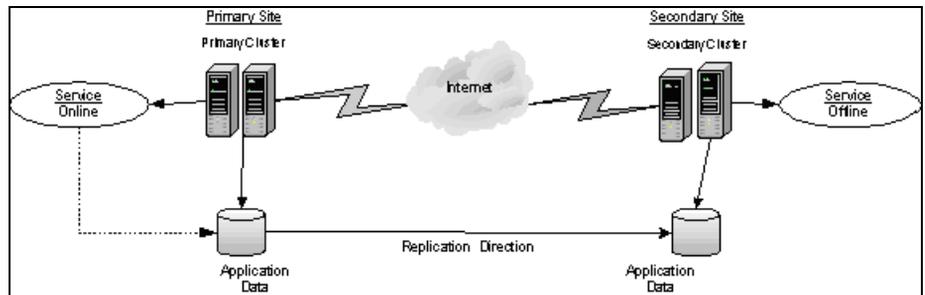
A strategic disaster recovery (DR) solution can provide businesses with ways to meet their service level agreements, comply with government regulations, and minimize their business risk.

## Overview of the recovery process

The illustrations that follow show the typical disaster recovery setup before and after a disaster.

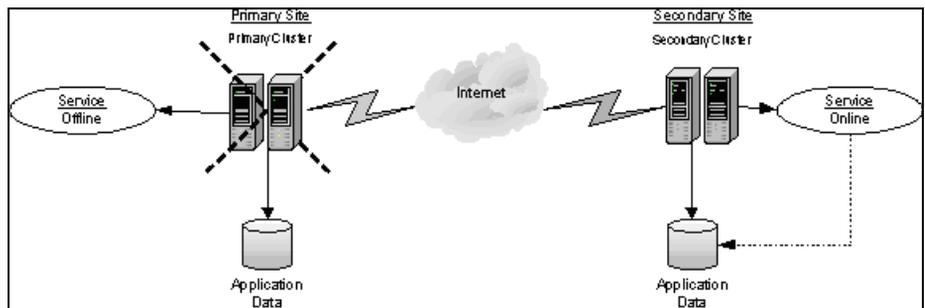
In the illustration before the disaster, the primary host replicates its application data to the secondary host. In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. Note that the primary and the secondary sites have clusters to make both the application and VVR highly available.

**Figure 11-2** Typical disaster recovery configuration setup



If a disaster, such as an earthquake, causes a failure at the primary site, a host on the secondary site can take over the role of the primary host to make the data accessible and restore the application services and data to users.

**Figure 11-3** Recovery situation after a disaster occurs



# Components of VVR that enable disaster recovery

This topic provides information about components of VVR that make the disaster recovery solution work.

## Understanding replication

The term “replication” generally refers to the use of a tool or service, or a combination of tools or services, to automate the process of regularly placing an up-to-date copy of data from a designated source, or primary, to one or more remote locations.

Replication can be used to provide solutions to problems in a variety of application environments. Any application that needs redundancy at multiple sites or can achieve better performance through geographic distribution can benefit from replication. Redundancy at multiple sites, where updates to the primary site are immediately reflected at remote sites, can be effectively used to manage disaster recovery with the use of a replication tool.

Veritas Veritas Storage Foundation™ for Windows (VVR) is a data replication service that helps to maintain a consistent copy of the application data at a remote site. It is built to contribute to an effective disaster recovery plan. If the primary data center is destroyed, the application data is readily available at the remote site, and the application can be restarted at the remote site. VVR works as an integrated component of Veritas Storage Foundation for Windows. Any application, even with existing data, can be configured to use VVR transparently. For more information on VVR, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator’s Guide*.

## Modes of replication

VVR replicates in synchronous, asynchronous, and synchronous override modes.

### Synchronous replication

The synchronous mode ensures that an update has been acknowledged by the secondary host before completing the update at the primary site. Thus, the primary site and the secondary site have the same data. If a disaster occurs on the primary site and its data is destroyed, the secondary site will already have an up-to-date copy of the data.

The synchronous mode of replication is most effective in application environments that have lower update rates but require all the hosts to always reflect the same data, or where a delay in updates between the primary and secondary hosts is not acceptable.

## Asynchronous replication

In the asynchronous mode of replication, the application updates are immediately reflected at the primary site and sent to the secondary site as soon as possible. The updates are stored in the Replicator Log until they are sent to the secondary site. This allows asynchronous replication to deal with temporary network or secondary host failures without affecting the performance of the application.

Asynchronous replication mode is most effective in application environments where it is not acceptable for the application performance to be impacted, only a minimal data loss can be tolerated, or the application has a high rate of updates.

## Synchronous override replication

The synchronous override mode of replication provides synchronous replication, as long as the network is available. If the network becomes unavailable, replication is continued in asynchronous mode.

The synchronous override replication mode is most effective in application environments where it is not acceptable for the primary site to be affected by a network failure.

---

**Note:** For additional information, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

---

## Features of VVR that help in disaster recovery

While many of the components described above are replicated at the disaster recovery site through conventional means, VVR solves the difficult problem of replicating the *user* database. Refer to the following information on how VVR helps with disaster recovery in any application environment:

- **Write Order Fidelity:** VVR guarantees that changes made to data on the primary host are made in the same sequence on the secondary host. This ensures that the data remains in a consistent state in the event of a disaster.
- **Synchronous Replication:** VVR guarantees that changes committed on the primary host are committed on the secondary host first. This ensures that the data on the secondary host matches the data on the primary host and minimizes data loss in the event of a disaster.
- **Asynchronous Replication:** VVR reflects the changes to the application immediately on the primary, and changes are then reflected on the secondary as soon as possible. Until the data is sent to the secondary, it is stored on the Replicator Log.

**Components of VVR that enable disaster recovery**

- **RVG Snapshot:** This provides the ability within VVR to take a point-in-time snapshot of a volume. This allows verification of the consistency of the data on the secondary host without impacting replication between the primary and secondary hosts.
- **Heterogeneous Storage Support:** VVR provides a replication technology that works with heterogeneous storage hardware. VVR allows replication to occur between similar or dissimilar storage arrays from a vendor or between different storage arrays from different vendors. This allows for maximum use of existing hardware and provides flexibility when adding new hardware.

# Deploying disaster recovery: New application installation

This chapter covers the following topics:

- Tasks for a new disaster recovery installation— additional applications
- Before you begin
- Setting up the secondary site: Configuring SFW HA and setting up a cluster
- Verifying that your application or server role is configured for HA at the primary site
- Configuring the VVR security service
- Configuring disaster recovery
- Assigning user privileges (secure clusters only)
- Cloning the storage on the secondary site using the DR wizard
- Installing and configuring the application or server role
- Cloning the service group configuration from the primary to the secondary site
- Configuring replication and global clustering
- Verifying the disaster recovery configuration
- Establishing secure communication within the global cluster (optional)
- Possible task after creating the DR environment: Adding a new failover node
- Maintaining: Normal operations and recovery procedures

## Tasks for a new disaster recovery installation— additional applications

This chapter provides the steps for setting up a disaster recovery (DR) solution, using SFW HA with the Veritas Volume Replicator (VVR) and Global Cluster Option (GCO) in a new installation. The chapter describes the process for any generic application or applications such as FileShare, PrintShare, IIS and MSVirtual Machines.

For examples of the SFW HA disaster recovery solution with specific applications, see the other Solutions Guides included with this release: *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL* and *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange*.

The steps for setting up the cluster described in the High Availability section of this guide are the basic foundation for this disaster recovery solution.

See [Chapter 7, “Deploying SFW HA for high availability: New installation”](#) on page 63.

After you complete configuring high availability on the primary site, you install and configure the high availability and application components on the secondary site, with the intent of creating an identical setup for the application service group on both sites. This environment involves an active/passive configuration with one to one failover capabilities.

The identical configuration can be achieved using the Disaster Recovery (DR) wizard. The DR wizard helps you to clone the storage configuration and the service group configuration from the primary site to the secondary site. The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [“Using the Solutions Configuration Center”](#) on page 23.

---

**Note:** If you want to create the identical configuration manually, without cloning the storage configuration or the service group, see [Appendix A, “Deploying Disaster Recovery: Manual implementation”](#) on page 487.

---

You can either choose to configure replication using VVR or an agent-supported array-based hardware replication and then use the DR wizard to configure global clustering. If you plan to use array-based hardware replication, then you must first complete configuring global clustering using the DR wizard and then proceed with configuring replication. Irrespective of the method you choose for

replication, you must set up Global Clustering to complete the disaster recovery configuration.

---

**Caution:** To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment, you must first run the wizard before configuring replication.

---

The main differences in the process of setting up the cluster for a disaster recovery, rather than for HA alone, are that you need to make sure that the VVR and the GCO options are selected during the SFW HA installation. You also need to configure the Veritas Volume Replicator Security Service (VxSAS) after the installation completes. For the secondary site, the cluster is set up in a similar manner as on the primary site.

[Table 12-1](#) outlines the high-level objectives and the tasks to complete each objective:

**Table 12-1** Task list

Objective	Tasks
<a href="#">“Before you begin”</a> on page 249	<ul style="list-style-type: none"> <li>■ Verifying hardware and software prerequisites</li> </ul>
<a href="#">“Reviewing the configuration”</a> on page 252	<ul style="list-style-type: none"> <li>■ Understanding Active/Passive configuration and site failover in a DR environment</li> </ul>
<a href="#">“Configuring the storage hardware and network”</a> on page 254	<ul style="list-style-type: none"> <li>■ Setting up the network and storage for a cluster environment</li> <li>■ Verifying the DNS entries for the systems on which the application will be installed</li> </ul>
<a href="#">“Setting up the secondary site: Configuring SFW HA and setting up a cluster”</a> on page 257	<ul style="list-style-type: none"> <li>■ Reviewing the prerequisites</li> <li>■ Reviewing the configuration</li> <li>■ Configuring the network and storage</li> <li>■ Installing SFW HA</li> <li>■ Configuring the cluster using the Veritas Cluster Server Configuration Wizard</li> </ul>
<a href="#">“Verifying that your application or server role is configured for HA at the primary site”</a> on page 278	<ul style="list-style-type: none"> <li>■ Verifying that the application has been configured for high availability at the Primary site</li> </ul>
<a href="#">“Configuring the VVR security service”</a> on page 278	<ul style="list-style-type: none"> <li>■ Configuring the VVR security service</li> </ul>

Objective	Tasks
“ <a href="#">Assigning user privileges (secure clusters only)</a> ” on page 282	<ul style="list-style-type: none"> <li>■ For secure clusters only, assigning user privileges</li> </ul>
“ <a href="#">Cloning the storage on the secondary site using the DR wizard</a> ” on page 283	<ul style="list-style-type: none"> <li>■ Cloning the storage configuration on the secondary</li> </ul>
“ <a href="#">Installing and configuring the application or server role</a> ” on page 287	<ul style="list-style-type: none"> <li>■ Reviewing the prerequisite checklist</li> <li>■ Installing the application</li> </ul>
“ <a href="#">Cloning the service group configuration from the primary to the secondary site</a> ” on page 289	<ul style="list-style-type: none"> <li>■ Cloning the service group configuration from the primary to the secondary site using the DR wizard</li> </ul>
“ <a href="#">Configuring replication and global clustering</a> ” on page 292	<ul style="list-style-type: none"> <li>■ Configuring VVR components and global clustering using the DR wizard</li> </ul>
“ <a href="#">Verifying the disaster recovery configuration</a> ” on page 298	<ul style="list-style-type: none"> <li>■ Reviewing required tasks when adding a new failover system to either the primary or secondary site</li> </ul>
“ <a href="#">Establishing secure communication within the global cluster (optional)</a> ” on page 300	<ul style="list-style-type: none"> <li>■ Adding secure communication between local clusters within the global cluster (optional task)</li> </ul>
“ <a href="#">Maintaining: Normal operations and recovery procedures</a> ” on page 305	<ul style="list-style-type: none"> <li>■ Monitor replication</li> <li>■ Perform planned migration</li> <li>■ Complete the recovery procedures after the primary site goes down</li> </ul>

## Before you begin

This disaster recovery solution requires a primary site and secondary site.

Before you begin setting up the secondary site for disaster recovery, you must first complete setting up the primary site for high availability.

See [Chapter 7, “Deploying SFW HA for high availability: New installation”](#) on page 63.

Review these product installation requirements for your systems before installation. Minimum requirements and Symantec recommended requirements may vary.

## Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Installation on a non-system drive requires space on both the system drive and the non-system drive.

[Table 12-2](#) estimates disk space requirements for SFW HA.

**Table 12-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
SFW HA + all options + client components	1675 MB	Non-system space: 1675 MB System space: 345 MB
SFW HA + all options	1230 MB	Non-system space: 1230 MB System space: 285 MB
Client components	630 MB	Non-system space: 630 MB System space: 115 MB

## Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.0 Hardware Compatibility List to confirm supported hardware at: <http://entsupport.symantec.com>

For a Disaster Recovery configuration select the Global Clustering Option and optionally select Veritas Volume Replicator.

## Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA)
- Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)  
or  
Windows Server 2003 Web Edition (limited to file share support for SFW HA), Windows Server 2003 Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 for all editions)  
or  
Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition  
or  
Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP1 required for all editions)  
or  
Windows Server 2003 for Intel Xeon (EM64T) or AMD Opteron: Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition  
or  
Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

## System requirements

Systems must meet the following requirements:

- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs.  
See [“Best practices”](#) on page 252.
- 1 GB of RAM for each system.
- All servers must run the same operating system, service pack level, and system architecture.

## Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows 2000 or Windows Server 2003 domain.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.
- Static IP addresses for the following purposes:
  - One static IP address available per site for each application virtual server
  - One IP address for each physical node in the cluster
  - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
  - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

## Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

## Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.

## Best practices

Symantec recommends that you perform the following tasks:

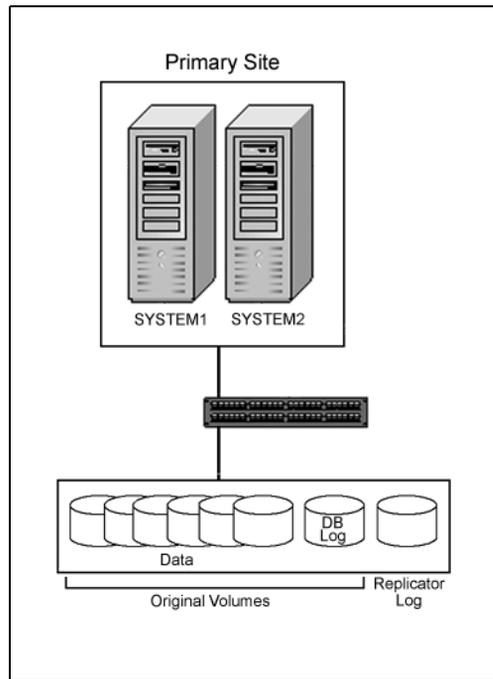
- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).  
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

## Reviewing the configuration

This configuration overview describes active/passive high availability within a cluster and disaster recovery between two sites. In an active/passive configuration, one or more application virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

Active/passive clusters involve one-to-one failover capabilities. For instance, if you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM5 and SYSTEM6 on the secondary site), then SYSTEM1 can fail over to SYSTEM2, and SYSTEM5 can fail over to SYSTEM6. The figure that follows illustrates the cluster configuration on the primary site. For a view of the DR configuration that includes both sites, see the illustration in the section “[About a disaster recovery solution](#)” on page 238.

**Figure 12-1** DR configuration primary site



## Supported disaster recovery configurations for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

Service group dependency configurations are described in detail in the VCS documentation.

See *Veritas Cluster Server Administrator's Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

The Disaster Recovery wizard supports only one level of dependency (one child). If you need to configure more levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

The wizard clones dependent service groups as global groups.

## 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 VCS 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.
  - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk.  
Use Windows Disk Management (**Start > Control Panel > Administrative Tools > Computer Management**) or Veritas Enterprise Administrator (VEA) on each system to verify 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 Right-click on 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 and Dial-up 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 VCS private network.
  - 5 From the status window, click **Properties**.
  - 6 In 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**.

## Managing disk groups and volumes

This section includes the following procedures:

- Importing a disk group and mounting a shared volume
- Unmounting a volume and deporting a disk group

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**.

# Setting up the secondary site: Configuring SFW HA and setting up a cluster

Before you begin setting up the secondary site for disaster recovery, you must first complete setting up the primary site for high availability.

See [Chapter 7, “Deploying SFW HA for high availability: New installation”](#) on page 63.

After completing the high availability configuration on the primary site, you repeat the appropriate tasks to complete the SFW HA installation at the secondary site.

Because the Disaster Recovery wizard is capable of cloning the storage, you must complete configuring SFW HA at the secondary site. Begin with reviewing the requirements on the secondary site, similar to the primary site:

- Reviewing the requirements  
 See “[Before you begin](#)” on page 249.

Then continue with the procedures given below.

## Installing SFW HA

The product installer enables you to install the software for Veritas Storage Foundation HA 5.0 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. For a disaster recovery configuration, select the option to install GCO. If you plan to use VVR for replication, you must also select the option to install VVR.

## Setting Windows driver signing options

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

[Table 12-3](#) 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

**Table 12-3** Installation behavior with unsigned drivers (Continued)

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.

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

**To install the product**

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

3 Click **Storage Foundation HA 5.0 for Windows**.



4 Do one of the following:

- Click **Complete/Custom** to begin installation.
- Click the **Administrative Console** link 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**.

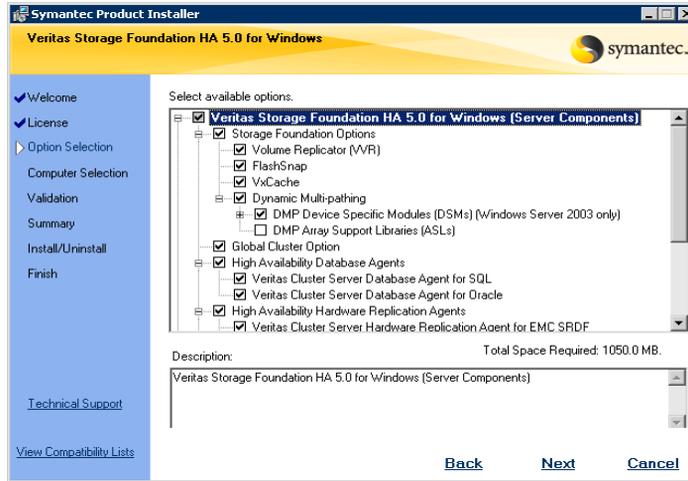
7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.

If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.

8 Repeat for additional license keys. Click **Next**

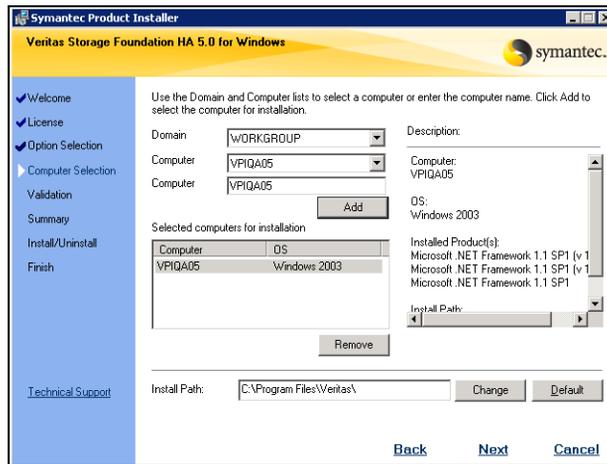
- To remove a license key, click the key to select it and click **Remove**.
- To see the license key's details, click the key.

9 Select the appropriate SFW product options and click **Next**.



- |                           |   |
|---------------------------|---|
| Client                    | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console.<br>Required to install the Solutions Configuration Center which provides information and wizards to assist configuration |
| Global Cluster Option     | Required for a disaster recovery configuration only.  |
| Veritas Volume Replicator | If you plan to use VVR for replication, you must also select the option to install VVR.   |

**10** Select the domain and the computers for the installation and click **Next**.



- |              |  |
|--------------|--|
| Domain       | Select a domain from the list.<br><br>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 <b>Add</b> .<br><br>To remove a computer after adding it, click the name in the Selected computers for installation field and click <b>Remove</b> .<br><br>Click a computer's name to see its description.  |
| Install Path | Optionally, change the installation path. <ul style="list-style-type: none"> <li>■ To change the path, select a computer in the Selected computers for installation field, type the new path, and click <b>Change</b>.</li> <li>■ To restore the default path, select a computer and click <b>Default</b>.</li> </ul> The default path is:<br>C:\Program Files\Veritas<br>For 64-bit installations, the default path is:<br>C:\Program Files (x86)\Veritas |

**11** 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.

- 12 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 list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
  - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
  - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs 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.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.  
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

## Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

### 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 the cluster

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

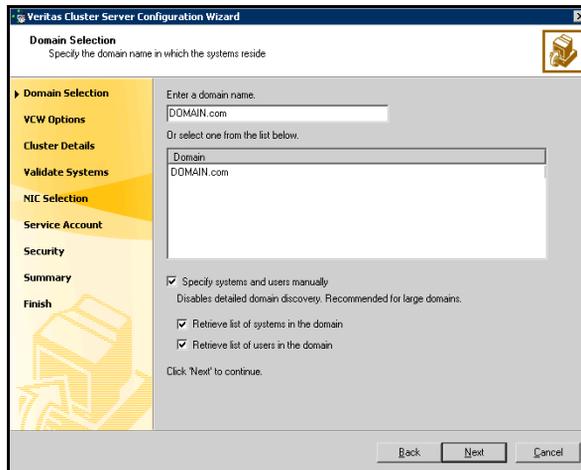
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
  - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
  - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
  - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.

#### To configure a VCS cluster

- 1 Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

Proceed to [step 7](#) on page 266.

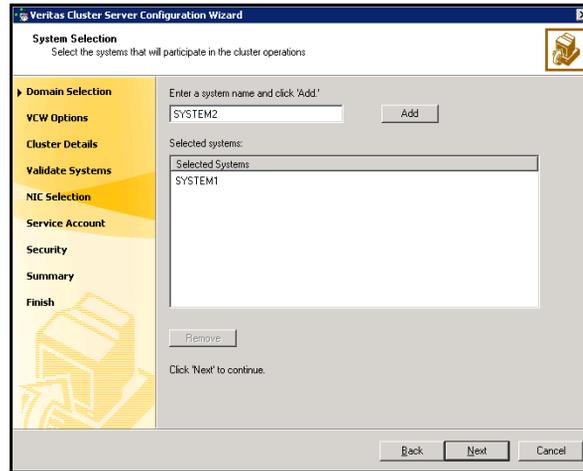
To specify systems and user names manually (recommended for large domains):

- Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.

- Click **Next**.

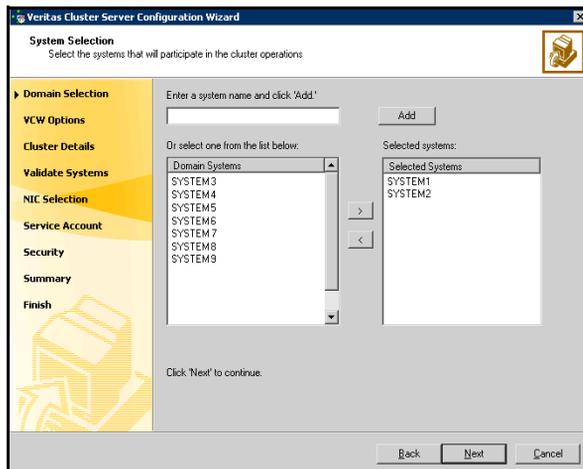
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 265. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



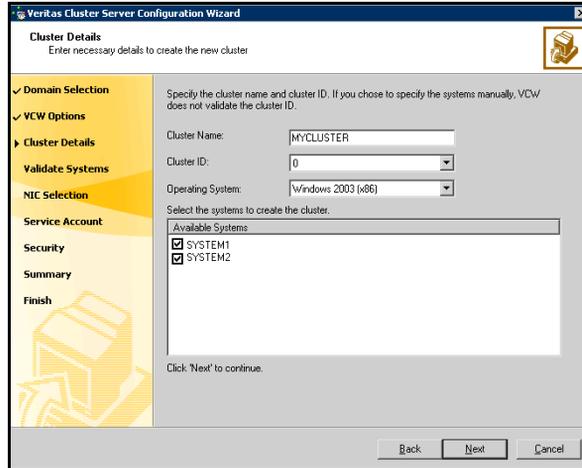
Proceed to [step 7](#) on page 266.

- 6 On the System Selection panel, specify the systems to form a cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.

- 7 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 8 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



**Cluster Name** Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

**Cluster ID** Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

---

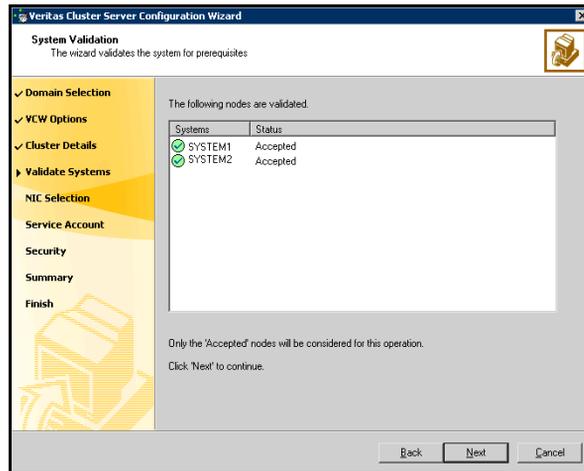
**Warning:** If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

---

**Operating System** From the drop-down list, select the operating system that the systems are running.

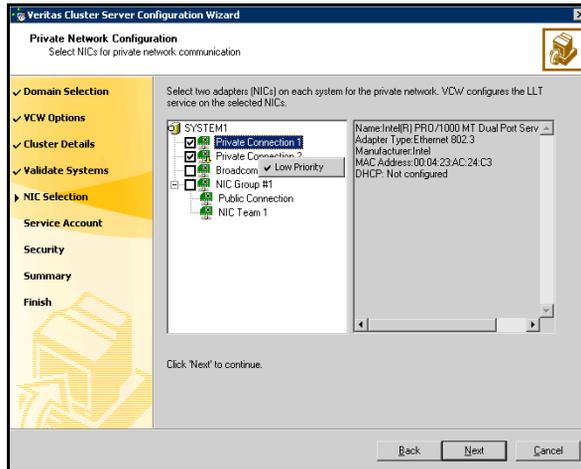
**Available Systems** Select the systems that will be part of the cluster. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.  
If you chose to configure a private link heartbeat in [step 8](#) on page 266, proceed to the next step. Otherwise, proceed to [step 11](#) on page 268.

- 10 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

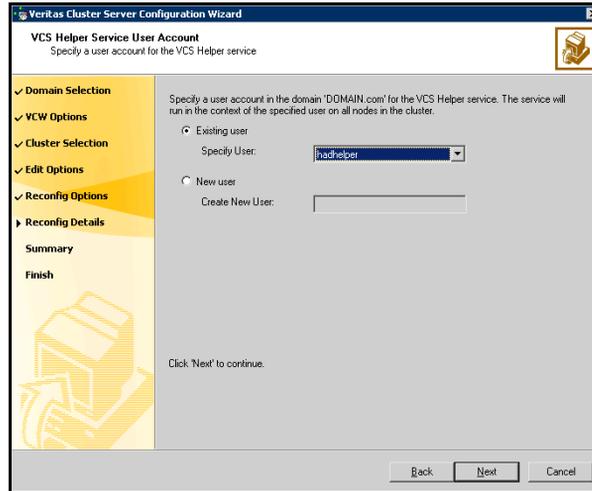


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of at least one NIC for that system. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.

If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

- 11 On the VCS Helper Service User Account panel, specify the name of a domain user context for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper

Service user context to access the network. Do not use the Administrator account for the VCS Helper service.



- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list,
  - If you chose not to retrieve the list of users in [step 4](#) on page 264, type the user name in the **Specify User** field, and then click **Next**.
  - To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

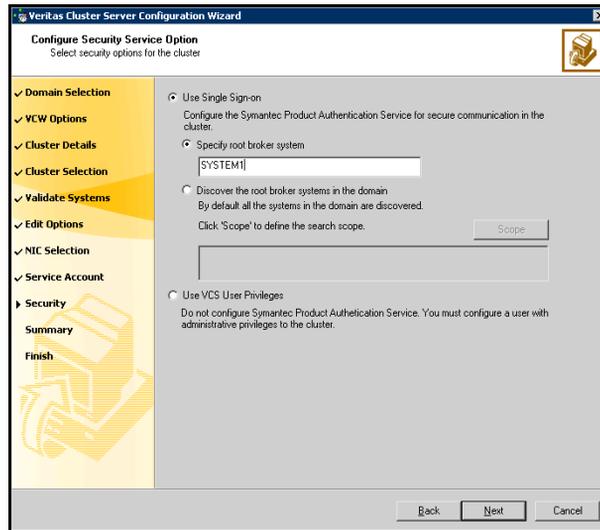
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

**12** On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

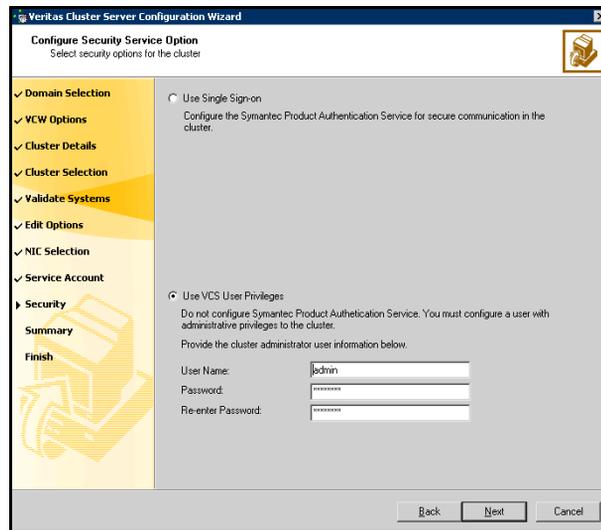
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition. For example, to search for systems managed by Administrator, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.
- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use VCS user privilege:



- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

The default user name for the VCS administrator is admin and the default password is password. Both are case-sensitive. Use this account

to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

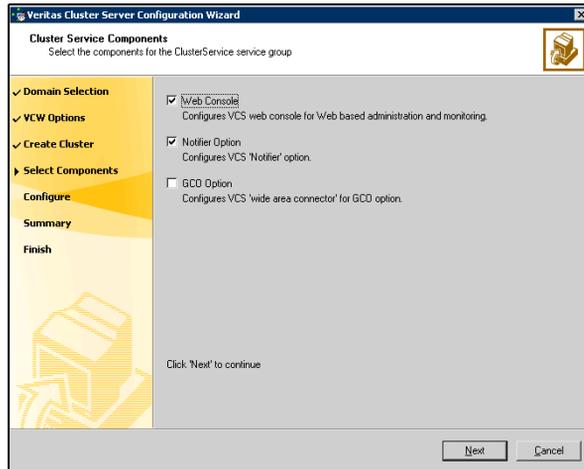
■ Click **Next**.

- 13 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network.  
If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 14 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.  
To configure the ClusterService group later, click **Finish**.  
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

- 15 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.



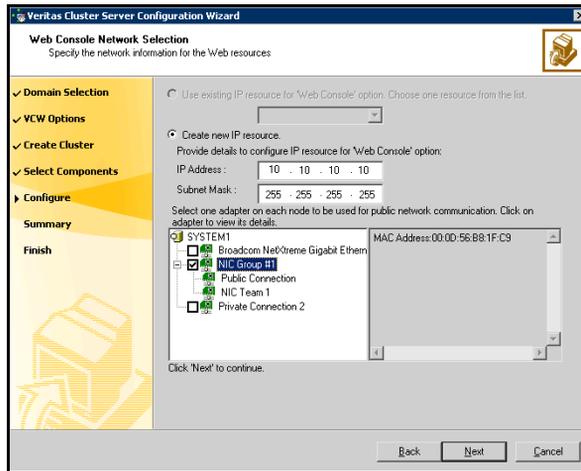
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See “[Configuring Web console](#)” on page 274.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 275.

## Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

### To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



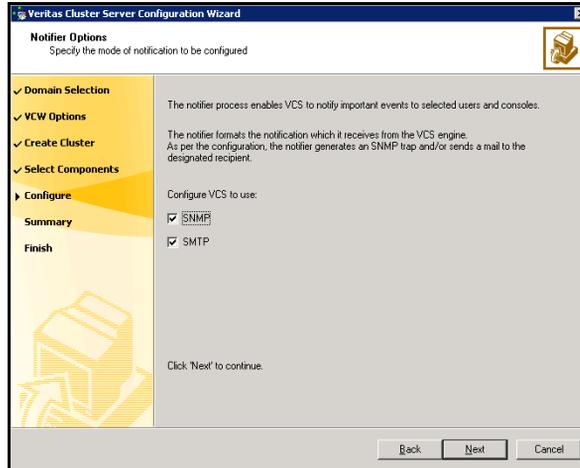
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
  - If you choose to configure a new IP address, type the IP address and associated subnet mask.
  - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
  - 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 275. Otherwise, click **Finish** to exit the wizard.

## Configuring notification

This section describes steps to configure notification.

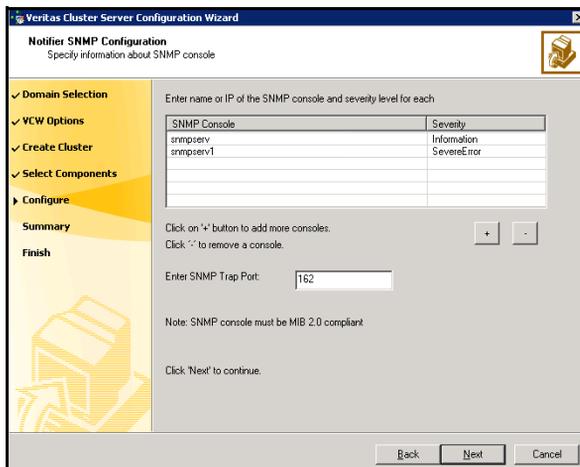
### To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.

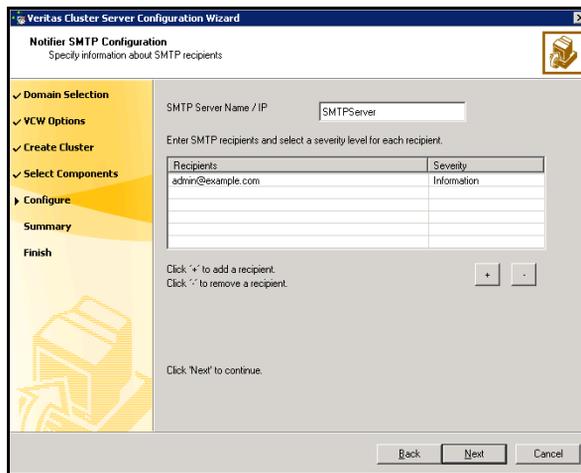


You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

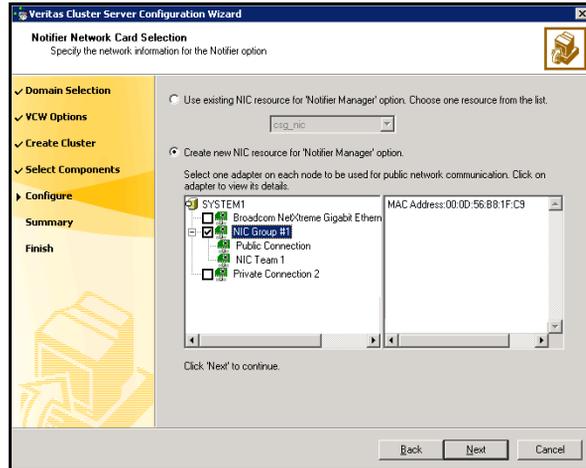


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
  - Click the corresponding field in the Severity column and select a severity level for the console.
  - Click '+' to add a field; click '-' to remove a field.
  - Enter an SNMP trap port. The default value is "162".
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
  - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
  - 6 Click **Configure**.
  - 7 Click **Finish** to exit the wizard.

## Verifying that your application or server role is configured for HA at the primary site

Make sure that your application has been configured for high availability at the primary site. If you have not yet configured the application for High Availability at the primary site, go to High Availability (HA) Configuration and follow the steps in the order shown.

See [Chapter 7, “Deploying SFW HA for high availability: New installation”](#) on page 63.

To verify the configuration, use the Cluster Manager (Java console) on the primary site and check the status of the service group in the tree view. Verify that all the resources are online.

## Configuring the VVR security service

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server 2003 environment, 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 <b>Password</b> 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 username and password when configuring the VxSAS service on the other hosts.

After providing the required information click **Next**.

- 3 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 <b>Add Domain</b> option. This displays a dialog that allows you to specify the domain name. Click <b>Add</b> to add the name to the Selected Domains list.

After specifying the domain click **Next**.

- 4 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 the **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 the **Configure** button to proceed with configuring the VxSAS service.

- 5 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.

- 6 Click **Finish** to exit the wizard.

## Configuring disaster recovery

The Disaster Recovery (DR) wizard clones the storage configuration and service group configuration from the primary site to the secondary site. It also configures VVR replication settings and connects the clusters into a global cluster. Although all the tasks can be performed using this single wizard, you may need to exit the wizard after cloning the storage to install the application, if not already done. Applications like FileShare and PrintShare are already installed as a part of the SFW HA installation process. The wizard allows you to exit the wizard, after the logical completion of each task. Launching the wizard again after you have exited the wizard brings up the Welcome page. However, when you click **Next** you are not required to re-enter information but instead proceed to the start page of the process following the one that you had last completed.

See the following topics:

- [Assigning user privileges \(secure clusters only\)](#)
- [Cloning the storage on the secondary site using the DR wizard](#)
- [Installing and configuring the application or server role](#)
- [Cloning the service group configuration from the primary to the secondary site](#)
- [Configuring replication and global clustering](#)
- [Verifying the disaster recovery configuration](#)

## Assigning user privileges (secure clusters only)

If you created secure clusters at the primary site and secondary site, in order to enable remote cluster operations you must configure a VCS user with the same name and privileges in each cluster.

When assigning privileges in secure clusters, you must specify fully-qualified user names, in the format `username@domain`. You cannot assign or change passwords for users when VCS is running in secure mode.

You must assign service group rights to the application service group as well as any dependent service groups except for the RVG service group.

See the *Veritas Cluster Server Administrator's Guide*.

### To assign user privileges at the primary site

- 1 Set the configuration to read/write mode:  
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.  
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Modify the attribute of the service group to add the user. Specify the application service group and any dependent service groups except for the RVG service group.  
`hauser -add user [-priv <Administrator|Operator> [-group service_groups]]`
- 4 Reset the configuration to read-only:  
`haconf -dump -makero`

### To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:  
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.  
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Reset the configuration to read-only:  
`haconf -dump -makero`

# Cloning the storage on the secondary site using the DR wizard

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, then the wizard can correct this and then complete the configuration.

## To clone the storage configuration from the primary site to the secondary site

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.

From the **Solutions Configurations Center** expand the **Solutions for Additional Applications** tab and from the display click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

---

**Note:** By design, the DR wizard requires specific settings for the Lanman attributes on the primary and secondary sites. Before beginning the DR configuration, the wizard checks for these values, and if they are not set as required, the wizard will automatically proceed with setting these values, both at the primary and secondary sites.

---

- 2 On the Welcome panel, read the introduction. Make sure your environment satisfies the required prerequisites and click **Next**.
- 3 In the System Selection panel, complete the requested information:

System Name	Enter the IP address or Fully Qualified Host Name (FQHN) of the primary system where the application is online. If you have launched the wizard on the system where the instance is online at the primary site, you can also specify <code>localhost</code> to connect to the system.
-------------	---

Click **Next**.

- 4 From the Service Group Selection panel, select the appropriate service group for the storage groups that you want to clone to the secondary site. You can choose to clone only the parent service group by not selecting the dependent service group. The wizard supports only one level of

dependency. In addition, only dependencies configured as online and local are supported, in soft, firm, or hard configurations.

- 5 In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.  
Click **Next**.
- 6 The Storage Validation Results panel, by default, displays detailed information about the configuration at the secondary site in comparison with that on the primary and a recommended action if required.

Disk Group	Displays the disk group name that needs to be created on the secondary site.
Volume	Displays the list of volumes, if necessary, that need to be created at the secondary site.
Size	Displays the size of the volume that needs to be created on the secondary site.
Recommended Action	Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary. <ul style="list-style-type: none"><li>■ If the volume does not exist then a new volume will be created.</li><li>■ If the volume exists but is of a smaller size than that on the primary then the volume will be expanded to the required size.</li><li>■ If the volume is of a greater size than that on the primary then the volume will be recreated using the appropriate size.</li><li>■ If the volume is the same as that on the primary then the message indicates that the volumes are identical and no action is required.</li></ul>

Click **Show Summary** to obtain summary information about the secondary storage configuration. This is a toggle button that is sensitive to the contents on the page. If the page is displaying the summary information then the button changes to **Show Details** and vice versa. The information displayed in the summary view includes:

Existing configuration	Displays the names of the disk groups that exist on the primary but do not exist on the secondary.
Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.

- If the panel displays a message indicating that the available disks are inadequate to clone the configuration at the primary site, then you can free some disks on the secondary or add more storage and then click **Refresh/Validate** to proceed with storage configuration cloning. You may also click **Refresh/Validate** to view any new components that may have got added while you were working through the wizard

---

**Note:** Before proceeding to the service group configuration, the wizard ensures that the configuration of the corresponding volumes, disk group, and the storage group component is the same at the primary site and the secondary site.

---

Click **Next**.

- 7 In the Disk selection for storage cloning panel, for each of the disk groups that does not exist or is not same as the corresponding disk group at the primary site, select disks that the wizard can use to create the respective disk groups at the secondary site.

Selecting Disks	For each of the disk groups that needs to be created, select the required disks from the Available disks pane. Either double-click on the host name or the >> option to move the hosts into the Selected disks pane.
-----------------	--

Click **Next**.

- 8 In the Volume Layout for Secondary Site Storage panel, complete the requested information:

Disk Group	Displays the disk group name to which the volume belongs.
Volume (Volume Size)	Displays the name and the size of the volume, corresponding to that on the primary, that needs to be created on the secondary.
Available Disks	Select the disks on which you want the wizard to create the volumes. From the Available Disks pane, either double-click on the host name or the >> option to move the hosts into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group.  Select disks for each unavailable volume that you want to clone on to the secondary.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click <b>Edit</b> to change the layout to suit your specific requirements.

Selected Disks	Displays the list of disks that have been moved in from the Available Disks pane.
View Primary Layout	Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 9 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.
- 10 On the Implementation panel, wait till all the tasks are implemented and the status for all the completed tasks is marked with a check (✓) symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (✗) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 11 On the Application Installation panel, do one of the following:
  - Click **Finish** to exit the wizard and proceed with installing the application. After completing the application installation, you can launch the DR wizard again.
  - Click **Next** to continue with service group cloning if the application is already installed on the system.
  - If the DR wizard is run from a remote node, then you can keep the wizard running on that node. You can then install the application locally on each of the required nodes and then click **Next** to continue.
  - If you are running the DR wizard from a local system and need to install the application on that system then you can keep the wizard running. Restart the wizard after the system gets restarted when the application installation is complete.

If you exit the wizard at any point, then after it is launched again, the wizard starts from the Welcome panel. Continue through the wizard, specifying the primary site system, the service group, and the secondary site system. The wizard proceeds to the storage cloning panel. If it detects that the storage is identical, it will proceed to the service group cloning.

# Installing and configuring the application or server role

This section provides important points which you must consider before you install the application.

The FileShare and PrintShare applications are installed as a part of the SFW HA installation. For any other application refer to the documentation provided with the application to complete the installation.

## Installing the FileShare application

Points to note when installing FileShare:

- Make sure the disk group and volumes that contain the file server shared directory exist on the shared storage.
- When installing and configuring a new file server shared directory, create the disk groups and volumes on the shared storage and subsequently create the directory structure for the file shares on the shared storage.
- If your configuration already has a file server with shares on the local storage, then move these shares to the shared storage using practices recommended by Microsoft.

## Installing the PrintShare application

Points to note when installing PrintShare:

- Make sure the required printer drivers have been added to all the systems in the cluster that are intended to be a part of the print share service group.
- For details, see *Veritas Cluster Server Administrator's Guide*.
- Make sure the printer is connected to the network and is configured with an IP address.

## Installing the IIS application

Points to note when installing IIS:

- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify that the sites to be monitored are on shared storage.
- Import the cluster disk groups and mount the volumes that contain the website data, on the first node.

- For a new IIS installation, while creating new websites, create the site folder on the shared storage and place the site content in that folder.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- For existing websites, stop the sites and then move the website content to volumes on the shared storage. You must also reconfigure the home directory location for the website in IIS and then restart the website again.
- Verify the port numbers assigned to IIS sites are not used by other applications or sites.
- Synchronize the IIS configuration on all nodes hosting the service group, as instructed in the next section.

## Installing the Microsoft Virtual Machine application

Points to note when installing MS Virtual Machine:

- Verify Microsoft Virtual Server is installed and configured identically on all nodes hosting the service group.
- Install the operating system and the applications that you want to make highly available on the virtual machine.
- Install and configure Virtual Machine Additions *on each virtual machine* if you plan to enable detailed monitoring for the virtual machine resources.
- Verify the Microsoft Virtual Server configuration files reside locally on each node.
- Make sure the name of the virtual machine is unique in the cluster.

## Installing additional applications

Following are some very generic points for installing any application:

- Make sure that the disk groups and volumes are mounted on the node. For the primary node, they should be mounted because they were just created.
- VCS requires the application program files to be installed on the same local drive on all nodes. For example, if you install the application program files on drive C of one node, installation of these same files on all other nodes must be on drive C. Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- The data files and any associated files, such as log files, should be installed on the shared storage.

# Cloning the service group configuration from the primary to the secondary site

Prior to cloning the service group on the secondary site verify that you have installed the application on the Primary and created the required service groups. You will also need to ensure that you have installed the application on the Secondary. After verifying, launch the DR wizard.

## To launch the DR configuration wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**, or click the shortcut for the **Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center**, expand the Solutions for Additional Applications tab and from the display click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

## To clone the service group configuration from the primary site to the secondary site

- 1 At the primary site, verify that you have brought the application service group online.
- 2 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.  
From the **Solutions Configurations Center** expand the Solution for Additional Applications tab and from the display click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 On the Welcome panel, click **Next** and continue through the wizard, providing the requested information for the primary site system, the service group, and the secondary site system. The wizard proceeds to the storage cloning panel. If it detects that the storage is identical it will proceed to the service group cloning.
- 4 Review the following information displayed on the Local Attributes for Service Group Cloning Analysis Results panel and click **Next** to continue with service group cloning.

Service Group Name Displays the list of application-related service groups present on the cluster at the primary site.

**Cloning the service group configuration from the primary to the secondary site**

Local attributes on Primary Cluster	Displays the service group attributes for the cluster at the Primary. These include: <ul style="list-style-type: none"> <li>■ IP Resource: consists of the IP address and the subnet mask</li> <li>■ NIC Resource: is the MAC address</li> </ul>
Local attributes on Secondary Cluster	Displays a message to indicate whether the service group or the corresponding attributes have been configured at the secondary site.

**5** In the Service Group Cloning panel, specify the requested system information for the secondary site.

Service Group Name	Depending on the application service group already created at the primary site, and subsequently selected on the Service Group Selection page, the wizard displays the names of the service groups that will be cloned at the secondary site.
Available Systems	Select the secondary systems on which you want the wizard to clone the application service group configuration.  Either double-click on the system name or use the > option to move the hosts into the Selected systems pane. The Available systems pane displays the list of all available systems on the secondary cluster.  <b>Note:</b> If you want to add systems to a service group after you finish cloning the service group configuration with the DR wizard, you cannot do so by running the DR wizard again. Instead, run the VCS configuration wizard and edit the system list of the existing service group.
Selected Systems	Displays the list of selected systems.

Click **Next**.

**6** In the Service Group Attribute Selection panel, complete the requested information to create the required resources on the secondary site. The panel also displays the service group resource name and the attribute information at the primary site.

Resource Name	Displays the list of resources that exist on the primary cluster.
Attribute Name	Displays the attribute name associated with each of the resources displayed in the Resource Name column.

**Cloning the service group configuration from the primary to the secondary site**

Primary Cluster	Displays the primary attribute values for each of the displayed attributes.
Secondary Cluster	Provides fields to specify values for the secondary attributes. For the MACAddress attribute select the appropriate public NIC from the drop-down list.

Click **Next**.

- 7 In the Service Group Summary, review the attribute information that will be cloned on to the secondary cluster. Click **Back** to change any of the secondary service group attributes. Otherwise, click **Next** to proceed with cloning the service group configuration on the Secondary.
- 8 In the Implementation panel, wait till all the tasks are implemented and the Status for all the completed tasks is marked with a check (✓) symbol indicating successful completion. Wait until the wizard creates the IP resource and the other resources. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 9 In the Service Group Cloning completion panel, click **Next** to continue with the replication configuration.
  - If you want to configure the replication settings later, click **Finish** to exit the wizard at this point; launch the wizard again whenever required.

When you launch the wizard again, continue through the wizard, specifying the primary site system, the service group, and the secondary site system. If the wizard detects that the storage is identical it will proceed to the service group cloning. If it detects that the service group has been configured, the wizard proceeds to the replication and GCO configuration panel.
  - If you plan to use VVR replication, click **Next** to continue with configuring VVR replication and global clustering.
  - If you plan to use hardware replication, click **Next** to configure the global clustering. Only after completing the global clustering configurations, configure hardware replication.

## Configuring replication and global clustering

You can choose to configure replication either using VVR or any other array-based hardware replication and then use this wizard to configure global clustering. If you plan to use array-based hardware replication, then you must first complete configuring global clustering using the DR wizard and then proceed with configuring replication.

---

**Caution:** To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment, you must first run the wizard before configuring replication. Not doing so can result in data corruption.

---

Irrespective of the method you choose for replication, you will still need to set up Global Clustering to complete the disaster recovery configuration.

### To start the Disaster Recovery Configuration Wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**, or click the shortcut for the **Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center**, expand the Solutions for Additional Applications tab and from the display click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

### To configure replication and GCO

- 1 Verify that you have brought the application server service group online at the primary site and imported the appropriate disk groups at the secondary site.
- 2 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.  
From the **Solutions Configurations Center** expand the Solution for Additional Applications tab and from the display click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 On the Welcome panel, click **Next** and continue through the wizard, providing the requested information. When the wizard reaches the storage cloning panel and detects that the storage is identical, it will proceed to the service group cloning panel. Similarly, when it finds the service group is

properly configured on the secondary, it will proceed to the Replication Options panel.

- 4 On the Replication Options panel, do one of the following:
  - If you want to configure both VVR replication and the Global Clustering Option (GCO), click **Configure VVR and the Global Cluster Option (GCO)**. Click **Next**.
  - If you plan to use hardware replication rather than VVR replication and therefore only want to configure the global clustering, click **Set Global Cluster Option**. Click **Next** to continue to [step 7](#).

---

**Note:** You must complete configuring global clustering before you configure hardware replication.

---

- To configure replication and global clustering later, click **Configure Replication and Global Cluster Option (GCO) later**. Click **Next** and then click **Finish**.
- 5 On the Replication Settings for Replicated Volume Group panel, specify the requested information.

Disk Group	The right pane displays the list of disk groups. By design, an RVG is created for each disk group.
RVG Name	Displays the default RVG name. If required, change this to a name of your choice.
RDS Name	Displays the default Replicated Data Set (RDS) name. If required, change this to a name of your choice.
Available Volumes	Displays the list of available volumes that have not been selected to be a part of the RVG.  Either double-click on the volume name or use the > option to move the volumes into the Selected RVG Volumes pane.
Selected RVG Volumes	Displays the list of volumes that have been selected to be a part of the RVG.  To remove a selected volume, either double-click the volume name or use the < option to move the volumes into the Available Volumes pane.

Primary SRL	Select the appropriate primary Replicator Log volume from the drop-down menu. <b>Size:</b> Enter an appropriate log size value in the corresponding <b>Size</b> field. If you did not create the Replicator Log volume earlier, click <b>Create New</b> on the drop-down menu.
Secondary SRL	Select the appropriate secondary Replicator Log volume from the drop-down menu. <b>Size:</b> Enter an appropriate size value in the corresponding <b>Size</b> field. If you did not create the Replicator Log volume earlier, click <b>Create New</b> on the drop-down menu.
Add RVG	Click this option to create a new RVG. This option is especially useful if you want to organize the volumes present in a disk group under separate RVGs. By default the wizard is designed to organize all the volumes under a disk group under one RVG. However, using the <b>Add RVG</b> option you can choose to organize them differently, based on your specific requirements.
Delete RVG	Click this option to delete any of the existing RVGs related to the DR set up that you are creating.
Start Replication after the wizard completes	Select this check box to start replication automatically after the wizard completes the necessary configurations.
<b>■</b> Click <b>Advanced Settings</b> to specify some additional replication properties. The options on the dialog box are described column-wise, from left to right; refer to the <i>Veritas Volume Replicator Administrator's Guide</i> for additional information on VVR replication options:	
Replication Mode	Select the required mode of replication; <b>Synchronous</b> , <b>Asynchronous</b> , or <b>Synchronous Override</b> . The default is synchronous override.
Log Protection	Select the appropriate log protection from the list.  The <b>Off</b> option disables Replicator Log Overflow protection.

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	Enter a name of your choice for the primary RLINK. If you do not specify any name then the wizard assigns a default name.
Secondary RLINK Name	Enter a name of your choice for the Secondary RLINK. If you do not specify any name then the wizard assigns a default name.
Bandwidth	By default, VVR replication uses the maximum available bandwidth. You can select <b>minimum</b> from the list to indicate that the minimum bandwidth should be used.  The default unit is Mega bits per second (Mbps) and the minimum allowed value is 1 Mbps.
Protocol	UDP/IP is the default replication protocol. Choose TCP/IP or UDP/IP for a regular Secondary.
Packet Size (Bytes)	Default is 1400 bytes. From the drop-down list, choose the required packet size for data transfer. The default unit for the packet size is Bytes. You can set the packet size only if the protocol is UDP/IP.
Latency Protection	By default, latency protection is set to <b>Off</b> .  When this option is selected the <b>High Mark Value</b> and the <b>Low Mark Value</b> are disabled. Select the <b>Fail</b> or <b>Override</b> option to enable Latency protection.  This <b>Override</b> option behaves like the <b>Off</b> option when the Secondary is disconnected and behaves like the <b>Fail</b> option when the Secondary is connected.

**High Mark Value** This option is enabled only when Latency Protection is set to **Override** or **Fail**. It specifies the maximum number of pending updates by which the secondary site can be behind the primary site. The default value is 10000, but you can specify the required limit.

To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.

**Low Mark Value** This option is enabled only when Latency Protection is set to **Override** or **Fail**. When the updates in the Replicator log reach the **High Mark Value**, then the writes to the system at the primary site continues to be stalled until the number of pending updates on the Replicator log falls back to the **Low Mark Value**. The default value is 9950, but you can specify the required limit.

**Initial Synchronization** If you are doing an initial setup, then use the **Auto Sync** option to synchronize the secondary site and start replication. This is the default.

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.

If you want to use the **Synchronize from Checkpoint** method then you must first create a checkpoint.

If you have a 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 the checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

Click **OK**. On the Replication Settings for Replicated Volume Group panel click **Next**.

- 6 On the Replication Attribute Settings panel, specify the requested replication attribute information for the cluster at the primary site and the secondary site. You can specify the replication attributes for each of the RVGs. Click the arrow icon present on each RVG row to expand the view, to display the required replication attribute fields.

**Disk Group** Displays the list of disk groups that have been configured.

RVG Name	Displays the Replicated Volume Groups corresponding to the disk groups.
IP Address	Enter replication IPs that will be used for replication, one for the primary site and another for the secondary site.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC from the drop-down list for the system at the primary and secondary site.
Copy	Enables you to copy the RLINK attributes across multiple RLINKs. You must have at least two RLINKs to be able to use this operation to copy RLINK attributes from the current to the other RLINKs.

After specifying the replication attributes for each of the RVGs, click **Next**.

- 7 On the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. If you do not to select this option then you may need to bring the service group online and start GCO manually, after the wizard completes.

- 8 On the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication settings, replication resource settings or the global cluster settings.  
Click **Next**.
- 9 On the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for Global Clustering and the Status for all the completed tasks is marked with a check (✓) symbol indicating successful completion. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 10 On the Finish panel, review the displayed information. If some task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

## Verifying the disaster recovery configuration

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- Confirm that the configuration of disk groups and volumes at the DR site have been created by the DR wizard storage cloning.
- Confirm that the application VCS service group has been created in the DR cluster including the same service group name, same resources, and same dependency structure as the primary site's application VCS service group.
- Confirm that the application service group is online at the primary site. The application service group should remain offline at the DR site.
- Ensure VVR replication configuration. This includes ensuring that the RVGs have been created at primary and secondary with the correct volume inclusion, replication mode, srl configuration, and any specified advanced options.
- Confirm that the replication state matches what was specified during configuration. If specified to start immediately, ensure that it is started. If specified to start later, ensure that it is stopped.
- Ensure that the VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.

- Confirm that the RVG service groups are online at the primary and secondary sites.
- Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- Ensure that the application service groups are configured as global.
- Check to ensure that the two clusters are communicating and that the status of communication between the two clusters has a state of Alive.
- If you are using VVR for replication and chose to start replication manually in the DR wizard, to avoid replicating large amounts of data over the network the first time, then you will need to start the process necessary to synchronize from checkpoint. This typically consists of
  - starting a VVR replication checkpoint
  - performing a block level backup
  - ending the VVR replication checkpoint
  - restoring the block level backup at the DR site
  - starting replication from the VVR replication checkpointTo learn more about the process of starting replication from a checkpoint, refer to the *Veritas Volume Replicator Administrator's Guide*.
- Do not attempt a wide area failover until data has been replicated and the state is consistent and up to date. The Solutions Configuration Center provides a Fire Drill Wizard to test wide area failover.

## Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. If the local clusters are running in secure mode, you may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters. The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.
- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

### To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

### To add the `-secure` option to the `StartProgram` resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add `-secure` switch to the path of the executable Scalar Value. For example:  

```
C:\Program Files\Veritas\Cluster Server\bin\wac.exe
-secure
```
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

### To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.

Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:

```
vssat setuptrust --broker <host:port> --securitylevel
<low|medium|high> [--hashfile <filename> | --hash <root
hash in hex>]
```

For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:

from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low
```

from RB2, type:

```
vssat setuptrust --broker RB1:14141 --securitylevel low
```

### To bring the `ClusterService-Proc (wac)` resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the **Process** agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

## Possible task after creating the DR environment: Adding a new failover node

The following procedure describes how to add an additional node to the cluster at either the primary or secondary site after your disaster recovery environment is in operation. The clusters at each site are not required to have the same number of nodes or the same failover configuration.

### Preparing the new node

Install SFW HA on the new system and then add the system to the cluster.

#### To install SFW HA and add the system to the cluster

- 1 Refer to “[Installing SFW HA](#)” on page 257 for installation instructions.
- 2 Start the Veritas Cluster Server Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center** or click the shortcut for the **Solutions Configuration Center**.

From the **Solutions Configurations Center** expand **Disaster Recovery Configuration > Configure the cluster at the Secondary site** and from the display click **Configure the cluster** to add the new system to the cluster. If necessary, refer to the *Veritas Cluster Server Administrator's Guide* for information on this procedure.

### Preparing the existing DR environment

If you plan to add a failover node to the secondary site, you must temporarily switch the roles of the Primary and Secondary sites so that the current site becomes the Primary. This action reverses the direction of replication.

#### To prepare the existing DR environment

- 1 If you are adding the failover node to the cluster at the primary site, proceed directly to [step 2](#). If you are adding a failover node to the secondary site, you must switch the roles of the primary and secondary sites. This action reverses the direction of replication.
  - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
  - Click **Switch To**, and click **Remote switch**.
  - In the **Switch global group** dialog box:
    - Click the cluster at the secondary site you want to switch the group to.

- Click the specific system where you want to bring the global application service group online.
- Click **OK**.

- 2 Take the global application service group offline at the current primary site.
- 3 Take the VVR replication service group offline.

## Modifying the replication and application service groups

Add the new failover node to the system lists in the Replication and application service groups.

### To add the failover node to the system lists

- 1 Bring the replication service group online on an existing cluster node of the current primary site.
- 2 Bring the MountV resources of the corresponding application service group online on the same node.
- 3 Use the **Modify an existing replication service group** option of the Volume Replicator Agent Configuration Wizard (**Start > All Programs > Symantec > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard**) to add a new node to the system list for the replication service group. If necessary, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for information on this procedure.
- 4 Use the **Modify service group** option of the FileShare, PrintShare, IIS, MSVirtual Server Machine Configuration Wizard or Application Configuration Wizard.  
 Start the appropriate Configuration Wizard from the Solutions Configuration Center. For example, for FileShare click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center High Availability (HA) Configuration > Configure the Service Group > FileShare Configuration Wizard** to add the new node to the system list for the respective application service group. Check the check box to bring the service group online after the wizard completes. If necessary, refer to the *Cluster Server Administrator's Guide* for information on this procedure.
- 5 After bringing the application service group online, configure all the application database stores to automatically mount on start-up.

## Reversing replication direction

If you added a failover node at the original secondary site and migrated the RVG in [“Preparing the existing DR environment”](#) on page 302, move the global application service group back to the original primary site and reverse the

direction of replication. These actions switch the Primary and Secondary sites back to their original roles.

**To reverse the replication direction**

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 In the **Switch global group** dialog box:
  - Click the cluster to switch the group to.
  - Click the specific system where you want to bring the global application service group online.
  - Click **OK**.

# Maintaining: Normal operations and recovery procedures

This section provides tasks during normal operations of the DR solutions 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 the following tools:

- The VEA GUI
- The Command Line Interface (CLI)
- Perfmon
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Storage Foundation 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. The following are a generic set of tasks that you may need to perform:

- Take the RVG resource offline on both the clusters.
- Transfer the primary role to the host at the secondary site by using the **Migrate** option.
  - a From the VEA screen, right-click the primary RVG and select **Migrate**.
  - b Select the secondary host and click **OK**. The replication role is migrated to the secondary host.
- Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- Bring the RVG resource online on the new secondary.
- Bring the application group online on the new primary.

You can now verify that the application functions properly on the new primary with the replicated data. After verifying its functioning, you can revert the roles to what they were originally by repeating the procedure.

---

**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.

---

## Disaster recovery procedures

This section provides information on bringing up an application 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 returned to normal functioning after a disaster.

### To bring up the application on the secondary host

- 1 From the left pane in the VEA GUI console on the secondary host, right-click the desired secondary RVG node inside the replication network.
- 2 Select **Takeover** and follow the instructions to perform the takeover operation. You can choose to perform takeover with the following options:
  - Perform **Takeover** *with* the **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 **Takeover** *without* the **fast-failback** option. In this case, you will 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 the takeover, the existing secondary becomes the new primary.

- 3 Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- 4 Bring the application group online.

Now you can start using the application on the new primary.

### Restoring the primary host

After a disaster, when the original primary becomes available again, you may want to revert the role of the primary back to this host.

### To restore the primary host

- 1 Take the RVG resource off-line on both the clusters.

- 2 Depending on whether you performed **Takeover** *with* or *without* the **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 new primary's context menu.
  - For **Takeover** *without* the **Fast-failback** option:

After performing a takeover without fast-failback, you must convert the original primary to a secondary by using the **Make Secondary** option.

Before performing the **Make Secondary** operation, the original primary's RVG and the new primary's RVG will be shown in separate RDS's. 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 the **Synchronize Automatically** option.
- 3 After the synchronization is complete, perform a migrate operation to transfer the primary role back to the original primary. Right-click on the primary RVG and select **Migrate** from the menu that appears.
- 4 Make sure that the volumes have retained the same drive letters as they had before the disaster.
- 5 Bring the RVG resource online on the secondary.
- 6 Bring the application group online on the original primary.

## Recovery procedures for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

See "[Supported disaster recovery configurations for service group dependencies](#)" on page 253.

The procedure and requirements for bringing service group dependencies online at the secondary site depends on their configuration: soft, firm, or hard.

In general, if a child or parent remains online at the primary site, you take it offline before you bring the child and parent service groups online in the correct order on the secondary site.

An exception is the RVG service group, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, soft dependency link.

**Table 12-4** Online, local, soft dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> <li>■ The parent remains online on the primary site.</li> <li>■ An alert notification at the secondary site occurs for the child service group only.</li> <li>■ The RVG group remains online.</li> </ul>	<ol style="list-style-type: none"> <li>1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.</li> <li>2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).</li> </ol>
The parent service group fails	<ul style="list-style-type: none"> <li>■ The child remains online on the primary site.</li> <li>■ An alert notification at the secondary site occurs for the parent only.</li> <li>■ The RVG group remains online.</li> </ul>	<ol style="list-style-type: none"> <li>1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.</li> <li>2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</li> </ol>

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, firm dependency link.

**Table 12-5** Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> <li>■ The parent goes offline on the primary site.</li> <li>■ An alert notification at the secondary site occurs for the child service group only.</li> <li>■ The RVG group remains online.</li> </ul>	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Leave the RVG group online at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> <li>■ The child remains online on the primary site.</li> <li>■ An alert notification at the secondary site occurs for the parent only.</li> <li>■ The RVG group remains online.</li> </ul>	<p>1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.</p> <p>2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p>

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, hard dependency link.

**Table 12-6** Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> <li>■ The parent goes offline on the primary site.</li> <li>■ An alert notification at the secondary site occurs for the child service group only.</li> <li>■ The RVG group remains online.</li> </ul>	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Do not take the RVG group offline at the primary site.</p>

**Table 12-6** Online, local, hard dependency link (Continued)

Failure condition	Results	Action required
The parent service group fails	<ul style="list-style-type: none"><li>■ The child remains online on the primary site.</li><li>■ An alert notification at the secondary site occurs for the parent only.</li><li>■ The RVG group remains online.</li></ul>	<ol style="list-style-type: none"><li>1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.</li><li>2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</li></ol>

# Testing fault readiness by running a fire drill

Topics in this chapter include:

- [About disaster recovery fire drills](#)
- [About the Fire Drill Wizard](#)
- [Tasks for configuring and running fire drills](#)
- [Prerequisites for a fire drill](#)
- [Fire Drill Wizard actions](#)
- [Preparing the fire drill configuration](#)
- [Running a fire drill](#)
- [Restoring the fire drill system to a prepared state](#)
- [Deleting the fire drill configuration](#)

## About disaster recovery fire drills

A disaster recovery plan should include regular testing of an environment to ensure that a DR solution is effective and ready should disaster strike. This testing is called a fire drill. SFW HA provides a Fire Drill Wizard to help you set up and run a fire drill.

## About the Fire Drill Wizard

The Fire Drill Wizard tests the fault readiness of a configuration by mimicking a failover from the primary site to the secondary site. The wizard does this without stopping the application at the primary site and disrupting user access.

The wizard prepares for the fire drill by completing the following steps:

- Creates a fire drill service group on the secondary site  
The fire drill service group is a copy of the application service group, using the same service group name with the prefix *FD<sub>nn</sub>*. The wizard renames the fire drill service group resources with a prefix *FD<sub>nn</sub>* and changes attribute values as necessary to refer to the FD resources.
- Prepares a copy (mirror) of the production data on the secondary site  
You assign one or more disks for the mirrored volumes while running the wizard. Mirror preparation can take some time, so you can exit the wizard once this step is started and let the preparation continue in the background.

Once these steps are complete, the wizard can run the fire drill. Running the fire drill detaches the mirrors from the original volumes to create point-in-time snapshots of the production data. It also brings the application online in the fire drill service group at the secondary site. Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to failover and come online at the secondary site should the need arise.

Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online on the primary site.

Running the fire drill creates a fire drill disk group on the secondary site with a snapshot of the application data to use for testing purposes. The wizard assigns the fire drill disk group name by prefixing the original disk group name with *FD<sub>nn</sub>*.

After running the fire drill, you can choose to restore the fire drill configuration to a prepared state for use in regularly testing the disaster recovery solution, or you can delete the fire drill configuration and recreate it as needed using the wizard.

# Tasks for configuring and running fire drills

The Fire Drill Wizard helps you configure and run a fire drill.

[Table 13-1](#) outlines the high-level objectives and the tasks to complete each objective.

**Table 13-1** Tasks for configuring and running fire drills

Objective	Tasks
<a href="#">“Prerequisites for a fire drill”</a> on page 313	<ul style="list-style-type: none"> <li>✓ Verifying hardware and software prerequisites</li> </ul>
<a href="#">“Preparing the fire drill configuration”</a> on page 315	<ul style="list-style-type: none"> <li>✓ Using the wizard to prepare the initial fire drill configuration</li> </ul>
<a href="#">“Running a fire drill”</a> on page 318	<ul style="list-style-type: none"> <li>✓ Using the wizard to run the fire drill</li> <li>✓ Performing your own tests of the application to confirm that it is operational</li> </ul>
<a href="#">“Restoring the fire drill system to a prepared state”</a> on page 319	<ul style="list-style-type: none"> <li>✓ Using the wizard to restore the fire drill system to a state of readiness for future fire drills or to prepare for removal of the fire drill configuration</li> </ul>
<a href="#">“Deleting the fire drill configuration”</a> on page 320	<ul style="list-style-type: none"> <li>✓ Using the wizard to remove the fire drill configuration</li> </ul>

## Prerequisites for a fire drill

Ensure that the following prerequisites are met before configuring and running a fire drill:

- ✓ The primary and secondary sites must be fully configured with VVR replication and the global cluster option.
- ✓ The Veritas FlashSnap option must be installed on all nodes of the clusters at the primary and secondary sites.

- ✓ The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- ✓ On the secondary site, empty disks must be available with enough disk space to create snapshot mirrors of the volumes. These disks must be in the same disk group that contains the RVG. If the disk group does not have empty disks available, you must use the VEA to add the disks to the disk group before you run the wizard. The secondary system must have access to the disks or LUNs.
- ✓ For each IP address in the application service group, an IP address must be available to use on the secondary site for the fire drill service group. The wizard can accept input for one IP address and Lanman resource. If the application service group has multiple IP addresses and Lanman resources, the wizard notifies you to edit the fire drill service group resources to supply these values. Information on editing service group resources is covered in the VCS administration guide.  
See *Veritas Cluster Server Administrator's Guide*.
- ✓ If you want the fire drill wizard to run a script that you supply, ensure that the script file is available on any secondary site nodes where you plan to run the fire drill.
- ✓ If the cluster is secured, the login you use to run the Fire Drill Wizard must have the appropriate permissions to make changes in the cluster.

In addition, for testing purposes, you may want to create and populate a new table from the active node at the primary site. After you run the fire drill to bring the fire drill service group online and create the fire drill snapshots, you can check that the table and its data were replicated and are available from the fire drill service group. You can automate this process with a script and when preparing to run the fire drill, specify it as a post-fire drill script.

You can run the Fire Drill Wizard from any node in the domain of the cluster, as long as the SFW HA client is installed on that node.

## Fire Drill Wizard actions

While running the Fire Drill Wizard, you select from a menu of fire drill wizard actions.

After an action is complete, if you proceed in the wizard, the menu is displayed so that you can select the next action. Therefore, you can execute all the actions sequentially without exiting the wizard. However, typically you perform the first two actions, run your own tests to verify the fire drill, and then later start the wizard again to complete one or both of the last two actions.

The actions consist of the following:

Prepare for Fire Drill	<p>Creates the configuration required to run a fire drill. This step takes some time as the wizard prepares the mirrors for the snapshots.</p> <p>If this option is unavailable, the fire drill configuration already exists on the specified system.</p> <p>See “<a href="#">Preparing the fire drill configuration</a>” on page 315.</p>
Run Fire Drill	<p>Runs the fire drill. The wizard creates the volume snapshots and brings the fire drill service group online. Optionally you can specify a script to be run once the fire drill is complete.</p> <p>If a fire drill has been run, you must restore the fire drill configuration to a prepared state before the wizard re-enables this option.</p> <p>See “<a href="#">Running a fire drill</a>” on page 318.</p>
Restore to Prepared State	<p>Restores the fire drill configuration for another fire drill or to prepare the fire drill configuration for deletion.</p> <p>This option becomes available once a fire drill has been run.</p> <p>The wizard snaps back the snapshot mirrors to reattach to the original volumes and takes the fire drill service group offline.</p> <p>See “<a href="#">Restoring the fire drill system to a prepared state</a>” on page 319.</p>
Delete Fire Drill Configuration	<p>Deletes the fire drill configuration to free up disk space. The wizard deletes the service group on the secondary site and performs a snap abort to delete the snapshot mirrors created on the secondary site for use in the fire drill.</p> <p>If a fire drill has been run, this option is disabled until you first restore the fire drill configuration to a prepared state. This ensures that mirrors are reattached and the fire drill service group is offline before the configuration is deleted.</p> <p>See “<a href="#">Deleting the fire drill configuration</a>” on page 320.</p>

## Preparing the fire drill configuration

Preparing the fire drill configuration creates a fire drill service group and snapshot mirrors of production data at the specified node on the secondary site. You specify the application service group and the secondary system to use. Only one service group can be prepared for a fire drill at one time.

---

**Note:** Preparing the snapshot mirrors takes some time to complete.

---

Before you prepare the fire drill configuration, you should verify that you meet the prerequisites.

See [“Prerequisites for a fire drill”](#) on page 313.

#### To prepare for the fire drill

- 1 Open the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**).
- 2 Start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, review the information and click **Next**.
- 4 In the System Selection panel, specify the primary site system on which the service group to be used for the fire drill is online and click **Next**.  
All systems containing online global service groups are available to select. The default system is the node where you launched the wizard. When selecting a system you can specify either a fully qualified host name or IP address.
- 5 In the Service Group Selection panel, select the service group that you want to use for the fire drill and click **Next**. (You can select only one service group at a time for a fire drill.)
- 6 In the Secondary System Selection panel, select the cluster and the system to be used for the fire drill at the secondary site, and then click **Next**.  
The selected system must have access to the replicated data and to disks for the snapshots that will be created for the fire drill.
- 7 In the Fire Drill Mode Selection panel, the available options depend on whether or not the fire drill service group already exists on this system and whether it is on or offline. Choose one of the following and click **Next**:

If the Prepare for Fire Drill option is available, a fire drill service group does not exist on this system. Click **Prepare for Fire Drill** and continue with the remaining steps in this procedure.

If the Run Fire Drill option is available, a fire drill service group has already been prepared. You can run the fire drill with no further preparation. Click **Run Fire Drill** and follow the procedure for running a fire drill.  
See [“Running a fire drill”](#) on page 318.

If the Restore to Prepared State option is available, the fire drill service group remains online from a previous fire drill. Click **Restore to Prepared State** and follow the procedure for restoring the fire drill configuration to a prepared state. See [“Restoring the fire drill system to a prepared state”](#) on page 319.

- 8 In the Fire Drill Service Group Settings panel, assign the virtual IP address and virtual name (Lanman name) to be used for the fire drill service group that will be created on the secondary site. These must be an address and name not currently in use.

If the service group contains more than one IP and Lanman resource, this panel does not display. After the fire drill service group is created, the wizard notifies you to manually update the IP and Lanman resources in the fire drill service group.

- 9 In the Disk Selection panel, review the information and make the selections as follows and click **Next**:

Volume	Select the volumes for the fire drill snapshots. By default all volumes associated with the service group are selected. If you deselect a volume that might result in the fire drill service group failing to come online, the wizard displays a warning message.
Disk Group	Shows the name of the disk group that contains the original volumes. This field is display only.
Fire Drill DG	Shows the name of the fire drill disk group that running the fire drill will create on the secondary system to contain the snapshots. This field is display only. For the fire drill disk group name, the wizard prefixes the original disk group name with <i>FDnn</i> .
Disk	Click the plus icon to the right of the Disk column and specify the disk to be used for the snapshot volume. Repeat for each row that contains a selected volume.  You can store multiple snapshot volumes on the same disk, if the production volumes reside on disks in the same disk group.
Mount Details	Shows the mount details for the snapshot volumes on the secondary system, which match the mounts for the production volumes. This field is display only.

- 10 Wait while the wizard completes the preparation tasks. First the fire drill service group is created on the secondary site (but remains offline). Next the

snapshot mirrors for the volumes are prepared; this can take some time. You may want to minimize the wizard while the task runs in the background. You can also track the mirror preparation progress in the VEA. When done, the wizard displays a message that the fire drill preparation is complete.

- 11 To run the fire drill now, choose **Next**, or click **Finish** to exit the wizard. If you choose **Finish** the fire drill preparation remains in place. The next time you run the wizard, you choose the primary and secondary systems and service group and then can continue with running the fire drill.

## Running a fire drill

After you complete the initial fire drill preparation step using the Fire Drill Wizard, you can run the fire drill immediately without exiting the wizard or run the wizard later to run the fire drill. Running the fire drill does the following:

- Creates the snapshots
  - Splits the fire drill disk group
  - Enables the firedrill resources
  - Brings the fire drill service group online
  - Optionally, executes a specified command to run a script
- For example, if you earlier created and populated a test table at the primary site, you could create a script to verify replication of the data. For the wizard to run the script, the script must exist on the secondary system that you are specifying for the fire drill.

### To run a fire drill

- 1 If you completed the initial preparation and have not exited the wizard, go to [step 7](#). Otherwise, if you need to restart the wizard, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify the primary site system that contains the service group for which you want to run the fire drill and click **Next**.
- 5 In the Service Group Selection panel, select the service group and click **Next**.
- 6 In the Secondary System Selection panel, specify the system previously prepared for the fire drill at the secondary site.

- 7 In the Fire Drill Mode Selection panel, click **Run Fire Drill** and click **Next**. If a fire drill has been run previously, you must restore the fire drill configuration to a prepared state before the wizard enables the option to run another fire drill.  
See “[Restoring the fire drill system to a prepared state](#)” on page 319.
- 8 In the Post Fire Drill Script panel, optionally specify the full path to a script for the wizard to run on the secondary system right after running the fire drill. The script must already exist on the secondary system. Click **Next**.
- 9 In the Fire Drill Implementation screen, wait until all fire drill tasks are performed and the Fire drill ran successfully message is displayed.
- 10 Click **Finish**.

---

**Warning:** After running the fire drill, the fire drill service group remains online. After you verify the fire drill results, remember to take the fire drill service group offline as soon as possible by running the wizard to restore the system to the prepared state. If the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

---

- 11 Run your own tests to verify the fire drill results.
- 12 Run the wizard again to restore the fire drill configuration to the prepared state.  
See “[Restoring the fire drill system to a prepared state](#)” on page 319.

## Restoring the fire drill system to a prepared state

After running a fire drill and verifying the results, use the Fire Drill Wizard to restore the fire drill system at the secondary site to a prepared state. When restoring the fire drill system to a prepared state, the wizard completes the following tasks:

- Takes the fire drill service group offline
- Disables the fire drill service group resources
- Imports the fire drill disk group
- Joins the fire drill disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

After running a fire drill, restoring the fire drill system to a prepared state is required to do any of the following:

- Run another fire drill.
- Restore the secondary system to a state where it can be used as failover for the application service group at the primary site.
- Delete the fire drill configuration.

#### To restore the fire drill system to a prepared state

- 1 If you completed running a fire drill and have not exited the wizard, go to [step 7](#). Otherwise, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify the system at the primary site that contains the service group on which the fire drill was run and click **Next**. The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 In the Fire Drill Mode Selection panel, click **Restore to Prepared State** and click **Next**.  
If you have run a fire drill but not yet restored the configuration, this is the only option available. If the option is unavailable, the configuration has already been restored or the fire drill has not yet been run.
- 8 In the Restore Fire Drill screen, wait until the screen shows the restoration tasks are completed. Then click **Next** if you want to delete the fire drill configuration or click **Finish** to exit the wizard, leaving the fire drill configuration in a prepared state.

## Deleting the fire drill configuration

If you no longer need a fire drill configuration you can delete it. Deleting a fire drill configuration deletes the fire drill service group on the secondary site and performs a snap abort of the snapshot mirrors created on the secondary site for use in the fire drill. It frees up the disk space used for the snapshot mirrors for other use.

If you have run a fire drill and want to delete the configuration, you must first restore the fire drill configuration to a prepared state before the wizard enables the option to delete the fire drill configuration.

See “[Restoring the fire drill system to a prepared state](#)” on page 319.

### To delete a fire drill configuration

- 1 If you have just used the wizard to restore the fire drill configuration and have not exited the wizard, go to [step 7](#). Otherwise continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify the system at the primary site that contains the service group on which the fire drill was run and click **Next**. The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 In the Fire Drill Mode Selection panel, click **Delete Fire Drill Configuration** and click **Next**.
- 8 In the Delete Fire Drill Configuration panel, wait until screen shows the deletion is complete and then click **Next** and **Finish**.



# MSCS Solutions

This section includes the following chapters:

- [Chapter 14, “MSCS solutions overview” on page 325](#)
- [Chapter 15, “Deploying SFW with MSCS” on page 329](#)
- [Chapter 16, “Deploying SFW with MSCS in a campus cluster” on page 367](#)
- [Chapter 17, “Deploying SFW and VVR with MSCS” on page 415](#)



# MSCS solutions overview

Microsoft Cluster Server (MSCS) may be used with Veritas Storage Foundation for Windows to provide high availability for your application. MSCS may be used with Veritas Storage Foundation for Windows and Veritas Volume Replicator to provide replication support for your application. Using VVR with MSCS provides a replicated backup of your application 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.

## About high availability

“High availability” maintains continued functioning of applications in the event of computer failure, where data and applications are continuously available using redundant software and hardware. “High availability” can refer to any software or hardware that provides fault tolerance, but generally it has become associated with clustering.

A cluster is a group of independent computers working together as a single system to ensure that mission-critical applications and resources are highly available. The cluster is managed as a single system, shares a common namespace, and is specifically designed to tolerate component failures and to support the addition or removal of components in a way that is transparent to users.

Keeping data and applications functioning 24 hours a day and seven days a week is the necessary for critical applications today. Clustered systems have several advantages, including fault tolerance, high availability, scalability, simplified management, and support for rolling upgrades.

## About campus clustering

Campus clusters are multiple-node clusters that provide protection against disasters. These clusters are in separate buildings (or sites) with mirrored SAN-attached storage located in each building. Typical campus clusters involve two sites; you can use more than two sites for additional redundancy.

In a typical configuration, each node has its own storage array with the same number of disks and contains mirrored data of the storage on the other array. Refer to [Chapter 16, “Deploying SFW with MSCS in a campus cluster” on page 367](#), for details on a typical active/passive configuration for a campus cluster.

This environment also provides a simpler solution for disaster recovery than a more elaborate Veritas disaster recovery environment with replication software; however, a campus cluster generally stretches a shorter distance than a replication-based solution depending on the hardware.

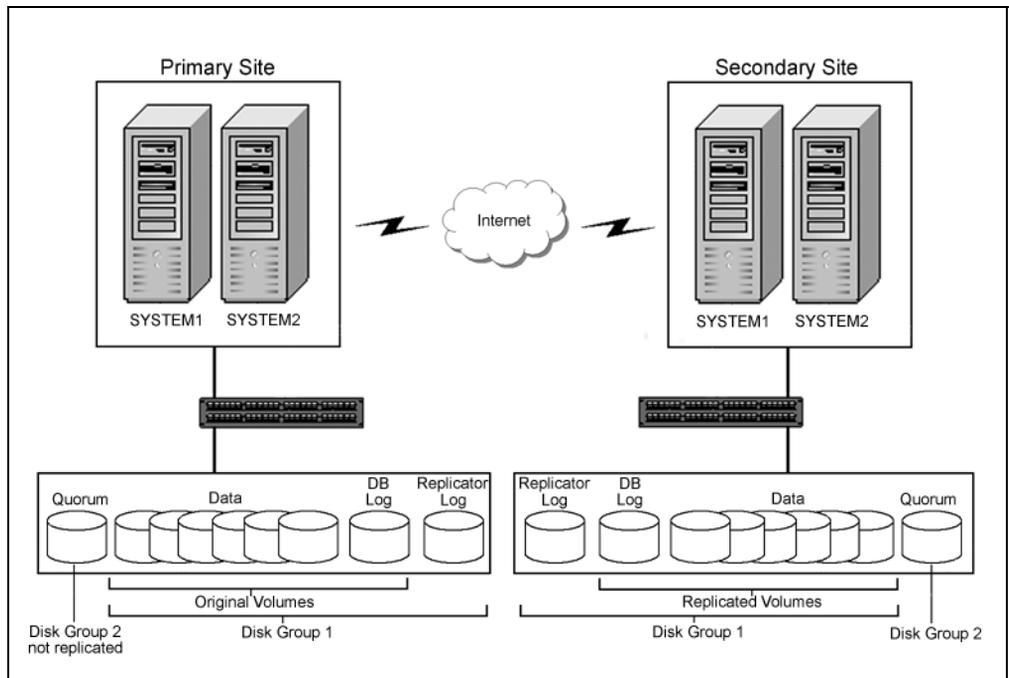
## About the SFW-MSCS-VVR configuration

A typical disaster recovery configuration requires that you have a source host on the primary site and a destination host on the secondary site. The application data is stored on the primary site and replicated to the secondary site by using a tool such as the Veritas Volume Replicator. The primary site provides data and services during normal operation. If a disaster occurs on the primary site and its data is destroyed, a secondary host can take over the role of the primary host to make the data accessible. The application can be restarted on that host.

This Disaster Recovery section includes a SFW-MSCS-VVR configuration. The configuration is described with a generic database application that includes both data and a database log.

The illustration below shows the SFW HA-VVR configuration with MSCS. For a SFW-MSCS-VVR configuration, at least two disk groups are necessary—one for the application and one for the quorum resource volume, which has to be in a separate disk group, as shown in the illustration that follows.

**Figure 14-1** SFW-MSCS-VVR configuration



The quorum volume is not replicated from the primary site to the secondary site. Each site has its own quorum volume. A two-way or four-way mirror is recommended for the quorum volume for redundancy.

## Configuring the quorum device for high availability

Either a single basic disk used as a physical disk resource or a volume located on a three-disk SFW cluster disk group can serve as the MSCS quorum device.

In general, a disk group containing a dedicated, three-way mirrored volume makes an ideal quorum device.

In Microsoft Cluster Service (MSCS) environments, the proper configuration of a quorum device is critical to providing the highest availability with SFW storage.

Using a single disk as the quorum device introduces a nonredundant component into an otherwise highly available system. A failure-tolerant volume used as a quorum device provides a level of availability that is consistent with that of the rest of the cluster.

An SFW cluster disk group containing a volume used as a quorum device should contain that volume only. Any other volumes in that disk group fail over whenever the quorum device changes ownership.

A disk group containing only a three-way mirrored volume makes an ideal quorum device. Such a device tolerates both disk failures, because it is mirrored, and server and interconnect failures, because SFW can import it when the disks and at least one server are running.

For a server to take ownership of a disk group containing the cluster quorum device, SFW must successfully import the disk group, and obtain SCSI reservations on more than half of its disks. Disk groups containing odd numbers of disks are best for use as quorum devices because of this behavior.

# Deploying SFW with MSCS

This chapter describes how to install and configure Storage Foundation for Windows with MSCS in a new installation, using an example two-node active/passive cluster configuration.

The example describes a generic database application in order to present general recommendations that apply to multiple applications. For specific examples of a SFW-MSCS clustering solution see the *Veritas Storage Foundation and High Availability Solutions, Solutions Guide for Microsoft Exchange* and the *Veritas Storage Foundation and High Availability Solutions, Solutions Guide for Microsoft SQL*.

[Table 15-1](#) outlines the high-level objectives for implementing the configuration and the tasks for each objective:

**Table 15-1** Task list

Objectives	Tasks
<a href="#">“Reviewing the requirements”</a> on page 332	<ul style="list-style-type: none"><li>■ Verify hardware and software prerequisites.</li><li>■ Review the configuration requirements.</li></ul>
<a href="#">“Configuring the network and storage”</a> on page 335	<ul style="list-style-type: none"><li>■ Install the operating system on both nodes.</li><li>■ Make necessary networking settings on both nodes.</li></ul>
<a href="#">“Establishing an MSCS cluster”</a> on page 337	<ul style="list-style-type: none"><li>■ Refer to Microsoft documentation for instructions on establishing a cluster under MSCS.</li></ul>

**Table 15-1** Task list (continued)

Objectives	Tasks
“Installing SFW” on page 337	<ul style="list-style-type: none"> <li>■ Verify the driver signing options for Windows 2003 systems</li> <li>■ Install SFW</li> <li>■ Install Cluster Option for Microsoft Cluster Service (MSCS)</li> <li>■ Restore driver signing options for the Windows 2003 systems</li> </ul>
“Creating SFW disk groups and volumes” on page 346	<ul style="list-style-type: none"> <li>■ In SFW on Node A, create at least two dynamic cluster disk groups on the storage—one or more for the application data files and one for the mirrored quorum.</li> <li>■ The disk group for the quorum can be created later, if desired.</li> </ul>
“Setting up a group for the application in MSCS” on page 356	<ul style="list-style-type: none"> <li>■ Create a group within MSCS for the application.</li> <li>■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.</li> </ul>
“Installing the application on cluster nodes” on page 358	<ul style="list-style-type: none"> <li>■ Install the application program files on the local drive of the first node.</li> <li>■ Install files relating to the data and logs on the shared storage.</li> <li>■ Use <b>Move Group</b> to move the cluster resources to the second node.</li> <li>■ Make sure that the volumes on the second node have the same drive letters or mount points that they had on the first node.</li> <li>■ Install the application on the second node.</li> </ul>
“Completing the setup of the application group in MSCS” on page 360	<ul style="list-style-type: none"> <li>■ Refer to the application documentation for help on creating its resource.</li> <li>■ Establish the appropriate dependencies.</li> <li>■ Test the application group by using the <b>Move Group</b> command to move the cluster resources to the other node.</li> </ul>

**Table 15-1** Task list (continued)

Objectives	Tasks
<a href="#">“Implementing a dynamic quorum resource”</a> on page 361	<ul style="list-style-type: none"> <li>■ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier.</li> <li>■ Make that disk group a Volume Manager Disk Group type resource in the default Cluster Group.</li> <li>■ Change the quorum resource to the dynamic mirrored quorum resource.</li> </ul>
<a href="#">“Verifying the cluster configuration”</a> on page 364	<ul style="list-style-type: none"> <li>■ Use the <b>Move Group</b> command to move the cluster resources to the second node. Move them back to the first node.</li> <li>■ Optionally, simulate a failure by turning off the power to the server that has control of the cluster resources.</li> </ul>

## Reviewing the requirements

Verify that the following requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation:

- [“Supported software”](#) on page 332
- [“Disk space requirements”](#) on page 332
- [“System requirements”](#) on page 333

### Supported software

- Veritas Storage Foundation 5.0 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS).
- Windows 2000 Advanced Server or Windows 2000 Datacenter Server (both require SP4 with Update Rollup 1)  
*or*  
 Windows Server 2003 Enterprise Edition or Datacenter Edition (SP1 supported but not required for all editions)  
*or*  
 Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)  
*or*  
 Windows Server 2003 for Intel Xeon (EM64T) or AMD Opteron: Enterprise x64 Edition or Datacenter x64 Edition

### Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Installation on a non-system drive requires space on both the system drive and the non-system drive.

The following table summarizes disk space requirements for SFW.

**Table 15-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
SFW + all options + client components	1240 MB	Non-system space: 1240 MB System space: 265 MB
SFW + all options	980 MB	Non-system Space: 980MB System space: 225 MB

**Table 15-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
Client components	420 MB	Non-system space: 420 MB System space: 80 MB

## System requirements

- The configuration described requires shared disks to support applications that migrate between nodes in the cluster.
- 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.
- Each system requires 1 GB of RAM.
- Systems to be clustered must be configured as part of a Windows 2000 or 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) are required. You also need a static IP address for the cluster itself.

---

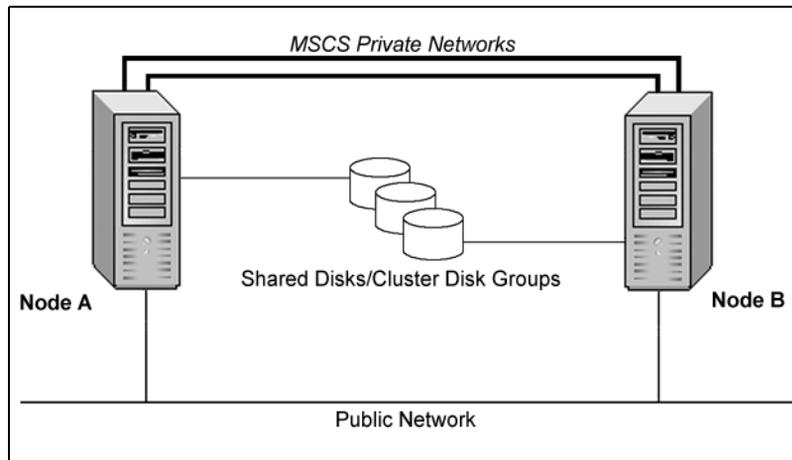
**Note:** Refer to the Hardware Compatibility List on the Symantec Support web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW.

---

## Reviewing the configuration

The example of a new installation with two servers and one storage array in an active/passive configuration is a typical configuration for a cluster. In an active/passive configuration the active node of the cluster hosts the virtual server and 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 example describes a generic database application.

**Figure 15-1** Storage Foundation configuration with MSCS and two servers



This configuration does not include DMP. For information about DMP and clustering, see “[Adding DMP to a clustering configuration](#)” on page 169.

Key points about the configuration:

- An MSCS cluster must be running to install SFW. Therefore, you need to set up the hardware and install the operating system and MSCS on both systems and establish the MSCS 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 active cluster resources to the other node, and install on the now inactive node.
- 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 single 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 a quorum disk fails, a mirror on another disk (another plex) takes over and the resource remains online. For 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 your environment is configured to move the quorum disk to a dynamic mirrored quorum volume; this enables you to verify that the application is working in the cluster before adding the dynamic quorum volume.

## Configuring the network and storage

Use the following procedures to configure the hardware and verify DNS settings. You must also install the operating system and MSCS on the systems. MSCS is automatically installed with Windows Server 2003 Enterprise Edition; Windows 2000 requires you to complete the additional step of installing MSCS from the Microsoft product CD.

Microsoft recommends waiting until after the cluster is established on the first node before connecting the second node to the storage array to avoid data corruption.

### 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.
- 5 Reboot each system. Verify that each system recognizes the attached shared disk.
- 6 Use Windows Disk Management on each system to verify that the attached shared disks are visible.

**To verify the DNS settings for all systems**

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Open **Network and Dial-up Connections**.
- 3 Ensure the public network adapter is the first bound adapter by selecting the **Advanced** menu and clicking **Advanced Settings**.
- 4 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.
- 5 Click **OK**.
- 6 In the Network and Dial-up 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.
- 7 From the status window, click **Properties**.
- 8 On the **General** tab, select the **Internet Protocol (TCP/IP)** check box and click **Properties**.
- 9 Select the **Use the following DNS server addresses** option.
- 10 Verify the correct value for the IP address of the DNS server.
- 11 Click **Advanced**.
- 12 On the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 13 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 14 Click **OK**.

## Establishing an MSCS cluster

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

### To establish an MSCS cluster (general guidelines)

- 1 Configure the shared storage and create a partition with a drive letter for the cluster quorum disk. 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.
- 2 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.
- 3 Connect the shared storage to the second node.
- 4 Add the second node (SYSTEM2) using Cluster Administrator on that system.
- 5 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

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 338.
- Installing the product  
See “[Installing Veritas Storage Foundation for Windows](#)” on page 339.

- Performing post-installation tasks  
See “[Post-installation tasks](#)” on page 344.

## Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options  
See “[Changing the driver signing options](#)” on page 338.
- Moving the Online Groups  
See “[Moving the online groups](#)” on page 339.

### Changing the driver signing options

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

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

**Table 15-3** 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

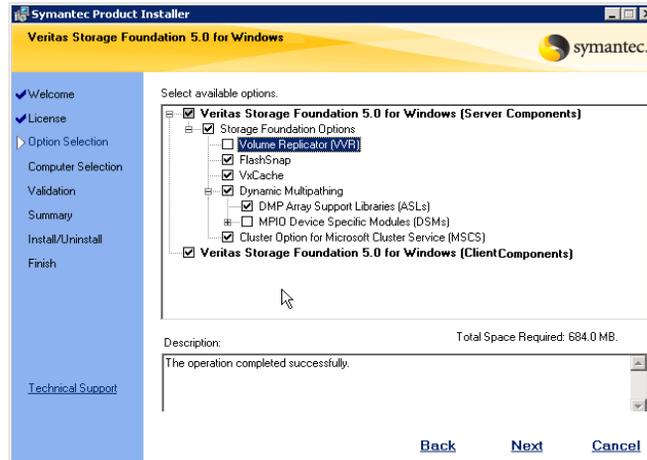
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the default language for your installation and click **OK**. The SFW Select Product screen appears.

- 3 Click **Storage Foundation 5.0 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)** and any additional options applicable to your environment.



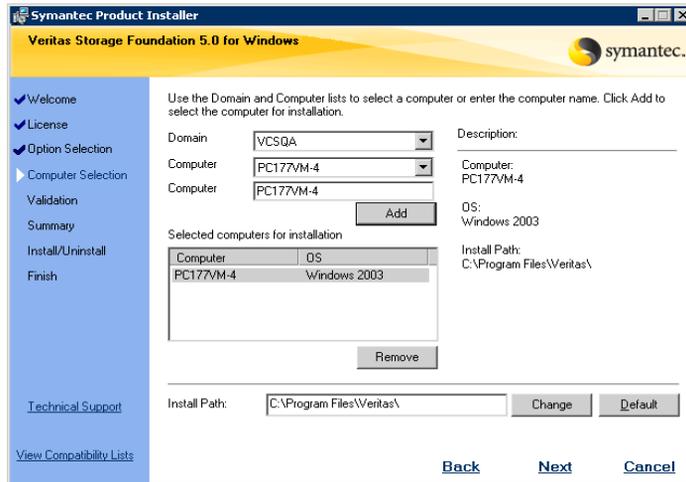
Displayed at the bottom of the screen is the total hard disk space required for the installation and a description of an option.

Note that under Veritas Dynamic Multi-pathing, you can select either:

- DMP Array Support Libraries (ASLs)
- DMP Device Specific Modules (DSMs)

- 11 Click **Next**.
- 12 Verify that the **Veritas Storage Foundation 5.0 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 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.

#### **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.

## About the PBX resource

The installation of Veritas Storage Foundation for Windows also includes the installation of the PBX resource. The PBX resource enables communication between system components such as the VEA and the StorageAgent. MSCS starts the PBX resource automatically and requires no manual intervention.

## Post-installation tasks

You must perform the following post-installation tasks:

- Moving the Online Groups  
See "[Moving the online groups](#)" on page 345.
- Completing the SFW Installation  
See "[Completing the SFW installation](#)" on page 345.
- Resetting the driver signing options  
See "[Resetting the driver signing options](#)" on page 345.

## 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 337.

## 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 SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and dynamic volumes for the application on the shared storage. A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different volume layouts. Configuring disk groups and volumes involves the following tasks:

- [Planning disk groups and volumes](#)
- [Creating dynamic cluster disk groups](#)
- [Creating dynamic volumes](#)

### Planning disk groups and volumes

Decide how you want to organize the disk groups and the number and type of volumes you want to create.

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 that are needed

The number of disk groups depends on your application and 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. Typically, a main organizational unit in your application, such as the storage group in Microsoft Exchange, is contained in a single disk group. You will also need a disk group with three disks for the mirrored quorum resource. If possible, use small disks.

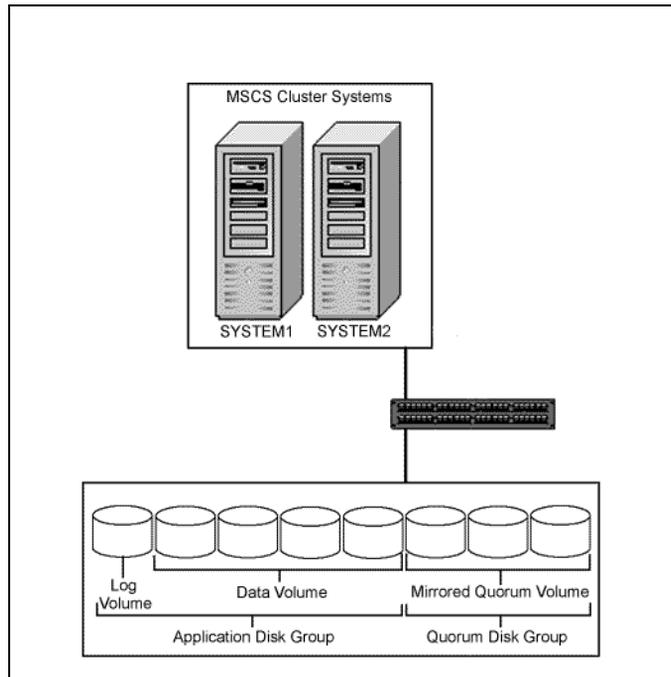
Microsoft recommends 500 MB for the quorum disk.

#### **Recommendations:**

- Use mirrored volumes for logs.
- Use striped or mirrored striped volumes for data.

The following illustration shows a typical setup of shared storage disks for a clustered database application and a dynamic mirrored quorum resource. The log volume is on a separate disk. The log and data volumes are in the application dynamic cluster disk group. The dynamic mirrored quorum is in a separate disk group and has a minimum of two disks, but three are recommended for added fault tolerance.

**Figure 15-2** MSCS clustered database with disks for data, the log, and the quorum resource

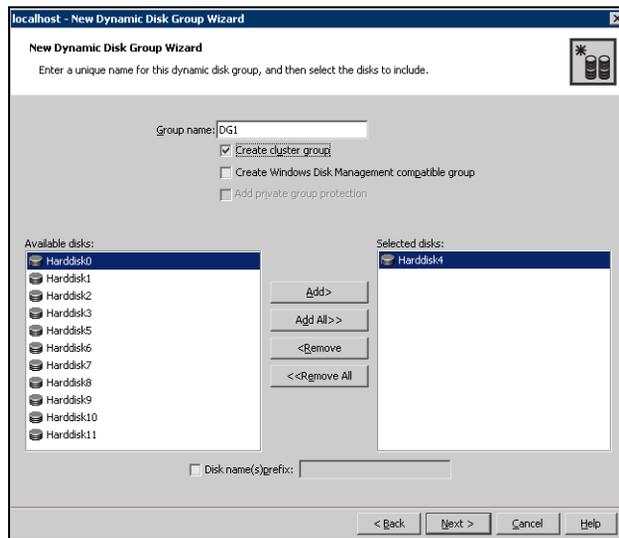


## Creating dynamic cluster disk groups

Follow the steps in this section to create one or more disk groups for your application.

### To create a dynamic (cluster) disk group

- 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:



- Enter the disk group name (for example, DG1).
- 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.
  - Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
  - 8 Click **Finish** to create the new disk group.

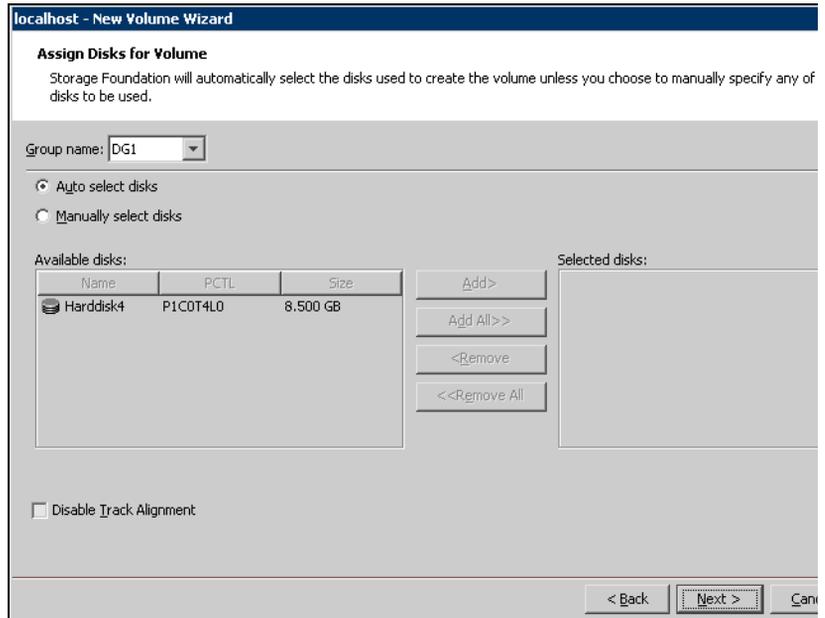
## Creating dynamic volumes

Once the disk groups are created, make the disks within them usable by creating the dynamic volumes that will store data.

### To create dynamic volumes

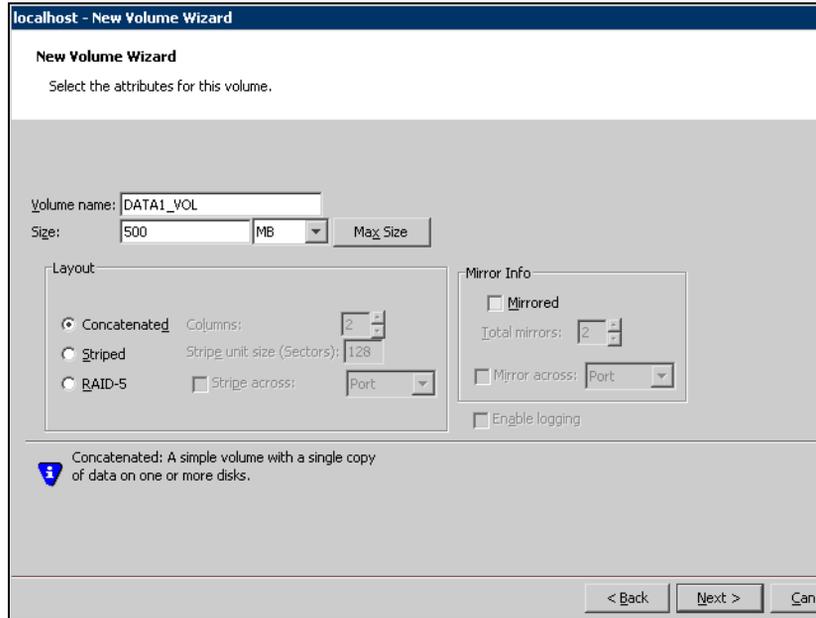
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog 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.
- 5 At the New Volume wizard opening screen, click **Next**.

6 Select the disks for the volume:



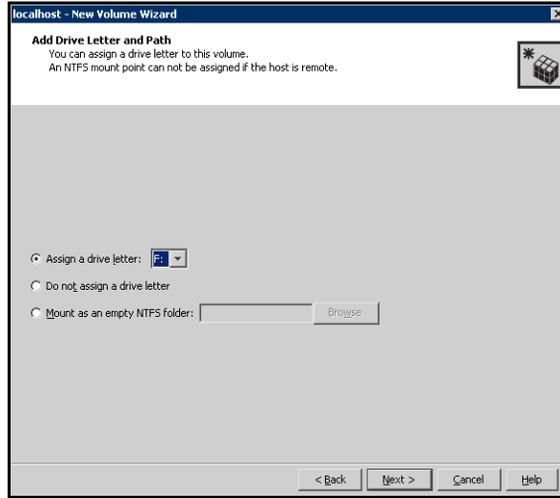
- a Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- b Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
- c 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.
- d Click **Next**.

7 Specify the volume attributes:



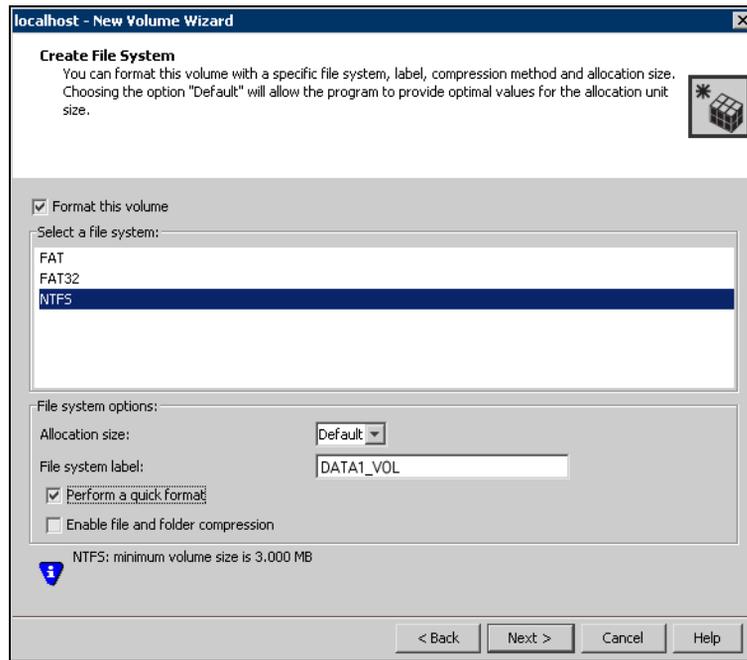
- a Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- b Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
- c Provide a size for the volume.  
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.
- d In the **Mirror Info** area, select the appropriate mirroring options.
- e Verify that **Enable Logging** is not selected.
- f Click **Next**.

- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume:



- a 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 a disaster recovery configuration, for the Replicator Log volume:  
Select **Do not assign a drive letter**.
- b Click **Next**.

9 Create an NTFS file system.



- a Make sure the **Format this volume** checkbox is checked, with the following exception: For the Replicator Log volume only in a disaster recovery configuration, clear the **Format this volume** check box.
  - b Click **NTFS**.
  - c Select an allocation size or accept the Default.
  - d The file system label is optional. SFW makes the volume name the file system label.
  - e Select **Perform a quick format** if you want to save time.
  - f 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.
  - g Click **Next**.
- 10 Click **Finish** to create the new volume.
- 11 Repeat these steps to create additional volumes.

---

**Note:** Create the cluster disk group and volumes on the first node of the cluster only.

---

## Setting up a group for the application in MSCS

The next task is to use Cluster Administrator in MSCS to set up a group for the application that will contain the SFW disk group or groups that were created for the application. The SFW disk groups will be added to the MSCS application group as Volume Manager Disk Group resources.

After the application is installed on both nodes and its accompanying files are placed on the shared storage, complete the setup of the application group by adding the application itself as a resource and any other resources that are required. Dependencies need to be set between the resources in the group. Information on this task is included in “[Completing the setup of the application group in MSCS](#)” on page 360.

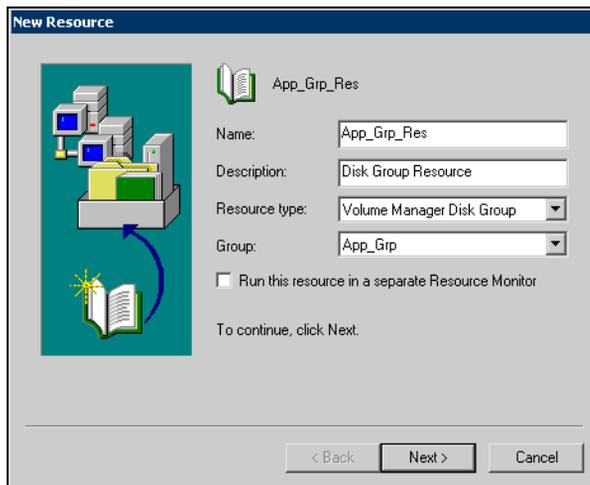
### To set up the application group

- 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**, click **New**, and click **Group**.
- 3 When the New Group dialog box appears, specify a name for the group (for example, App\_Grp).
- 4 Click **Next** to continue.
- 5 When the Preferred Owners dialog box appears, make sure that all the preferred owners are added to the **Preferred Owners** list.
- 6 Click **Finish** to create the group.

### To create a disk group resource

- 1 From the Cluster Administrator configuration tree, right-click on the MSCS group that you have created (App\_Grp) and click **New > Resource**.

- 2 In the New Resource dialog box, specify a name for the disk group resource and a description for the resource, if necessary.



- 3 Select **Volume Manager Disk Group** from the **Resource type** list.
- 4 Select **App\_Grp** from the **Group** list.
- 5 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 Disk Group resource.
- 8 When the Volume Manager Disk Group Parameters dialog box appears, select the disk group.
- 9 Click **Finish**.
- 10 Click **OK**.
- 11 Bring the resources online.

## Installing the application on cluster nodes

The application program files need to be installed on the same local drive on all the cluster nodes. The application data and log files or other files related to the application data are installed on the shared storage.

### 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, such as Microsoft SQL Server, install on both nodes at once.
- All nodes of the clustered application need to share the same virtual name and IP address.
- Do not to accept the default locations for the application data and log files. Instead, set the paths for these files to the drive letters or mount points of the volumes created in [“Creating dynamic volumes”](#) on page 350.

### Pointers for installing the application on the second node

- Use the **Move Group** command to move the cluster resources to the second node.
- Verify that the volumes on shared storage can be accessed from the second node using the same drive letters or mount points that were assigned when they were created on the first node. To change a drive letter or mount point, see [“To add or change a drive letter or mount point”](#) on page 358.
- 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. You 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.
- 3 To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.

- 4 To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Change the drive letter and click **OK**.
- 5 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, and click **OK** to mount the volume.

---

**Note:** A mount point is also referred to as a “drive path.”

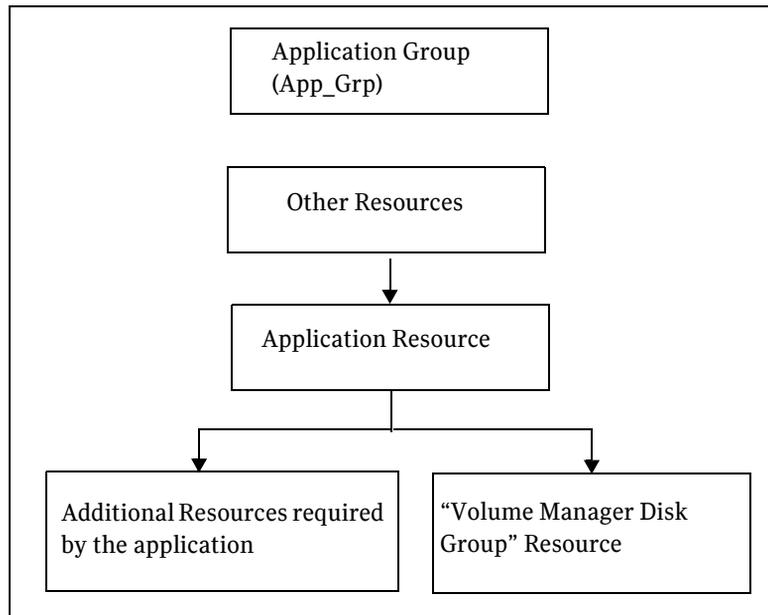
---

- 6 To change a mount point, you must remove it and recreate it ([step 5](#)). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.

## Completing the setup of the application group in MSCS

The additional steps in this section make the application group functional in MSCS. The application resource needs to be added, as well as any other resources that are associated with the application. Also, dependencies need established for the resources. This section presents a high-level summary of the process for completing the application group setup.

- Before creating the application resource, make sure that all the other resources, including the disk group resource and any additional application resources, are online.
- Refer to the application documentation for information on creating its resource and additional resources that may be required. You will need to create an IP address resource and a Network name resource in addition to the Volume Manager Disk Group resource you created earlier.  
Note that when creating the application resource, on the Dependencies screen, select the **Volume Manager Disk Group** resource from “Available Resources” and add it to “Resource Dependencies.”
- The following dependency chart indicates the dependencies that are established.

**Figure 15-3** Application group dependencies

- **Testing:** After the application group is set up, test it by using the **Move Group** command to move the cluster resources to another node and then move them back.

## Implementing a dynamic quorum resource

Although Symantec recommends implementing a dynamic quorum resource in order to take full advantage of the Storage Foundation functionality, it is not a required task:

- [“Creating a dynamic cluster disk group for the Quorum Resource with mirrored volume”](#) on page 362
- [“Creating the quorum resource for the cluster group”](#) on page 362
- [“Changing the quorum resource to a dynamic mirrored quorum resource”](#) on page 363

---

**Note:** If you are using DMP, you must create a dynamic quorum resource in order for the groups to fail over properly.

---

## Creating a dynamic cluster disk group for the Quorum Resource with mirrored volume

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** checkbox, and specify three mirrors. For full details about creating cluster disk groups and volumes, see “[Creating SFW disk groups and volumes](#)” on page 346.

---

**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.

---

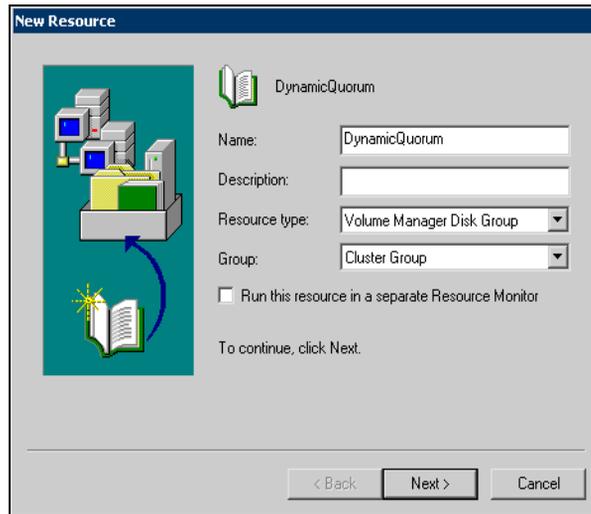
## Creating the quorum resource for the cluster group

Create a quorum resource for the cluster group to protect the quorum during failover.

### To create the quorum resource

- 1 From Cluster Administrator (**Start > Control Panel > Administrative Tools > Cluster Administrator**), verify that the Cluster Group is online on the same node where you created the disk group.

- 2 Create the quorum resource, right-click the **Cluster Group**, click **New**, and click **Resource**.



- 3 When the New Resource dialog box appears, specify a name for the quorum resource (QuorumDG) and, if necessary, add a description about the resource.
- 4 Select **Volume Manager Disk Group** from the **Resource type** list.
- 5 Select **Cluster Group** from the **Group** list.
- 6 Click **Next**.
- 7 In the Possible Owners dialog box, click **Next**.
- 8 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a quorum resource.
- 9 When the Volume Manager Disk Group Parameters dialog box appears, select the disk group.
- 10 Click **Finish**.
- 11 Bring the newly added resource online.

## Changing the quorum resource to a dynamic mirrored quorum resource

The last step in this process is to change the quorum resource to a dynamic mirrored quorum.

#### To change the quorum resource to a dynamic mirrored quorum

- 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 added in “[Creating the quorum resource for the cluster group](#)” on page 362.
- 4 Click **OK**.

## 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 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 in a campus cluster

This chapter covers the following topics:

- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Installing and configuring the hardware](#)
- [Establishing the cluster under MSCS](#)
- [Installing SFW](#)
- [Creating disk groups and volumes](#)
- [Setting up a group for the application in MSCS](#)
- [Installing the application on the cluster nodes](#)
- [Completing the setup of the application group in MSCS](#)
- [Changing the quorum resource to a dynamic quorum resource](#)
- [Verifying the cluster configuration](#)

This chapter presents an MSCS campus clustering example with a two-node cluster.

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

**Table 16-1** Task list

Objectives	Tasks
<a href="#">“Reviewing the prerequisites”</a> on page 370	■ Verify hardware and software prerequisites.

**Table 16-1** Task list

Objectives	Tasks
“ <a href="#">Reviewing the configuration</a> ” on page 372	<ul style="list-style-type: none"> <li>■ Review the configuration requirements.</li> <li>■ Overview of MSCS campus cluster, and recovery scenarios</li> </ul>
“ <a href="#">Installing and configuring the hardware</a> ” on page 381	<ul style="list-style-type: none"> <li>■ 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.</li> <li>■ Install the hardware in the same manner for Site B.</li> </ul>
“ <a href="#">Establishing the cluster under MSCS</a> ” on page 382	<ul style="list-style-type: none"> <li>■ Install and configure the operating system and MSCS on Server A.</li> <li>■ On Site A, configure the storage and create a partition for the cluster quorum disk.</li> <li>■ Create the first node of the cluster on Server A.</li> <li>■ Install and configure the operating system and MSCS on Server B.</li> <li>■ Connect the two nodes.</li> <li>■ Create the second node of the cluster on Server B.</li> <li>■ 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.</li> </ul>
“ <a href="#">Installing SFW</a> ” on page 385	<ul style="list-style-type: none"> <li>■ Install SFW on Node A (Node B active).</li> <li>■ Install SFW on Node B (Node A active).</li> </ul>
“ <a href="#">Creating disk groups and volumes</a> ” on page 393	<ul style="list-style-type: none"> <li>■ 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.</li> </ul>
“ <a href="#">Setting up a group for the application in MSCS</a> ” on page 403	<ul style="list-style-type: none"> <li>■ Create a group within MSCS for the application.</li> <li>■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.</li> </ul>

**Table 16-1** Task list

Objectives	Tasks
<p>“Installing the application on the cluster nodes” on page 406</p>	<ul style="list-style-type: none"> <li>■ Install the application program files on the local drive of the first node.</li> <li>■ Install files relating to the data and logs on the shared storage.</li> <li>■ Use <b>Move Group</b> to move the cluster resources to the second node.</li> <li>■ Make sure that the volumes on the second node have the same drive letters or mount points as they had on the first node.</li> <li>■ Install the application on the second node.</li> </ul>
<p>“Completing the setup of the application group in MSCS” on page 408</p>	<ul style="list-style-type: none"> <li>■ Refer to the application documentation for help on creating its resource.</li> <li>■ Establish the appropriate dependencies.</li> <li>■ Test the application group by using <b>Move Group</b> to move the cluster resources to the other node.</li> </ul>
<p>“Changing the quorum resource to a dynamic quorum resource” on page 410</p>	<ul style="list-style-type: none"> <li>■ Create a dynamic disk group for the quorum with a mirrored volume.</li> <li>■ Make that disk group into a Volume Manager Disk Group type resource in the default Cluster Group.</li> <li>■ Change the quorum resource to the dynamic mirrored quorum resource.</li> </ul>
<p>“Verifying the cluster configuration” on page 413</p>	<ul style="list-style-type: none"> <li>■ Verify the cluster configuration by switching service groups or shutting down an active cluster node</li> </ul>

## Reviewing the prerequisites

Reviewing the prerequisites and the configuration allows you to gain an overall understanding of the configuration and its requirements.

### Supported software

- Veritas Storage Foundation 5.0 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS).
- Windows 2000 Advanced Server or Datacenter Server (all require Service Pack 4 with Update Rollup1)  
*or*  
Windows Server 2003 Enterprise Edition or Datacenter Edition (SP1 supported but not required for all editions)  
*or*  
Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)  
*or*  
Windows Server 2003 for Intel Xeon (EM64T) or AMD Opteron: Enterprise x64 Edition or Datacenter x64 Edition

### System 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 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 2000 or 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:** Refer to the Hardware Compatibility List on the Symantec Support web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW.

---

## Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Installation on a non-system drive requires space on both the system drive and the non-system drive.

The following table summarizes disk space requirements for SFW.

**Table 16-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
SFW + all options + client components	1240 MB	Non-system space: 1240 MB System space: 265 MB
SFW + all options	980 MB	Non-system Space: 980MB System space: 225 MB
Client components	420 MB	Non-system space: 420 MB System space: 80 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.

---

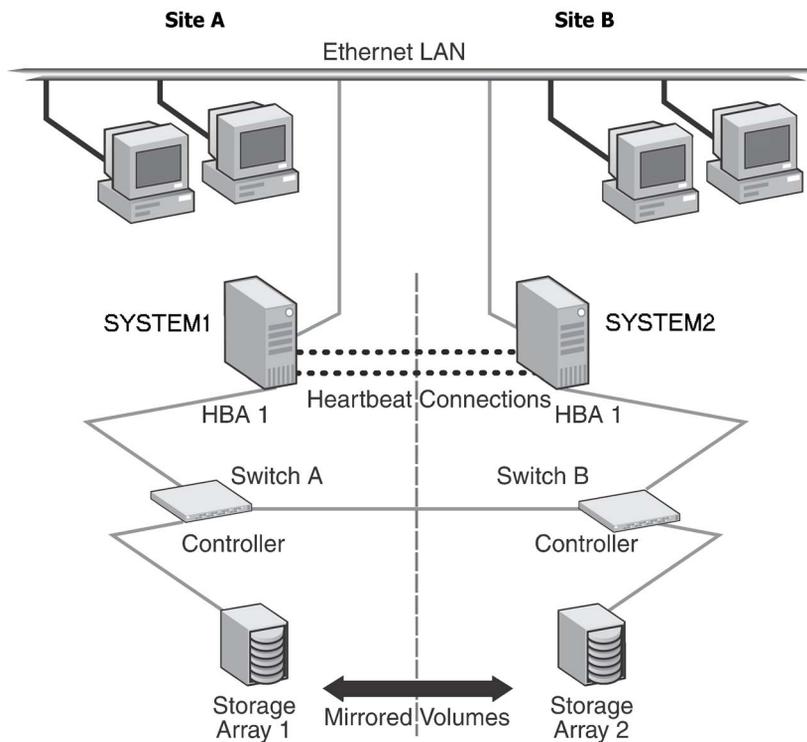
## Reviewing the configuration

This configuration example describes the most common configuration, 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 374
- “[MSCS campus cluster failure scenarios](#)” on page 375

**Figure 16-1** 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 with the same number of disks and contains mirrored data of the storage on the other array. The example describes a generic database application.

Plan for an equal number and size of disks on the two sites, because 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 are using 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 configuration does not include DMP. For instructions on how to add DMP to a clustering configuration, see the DMP chapter, “[Adding DMP to a clustering configuration](#)” on page 169.

When you are installing SFW and MSCS together, remember the following:

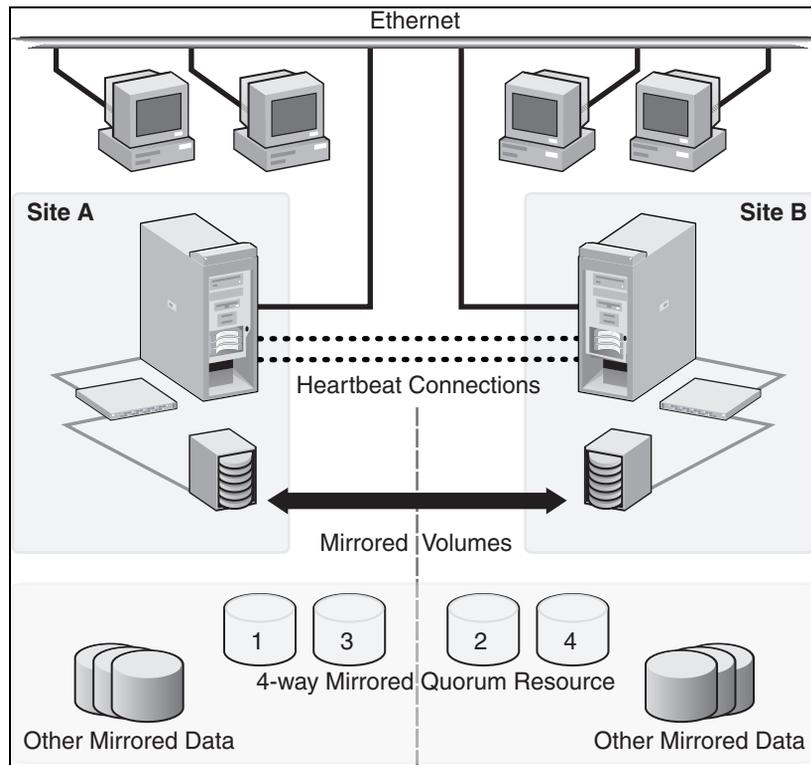
- An MSCS cluster must be running to install SFW.  
You need to set up the hardware and install the operating system and MSCS on all systems and establish the MSCS cluster before installing SFW. Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Use a “rolling 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.
- Wait until the end of the configuration process to enable the dynamic mirrored quorum.  
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. Wait until after you have your application installed, running, and tested with the cluster to actually convert the original basic quorum disk to the dynamic quorum volume. By waiting until the end of the process to convert from the basic physical disk quorum to the dynamic mirrored volume, you can make sure that the application is working first with the cluster and then add the dynamic quorum volume.
- Using SFW also offers other advantages over using MSCS alone because SFW has superior features for managing your disks and volumes, compared with Disk Management, the default disk and volume management utility that comes with Windows 2000 and Windows Server 2003.  
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 MSCS

The following example of an MSCS campus cluster configuration features mirrored storage across clusters and a mirrored quorum resource. The illustration shows a 4-way mirrored quorum that 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 16-2 Typical MSCS campus clustering configuration



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 will fail, because none of the cluster servers will be able to 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 378.

The following table lists failure situations and the outcomes that occur:

**Table 16-3** List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
<b>1) Application fault</b> 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.
<b>2) Server failure (Site A)</b> 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.
<b>3) Server failure (Site B)</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) results in no interruption of service to the active site (Site A).

**Table 16-3** List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
<p><b>4) Partial SAN network failure</b></p> <p>May mean that SAN fiber channel cables were disconnected to Site A or Site B Storage.</p>	<p>No interruption of service.</p>	<p>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 have no effect to any cluster resources running on that node, because the underlying DMP solution should seamlessly handle the SAN fiber path failover.</p>
<p><b>5) Private IP Heartbeat Network Failure</b></p> <p>May mean that the private NICs or the connecting network cables failed.</p>	<p>No interruption of service.</p>	<p>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 have no effect on the cluster software and the cluster resources, because the cluster software will simply route the heartbeat packets through the public network.</p>
<p><b>6) Public IP Network Failure</b></p> <p>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><b>7) Public and Private IP or Network Failure</b></p> <p>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 had ownership of 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 will self-terminate because it has lost the cluster arbitration for the quorum resource.</p>

**Table 16-3** List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
<p><b>8) Lose Network Connection (SAN &amp; LAN), failing both heartbeat and connection to storage</b></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. Disks on the same node are functioning. Mirroring is not working.</p>	<p>The node/site that had ownership of 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 will self-terminate 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><b>9) Storage Array failure on Site A, or on Site B</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. 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 have no 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><b>10) Site A failure (power)</b></p> <p>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 onlined on the remaining live site without manual intervention.</p>
<p><b>11) Site B failure (power)</b></p> <p>Means that all access to site B, including server and storage, is lost.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>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.</p>

## Dealing with a failover situation

In summary, the site scenarios that can occur when there is a cluster server failure include the following two possibilities:

- If the site not owning the quorum volume and the cluster goes offline, the quorum and data volumes will stay online at the other site and other cluster resources will stay online or move to that site. Storage Foundation for Windows will allow 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 will not be able to 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 a cluster disk group containing the MSCS quorum is manually imported to the secondary (failover) server when the primary server is still active, this will cause a split-brain situation. 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 provides an explanation of 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 has to 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 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 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 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 MSCS Cluster Administrator, 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.

## Installing and configuring the hardware

This topic gives the general steps for the hardware installation. For complete details on installing the hardware, refer to the manufacturers' hardware documentation.

### To set up the hardware

- 1 Install three network interface cards on Site A, one for the public network and two for the private network.
- 2 Install the host adapter on Site A.
- 3 Install the switch and the storage array on Site A.
- 4 Connect the server and the storage array to the SAN. Do not connect the switch to the other switch at Site B.

It is recommended that you wait until after the MSCS cluster is established on the first node before making the hardware connections to the second node.

Install the hardware in the same manner for Site B.

## Establishing the cluster under MSCS

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 382
- [“Configuring the shared storage and creating a partition for the Cluster quorum disk”](#) on page 383
- [“Creating the first node of the cluster on server A”](#) on page 383
- [“Installing and configuring the operating system and MSCS on server B”](#) on page 383
- [“Connecting the two nodes”](#) on page 383
- [“Creating the second node of the cluster on server B”](#) on page 384
- [“Verifying the cluster configuration”](#) on page 384

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 2000 or Windows Server 2003 operating system on Server A.
- 2 Identify the static Server A network addresses for the public and private networks in the cluster. This task is done through the Internet Protocol (TCP/IP) window.
- 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 Set up a cluster account for the cluster using **Administrative Tools > Active Directory > Users and Computers**. Microsoft recommends having a separate user account under which the cluster can run.
- 5 If you are running Windows 2000, install MSCS from the Windows 2000 installation media.  
If you are using Windows Server 2003, MSCS will be already installed.

## Configuring the shared storage and creating a partition for the Cluster quorum disk

- 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

Use **Cluster Administrator** to create the first node of the cluster on Server A. Refer to the Microsoft documentation for details.

Once 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 installation steps for Server B, using [step 1](#) through [step 5](#) that are listed for installing and configuring the operating system on Server A. See the section "[Installing and configuring the operating system and MSCS on server A](#)" on page 382.

## Connecting the two nodes

Make the necessary connections between the two sites. The cluster is already active on Server A, so MSCS 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.

## Creating the second node of the cluster on server B

Use **Cluster Administrator** to 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.

## 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 385.
- Installing the product  
See “[Installing Veritas Storage Foundation for Windows](#)” on page 387.
- Performing post-installation tasks  
See “[Post-installation tasks](#)” on page 391.

## Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options  
See “[Changing the driver signing options](#)” on page 385.
- Moving the Online Groups  
See “[Moving the online groups](#)” on page 386.

## Changing the driver signing options

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

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

**Table 16-4** Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
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.

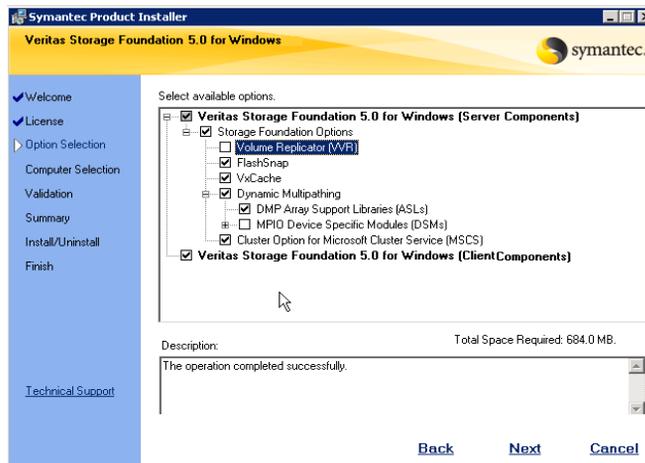
### To install the product

- 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 5.0 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)** 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 either:

- DMP Array Support Libraries (ASLs)
- DMP Device Specific Modules (DSMs)

- 11 Click **Next**.
- 12 Verify that the **Veritas Storage Foundation 5.0 for Windows (Client Components)** check box is checked, to install the client component and click **Next**.



- 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 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.

#### **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.

## About the PBX resource

The installation of Veritas Storage Foundation for Windows also includes the installation of the PBX resource. The PBX resource enables communication between system components such as the VEA and the StorageAgent. MSCS starts the PBX resource automatically and requires no manual intervention.

## Post-installation tasks

You must perform the following post-installation tasks:

- Moving the Online Groups  
See "[Moving the online groups](#)" on page 392.
- Completing the SFW Installation  
See "[Completing the SFW installation](#)" on page 392.
- Resetting the driver signing options  
See "[Resetting the driver signing options](#)" on page 392.

## 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 385.

## 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
- Types of volumes required and location of the plex of each volume in the storage array

---

**Note:** For campus clusters, each disk group *must* contain an equal number of disks on each site.

---

---

**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.

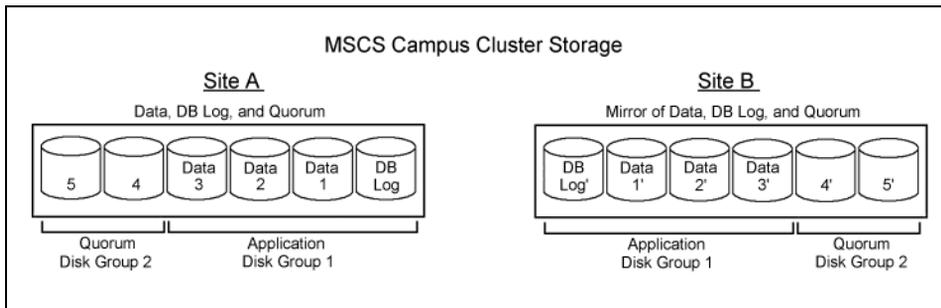
---

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.

The illustration that follows 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 the example, a four-way mirror 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 16-3 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 395
- [“Creating a volume”](#) on page 397

### Considerations when creating new volumes

- For campus clusters, when creating 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

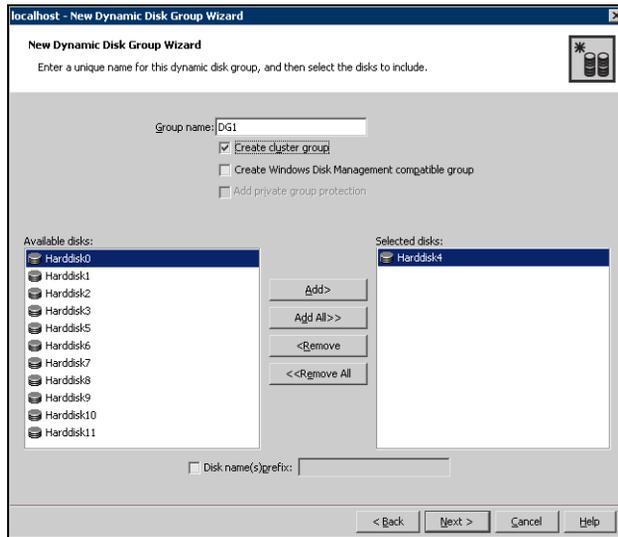
- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.  
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

### To create a dynamic (cluster) disk group

- 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:



- Enter the disk group name (for example, DG1).
  - 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.
  - 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.

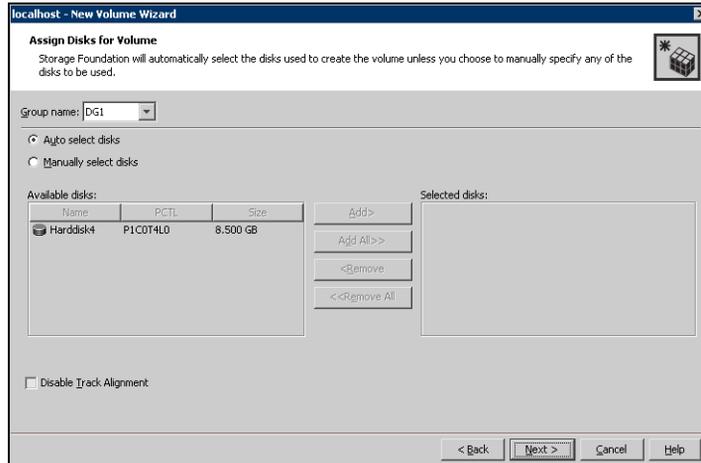
## Creating a volume

Use the following procedure to create dynamic volumes.

### To create dynamic volumes

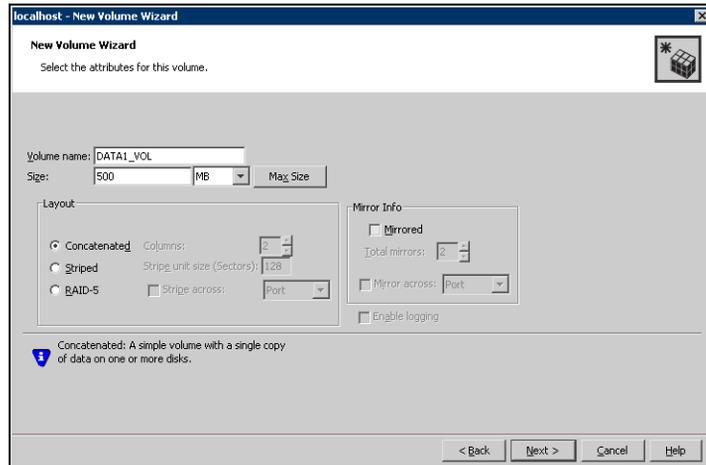
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog 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.
- 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.

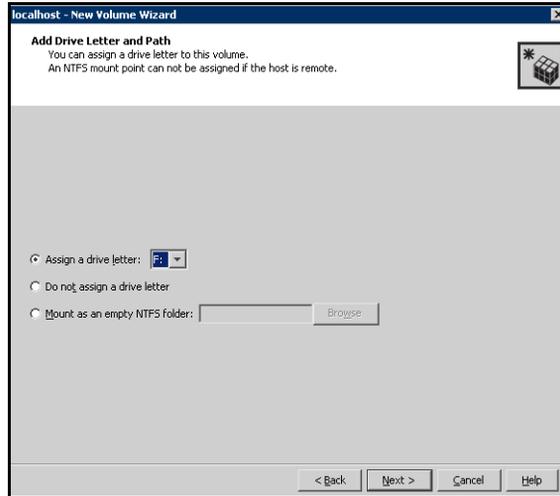
- 8 Click **Next**.
- 9 Specify the volume attributes.



- 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

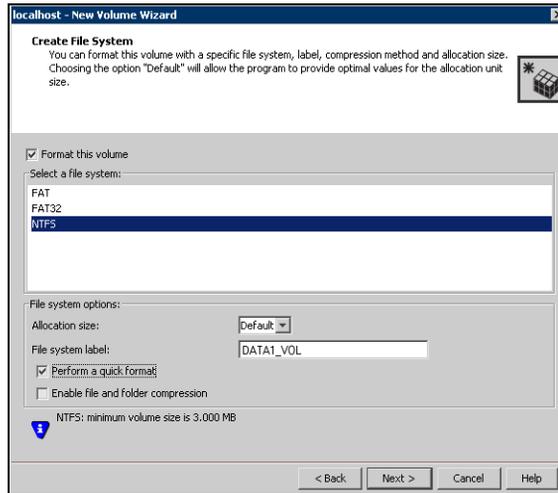
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

- 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.

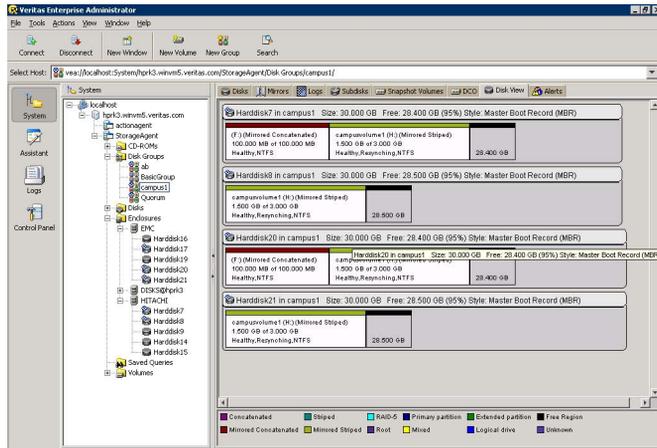
14 Repeat these steps to create additional volumes.

---

**Note:** Create the cluster disk group and volumes on the first node of the cluster only.

---

Figure 16-4 View of disks with volumes in VEA Console



## Setting up a group for the application in MSCS

Use Cluster Administrator in MSCS to set up a group for the application that will contain the SFW disk group or groups that were created for the application. The SFW disk groups will be added to the MSCS application group as Volume Manager Disk Group resources.

After the application is installed on both nodes and its accompanying files are placed on the shared storage (which in this case is shared across two storage arrays), complete the setup of the application group by adding the application itself as a resource and any other resources that are required. Dependencies need to be set between the resources in the group as described in [“Completing the setup of the application group in MSCS”](#) on page 408.

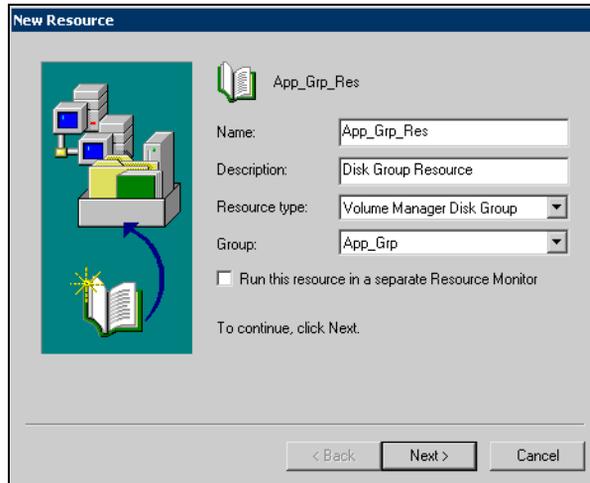
### To set up the application group

- 1 Click **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator** to start Cluster Administrator in MSCS.
- 2 Make sure you are connected to the appropriate cluster.
- 3 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** from the menu.  
The New Group window appears.
- 4 Specify a name for the group in the **Name** field (for example, App\_Grp). Click **Next** to continue.
- 5 The Preferred Owners page appears. Make sure that all the preferred owners are added to the **Preferred Owners** list.
- 6 Click **Finish** to create the group.

Proceed to add resources to the group in the next task.

### To add SFW disk groups as resources to the application group

- 1 Right-click on the MSCS group that you have created for the application and click **New > Resource**. The New Resource window appears.



- a Specify a name for the disk group resource in the **Name** field.
- b If required, you can add a description about the resource in the **Description** field. “Disk Group Resource” is an appropriate description.
- c 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.

---

- d If necessary, use the drop-down list to select the appropriate MSCS group; the group should already be selected.
  - e Generally, make sure that **Run this resource in a separate Resource Monitor** is not checked.
  - f Click **Next**.  
The Possible Owners screen appears.
- 2 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.  
The Dependencies screen appears.
  - 3 On the Dependencies screen, click **Next**. (You do not need to set any dependencies for a disk group resource.)

- 4 Make sure the appropriate SFW cluster disk group is selected from the drop-down list for the resource, and click **Finish**.

If there is more than one disk group for the application, repeat the steps in this process. You can also add more resources at this time, as required by the application, or wait until after the application is installed. You will not be able to add the application resource until after the application is installed on both nodes.

---

**Note:** If you have the cluster disk group created for the quorum and want to do the necessary steps for converting the dynamic mirrored volume to the quorum resource for the cluster now, you can do so. Refer to the section “[Changing the quorum resource to a dynamic quorum resource](#)” on page 410. It is recommended, but not necessary, that you wait until the application is installed and working with the cluster before completing the quorum steps.

---

## 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.

### 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 need to share the same virtual name and IP address.
- Remember not to accept the default location for the application data and log files when installing the application. 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, use the **Move Group** command to 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”](#) in the next section.
- 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.
  - a To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter.

- b To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Change the drive letter.
- c 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.”

---

- d To change a mount point, you must remove it and recreate it ([step c](#)). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.
- e Click **OK**.

## Completing the setup of the application group in MSCS

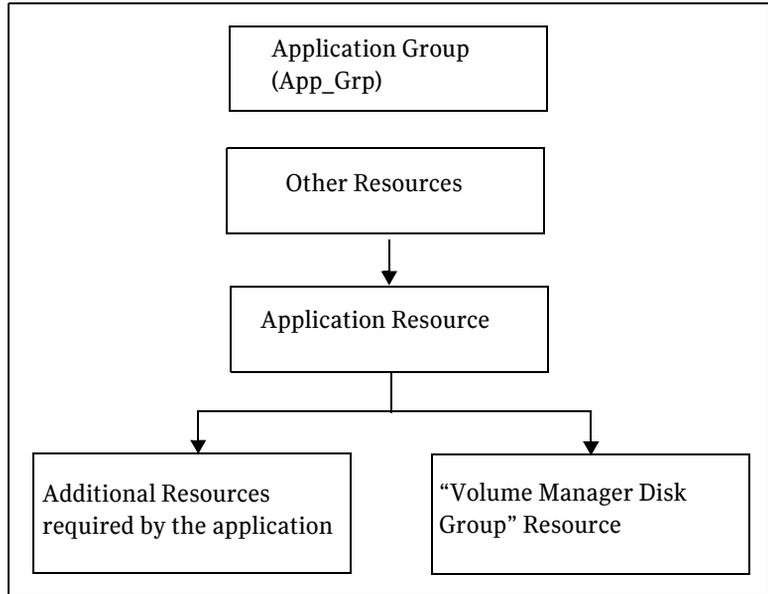
To make the application group functional in MSCS, the application resource needs to be added, as well as any other resources that are associated with the application. Also, dependencies need to be established for the resources. This section presents a summary of the process for completing the application group setup.

- Before creating the application resource, make sure that all the other resources that you created are online, including the disk group resource and any additional application resources.
- Refer to the application documentation for help on creating its resource and additional resources that may be required. You may need to create an IP address resource and a network name resource in addition to the Volume Manager Disk Group resource you created earlier.

When creating the application resource, on the Dependencies screen, select the **Volume Manager Disk Group** resource from “Available Resources” and add it to “Resource Dependencies.”

- The following dependency chart indicates the dependencies that are established.

Figure 16-5 Application group dependencies



- **Testing:** After the application group is set up, test it by using the **Move Group** command to move the cluster resources to another node and then move them back.

## Changing the quorum resource to 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 will 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 410
- [“Making the quorum cluster disk group an MSCS resource”](#) on page 411
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 412

### Creating a dynamic cluster disk group for the quorum, mirrored

If you have not already completed this step, use SFW to create a dynamic disk group for the quorum disks. 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, since the disk group will be used only 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 393.

---

**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.

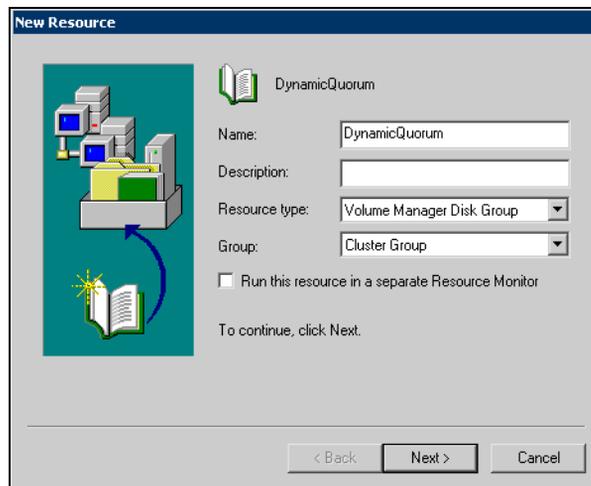
---

## 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.



- a Specify a name for the disk group resource in the **Name** field, such as “QuorumDG.”
- b If necessary, you can add a description about the resource in the **Description** field.
- c Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource type** field drop-down list.

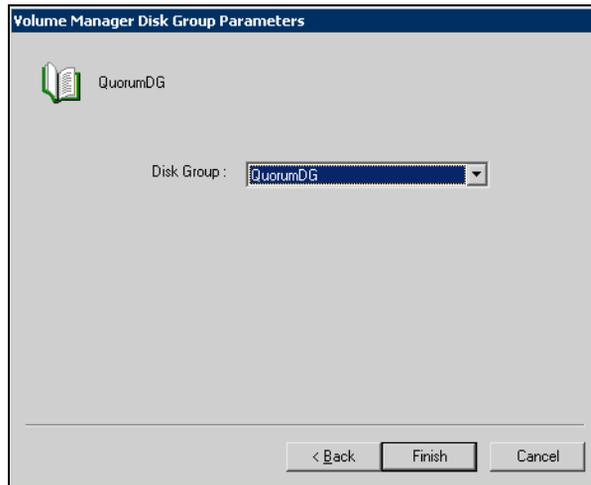
---

**Note:** The resource name has not been changed to Storage Foundation Disk Group.

---

- d Generally, make sure that **Run this resource in a separate Resource Monitor** is not checked.
- e Click **Next**.

- 3 On the Possible Owners screen, by default all the nodes in the cluster are listed as possible owners. Click **Next**.
- 4 On the Dependencies screen, click **Next**. (You do not need to set any dependencies for a disk group resource.)
- 5 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**.

## 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 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.

**Verifying the cluster configuration**

# Deploying SFW and VVR with MSCS

This chapter provides the steps for setting up a disaster recovery (DR) solution, using SFW with an MSCS cluster and VVR in a new installation. The example describes a generic database application in order to present general recommendations that apply to applications in a DR solution.

The process for setting up and working with the SFW-MSCS-VVR disaster recovery solution has four main parts:

- [Part 1: Setting up the cluster on the primary site](#)
- [Part 2: Setting up the cluster on the secondary site](#)
- [Part 3: Adding the VVR components for replication](#)
- [Part 4: Maintaining normal operations and recovery procedures](#)

The steps for setting up the MSCS cluster that were described in the High Availability section of this guide are the basic foundation on which this disaster recovery solution is built (see [Chapter 15, “Deploying SFW with MSCS” on page 329](#)). The main differences in the process of setting up the cluster for a disaster recovery rather than for HA alone are that you need to make sure that the VVR option is selected during the SFW installation and to configure the Veritas Volume Replicator Security Service (VxSAS) after the installation completes. In setting up the secondary site, the cluster process is similar.

Once the two clusters are set up, one at the primary site and the other at the secondary site, VVR is used to enable replication from the primary site to the secondary site.

The table that starts on the next page outlines the process for this configuration in more detail.

The high-level objectives and the tasks to complete each objective for the configuration are as follows:

**Table 17-1**

Objective	Tasks
“Reviewing the prerequisites and the configuration” on page 419	<ul style="list-style-type: none"> <li>■ Verify hardware and software prerequisites.</li> <li>■ Review configuration requirements.</li> </ul>
<b>Part 1: Setting up the cluster on the primary site</b>	
“Installing and configuring the hardware” on page 423	<ul style="list-style-type: none"> <li>■ Set up and configure the hardware according to the manufacturers’ instructions.</li> </ul>
“Installing Windows and configuring network settings” on page 423	<ul style="list-style-type: none"> <li>■ Install the operating system on both nodes.</li> <li>■ Make necessary networking settings on both nodes.</li> </ul>
“Establishing the cluster under MSCS (Primary site)” on page 423	<ul style="list-style-type: none"> <li>■ Refer to Microsoft documentation for instructions on establishing a cluster under MSCS.</li> </ul>
“Installing SFW (Primary site)” on page 424	<ul style="list-style-type: none"> <li>■ In the Options screen of the installer, select the VVR option.</li> </ul>
“Installing Veritas Volume Replicator Security Services (VxSAS)” on page 424	<ul style="list-style-type: none"> <li>■ Complete the steps to configure VxSAS.</li> </ul>
“Creating SFW disk groups and volumes” on page 427	<ul style="list-style-type: none"> <li>■ In SFW on the primary cluster node, 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.</li> <li>■ The disk group for the quorum can be created later, if desired.</li> </ul>
“Setting up a group for the application in MSCS” on page 428	<ul style="list-style-type: none"> <li>■ Create a group within MSCS for the application.</li> <li>■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.</li> </ul>

Table 17-1

Objective	Tasks
“Installing the application (Primary site)” on page 430	<ul style="list-style-type: none"> <li>■ Install the application program files on the local drive of the first node.</li> <li>■ Install files relating to the data and logs on the shared storage.</li> <li>■ Use <b>Move Group</b> to move the cluster resources to the second node.</li> <li>■ Make sure that the volumes on the second node have the same drive letters or mount points that they had on the first node.</li> <li>■ Install the application on the second node.</li> </ul>
“Completing the setup of the application group in MSCS” on page 430	<ul style="list-style-type: none"> <li>■ Refer to the application documentation for help on creating its resource.</li> <li>■ Establish the appropriate dependencies.</li> <li>■ Test the application group by using the <b>Move Group</b> command to move the cluster resources to the other node.</li> </ul>
“Changing the quorum resource to the dynamic mirrored quorum resource” on page 434	<ul style="list-style-type: none"> <li>■ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier.</li> <li>■ Make the disk group a Volume Manager Disk Group type resource in the default Cluster Group.</li> <li>■ Change the quorum resource to the dynamic mirrored quorum resource.</li> </ul>
“Testing of the cluster on the primary site” on page 434	<ul style="list-style-type: none"> <li>■ Use the <b>Move Group</b> command to move the cluster resources to the second node. Move them back to the first node.</li> <li>■ Optionally, simulate a failure by turning off the power to the server that has control of the cluster resources.</li> </ul>

---

**Part 2: Setting up the cluster on the secondary site**

---

Table 17-1

Objective	Tasks
“Repeating cluster configuration steps for the secondary site” on page 436	<p>The tasks are:</p> <ul style="list-style-type: none"> <li>■ Installing and configuring hardware</li> <li>■ Installing Windows and configuring network settings</li> <li>■ Establishing the cluster under MSCS</li> <li>■ Installing SFW</li> <li>■ Installing Veritas Volume Replicator Security Services (VxSAS)</li> <li>■ Creating SFW disk groups and volumes</li> <li>■ Setting up a group for the application in MSCS</li> <li>■ Installing the application on cluster nodes</li> <li>■ Completing the setup of the MSCS application group</li> <li>■ Changing the quorum resource to the dynamic quorum resource</li> <li>■ Final testing of the cluster</li> </ul>
<b>Part 3: Adding the VVR components for replication</b>	
“Configuring the Replicator Log volumes for VVR” on page 440	<ul style="list-style-type: none"> <li>■ Use SFW to create Replicator Log volumes for the primary and secondary sites.</li> </ul>
“Setting up the Replicated Data Sets (RDS) for VVR” on page 442	<ul style="list-style-type: none"> <li>■ Create Replicated Data Sets with VVR’s Replicated Data Set wizard and start replication for the primary and secondary sites.</li> </ul>
“Creating an RVG resource and setting the dependencies” on page 450	<ul style="list-style-type: none"> <li>■ In MSCS Cluster Administrator, create an RVG resource for replication.</li> <li>■ Set the application resource dependency on the RVG resource.</li> <li>■ Remove the direct dependency of the application resource on the Volume Manager Disk Group resource, VMDg.</li> </ul>
<b>Part 4: Maintaining normal operations and recovery procedures</b>	
“Normal operations: Monitoring the status of the replication” on page 453	<ul style="list-style-type: none"> <li>■ Monitor replication.</li> <li>■ Perform planned migration.</li> </ul>
“Disaster recovery procedures” on page 454	<ul style="list-style-type: none"> <li>■ Complete the recovery procedures after the primary site goes down.</li> </ul>

# Part 1: Setting up the cluster on the primary site

This section provides more information on the steps for creating the cluster on the primary site.

## Reviewing the prerequisites and the configuration

This topic describes the hardware and software requirements and gives an overview of the configuration.

---

**Note:** Before configuring the cluster, refer to the Microsoft documentation for MSCS requirements and the application documentation for application-specific requirements.

---

### Supported software

- Veritas Storage Foundation 5.0 for Windows (SFW) with the Cluster Option for Microsoft Cluster Service (MSCS) and the Veritas Volume Replicator Option.
- Windows 2000 Advanced Server or Windows 2000 Datacenter Server (both require SP4 with Update Rollup 1)  
*or*  
Windows Server 2003 Enterprise Edition or Datacenter Edition (SP1 supported but not required for all editions)  
*or*  
Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)  
*or*  
Windows Server 2003 for Intel Xeon (EM64T) or AMD Opteron: Enterprise x64 Edition or Datacenter x64 Edition

### System requirements

- One CD-ROM drive accessible to each system on which you are installing MSCS.
- The configuration described requires shared disks to support applications that migrate between nodes in each cluster.
- 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 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.
- Each system requires 1 GB of RAM.
- Systems to be clustered must be configured as part of a Windows 2000 or 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 Domain 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) are required. You also need a static IP address for the cluster itself. Thus, you need seven IP addresses for each cluster.
- In addition, you need two more IP addresses for replication, one for the primary site and one for the secondary site.

---

**Note:** Refer to the Hardware Compatibility List on the Symantec Support web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW.

---

## Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Installation on a non-system drive requires space on both the system drive and the non-system drive.

The following table summarizes disk space requirements for SFW.

**Table 17-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
SFW + all options + client components	1240 MB	Non-system space: 1240 MB System space: 265 MB

**Table 17-2** Disk space requirements

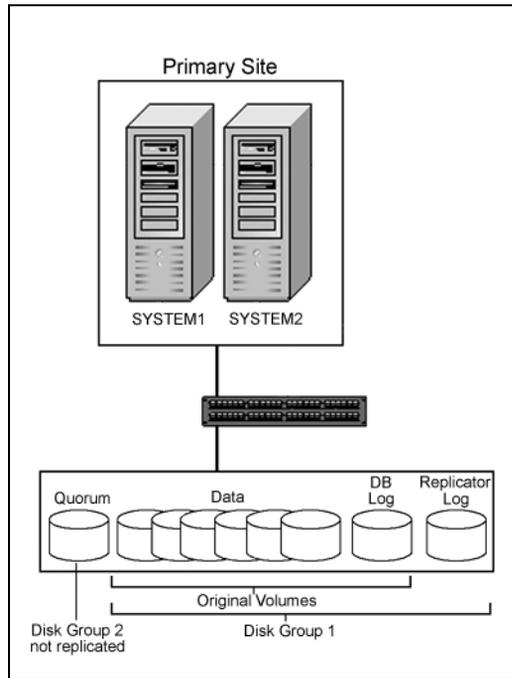
Installation options	Installation on system drive	Installation on non-system drive
SFW + all options	980 MB	Non-system Space: 980MB System space: 225 MB
Client components	420 MB	Non-system space: 420 MB System space: 80 MB

### Reviewing the configuration

This configuration overview highlights the active/passive high availability within a cluster and disaster recovery between two sites. In an active/passive configuration, one or more application virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

Active/passive clusters involve one-to-one failover capabilities. For instance, if you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM5 and SYSTEM6 on the secondary site), SYSTEM1 can fail over to SYSTEM2, and SYSTEM5 can fail over to SYSTEM6. The figure that follows illustrates the cluster configuration on the primary site. For a view of the configuration that includes both sites, see the illustration in the section [“About a disaster recovery solution”](#) on page 238.

Figure 17-1 DR configuration primary site



In addition, other items may be needed for this configuration:

- An MSCS cluster must be running to install SFW.  
Thus, you need to set up the hardware and install the operating system and MSCS on both systems in each cluster and establish the MSCS cluster before installing SFW.  
Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Thus, Symantec recommends that you use a “rolling install” procedure to install SFW first on the inactive cluster node, then move the active cluster resources to the other node, and install on the now inactive node.
- SFW adds the advantage of the dynamic mirrored quorum.  
The main advantage of using SFW instead of MSCS alone is that MSCS by itself does not support mirroring the quorum resource; thus, with MSCS alone, the quorum disk itself is a point of failure for the cluster. SFW provides a dynamic mirrored quorum resource for MSCS.  
After SFW is installed on the cluster nodes, the next task is to 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, but you may want to wait until after you have your application installed, running, and tested with the cluster to convert the original basic quorum disk to the dynamic quorum volume. By waiting until the end of the process to convert from the basic physical disk quorum to the dynamic mirrored volume, you can make sure that the application is working first with the cluster and then add the dynamic quorum volume. The quorum disk group on each site does not get replicated because each cluster has its own quorum.

## Installing and configuring the hardware

Refer to the hardware documentation and Microsoft documentation for specific details of your hardware setup.

As a best practice, Microsoft recommends that you wait until after the cluster is established on the first node before connecting the second node to the storage array in order to avoid corruption of data on the disks.

## Installing Windows and configuring network settings

This topic summarizes the steps for installing the operating system and configuring the network settings. For specific details, refer to the Microsoft documentation.

### To install Windows and configure network settings

- 1 Install the operating system and MSCS on both servers.  
MSCS is automatically installed with Windows Server 2003. With Windows 2000, you must do the additional step of installing MSCS.
- 2 Establish the network settings for the NICs and the domain on both servers.  
You need to establish static IP addresses for all six NICs—two private NICs and one public NIC for each system.

## Establishing the cluster under MSCS (Primary site)

Before you install SFW, you must establish an MSCS cluster. This section summarizes the tasks involved. For full details, refer to the Microsoft documentation.

### Tasks summary

- Configure the shared storage and create a partition for the cluster quorum disk.

- You must have a basic disk reserved for this purpose on your shared storage. Microsoft recommends 500 MB for the quorum disk; refer to the Microsoft documentation for specific instructions.
- Create the first node of the cluster on Server A, using MSCS Cluster Administrator, and make sure that it can access the shared storage.
- Connect the shared storage to the second node.
- Add the second node of the cluster on Server B with Cluster Administrator.
- Test the cluster by using the **Move Group** command to move the cluster resources to the second node.  
Server B becomes the active cluster node.

---

**Note:** To prepare for a rolling installation of SFW on Node A, Symantec recommends that you leave the cluster resources on Node B at this point.

---

## Installing SFW (Primary site)

The procedure for adding SFW support to the cluster on the primary site involves the same installation steps that were described earlier in the chapter on setting up a cluster with SFW and MSCS with one important difference: that you select the VVR option from the product installer Options screen.

Refer to [“Installing SFW”](#) on page 337 for detailed steps.

## Installing Veritas Volume Replicator Security Services (VxSAS)

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server 2003 environment, 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 <b>Password</b> 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 username and password when configuring the VxSAS service on the other hosts.

After providing the required information click **Next**.

- 3 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 <b>Add Domain</b> option. This displays a dialog that allows you to specify the domain name. Click <b>Add</b> to add the name to the Selected Domains list.

After specifying the domain click **Next**.

- 4 Select the required hosts from the Host Selection page.

Selecting Hosts	<p>The Available Hosts pane lists the hosts that are present in the specified domain.</p> <p>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.</p>
Adding a Host	<p>If the host name you require is not displayed, then add it using the <b>Add Host</b> option. In the Add Host dialog specify the required host name or IP in the <b>Host Name</b> field. Click <b>Add</b> to add the name to the Selected Hosts list.</p>

After you have selected a host name the **Configure** button is enabled. Click the **Configure** button to proceed with configuring the VxSAS service.

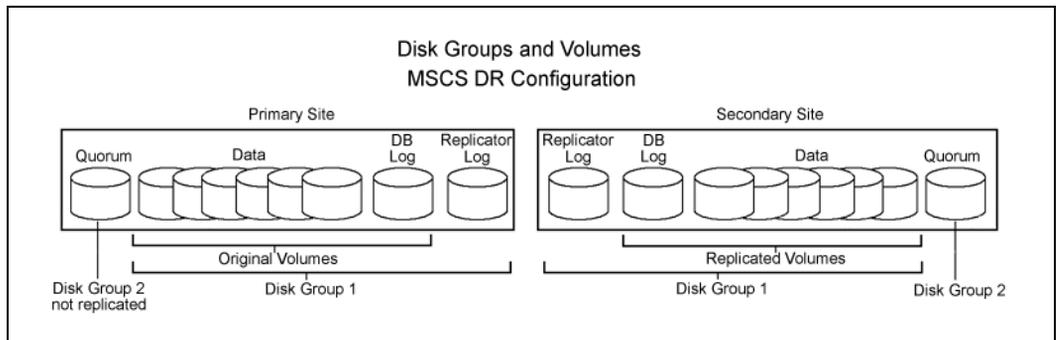
- 5 After the configuration completes, the Configuration Results page is displayed. If the operation is successful then the Status column displays the appropriate message to indicate that the operation was successful.  
If the operation was not successful then the page displays the status as failed and 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.
- 6 Click **Finish** to exit the wizard.

## Creating SFW disk groups and volumes

The following figure shows a typical setup of volumes for an MSCS VVR configuration with a database application. The example has one disk group for the application on each site.

If there are more application disk groups in your configuration, note that each disk group requires an additional Replicator Log volume. In the procedures described in this chapter, the Replicator Log volume will be created later; but you will need to allow sufficient disk space for the number of Replicator Log volumes required by your configuration. The quorum volume is not replicated to the second site and is in a separate disk group. It has to be created on each site and functions only on that site. The minimum number of disks for the mirrored quorum is two disks. Symantec recommends using three disks for the mirrored quorum for additional redundancy.

**Figure 17-2** MSCS clustered database with disks for data, logs, and the quorum resource



Do not use the following types of volumes for the data and Replicator Log volumes; VVR does not support these types of volumes:

- SFW (software) RAID 5 volumes
- Volumes with commas in the names

For detailed steps in creating disk groups and volumes, see [“Creating SFW disk groups and volumes”](#) on page 346 in Chapter 8.

## Setting up a group for the application in MSCS

Use Cluster Administrator in MSCS to set up a group for the application that will contain the SFW disk group or groups that were created for the application. The SFW disk groups are added to the MSCS application group as Volume Manager Disk Group resources.

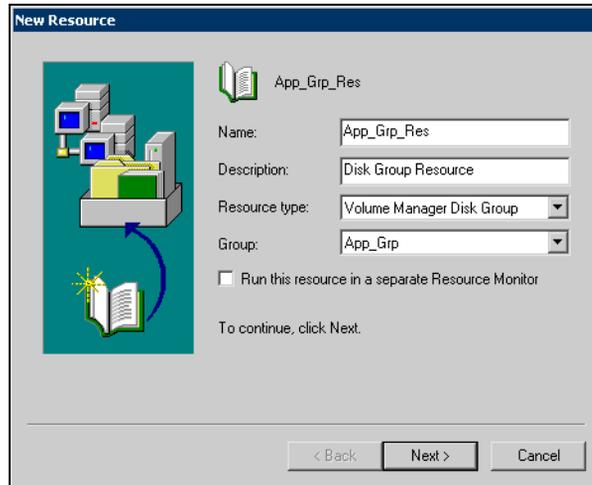
After the application is installed on both nodes and its accompanying files are placed on the shared storage, complete the setup of the application group by adding the application itself as a resource and any other resources that are required. Dependencies need to be set between the resources in the group. Information on this task is included in [“Completing the setup of the application group in MSCS”](#) on page 430.

### To set up the application group

- 1 Start Cluster Administrator in MSCS by selecting **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator**.  
Make sure you are connected to the appropriate cluster.
- 2 Create a new group by selecting the **Groups** node from the tree that is displayed in the left pane. Right-click to display the **Groups** menu. Select **New > Group** from the menu.  
The New Group window appears.
- 3 Specify a name for the group in the **Name** field (for example, App\_Grp). Click **Next** to continue.
- 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 add resources to the group.

**To add SFW disk groups as resources to the application group**

- 1 Right-click on the MSCS group that you have created for the application and select **New > Resource**. The New Resource window appears.



- Specify a name for the disk group resource in the **Name** field.
- If required, you can add a description about the resource in the **Description** field. “Disk Group Resource” is an appropriate description.
- Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource type** field drop-down list.

Note that the resource name has not been changed to Storage Foundation Disk Group.

- If necessary, use the drop-down list to select the appropriate MSCS group; the group should already be selected.
- Generally, make sure that **Run this resource in a separate Resource Monitor** is not checked. Click **Next**.

The Possible Owners screen appears.

- 2 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.

The Dependencies screen appears.

- 3 On the Dependencies screen, click **Next**. You do not need to set any dependencies for a disk group resource.
- 4 In the next screen, make sure the appropriate SFW cluster disk group is selected from the drop-down list for the resource, and click **Finish** to complete the operation.

If there is more than one disk group for the application, you need to repeat the process. You can also add more resources at this time, as required by the application, or wait until after the application is installed. You will not be able to add the application resource until after the application is installed on both nodes.

## Installing the application (Primary site)

This section gives more information on installing the application software on the two nodes of the primary site. Refer to the application documentation for any instructions on installing the application in a cluster.

### Requirements

- The application program files need to be installed on the same local drive on all nodes. For example, if you install the application program files at **C:\Program Files\<application>** on one node, then these files must be installed at **C:\Program Files\<application>** on all the other nodes. Make sure that the same drive letter is available on all nodes and that there is adequate space for the installation.
- The data files and any associated files, such as log files, should be installed on the volumes under the clustered disk group or groups on the shared storage.
- To install the application on the second node, use the **Move Group** command to move the cluster resources to the second node.

Refer to “[Installing the application on cluster nodes](#)” on page 358 for more information on this task.

## Completing the setup of the application group in MSCS

At this point in the process, additional steps make the application group functional in MSCS. The application resource is added, as well as any other resources that are associated with the application. Also, dependencies are established for the resources.

This section presents a high-level summary of the process for completing the application group setup.

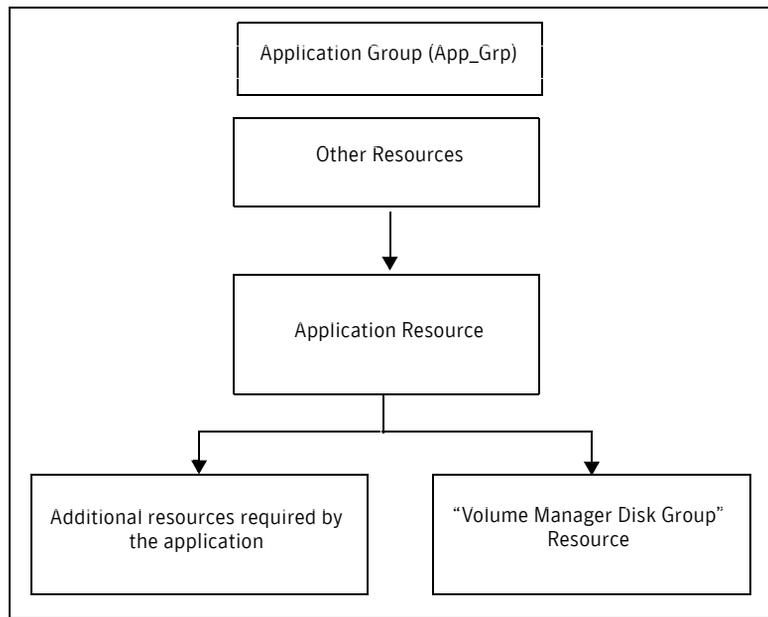
- Before creating the application resource, make sure that all the other resources that you created—that is, the disk group resource and any additional application resources—are online.
- Refer to the application documentation for help on creating its resource and additional resources that may be required. You may need to create an IP

Address resource and a Network Name resource in addition to the Volume Manager Disk Group resource you created earlier.

When creating the application resource, on the Dependencies screen select the **Volume Manager Disk Group** resource from “Available Resources” and add it to “Resource Dependencies.”

- The following dependency chart indicates the Dependencies that are established.

Application group dependencies



- **Testing:** After the application group is set up, test it by using the **Move Group** command to move the cluster resources to another node and then move them back.

## Changing the quorum resource to a dynamic quorum resource

One of the key advantages of using SFW with MSCS to create a mirrored quorum resource to add fault tolerance to the quorum, thus protecting the cluster from failure, if the disk with the quorum fails. In the following procedure, transfer the cluster’s quorum resource from a physical disk resource to a mirrored dynamic quorum resource. The tasks involved are:

- [“Creating a dynamic cluster disk group for the quorum with a mirrored volume”](#) on page 432
- [“Making the quorum cluster disk group an MSCS resource”](#) on page 432
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 434

## Creating a dynamic cluster disk group for the quorum with a mirrored volume

If you have not completed this step earlier, use SFW to create a dynamic disk group for the quorum disks. The minimum number of disks for the mirrored quorum is two disks. Symantec recommends using three small disks for the mirrored quorum for additional redundancy.

If possible, use small disks, because the disk group will be used only for the quorum volume, which Microsoft recommends to be 500 MB. To create a three-way mirrored volume in the New Volume wizard, select the **Concatenated** layout, click the **Mirrored** checkbox, and specify three mirrors. For full details on creating cluster disk groups and volumes, see [“Creating SFW disk groups and volumes”](#) on page 346.

---

**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.

---

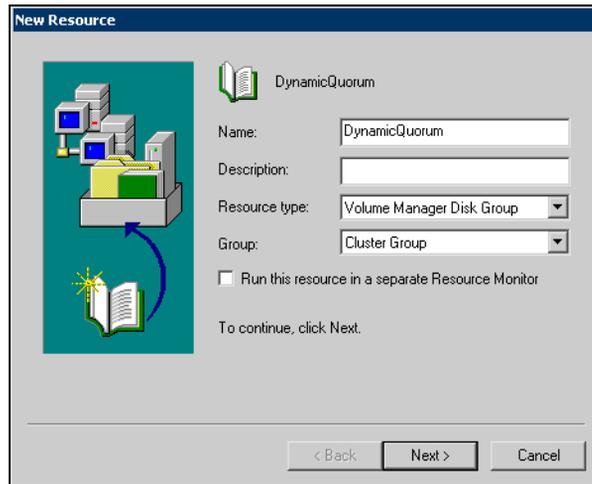
## Making the quorum cluster disk group an MSCS resource

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.



- Specify a name for the disk group resource in the **Name** field, such as “DynamicQuorum.”
- If necessary, 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.  
 Note that 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**.

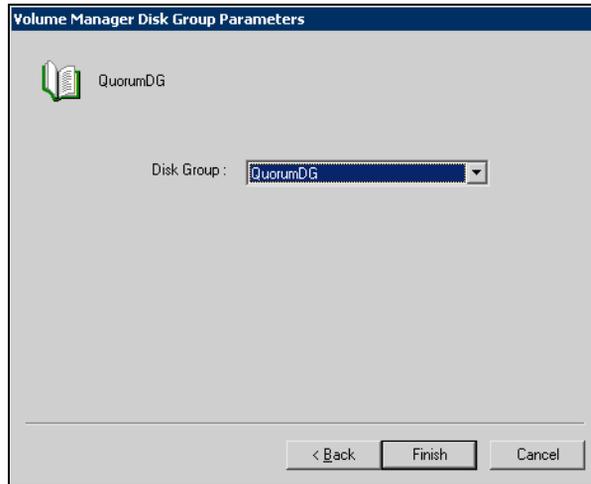
The Possible Owners screen appears.

- 3 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.

The Dependencies screen appears.

- 4 On the Dependencies screen, click **Next**. You do not need to set any dependencies for a disk group resource.

- 5 In the next screen, 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 the prepared 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**.

### Testing of the cluster on the primary site

After the application is installed and the dynamic mirrored quorum is installed, test the cluster to make sure that it functions properly before adding the VVR components.

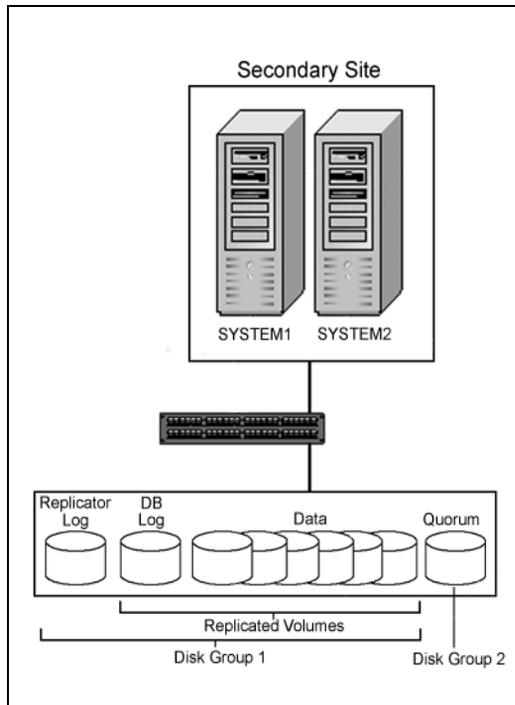
Refer to the section on testing an MSCS cluster, “[Verifying the cluster configuration](#)” on page 364.

## Part 2: Setting up the cluster on the secondary site

On the secondary site, repeat the tasks performed on the primary site to create a cluster that duplicates the primary site's disk groups and volumes.

The secondary disk groups and volumes should have the same names as those on the primary site. The data volumes should be the same sizes as the corresponding data volumes on the primary site. The log volume on the secondary site can be a different size, but Symantec recommends that the sizes be the same. Install the application on the secondary cluster nodes the same as on the primary cluster.

Figure 17-3 DR configuration secondary site



### Repeating cluster configuration steps for the secondary site

Refer to the guidelines provided earlier in this chapter to complete the same tasks on the secondary site prior to the application installation.

- [“Installing and configuring the hardware”](#) on page 423

- [“Installing Windows and configuring network settings”](#) on page 423
- [“Establishing the cluster under MSCS \(Primary site\)”](#) on page 423
- [“Installing SFW \(Primary site\)”](#) on page 424
- [“Installing Veritas Volume Replicator Security Services \(VxSAS\)”](#) on page 424
- [“Creating SFW disk groups and volumes”](#) on page 427
- [“Setting up a group for the application in MSCS”](#) on page 428

After completing these tasks, continue with the next section for instructions on installing the application on the secondary site.

## Installing the application (Secondary site)

---

**Note:** Before installing the application on the secondary site, offline all the resources in the MSCS application group on the primary site, except the disk group resource.

---

Installing the application on the secondary site is similar to installing it on the primary site. See the section [“Installing the application \(Primary site\)”](#) on page 430.

## Completing the setup of the MSCS application group

Refer to [“Completing the setup of the application group in MSCS”](#) on page 430.

## Changing the quorum resource to the dynamic quorum resource

Refer to [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 434.

## Final testing of the cluster

Refer to [“Testing of the cluster on the primary site”](#) on page 434.

## Before configuring VVR components

After both clusters are running, one on the primary site and one on the secondary site, you can add the VVR components to the configuration.

---

**Note:** Before configuring the VVR components, on the secondary site, offline all the resources in the application group, except the disk group resource.

---

## Part 3: Adding the VVR components for replication

This section provides information on configuring the VVR components for replication. Topics include:

- [“VVR components overview”](#) on page 439
- [“Configuring the Replicator Log volumes for VVR”](#) on page 440
- [“Setting up the Replicated Data Sets \(RDS\) for VVR”](#) on page 442
- [“Creating an RVG resource and setting the dependencies”](#) on page 450

### VVR components overview

The terms Replicated Volume Group (RVG), Replicator Log, and Replicated Data Set (RDS) are used frequently in this section. Here are their definitions:

#### **Replicated Volume Group (RVG)**

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, 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, but Replicator Log volume sizes can differ. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG on the primary host.

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.

#### **Replicated Data Set (RDS)**

An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).

#### **Replicator Log**

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 logs at the two sites must have the same name; however, the sizes of the logs can vary. Symantec recommends having Replicator Log volumes of the same size at the primary site and the secondary site.

The process described in this section involves setting up an RDS for each SFW disk group on the primary site that will have replicated volumes, and then creating a VVR service group that is linked to the application service group.

## Configuring the Replicator Log volumes for VVR

---

**Note:** Before configuring the Replicator Log volumes, make sure that all the resources in the MSCS application group are offline, except the disk group resource. This task must be done on the primary site as well as the secondary site.

---

Create the volume for the Replicator Log at each site. The task of creating the logs can also be done during the RDS creation process, but some storage administrators may prefer to do it manually (as is being done here) as a preparatory step to setting up the RDS.

---

**Note:** To improve write performance, Symantec recommends that you create the Replicator Log volume on a different disk from the disks used for your application data volumes.

---

### To configure the Replicator Log volumes for VVR

- 1 Click **Start > All Programs > Symantec > Veritas Enterprise Administrator** on the desktop to open the VEA console on the active node of the primary site.
- 2 Create a volume for the disk group that contains the storage group data:
  - On the System configuration tree, click the disk group where the log volume will be created (*Hostname>Disk Groups>Diskgroupname*).
  - Right-click on a disk group that has the volumes to be replicated, and click **New Volume**.
- 3 On the Welcome page of the New Volume wizard, click **Next**.
- 4 Select the disks for the volume:
  - Select the group name.
  - Select **Manually select disks**.
  - Click the disk name.
  - Click **Add**.
  - After selecting all the necessary disks, click **Next**.
- 5 Specify the parameters of the volume:
  - Enter the volume name.
  - Enter the size. The size of the Replicator Log volume varies for different environments. To determine the appropriate size for your

environment, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

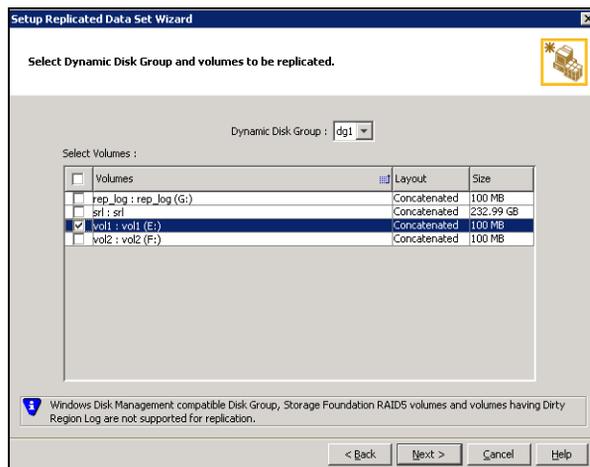
- Select the volume layout.
  - Select the appropriate mirror options.
  - Click **Next**.
- 6 On the Add Drive Letter and Paths dialog box:
    - Click **Do not assign a drive letter**.
    - Click **Next**.
  - 7 When prompted to format the volume:
    - Deselect **Format this volume**.
    - Click **Next**.
  - 8 Click **Finish** to create the new volume.
  - 9 If necessary, repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for any additional RVGs on the primary site.
  - 10 Repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for additional disk groups on the secondary site.

## Setting up the Replicated Data Sets (RDS) for VVR

Configuring VVR involves setting up the Replicated Data Sets on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

### To create the Replicated Data Set

- 1 From the cluster node on the primary site where the cluster disk group is imported, use the VEA console to launch the Setup Replicated Data Set Wizard. Right-click **Replication Network** on the Management Host configuration tree, and click **Setup Replicated Data Set**.
- 2 Read the Welcome page and click **Next**.
- 3 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

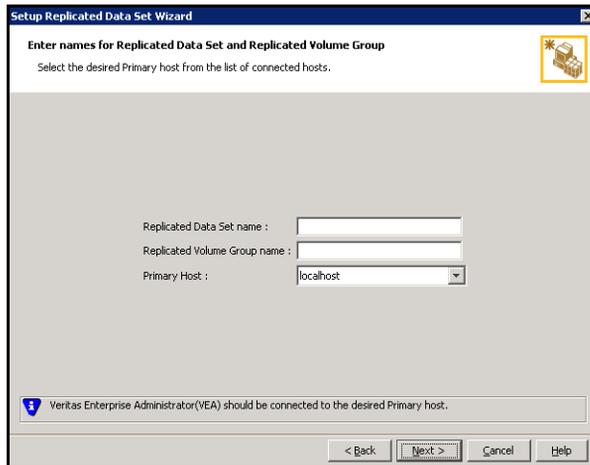


By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

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.

- 4 Click **Next**.

- 5 Select from the table the dynamic disk group and data volumes that will undergo replication.

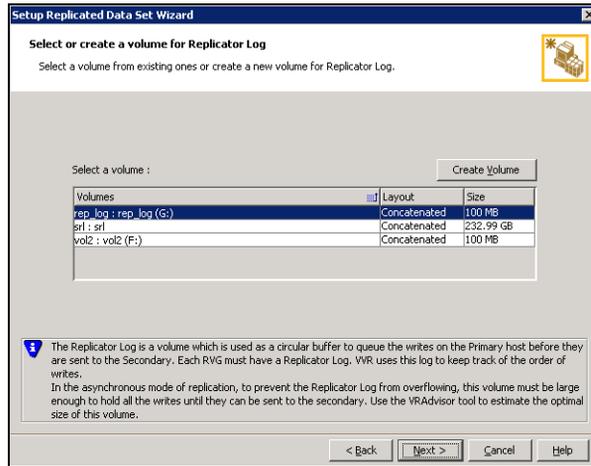


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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.

- 6 Click **Next**.

7 Select or create a volume for the Replicator Log:



**To select an existing volume**

- Select the volume for the Replicator Log in the table (APP\_REPL\_LOG). 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.

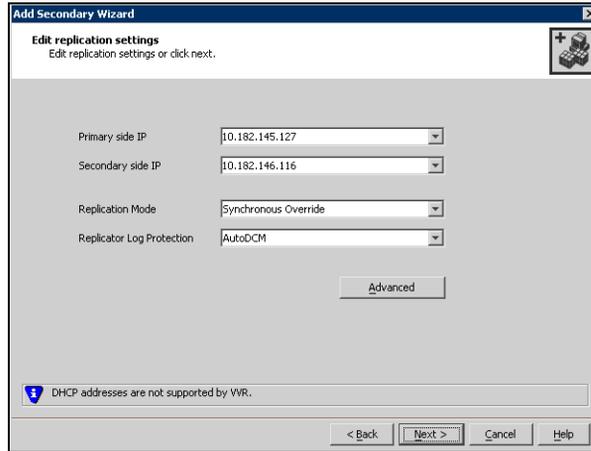
- |                       |  |
|-----------------------|--|
| <b>Name</b>           | Enter the name for the volume in the <b>Name</b> field.  |
| <b>Size</b>           | Enter a size for the volume in the <b>Size</b> field.  |
| <b>Layout</b>         | Select the desired volume layout.  |
| <b>Disk Selection</b> | <ul style="list-style-type: none"><li>■ Choose <b>Select disks automatically</b> if you want VVR to select the disks for the Replicator Log.</li><li>■ Choose <b>Select disks manually</b> to use specific disks from the available disks pane for creating the Replicator Log volume.<br/>Either double-click the disk to select it, or select <b>Add</b> to move the disks into the selected disks pane.</li></ul> |
- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 8 Review the information on the summary page and click **Create Primary RVG**.
  - 9 After the RVG for the primary site is successfully created, click **Yes** to add the secondary host to the RDS for replication.
  - 10 Specify the name of the host where the disk group is imported on the secondary site. If necessary, specify the fully qualified domain name.
  - 11 Click **Next**.
  - 12 If the Veritas Enterprise Administrator console is not already connected to the secondary host, the connection process starts when you click **Next**. Enter valid user credentials, click **OK**, and click **Next** again.
  - 13 The configuration for these volumes on the primary and secondary sites must be identical and meet VVR configuration requirements. If a Replicator Log volume does not exist on the secondary site, it can be created with this procedure.
    - 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 If necessary, edit the replication settings for a secondary host.



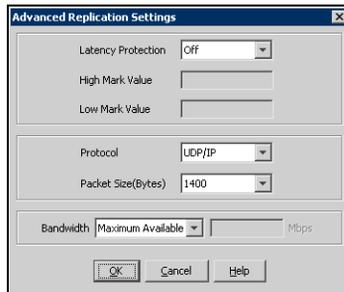
- Enter the virtual IP address for the Primary IP resource that will be used for replication.
- Select or specify an IP address for the Secondary IP resource.
- Specify the replication mode.

<b>Synchronous Override</b>	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 <b>Asynchronous</b> .
<b>Synchronous</b>	Determines updates from the application on the primary site are completed only after the secondary site successfully receives the updates.
<b>Asynchronous</b>	Determines updates from the application on the primary site are completed after VVR stores the updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously

- Specify the replicator log overflow protection property.

<b>AutoDCM</b>	Is the default option and enables the DCM when the Replicator Log overflows even though the secondary site is connected. All data volumes in the primary RVG must have a DCM log to use this option. The wizard attempts to ensure all volumes under the primary RVG have a DCM log.
<b>DCM</b>	Enables Replicator Log protection for the secondary site. DCM is enabled when the Replicator Log overflows and the secondary site is disconnected from the primary site. All data volumes in the primary RVG must have a DCM log to use this option. The wizard attempts to ensure all volumes under the primary RVG have a DCM log.
<b>Off</b>	Disables Replicator Log overflow protection.
<b>Override</b>	<p>Enables log protection. If the secondary site is still connected and the Replicator Log is about to overflow, VVR stalls write operations until five percent or 20 megabytes (whichever is smaller) of the space on the Replicator Log becomes available.</p> <p>If the secondary site becomes inactive because of a connection failure or administrative action, VVR disables Replicator Log protection and causes the Replicator Log to overflow.</p>
<b>Fail</b>	Enables log protection. When the Replicator Log is about to overflow, VVR stalls write operations until five percent or 20 megabytes (whichever is smaller) of the space on the Replicator Log becomes available. If the connection between the primary RVG and secondary RVG is broken, subsequent write operations to the primary RVG fail.

- 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.

16 Click **OK** to close the dialog box.

17 Click **Next**.

18 On the **Start Replication** page, accept the **Synchronize Automatically** option, which is the default recommended for initial setup.

19 Select **Start Replication**, which is the default.

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.

If the virtual IPs have been created, select **Start Replication** to start synchronization immediately.

If replication must be started later, use the **Start Replication** option of VEA to begin replication. Refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for additional details.

20 Click **Next**.

21 Review the specifications and click **Finish** to add the host on the secondary site to the RDS. Click **Back** to change any information. Replication physically starts when the IP address is created.

## Creating an RVG resource and setting the dependencies

This section describes additional tasks that must be done to complete the configuration of the MSCS application service group at both the primary and secondary sites. The tasks are:

- [Creating a replicated volume group \(RVG\) resource](#)
- [Setting the application resource dependency on the RVG resource](#)

### Creating a replicated volume group (RVG) resource

To create an RVG resource

- 1 On the primary site, access Cluster Administrator, right-click on the application cluster group that you created, and select **New > Resource**. The New Resource screen appears.
  - Specify a name for the RVG resource in the **Name** field. For example, **RVGResource** is an appropriate name.
  - If desired, 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.
  - Configure a separate resource monitor process for the RVG resource by selecting the **Run this resource in a separate Resource Monitor** checkbox.
  - Click **Next**.  
The Possible Owners screen appears.
- 2 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.  
The Dependencies screen appears.
- 3 On the Dependencies screen, select the IP resource created earlier for VVR and the Volume Manager Disk Group resource from the “Available Resources” list in the left pane and add it to “Resource Dependencies” list in the right pane of the screen. Click **Next**.
- 4 In the Replicated Volume Group Parameters screen, select the disk group for the RVG resource. Click **Finish**.
- 5 Bring the RVG resource online.
- 6 Repeat these same steps on the secondary site.

## Setting the application resource dependency on the RVG resource

The application resource has a direct dependency on the Volume Manager Disk Group resource. With the addition of the RVG resource to the application group, the application's dependency will change. The application will have a direct dependency on the RVG resource, which in turn depends on the Volume Manager Disk Group resource.

---

**Note:** The Volume Manager Disk Group resource represents the cluster disk groups created and managed by SFW.

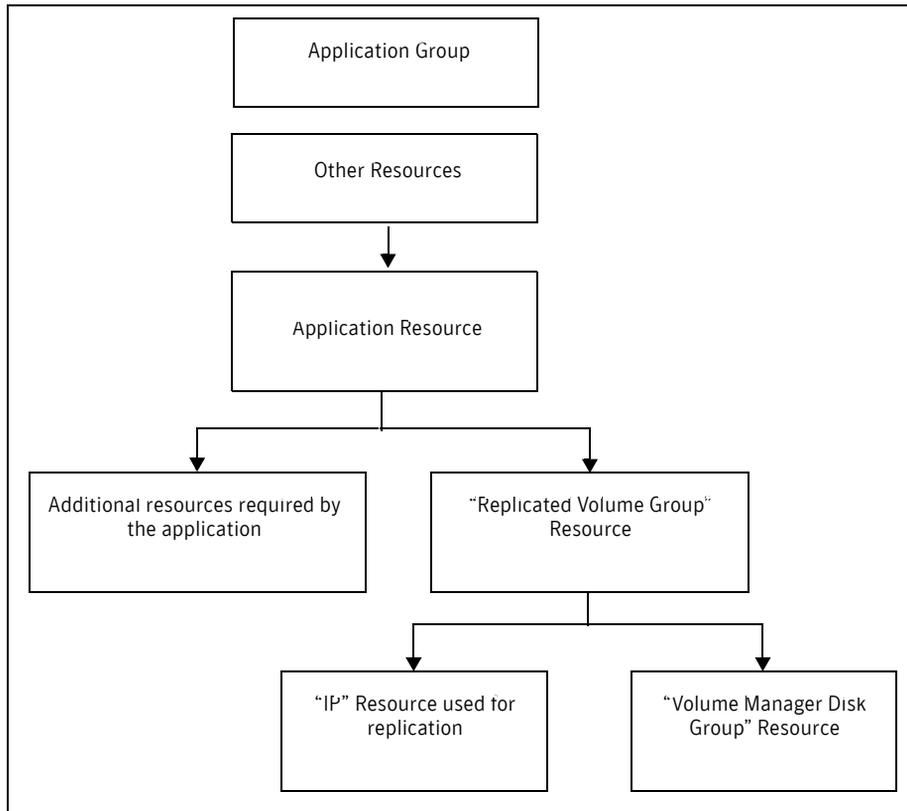
---

### To set the application resource dependency on the RVG resource

- 1 Make sure the application resource is off-line before attempting to modify the dependencies.
- 2 Right-click on the application resource and select the **Properties > Dependencies** tab. This will display the Dependencies screen.
- 3 Click **Modify**.
- 4 Select the **Replicated Volume Group** resource from the "Available Resources" list and move it to the "Resource Dependencies" list.
- 5 Remove the **Volume Manager Disk Group** resource from the "Resource Dependencies" list. Click **OK**.  
There is no longer a direct dependency between the application resource and the Volume Manager Disk Group resource.
- 6 The cluster configuration is now complete. Online the entire MSCS application group on the primary cluster.

The dependency chart that follows indicates the dependencies that have been established.

Figure 17-4 Dependencies of VVR-related resources



The chart shows only the VVR-related resources. Normally, there would be other resources involved in any clustered application. The main point of the chart is to show that the RVG resource is now dependent on the Volume Manager Disk Group resource and the VVR virtual IP resource. The dependencies relationship has changed. The application resource is no longer directly dependent on the Volume Manager Disk Group resource.

## Part 4: Maintaining normal operations and recovery procedures

This section provides tasks during normal operations of this solution 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:

- The VEA GUI
- The command line interface (CLI)
- Perfmon
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Storage Foundation 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. The following are a generic set of tasks that you may need to perform:

- Take the RVG resource offline on both the clusters.
- Transfer the primary role to the host at the secondary site by using the **Migrate** option.
  - From the VEA screen, right-click the primary RVG and select **Migrate**.
  - Select the secondary host and click **OK**. The replication role is migrated to the secondary host.
- Assign drive letters to the volumes on the new primary.  
Make sure that these drive letters are the same as those of the original primary.
- Bring the RVG resource online on the new secondary.
- Bring the application group online on the new primary.

You can now verify that the application functions properly on the new primary with the replicated data. After verifying its functioning, you can revert the roles to what they were originally by repeating the procedure.

---

**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.

---

## Disaster recovery procedures

This section provides information on bringing up an application 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 returned to normal functioning after a disaster.

### Bringing up the application on the secondary host

To bring up the application on the secondary host

- 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 to perform the takeover operation. You can choose to perform takeover with the following options:
  - Perform **Takeover** with the **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 **Takeover** without the **fast-failback** option. In this case, you will 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 the takeover, the existing secondary becomes the new primary.

- 3 Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- 4 Bring the application group online.

Now you can start using the application 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 restore the primary host

- 1 Take the RVG resource off-line on both the clusters.
- 2 Depending on whether you performed **Takeover** with or without the **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 new primary's context menu.
  - For Takeover without the Fast-failback option:

After performing a takeover without fast-failback, you must convert the original primary to a secondary by 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 the **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** from the menu that appears.
- 4 Make sure that the volumes have retained the same drive letters as they had before the disaster.
- 5 Bring the RVG resource online on the secondary.
- 6 Bring the application group online on the original primary.



# Server Consolidation

This section highlights server consolidation, the practice of consolidating server hardware, software, and data from multiple smaller servers to fewer, larger servers. This section also includes two sample configurations.

This section has the following chapters:

- [Chapter 18, “Server consolidation overview” on page 459](#)
- [Chapter 19, “Server consolidation configurations” on page 463](#)



# Server consolidation overview

This overview chapter describes server consolidation and focuses on how Veritas Storage Foundation for Windows (SFW) supports a server consolidation solution. The chapter's topics are:

- [“Server consolidation definition”](#) on page 459
- [“Need for implementing server consolidation”](#) on page 459
- [“Advantages of using SFW with server consolidation”](#) on page 460
- [“Overview of the server consolidation process”](#) on page 462

## Server consolidation definition

Server consolidation is the consolidation of server hardware, applications, and data from multiple smaller, less powerful machines to fewer, more powerful servers. It involves sharing data in storage pools, usually in a storage area network (SAN).

## Need for implementing server consolidation

Server consolidation provides the benefit of overall cost reduction by reducing the number of servers and their maintenance and administrative costs. Server consolidation also frees up space in the data center and improves security by reducing virus or software gateway risks, while improving service and availability. The larger, more powerful servers are better able to provide the computing power necessary to keep businesses competitive for the future.

## Advantages of using SFW with server consolidation

Storage Foundation for Windows is ideally suited to support a server consolidation environment. Once servers are consolidated, SFW provides key features that assure fault tolerance and improve storage utilization. SFW's fault-tolerant features, such as software mirroring and RAID-5, Dynamic Multipathing (DMP), and clustering support assure high availability for consolidated storage, when business continuity is a requirement in a competitive business environment.

The SFW features that support server consolidation are:

- Ability to work in a heterogeneous storage environment  
You are not tied to a solution offered by a single hardware vendor.
- Simple migration of data with disk group import and deport commands  
If you have SFW disk groups already set up on multiple servers, you deport them on the source server, disconnect the attached storage, reattach the storage on the new larger server, and use the disk group import command to import the disk groups on the new server.
- Storage virtualization with software RAID volumes  
Once the applications and data are consolidated on the new server, mirrored and RAID-5 volumes provide fault tolerance for critical data. Striped volumes add performance capabilities. Volumes that are both striped and mirrored offer both better performance and fault tolerance. Logical RAID volumes overcome the limitations of physical disks because these RAID volumes can span across disks and even disk arrays, thus assuring more efficient use of storage. Volumes can be configured online without restarting the server.
- Capacity management and online volume growth  
Managing the space allocated for different functions is an important task that a system administrator must do on a consolidated server. SFW has a capacity monitoring function that alerts administrators when used space on a volume is near its capacity so that the volume can grow while it remains online. With this feature, you do not have to preallocate set amounts of storage for different purposes. More storage can be held in reserve in a pool for use only when it is needed. SFW volumes can be configured to increase capacity automatically when they pass a certain threshold.
- Online storage migration  
If you need to take down a disk or even a whole disk array for maintenance, you can migrate the data online through the **Move Subdisk** command.
- Special features that support storage in a SAN

The importing and deporting of disk groups with host ID protection and private disk group protection can support storage in a SAN.

- **Dynamic Multi-pathing (DMP)**  
The DMP software option increases performance of SAN-based disk arrays by spreading I/O between multiple paths to an array. Each path has a separate host adapter and cabling connecting the array and the server. If one path goes down, the DMP software automatically switches the storage associated with the failed path to an alternate path. Thus, the DMP software provides both fault tolerance for path failure and increases in performance through load balancing.
- **Clustering**  
Storage Foundation for Windows supports clustering with MSCS and Storage Foundation HA for Windows includes Veritas Cluster Server. Clustering adds fault tolerance for servers. If one server in a clustered group of servers goes down, the storage of that server is taken over by another server in the cluster.
- **Additional fault tolerance features**  
RAID-5 logging, dirty region logging, Hot Relocation, and FastResync (FR) increase the efficiency of the mirroring and RAID-5 functions in SFW.
- **Performance monitoring**  
Online performance monitoring and tuning tools provide easy identification and minimization of I/O bottlenecks. These features allow you to increase throughput of the I/O in your system.

## Overview of the server consolidation process

The server consolidation process involves more than just implementing the consolidation itself. It requires advance planning and approval of upper management. Here are some high-level steps:

- Preliminary analysis: Determine what servers need to be consolidated. Take into account the applications being used and the departments involved. Research the hardware and software needs and costs.
- Design a plan for the consolidation and secure approval and budget from upper management. The primary justifications in the plan are cost savings and the need to remain competitive in today's business environment. The plan should also address IT management of the servers after the consolidation takes place.
- Communicate with users about the proposed plan and identify the advantages of the plan before implementing the consolidation. Involve users in the planning process.
- Do a proof of concept for the consolidation. Prototype the consolidation with a smaller number of servers that are not in production to see if your plan works. In the next section, two sample configurations are provided for demonstrating a proof of concept for consolidation.
- Implement the consolidation on actual production servers.
  - Purchase, install, and configure the new hardware and software for the migration.
  - Migrate the data.
  - Test to see that everything is working properly.
  - Put into effect new IT management processes for the consolidated servers.
- In the months following the consolidation, implement a procedure to evaluate its effectiveness and the effectiveness of the IT management processes for the consolidated servers.

# Server consolidation configurations

The chapter's topics are:

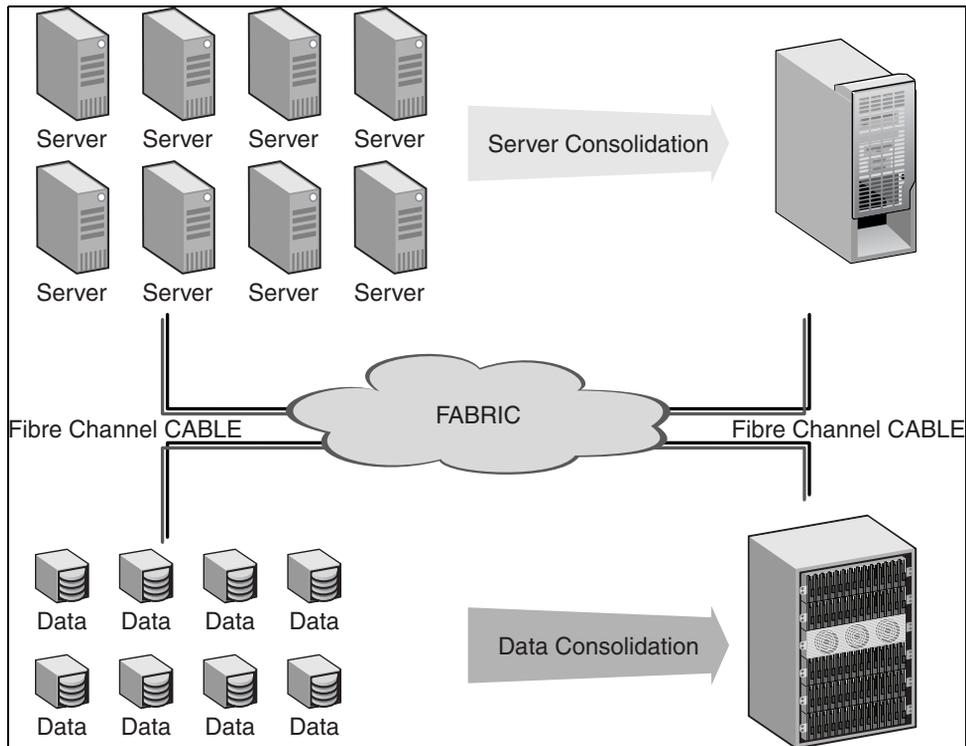
- [“Typical server consolidation configuration”](#) on page 464
- [“Server consolidation configuration 1 – many to one”](#) on page 465
- [“Server consolidation configuration 2 – many to two: Adding clustering and DMP”](#) on page 473
- [“SFW features that support server consolidation”](#) on page 481
- [“Server consolidation customer success story”](#) on page 482

## Typical server consolidation configuration

This chapter provides two sample configurations that can be used as proof of concept for a consolidation.

The example shows a typical server consolidation situation. The consolidation could involve consolidating as many as 20 to 40 servers to one or two servers.

Figure 19-1 General server consolidation configuration



### Proof of concept

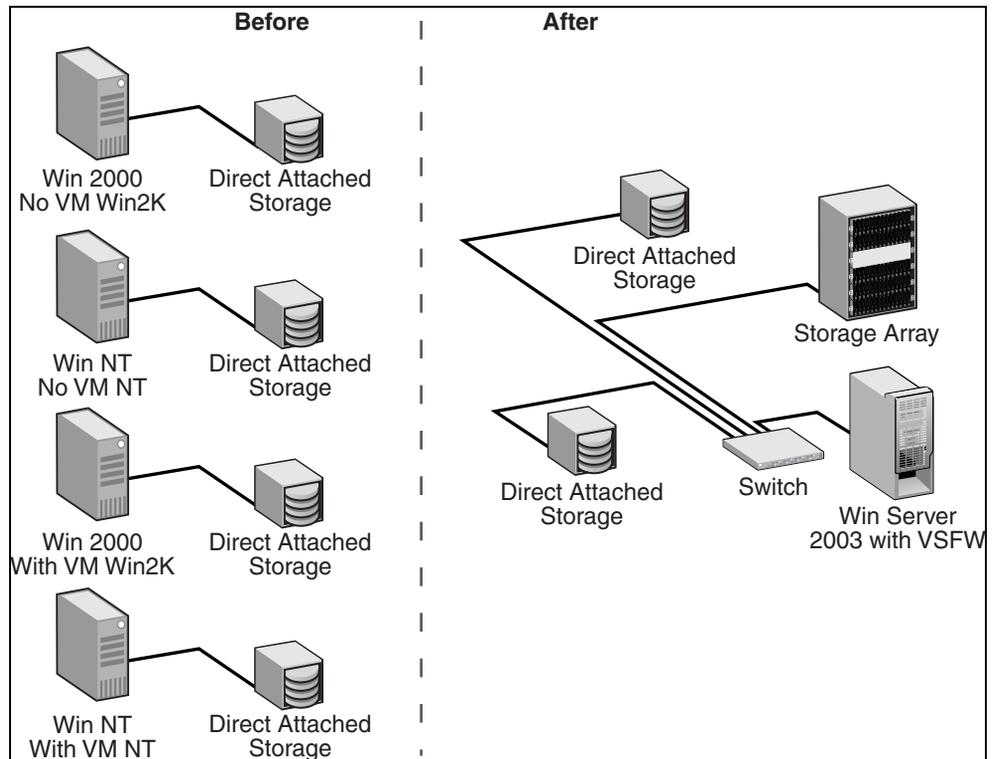
Testing the consolidation steps on a smaller number of servers provides an overview of the issues involved and how the process would work. In the configurations presented in this chapter, four servers are consolidated into one or two servers. The first configuration, which consolidates four smaller servers to one large server, provides fault tolerance through mirroring. In the second configuration, clustering and DMP are added to improve the fault tolerance, and an additional server is needed to support clustering. New, larger, more powerful

servers can be used in this proof of concept testing. Once the concept is tested, the main task is to migrate the data from the production servers to the new larger servers.

## Server consolidation configuration 1 — many to one

The following configuration illustrates consolidating many servers to one.

**Figure 19-2** Proof of concept: Consolidating four small servers to one large server



### About this configuration

In this configuration, four small servers are consolidated into a single larger server. The configuration also demonstrates that a server consolidation does not require that you eliminate all existing direct-attached storage units and replace them with large storage arrays. Setting up the storage on a SAN allows

you to use different combinations of storage devices and still derive the benefits from SFW's storage management features once the storage has been migrated from the small servers to a SAN.

## Proof of concept

The four servers represent different Windows operating systems and Storage Foundation for Windows software combinations, which might be present in a production environment. The steps demonstrate that slightly different procedures are needed in preparing the storage for migration in each of these combinations.

In setting up your server consolidation configuration for proof of concept, select servers to migrate that have different combinations of typical hardware and software to determine the special requirements of such cases.

## Phased approach: Flexible use of storage devices

In this example configuration, the steps are organized in phases:

- Preparing to consolidate
- Migrating the data to the large server
- Migrating data from the direct-attached storage to the storage array
- Adding the storage array
- Completing the consolidation process by migrating the storage from remaining servers

After the second phase in this example, all the direct-attached storage units have been detached from the small servers and are a storage pool on a SAN that is under the control of the new, large Windows Server 2003 system that is running SFW. You could stop at this point and still have many benefits from the storage that is now under SFW's management. If circumstances do not permit the purchase of a large storage array, you can simply use the existing direct-attached storage. Another alternative is to use both a storage array and some of the direct-attached storage. In this configuration and in Server Consolidation Configuration 2, using both a storage array and some of the direct-attached storage is shown.

The table below outlines the high-level objectives for implementing the configuration and the tasks for each objective:

**Table 19-1** Tasks for server consolidation for many to one configuration

Objectives	Tasks
<a href="#">“Reviewing the configuration requirements”</a> on page 468	<ul style="list-style-type: none"> <li>✓ Verify hardware and software requirements.</li> </ul>
<a href="#">“Preparing to consolidate”</a> on page 470	<ul style="list-style-type: none"> <li>✓ Make sure the data is backed up from the smaller servers before proceeding.</li> <li>✓ Set up the new large server and install the operating system and SFW. Connect it to the switch.</li> <li>✓ Prepare the data from each smaller server for consolidation by upgrading the server’s disks to dynamic disk groups, using either Disk Management or a version of Volume Manager for Windows.</li> <li>✓ Power down all the smaller servers and detach the storage.</li> </ul>
<a href="#">“Migrating the data to the large server”</a> on page 471	<ul style="list-style-type: none"> <li>✓ Reattach the direct-attached storage to the switch.</li> <li>✓ From the large server, import the disk groups from the direct-attached storage.            The direct-attached storage is now attached to the SAN and is under the management of the large server that is running SFW. You could stop at this point if a large storage array is not available.</li> </ul>
<a href="#">“Adding the storage array”</a> on page 472	<ul style="list-style-type: none"> <li>✓ If you want to use a large storage array, set up the hardware array and connect it to the switch.</li> <li>✓ Migrate the data to the large storage array.</li> </ul>

**Table 19-1** Tasks for server consolidation for many to one configuration

Objectives	Tasks
“ <a href="#">Completing the consolidation process</a> ” on page 472	✓ Migrate the data from the remaining servers.

## Reviewing the configuration requirements

Reviewing the prerequisites and the configuration allows you to gain an overall understanding of the configuration and its requirements.

### Prerequisites

These procedures assume:

- Experience in setting up computer hardware, switches, and storage arrays
- Familiarity with the Windows operating systems and SFW/Volume Manager for Windows commands

### Supported software

- Veritas Storage Foundation 5.0 for Windows (SFW)  
*or*  
Veritas Storage Foundation HA 5.0 for Windows (SFW HA)
- One of the following operating systems:
  - Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
  - Windows Server 2003 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition (SP 1 required for all editions)
  - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
  - Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)
  - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition
  - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

## Hardware setup

- 4 smaller servers with direct-attached storage
- 1 larger, more powerful server
- Fibre switch and appropriate cabling for the SAN
- Hardware storage array

---

**Note:** Refer to the Hardware Compatibility List on the Symantec Support web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW or SFW HA.

---

## Preparing to consolidate

In this phase, set up the large server and prepare the data for migration.

### To prepare for consolidation

- 1 Identify the applications and data on the smaller servers that are a subset of the applications and data to be moved to the large server. You may want to have the users delete unnecessary files before the consolidation takes place.
- 2 Back up the data from the small servers.

---

**Caution:** back up the data from the small servers before proceeding.

---

- 3 Set up the large server and connect it to the switch.
- 4 Install the Windows Server 2003 operating system and Storage Foundation for Windows on the large server.
- 5 Prepare the data from each smaller server for migration by upgrading the server's disks to dynamic disk groups and powering down the server.

#### **For Windows 2000 or Windows Server 2003 (no VM or SFW installed)**

- Use Disk Management to upgrade basic disks to dynamic disks.
- Power down the server.

#### **For Windows 2000 with SFW installed**

- Deport all dynamic disk groups.
- Power down the server.

## Migrating the data to the large server

Migrate the data to the large server. Perform the steps for each smaller server, one at a time.

### To migrate the data to the large server

- 1 Disconnect the direct-attached storage from the small server.
- 2 Connect the direct-attached storage to the switch to make it accessible to the large server.

---

**Note:** All the direct-attached storage devices and the large server need to be in the same zone on the switch.

---

- 3 Using SFW on the large server, rescan the disks.
- 4 In SFW, import the disk groups from the direct-attached storage to make them a part of the storage that the large server manages.  
Clear the host ID during the import process, if the source disk group was not created with SFW. A dialog box will come up for this purpose during the import command.
- 5 Assign drive letters to the imported disk groups.  
On a Windows 2000 server, drive letters are automatically assigned after a disk group is imported. On a Windows Server 2003 system, the default operating system setting requires the manual assignment of drive letters. Many administrators prefer to set drive letters manually rather than have the operating system do it.

---

**Note:** If you want the drive letters to be assigned automatically after a disk group is imported, use the `mountvol` command to change the default setting. Refer to the Microsoft documentation about the `mountvol` command for information on how to set up the automatic assignment of drive letters.

---

- 6 If desired, update the imported disk groups to the latest version of dynamic disk group type.  
This is recommended to take advantage of the Windows Server 2003 features in SFW. Use the **Upgrade Dynamic Disk Group Version** command.
- 7 Test the data on the Windows Server 2003 system.

At this point, you can stop if you do not have a large storage array available. You can still take advantage of SFW's storage management features by having the direct-attached storage on the SAN. It is not necessary to have a large storage array to have these benefits.

## Adding the storage array

If you have a large storage array available, the data may also be migrated to a hardware storage array on the SAN. You can eliminate all the direct-attached storage devices or keep them to increase your storage capacity. They can also be added into the configuration when needed.

### To add the storage array

- 1 Set up and connect the hardware storage array to the switch.  
On the switch, the hardware storage array must be in the same zone as the direct-attached storage devices and the large server.
- 2 Configure the array so that half of its disks are a mirror to the other half, using RAID-1. This provides fault tolerance to the storage.
- 3 Join the disk groups on the array storage and the direct-attached storage. This is done through the **Join Dynamic Disk Group** command.
- 4 Use the **Move Subdisk** command to move the volumes with data from the direct-attached storage to the array storage. You may want to keep some of the direct-attached storage on the SAN under the control of the large server.

To access the **Move Subdisk** command:

- Select the volume that contains the subdisk you want to move.
- Click on the **Subdisks** tab in the right pane of the window.
- Right-click the desired subdisk in the **Subdisks** tab and select **Move Subdisk** from the context menu.

The **Move Subdisk** command also can be done by dragging and dropping the subdisks between disks in the Disk View. This method should be used with care to make sure that you do not move the subdisk to the wrong disk.

- 5 Test the data on the Windows Server 2003 system.  
At this point, the migration of the storage from the four smaller servers is complete.

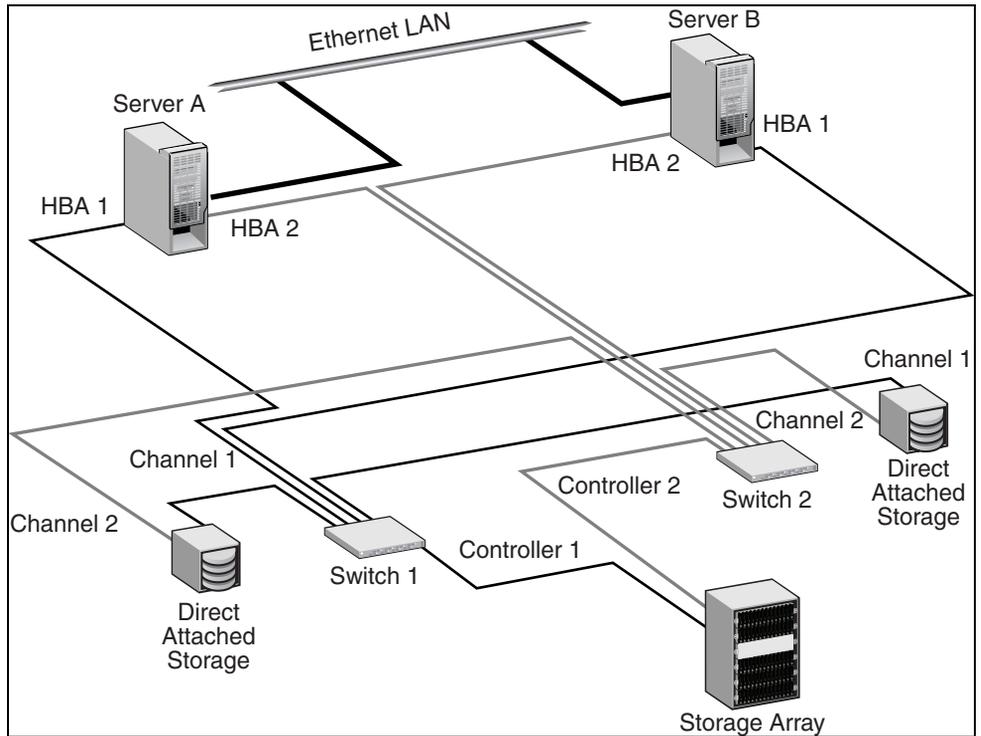
## Completing the consolidation process

When you are satisfied that everything is working properly, migrate data from the remaining servers, using the methods shown in this configuration example.

## Server consolidation configuration 2 — many to two: Adding clustering and DMP

The following configuration consolidates many servers to two with MSCS clustering and DMP.

Figure 19-3 Adding fault tolerance with MSCS and DMP — requires two servers



### About this configuration

This configuration is an upgrade to Server Consolidation Configuration 1, to add MSCS and DMP. Add a new server and host adapters, NICs, and a new switch.

The table below outlines the high-level objectives for implementing the configuration and the tasks for each objective:

**Table 19-2** Tasks for server consolidation adding MSCS and DMP

Objectives	Tasks
“ <a href="#">Reviewing the configuration requirements</a> ” on page 476	<ul style="list-style-type: none"> <li>✓ Verify hardware and software requirements.</li> </ul>
“ <a href="#">Adding the new hardware</a> ” on page 477	<ul style="list-style-type: none"> <li>✓ Add the new server, HBAs, network cards, and fibre switch.</li> <li>✓ Leave the second path for DMP unconnected on the existing server and the new server. It does not get connected until the end of the installation process.</li> </ul>
“ <a href="#">Establishing the MSCS cluster</a> ” on page 478	<ul style="list-style-type: none"> <li>✓ Refer to Microsoft instructions for establishing the cluster under MSCS.</li> </ul>
“ <a href="#">Adding SFW support to the cluster</a> ” on page 478	<ul style="list-style-type: none"> <li>✓ With Server B as the active cluster node, use <b>Add or Remove Programs</b> to add DMP and the MSCS support option to the first server.</li> <li>✓ With Server A as the active node, install SFW with the DMP and MSCS options to Server B.</li> <li>✓ Change the existing disk groups to cluster disk groups.</li> <li>✓ Prepare a disk group for the dynamic mirrored quorum.</li> </ul>
“ <a href="#">Setting up MSCS cluster groups for the applications</a> ” on page 479	<ul style="list-style-type: none"> <li>✓ If you have applications on the server that you want to cluster, create MSCS cluster groups for them.</li> </ul>
“ <a href="#">Installing applications on the second computer</a> ” on page 479	<ul style="list-style-type: none"> <li>✓ Install the applications’ program files on the local drive of Server B.</li> </ul>

**Table 19-2** Tasks for server consolidation adding MSCS and DMP

Objectives	Tasks
<a href="#">“Completing the setup of the application group in MSCS”</a> on page 480	<ul style="list-style-type: none"> <li>✓ Complete the cluster application group by adding resources and setting dependencies.</li> </ul>
<a href="#">“Changing the quorum resource to the dynamic quorum resource”</a> on page 480	<ul style="list-style-type: none"> <li>✓ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier.</li> <li>✓ Make that disk group a Volume Manager Disk Group type resource in the default Cluster Group.</li> <li>✓ Change the quorum resource to the dynamic mirrored quorum resource.</li> </ul>
<a href="#">“Verifying the cluster configuration”</a> on page 480	<ul style="list-style-type: none"> <li>✓ Test the cluster by moving the cluster resources to the other node.</li> </ul>
<a href="#">“Enabling DMP”</a> on page 480	<ul style="list-style-type: none"> <li>✓ Using DMP, include the main storage array and, optionally, the direct-attached storage devices. Now attach the second path to the configuration and rescan.</li> </ul>

### More on DMP paths

In this configuration, there are two DMP paths, one going through Switch 1, which includes HBA 1 from Server A, HBA 1 from Server B, Channel 1 from the first direct-attached storage device, and Channel 1 from the second direct-attached storage device. The second path includes HBA 2 from Server A, HBA 2 from Server B, Channel 2 from the first direct-attached storage device, and Channel 2 from the second direct-attached storage device.

**Caution:** Do not have the second path to the storage connected to the SAN until DMP is installed and the storage array is included under DMP. If you allow two paths to the storage without DMP control, data can become corrupted.

The two switches keep the paths separate. You could use one large switch and zone it with two zones, one for each path.

In most DMP configurations, direct-attached storage is not included along with a storage array, but it is shown in this example to demonstrate that you can use direct-attached storage with DMP.

---

**Note:** When Storage Foundation for Windows is first installed, DMP control is not in effect. All arrays attached to the system come up as excluded. You must include the storage array and any direct-attached storage devices to enable DMP.

---

## Reviewing the configuration requirements

Reviewing the prerequisites and the configuration allows you to gain an overall understanding of the configuration and its requirements.

### Prerequisites

These procedures assume:

- Experience in setting up computer hardware, switches, and storage arrays
- Familiarity with the Windows operating systems and SFW commands

### Supported software

- Veritas Storage Foundation 5.0 for Windows (SFW)  
*or*  
Veritas Storage Foundation HA 5.0 for Windows (SFW HA)
- One of the following operating systems:
  - Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
  - Windows Server 2003 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition (SP 1 required for all editions)
  - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
  - Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)
  - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition
  - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

or

Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

## Hardware setup

Assume an original configuration of one large already consolidated server, one storage array, and two direct-attached storage devices that are all connected on a SAN.

Add:

- 1 large server of the same type as the first server
- Fibre switch
- 2 HBAs, one for each computer, required for DMP
- 6 network interface cards, 3 for each server (2 each for the private network and 1 for the public network)

---

**Note:** Refer to the Hardware Compatibility List on the Symantec Support web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW and SFW HA.

---

### Recommendations

It is acceptable to use one NIC for the private network, but using two cards is strongly recommended to avoid making the private network a single point of failure in the configuration.

Refer to the Microsoft documentation for the specific requirements for the MSCS cluster. For example, you will need a static IP address for each network interface card and a static IP address for the cluster. The two clusters need to be members of the same domain.

Refer to [Chapter 15, “Deploying SFW with MSCS” on page 329](#) for more information on the process of setting up an MSCS cluster to work with SFW.

## Adding the new hardware

Install the necessary hardware on both Server A and Server B.

### To add the new hardware

- 1 Verify that your data from the large server is backed up before proceeding.
- 2 Install two host adapters in each server.

---

**Caution:** Do not connect the second path through HBA 2 on each server at this time.

---

- 3 Install the three network interface cards in each server. Do not make the connections between the two servers at this time.
- 4 Do any necessary configuration of the second switch without actually connecting it to the servers.

## Establishing the MSCS cluster

Complete the steps necessary to install a cluster on Server A and Server B, using MSCS. Refer to the Microsoft documentation for the detailed instructions. The general steps are:

- ✓ Do the necessary network configuration steps on Server A.  
For example, establish the static IP addresses of the network cards and make sure a domain is set up that can be used by the two servers on the cluster.
- ✓ On Server A, access SFW and create a 500 MB partition on a disk that will be used as the quorum disk when the first node of the cluster is created. You may need to revert a dynamic disk to basic to implement this step.
- ✓ Create the first node of the cluster on Server A, using Cluster Administrator.
- ✓ Install the Windows Server 2003 operating system on Server B and do the networking configuration steps for Server B.
- ✓ Connect the networks between the two sites and verify their connectivity.
- ✓ Add the second node of the cluster to Server B.
- ✓ Test the cluster by moving the cluster resources from Server A to Server B. Server B becomes the active node. At this point, keep the control of the cluster with Server B.

## Adding SFW support to the cluster

Use the following procedure to add SFW support to the cluster.

### To add SFW support to the cluster

- 1 With the active node of the cluster on Server B, use **Add/Remove Programs** on Server A to add the MSCS and DMP options to SFW on that server and reboot. Then move the cluster resources back to server A. Server A is now the active node.

---

**Note:** If you reboot a server that has the active node of the cluster, it will fail over to the other node. You have more control of the situation by moving the resources to the other node before doing a reboot.

---

- 2 On Server B, install SFW with the DMP and MSCS options and reboot.
- 3 On Server A, which is now the active node of the cluster, use SFW to create a dynamic cluster disk group that will be used for the dynamic quorum. The disk group should contain three disks, and the disk size is recommended to be 500 MB. You need to create a three-way mirrored volume on the three disks with SFW. You can also use two disks, but three disks provide added redundancy.
- 4 Change the existing regular SFW dynamic disk groups on Server A to cluster disk groups.

A regular dynamic disk group is converted to a cluster disk group through the command line by using the command to import a disk group, `vxdg import`, with the `-s` option, the option that does the conversion. You will need to deport the disk groups first before you can import them. You can deport them through the GUI **Deport Dynamic Disk Group** command.

## Setting up MSCS cluster groups for the applications

If you have applications on the server that you want to cluster, you need to set up an MSCS cluster group for each application. Set up the groups first before the application is installed because if the application is cluster-aware, it may need to reference the cluster group. For detailed steps on setting up MSCS cluster groups, see [“Setting up a group for the application in MSCS”](#) on page 356.

Note that you will not be able to finish setting up the resources for the group until the application is installed on the second node.

## Installing applications on the second computer

If you have one or more applications on the existing computer and you want their data and associated files to be clustered, you need to install the applications on the local drive of the new computer. The applications may be cluster-aware and require specific procedures to install. Refer to the application documentation.

For tips on installing applications in an MSCS environment, see [“Installing the application on cluster nodes”](#) on page 358.

## Completing the setup of the application group in MSCS

Once the application is installed, complete the configuration of the application group in MSCS. For details, see [“Completing the setup of the application group in MSCS”](#) on page 360.

## Changing the quorum resource to the dynamic quorum resource

For details about changing the quorum resource to the dynamic quorum resource, see [“Implementing a dynamic quorum resource”](#) on page 361.

## Verifying the cluster configuration

Verify that the cluster can fail over by moving the cluster group manually between the nodes to make sure it works properly. For details, see [“Verifying the cluster configuration”](#) on page 364.

## Enabling DMP

These steps assume that SFW with the DMP option has been installed. See [“Adding SFW support to the cluster”](#) on page 478 in this example.

### To enable DMP

- 1 With SFW on the first server, bring up DMP and include the disks on the storage array and optionally the two direct-attached storage devices. To include each storage array or direct-attached storage device under DMP control:
  - a Display the Array Settings screen for the device you are including by doing the following:
    - In the tree view under the **Disks** icon, select a disk from the storage array.
    - In the right pane, click the **Paths** tab for the disk. Only one path should display in the **Paths** tab, since the disk is not yet under DMP control.
    - Right-click the path and select **Array Settings** from the path context menu that comes up.
    - The Array Settings window comes up. The **Exclude** checkbox is checked.
  - b Uncheck the **Exclude** checkbox.
- 2 Using appropriate cables, connect the second path on Server A to Switch 2.

- a Connect the path through Server A, HBA 2, Channel 2 of the direct-attached storage, and Controller 2 of the large storage array.
  - b Complete any necessary configuration of the switch.
- 3 Go to **Actions** and select **Rescan** to verify that two paths are shown under the **Paths** tab. This indicates that one set of disks has two paths and that DMP is installed correctly.
- 4 Complete [step 1](#) to [step 3](#) on Server B.

MSCS and DMP are now set up, and the upgraded configuration steps are complete.

## SFW features that support server consolidation

With consolidated servers, Storage Foundation for Windows has multiple features that assure fault tolerance and improve storage utilization. Many of those features are highlighted in the section “[Advantages of using SFW with server consolidation](#)” on page 460.

The following section adds more information about some of the features. It describes how to create a script for Automatic Volume Growth based on capacity and gives a high-level view of SFW features for supporting storage in a SAN and for performance management. Topics in this section include:

- [Automatic volume growth](#)
- [Features that support storage in a SAN](#)
- [Performance monitoring](#)

### Automatic volume growth

Storage Foundation for Windows comes with an Automatic Volume Growth feature that monitors the capacity of dynamic volumes and automatically increases the size of the volume when used space on it reaches a predetermined size.

With this procedure, you can conserve disk space on your servers because space is distributed automatically on an as-needed basis. You do not have to be available to allocate the additional disk space when it is required.

An example of this feature is presented in the section “[Example 2: Automatic volume growth](#)” on page 69.

## Features that support storage in a SAN

In a SAN environment, it is important to protect storage so that it cannot be accessed by more than one host at a time. SFW provides the feature of private dynamic disk group protection that protects a disk group with a SCSI reservation so that other hosts cannot access the data. For more information on this feature, see the *Veritas Storage Foundation Administrator's Guide*. Clustering is another way to protect the storage in a SAN. It also uses a SCSI reservation to keep the disk group from being accessed by other hosts in a SAN.

## Performance monitoring

The statistics feature of SFW provides I/O statistics to allow performance tuning to improve overall disk and system performance. Through the Online Monitoring window, hot spots are identified. A hot spot is an area of high I/O activity that may cause bottlenecks in I/O throughput. If a disk has these hot spots, consider moving one or more of its subdisks to another disk that shows below-average I/O activity. For more information on this topic, see the *Veritas Storage Foundation Administrator's Guide*.

## Server consolidation customer success story

At El Camino Hospital (ECH) in Mountain View, California—known as “The Hospital of Silicon Valley” and a world leader in the use of IT innovations and high-tech devices—there is never a good time for a problem to occur or to take systems offline for maintenance or upgrades. Even 99% uptime may not be good enough when patient care is at stake.

“We aim for 99.999%, because those clinical systems have to be there for physicians and nurses all the time,” said Joe Wagner, CTO of El Camino Hospital. “With Veritas, we deliver all departments, all applications, to all end users, 24 x 7. We've seen instant results and tangible cost savings, for a 24-month return on investment of about \$3.4 million.”

El Camino Hospital had been using Veritas backup products for years when it began researching software-based approaches that could cost-effectively deliver high availability. Like most hospitals, ECH was also looking for ways to consolidate servers to save money on hardware, software, and maintenance. “We came to the conclusion that Veritas's high-availability solution could also deliver the server consolidation we were looking for,” said Wagner.

Veritas Consulting helped ECH design and implement a solution based on a many-to-one (many active to one standby) cluster strategy. “With Veritas, we have 15 live clusters or systems able to fail over to a single machine,” explains Wagner. “This means huge cost savings compared to hardware-based high-

availability solutions.” The fact that Veritas is not platform-specific was also key to selecting this approach.

The return on the hospital’s investment has been “tremendous,” Wagner reported. “For more than a year now, we’ve been seeing cost savings in servers, storage, operating systems, network interface cards, host adapters, ports, power supplies, tape drives, and licenses, for a 24-month ROI of about \$3.4 million,” he explained. Significant savings have also been realized through cost avoidance and staff redeployments.

In addition to optimizing the ROI on the Veritas products the hospital had purchased, Veritas Consulting also worked closely with hospital staff to create a comprehensive business continuity plan for an upcoming data center move, without disrupting the critical ongoing activity in the hospital. “Veritas has solved our immediate need for high availability, and now it’s enabling us to address the broader issues of business continuity planning,” said Wagner. “Veritas is our one-stop shop for continuous, cost-effective computing at El Camino Hospital.”



# Appendix

- [Deploying Disaster Recovery: Manual implementation](#)



# Deploying Disaster Recovery: Manual implementation

This chapter provides the steps for setting up a disaster recovery (DR) solution, using SFW HA with the Veritas Volume Replicator (VVR) and Global Cluster Option (GCO) in a new installation. The example describes a generic database application.

For examples of the SFW HA disaster recovery solution with specific applications, see the other Solutions Guides included with this release: *Veritas Storage Foundation and High Availability Solutions, Solutions Guide for Microsoft Exchange* and *Veritas Storage Foundation and High Availability Solutions, Solutions Guide for Microsoft SQL*.

The process of setting up and working with the SFW-VVR disaster recovery solution has five parts:

- [“Part 1: Setting up the cluster on the primary site”](#) on page 492
- [“Part 2: Setting up the cluster on the secondary site”](#) on page 544
- [“Part 3: Adding the VVR components for replication”](#) on page 546
- [“Part 4: Adding GCO components for wide-area recovery”](#) on page 561
- [“Part 5: Maintaining: Normal Operations and recovery procedures”](#) on page 568

The steps for setting up the cluster described in the High Availability section of this guide are the basic foundation for this disaster recovery solution (see [Chapter 7, “Deploying SFW HA for high availability: New installation” on page 63](#)). The main differences in the process of setting up the cluster for a disaster recovery, rather than for HA alone, are that you need to make sure that the VVR and the GCO options are selected during the SFW HA installation. You

also need to configure the Veritas Volume Replicator Security Service (VxSAS) after the installation completes. For the secondary site, the cluster is set up in a similar manner as on the primary site.

Once the two clusters are set up, one at the primary site and the other at the secondary site, VVR is used to enable replication from the primary site to the secondary site.

The Global Cluster Option allows the two clusters to become global clusters and to be able to fail over to one another. Normally, two independent clusters cannot fail over to each other.

The following table outlines the DR process for this configuration in more detail. The high-level objectives and the tasks to complete each objective for the configuration are as follows:

**Table A-1**

Objective	Tasks
“ <a href="#">Reviewing the requirements</a> ” on page 492	<ul style="list-style-type: none"> <li>✓ Verify hardware and software prerequisites.</li> <li>✓ Review configuration requirements.</li> </ul>
<b>“Part 1: Setting up the cluster on the primary site.”</b>	
“ <a href="#">Installing and configuring the hardware</a> ” on page 497	<ul style="list-style-type: none"> <li>✓ Set up and configure the hardware according to the manufacturer’s instructions.</li> </ul>
“ <a href="#">Installing Windows and configuring network settings</a> ” on page 497	<ul style="list-style-type: none"> <li>✓ Install the operating system on both nodes.</li> <li>✓ Make necessary networking settings on both nodes.</li> </ul>
“ <a href="#">Installing SFW HA (Primary site)</a> ” on page 499	<ul style="list-style-type: none"> <li>✓ Verify the driver signing option for the systems.</li> <li>✓ Install SFW HA (automatic installation). <ul style="list-style-type: none"> <li>■ Select the option to install VVR.</li> <li>■ Select the Global Cluster Option for VCS to enable wide-area failover.</li> </ul> </li> </ul>

Table A-1

Objective	Tasks
“Configuring VVR security service” on page 504	<ul style="list-style-type: none"> <li>✓ Complete the steps to configure VxSAS. The VxSAS wizard is launched automatically after the reboot of the local node at the end of the installation.</li> </ul>
“Configuring the cluster (Primary site)” on page 506	<ul style="list-style-type: none"> <li>✓ Verify static IP addresses and name resolution configured for each node.</li> <li>✓ Configure cluster components using the Veritas Cluster Server Configuration wizard.</li> <li>✓ Set up secure communication for the cluster.</li> </ul>
“Configuring disk groups and volumes (Primary site)” on page 523	<ul style="list-style-type: none"> <li>✓ Create disk groups.</li> <li>✓ Create volumes.</li> </ul>
“Installing the application on cluster nodes (Primary site)” on page 530	<ul style="list-style-type: none"> <li>✓ Install the application program files on the local drive of the first node.</li> <li>✓ Install files relating to the data and logs on the shared storage.</li> <li>✓ Deport the disk groups on the first node and import them on the second node.</li> <li>✓ Make sure that the volumes on the second node have the same drive letters or mount points as on the first node.</li> <li>✓ Install the application on the second node.</li> </ul>
“Creating VCS service groups (primary site)” on page 532	<ul style="list-style-type: none"> <li>✓ Use an appropriate method to create the VCS service group for the application.</li> </ul>
“Verifying the cluster configuration” on page 542	<ul style="list-style-type: none"> <li>✓ Switch the service group to the second node.</li> <li>✓ Switch it back to the first node</li> </ul>

Table A-1

Objective	Tasks
<b>“Part 2: Setting up the cluster on the secondary site.”</b>	
“Creating a parallel environment on the secondary site” on page 544	<ul style="list-style-type: none"> <li>✓ Setting up a parallel environment on the secondary site involves: <ul style="list-style-type: none"> <li>■ Installing and configuring hardware</li> <li>■ Installing Windows and configuring network settings</li> <li>■ Installing SFW HA</li> <li>■ Configuring the cluster</li> </ul> </li> </ul>
“Configuring disk groups and volumes (Secondary site)” on page 545	<ul style="list-style-type: none"> <li>✓ Create exactly the same disk groups and volumes on the secondary site as on the primary site.</li> </ul>
“Installing the application (Secondary site)” on page 545	<ul style="list-style-type: none"> <li>✓ Install the application program files on the local drive of the first node.</li> <li>✓ Install files relating to the data and logs on the shared storage.</li> <li>✓ Install the second node like the first node.</li> </ul>
“Configuring the Service group for VCS (Secondary site)” on page 545	<ul style="list-style-type: none"> <li>✓ Use an appropriate method to create the VCS service group for the application.</li> </ul>
<b>“Part 3: Adding the VVR components for replication.”</b>	
“Configuring the Replicator Log volumes for VVR” on page 547	<ul style="list-style-type: none"> <li>✓ Use SFW to create Replicator Log volumes for the primary and secondary sites.</li> </ul>
“Setting up the replicated data sets (RDS) for VVR” on page 549	<ul style="list-style-type: none"> <li>✓ Create Replicated Data Sets with VVR’s Replicated Data Set wizard and start replication for the primary and secondary sites.</li> </ul>
“Creating the VVR RVG Service group” on page 557	<ul style="list-style-type: none"> <li>✓ Create a VVR RVG service group for the replicated volume group.</li> </ul>

Table A-1

Objective	Tasks
<b>“Part 4: Adding GCO components for wide-area recovery.”</b>	
“Linking clusters by adding a remote cluster” on page 562	✓ Create a global cluster by adding the first cluster to the second one through the command <b>Add/Delete Remote Cluster</b> .
“Converting a local Service group to a global group” on page 564	✓ Convert service groups that are common to all clusters to global service groups.
<b>“Part 5: Maintaining: Normal Operations and recovery procedures.”</b>	
“Normal operations: Monitoring the status of the replication” on page 568	<ul style="list-style-type: none"> <li>✓ Monitor replication</li> <li>✓ Perform planned migration</li> </ul>
“Disaster recovery procedures” on page 569	✓ Complete the recovery procedures after the primary site goes down.

## Part 1: Setting up the cluster on the primary site

This section details the steps for creating the cluster on the primary site.

This disaster recovery solution requires a primary site and a secondary site.

### Reviewing the requirements

Review these product installation requirements for your systems before installation. Minimum requirements and Veritas recommended requirements may vary.

---

**Note:** Refer to the Hardware Compatibility List on the Symantec Support Web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW and SFW HA.

---

### Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Installation on a non-system drive requires space on both the system drive and the non-system drive.

[Table A-2](#) estimates disk space requirements for SFW HA.

**Table A-2** Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
SFW HA + all options + client components	1675 MB	Non-system space: 1675 MB System space: 345 MB
SFW HA + all options	1230 MB	Non-system space: 1230 MB System space: 285 MB
Client components	630 MB	Non-system space: 630 MB System space: 115 MB

## Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.0 Hardware Compatibility List to confirm supported hardware at: <http://entsupport.symantec.com>

### Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA) with the Veritas Volume Replicator and Global Clustering Options
- Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)  
*or*  
Windows Server 2003 Web Edition (limited to file share support for SFW HA), Windows Server 2003 Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 for all editions)  
*or*  
Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition  
*or*  
Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)  
*or*  
Windows Server 2003 for Intel Xeon (EM64T) or AMD Opteron: Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition  
*or*  
Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

### System requirements

Systems must meet the following requirements:

- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs.

See “[Best practices](#)” on page 495.

- 1 GB of RAM for each system.
- All servers must run the same operating system, service pack level, and system architecture.

## Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows 2000 or Windows Server 2003 domain.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.
- Static IP addresses for the following purposes:
  - One static IP address available per site for each application virtual server
  - One IP address for each physical node in the cluster
  - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
  - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers. See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information.
- For a DR environment, all sites must reside in the same Active Directory domain.

## Permission requirements

- You must have administrative access to all systems selected for cluster operations. For this, add the domain user to the local Administrators group of each system.
- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for objects corresponding to these nodes in the Active Directory.
- Domain Administrator or Account Operator privileges: if you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

## Additional requirements

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on C : \WINDOWS of one node, installations on all other nodes must be on C : \WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- When installing, install only in a single domain or workgroup.

## Best practices

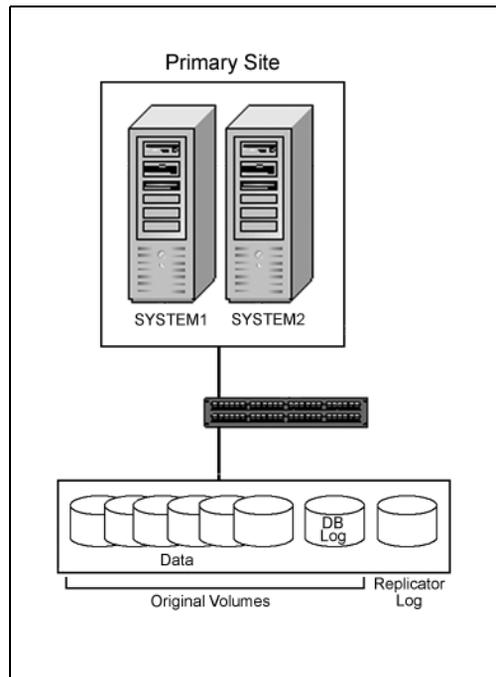
- Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Symantec recommends three network adapters (two NICs exclusively for the private network and one for the public network).  
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

## Reviewing the configuration

This configuration overview describes active/passive high availability within a cluster and disaster recovery between two sites. In an active/passive configuration, one or more application virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

Active/passive clusters involve one-to-one failover capabilities. For instance, if you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM5 and SYSTEM6 on the secondary site), then SYSTEM1 can fail over to SYSTEM2, and SYSTEM5 can fail over to SYSTEM6. The figure that follows illustrates the cluster configuration on the primary site. For a view of the DR configuration that includes both sites, see the illustration in the section “[About a disaster recovery solution](#)” on page 238.

**Figure A-1** DR configuration primary site



## Installing and configuring the hardware

This section summarizes the steps for the hardware installation. For specific details on installing the hardware, refer to the hardware documentation.

### To set up the hardware

- 1 Install the required network adapters and SCSI controllers or Fibre Channel HBA on each system on the primary site.  
Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- 2 Connect the network adapters on each system.  
To prevent lost heartbeats on the private networks and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Symantec also recommends removing TCP/IP from private NICs to lower system overhead. Contact the NIC manufacturer for details on this process.
- 3 Verify that each system can access the storage devices.
- 4 Reboot each system. Verify that each system recognizes the attached shared disks.

## Installing Windows and configuring network settings

This section focuses on network configuration procedures that are requirements for a VCS setup. Refer to Microsoft documentation for complete information on installing the operating system and network configuration.

### To install Windows and configure network settings

- 1 Install the Windows 2000 or Windows Server 2003 operating system on each node.
- 2 Configure the necessary network settings for the network cards and the domain setup on each node.
- 3 Verify DNS settings for all systems on which the application will be installed.
  - Open the Control Panel (**Start > Control Panel**).
  - Open **Network and Dial-up Connections**.
  - Make sure that the public network adapter is the first bound adapter:

- From the **Advanced** menu, click **Advanced Settings**.
- On the **Adapters and Bindings** tab, verify that 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 and click **OK**.
- In the Network and Dial-up Connections window, double-click the adapter for the public network.

When enabling DNS name resolution, make sure that you use the public network adapters, not those configured for the VCS private network.

- From the status window, click **Properties**.
  - On the **General** tab:
    - Select the **Internet Protocol (TCP/IP)** checkbox.
    - Click **Properties**.
  - Select the **Use the following DNS server addresses** option.
  - Verify that the value for the IP address of the DNS server is correct.
  - Click **Advanced**.
  - On the **DNS** tab, make sure the **Register this connection's address in DNS** checkbox is selected.
  - Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field, and click **OK**.
- 4 Use Windows Disk Management on each system to verify that the attached shared disks are visible.

# Installing SFW HA (Primary site)

The product installer enables you to install the software for Veritas Storage Foundation HA 5.0 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. For a disaster recovery configuration, select the option to install GCO. If you plan to use VVR for replication, you must also select the option to install VVR.

## Setting Windows driver signing options

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

[Table A-3](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

**Table A-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.

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

### To install the product

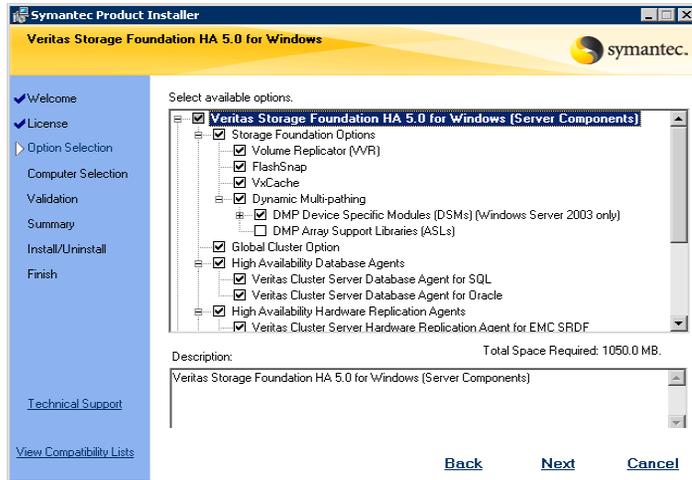
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.
- 3 Click **Storage Foundation HA 5.0 for Windows**.



- 4 Do one of the following:
  - Click **Complete/Custom** to begin installation.
  - Click the **Administrative Console** link 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**.

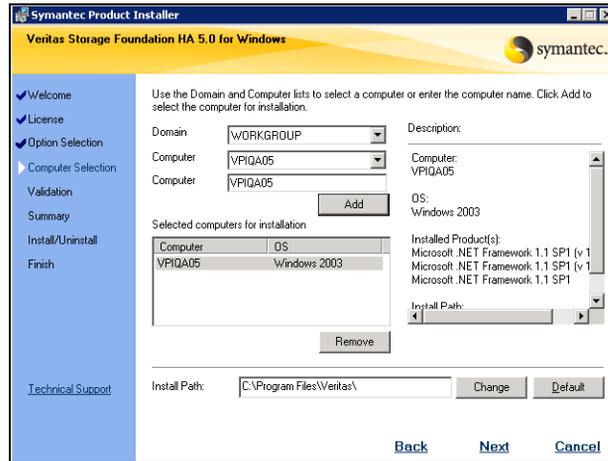
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.

- 8 Repeat for additional license keys. Click **Next**
  - 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 Select the appropriate SFW product options and click **Next**.



- |                           |  |
|---------------------------|--|
| Client                    | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console. Required to install the Solutions Configuration Center which provides information and wizards to assist configuration |
| Global Cluster Option     | Required for a disaster recovery configuration only.   |
| Veritas Volume Replicator | If you plan to use VVR for replication, you must also select the option to install VVR.  |

## 10 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

- 11 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.
- 12 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 list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
  - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
  - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs 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.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.  
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

## Configuring VVR security service

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server 2003 environment, 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 <b>Password</b> 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 username and password when configuring the VxSAS service on the other hosts.

After providing the required information click **Next**.

- 3 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**.

- 4 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 the **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 the **Configure** button to proceed with configuring the VxSAS service.

- 5 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.
- 6 Click **Finish** to exit the wizard.

## Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

### 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 the cluster (Primary site)

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters. The GCO option is needed for a disaster recovery solution. Thus, once you have completed the cluster configuration steps described in this section, complete the steps in the section that follows to configure the GCO option.

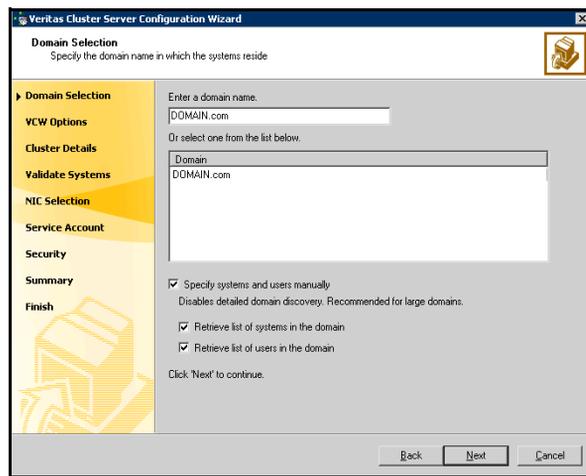
Complete the following tasks before creating a cluster:

- Verify that each node uses a static IP address (DHCP is not supported) and that name resolution is configured for each node.
- Set the required permissions:
  - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
  - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
  - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

**To configure a VCS cluster**

- 1 Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



To discover information about all systems and users in the domain:

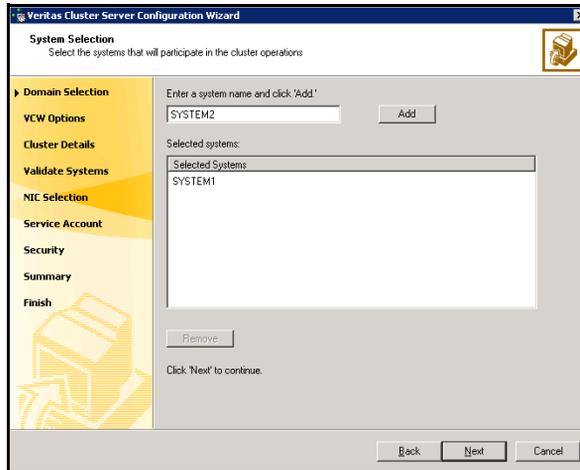
- Clear the **Specify systems and users manually** check box.
- Click **Next**.  
Proceed to [step 7](#) on page 509.

To specify systems and user names manually (recommended for large domains):

- Check the **Specify systems and users manually** check box.  
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

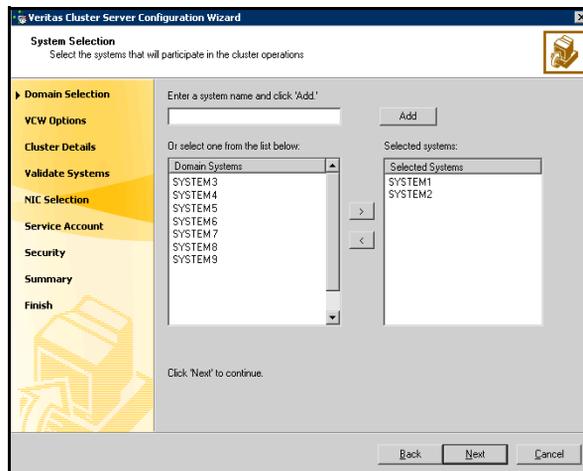
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 508. Otherwise proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



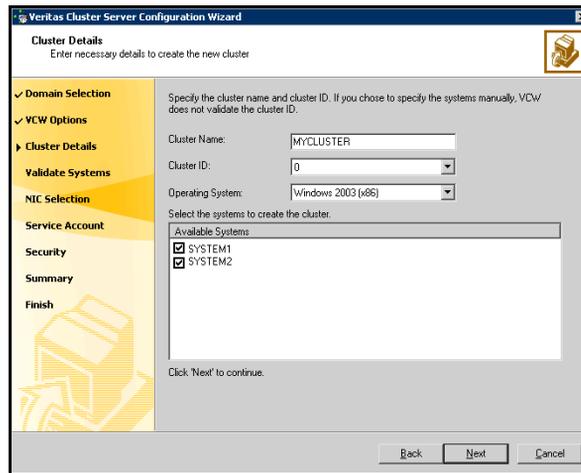
Proceed to [step 7](#) on page 509.

- 6 On the System Selection panel, specify the systems to form a cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.

- 7 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 8 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



**Cluster Name** Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

**Cluster ID** Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

---

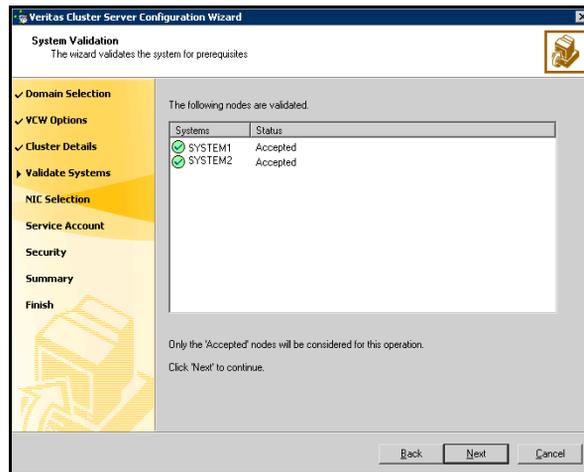
**Warning:** If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

---

**Operating System** From the drop-down list, select the operating system that the systems are running.

**Available Systems**      Select the systems that will be part of the cluster. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

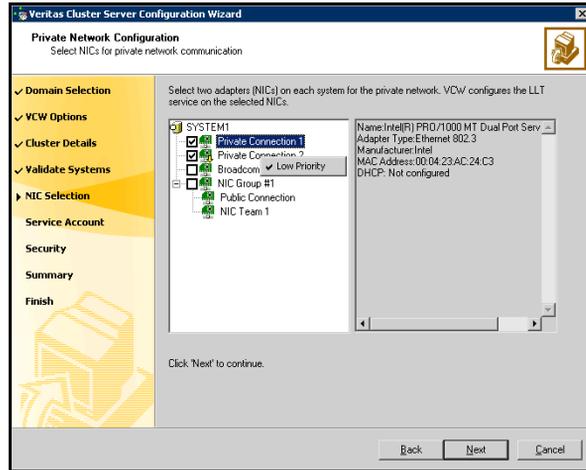
- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

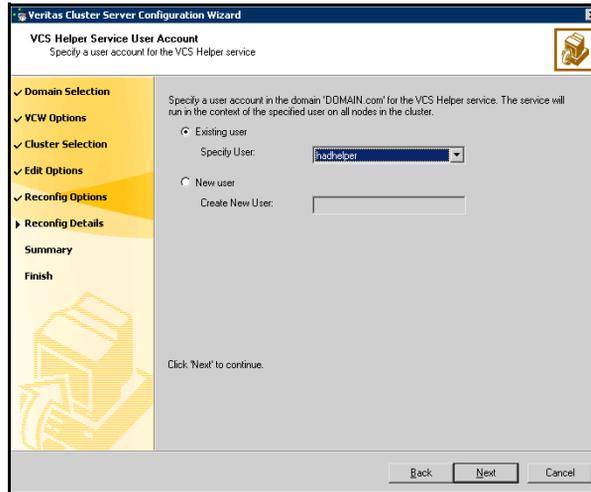
If you chose to configure a private link heartbeat in [step 8](#) on page 509, proceed to the next step. Otherwise, proceed to [step 11](#) on page 511.

- 10 On the Private Network Configuration panel, configure the VCS private network and click **Next**.



- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
  - If you have only two NICs on a selected system, make sure you lower the priority of the NIC that is used for public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
  - If your configuration contains teamed NICs, the wizard groups them as NIC Group #N where N is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.
- 11 On the VCS Helper Service User Account panel, specify the name of a domain user context for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper

Service user context to access the network. Do not use the Administrator account for the VCS Helper service.



- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list
  - If you chose not to retrieve the list of users in [step 4](#) on page 507, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

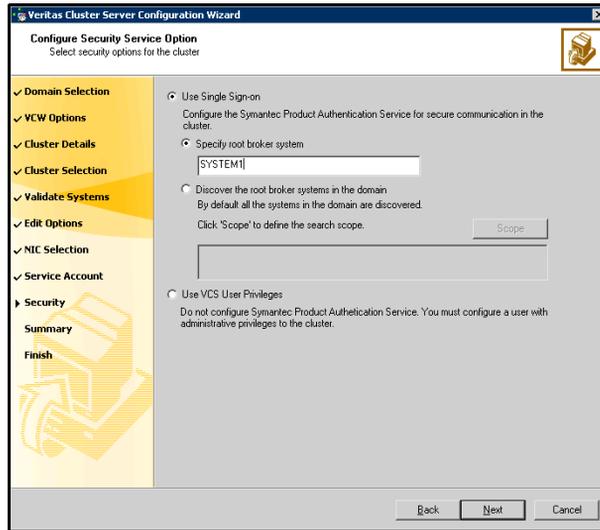
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 12 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition. For example, to search for systems managed by Administrator, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.
- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use VCS user privilege

Veritas Cluster Server Configuration Wizard

Configure Security Service Option  
Select security options for the cluster

Domain Selection  
 VCS Options  
 Cluster Details  
 Cluster Selection  
 Validate Systems  
 Edit Options  
 NIC Selection  
 Service Account  
 Security  
 Summary  
 Finish

Use Single Sign-on  
 Configure the Symantec Product Authentication Service for secure communication in the cluster.

Use VCS User Privileges  
 Do not configure Symantec Product Authentication Service. You must configure a user with administrative privileges to the cluster.  
 Provide the cluster administrator user information below.

User Name:   
 Password:   
 Re-enter Password:

Back Next Cancel

- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

The default user name for the VCS administrator is admin and the default password is password. Both are case-sensitive. Use this account

to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

**13 Review the summary information on the Summary panel, and click **Configure**.**

The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.

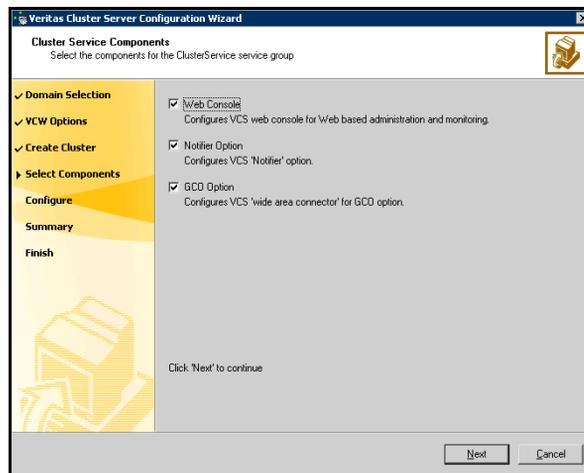
The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

**14 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Web Console, notification, and for global clusters.**

To configure the ClusterService group later, click **Finish**.

At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

**15 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.**



- Check the **Web Console** check box to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.

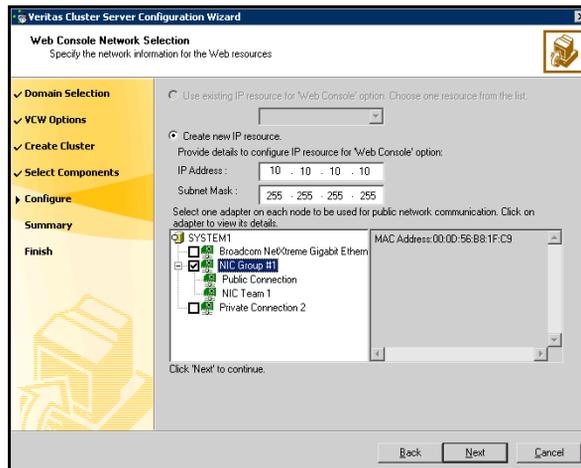
The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

## Configuring the Web Console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

### To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
- If you choose to configure a new IP address, type the IP address and associated subnet mask.

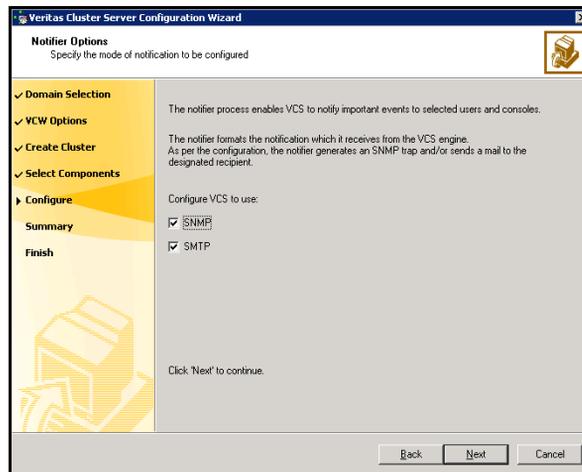
- Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
  - 3 If you chose to configure a Notifier resource, proceed to “[Configuring notification](#)” on page 517.  
If you chose to configure global cluster components, proceed to “[Configuring the wide-area connector process for global clusters](#)” on page 521.  
Otherwise, click **Finish** to exit the wizard.

## Configuring notification

This section describes steps to configure notification.

### To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.



You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

The screenshot shows the 'Notifier SNMP Configuration' window in the Veritas Cluster Server Configuration Wizard. The window title is 'Veritas Cluster Server Configuration Wizard' and the subtitle is 'Notifier SNMP Configuration'. Below the subtitle, it says 'Specify information about SNMP console'. On the left, there is a navigation pane with options: 'Domain Selection', 'VCW Options', 'Create Cluster', 'Select Components', 'Configure', 'Summary', and 'Finish'. The main area contains a table for entering console information and a text input for the trap port.

SNMP Console	Severity
snmpserv	Information
snmpserv1	SevereError

Enter name or IP of the SNMP console and severity level for each

Click on "+" button to add more consoles.  
Click "-" to remove a console.

Enter SNMP Trap Port:

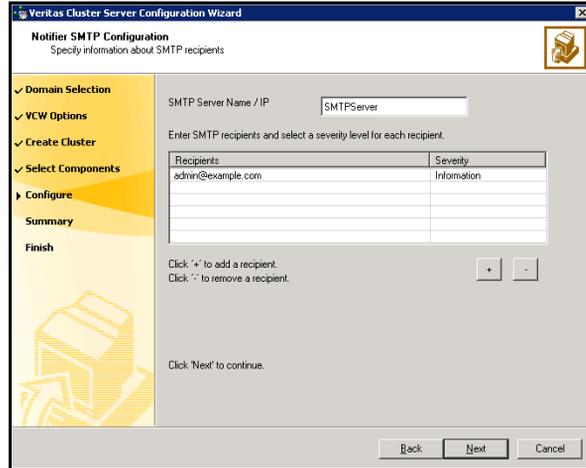
Note: SNMP console must be MIB 2.0 compliant

Click 'Next' to continue.

Buttons: Back, Next, Cancel

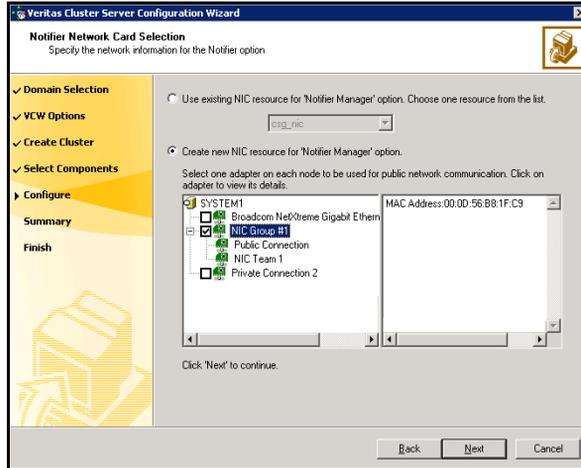
- Click a field in the SNMP Console column and enter the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click + to add a field; click - to remove a field.
- Enter an SNMP trap port. The default value is 162.

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



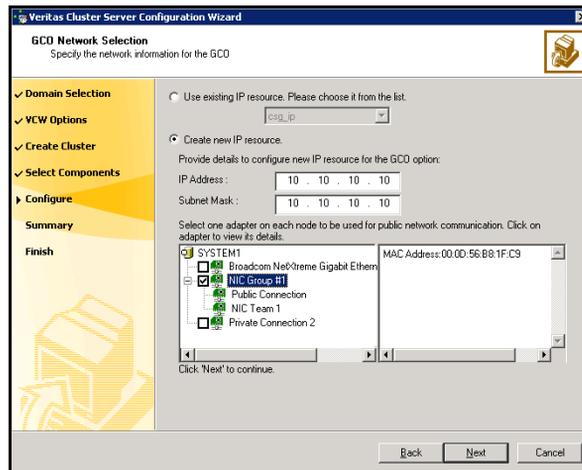
- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
  - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started and click **Configure**.
  - 6 If you chose to configure global cluster components, proceed to [“Configuring the wide-area connector process for global clusters”](#) on page 521. Otherwise, click **Finish** to exit the wizard.

## Configuring the wide-area connector process for global clusters

This section describes steps to configure the wide-area connector resource required for global clusters.

### To configure the wide-area connector process for global clusters

- 1 On the GCO Network Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
  - If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
  - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.
  - 3 Click **Finish** to exit the wizard.

## Configuring global cluster components

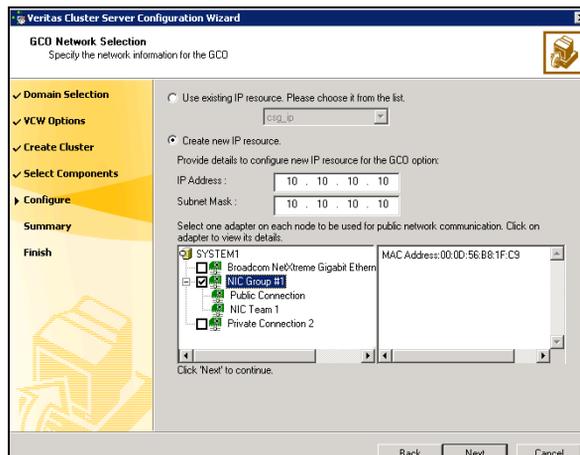
The next task is to identify an IP resource for the wide-area connector that is required for inter-cluster communication. If the cluster has a ClusterService

group configured, you can use the IP address configured in the service group or configure a new IP address.

This task does not set up a global cluster environment. That process is done later and is described in “[Part 4: Adding GCO components for wide-area recovery](#)” on page 561.

### To configure an IP resource for GCO

- 1 In the GCO Network Selection panel, specify the network information:

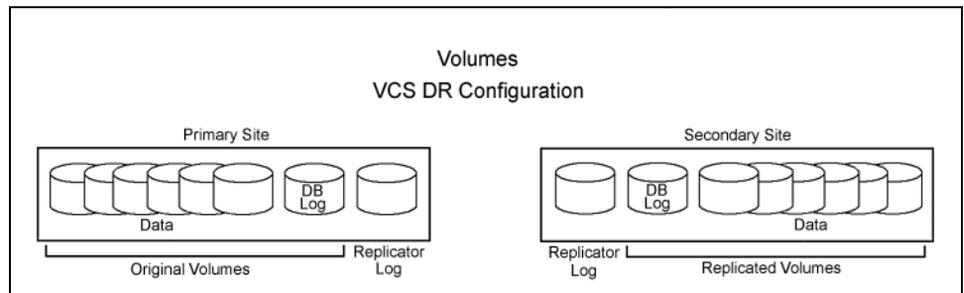


- If the cluster has a ClusterService group configured, you can use the IP address configured in the service group or configure a new IP address.
  - If you choose to configure a new IP address, enter the IP address and associated subnet mask.  
Make sure the specified IP address has a DNS entry.
  - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
  - Click **Next**.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts.
  - 3 Click **Configure**.
  - 4 Click **Finish** to exit the wizard.

# Configuring disk groups and volumes (Primary site)

The following figure shows a typical setup of volumes for a VCS disaster recovery configuration with a database application. The example has one disk group on each site.

**Figure A-2** VCS clustered database volumes, DB log, and Replicator Log



Use Veritas Storage Foundation for Windows to create cluster disk groups and dynamic volumes for the application on the shared storage. A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different RAID layouts. Configuring disk groups and volumes involves the following tasks:

- [“Planning disk groups and volumes”](#) on page 523
- [“Configuring disk groups and volumes \(Primary site\)”](#) on page 523
- [“Creating dynamic volumes”](#) on page 527

## Planning disk groups and volumes

Decide how you want to organize the disk groups and the number and type of volumes you want to create. Some considerations are:

- The number of disk groups that are needed  
The number of disk groups depends on your application and the planned organization of the data. VCS requires that the application program files be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage. Typically, a main organizational unit in your application, such as the storage group in Microsoft Exchange, would be contained in a single disk group.
- The type of volumes you want to create

- Mirrored and RAID-5 volumes provide fault tolerance for critical data.
- Striped volumes add performance capability.
- Volumes that are both mirrored and striped offer both performance and fault tolerance.

---

**Note:** If you plan to use replication software, such as VVR, do not use RAID-5 volumes. This does not apply to hardware RAID-5. VVR also does not support volumes with commas in the names.

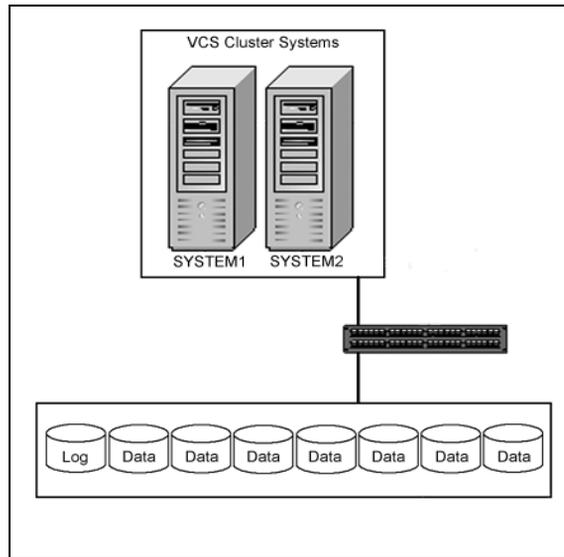
---

Recommendations:

- Use mirrored volumes for logs.
- Use striped or mirrored striped volumes for data.
- The implications of backup and restore operations for the disk group setup.
- The sizes of databases and logs, which depend on the traffic load.

The following illustration shows a typical setup of disks for a clustered database application with shared storage. The log volume takes a single disk. Volumes for data and associated files take the remaining disks. Because you have dynamic volumes, the volumes can span multiple disks. You can have a mirrored striped volume that uses the disks other than the log disk.

**Figure A-3** VCS clustered database with disks for data and the log



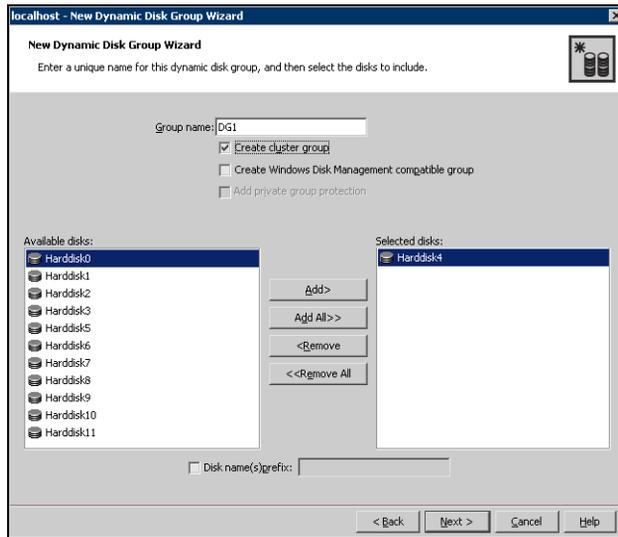
## Creating dynamic cluster disk groups

Follow the steps in this section to create one or more disk groups for your application.

### To create a dynamic (cluster) disk group

- 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:



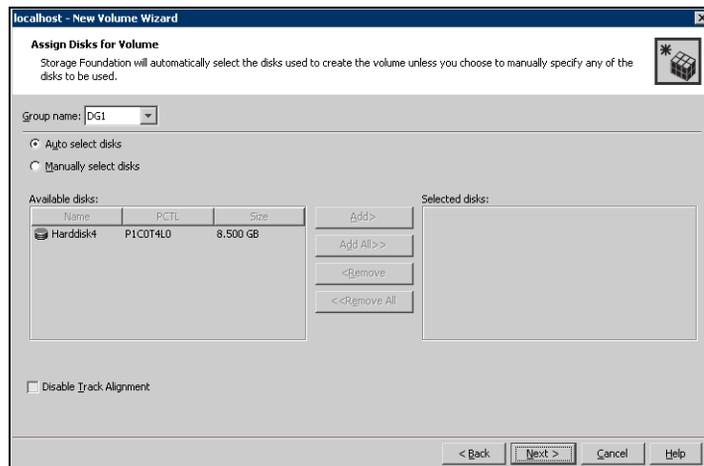
- Enter the disk group name (for example, DG1).
  - Click the checkbox for **Create cluster group**.
  - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter 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.
  - 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

Once the disk groups are created, make the disks within them usable by creating the dynamic volumes that will store data.

### 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.
- 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.

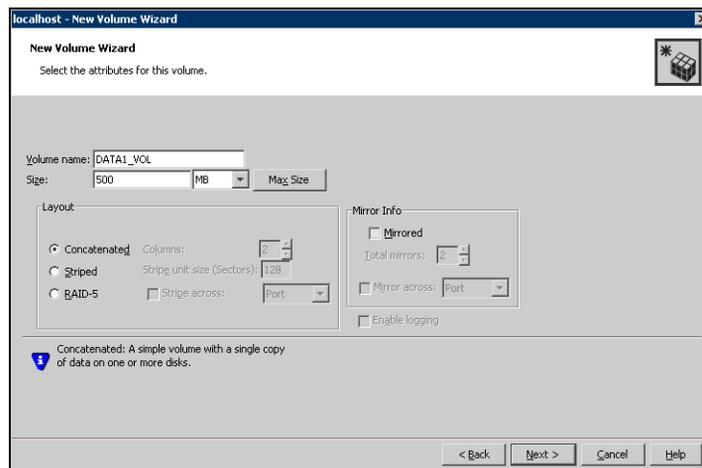


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

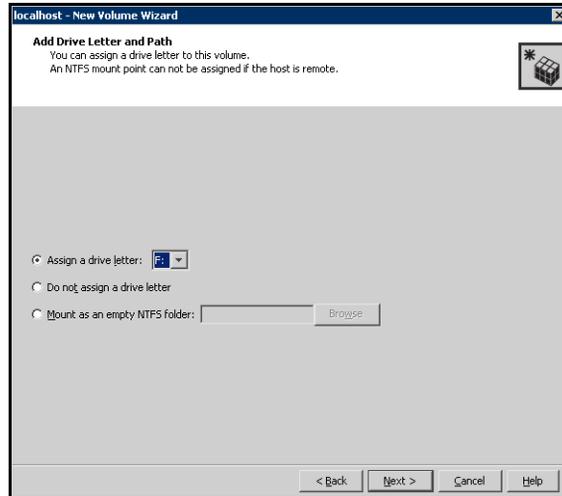
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

- 8 Click **Next**.
- 9 Specify the volume attributes.



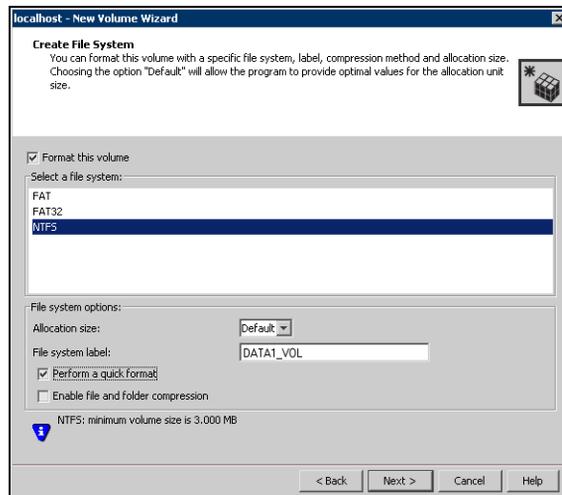
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
  - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
  - If you are creating a striped volume, the **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.

## Installing the application on cluster nodes (Primary site)

VCS requires that the application program files be installed on the same local drive of all cluster nodes and that the application data and log files or other files related to the application data be installed on the shared storage.

### 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.
- Make sure that the disk groups and volumes are imported and thus mounted on the server before you install the application.
- If you have just created the disk groups and volumes, they will be mounted and accessible. When a disk group is created, it is automatically imported on that node. You can verify that the disk group and volumes are accessible if you can see the disk group and volume icons in the VEA GUI for the server.
- 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 when installing the application. Instead, click to browse to the dynamic volumes that were prepared previously.

## Pointers for installing the application on the second node

- To install the application on the second node, deport any disk groups from the first node and import them on the second node. Steps for deporting and importing disk groups are in the section “[Deporting and importing a disk group](#)” on page 531.
- You need to 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 instructions in the section “[To add or change a drive letter or mount point](#)” on page 532.
- 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. You would then restart the service after the application is installed.

## Deporting and importing a disk group

This topic describes the steps for deporting and importing a disk group in order to install the application on the second node.

### To deport a disk group on the first node

- 1 If SFW is not already running, start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**). If the Storage Foundation Assistant automatically opens, close it.
- 2 Navigate to **dynamic disk groups** on the node on which the dynamic disk group is currently imported.
- 3 Right-click the dynamic disk group to be deported and click **Deport**.

### To import the dynamic disk group on the second node

- 1 Start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**). If the Storage Foundation Assistant automatically opens, close it.
- 2 Navigate to **dynamic disk groups** on the node to which you will import the dynamic disk group.
- 3 Right-click the dynamic disk group to be imported and click **Import**. There may be no drive letter associated with an existing dynamic volume when it is imported to a computer for the first time. Use SFW to add or change drive letters, as needed. Make sure that drive letters or mount

points for the volumes on the second node are the same as were used on the first node.

#### To add or change a drive letter or mount point

- 1 In SFW, 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.
- 3 To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.
- 4 To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Select the new drive letter and click **OK**.
- 5 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, and click **OK** to mount the volume.

A mount point is also referred to as a “drive path.”

To change a mount point, you must remove it and recreate it ([step 5](#)). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.

## Creating VCS service groups (primary site)

This section on VCS service groups first describes the VCS service group and then presents an example of creating a service group with a generic database application.

### About VCS service groups

In order for VCS to be able to monitor and fail over an application in a cluster, the application must be included in a VCS service group.

A service group is a collection of resources working together to provide application services to clients. It can also relate to a print or a file share that does not contain a specific application. A service group’s resources fail over as a group to another cluster node when there is an application failure or server failure on the active node.

VCS has a collection of defined resource types. For each type, VCS has a corresponding agent that provides a type-specific logic to control resources. The

bundled agents come with the program and apply to a group of resources that are commonly used with applications.

VCS provides multiple methods for creating a service group. If you are using Microsoft Exchange Server, Microsoft SQL Server, or Oracle, VCS provides a wizard for each of these, but you need to purchase the VCS agents for these programs. There are also separate wizards for file and print servers. In addition, there are several ways to create a service group such as using the VCS Java Console, as well as an Application Configuration wizard. A command line interface can be used to create a service group as well.

All these different methods accomplish the same purposes:

- Defining the cluster resources and their attributes.
- Setting their dependencies; for example, a NIC resource depends on an IP resource.
- Logically grouping the resources together.
- Providing capabilities for monitoring the service group and taking it online or offline.

For more information refer to the *Veritas Cluster Server Administrator's Guide*.

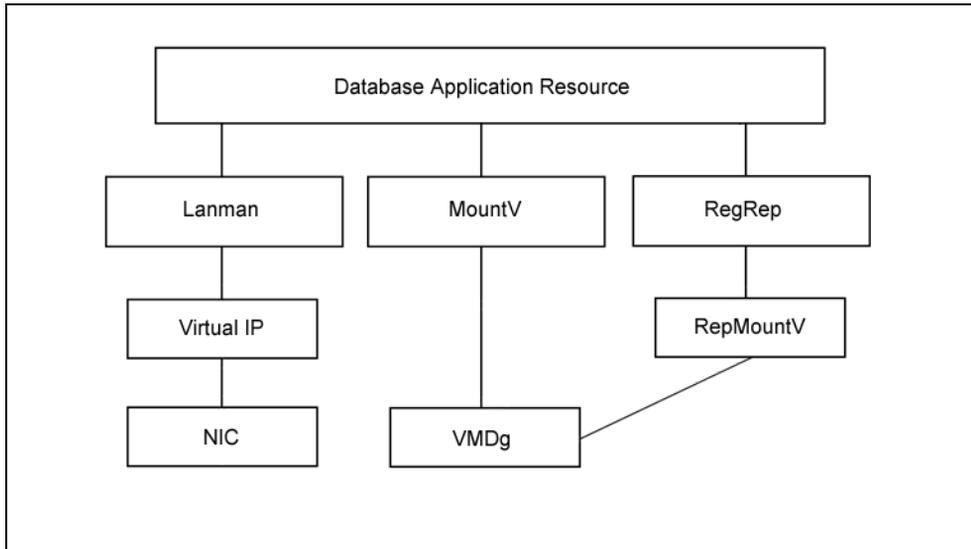
The next section illustrates how to create a VCS service group using the Application Configuration wizard.

## Service group example with a generic database application

The following steps show an example of creating a service group for a generic database application.

### Application description

Assume that the application is a database application that runs as a Windows service. Its resources with their dependencies are shown in the chart that follows:

**Figure A-4** Database application resources for the Service group example

- The Lanman resource makes the application available to clients. It depends on other resources that are associated with it.
- The virtual IP resource identifies the cluster and allows the cluster to communicate across the network. It depends on the NIC being configured for it to function.
- The MountV resource mounts the SFW disk group volumes and depends on the VMDg resource, which includes the SFW disk groups.
- The RegRep resource replicates the registry of the active cluster node and depends on the RepMountV resource and the VMDg resource.

The resources at the bottom of the chart have to be made available or brought online before the linked resources above them. When the cluster is shut down, the resources need to be brought down in the opposite order, from top to bottom. When you use the Application Configuration wizard to create a service group, it establishes these resources.

### Prerequisites

- Verify that the binaries of the application to be configured are present on the nodes on which the service group will be configured.
- Verify that the shared drives required by the applications are mounted.

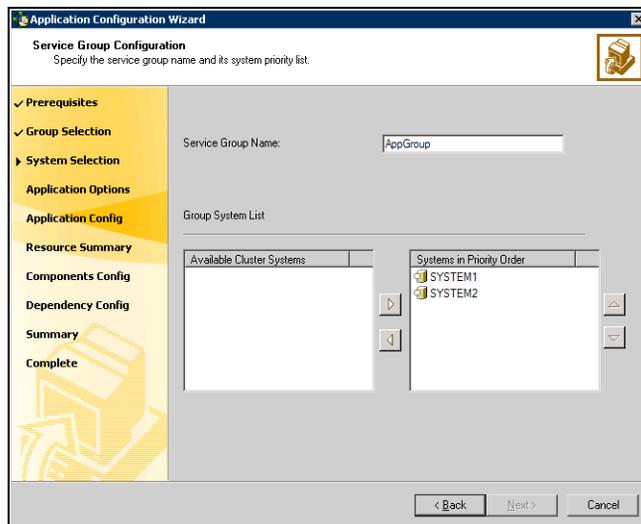
- Before running the wizard, make sure you have the following information ready:
  - Type of applications for which resources are to be configured.
  - Shared storage used by the applications.
  - Registry replication information.
  - Network information.

## Creating the service group with the application configuration wizard

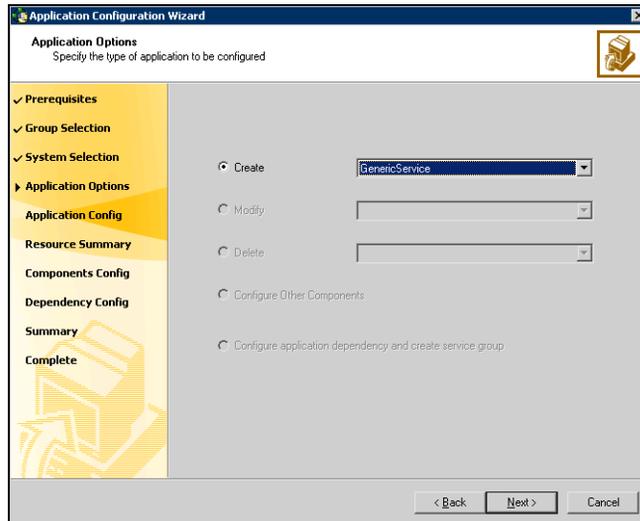
The steps given in this section may not be the exact steps needed to create a service group for every database application. They are presented to give a typical example of what may be involved in creating a VCS service group for a database application.

### To establish a service group for an application database

- 1 Select **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Configuration Wizard**.
- 2 Read the information in the Welcome panel and click **Next**.
- 3 In the **Wizard Options** panel, click **Create service group** to add a new service group to the cluster. Click **Next** to continue.
- 4 Complete the following in the Service Group Configuration panel:



- Enter a name for the service group in the **Service Group Name** field. Specify a name that conveys information about the service group you are creating.
  - Select systems on which to configure the service group from the “Available Cluster Systems” list and add them to the “Systems in Priority Order” list. Arrange them in the order that matches their position in the cluster. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority.
  - Click **Next** to continue.
- 5 The wizard validates the configuration.
- 6 In the Applications Options panel that appears next, do the following:

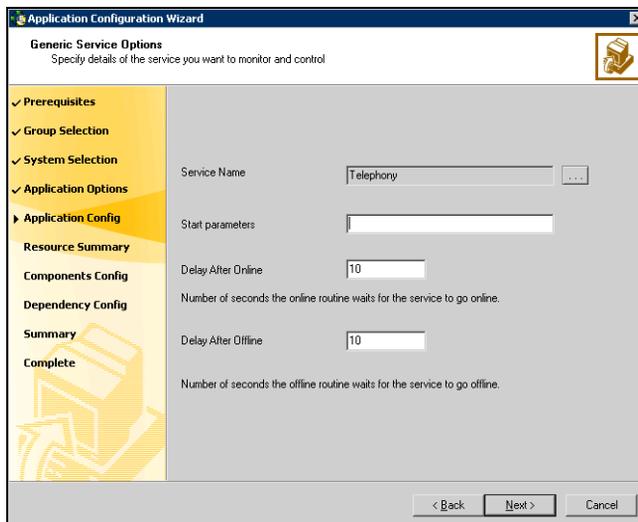


- Select **Create** to create a service group for an application.
- Select the agent type used to bring the application online or offline and to monitor its status.  
For this example, GenericService would be selected because the database application runs as a Windows service.  
If the application does not run as a generic service, you should select the Process agent at this point.  
The Service Monitor agent does not control an application. It can be used to monitor a resource that the application resource may depend on. It does not bring the application online or offline. It monitors a

service, starts a user-defined script, and interprets the exit code of the script.

- Click **Next** to continue.

7 In the Generic Service Options panel, do the following:



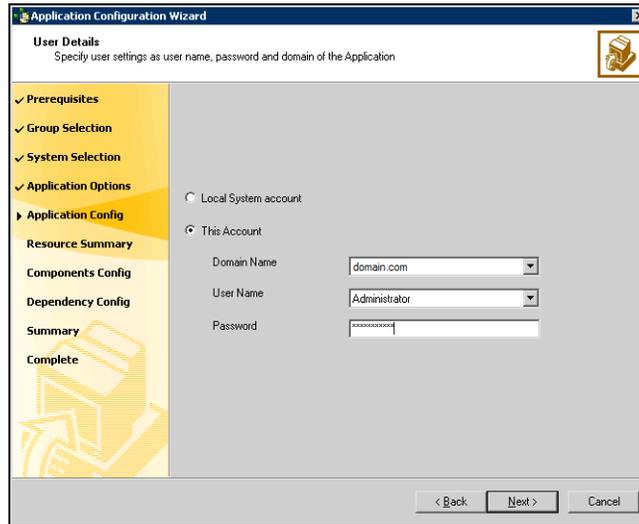
The screenshot shows the 'Application Configuration Wizard' window, specifically the 'Generic Service Options' step. The window title is 'Application Configuration Wizard' and the subtitle is 'Generic Service Options - Specify details of the service you want to monitor and control'. On the left, a navigation pane lists several steps: Prerequisites, Group Selection, System Selection, Application Options, Application Config (selected), Resource Summary, Components Config, Dependency Config, Summary, and Complete. The main area contains the following fields and labels:

- Service Name:** A text box containing 'Telephony' and a button with three dots to its right.
- Start parameters:** An empty text box.
- Delay After Online:** A text box containing '10', with the label 'Number of seconds the online routine waits for the service to go online.' below it.
- Delay After Offline:** A text box containing '10', with the label 'Number of seconds the offline routine waits for the service to go offline.' below it.

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

- Select the service name.  
Click the icon to the right of the **Service Name** entry box to bring up a list of Windows services and click the desired service. In this example, the service for the database would be selected.
- Provide the start parameters for the service, if applicable.
- In the **Delay After Online** field, specify the number of seconds the agent waits after the service is brought online before starting the monitor routine.
- In the **Delay After Offline** field, specify the number of seconds the agent waits after the service is taken offline before starting the monitor routine.
- Click **Next**.

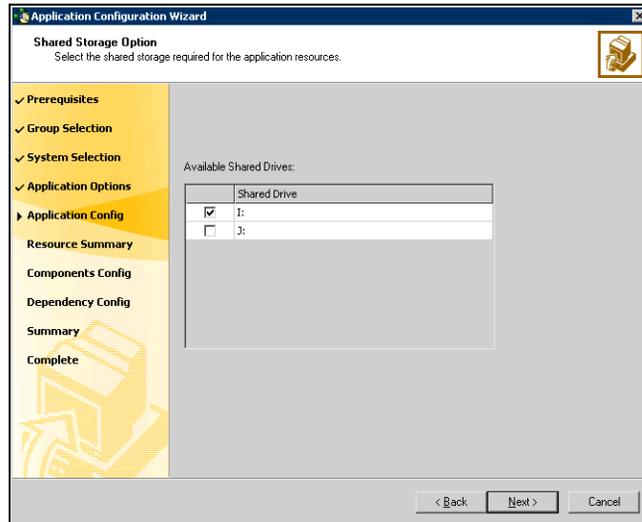
## 8 Make the necessary settings in the User Details window:



The screenshot shows the 'User Details' window of the 'Application Configuration Wizard'. The window title is 'Application Configuration Wizard' and the subtitle is 'User Details'. Below the subtitle is the instruction: 'Specify user settings as user name, password and domain of the Application'. The window is divided into a left sidebar and a main content area. The sidebar contains a list of steps: 'Prerequisites', 'Group Selection', 'System Selection', 'Application Options', 'Application Config' (which is expanded), 'Resource Summary', 'Components Config', 'Dependency Config', 'Summary', and 'Complete'. The main content area has two radio buttons: 'Local System account' (unselected) and 'This Account' (selected). Below these are three input fields: 'Domain Name' with a dropdown menu showing 'domain.com', 'User Name' with a dropdown menu showing 'Administrator', and 'Password' with a masked input field. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

- Select **This Account** to have the service group run in the cluster.
- Specify the following details about the user.
  - Select the domain in the **Domain Name** list box.
  - Specify the user in the **User Name** list box.
  - Enter the password for the user in the **Password** field.
- Click **Next** to continue.

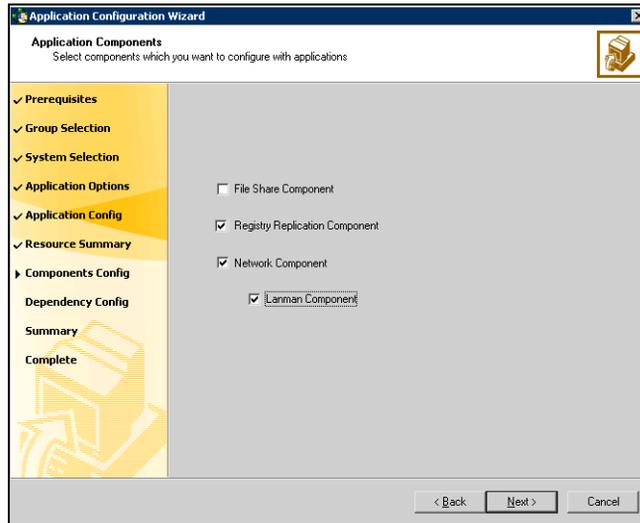
- 9 In the Shared Storage Option window, select the shared storage required for the GenericService resource by clicking the check box adjacent to the shared drive or drives. Click **Next**.



The wizard determines from the storage indicated that SFW cluster disk groups are involved, thus adding the VMDg resource and the MountV resource to the service group.

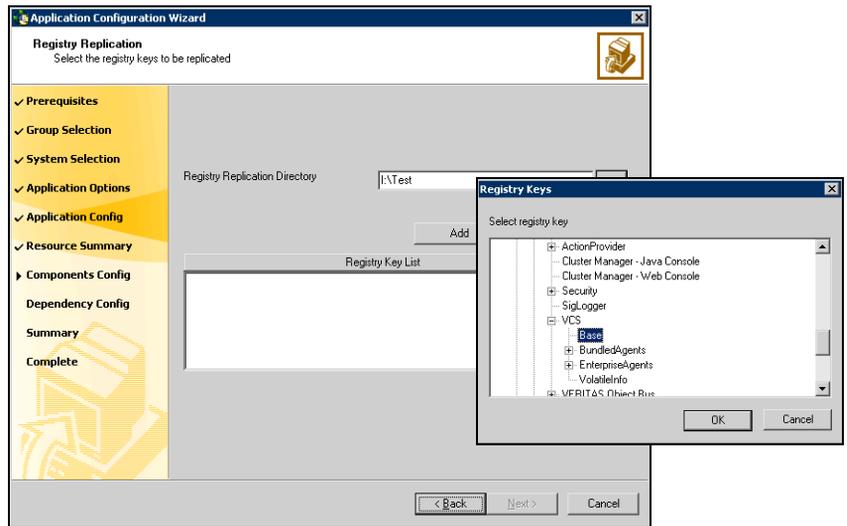
- 10 In the Application Resource Summary panel, review the summary of the GenericService resource. Click **Back** to make changes. Otherwise, click **Next**.
- 11 The Application Options panel appears. Select the **Configure Other Components** radio button to configure additional resources for the service group and click **Next**.
- 12 In the Application Components window, check the **Registry Replication Component** and **Network Component** check boxes to add the resources for

replicating the registry as well as the NIC, IP, and Lanman resources to the service group. Click **Next**.



- 13 In the Registry Replication window, specify the registry keys to be replicated by doing the following:
- Specify the directory on the shared disk in which the registry changes are logged.
  - Click **Add**.

- In the **Registry Keys** panel, select the registry key to be replicated:



- Click **OK**.
- The selected registry key is added to “Registry Key List” box. Click **Next**.  
The RegRep and MountV resources are added to the service group.

14 In the Virtual Computer Configuration panel, specify the information related to your network:

- Enter a unique virtual computer name by which the node will be visible to the other nodes. Note that the virtual name must not exceed 15 characters.

The **Virtual Computer Name** field will not be displayed if you did not select **Lanman Component** in the Application Components panel.

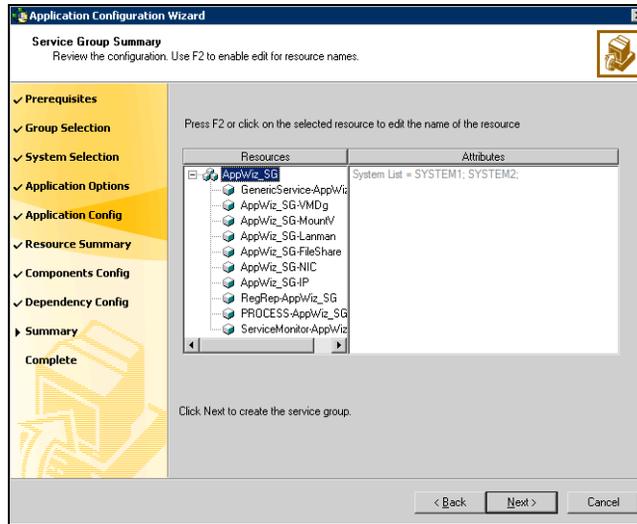
- Enter a unique virtual IP address for the virtual server.
- Enter the subnet to which the virtual server belongs.
- For each system in the cluster, select the public network adapter name. To view the adapters associated with a system, click the **Adapter Name** field and click the arrow.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if applicable. Verify that you select the adapters assigned to the public network, not the private network.

- Click **Next**.

The Lanman resource, the virtual IP resource, and the NIC resource are added to the service group.

- 15 In the Application Options panel, click **Configure application dependency and create service group** and click **Next**.
- 16 In the Service Group Summary panel, review your configuration:



- Change the names of resources, if required. The wizard assigns unique names to resources. Click a resource name to edit it. Review your configuration and click **Next**.
  - In the Confirmation dialog box, click **No** to review your settings. Otherwise, click **Yes**. The wizard starts running commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion panel appears.
- 17 In the completion dialog box, click **Bring the service group online** check box if you want to bring the service group online on the local system. Click **Finish** to exit the Application Configuration wizard.

## Verifying the cluster configuration

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Veritas Cluster Manager (Java Console) to switch all the service groups from one node to another.
- Simulate a local cluster failover by shutting down an active cluster node.

#### To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
  - Click **Switch To**, and click the appropriate node from the menu.
  - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

#### To shut down an active cluster node

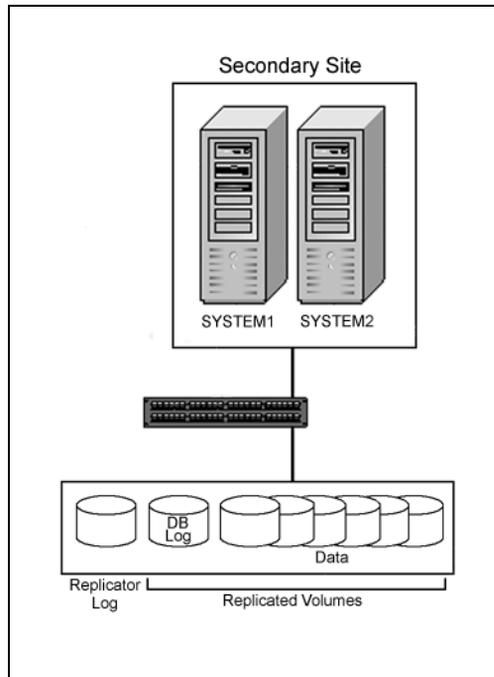
- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node:
  - Restart the node you shut down in [step 1](#).
  - Click **Switch To**, and click the appropriate node from the menu.
  - In the dialog box, click **Yes**.

The service group you selected is taken offline and brought online on the node that you selected.

## Part 2: Setting up the cluster on the secondary site

The tasks documented in this section enable you to set up a nearly identical cluster at the secondary site, with disk groups and volumes with the same names and with the same application files. The data volumes on the secondary site should be the same size as the corresponding data volumes on the primary site. The log volume on the secondary site can be a different size, but Symantec recommends that the sizes be the same. Almost all the steps for creating the cluster on the primary site are repeated on the secondary site.

**Figure A-5** DR configuration secondary site



### Creating a parallel environment on the secondary site

After setting up a SFW HA environment on the primary site, use the guidelines provided in this chapter to complete the same tasks on the secondary site prior to the application installation:

- [“Reviewing the requirements”](#) on page 492
- [“Installing and configuring the hardware”](#) on page 497
- [“Installing Windows and configuring network settings”](#) on page 497

- [“Installing SFW HA \(Primary site\)”](#) on page 499
- [“Configuring the cluster \(Primary site\)”](#) on page 506

After completing these tasks, continue with the next section for instructions on installing the application on the secondary site.

## Configuring disk groups and volumes (Secondary site)

Create an identical disk group and volume setup on the secondary site, as on the primary site. The disks, disk groups, and volumes must be the same sizes and have the same names. The volumes need to be the same volume types.

Refer to [“Configuring disk groups and volumes \(Primary site\)”](#) on page 523 for further directions.

## Installing the application (Secondary site)

Installing the application on the secondary site is similar to installing it on the primary site. See the section [“Installing the application on cluster nodes \(Primary site\)”](#) on page 530.

## Configuring the Service group for VCS (Secondary site)

To configure the service group for VCS, refer to the section, [“Creating VCS service groups \(primary site\)”](#) on page 532.

---

**Note:** Do not bring the service group online if the service group on the primary site is online.

---

## Verifying the cluster configuration (Secondary site)

Refer to the section [“Verifying the cluster configuration”](#) on page 542.

## Part 3: Adding the VVR components for replication

This section provides information on configuring the VVR components for replication. Topics include:

- [“VVR components overview”](#) on page 546
- [“Configuring the Replicator Log volumes for VVR”](#) on page 547
- [“Setting up the replicated data sets \(RDS\) for VVR”](#) on page 549
- [“Creating the VVR RVG Service group”](#) on page 557

### VVR components overview

The terms Replicated Volume Group (RVG), Replicator Log, and Replicated Data Set (RDS) are used frequently in this section. Here are their definitions:

#### Replicated Volume Group (RVG)

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, 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, but Replicator Log volume sizes can differ. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG on the primary host.

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.

#### Replicated Data Set (RDS)

An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).

#### Replicator Log

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 logs at the two sites must have the same name; however, the sizes of the logs can vary. Symantec recommends

having Replicator Log volumes of the same size at the primary site and the secondary site.

The process described in this section involves setting up an RDS for each SFW disk group on the primary site that will have replicated volumes, and then creating a VVR service group that is linked to the application service group.

## Configuring the Replicator Log volumes for VVR

Create the volume for the Replicator Log at each site. The task of creating the logs can also be done during the RDS creation process, but some storage administrators may prefer to do it manually (as is being done here) as a preparatory step to setting up the RDS.

---

**Note:** To improve write performance, Symantec recommends that you create the Replicator Log volume on a different disk from the disks used for your application data volumes.

---

### To configure the Replicator Log volumes for VVR

- 1 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop to open the VEA console on the active node of the primary site.
- 2 Create a volume for the disk group that contains the storage group data:
  - On the System configuration tree, click the disk group where the log volume will be created (*Hostname > Disk Groups > Diskgroupname*).
  - Right-click on the disk group that has the volumes to be replicated, and click **New Volume**.
- 3 On the Welcome panel of the New Volume wizard, click **Next**.
- 4 Select the disks for the volume:
  - Select the group name.
  - Select **Manually select disks**.
  - Click the disk name.
  - Click **Add**.
  - After selecting all the necessary disks, click **Next**.
- 5 Specify the parameters of the volume:
  - Enter the volume name.
  - Enter the size. The size of the Replicator Log volume varies for different environments. To determine the appropriate size for your

environment, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

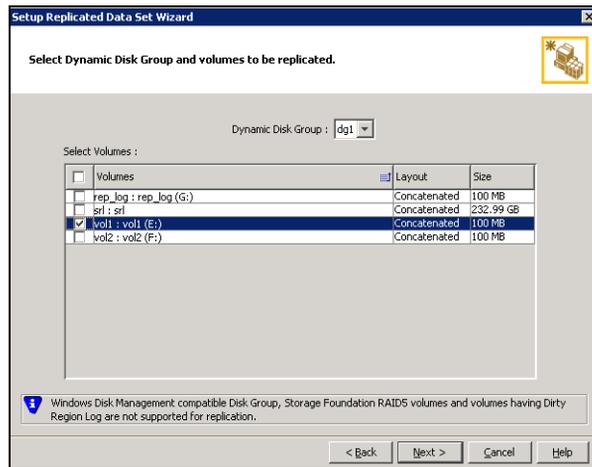
- Select the volume layout.
  - Select the appropriate mirror options.
  - Click **Next**.
- 6 On the Add Drive Letter and Paths dialog box:
    - Click **Do not assign a drive letter**.
    - Click **Next**.
  - 7 When prompted to format the volume:
    - Deselect **Format this volume**.
    - Click **Next**.
  - 8 Click **Finish** to create the new volume.
  - 9 If necessary, repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for any additional RVGs on the primary site.
  - 10 Repeat [step 2](#) through [step 8](#) to create a Replicator Log volume for each RVG on the secondary site.

## Setting up the replicated data sets (RDS) for VVR

Configuring VVR involves setting up the Replicated Data Sets on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

### To create the Replicated Data Set

- 1 From the cluster node on the primary site where the cluster disk group is imported, use the VEA console to launch the Setup Replicated Data Set Wizard. Right-click **Replication Network** on the Management Host configuration tree, and click **Setup Replicated Data Set**.
- 2 Read the Welcome page and click **Next**.
- 3 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

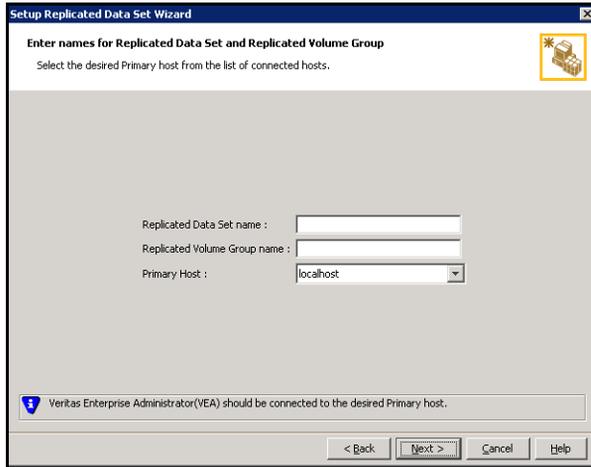


By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

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.

- 4 Click **Next**.

- 5 Select from the table the dynamic disk group and data volumes that will undergo replication.



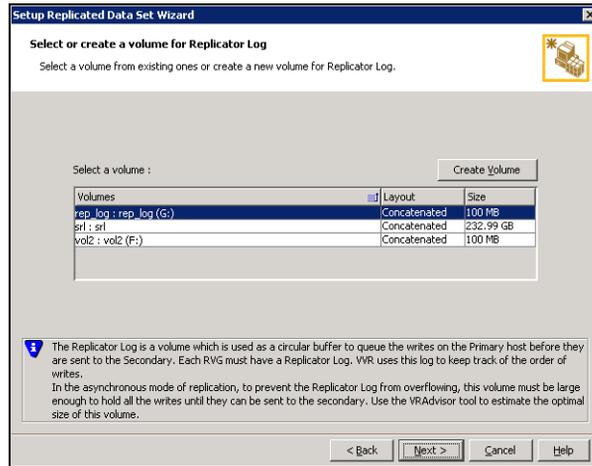
The screenshot shows a dialog box titled "Setup Replicated Data Set Wizard". The main heading is "Enter names for Replicated Data Set and Replicated Volume Group". Below this, a sub-heading reads "Select the desired Primary host from the list of connected hosts." The dialog contains three input fields: "Replicated Data Set name :", "Replicated Volume Group name :", and "Primary Host :". The "Primary Host" dropdown menu is currently set to "localhost". At the bottom left, there is an information icon and a message: "Veritas Enterprise Administrator(VEA) should be connected to the desired Primary host." At the bottom right, there are four buttons: "< Back", "Next > ||", "Cancel", and "Help".

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.

- 6 Click **Next**.

## 7 Select or create a volume for the Replicator Log:



### To select an existing volume

- Select the volume for the Replicator Log in the table (APP\_REPL\_LOG). 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.

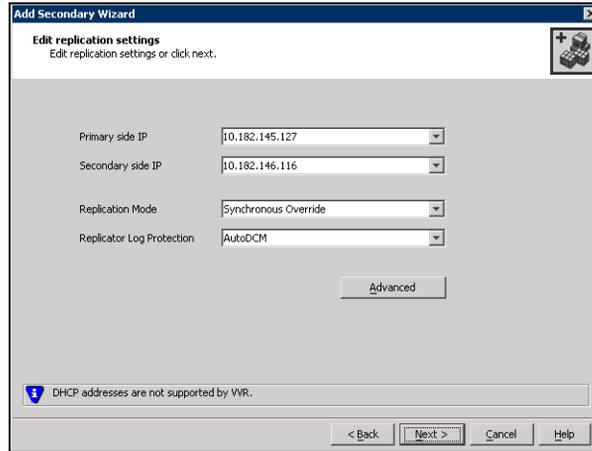
- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 8 Review the information on the summary page and click **Create Primary RVG**.
  - 9 After the RVG for the primary site is successfully created, click **Yes** to add the secondary host to the RDS for replication.
  - 10 Specify the name of the host where the disk group is imported on the secondary site. If necessary, specify the fully qualified domain name.
  - 11 Click **Next**.
  - 12 If the Veritas Enterprise Administrator console is not already connected to the secondary host, the connection process starts when you click **Next**. Enter valid user credentials, click **OK**, and click **Next** again.
  - 13 The configuration for these volumes on the primary and secondary sites must be identical and meet VVR configuration requirements. If a Replicator Log volume does not exist on the secondary site, it can be created with this procedure.
    - 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 If necessary, edit the replication settings for a secondary host.



- Enter the virtual IP address for the Primary IP resource that will be used for replication.
- Select or specify an IP address for the Secondary IP resource.
- Specify the replication mode.

**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**.

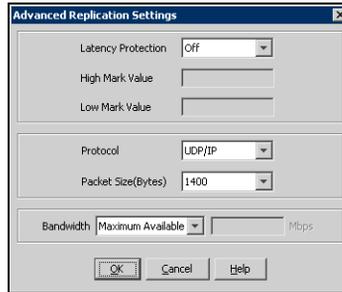
**Synchronous** Determines updates from the application on the primary site are completed only after the secondary site successfully receives the updates.

**Asynchronous** Determines updates from the application on the primary site are completed after VVR stores the updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously

- Specify the replicator log overflow protection property.

<b>AutoDCM</b>	Is the default option and enables the DCM when the Replicator Log overflows even though the secondary site is connected. All data volumes in the primary RVG must have a DCM log to use this option. The wizard attempts to ensure all volumes under the primary RVG have a DCM log.
<b>DCM</b>	Enables Replicator Log protection for the secondary site. DCM is enabled when the Replicator Log overflows and the secondary site is disconnected from the primary site. All data volumes in the primary RVG must have a DCM log to use this option. The wizard attempts to ensure all volumes under the primary RVG have a DCM log.
<b>Off</b>	Disables Replicator Log overflow protection.
<b>Override</b>	<p>Enables log protection. If the secondary site is still connected and the Replicator Log is about to overflow, VVR stalls write operations until five percent or 20 megabytes (whichever is smaller) of the space on the Replicator Log becomes available.</p> <p>If the secondary site becomes inactive because of a connection failure or administrative action, VVR disables Replicator Log protection and causes the Replicator Log to overflow.</p>
<b>Fail</b>	Enables log protection. When the Replicator Log is about to overflow, VVR stalls write operations until five percent or 20 megabytes (whichever is smaller) of the space on the Replicator Log becomes available. If the connection between the primary RVG and secondary RVG is broken, subsequent write operations to the primary RVG fail.

- 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.

16 Click **OK** to close the dialog box.

17 Click **Next**.

18 On the **Start Replication** page, accept the **Synchronize Automatically** option, which is the default recommended for initial setup.

19 Select **Start Replication**, which is the default.

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.

If the virtual IPs have been created, select **Start Replication** to start synchronization immediately.

If replication must be started later, use the **Start Replication** option of VEA to begin replication. Refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator’s Guide* for additional details.

20 Click **Next**.

21 Review the specifications and click **Finish** to add the host on the secondary site to the RDS. Click **Back** to change any information. Replication physically starts when the IP address is created.

## Creating the VVR RVG Service group

Run the wizard from the system that contains the application service group.

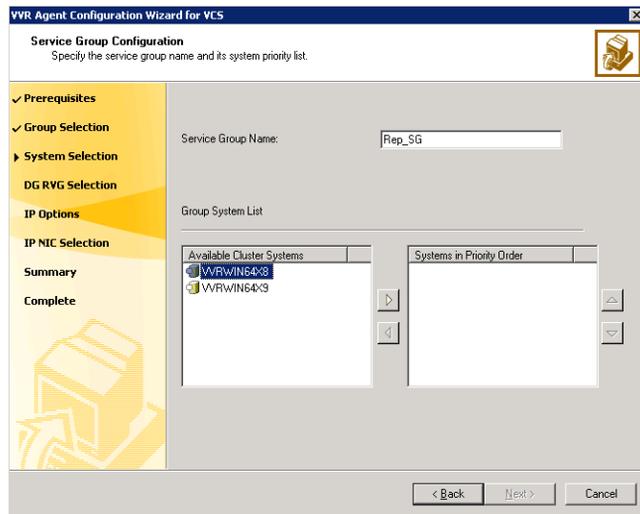
Prerequisites:

- ✓ Verify that the disk group is imported on the node on which you want to create the Replication Service Group.
- ✓ Verify VCS is running, by running the following command on the host on which the you intend to run the Volume Replicator Agent Configuration Wizard.

```
> hasys -state
```

To create a replication service group

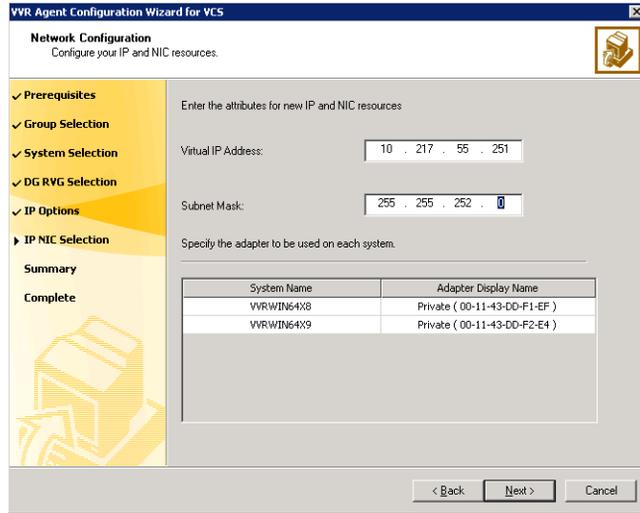
- 1 From the active node of the cluster at the primary site, click **Start > All Programs > Symantec > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the Welcome page, click **Next**.
- 3 In the **Wizard Options** dialog box:
  - a Click **Create a new replication service group**.
  - b Click **Next**.
- 4 Specify the service group name and system priority list:



- a Enter the service group name.

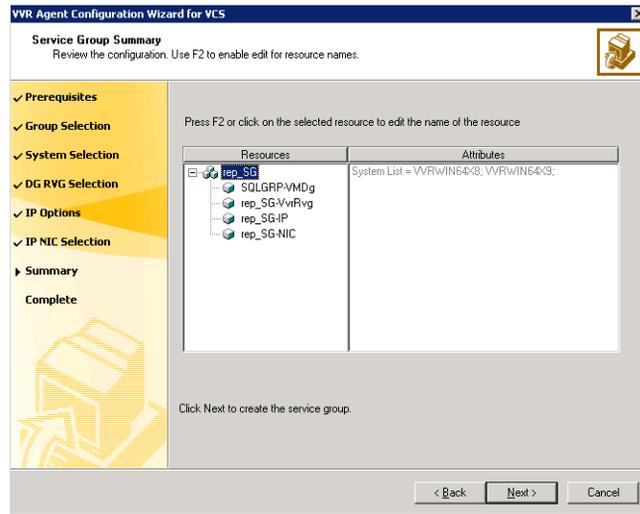


8 Enter the network information:



- a Verify or enter the virtual IP address; use the IP address specified as the primary IP address when you configured the RDS.
- b Specify the subnet mask.
- c Specify the adapters for each system in the configuration.
- d Click **Next**.

## 9 Review the summary of the service group configuration:



The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the Attributes box.

- a If necessary, change the resource names; the wizard assigns unique names to resources based on their respective name rules.  
To edit a resource name, click the resource name and modify it. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.
  - b Click **Next** to create the replication service group.
- 10 A warning informing you that the service group will be created is displayed. When prompted, click **Yes** to create the service group.
  - 11 Click **Finish** to bring the replication service group online.
  - 12 Check the prerequisites, then repeat the wizard at the secondary site, specifying the appropriate values.

---

**Note:** The name for the application service group must be the same on both sites.

---

## Part 4: Adding GCO components for wide-area recovery

The Global Cluster Option is required to manage global clustering for wide-area disaster recovery. The process of creating a global cluster environment involves the following tasks:

- Connecting standalone clusters by adding a remote cluster to a local cluster.
- Converting the local service group that is common to all the clusters to a global service group.

Use the VCS Java Console or Cluster Management Console (Single Cluster Mode) also referred to as Web Console, to perform global cluster operations; this guide provides procedures only for the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more information on GCO operations available from the Java and Web Consoles.

Topics in this section include:

- [“Prerequisites for a global cluster environment”](#) on page 561
- [“Linking clusters by adding a remote cluster”](#) on page 562
- [“Converting a local Service group to a global group”](#) on page 564
- [“Additional global cluster administration tasks”](#) on page 566

### Prerequisites for a global cluster environment

Creating a global cluster environment requires the following conditions:

- All service groups are properly configured and able to come online.
- The service group that will serve as the global group has the same unique name across all applicable clusters.
- The clusters must use the same version of VCS.
- The clusters must use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.

## Linking clusters by adding a remote cluster

The VCS Java Console provides a wizard to create global clusters by linking standalone clusters or bringing a standalone cluster into an existing global environment.

- If you are creating a global cluster environment for the first time with two standalone clusters, run the wizard from either the cluster on the primary site or the cluster on the secondary site.
- If you are adding a standalone cluster to an existing global cluster environment, run the wizard from a cluster already in the global cluster environment.

The following information is required for the Remote Cluster Configuration Wizard in Cluster Explorer:

- The active host name or IP address of each cluster in the global configuration and of the cluster being added to the configuration.
- The user name and password of the administrator for each cluster in the configuration.
- The user name and password of the administrator for the cluster being added to the configuration.

---

**Note:** Symantec does not support adding a cluster that is already part of a global cluster environment. To merge the clusters of one global cluster environment (for example, cluster A and cluster B) with the clusters of another global environment (for example, cluster C and cluster D), separate cluster C and cluster D into standalone clusters and add them one by one to the environment containing cluster A and cluster B.

---

### To add a remote cluster in Cluster Explorer

- 1 From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.  
*or*  
From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.
- 2 Review the required information for the **Remote Cluster Configuration Wizard** and click **Next**.
- 3 In the **Wizard Options** dialog box:
  - a Click **Add Cluster**.
  - b Click **Next**.
- 4 Enter the details of the new cluster:

If the cluster is not running in secure mode:

- a Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- b If necessary, change the default port number.
- c Enter the user name.
- d Enter the password.
- e Click **Next**.

If the cluster is running in secure mode:

- a Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- b Verify the port number.
- c Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.  
If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
- d Click **Next**.

5 Click **Finish**. After running the wizard, the configurations on all the relevant clusters are in read-write mode; the wizard does not close the configurations.

6 Verify that the heartbeat connection between clusters is alive. From the command window enter `hahb -display`. The state attribute in the output should show **alive**.

If the state is **unknown**, then offline and online the ClusterService group.

## Converting a local Service group to a global group

After linking the clusters, use the Global Group Configuration wizard to convert a local service group that is common to the global clusters to a global group. This wizard also enables you to convert global groups into local groups.

Administering global groups requires the following conditions:

- A group that will serve as the global group must have the same name across all applicable clusters.
- You must know the user name and password for the administrator to each cluster in the configuration.

Use the VCS Java Console or Cluster Management Console (Single Cluster Mode) also referred to as Web Console, to bring a global group online, take a global group offline, or switch a global group on a remote cluster; the section below provides procedures from the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more information on global cluster operations from the Java Console and Web Consoles.

---

**Note:** For remote cluster operations, the user must have the same name and privilege as the user logged on to the local cluster.

---

### To convert a local service group to a global group

- 1 From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.  
or  
From the Cluster Explorer configuration tree, right-click the cluster, and click **Configure Global Groups**.  
or  
From the Cluster Explorer configuration tree, right-click the service group, click **Configure As Global**, and proceed to step 3b.
- 2 Review the information required for the Global Group Configuration wizard and click **Next**.
- 3 Enter the details of the service group to modify:
  - a Click the name of the service group that will be converted from a local group to a global group, or vice versa.
  - b From the **Available Clusters** box, click the clusters on which the group can come online. Click the right arrow to move the cluster name to the **Clusters for Service Group** box; for global to local cluster conversion, click the left arrow to move the cluster name back to the **Available Clusters** box. A priority number (starting with 0) indicates the cluster

on which the group will attempt to come online. If necessary, double-click the entry in the **Priority** column and enter the new value.

- c Select the policy for cluster failover:
    - **Manual** prevents a group from automatically failing over to another cluster.
    - **Auto** enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster fails.
    - **Connected** enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster.
  - d Click **Next**.
- 4 Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:
- Cluster not in secure mode:
- a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
  - b Verify the port number.
  - c Enter the user name.
  - d Enter the password.
  - e Click **OK**.
  - f Repeat these steps for each cluster in the global environment.
- Cluster in secure mode:
- a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
  - b Verify the port number.
  - c Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and domain.

If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
  - d Click **OK**.
  - e Repeat these steps for each cluster in the global environment.
- 5 Click **Next**.
- 6 Click **Finish**.

At this point, you must bring the global service group online from Cluster Explorer.

#### To bring a remote global service group online from Cluster Explorer

- 1 In the **Service Groups** tab of the configuration tree, right-click the service group.  
*or*  
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Online**, and click **Remote online**.
- 3 In the Online global group dialog box:
  - a Click the remote cluster to bring the group online.
  - b Click the specific system, or click **Any System**, to bring the group online.
  - c Click **OK**.

## Additional global cluster administration tasks

This section provides the following global cluster administration tasks:

- [“Taking a remote global Service group offline”](#) on page 566
- [“Switching a remote global Service group”](#) on page 567

For further information and procedures relating to global clustering, see the “Global Clustering” section in the *Veritas Cluster Server Administrator’s Guide*.

### Taking a remote global Service group offline

#### To take a remote global service group offline

- 1 On the **Service Groups** tab of the configuration tree, right-click the service group.  
*or*  
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Offline**, and click **Remote offline**.
- 3 In the Offline global group dialog box:
  - Click the remote cluster to take the group offline.
  - Click the specific system, or click **All Systems**, to take the group offline and click **OK**.

## Switching a remote global Service group

### To switch a remote global service group

- 1 On the **Service Groups** tab of the configuration tree, right-click the service group.  
*or*  
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 On the Switch global group dialog box:
  - Click the cluster to switch the group.
  - Click the specific system, or click **Any System**, to take the group offline and click **OK**.

## Part 5: Maintaining: Normal Operations and recovery procedures

This section provides tasks during normal operations of the DR solutions 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 the following tools:

- The VEA GUI
- The Command Line Interface (CLI)
- Perfmon
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Storage Foundation 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. The following are a generic set of tasks that you may need to perform:

- Take the RVG resource offline on both the clusters.
- Transfer the primary role to the host at the secondary site by using the **Migrate** option.
  - From the VEA screen, right-click the primary RVG and select **Migrate**.
  - Select the secondary host and click **OK**. The replication role is migrated to the secondary host.
- Assign drive letters to the volumes on the new primary.  
Make sure that these drive letters are the same as those of the original primary.
- Bring the RVG resource online on the new secondary.
- Bring the application group online on the new primary.

You can now verify that the application functions properly on the new primary with the replicated data. After verifying its functioning, you can revert the roles to what they were originally by repeating the procedure.

---

**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.

---

## Disaster recovery procedures

This section provides information on bringing up an application 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 returned to normal functioning after a disaster.

### To bring up the application on the secondary host

- 1 From the left pane in the VEA GUI console on the secondary host, right-click the desired secondary RVG node inside the replication network.
- 2 Select **Takeover** and follow the instructions to perform the takeover operation. You can choose to perform takeover with the following options:
  - Perform **Takeover** *with* the **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 **Takeover** *without* the **fast-failback** option. In this case, you will 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 the takeover, the existing secondary becomes the new primary.

- 3 Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- 4 Bring the application group online.

Now you can start using the application on the new primary.

### Restoring the primary host

After a disaster, when the original primary becomes available again, you may want to revert the role of the primary back to this host.

### To restore the primary host

- 1 Take the RVG resource off-line on both the clusters.

- 2 Depending on whether you performed **Takeover** *with* or *without* the **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 new primary's context menu.
  - For **Takeover** *without* the **Fast-failback** option:

After performing a takeover without fast-failback, you must convert the original primary to a secondary by using the **Make Secondary** option.

Before performing the **Make Secondary** operation, the original primary's RVG and the new primary's RVG will be shown in separate RDS's. 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 the **Synchronize Automatically** option.
- 3 After the synchronization is complete, perform a migrate operation to transfer the primary role back to the original primary. Right-click on the primary RVG and select **Migrate** from the menu that appears.
- 4 Make sure that the volumes have retained the same drive letters as they had before the disaster.
- 5 Bring the RVG resource online on the secondary.
- 6 Bring the application group online on the original primary.

# Index

## A

- automatic volume growth 460
  - definition 481

## C

- campus cluster
  - changing MSCS quorum resource to dynamic quorum resource 410
  - completing setup of application group in MSCS 408
  - configuration 191, 372
  - create SFW cluster disk groups and volumes 393
  - creating VCS service groups 232
  - installing and configuring hardware 381
  - installing application on cluster nodes 229, 406
  - installing Windows and configuring network settings 197
  - making quorum cluster disk group an MSCS resource 411
  - overview 326
  - requirements 187
  - setting up a group for the application 403
  - verifying cluster configuration 233
  - verifying MSCS cluster configuration 413
- campus clustering
  - forceimport attribute of the vmdg resource 195
  - introduction 181
  - MSCS configuration 370, 372
  - MSCS failure scenarios 378
  - prerequisites 370
  - SFW HA configuration 187, 191
  - VCS failure scenarios 192
  - verifying cluster configuration 413
  - Vxclus 379
- cluster configuration
  - steps for a new cluster 174
  - steps for an existing cluster 176

- cluster node
  - installing application 229, 358, 406
- clustering concepts
  - ownership of quorum 379
  - quorum 378
- clusters
  - assigning user privileges 282
  - verifying configuration 151, 542
- creating new volumes 394

## D

- disaster recovery
  - about 238, 240
  - adding GCO components for wide-area recovery 561
  - adding VVR components for replication 439, 546
  - changing MSCS quorum resource to dynamic quorum resource 431
  - completing setup of application group in MSCS 430
  - components of VVR that enable disaster recovery 242
  - configuring replicator log volumes for VVR 440, 547
  - creating VVR RVG service group 557
  - defined 238
  - disk space requirements 249
  - establishing cluster under MSCS (primary site) 423
  - illustrated 241
  - installing and configuring hardware 423
  - installing SFW (primary site) 424
  - installing SFW HA (primary site) 499
  - installing windows and configuring network settings 497
  - normal operations and recovery
    - procedures 298, 302, 453, 568
  - overview 237, 241
  - setting up cluster on primary site 419, 492
  - setting up cluster on secondary site 436, 544

## disaster recovery (continued)

- SFW-MSCS-VVR configuration 415
- solution 240

## Disaster Recovery Procedures 306, 454, 569

## disk groups

- deporting 84
- importing 84
- overview 77

## disk groups and volumes

- configuring 77
- creating SFW cluster disk groups 220, 393
- managing 84

## disk space requirements 249, 332, 371, 420

## DMP ASLs 174, 176

## DMP DSMs 175, 177

## driver signing options

- resetting 77, 204, 263, 506

## dynamic multi-pathing

- about 170
- adding to a clustering configuration 169
- configuration 173
- configuration tasks 171
- existing cluster configuration 176
- more on DMP paths 475
- new cluster configuration 174
- prerequisites 171
- server consolidation 475

## dynamic quorum

- implementation 361

**F**

## FastResync 47

## FastResync (FR) 47

## Fire Drill Wizard

- actions 314
- deleting the configuration 320
- overview 311
- preparing the configuration 315
- prerequisites for a fire drill 313
- restoring the prepared configuration 319
- running a fire drill 318

## FlashSnap 46

- quick recovery 46
- reporting and analysis 51
- tips and references 58

## forceimport attribute of vmdg resource 195

**G**

## Global Cluster Option

- secure configuration 300

## global clusters

- adding 562
- prerequisites 561

**H**

## high availability

- about 61, 325
- local clustering and high availability 61
- overview 326
- solution 61, 325

**I**

## IIS configuration, synchronizing 102

## installing and configuring hardware 196, 381

**L**

## local clustering

- high availability 61
- MSCS 329
- overview 61

**M**

## Microsoft volume shadow copy service (VSS) 47

## modes of replication

- asynchronous 243
- synchronous 242
- synchronous override 243

## MSCS

- campus clustering configuration 329
- changing to a dynamic quorum 410
- completing setup of application group 360
- configuration 334
- configuring the network and storage 335
- creating dynamic cluster disk groups 348
- creating dynamic volumes 350
- disaster recovery procedures 454
- establishing cluster 382
- establishing cluster under MSCS 382
- installing application on cluster nodes 358
- installing SFW 337
- local clustering configuration 329
- prerequisites 332
- setting up a group for application 356, 403

- MSCS (continued)
  - setting up a group for the application 356, 428
  - setup of application group 360, 408
  - verifying cluster configuration 364
- multiple disk group best practices
  - disk group clusters 34
  - disk group structure 33
  - quorum device configuration 328

## N

- network and storage
  - configuration 335
- new volumes
  - creating 394

## O

- off-host backup, defined 50
- online storage migration 460
- options
  - driver signing 77, 204, 263, 506

## P

- performance monitoring
  - data-transfer intensive applications 37
  - request intensive applications 36
- print share groups
  - modifying using wizard 164

## Q

- quick recovery
  - about 44
  - components 46
  - definition 44
  - example 54
  - Oracle database example 54
  - overview 48
  - recover database using split-mirror snapshot
    - and database logs 55
  - solution 44
  - understanding components of quick
    - recovery 46

## R

- RAID best practices
  - hardware RAID 39
  - mirroring 31, 32

- RAID best practices (continued)
  - RAID-5 35
  - read performance and failure tolerance 34
  - striping across hardware 35
  - striping and mirroring 32
- RAID Configurations Using Logical Volume
  - Management 31, 34
- replicated data set (RDS) 546
- replicated volume group (RVG) 546, 557
- replication
  - adding VVR components 439, 546
  - asynchronous 243
  - general definition 242
  - replicator log 546
  - RVG snapshot 244
  - synchronous 242
  - synchronous override 243
  - write order fidelity 243
- requirements
  - disk space 332, 371, 420
- resetting
  - driver signing options 77, 204, 263, 506

## S

- SAN
  - setting up storage 465
  - SFW features 466, 482
- secure clusters
  - assigning user privileges 282
- secure GCO, establishing 300
- Security Services
  - configuring 92, 210, 269, 512
- server consolidation 457
  - about 459
  - adding storage array 472
  - advantages 460
  - advantages of using SFW 460
  - configuration 464, 468, 476
  - configuration 1 – many to one 465
  - configuration 2 – many to two
    - adding clustering and DMP 473
  - customer success story 482
  - definition 459
  - general configuration 464
  - many-to-one configuration 465
  - many-to-two configuration with MSCS and
    - DMP 473
  - migrating the data to the large server 471
  - overview 459, 462

- server consolidation 457 (continued)
  - performance monitoring 461, 481, 482
  - preparing to consolidate 470
  - process to implement 462
  - SFW features 481
- service groups
  - administering global groups 564
  - creating 532
  - dependencies 253, 307
- setting bandwidth
  - using RDS wizard 295
- setting up cluster
  - primary site 249, 419, 492
  - secondary site 436, 544
- SFW
  - best practices 31
  - high availability configuration 61
  - implementing dynamic MSCS quorum resource 361
  - installing 337
  - SFW-specific solutions 17, 179
  - typical high availability configuration 61
- SFW HA
  - creating cluster disk groups and volumes 220
  - creating dynamic volumes 81, 527
  - disaster recovery procedures 306, 569
  - installing 198, 499
  - installing and configuring hardware 196, 253, 497
  - installing application (primary site) 530
  - service group example 533
  - verifying cluster configuration 233
- snapshot
  - commands 46
  - other applications for point-in-time snapshots 50
- Solutions Configuration Center
  - context sensitivity 24
  - overview 23
  - running wizards remotely 28
  - starting 24
  - wizard descriptions 28
- SQL Server Virtual Device Interface (VDI)
  - quiescing the database 44
- storage capacity best practices
  - allocation planning 38
  - failure-tolerant volume recovery 39
  - location of data objects 32
  - manage unallocated space 39

**T**

- tips
  - FlashSnap 58

**U**

- user privileges
  - assigning 282

**V**

- VCS
  - campus clustering configuration 185
  - VVR configuration 245
- VCS service groups
  - creating 232
- verifying
  - cluster configuration for HA 151, 542
- volumes
  - creating 397
  - creating on primary 223, 397
  - mounting 84
  - overview 77
  - unmounting 84
- VSS 47
- VVR
  - components that enable disaster recovery 242
  - definition 242
  - setting up Replicated Data Sets (RDS) 549
  - SFW HA-VCS configuration 245, 487
- Vxclus utility 379
- Vxsnap command 47

**W**

- wide-area recovery
  - adding GCO components 561
- Windows
  - network settings 197, 497