

Veritas Storage Foundation™ and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL

Windows 2000, Windows Server 2003

5.0

Veritas SFW HA HA and Disaster Recovery Solutions Guide for Microsoft SQL

Copyright © 2007 Symantec Corporation. All rights reserved.

Storage Foundation 5.0 for Windows HA

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

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

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

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

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

Third-party legal notices

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

Windows is a registered trademark of Microsoft Corporation.

Licensing and registration

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

Technical support

For technical assistance, visit <http://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 for Microsoft SQL

About the solutions guides	21
About high availability	22
About disaster recovery	22
How this guide is organized	22

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

Section 2 High Availability

Chapter 3 High availability for SQL: Overview

What is high availability?	33
Why implement a high availability solution?	34
How the agent makes SQL Server highly available	34

Chapter 4 Deploying SFW HA for high availability: New SQL Server 2000 installation

Tasks for a new HA installation of SQL Server 2000	36
Tasks for an active/passive configuration	37
Tasks for an active/active configuration	39
Reviewing the requirements	43
Disk space requirements	43
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	43

Reviewing the configuration	47
Active/passive configuration	47
Active/active configuration	48
Disaster recovery configuration	51
Reviewing considerations for active/active configurations	54
Key information for active/active configurations	54
Following the workflow in the Solutions Configuration Center	55
Configuring the storage hardware and network	57
Installing Veritas Storage Foundation HA for Windows	59
Setting Windows driver signing options	59
Installing Storage Foundation HA for Windows	60
Resetting the driver signing options	65
Configuring cluster disk groups and volumes for SQL Server 2000	65
About cluster disk groups and volumes	65
Prerequisites for configuring cluster disk groups and volumes	66
Sample disk group and volume configuration	66
Considerations for disk groups and volumes for multiple instances	67
Considerations for volumes for a disaster recovery configuration	68
Creating a cluster disk group	68
Creating volumes	69
Configuring the cluster	74
Configuring Web console	85
Configuring notification	86
About installing multiple instances	88
Installing and configuring SQL Server 2000 on the first node	89
Installing Microsoft SQL Server 2000	89
Setting SQL Server 2000 services to manual start	92
Preparing to install SQL Server 2000 on the second node	92
Stopping the SQL Server 2000 service	93
Deporting the cluster disk group	93
Importing the cluster disk group	94
Adding drive letters to mount the volumes	94
Renaming shared SQL Server 2000 files	96
Installing and configuring SQL Server 2000 on the second node	96
Installing SQL Server 2000 on the second node	96
Removing shared SQL Server 2000 files	99
Setting the internal name of the clustered instance	100
Creating a SQL Server user-defined database	102
Completing configuration steps in SQL Server	102
Moving the tempdb database if using VVR for disaster recovery	103
Assigning the port for multiple instances	103

Configuring the VCS SQL Server 2000 service group	103
Service group requirements for active/active configurations	104
Prerequisites for configuring the service group	104
Creating the SQL Server 2000 service group	105
Verifying the SQL Server 2000 cluster configuration	110
Determining additional steps needed	111
Configuring the Cluster Management Console connection	111
Modifying the SQL 2000 service group to add VMDg and MountV resources	116

Chapter 5 Deploying SFW HA for high availability: Standalone SQL Servers

Tasks for converting a standalone SQL 2000 Server for high availability	119
Reviewing the prerequisites	122
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	122
Reviewing the configuration	126
Sample configuration	127
Configuring the storage hardware and network	128
Preparing the standalone SQL Server	130
Backing up existing SQL data	130
Setting SQL Server services to manual start	130
Installing and configuring SFW HA	131
Installing and configuring SQL Server 2000 on additional nodes	131
Moving the existing SQL Server 2000 data files and user databases	132
Configuring the VCS SQL Server service group	133
Creating a SQL Server user-defined database	133
Verifying the installation	133
Additional instructions for disaster recovery	134

Chapter 6 Deploying SFW HA for high availability: New SQL Server 2005 installation

Tasks for a new HA installation of SQL Server 2005	136
Tasks for an active/passive configuration	137
Tasks for an active/active configuration	139
Reviewing the requirements	142
Disk space requirements	142
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	143

Reviewing the configuration	146
Active/passive configuration	146
Active/active configuration	148
Disaster recovery configuration	151
Reviewing considerations for active/active configurations	154
Key information for active/active configurations	154
Following the workflow in the Solutions Configuration Center	155
Configuring the storage hardware and network	156
Installing Veritas Storage Foundation for Windows	158
Setting Windows driver signing options	158
Installing Storage Foundation HA for Windows	159
Resetting the driver signing options	163
Configuring cluster disk groups and volumes for SQL Server 2005	164
About cluster disk groups and volumes	164
Prerequisites for configuring cluster disk groups and volumes	164
Sample disk group and volume configuration	165
Considerations for disk groups and volumes for multiple instances	165
Considerations for volumes for a disaster recovery configuration ...	166
Creating a cluster disk group	167
Creating volumes	168
Configuring the cluster	173
Configuring Web console	184
Configuring notification	185
About installing multiple instances	187
Installing and configuring SQL Server 2005 on the first node	188
Installing Microsoft SQL Server 2005 on the first node	188
Setting SQL Server 2005 services to manual start	191
Preparing to install SQL Server 2005 on the second node	193
Stopping the SQL Server 2005 service	193
Deporting the cluster disk group	194
Importing the cluster disk group	194
Adding drive letters to mount the volumes	195
Renaming shared SQL Server 2005 files	196
Installing and configuring SQL Server 2005 on the second node	196
Installing Microsoft SQL Server on the second node	196
Removing shared SQL Server files	199
Setting the internal name of the clustered instance	200
Creating a SQL Server user-defined database	203
Completing configuration steps in SQL Server	204
Moving the tempdb database if using VVR for disaster recovery	204
Assigning the port for multiple instances	204

	Configuring the VCS SQL Server 2005 service group	205
	Service group requirements for active/active configurations	205
	Prerequisites for configuring the service group	206
	Creating the SQL Server 2005 service group	207
	Verifying the SQL Server 2005 cluster configuration	211
	Determining additional steps needed	212
	Configuring the Cluster Management Console connection	212
	Modifying a SQL 2005 service group to add VMDg and MountV resources	217
Chapter 7	Configuring an MSDTC service group for high availability	
	Tasks for configuring MSDTC for high availability	219
	Reviewing the prerequisites	221
	Reviewing the configuration	221
	Configuring cluster disk groups and volumes	224
	Creating a cluster disk group	225
	Creating volumes	226
	Mounting volumes used by the MSDTC service group	230
	Creating an MSDTC service group	231
	Creating an MSDTC client	233
Section 3	Replicated Data Clusters	
Chapter 8	About Replicated Data Clusters	
	About Replicated Data Clusters	237
	How VCS Replicated Data Clusters work	239
	Setting up a Replicated Data Cluster configuration	240
	Setting up replication	240
	Configuring the service groups	241
	Migrating the service group	242
Chapter 9	Configuring Replicated Data Clusters for SQL 2000	
	Tasks for configuring Replicated Data Clusters for SQL 2000	246
	Reviewing the prerequisites	250
	Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	250
	Reviewing the configuration	255
	Sample configuration	255

Configuring the storage hardware and network	256
Installing Veritas Storage Foundation for Windows	258
Setting Windows driver signing options	258
Installing Storage Foundation HA for Windows	258
Configuring VxSAS	263
Resetting the driver signing options	265
Configuring the cluster	265
Configuring Web console	276
Configuring notification	277
Configuring cluster disk groups and volumes	280
Creating a cluster disk group	281
Creating volumes	282
Installing and configuring SQL Server 2000 on the first node	288
Installing Microsoft SQL Server	288
Setting SQL Server 2000 services to manual start	291
Preparing to install SQL Server on the second node	291
Stopping the SQL Server 2000 service	292
Deporting the cluster disk group	292
Importing the cluster disk group	293
Adding drive letters to mount the volumes	293
Renaming shared SQL Server 2000 files	295
Installing SQL Server 2000 on the second node	295
Installing SQL Server	295
Removing shared SQL Server files	298
Setting the internal name of the clustered instance	299
Configuring the VCS SQL Server service group	301
Creating the primary system zone	306
Registering the virtual server in the SQL Server	307
Creating a SQL Server user-defined database	309
Creating new volumes	309
Creating a new SQL Server database	309
Adding VMDg and MountV resources	310
Verifying the installation in the primary zone	312
Creating a parallel environment in the secondary zone	313
Adding the IP addresses for the primary and secondary zones	314
Adding the systems in the secondary zone to the cluster	314
Modifying values for ClusterService group attributes	320
Setting up the Replicated Data Sets (RDS)	321
Configuring a hybrid RVG service group for replication	329
Creating the RVG service group	329
Configuring the RVG service group for RDC replication	330
Configuring the RVG Primary resources	340
Configuring the primary system zone for the RVG	343

Setting a dependency between the service groups	343
Adding the nodes from the secondary zone to the RDC	344
Adding the nodes from the secondary zone to the RVG service group	344
Configuring secondary zone nodes in the RVG service group	347
Configuring the IP resources for failover	347
Adding the nodes from the secondary zone to the SQL Server service group	350
Configuring the zones in the SQL Server service group	351
Verifying the RDC configuration	351
Switching online nodes	351

Chapter 10 Configuring Replicated Data Clusters for SQL 2005

Tasks for configuring Replicated Data Clusters for SQL 2005	354
Reviewing the prerequisites	358
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	358
Reviewing the configuration	363
Sample configuration	363
Configuring the storage hardware and network	364
Installing Veritas Storage Foundation for Windows	366
Setting Windows driver signing options	366
Installing Storage Foundation HA for Windows	366
Configuring VxSAS	371
Resetting the driver signing options	373
Configuring the cluster	373
Configuring Web console	384
Configuring notification	385
Configuring cluster disk groups and volumes	388
Creating a cluster disk group	389
Creating volumes	390
Installing and configuring SQL Server 2005 on the first node	396
Installing Microsoft SQL Server	396
Setting SQL Server 2005 services to manual start	399
Preparing to install SQL Server 2005 on the second node	400
Stopping the SQL Server 2005 Service	401
Deporting the cluster disk group	401
Importing the cluster disk group	401
Adding drive letters to mount the volumes	402
Renaming shared SQL Server 2005 files	403
Installing SQL Server 2005 on the second node	403
Installing SQL Server	403
Removing shared SQL Server files	406

Setting the internal name of the clustered instance	407
Configuring the VCS SQL Server service group	409
Creating the primary system zone	414
Registering the virtual server in the SQL Server	415
Creating a SQL Server user-defined database	417
Creating new volumes	417
Creating a new SQL Server database	417
Adding VMDg and MountV resources	418
Verifying the installation in the primary zone	420
Creating a parallel environment in the secondary zone	421
Adding the IP addresses for the primary and secondary zones	422
Adding the systems in the secondary zone to the cluster	422
Modifying values for ClusterService group attributes	428
Setting up the Replicated Data Sets (RDS)	429
Configuring a hybrid RVG service group for replication	437
Creating the RVG service group	438
Configuring the RVG service group for RDC replication	438
Configuring the RVG Primary resources	450
Configuring the primary system zone for the RVG	453
Setting a dependency between the service groups	454
Adding the nodes from the secondary zone to the RDC	454
Adding the nodes from the secondary zone to the RVG service group	455
Configuring secondary zone nodes in the RVG service group	457
Configuring the IP resources for failover	457
Adding the nodes from the secondary zone to the SQL Server service group	460
Configuring the zones in the SQL Server service group	461
Verifying the RDC configuration	462
Switching online nodes	462

Section 4 Disaster Recovery

Chapter 11 Disaster recovery for SQL: Overview

What is a disaster recovery solution?	465
Why implement a disaster recovery solution	466
Understanding replication	466
What needs to be protected in a SQL Server environment?	466
Typical disaster recovery configuration	467

Chapter 12	Deploying disaster recovery: New SQL Server 2000 installation	
	Tasks for a new disaster recovery installation of Microsoft SQL	470
	Before you begin	473
	Disk space requirements	473
	Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	474
	Reviewing the configuration	477
	Supported disaster recovery configurations for service group dependencies	479
	Configuring the storage hardware and network	480
	Setting up the secondary site: Configuring SFW HA and setting up a cluster	482
	Installing SFW HA	482
	Setting Windows driver signing options	482
	Installing Storage Foundation HA for Windows	483
	Resetting the driver signing options	488
	Configuring the cluster	488
	Configuring Web console	499
	Configuring notification	500
	Verifying your primary site configuration	503
	Setting up security for VVR	503
	Configuring disaster recovery	506
	Assigning user privileges (secure clusters only)	507
	Cloning the storage on the secondary site using the DR wizard	508
	Installing and configuring SQL Server 2000 on the first node (secondary site)	512
	Installing Microsoft SQL Server	512
	Setting SQL Server 2000 services to manual start	515
	Preparing to install SQL Server on the second node (secondary site)	515
	Stopping the SQL Server 2000 Service	516
	Deporting the cluster disk group	516
	Importing the cluster disk group	517
	Adding drive letters to mount the volumes	517
	Renaming shared SQL Server 2000 files	519
	Installing SQL Server 2000 on the second node (secondary site)	519
	Installing SQL Server	519
	Removing shared SQL Server files	522
	Setting the internal name of the clustered instance	523
	Cloning the service group configuration from the primary to the secondary site	525
	Configuring replication and global clustering	527
	Verifying the disaster recovery configuration	534

Establishing secure communication within the global cluster (optional)	535
Recovery procedures for service group dependencies	537

Chapter 13 Deploying disaster recovery: New SQL Server 2005 installation

Tasks for a new disaster recovery installation of Microsoft SQL	542
Before you begin	546
Disk space requirements	546
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	546
Reviewing the configuration	550
Supported disaster recovery configurations for service group dependencies	553
Configuring the storage hardware and network	553
Setting up the secondary site: Configuring SFW HA and setting up a cluster	555
Installing SFW HA	555
Setting Windows driver signing options	555
Installing Storage Foundation HA for Windows	556
Resetting the driver signing options	561
Configuring the cluster	561
Configuring Web console	572
Configuring notification	573
Verifying your primary site configuration	576
Setting up security for VVR	576
Configuring disaster recovery	579
Assigning user privileges (secure clusters only)	580
Cloning the storage on the secondary site using the DR wizard	580
Installing and configuring SQL Server 2005 on the first node (secondary site)	585
Installing SQL Server 2005 on the first node	585
Setting SQL Server 2005 services to manual start	587
Preparing to install SQL Server 2005 on the second node (secondary site)	589
Stopping the SQL Server 2005 Service	589
Deporting the cluster disk group	590
Importing the cluster disk group	590
Adding drive letters to mount the volumes	590
Renaming shared SQL Server 2005 files	591
Installing SQL Server 2005 on the second node (secondary site)	592
Installing SQL Server on the second node	592
Removing shared SQL Server files	595

	Setting the internal name of the clustered instance	596
	Cloning the service group configuration from the primary to the secondary site	598
	Configuring replication and global clustering	601
	Verifying the disaster recovery configuration	607
	Establishing secure communication within the global cluster (optional)	608
	Recovery procedures for service group dependencies	610
Chapter 14	Testing fault readiness by running a fire drill	
	About disaster recovery fire drills	615
	About the Fire Drill Wizard	615
	Tasks for configuring and running fire drills	617
	Prerequisites for a fire drill	617
	Fire Drill Wizard actions	618
	Preparing the fire drill configuration	619
	Running a fire drill	622
	Restoring the fire drill system to a prepared state	623
	Deleting the fire drill configuration	624
Section 5	Appendices	
Appendix A	Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation	
	Reviewing the prerequisites	633
	Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	633
	Reviewing the configuration	637
	Sample configuration	639
	Configuring the storage hardware and network	640
	Installing Veritas Storage Foundation for Windows	642
	Setting Windows driver signing options	642
	Installing Storage Foundation HA for Windows	643
	Configuring VxSAS	647
	Resetting the driver signing options	649
	Configuring the cluster	649
	Configuring cluster disk groups and volumes	665
	Creating a cluster disk group	666
	Creating volumes	668
	Installing and configuring SQL Server 2000 on the first node	673
	Installing Microsoft SQL Server	673
	Setting SQL Server 2000 services to manual start	676

Preparing to install SQL Server on the second node	676
Stopping the SQL Server 2000 Service	677
Deporting the cluster disk group	677
Importing the cluster disk group	678
Adding drive letters to mount the volumes	678
Renaming shared SQL Server 2000 files	680
Installing SQL Server 2000 on the second node	680
Installing SQL Server	680
Removing shared SQL Server files	683
Setting the internal name of the clustered instance	684
Configuring the VCS SQL Server service group	686
Creating a SQL Server user-defined database	690
Creating new volumes	691
Creating a new SQL Server database	691
Adding VMDg and MountV resources	691
Verifying the cluster configuration	693
Creating a parallel environment on the secondary site	694
Installing DR components on primary and secondary sites	695

Appendix B Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation

Reviewing the prerequisites	701
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	701
Reviewing the configuration	705
Sample configuration	707
Configuring the storage hardware and network	708
Installing Veritas Storage Foundation for Windows	710
Setting Windows driver signing options	710
Installing Storage Foundation HA for Windows	711
Configuring VxSAS	715
Resetting the driver signing options	717
Configuring the cluster	717
Configuring cluster disk groups and volumes	732
Creating a cluster disk group	734
Creating volumes	736
Installing and configuring SQL Server 2005 on the first node	741
Installing SQL Server 2005 on the first node	741
Setting SQL Server 2005 services to manual start	743

Preparing to install SQL Server 2005 on the second node	745
Stopping the SQL Server 2005 Service	745
Deporting the cluster disk group	746
Importing the cluster disk group	746
Adding drive letters to mount the volumes	746
Renaming shared SQL Server 2005 files	747
Installing SQL Server 2005 on the second node	748
Installing SQL Server on the second node	748
Removing shared SQL Server files	751
Setting the internal name of the clustered instance	752
Configuring the VCS SQL Server service group	754
Creating a SQL Server user-defined database	759
Creating new volumes	759
Creating a new SQL Server database	760
Adding VMDg and MountV resources	760
Verifying the cluster configuration	762
Creating a parallel environment on the secondary site	763
Installing DR components on the primary and secondary sites	764

Appendix C Configuring the DR components (VVR and GCO) without using the DR wizard

Reviewing the prerequisites	767
Setting up the Replicated Data Sets (RDS)	767
Creating the VVR RVG service group	776
Configuring the global cluster option for wide-area failover	780
Prerequisites	780
Linking clusters: Adding a remote cluster to a local cluster	781
Converting a local service group to a global service group	782
Bringing a global service group online	784
Establishing secure communication within the global cluster (optional)	785

Appendix D Configuring an MSDTC service group for disaster recovery

Tasks for configuring MSDTC for disaster recovery	787
Reviewing the prerequisites	790
Reviewing the configuration	790
Configuring cluster disk groups and volumes	794
Creating a cluster disk group	795
Creating volumes	795
Mounting drives used by the MSDTC service group	798
Creating an MSDTC service group	799

Creating an MSDTC client 800

Setting up the secondary site: Creating a parallel environment 802

Installing DR components on the primary and secondary sites 802

Index 803

Introduction

This section contains the following chapters:

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

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL

This chapter includes the following topics:

- [About the solutions guides](#)
- [About high availability](#)
- [About disaster recovery](#)
- [How this guide is organized](#)

About the solutions guides

The *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL* contains solutions for the following:

- High availability (HA)
- Disaster recovery (DR)

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

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

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.

Information about high availability for SQL Server includes procedures for installing and configuring clustered Microsoft SQL Server environments using SFW HA.

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.

Information about the disaster recovery solution for SQL Server includes procedures for installing, configuring, and testing clustered and replicated Microsoft SQL Server environments for disaster recovery using SFW HA.

How this guide is organized

The *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL* 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



High Availability

This section includes the following chapters:

- [High availability for SQL: Overview](#)
- [Deploying SFW HA for high availability: New SQL Server 2000 installation](#)
- [Deploying SFW HA for high availability: New SQL Server 2005 installation](#)
- [Deploying SFW HA for high availability: Standalone SQL Servers](#)
- [Configuring an MSDTC service group for high availability](#)

High availability for SQL: Overview

This chapter includes the following topics:

- [What is high availability?](#)
- [Why implement a high availability solution?](#)
- [How the agent makes SQL Server highly available](#)

What is high availability?

The term *High Availability* 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. This section focuses on configurations that use Veritas Storage Foundation HA for Windows. Storage Foundation HA for Windows (SFW HA) includes Veritas Storage Foundation and Veritas Cluster Server.

Local clustering provides high availability (HA) through database and application failover. This solution provides local recovery in the event of application, operating system, or hardware failure, and minimizes planned and unplanned application downtime. The High Availability section includes procedures for installing and configuring clustered Microsoft SQL Server environments using Veritas Storage Foundation HA for Windows.

Setting up the clustered environment is also the first step in creating a disaster recovery solution. Some installation and configuration options in this section are identified as required for disaster recovery only. These options apply only if you intend to set up a secondary site for wide-area disaster recovery.

A cluster is a group of independent computers working together to ensure that mission-critical applications and resources are as highly available as possible. The group 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.

Why implement a high availability solution?

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

Using Veritas Storage Foundation HA for Windows as a local high availability solution paves the way for a wide-area disaster recovery solution in the future. A high availability solution is built on top of a backup strategy and provides the following benefits:

- Reduces planned and unplanned downtime.
- Serves as a local and wide-area failover (rather than load-balancing) solution. Enables failover between sites or between clusters.
- Manages applications and provides an orderly way to bring processes online and take them offline.
- Consolidates hardware in larger clusters. The HA environment accommodates flexible fail over policies, active-active configurations, and shared standby servers for SQL Server.

How the agent makes SQL Server highly available

The VCS enterprise agent for SQL Server detects an application failure if a configured virtual server becomes unavailable. When this occurs, the SQL Server service group is failed over to the next available system in the service group's system list. The configured SQL services and virtual server are started on the new system; ensuring continuous availability to SQL data.

Deploying SFW HA for high availability: New SQL Server 2000 installation

This chapter covers the following topics:

- [Tasks for a new HA installation of SQL Server 2000](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Reviewing considerations for active/active configurations](#)
- [Configuring the storage hardware and network](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Configuring cluster disk groups and volumes for SQL Server 2000](#)
- [Configuring the cluster](#)
- [About installing multiple instances](#)
- [Installing and configuring SQL Server 2000 on the first node](#)
- [Preparing to install SQL Server 2000 on the second node](#)
- [Installing and configuring SQL Server 2000 on the second node](#)
- [Creating a SQL Server user-defined database](#)
- [Completing configuration steps in SQL Server](#)
- [Configuring the VCS SQL Server 2000 service group](#)

- [Verifying the SQL Server 2000 cluster configuration](#)
- [Determining additional steps needed](#)
- [Modifying the SQL 2000 service group to add VMDg and MountV resources](#)

Tasks for a new HA installation of SQL Server 2000

You can install and configure a new Veritas Storage Foundation HA environment for SQL Server 2000 in the following ways:

active/passive	One SQL 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 SQL instances per cluster node. For example, in a two-node cluster with two SQL 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.

Note: Some installation and configuration options in this section are identified as required “for a disaster recovery configuration.” These options apply only if you intend to set up a secondary site for disaster recovery.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA for SQL Server 2000.

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

To configure MSDTC service groups, see “[Configuring an MSDTC service group for high availability](#)” on page 219.

If you have an existing SQL Server 2000 deployment on a standalone server, refer to “[Deploying SFW HA for high availability: Standalone SQL Servers](#)” on page 119.

Tasks for an active/passive configuration

Table 4-1 outlines the high-level objectives and the tasks to complete each objective for an active/passive configuration.

Table 4-1 Tasks for installing and configuring HA for SQL Server 2000

Objective	Tasks
“Reviewing the requirements” on page 43	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 47	✓ Understanding active/passive configuration ✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 57	✓ Setting up the storage hardware for a cluster environment ✓ Verifying the DNS entries for the systems on which SQL will be installed
“Installing Veritas Storage Foundation HA for Windows” on page 59	✓ Verifying the driver signing option for the system ✓ Installing Veritas Storage Foundation HA for Windows (automatic installation) ✓ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“Configuring cluster disk groups and volumes for SQL Server 2000” on page 65	✓ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ✓ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator

Table 4-1 Tasks for installing and configuring HA for SQL Server 2000

Objective	Tasks
“Configuring the cluster” on page 74	✓ Verifying static IP addresses and name resolution configured for each node ✓ Running the Veritas Cluster Server Configuration Wizard to configure cluster components and set up secure communication for the cluster
“Installing and configuring SQL Server 2000 on the first node” on page 89	✓ Installing and configuring SQL Server 2000 ✓ Setting SQL Server services to manual start
“Preparing to install SQL Server 2000 on the second node” on page 92	✓ Stopping the SQL service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Mounting the volumes (adding drive letters) ✓ Removing shared SQL files
“Installing and configuring SQL Server 2000 on the second node” on page 96	✓ Installing and configuring SQL Server 2000
“Setting the internal name of the clustered instance” on page 100	✓ Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 102	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server

Table 4-1 Tasks for installing and configuring HA for SQL Server 2000

Objective	Tasks
“Completing configuration steps in SQL Server” on page 102	✓ Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2000 service group” on page 103	✓ Creating a SQL Server service group using the SQL Server Configuration Wizard
“Verifying the SQL Server 2000 cluster configuration” on page 110	✓ Simulating failover ✓ Switching online nodes

Tasks for an active/active configuration

Table 4-2 outlines the high-level objectives and the tasks to complete each objective for an active/active configuration.

Table 4-2 Tasks for installing and configuring active/active HA for SQL Server 2000

Objective	Tasks
“Reviewing the requirements” on page 43	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 47	✓ Understanding active/active configuration ✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 57	✓ Setting up the storage hardware for a cluster environment ✓ Verifying the DNS entries for the systems on which SQL will be installed

Table 4-2

Tasks for installing and configuring active/active HA for SQL Server 2000 (continued)

Objective	Tasks
“Installing Veritas Storage Foundation HA for Windows” on page 59	✓ Verifying the driver signing option for the system ✓ Installing Veritas Storage Foundation HA for Windows (automatic installation) ✓ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“Configuring cluster disk groups and volumes for SQL Server 2000” on page 65	✓ Creating a dynamic disk group for each instance ✓ Creating dynamic volumes for the SQL system database, user databases, transaction logs, and replicated registry keys for each instance
“Configuring the cluster” on page 74	✓ Verifying static IP addresses and name resolution configured for each node ✓ Running the Veritas Cluster Server Configuration Wizard to configure cluster components and set up secure communication for the cluster
“Installing and configuring SQL Server 2000 on the first node” on page 89	✓ Installing and configuring SQL Server 2000 ✓ Setting SQL Server services to manual start

Table 4-2 Tasks for installing and configuring active/active HA for SQL Server 2000 (continued)

Objective	Tasks
“Preparing to install SQL Server 2000 on the second node” on page 92	✓ Stopping the SQL service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Mounting the volumes (adding drive letters) ✓ Renaming or removing shared SQL files
“Installing and configuring SQL Server 2000 on the second node” on page 96	✓ Installing SQL Server on the second node
“Setting the internal name of the clustered instance” on page 100	✓ Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 102	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server
“Completing configuration steps in SQL Server” on page 102	✓ Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2000 service group” on page 103	✓ Creating a SQL Server service group using the SQL Server Configuration Wizard ✓ Ensuring the priority order of the systems is set up in reverse order for each instance

Table 4-2 Tasks for installing and configuring active/active HA for SQL Server 2000 (continued)

Objective	Tasks
“Verifying the SQL Server 2000 cluster configuration” on page 110	✓ Simulating fail over ✓ Switching online nodes
“Determining additional steps needed” on page 111	✓ Repeating the installation and configuration steps for the next instance

Reviewing the requirements

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

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 4-3 estimates disk space requirements for SFW HA.

Table 4-3 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 Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the following options:
Veritas Volume Replicator and the Global Clustering Option
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with

Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 required for all editions)

or

Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition

or

Windows 2000 Server, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server (all require SP4 with Update Rollup 1)

or

- Microsoft SQL Server 2000 (64-bit) Enterprise Edition

with

Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required for both)

or

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

with

Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): 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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 46.
- 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 SQL 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 *VCS 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.
- 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

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 doing the following tasks:

- 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

Review the information for the configurations you have planned:


- [Active/passive configuration](#)
- [Active/active configuration](#)
- [Disaster recovery configuration](#)

Active/passive configuration

In a typical example of a high availability cluster, you create a virtual server in an active/passive SQL configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node.

[Figure 4-1](#) illustrates a typical active/passive configuration. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured on the active node (SYSTEM1). If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 4-1 Active/passive configuration



Sample active/passive configuration

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

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

Name	Object
SYSTEM1 & SYSTEM2	servers
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	virtual SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server
INST1_SG	SQL Server service group

IP addresses for sample active/passive configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there is one SQL Server virtual server. Therefore you would need one virtual server IP address. If you want to use the VCS Web Console or the notification service, you require a cluster IP address (the cluster IP address is also used by the Global Cluster Option for disaster recovery).

See “Network requirements” in “[Reviewing the requirements](#)” on page 43.
See “[Disaster recovery configuration](#)” on page 51.

Active/active configuration


In an active/active SQL Server configuration, several instances are intended to run on a single node when necessary. A SQL Server instance is a completely independent SQL Server installation, with its own services, master database, storage, and memory resources. Each instance is defined uniquely by a separate SQL Server virtual server and service group.

SQL Server 2000 supports up to 16 independent instances of SQL Server 2000 to run on a single machine. A SQL Server 2000 instance can fail over to any of the

other nodes on its system list. You can choose an active/active SQL Server 2000 configuration where several instances are intended to run on a single node. However, remember that you must configure failover nodes so that a single node can never host more than 16 instances.

The following figure illustrates a two node active/active configuration. The SQL Server databases are configured on the shared storage on volumes contained in cluster disk groups. Each SQL Server virtual server is configured in a separate SQL Server service group. Each service group can fail over to the other node in the cluster.

Figure 4-2 Active/active configuration



For example, consider a two-node cluster hosting two SQL Server virtual servers, BILLING_VS and PAYROLL_VS. The table below and the sample configuration illustrate that the virtual servers are configured in two separate service groups with BILLING_VS online on SYSTEM1 but able to fail over to SYSTEM2, and PAYROLL_VS online on SYSTEM2 but able to fail over to SYSTEM1.

SQL Virtual Server	Service Group	System List
BILLING_VS	BILLING_SG	SYSTEM1, SYSTEM2
PAYROLL_VS	PAYROLL_SG	SYSTEM2, SYSTEM1

Sample active/active configuration

A sample setup is used to illustrate the installation and configuration tasks for two instances of SQL Server, Billing and Payroll. During normal operation, one instance will be online on each of the two servers. If a failure occurs, the

instance on the failing node will be brought online on the other server, resulting in two instances running on one server.

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

Name	Description
SYSTEM1 & SYSTEM2	server names
BILLING_DG	cluster disk group for the billing instance
PAYROLL_DG	cluster disk group for the payroll instance
BILLING_VS_SYS_FILES	volume for the SQL Server system data files for the billing instance
PAYROLL_VS_SYS_FILES	volume for the SQL Server system data files for the payroll instance
BILLING_DATA	volume for a SQL Server user-defined database for the billing instance
PAYROLL_DATA	volume for a SQL Server user-defined database for the payroll instance
BILLING_LOG	volume for a SQL Server user-defined database log file for the billing instance
PAYROLL_LOG	volume for a SQL Server user-defined database log file for the payroll instance
BILLING_REGREP	volume for the list of registry keys replicated among the nodes for the billing instance
PAYROLL_REGREP	volume for the list of registry keys replicated among the nodes for the payroll instance
SQL_CLUS1	virtual SQL Server cluster
BILLING_INST	instance name for the billing instance
PAYROLL_INST	instance name for the payroll instance
BILLING_VS	virtual SQL Server name for the billing instance
PAYROLL_VS	virtual SQL Server name for the payroll instance
BILLING_SG	SQL Server service group for the billing instance
PAYROLL_SG	SQL Server service group for the payroll instance

IP addresses for sample active/active configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each

SQL Server virtual server requires its own virtual IP address. In the sample configuration there are two virtual servers: BILLING-VS and PAYROLL-VS. Therefore, you would need two virtual server IP addresses. If you want to use the VCS Web Console or the notification service, you require a cluster IP address (the cluster IP address is also used by the Global Cluster Option for disaster recovery).

See “Network requirements” under “[Reviewing the requirements](#)” on page 43.

See “[Disaster recovery configuration](#)” on page 51.

Disaster recovery configuration


You may be preparing to configure both a primary site and a secondary site for disaster recovery.

[Figure 4-3](#) illustrates a typical active/passive disaster recovery configuration using Veritas Volume Replicator (VVR).

In the example illustration, the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the Replicated Volume Group (RVG). The SQL Server application data is stored on the volumes that are under the control of the RVG.

If the Microsoft SQL Server server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Figure 4-3 Typical disaster recovery configuration with VVR



Sample disaster recovery configuration

The sample disaster recovery setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

A sample setup is used to illustrate the installation and configuration tasks:

Primary Site

SYSTEM1 & servers
SYSTEM2

INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_REPLOG	replicator log volume required by VVR
INST1_DB1_RDS	RDS (Replicated Data Set) name
INST1_DB1_RVG	RVG (Replicated Volume Group) name
INST1_DB1_RVG_SG	Replication service group

IP addresses for sample disaster recovery configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration.

See “Network requirements” under “[Reviewing the requirements](#)” on page 43.

You specify the following addresses during the replication process:

SQL virtual server IP address	For a disaster recovery configuration, the virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

Cluster IP address	You need one for the primary site cluster and one for the secondary site cluster.
Replication IP address	You need two IP addresses per application instance, one for the primary site and one for the secondary site.

Reviewing considerations for active/active configurations

For an active/active configuration or in other cases when you are installing multiple instances of SQL Server on the same system, there are special considerations.

To assist you in planning your deployment, the considerations are summarized in the following topics:

- [“Key information for active/active configurations”](#) on page 54
- [“Following the workflow in the Solutions Configuration Center”](#) on page 55

Key information for active/active configurations

The following table summarizes key information about active/active configurations and multiple instances and cross-references additional information:

Configuring disk groups and volumes	Create a separate set of cluster disk groups and volumes for each instance. You can create all the disk groups and volumes at one time or create them as a separate step for each instance. See “Considerations for disk groups and volumes for multiple instances” on page 67.
Configuring the cluster	If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. That way, you do not need to run the wizard again later to add the nodes. See “Configuring the cluster” on page 74.

Installing and configuring SQL Server	Assign a unique instance name, virtual server name, and port to each instance. “About installing multiple instances” on page 88. “Setting the internal name of the clustered instance” on page 100. “Assigning the port for multiple instances” on page 103.
Configuring the service group	For an active/active configuration, create a separate service group for each instance. Each service group must have a unique name and virtual IP address. There are also special considerations for specifying the priority order of systems for failover. See “Service group requirements for active/active configurations” on page 104.

Following the workflow in the Solutions Configuration Center


The Solutions Configuration Center helps you through the process of installing and configuring a new Veritas Storage Foundation HA environment for one or more instances of SQL Server 2000, in either an active/passive or active/active configuration.

[Figure 4-4](#) shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

If you are setting up multiple instances of SQL in the cluster, you may find it helpful to use the Configuration Center as follows:

- Under High Availability (HA) Configuration, complete all the steps for the first instance.
- For the next instance:
 - For step 1, Configure the storage: If you configured disk groups and volumes for the instance earlier, verify that they are available and continue with step 2.
 - For step 2, Configure the cluster: If you configured the nodes as part of the cluster earlier, as recommended, continue with step 3 and complete all subsequent steps.

Figure 4-4 Configuration steps in the Solutions Configuration Center



See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 23.

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 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. You must select the option to install the Veritas Cluster Server Database Agent for SQL. 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 4-4 describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 4-4 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.

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**. Be sure to select the following as appropriate for your installation:




Veritas Cluster Server Data- Required to configure high availability for SQL Server.
base Agent for SQL

Client Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also 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 cluster disk groups and volumes for SQL Server 2000

You must create a cluster disk group and volumes to manage your SQL Server database and logs, as covered in the following topics:

- [“About cluster disk groups and volumes”](#) on page 65
- [“Prerequisites for configuring cluster disk groups and volumes”](#) on page 66
- [“Sample disk group and volume configuration”](#) on page 66
- [“Considerations for disk groups and volumes for multiple instances”](#) on page 67
- [“Considerations for volumes for a disaster recovery configuration”](#) on page 68
- [“Creating a cluster disk group”](#) on page 68
- [“Creating volumes”](#) on page 69

About cluster disk groups and volumes

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes on only one node of a cluster. You make the volumes accessible by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Prerequisites for configuring cluster disk groups and volumes

Complete the following tasks before you create the cluster disk group and volumes:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

Sample disk group and volume configuration

On the first node of the cluster where the SQL instance is going to be installed, you first create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

For additional disk group and volume requirements, see the following topics:

- [Considerations for disk groups and volumes for multiple instances](#)
- [Considerations for volumes for a disaster recovery configuration](#)

Considerations for disk groups and volumes for multiple instances

For an active/active configuration or other cases where you are setting up multiple SQL instances in the cluster, you create a separate set of cluster disk groups and volumes for each instance.

For example, if you have a Billing instance and a Payroll instance, you could create the following disk groups and volumes.

For the Billing instance, create the following:

- BILLING_DG: a cluster disk group for the volumes related to the Billing instance
- BILLING_DATA_FILES: volume for the SQL Server system data files
- BILLING_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Billing instance
- BILLING_DB1_VOL: volume for the user database files
- BILLING_DB1_LOG: volume for the user database log files

For the Payroll Instance, create the following:

- PAYROLL_DG: a cluster disk group for the volumes related to the Payroll instance
- PAYROLL_DATA_FILES: volume for the SQL Server system data files
- PAYROLL_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Payroll instance
- PAYROLL_DB1_VOL: volume for the user database files
- PAYROLL_DB1_LOG: volume for the user database log files

You can choose either of the following:

- Set up disk groups and volumes for all instances at one time.

- Set up disk groups and volumes for the current instance only and complete all configuration steps for this instance. Then return to this step for the next instance.

Considerations for volumes for a disaster recovery configuration

For a disaster recovery configuration using Veritas Volume Replicator, note the following:

- When you run the Disaster Recovery Wizard, a VVR Storage Replicator Log volume is created automatically for each disk group that contains volumes that are replicated.
- Symantec recommends that for replication considerations, you create a separate volume for tempdb within the system database disk group, for example, INST1_TEMPDB. When you later configure replication for disaster recovery, you can exclude the tempdb volume from replication. You can create the volume now and later, after the SQL installation is complete and before configuring replication, move tempdb to the volume.
See [“Moving the tempdb database if using VVR for disaster recovery”](#) on page 103.
- VVR does not support the following types of volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.


Creating a cluster disk group

Create a cluster disk group on the first node where the SQL instance is being installed. Repeat the procedure if you want to create additional disk groups.

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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list. Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- **INST1_DATA_FILES**: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

Additional considerations apply to configuring volumes for a disaster recovery configuration using VVR.


See “[Considerations for volumes for a disaster recovery configuration](#)” on page 68

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.




- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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


- 8 Click **Next**.

9 Specify the parameters of the volume.




- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.
- 12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

13 Click **Finish** to create the new volume.

14 Repeat these steps to create additional volumes.

Create the cluster disk group and volumes on the first node of the cluster only.

Configuring the cluster

After installing SFW HA, configure the 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 provides the option to configure 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 a static IP address and that name resolution is configured for each node.
- ✓ Verify that you have the required privileges.
See “[Reviewing the requirements](#)” on page 43.


Note: If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. If you do that, you do not need to run the wizard again later to add the nodes.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS, including instructions on adding cluster nodes or 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 76.

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 76.
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 76.


- 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 77, proceed to the next step. Otherwise, proceed to [step 11](#) on page 79.

- 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 75, 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 85.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 86.

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 86. 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**.

Veritas Cluster Server Configuration Wizard
Notifier SMTP Configuration
Specify information about SMTP recipients

✓ Domain Selection
✓ VCW Options
✓ Create Cluster
✓ Select Components
▶ **Configure**
Summary
Finish

SMTP Server Name / IP:

Enter SMTP recipients and select a severity level for each recipient.

Recipients	Severity
admin@example.com	Information


Click '+' to add a recipient.
Click '-' to remove a recipient.

Click 'Next' to continue.

Back Next Cancel

- 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.

About installing multiple instances

If you are installing multiple instances of SQL Server on the same system, as in an active/active cluster configuration, some additional requirements apply. Procedures include these requirements where necessary. The following summary is provided for your review to assist you in planning the installation:

- Symantec recommends that you follow all steps for installing and setting up high availability for the first instance before you begin installing the next instance.
- Multiple instances of SQL Server must be installed in the same order on every node in the cluster. For example, if you install INST1 on SYSTEM1 and then on SYSTEM2, then install INST2 on SYSTEM1 and then on SYSTEM2.

- Assign a unique name to each instance. When installing SQL Server on additional nodes for the same instance, use the same instance name.
- Set a unique virtual server name as the internal name for each clustered instance.
- Assign a unique port number for each instance.

Installing and configuring SQL Server 2000 on the first node

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node for this SQL instance and the volumes are mounted (are assigned drive letters). See “[Importing the cluster disk group](#)” on page 94 and “[Adding drive letters to mount the volumes](#)” on page 94.

Complete the following procedures to install and configure this instance of Microsoft SQL Server 2000:

- ✓ [Installing Microsoft SQL Server 2000](#)
- ✓ [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server 2000

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.


To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.

- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.




- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.

- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
 - b In the Services box, select the **SQL Server** option.
 - c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
 - d If you chose to run the service in the context of a domain user, specify the information for the user.
 - e Clear the **Auto Start Service** option.
 - f Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g Click **Next**.
- 11 Follow the wizard instructions to complete the installation.
 - 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
 - 13 If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.


Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- a Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - b Select a service from the **Services** list.
 - c Clear the **Auto-start service when OS starts** check box.
- 2 Repeat [step b](#) and [step c](#) for all other SQL Server services that are running on the server.

Preparing to install SQL Server 2000 on the second node

Follow the following procedures before installing SQL Server on the second or additional nodes for the SQL instance:


- [“Stopping the SQL Server 2000 service”](#) on page 93
- [“Deporting the cluster disk group”](#) on page 93
- [“Importing the cluster disk group”](#) on page 94
- [“Adding drive letters to mount the volumes”](#) on page 94
- [“Renaming shared SQL Server 2000 files”](#) on page 96

Stopping the SQL Server 2000 service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.

- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.


Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2000 on the second node

Use the following procedures to install and configure SQL Server on additional nodes for this instance:

- [“Installing SQL Server 2000 on the second node”](#) on page 96
- [“Removing shared SQL Server 2000 files”](#) on page 99

Installing SQL Server 2000 on the second node

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 94 and [“Adding drive letters to mount the volumes”](#) on page 94.

Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.


Multiple instances of SQL Server must be installed in the same order on every node of the cluster.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.


Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000


- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
 Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
- b In the Services box, select the **SQL Server** option.
- c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
- d If you chose to run the service in the context of a domain user, specify the information for the user.

- e Clear the **Auto Start Service** option.
 - f Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g Click **Next**.
- 11 Follow the wizard instructions to complete the installation.
 - 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
 - 13 If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server 2000 on the second node](#)” on page 92 and “[Installing and configuring SQL Server 2000 on the second node](#)” on page 96 on any additional nodes.

Removing shared SQL Server 2000 files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.


Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.


To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example `SYSTEM2\INST1`.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
“sp_dropserver ‘*System_Name\Instance_Name*.’”
where *System_Name\Instance_Name* is the name noted in step 3.
For example, for named instance:
“sp_dropserver ‘SYSTEM2\INST1.’”
For example, for a default instance:
“sp_dropserver ‘SYSTEM1.’”
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
- In the top pane, enter
“sp_addserver ‘*Virtual_Server_Name\Instance_Name*’,
local”
For example ‘INST1-VS\INST1’, local for a named instance, or
‘INST1-VS’, local for a default instance.

- Press F5.

Creating a SQL Server user-defined database

You can use SFW HA to manage a SQL Server user-defined database.

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them first.

See [“Creating volumes”](#) on page 69.

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.
- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.
- 8 Configure whatever other options are required for your database.
- 9 Depending on your configuration plans, you may have additional steps to complete in SQL Server.
See [“Completing configuration steps in SQL Server”](#) on page 102.
- 10 If the SQL Server service group has already been configured, you need to add the resources for the new database to the service group.
See [“Modifying the SQL 2000 service group to add VMDg and MountV resources”](#) on page 116.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Veritas Volume Replicator (VVR), Symantec recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume.

See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 103.

If you are running multiple SQL Server instances, you must assign a different port to each SQL Server instance.

See “[Assigning the port for multiple instances](#)” on page 103.

Moving the tempdb database if using VVR for disaster recovery

If you plan to implement a disaster recovery configuration using VVR, Symantec recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See “[Creating volumes](#)” on page 69.

Then, refer to the Microsoft Knowledge Base for the instructions on moving the tempdb database. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 224071: How to move SQL Server databases to a new location by using Detach and Attach functions in SQL Server

Refer to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Assigning the port for multiple instances

When you are running multiple SQL Server instances, you must assign a different port to each instance.

Refer to the Microsoft Knowledge Base for the instructions on assigning static ports. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 823938: How to configure an instance of SQL Server to listen on a specific TCP port or a dynamic port

Refer to:

<http://support.microsoft.com/kb/823938/en-us>

Configuring the VCS SQL Server 2000 service group

A VCS SQL Server service group is used to bring a SQL Server 2000 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the

service group. You use the VCS SQL Configuration Wizard to configure the service group.

Read the following topics:

- [Service group requirements for active/active configurations](#)
- [Prerequisites for configuring the service group](#)
- [Creating the SQL Server 2000 service group](#)

Service group requirements for active/active configurations

Note the following requirements for active/active configurations:

- For an active/active configuration, you must create a separate service group for each instance.
- Each service group that you create must have a unique service group name and virtual IP address.
- For an active/active configuration, when you specify the priority order of systems, reverse the order for each service group so that the active system and failover system are opposite for each instance. For example, if you have two instances and two systems, you would set the priority order as follows:

INSTANCE 1	Priority order:
	SYSTEM1
	SYSTEM2
INSTANCE2	Priority order:
	SYSTEM2
	SYSTEM1

Prerequisites for configuring the service group

Complete the following tasks before configuring the service group:

- ✓ Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2000, is installed on all cluster nodes. See [“Installing Veritas Storage Foundation HA for Windows”](#) on page 59.
- ✓ Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See [“Configuring the cluster”](#) on page 74.
- ✓ Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.

- ✓ Verify that the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group. See [“Importing the cluster disk group”](#) on page 94 and [“Adding drive letters to mount the volumes”](#) on page 94.
- ✓ Verify that the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- ✓ To avoid having to modify the service group later, create the user-defined databases before creating the service group. See [“Creating a SQL Server user-defined database”](#) on page 102.
- ✓ Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name when configuring the service group. See [“Setting the internal name of the clustered instance”](#) on page 100.

Note: For a disaster recovery configuration, the SQL Server virtual server name on the secondary site cluster must match the one on the primary site cluster.

- ✓ Assign a unique virtual IP address for the SQL Server 2000 instance. You specify this IP address when configuring the service group.
- ✓ Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists on all the cluster nodes. A sample script is supplied in C:\Program Files\Veritas\cluster server\bin\SQLServer2000\sample_script.sql. Detailed monitoring is often not necessary.
- ✓ Stop the SQL 2000 Server service for the SQL instance. See [“Stopping the SQL Server 2000 service”](#) on page 93.

Creating the SQL Server 2000 service group


The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server >**


Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard.

- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:




- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- 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 Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
- Click **Next**.

- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment. Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 107 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
- Click **Apply**.

- 9
- If you want to configure detail monitoring for additional instances, repeat [step 8](#) for all the instances for which detail monitoring will be configured.
- 10
- Click **Next**.
- 11
- In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12
- Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.

- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review 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.
 - 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
 - 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
 - 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
 You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for high availability”](#) on page 219.

Verifying the SQL Server 2000 cluster configuration

Failover simulation is an important part of configuration testing.

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.

Determining additional steps needed

Determine the next step as follows:

- If you have no more instances of SQL Server that you plan to install and configure for high availability, your high availability configuration is complete.
- If you have additional instances of SQL Server to install and configure in the cluster, start the installation and configuration sequence again for the next instance.
 - If you are using the Solutions Configuration Center:
 See [“Following the workflow in the Solutions Configuration Center”](#) on page 55.
 - If you are using the printed document as a guide:
 See [“Tasks for a new HA installation of SQL Server 2000”](#) on page 36.
- If you are planning to set up a secondary site for disaster recovery:
 See [Chapter 12, “Deploying disaster recovery: New SQL Server 2000 installation”](#) on page 469.
- If you want to configure a Cluster Management Console connection:
 See [“Configuring the Cluster Management Console connection”](#) on page 111.
- If you need to modify the SQL Server service group:
 See [“Modifying the SQL 2000 service group to add VMDg and MountV resources”](#) on page 116.
- If you need to make other changes to the cluster:
 See *Veritas Cluster Server Administrator’s Guide*.

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 server

The 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: 9dfde3d9aaebec084f8e35819c1fed7e6b01d2ae

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 SQL 2000 service group to add VMDg and MountV resources

If you add a new SQL Server database after configuring the SQL Server service group, you must add the VMDg and MountV resources to the SQL Server service group.

Before running the SQL Server Agent wizard to add the VMDg and MountV resources:

- ✓ Make sure the SQL Server resources are online.
- ✓ Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.

- 11 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for high availability”](#) on page 219.

Deploying SFW HA for high availability: Standalone SQL Servers

This chapter includes the following topics:

- [Tasks for converting a standalone SQL 2000 Server for high availability](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Preparing the standalone SQL Server](#)
- [Installing and configuring SFW HA](#)
- [Installing and configuring SQL Server 2000 on additional nodes](#)
- [Moving the existing SQL Server 2000 data files and user databases](#)
- [Configuring the VCS SQL Server service group](#)
- [Creating a SQL Server user-defined database](#)
- [Verifying the installation](#)
- [Additional instructions for disaster recovery](#)

Tasks for converting a standalone SQL 2000 Server for high availability

This chapter describes the procedure to convert a standalone SQL 2000 Server into a “clustered” SQL Server in a new Veritas Storage Foundation HA

environment. This environment involves an active/passive configuration with one to one failover capabilities.

If you are planning a new SQL Server 2000 deployment, or to review additional considerations for an active/active configuration, refer to “[Deploying SFW HA for high availability: New SQL Server 2000 installation](#)” on page 35.

Note: In addition to the information contained in this chapter, the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach are required. Refer to: <http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

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

Table 5-1 Tasks for converting a standalone SQL 2000 server for high availability

Objective	Tasks
“ Reviewing the prerequisites ” on page 122	✓ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 126	✓ Understanding active/passive configuration ✓ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 128	✓ Setting up the storage hardware for a cluster environment ✓ Verifying the DNS entries for the systems on which SQL will be installed
“ Preparing the standalone SQL Server ” on page 130	✓ Backing up existing data ✓ Setting SQL Server services to manual start

Table 5-1 Tasks for converting a standalone SQL 2000 server for high availability (continued)

Objective	Tasks
“Installing and configuring SFW HA” on page 131	✓ Installing Veritas Storage Foundation HA for Windows (automatic installation) ✓ Configuring the cluster ✓ Configuring cluster disk groups and volumes
“Installing and configuring SQL Server 2000 on additional nodes” on page 131	✓ Preparing for installation on additional nodes ✓ Installing SQL Server on additional nodes
“Moving the existing SQL Server 2000 data files and user databases” on page 132	✓ Stopping SQL Server service ✓ Modifying data file and user database locations
“Configuring the VCS SQL Server service group” on page 133	✓ Creating a SQL Server service group using the VCS SQL Configuration wizard
“Creating a SQL Server user-defined database” on page 133	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server ✓ Adding resources for a user-defined database in VCS
“Verifying the installation” on page 133	✓ Simulating failover ✓ Switching online nodes
“Additional instructions for disaster recovery” on page 134	✓ Locating the instructions for continuing with a disaster recovery configuration (optional)

Reviewing the prerequisites

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

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 Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the following options:
Veritas Volume Replicator and the Global Clustering Option
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with
Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 required for all editions)
or
Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition
or
Windows 2000 Server, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server (all require SP4 with Update Rollup 1)
or
 - Microsoft SQL Server 2000 (64-bit) Enterprise Edition
with
Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required for both)
or
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with

Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T):
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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 125.
- 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 SQL 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 *VCS 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.
- 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

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 doing the following tasks:

- 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.

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 5-2](#) estimates disk space requirements for SFW HA.

Table 5-2 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

Reviewing the configuration

This chapter describes the tasks needed to incorporate an existing standalone SQL Server into a high available environment in order to ensure that the mission critical SQL resource is always available.

This chapter describes the tasks necessary to create a virtual server in an active/passive SQL configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node. At the end of this process, their environment will look like this:

Figure 5-1 Active/passive configuration



The virtual SQL Server is online on SYSTEM1, serving client requests. The shared disks provide storage for the SQL Server databases. SYSTEM2 waits in a warm standby state as a backup node, prepared to begin handling client requests if SYSTEM1 becomes unavailable. From the user's perspective there will be a small delay as the backup node comes online, but the interruption in effective service is minimized.

Sample configuration

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

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

- ✓ SQL virtual server: the IP address should be the same on all nodes
- ✓ Cluster IP address: used by Veritas Cluster Management Console (Single Cluster Mode) also referred to as Web Console

You should have these IP addresses available before you start deploying your environment. The following names describe the objects created and used during the installation and configuration:

Name	Object
SYSTEM1 & SYSTEM2	server names; SYSTEM1 is the existing standalone SQL Server
INST1_SG	Microsoft SQL Server 2000 service group
SQL_CLUS1	virtual SQL Server cluster
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1	SQL Instance Name
INST1-VS	name of the SQL Virtual Server

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings.

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.

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. Contact the NIC manufacturer for details on this process. Symantec recommends removing Internet Protocol 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. GAB 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.
- 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 on which SQL will be installed

- 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:
 - a From the **Advanced** menu, click **Advanced Settings**.
 - b On 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.
 - c Click **OK**.
- 4 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, and not those configured for the VCS private network.

- 5 From the status window, click **Properties**.
- 6 On the **General** tab:
 - a Select the **Internet Protocol (TCP/IP)** check box.
 - b Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify that the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 On 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**.

Preparing the standalone SQL Server

Complete the following tasks before you begin the process of installing Veritas Storage Foundation HA for Windows and creating a clustered environment:

- ✓ Backing up existing SQL data
- ✓ Setting SQL Server services to manual start

Backing up existing SQL data


Create a backup of the data on the existing standalone SQL Server.

Setting SQL Server services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- a Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - b Select a service from the **Services** list.
 - c Clear the **Auto-start service when OS starts** check box.
- 2 Repeat [step b](#) and [step c](#) for all other SQL Server services that are running on the server.

Installing and configuring SFW HA

Before converting the existing standalone SQL into a “clustered” server, complete the following procedures:

- ✓ [“Installing Veritas Storage Foundation HA for Windows”](#) on page 59 to install SFW HA on all nodes of the cluster.
- ✓ [“Configuring cluster disk groups and volumes for SQL Server 2000”](#) on page 65 to create the disk group and volumes necessary to manage the SQL Server storage.
- ✓ [“Configuring the cluster”](#) on page 74 to configure VCS within the cluster.

Installing and configuring SQL Server 2000 on additional nodes

Ensure the cluster disk group is imported on the node where you are installing SQL Server 2000. If necessary, complete the procedures [“Deporting the cluster disk group”](#) on page 93 on the first node and then [“Importing the cluster disk group”](#) on page 94 on the second node so the cluster disk group is imported to the second node.

SQL Server 2000 must have the same configuration on all nodes in the cluster. You will need the following information:

- ✓ Instance name (if applicable)
- ✓ Destination Folder for Program Files and Data Files
- ✓ Authentication Mode

Complete the procedure: [“Installing and configuring SQL Server 2000 on the second node”](#) on page 96

Moving the existing SQL Server 2000 data files and user databases

After completing the SQL installation and configuration on the additional nodes, move the existing standalone SQL Server 2000 data files and user databases from the local drive to the shared drives to ensure proper failover operations in the cluster. Complete the following tasks to move the databases:

- 1 Stop the SQL Server service as described in the section “[Stopping the SQL Server 2000 service](#)” on page 93.
- 2 Verify that you have backed up your existing data.
- 3 Import the dynamic disk group to the node where the original files are located on the local drives and mount the volumes (add drive letters). Refer to “[Deporting the cluster disk group](#)” on page 93, “[Importing the cluster disk group](#)” on page 94, and “[Adding drive letters to mount the volumes](#)” on page 94 for instructions.
- 4 Modify the SQL Server 2000 Data File and User Database locations. Follow the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach.
<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>
- 5 Restart SQL Server 2000.

To start the SQL service

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).
- 2 Click **Start**.

Configuring the VCS SQL Server service group

The Enterprise Agent Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

Complete the procedure “[Creating a SQL Server user-defined database](#)” on page 102 to configure the SQL Server service group.

Creating a SQL Server user-defined database

Complete the procedure “[Creating a SQL Server user-defined database](#)” on page 102 to create and manage a SQL Server user-defined database. The procedure includes the following tasks:

- ✓ Create volumes for a user-defined SQL Server database and its transaction log.
- ✓ Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- ✓ Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Verifying the installation

Verify your installation by switching online nodes or by shutting down the computer that is currently online. Either process will test that the service group can be smoothly transferred between nodes.

Shutting down a node creates an actual failure, stressing your system, but more truly testing its high availability than by switching nodes. If you do shut down the online computer in your cluster, remember to bring it back up after you have confirmed that the service group successfully failed over to another node.

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.

Additional instructions for disaster recovery

After completing the tasks for setting up a SFW HA environment for SQL Server 2000 on a primary site, you can optionally create a secondary or “failover” site for disaster recovery.

See [Chapter 12, “Deploying disaster recovery: New SQL Server 2000 installation”](#) on page 469.

Deploying SFW HA for high availability: New SQL Server 2005 installation

This chapter covers the following topics:

- [Tasks for a new HA installation of SQL Server 2005](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Reviewing considerations for active/active configurations](#)
- [Configuring the storage hardware and network](#)
- [Installing Veritas Storage Foundation for Windows](#)
- [Configuring cluster disk groups and volumes for SQL Server 2005](#)
- [Configuring the cluster](#)
- [About installing multiple instances](#)
- [Installing and configuring SQL Server 2005 on the first node](#)
- [Preparing to install SQL Server 2005 on the second node](#)
- [Installing and configuring SQL Server 2005 on the second node](#)
- [Setting the internal name of the clustered instance](#)
- [Creating a SQL Server user-defined database](#)
- [Completing configuration steps in SQL Server](#)

- [Configuring the VCS SQL Server 2005 service group](#)
- [Verifying the SQL Server 2005 cluster configuration](#)
- [Determining additional steps needed](#)

Tasks for a new HA installation of SQL Server 2005

You can install and configure a new Veritas Storage Foundation HA environment for SQL Server 2005 in the following ways:

active/passive	One SQL 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 SQL instances per cluster node. For example, in a two-node cluster with two SQL 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.

Note: Some installation and configuration options are identified as required “for a disaster recovery configuration.” These options apply only if you intend to set up a secondary site for disaster recovery.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA for SQL Server 2005.

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

To configure MSDTC service groups, see “[Configuring an MSDTC service group for high availability](#)” on page 219.

Tasks for an active/passive configuration

Table 6-1 outlines the high-level objectives and the tasks to complete each objective for an active/passive configuration.

Table 6-1 Tasks for installing and configuring active/passive HA for SQL Server 2005

Objective	Tasks
“ Reviewing the requirements ” on page 142	✓ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 146	✓ Understanding active/passive configuration ✓ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 156	✓ Setting up the storage hardware for a cluster environment ✓ Verifying the DNS entries for the systems on which SQL will be installed
“ Installing Veritas Storage Foundation for Windows ” on page 158	✓ Verifying the driver signing option for the system ✓ Installing Veritas Storage Foundation HA for Windows (automatic installation) ✓ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“ Configuring cluster disk groups and volumes for SQL Server 2005 ” on page 164	✓ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ✓ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator

Table 6-1

Tasks for installing and configuring active/passive HA for SQL Server 2005 (continued)

Objective	Tasks
“Configuring the cluster” on page 173	✓ Verifying static IP addresses and name resolution configured for each node ✓ Running the Veritas Cluster Configuration Wizard to configure cluster components and set up secure communication for the cluster
“Installing and configuring SQL Server 2005 on the first node” on page 188	✓ Installing and configuring SQL Server 2005 ✓ Setting SQL Server services to manual start
“Preparing to install SQL Server 2005 on the second node” on page 193	✓ Stopping the SQL service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Mounting the volumes (adding drive letters) ✓ Renaming or removing shared SQL files
“Installing and configuring SQL Server 2005 on the second node” on page 196	✓ Installing and configuring SQL Server 2005 on the second and any additional nodes
“Setting the internal name of the clustered instance” on page 200	✓ Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 203	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server

Table 6-1 Tasks for installing and configuring active/passive HA for SQL Server 2005 (continued)

Objective	Tasks
“Completing configuration steps in SQL Server” on page 204	✓ Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2005 service group” on page 205	✓ Creating a SQL Server service group using the SQL Server Configuration Wizard
“Verifying the SQL Server 2005 cluster configuration” on page 211	✓ Simulating failover ✓ Switching online nodes

Tasks for an active/active configuration

Table 6-2 outlines the high-level objectives and the tasks to complete each objective for an active/active configuration.

Table 6-2 Tasks for installing and configuring active/active HA for SQL Server 2005

Objective	Tasks
“Reviewing the requirements” on page 142	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 146	✓ Understanding active/active configuration ✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 156	✓ Setting up the storage hardware for a cluster environment ✓ Verifying the DNS entries for the systems on which SQL will be installed

Table 6-2

Tasks for installing and configuring active/active HA for SQL Server 2005

Objective	Tasks
“Installing Veritas Storage Foundation for Windows” on page 158	✓ Verifying the driver signing option for the system ✓ Installing Veritas Storage Foundation HA for Windows (automatic installation) ✓ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“Configuring cluster disk groups and volumes for SQL Server 2005” on page 164	✓ Creating a dynamic disk group for each instance ✓ Creating dynamic volumes for the SQL system database, user databases, transaction logs, and replicated registry keys for each instance
“Configuring the cluster” on page 173	✓ Verifying static IP addresses and name resolution configured for each node ✓ Running the Veritas Cluster Configuration Wizard to configure cluster components and set up secure communication for the cluster
“Installing and configuring SQL Server 2005 on the first node” on page 188	✓ Installing and configuring SQL Server on the first node ✓ Setting SQL Server services to manual start

Table 6-2 Tasks for installing and configuring active/active HA for SQL Server 2005

Objective	Tasks
“Preparing to install SQL Server 2005 on the second node” on page 193	✓ Stopping the SQL service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Mounting the volumes (adding drive letters) ✓ Renaming or removing shared SQL files
“Installing and configuring SQL Server 2005 on the second node” on page 196	✓ Installing SQL Server on the second node
“Setting the internal name of the clustered instance” on page 200	✓ Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 203	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server
“Completing configuration steps in SQL Server” on page 204	✓ Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2005 service group” on page 205	✓ Creating a SQL Server service group using the VCS SQL Configuration Wizard ✓ Ensuring the priority order of the systems is set up in reverse order for each instance

Table 6-2 Tasks for installing and configuring active/active HA for SQL Server 2005

Objective	Tasks
“ Verifying the SQL Server 2005 cluster configuration ” on page 211	✓ Simulating fail over ✓ Switching online nodes
“ Determining additional steps needed ” on page 212	✓ Repeating the installation and configuration steps for the next instance

Reviewing the requirements

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

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 6-3](#) estimates disk space requirements for SFW HA.

Table 6-3 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) with the Veritas Cluster Server Database Agent for Microsoft SQL
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SP1 required)
with
 - Windows Server 2003 (32-bit) Enterprise Edition or Datacenter Edition (SP1 required)
or
 - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
or
 - Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SP1 required)
with
 - Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required)
or
 - Windows Server 2003 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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 146.
- 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 SQL 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.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

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

Review the information for the configurations you have planned:


- [Active/passive configuration](#)
- [Active/active configuration](#)
- [Disaster recovery configuration](#)

Active/passive configuration

In a typical example of a high availability cluster, you create a virtual SQL server in an active/passive configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node.

[Figure 6-1](#) illustrates a typical active/passive configuration. SQL Server is installed on both SYSTEM1 and SYSTEM2 and configured as a virtual server (INST1-VS) with a virtual IP address. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured to come online on SYSTEM1 first. If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 6-1 Active/passive configuration



Sample active/passive configuration

A sample setup is used to illustrate the installation and configuration tasks. The following names describe the objects created and used during the installation and configuration:

Name	Object
SYSTEM1 & SYSTEM2	servers
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server
INST1_SG	SQL Server service group

IP addresses for sample active/passive configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there is one SQL Server virtual server. Therefore you would need one virtual server IP address. If you want to use the VCS Web Console or the notification service, you require a cluster IP address. The cluster IP address is also used by the Global Cluster Option for disaster recovery.

See “Network requirements” under “[Reviewing the requirements](#)” on page 142.

See “[Disaster recovery configuration](#)” on page 151.


Active/active configuration

In an active/active SQL Server configuration, several instances are intended to run on a single node when necessary. A SQL Server instance is a completely independent SQL Server installation, with its own services, master database, storage, and memory resources. Each instance is defined uniquely by a separate SQL Server virtual server and service group.

A SQL Server instance can fail over to any of the other cluster nodes that you specify when you configure the SQL Server service group.

The following figure illustrates a two node active/active configuration. The SQL Server databases are configured on the shared storage on volumes contained in cluster disk groups. Each SQL Server virtual server is configured in a separate SQL Server service group. Each service group can fail over to the other node in the cluster.

Figure 6-2 Active/active configuration



For example, consider a two-node cluster hosting two SQL Server virtual servers, BILLING_VS and PAYROLL_VS. The table below and the sample configuration illustrate that the virtual servers are configured in two separate service groups with BILLING_VS online on SYSTEM1 but able to fail over to SYSTEM2, and PAYROLL_VS online on SYSTEM2 but able to fail over to SYSTEM1.

SQL Virtual Server	Service Group	System List
BILLING_VS	BILLING_SG	SYSTEM1, SYSTEM2
PAYROLL_VS	PAYROLL_SG	SYSTEM2, SYSTEM1

Sample active/active configuration

A sample setup is used to illustrate the installation and configuration tasks for two instances of SQL Server, Billing and Payroll. During normal operation, one instance will be online on each of the two servers. If a failure occurs, the instance on the failing node will be brought online on the other server, resulting in two instances running on one server.

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

Name	Description
SYSTEM1 & SYSTEM2	server names

BILLING_DG	cluster disk group for the billing instance
PAYROLL_DG	cluster disk group for the payroll instance
BILLING_VS_SYS_FILES	volume for the SQL Server system data files for the billing instance
PAYROLL_VS_SYS_FILES	volume for the SQL Server system data files for the payroll instance
BILLING_DATA	volume for a SQL Server user-defined database for the billing instance
PAYROLL_DATA	volume for a SQL Server user-defined database for the payroll instance
BILLING_LOG	volume for a SQL Server user-defined database log file for the billing instance
PAYROLL_LOG	volume for a SQL Server user-defined database log file for the payroll instance
BILLING_REGREP	volume for the list of registry keys replicated among the nodes for the billing instance
PAYROLL_REGREP	volume for the list of registry keys replicated among the nodes for the payroll instance
SQL_CLUS1	virtual SQL Server cluster
BILLING_INST	instance name for the billing instance
PAYROLL_INST	instance name for the payroll instance
BILLING_VS	virtual SQL Server name for the billing instance
PAYROLL_VS	virtual SQL Server name for the payroll instance
BILLING_SG	SQL Server service group for the billing instance
PAYROLL_SG	SQL Server service group for the payroll instance

IP addresses for sample active/active configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there are two virtual servers: BILLING-VS and PAYROLL-VS. Therefore, you would need two virtual server IP addresses. If you want to use the VCS Web Console or the notification service, you require a cluster IP address. The cluster IP address is also used by the Global Cluster Option for disaster recovery.

See “Network requirements” under “[Reviewing the requirements](#)” on page 142.

See “[Disaster recovery configuration](#)” on page 151.

Disaster recovery configuration


You may be preparing to configure both a primary site and a secondary site for disaster recovery.

[Figure 6-3](#) illustrates a typical active/passive disaster recovery configuration using Veritas Volume Replicator (VVR).

In the example illustration, the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the Replicated Volume Group (RVG). The SQL Server application data is stored on the volumes that are under the control of the RVG.

If the Microsoft SQL Server server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Figure 6-3 Typical active/passive disaster recovery configuration with VVR



Sample disaster recovery configuration

The sample disaster recovery setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

A sample setup is used to illustrate the installation and configuration tasks:

Primary Site

SYSTEM1 & servers
SYSTEM2

INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_REPLOG	replicator log volume required by VVR
INST1_DB1_RDS	RDS (Replicated Data Set) name
INST1_DB1_RVG	RVG (Replicated Volume Group) name
INST1_DB1_RVG_SG	Replication service group

IP addresses for sample disaster recovery configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration.

See “Network requirements” under “[Reviewing the requirements](#)” on page 142.

You specify the following addresses during the replication process:

SQL virtual server IP address	For a disaster recovery configuration, the virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

Cluster IP address	You need one for the primary site cluster and one for the secondary site cluster.
Replication IP address	You need two IP addresses per application instance, one for the primary site and one for the secondary site.

Reviewing considerations for active/active configurations

- For an active/active configuration or in other cases when you are installing multiple instances of SQL Server on the same system, there are special considerations.
- To assist you in planning your deployment, the considerations are summarized in the following topics:
- [“Key information for active/active configurations”](#) on page 154
 - [“Following the workflow in the Solutions Configuration Center”](#) on page 155

Key information for active/active configurations

The following table summarizes key information about active/active configurations and multiple instances and cross-references additional information:

Configuring disk groups and volumes	Create a separate set of cluster disk groups and volumes for each instance. You can create all the disk groups and volumes at one time or create them as a separate step for each instance. See “Considerations for disk groups and volumes for multiple instances” on page 165.
Configuring the cluster	If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. That way, you do not need to run the wizard again later to add the nodes. See “Configuring the cluster” on page 173.

Installing and configuring SQL Server	Assign a unique instance name, virtual server name, and port to each instance. “About installing multiple instances” on page 187. “Setting the internal name of the clustered instance” on page 200. “Assigning the port for multiple instances” on page 204.
Configuring the service group	For an active/active configuration, create a separate service group for each instance. Each service group must have a unique name and virtual IP address. There are also special considerations for specifying the priority order of systems for failover. See “Service group requirements for active/active configurations” on page 205.

Following the workflow in the Solutions Configuration Center


The Solutions Configuration Center helps you through the process of installing and configuring a new Veritas Storage Foundation HA environment for one or more instances of SQL Server 2005, in either an active/passive or active/active configuration.

[Figure 6-4](#) shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

If you are setting up multiple instances of SQL in the cluster, you may find it helpful to use the Configuration Center as follows:

- Under High Availability (HA) Configuration, complete all the steps for the first instance.
- For the next instance:
 - For step 1, Configure the storage: If you configured disk groups and volumes for the instance earlier, verify that they are available and continue with step 2.
 - For step 2, Configure the cluster: If you configured the nodes as part of the cluster earlier, as recommended, continue with step 3 and complete all subsequent steps.

Figure 6-4 Configuration steps in the Solutions Configuration Center



See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 23.

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 Veritas Storage Foundation 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. You must select the option to install the Veritas Cluster Server Database Agent for SQL. 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 6-4 describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 6-4 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.

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**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL	Required to configure high availability for SQL Server.
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also 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 cluster disk groups and volumes for SQL Server 2005

You must create a cluster disk group and volumes to manage your SQL Server database and logs, as covered in the following topics:

- [About cluster disk groups and volumes](#)
- [Prerequisites for configuring cluster disk groups and volumes](#)
- [Sample disk group and volume configuration](#)
- [Considerations for disk groups and volumes for multiple instances](#)
- [Considerations for volumes for a disaster recovery configuration](#)
- [Creating a cluster disk group](#)
- [Creating volumes](#)

About cluster disk groups and volumes

SFW uses disk groups to organize disks or LUNs for management purposes. A dynamic disk group is a collection of disks that is imported or deported as a single unit. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes for the SQL instance on only one node of a cluster. You make the volumes accessible by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Prerequisites for configuring cluster disk groups and volumes

Complete the following tasks before you create the cluster disk group and volumes for the SQL instance:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.

- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

Sample disk group and volume configuration

On the first node where the SQL instance is going to be installed, you first create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process:

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

For additional disk group and volume requirements, see the following topics:

- [Considerations for disk groups and volumes for multiple instances](#)
- [Considerations for volumes for a disaster recovery configuration](#)

Considerations for disk groups and volumes for multiple instances

For an active/active configuration or other cases where you are setting up multiple SQL instances in the cluster, you create a separate set of cluster disk groups and volumes for each instance.

For example, if you have a Billing instance and a Payroll instance, you could create the following disk groups and volumes.

For the Billing instance, create the following:

- BILLING_DG: a cluster disk group for the volumes related to the Billing instance
- BILLING_DATA_FILES: volume for the SQL Server system data files
- BILLING_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Billing instance
- BILLING_DB1_VOL: volume for the user database files
- BILLING_DB1_LOG: volume for the user database log files

For the Payroll Instance, create the following:

- PAYROLL_DG: a cluster disk group for the volumes related to the Payroll instance
- PAYROLL_DATA_FILES: volume for the SQL Server system data files
- PAYROLL_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Payroll instance
- PAYROLL_DB1_VOL: volume for the user database files
- PAYROLL_DB1_LOG: volume for the user database log files

You can choose either of the following:

- Set up disk groups and volumes for all instances at one time.
- Set up disk groups and volumes for the current instance only and complete all configuration steps for this instance. Then return to this step for the next instance.

Considerations for volumes for a disaster recovery configuration

For a disaster recovery configuration using Veritas Volume Replicator, note the following:

- When you run the Disaster Recovery Wizard, a VVR Storage Replicator Log volume is created automatically for each disk group that contains volumes that are replicated.
- Symantec recommends that for replication considerations, you create a separate volume for tempdb within the system database disk group, for example, INST1_TEMPDB. When you later configure replication for disaster recovery, you can exclude the tempdb volume from replication. You can create the volume now and later, after the SQL installation is complete and before configuring replication, move tempdb to the volume.

See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 204.

- VVR does not support the following types of volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.


Creating a cluster disk group

Use the Veritas Enterprise Administrator (VEA) to create a cluster disk group on the first node where the SQL instance is being installed. Repeat the procedure if you want to create additional disk groups.

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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Creating volumes

This procedure will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure to create additional volumes. For example, create the following volumes on the INST1-DG disk group:

- **INST1_DATA_FILES**: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)

- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

You could create the following volumes on the same disk group as the system volumes or on a separate disk group, depending on how you prefer to configure your storage:

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

Additional considerations apply to configuring volumes for a disaster recovery configuration using VVR.


See [“Considerations for volumes for a disaster recovery configuration”](#) on page 166

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
 You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6
- Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.




- 7
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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


- 8
- Click **Next**.

9 Specify the parameters of the volume.




- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
 - 14 Repeat these steps to create additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Configuring the 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 provides the option to configure the ClusterService group, which can contain 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 and that name resolution is configured for each node.
- ✓ Verify that you have the required privileges.
See “[Reviewing the requirements](#)” on page 142.


Note: If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. If you do that, you do not need to run the wizard again later to add the nodes.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS, including instructions on adding cluster nodes or 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 176.

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 175. 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 176.

- 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 176, proceed to the next step. Otherwise, proceed to [step 11](#) on page 178.

- 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 174, 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:

The screenshot shows the 'Veritas Cluster Server Configuration Wizard' window. The title bar reads 'Veritas Cluster Server Configuration Wizard'. The main window has a left-hand pane with a tree view containing the following items: 'Domain Selection', 'VCW Options', 'Cluster Details', 'Cluster Selection', 'Validate Systems', 'Edit Options', 'NIC Selection', 'Service Account', 'Security', 'Summary', and 'Finish'. The 'Security' item is currently selected and highlighted in yellow. The main area of the wizard is titled 'Configure Security Service Option' with the subtitle 'Select security options for the cluster'. It contains two radio button options. The first option is 'Use Single Sign-on' with the description 'Configure the Symantec Product Authentication Service for secure communication in the cluster.' The second option is 'Use VCS User Privileges' with the description '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.' The 'Use VCS User Privileges' option is selected. Below this description are three text input fields: 'User Name:' (containing 'admin'), 'Password:' (containing 'password'), and 'Re-enter Password:' (containing 'password'). At the bottom right of the window are three buttons: 'Back', 'Next', and '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 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 184.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 185.

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 185. 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.

About installing multiple instances

If you are installing multiple instances of SQL Server on the same system, as in an active/active cluster configuration, some additional requirements apply. Procedures include these requirements where necessary. The following summary is provided for your review to assist you in planning the installation:

- Symantec recommends that you follow all steps for installing and setting up high availability for the first instance before you begin installing the next instance.

- Multiple instances of SQL Server must be installed in the same order on every node in the cluster. For example, if you install INST1 on SYSTEM1 and then on SYSTEM2, then install INST2 on SYSTEM1 and then on SYSTEM2.
- Assign a unique name to each instance. When installing SQL Server on additional nodes for the same instance, use the same instance name.
- Set a unique virtual server name as the internal name for each clustered instance.
- Assign a unique port number for each instance.

Installing and configuring SQL Server 2005 on the first node

In preparation for installing Microsoft SQL Server 2005, ensure that the cluster disk group is imported to the first node for this SQL instance and the volumes are mounted. See [“Importing the cluster disk group”](#) on page 194 and [“Adding drive letters to mount the volumes”](#) on page 195.

Complete the following procedures to install and configure this instance of Microsoft SQL Server 2005:

- ✓ [Installing Microsoft SQL Server 2005 on the first node](#)
- ✓ [Setting SQL Server 2005 services to manual start](#)

Installing Microsoft SQL Server 2005 on the first node

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.


Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.


To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
 Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
 If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.


Setting SQL Server 2005 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.

- 4
- In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5
- Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node

Complete the following procedures before installing SQL Server on the second or additional nodes for the SQL instance:

- [“Stopping the SQL Server 2005 service”](#) on page 193
- [“Deporting the cluster disk group”](#) on page 194
- [“Importing the cluster disk group”](#) on page 194
- [“Adding drive letters to mount the volumes”](#) on page 195
- [“Renaming shared SQL Server 2005 files”](#) on page 196

Stopping the SQL Server 2005 service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

To install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name where the disk group is current imported, expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (for example, INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (for example, INST1_DG) to the next node in the cluster.

To import a cluster disk group


- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name, right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (for example, INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
 Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2005 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes for this instance:

- [“Installing Microsoft SQL Server on the second node”](#) on page 196
- [“Setting the internal name of the clustered instance”](#) on page 200

Installing Microsoft SQL Server on the second node

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the additional node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 194 and [“Adding drive letters to mount the volumes”](#) on page 195.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

Multiple instances of SQL Server 2005 must be installed in the same order on every node of the cluster.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.


To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
 - 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.

Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.

- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.

Refer to [“Setting SQL Server 2005 services to manual start”](#) on page 191 for the procedure.

Repeat the procedures described in [“Preparing to install SQL Server 2005 on the second node”](#) on page 193 and [“Installing and configuring SQL Server 2005 on the second node”](#) on page 196 on any additional nodes for the same SQL instance.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, `INST1-VS\INST1`).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.


Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service


- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver
 - Press **F5**. The right pane divides into an upper and lower panel.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, SYSTEM1\INST1. (For a default instance, you see only *System_Name*.)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
sp_dropserver "System_Name\Instance_Name"
where **System_Name\Instance_Name** is the name noted in [step 3](#).
For example, for a named instance:
sp_dropserver "SYSTEM1\INST1"
For example, for a default instance:
sp_dropserver "SYSTEM1"
 - Press F5.
 - 6 Delete the contents in the upper pane.
 - 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
sp_addserver "INST1-VS\INST1", local
For example, for a default instance:
sp_addserver "INST1-VS", local
 - Press F5.
 - 8 Exit the SQL Server Management Studio.

Creating a SQL Server user-defined database

You can use SFW HA to manage a SQL Server user-defined database.

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them first.

See [“Creating volumes”](#) on page 168.

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2005 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 Connect to the SQL Server virtual server instance for which you want to create the database (for example, `INST1-VS\INST1`).
- 3 Expand the icon associated with your server.
- 4 Right-click on **Databases** and select **New Database**.
- 5 In the New Database page, enter a name for the new database.
- 6 Click the browse button (...) in the **Path** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 7 Select and edit other file properties as desired.
- 8 Click the browse button (...) in the **Path** column for the **Transaction Log** row and browse to the location of the volume you want to create for the transaction log, and click **OK**.
- 9 Add more data files if required and configure whatever options are required for your database.
- 10 Click **OK**.
- 11 Depending on your configuration plans, you may have additional steps to complete in SQL Server.
See [“Completing configuration steps in SQL Server”](#) on page 204.
- 12 If the SQL Server service group has already been configured, you need to add the resources for the new database to the service group.
See [“Modifying a SQL 2005 service group to add VMDg and MountV resources”](#) on page 217.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Veritas Volume Replicator (VVR), Symantec recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume within the system database disk group.

See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 204.

If you are running multiple SQL Server instances, you must assign a different port to each SQL Server instance.

See “[Assigning the port for multiple instances](#)” on page 204.

Moving the tempdb database if using VVR for disaster recovery

If you plan to implement a disaster recovery configuration using VVR, Symantec recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See “[Creating volumes](#)” on page 168.

Then, refer to the Microsoft Knowledge Base for the instructions on moving the tempdb database. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 224071: How to move SQL Server databases to a new location by using Detach and Attach functions in SQL Server

Refer to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Assigning the port for multiple instances

When you are running multiple SQL Server instances, you must assign a different port to each instance.

Refer to the Microsoft Knowledge Base for the instructions on assigning static ports. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 823938: How to configure an instance of SQL Server to listen on a specific TCP port or a dynamic port

Refer to:

<http://support.microsoft.com/kb/823938/en-us>

Configuring the VCS SQL Server 2005 service group

A VCS SQL Server service group is used to bring a SQL Server 2005 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. You use the VCS SQL Configuration Wizard to configure the service group.

Read the following topics:

- [Service group requirements for active/active configurations](#)
- [Prerequisites for configuring the service group](#)
- [Creating the SQL Server 2005 service group](#)

Service group requirements for active/active configurations

Note the following requirements for active/active configurations:

- For an active/active configuration, you must create a separate service group for each instance.
- Each service group that you create must have a unique service group name and virtual IP address.
- For an active/active configuration, when you specify the priority order of systems, reverse the order for each service group so that the active system and failover system are opposite for each instance. For example, if you have two instances and two systems, you would set the priority order as follows:

INSTANCE 1	Priority order:
	SYSTEM1
	SYSTEM2
INSTANCE2	Priority order:
	SYSTEM2
	SYSTEM1

Prerequisites for configuring the service group

Complete the following tasks before configuring the service group:

- ✓ Verify that SFW HA, along with the VCS database agent for SQL Server 2005, is installed on all cluster nodes. See [“Installing Veritas Storage Foundation for Windows”](#) on page 158.
- ✓ Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See [“Configuring the cluster”](#) on page 173.
- ✓ Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- ✓ Verify that the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- ✓ To avoid having to modify the service group later, create the user-defined databases before creating the service group. See [“Creating a SQL Server user-defined database”](#) on page 203.
- ✓ Verify that the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group. See [“Importing the cluster disk group”](#) on page 194 and [“Adding drive letters to mount the volumes”](#) on page 195.
- ✓ Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. You specify the virtual server name when configuring the service group. See [“Setting the internal name of the clustered instance”](#) on page 200.

Note: For a disaster recovery configuration, the SQL Server virtual server name on the secondary site cluster must match the one on the primary site cluster.


- ✓ Assign a unique virtual IP address for the SQL Server 2005 instance. You specify this IP address when configuring the service group.
- ✓ Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists in the same location on all the cluster nodes. A sample script is supplied in C:\Program Files\Veritas\cluster server\bin\SQLServer2005\sample_script.sql. Detailed monitoring is often not necessary.
- ✓ Stop the SQL 2005 Server service for the SQL instance. See [“Stopping the SQL Server 2005 service”](#) on page 193.

Creating the SQL Server 2005 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.


To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- 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 Systems in Priority Order list.

- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment. Click **Next**.
 - 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
 - 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 208 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.

- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.
 - 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
 - 12 Configure the virtual server as follows:

SQL Server Agent Configuration Wizard
Virtual Server Configuration
Enter a virtual server name for the application and specify the virtual IP information.

✓ Prerequisites
✓ System List
✓ Instance Selection
✓ Databases
✓ Detail Monitoring
✓ Registry Replication
▶ Virtual Computer

Summary
Configuration Status

Virtual Server Name:

Virtual IP Address:

Subnet Mask:

Specify the adapter to be used on each system.

System Name	Adapter Display Name
PC177VM-3	Public
PC177VM-4	Public

Advanced Settings...

< Back Next > Cancel

- Enter the virtual name for the server, for example **INST1-VS**. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.

- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
 - Enter the subnet mask to which the virtual IP address belongs.
 - For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review 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.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
- To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.

You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.
- The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.
- To configure an MSDTC service group, see [“Configuring an MSDTC service group for high availability”](#) on page 219.

Verifying the SQL Server 2005 cluster configuration

Failover simulation is an important part of configuration testing.

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.

Determining additional steps needed

Determine the next step as follows:

- If you have no more instances of SQL Server that you plan to install and configure for high availability, your high availability configuration is complete.
- If you have additional instances of SQL Server to install and configure in the cluster, start the installation and configuration sequence again for the next instance.
 - If you are using the Solutions Configuration Center:
See [“Following the workflow in the Solutions Configuration Center”](#) on page 155.
 - If you are using the printed document as a guide:
See [“Tasks for a new HA installation of SQL Server 2005”](#) on page 136.
- If you are planning to set up a secondary site for disaster recovery:
See [Chapter 13, “Deploying disaster recovery: New SQL Server 2005 installation”](#) on page 541.
- If you want to configure a Cluster Management Console connection:
See [“Configuring the Cluster Management Console connection”](#) on page 212.
- If you need to modify the SQL Server service group:
See [“Modifying a SQL 2005 service group to add VMDg and MountV resources”](#) on page 217.
- If you need to make other changes to the cluster:
See *Veritas Cluster Server Administrator’s Guide*.

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 server

The 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:          9dfde3d9aaeb0084f8e35819c1fed7e6b01d2ae
```

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 a SQL 2005 service group to add VMDg and MountV resources

If you add a new SQL Server database after configuring the SQL Server service group, you must add the VMDg and MountV resources to the service group.

Before running the SQL Server Agent wizard to add the VMDg and MountV resources:

- ✓ Make sure the SQL Server resources are online.
- ✓ Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.

12 To complete the user database configuration, choose one of the following:

- Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.
- Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for high availability](#)” on page 219.

Configuring an MSDTC service group for high availability

This chapter includes the following topics:

- [Tasks for configuring MSDTC for high availability](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring cluster disk groups and volumes](#)
- [Mounting volumes used by the MSDTC service group](#)
- [Creating an MSDTC service group](#)
- [Creating an MSDTC client](#)

Tasks for configuring MSDTC for high availability

The VCS Database Agent for SQL, which you install during SFW HA installation for Microsoft SQL high availability, includes an MSDTC agent. The MSDTC agent can provide high availability for the Microsoft Data Transaction Coordinator (MSDTC) service. The MSDTC agent comprises two parts: MSDTC client and MSDTC server.

You can configure high availability for MSDTC either before or after configuring high availability for Microsoft SQL. To configure high availability for MSDTC, you use the SQL Server Configuration Wizard to configure an MSDTC service group for the MSDTC server and to configure the MSDTC client.

To modify an existing MSDTC service group, see the *Veritas Cluster Server Database Agent for Microsoft SQL Configuration Guide*.

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

Table 7-1 Tasks for configuring MSDTC for high availability

Objective	Tasks
“ Reviewing the prerequisites ” on page 221	✓ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 221	✓ Understanding MSDTC service group configuration ✓ Reviewing the sample configuration
“ Configuring cluster disk groups and volumes ” on page 224	✓ Configuring cluster disk groups for an MSDTC service group ✓ Configuring volumes for an MSDTC service group
“ Mounting volumes used by the MSDTC service group ” on page 230	✓ Mounting the drives used by the service group on the node on which you are configuring the service group
“ Creating an MSDTC service group ” on page 231	✓ Creating an MSDTC service group
“ Creating an MSDTC client ” on page 233	✓ Creating an MSDTC client.

Reviewing the prerequisites

You must meet the following prerequisites before creating and configuring the MSDTC service group:

- ✓ You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- ✓ You must be a Local Administrator on the node where you run the wizard.
- ✓ Verify that the VCS Database Agent for SQL is installed on all cluster nodes.
- ✓ Verify that the VCS cluster is configured using the VCS Configuration Wizard (VCW).
- ✓ Verify that the drives containing the MSDTC logs and registry replication directory are mounted on one node (the node on which you are configuring the service group) and unmounted on all other nodes.
- ✓ Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC service group.
- ✓ Assign a unique virtual server name and virtual IP address for the MSDTC server.
- ✓ Verify that the Distributed Transaction Coordinator service is stopped.

Reviewing the configuration

MSDTC servers can coexist with SQL servers on the same cluster nodes. If the MSDTC server and the SQL server are running on the same node, the MSDTC client is left in the default configuration. If the MSDTC server is not configured on the same node as the SQL server, then the MSDTC client must be configured on that node using the SQL Server Configuration Wizard. The MSDTC client and the MSDTC server must not run on the same cluster node.

For instance, a SQL Server configuration in a VCS cluster might span four nodes and two sets of shared storage. The shared storage can be managed using Veritas Storage Foundation for Windows (SFW). Two configurations are possible:

- SQL server is configured on different nodes than the MSDTC server
- SQL server is configured on the same node as the MSDTC server

Figure 7-1 MSDTC server configured on different nodes than SQL server





Figure 7-2 MSDTC server configured on the same node as SQL server



Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your MSDTC service group.

A disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the cluster disk group and volumes:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place MSDTC files and log files on separate volumes.
- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (MSDTC_DG) on shared disks and then create the following volumes:


- MSDTC_LOG: contains the MSDTC log files.
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the MSDTC service group. Create a 100 MB volume for this purpose.

Creating a cluster disk group

Create a cluster disk group on the first node of the cluster.

To create a 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 On the **Welcome** page of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group you want to create.



- Enter the name of the disk group (for example, MSDTC_DG).
- Click **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and click the **Add** button to move them to the **Selected disks** list.

- Click **Next**.
- 7 Review the selected disks and click **Next**.
- 8 Review the summary information and click **Finish**.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:


- MSDTC_LOG: contains the MSDTC log files
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the MSDTC service group.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.


To create a volume

- 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 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 the disk group you created for the MSDTC volumes and select **New Volume**.
- 5 On the **Welcome** page of the New Volume Wizard, click **Next**.


- 6 On the **Assign Disks for Volume** page, make sure the name of the disk group that you created for MSDTC appears in the Group name field.



- 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. After selecting the necessary disks, click **Next**.
- 8 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.
 - 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.
- 9 In the **Add Drive Letter and Path** page, 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.

- 10 Specify information about the attributes for the file system that will be created:



- 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**.
- 11 Review the Summary dialog box and click **Finish**.
 - 12 Repeat the procedure as necessary to create the additional volumes needed for your deployment. Create the volumes on the first node of the cluster only.

Mounting volumes used by the MSDTC service group

Ensure that the shared volumes created for the MSDTC log files and registry replication information are mounted on the node where you run the SQL Server Configuration Wizard to configure the service group. The volumes must be unmounted on all other nodes in the cluster.

If you created the MSDTC disk group and volumes on another node, mounting a volume involves deporting the disk group from the other node and then importing the disk group to the node where you are configuring the service group.

Occasionally, when a disk group is imported, a drive letter may not be associated with an existing volume. If this occurs, use VEA to add a drive letter or folder path and mount the volume so that it can be seen by the operating system.

To mount volumes

- 1 If the disk group is not imported, deport it from the other node and then import it on the current node:
 - Stop all processes accessing the volumes in the disk group, and from the VEA console's tree view, right-click the disk group and select **Deport Dynamic Disk Group**. Then select **Yes**.
 - In the VEA, connect to the node where you want to import the cluster disk group. In the tree, right-click **Storage Agent**, and click **Rescan** to update the disk information on the node. Then right-click the disk group and select **Import Dynamic Group** and click **OK**.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.
- 3 In the Drive Letter and Paths dialog box, select **Add**.
- 4 Select one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder. Assign the same drive letter or mount path that was assigned when the volume was created
 - *To assign a drive letter*
In the Assign Drive Letter dialog box, select the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder option** and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.
- 6 Repeat [step 2](#) through [step 5](#) for all the volumes to be mounted.


You can now create the MSDTC service group.

Creating an MSDTC service group

MSDTC is a global resource and can be accessed by more than one SQL Server service group. Symantec recommends that you configure one MSDTC service group in a VCS cluster.

To configure an MSDTC service group

- 1 If you have just configured a SQL service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Select **MSDTC Server - Service Group Configuration** and **Create**. Click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 Specify the service group name and system list.




- Enter a name for MSDTC 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. Make sure you select the

systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.

- To change a system’s priority in the **Systems in Priority Order** list, select the system and click the up and down arrows. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
- Click **Next**. If the configuration is in read-only mode, the wizard prompts you before changing it to read-write mode. The wizard starts validating your configuration. Various messages indicate the validation status.

5 Specify the information related to the virtual server.



- Enter a virtual server name for the node on which the DTC service is running. Ensure that the virtual server name you enter is unique in the cluster.
- Enter a unique virtual IP address for the MSDTC server.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Click the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.

- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop down list.
 - Click **Next**.
- 6 Specify the MSDTC log and replication directory and click **Next**. Symantec recommends using different paths for these directories. If the directory does not exist, the wizard creates it.
 - 7 Review the service group configuration, change the resource names if desired and click **Next** to create the service group.
 - 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 the resources, if desired. 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**.
 - 8 A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes** to create the service group.
 Various messages indicate the status of these commands.
 - 9 In the Configuration Complete panel, check **Bring the service group online** to bring the configured service group online. To bring the service group online later, uncheck the option.
 - 10 Click **Next** to create an MSDTC client, or click **Finish** to exit the wizard.

Creating an MSDTC client

Set the MSDTC client to run on nodes where:

- A SQL instance is configured to run
- The MSDTC server is not configured to run.

Before running the wizard to configure an MSDTC client:


- ✓ Verify that the MSDTC service group is online in the cluster.
- ✓ Verify that the node on which you run the wizard is not a part of an MSDTC service group system list.

To configure an MSDTC client

- 1 If you have just configured the MSDTC service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL

Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.

- 2 In the Select Configuration Option panel, select **MSDTC Client Configuration** and click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 In the System Selection panel, specify the nodes on which the MSDTC client will be configured and click **Next**:



- Select the nodes in the **Available Cluster Systems** list. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
 - Click the right arrow to add them to the **Selected systems** list.
 - The **Unavailable Cluster Systems** lists the nodes that have an MSDTC service group configured and are therefore not available for setting up an MSDTC client.
- 5 If the MSDTC service group is not online in the cluster, an informational message appears informing you that the wizard will bring the MSDTC service group online. Click **Yes**.
 - 6 On the Specify DTC Node panel, specify the virtual DTC server name and click **Next**.
 - 7 On the Configuration Complete panel, click **Finish** to exit the wizard.

Replicated Data Clusters

This section includes the following chapters:

- [About Replicated Data Clusters](#)
- [Configuring Replicated Data Clusters for SQL 2000](#)
- [Configuring Replicated Data Clusters for SQL 2005](#)

About Replicated Data Clusters

This chapter includes the following topics:

- [“About Replicated Data Clusters”](#) on page 237
- [“How VCS Replicated Data Clusters work”](#) on page 239
- [“Setting up a Replicated Data Cluster configuration”](#) on page 240
- [“Migrating the service group”](#) on page 242

About Replicated Data Clusters

A Replicated Data Cluster (RDC) uses data replication, instead of shared storage, to assure data access to all the nodes in a cluster.

The Replicated Data Cluster configuration provides both local high availability and disaster recovery functionality in a single VCS cluster. You can set up RDC in a VCS environment using Veritas Volume Replicator (VVR.)

An RDC exists within a single VCS cluster with a primary zone and a secondary zone, which can stretch over two buildings or data centers connected with Ethernet. In an RDC configuration, if an application or a system fails, the application is failed over to another system within the current primary zone. If the entire primary zone fails, the application is migrated to a system in the secondary zone (which then becomes the new primary).

For VVR replication to occur, the disk groups containing the Replicated Volume Group (RVG) must be imported at the primary and secondary zones. The replication service group must be online at both zones simultaneously, and must be configured as a hybrid VCS service group.

The SQL Server service group is configured as a failover service group. The SQL Server service group must be configured with an online local hard dependency on the replication service group.

Note: VVR supports multiple replication secondary targets for any given primary. However, RDC for VCS supports only one replication secondary for a primary.

An RDC configuration is appropriate in situations where dual dedicated LLT links are available between the primary zone and the secondary zone but lacks shared storage or SAN interconnect between the primary and secondary data centers. In an RDC, data replication technology is employed to provide node access to data in a remote zone.


You must use dual dedicated LLT links between the replicated nodes.

How VCS Replicated Data Clusters work

To understand how a RDC configuration works, let us take the example of SQL 2000 configured in a VCS replicated data cluster. The configuration has two system zones:

- Primary zone (zone 0) comprising nodes located at the primary site and attached to the primary storage
- Secondary zone (zone 1) comprising nodes located at the secondary site and attached to the secondary storage

SQL 2000 is installed and configured on all nodes in the cluster. The SQL 2000 data is located on shared disks within each RDC zone and is replicated across RDC zones to ensure data concurrency. The SQL Server service group is online on a system in the current primary zone and is configured to fail over in the cluster.



In the event of a system or SQL 2000 failure, VCS attempts to fail over the SQL Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone (zone 1). VCS also redirects clients once the application is online on the new location.

Setting up a Replicated Data Cluster configuration

In the example, SQL Server is configured as a VCS service group in a four-node cluster, with two nodes in the primary RDC zone and two in the secondary RDC zone. In the event of a failure on the primary node, VCS can fail over the SQL 2000 instance to the second node in the primary zone.

The process involves the following steps:

- [Setting up replication](#)
- [Configuring the service groups](#)

Setting up replication


Set up replication between the shared disk groups. Use VVR to group the shared data volumes into a Replicated Volume Group, and creating the VVR Secondary on hosts in your secondary zone.

Create a Replicated Data Set (RDS) with the Primary RVG consisting of the shared volumes between the nodes in the first zone and Secondary RVG consisting of shared volumes between nodes in the second zone. Therefore, use the same Disk Group and RVG name in both zones so that the MountV resources will mount the same block devices.


Configuring the service groups

For a successful wide-area failover, the mount points and applications must fail over to the secondary RDC zone. Additionally, the VVR secondary disk group and RVG must be imported and started on the secondary RDC zone.

The following screen from the VCS Cluster Manager (Java Console) depicts a typical SQL Server service group RDC configuration:



The following screen from the VCS Cluster Manager (Java Console) depicts a typical SQL replication service group (RVG) configuration:



Migrating the service group

In the RDC configuration, consider a case where the primary RDC zone suffers a total failure of the shared storage. In this situation, none of the nodes in the primary zone see any device.

The service group cannot fail over locally within the primary RDC zone, because the shared volumes cannot be mounted on any node. So, the service group must fail over, to a node in the current secondary RDC zone.

The RVGPrimary agent ensures that VVR volumes are made writable. The application can be started at the secondary RDC zone and run there until the problem with the local storage is corrected.

If the storage problem is corrected, you can switch the application back to the primary zone using VCS.

To switch the service group

Before switching the application back to the original primary RDC zone, you must resynchronize any changed data from the active secondary RDC zone since the failover. Once the resynchronization completes, switch the service group to the primary zone.

In the **Service Groups** tab of the of the Cluster Explorer configuration tree, right-click the service group. Click **Switch To** and select the system in the primary RDC zone to switch to and click OK.

Configuring Replicated Data Clusters for SQL 2000

This chapter includes the following topics:

- [“Tasks for configuring Replicated Data Clusters for SQL 2000”](#) on page 246
- [“Reviewing the prerequisites”](#) on page 250
- [“Reviewing the configuration”](#) on page 255
- [“Configuring the storage hardware and network”](#) on page 256
- [“Installing Veritas Storage Foundation for Windows”](#) on page 258
- [“Configuring the cluster”](#) on page 265
- [“Configuring cluster disk groups and volumes”](#) on page 280
- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 288
- [“Preparing to install SQL Server on the second node”](#) on page 291
- [“Installing SQL Server 2000 on the second node”](#) on page 295
- [“Setting the internal name of the clustered instance”](#) on page 299
- [“Configuring the VCS SQL Server service group”](#) on page 301
- [“Creating the primary system zone”](#) on page 306
- [“Registering the virtual server in the SQL Server”](#) on page 307
- [“Creating a SQL Server user-defined database”](#) on page 309
- [“Verifying the installation in the primary zone”](#) on page 312
- [“Creating a parallel environment in the secondary zone”](#) on page 313

- [“Adding the IP addresses for the primary and secondary zones”](#) on page 314
- [“Adding the systems in the secondary zone to the cluster”](#) on page 314
- [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 321
- [“Configuring a hybrid RVG service group for replication”](#) on page 329
- [“Setting a dependency between the service groups”](#) on page 343
- [“Adding the nodes from the secondary zone to the RDC”](#) on page 344
- [“Verifying the RDC configuration”](#) on page 351

Tasks for configuring Replicated Data Clusters for SQL 2000

Configure the high availability and SQL components on the primary and secondary zones. Then complete the Replicated Data Set solution by configuring the components for both zones.

Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details on VVR.

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

Table 9-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Reviewing the prerequisites” on page 250	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 255	✓ Understanding active/passive configuration and zone failover in a RDC environment
	✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 256	✓ Setting up the storage hardware for a cluster environment
	✓ Verifying the DNS entries for the systems on which SQL will be installed

Table 9-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Installing Veritas Storage Foundation for Windows” on page 258	✓ Verifying the driver signing option for the system ✓ Installing Veritas Storage Foundation for Windows HA (automatic installation) ✓ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ✓ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server ✓ Configuring VxSAS
“Configuring the cluster” on page 265	✓ Verifying static IP addresses and name resolution configured for each node ✓ Configuring cluster components using the Veritas Cluster Server Configuration Wizard ✓ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 280	✓ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ✓ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2000 on the first node” on page 288	✓ Installing and configuring SQL Server 2000 ✓ Configuring SQL services

Table 9-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Preparing to install SQL Server on the second node” on page 291	✓ Stopping the SQL Service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Adding drive letters ✓ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2000 on the second node” on page 295	✓ Installing and configuring SQL Server 2000
“Setting the internal name of the clustered instance” on page 299	✓ Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 301	✓ Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating the primary system zone” on page 306	✓ Creating the primary system zone ✓ Adding the nodes to the primary zone
“Registering the virtual server in the SQL Server” on page 307	✓ Registering the virtual server in the SQL Server
“Creating a SQL Server user-defined database” on page 309	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server ✓ Adding resources for a user-defined database in VCS

Table 9-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Verifying the installation in the primary zone” on page 312	✓ Simulating failover ✓ Switching online nodes
“Creating a parallel environment in the secondary zone” on page 313	✓ Reviewing the prerequisites ✓ Reviewing the configuration ✓ Configuring the network and storage ✓ Installing SFW HA ✓ Configuring disk groups and volumes for SQL
“Setting up the Replicated Data Sets (RDS)” on page 321	✓ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones
“Configuring a hybrid RVG service group for replication” on page 329	✓ Creating a hybrid Replicated Volume Group (RVG) service group ✓ Configuring the hybrid RVG service group
“Setting a dependency between the service groups” on page 343	✓ Setting up a dependency from the VVR RVG Service Group to the SQL Server Service Group
“Adding the nodes from the secondary zone to the RDC” on page 344	✓ Adding the nodes from the secondary zone to the RVG service group ✓ Configuring the IP resources for failover ✓ Adding the nodes from the secondary zone to the SQL Server service group
“Verifying the RDC configuration” on page 351	✓ Verifying that failover occurs first within zones and then from the primary to the secondary zone

Reviewing the prerequisites

Verify that the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary zone and a secondary zone.

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 Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the following options:
Veritas Volume Replicator and the Global Clustering Option
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with
Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 required for all editions)
or
Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition
or
Windows 2000 Server, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server (all require SP4 with Update Rollup 1)
or
 - Microsoft SQL Server 2000 (64-bit) Enterprise Edition
with
Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required for both)
or
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)

with

Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T):
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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 253.
- 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 SQL 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 *VCS 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.
- 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

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 doing the following tasks:

- 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.

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 9-2](#) estimates disk space requirements for SFW HA.

Table 9-2 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

Reviewing the configuration

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

- ✓ SQL virtual server: the IP address should be the same on all nodes at the primary and secondary zones
- ✓ Replication IP address for the primary zone
- ✓ Replication IP address for the secondary zone

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

Sample configuration

The sample setup has four servers, two for the primary zone and two for the secondary zone. The nodes will form two separate clusters, one at the primary zone and one at the secondary zone.

The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary zone

SYSTEM1 & SYSTEM2	First and second nodes of the primary zone
INST1_SG	Microsoft SQL Server 2000 service group
INST1-VS	Virtual SQL Server cluster
INST1	SQL Instance Name
INST1_DG	Cluster disk group for SQL system database and files
INST1_DATA_FILES	Volume for Microsoft SQL Server system data files
INST1_REGREP_VOL	Volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	Replicator log volume required by VVR
INST1_DB1_DG	Cluster disk group for SQL Server user-defined database and files
INST1_DB1_VOL	Volume for storing a Microsoft SQL Server user-defined database

INST1_DB1_LOG	Volume for storing a Microsoft SQL Server user-defined database log file
INST1_DB1_REPLOG	Replicator log volume required by VVR for SQL user-defined database

Secondary zone

SYSTEM3 & SYSTEM4 First and second nodes of the secondary zone

All the other parameters are the same as on the primary zone.

RDS and VVR
Components

INST1_RDS	RDS name for SQL system database and files
INST1_RVG	RVG name for SQL system database and files
INST1_RVG_SG	Replication service group for SQL system database and files
INST1_DB1_RDS	RDS name for SQL Server user-defined database and files
INST1_DB1_RVG	RVG name for SQL Server user-defined database and files
INST1_DB1_RVG_SG	Replication service group for SQL Server user-defined database and files

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 Veritas Storage Foundation 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. You must select the options to install VVR and the Veritas Cluster Server Database Agent for SQL. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation.

Setting Windows driver signing options

When installing on remote or local systems running Windows Server 2003, you must set the Windows driver signing options to allow the installation.

To change the driver signing options on each local system

- 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, note the current setting, and select **Ignore** or one of the other options from the table, to allow installation to proceed.
- 4 Click **OK**.
- 5 Repeat for each computer.

If you do not change these options, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing options to their previous states.

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**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Data- base Agent for SQL	Required to configure high availability for SQL Server.
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability.
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 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 `vxsasscfg.exe` from the command prompt of the required machine.

The Welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
Password	Specify a password in the Password field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same 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

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 that name resolution is configured for each node.
- ✓ Set the required privileges:
 - 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.

Note: Add only systems in the primary zone (zone 0) to the cluster at this time.

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 267.

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 267. 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 267.


- 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 268, proceed to the next step. Otherwise, proceed to [step 11](#) on page 270.

- 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 266, 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:

The screenshot shows the 'Veritas Cluster Server Configuration Wizard' window. The title bar reads 'Veritas Cluster Server Configuration Wizard'. The main window has a left-hand pane with a tree view containing the following items: 'Domain Selection', 'VCW Options', 'Cluster Details', 'Cluster Selection', 'Validate Systems', 'Edit Options', 'NIC Selection', 'Service Account', 'Security', 'Summary', and 'Finish'. The 'Security' item is currently selected and highlighted. The main area of the wizard is titled 'Configure Security Service Option' with the subtitle 'Select security options for the cluster'. It contains two radio button options. The first option is 'Use Single Sign-on' with the description 'Configure the Symantec Product Authentication Service for secure communication in the cluster.' The second option is 'Use VCS User Privileges' with the description '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.' The 'Use VCS User Privileges' option is selected. Below this description are three text input fields: 'User Name:' (containing 'admin'), 'Password:' (containing 'password'), and 'Re-enter Password:' (containing 'password'). At the bottom right of the window are three buttons: 'Back', 'Next', and '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 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 276.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 277.

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 277. 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.

Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the disk group are shared and are available from all nodes in each zone. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

- INST1_REPLOG contains the VVR Replicator Log.
You can create this volume later while setting up the Replicated Data Sets.
See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 321.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and files.

The following disk group and volumes may be created now or later in the configuration process.

- INST1_DB1_DG is the disk group for the SQL Server user-defined database and files
- INST1_DB1_VOL contains the user database files
- INST1_DB1_LOG contains the user database log files
- INST1_DB1_REPLOG contains the VVR Storage Replicator Log.

Warning: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.


Creating a cluster disk group

Create a cluster disk group on the first node of the cluster. You can repeat the process to create the disk group for the SQL Server user-defined database and files at this time or later.

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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes in the cluster disk group (INST1_DG) for the system files on the first node of the cluster:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)

- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_REPLOG** contains the VVR Storage Replicator Log.
You can create the **INST1_REPLOG** volume at this time or during the process of creating Replicated Data Sets.
See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 321.

You can create the **INST1_DB1_VOL** and **INST1_DB1_LOG** volumes at this time or when you create a SQL user-defined database.

See “[Creating a SQL Server user-defined database](#)” on page 309.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and the log files. The following volumes in the disk group (**INST1_DB1_DG**) for the SQL Server user-defined database and files may be created now or later in the configuration process:

- **INST1_DB1_VOL** contains the user database files
- **INST1_DB1_LOG** contains the user database log files
- **INST1_DB1_REPLOG** contains the VVR Storage Replicator Log.

Warning: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.


Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.

You can right-click the disk group you have just created, for example INST1_DG.

- 5
- At the New Volume wizard opening screen, click **Next**.
- 6
- Select the disks for the volume.




-
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
-
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check Disable Track Alignment to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
-
- Click **Next**.

7 Specify the parameters of the volume.

- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the **Mirror Info** area, select the appropriate mirroring options.
 - Verify that **Enable Logging** is not selected.
 - Click **Next**.
- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on


all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
 - For the Replicator Log volume only: Clear the Format this volume check box.
 - Click **NTFS**.
 - Select an allocation size or accept the Default.
 - The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 11 Click **Finish** to create the new volume.**
- 12 Repeat these steps to create additional volumes.**

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Installing and configuring SQL Server 2000 on the first node

Before installing Microsoft SQL Server 2000, ensure the cluster disk group is imported to the first node and the volumes are mounted.

Complete the following procedures to install and configure Microsoft SQL Server 2000:

- ✓ [Installing Microsoft SQL Server](#)
- ✓ [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters). See “[Importing the cluster disk group](#)” on page 293 and “[Adding drive letters to mount the volumes](#)” on page 293.

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.


To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.

- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.




- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

- 10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
 - b In the Services box, select the **SQL Server** option.
 - c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
 - d If you chose to run the service in the context of a domain user, specify the information for the user.
 - e Clear the **Auto Start Service** option.
 - f Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g Click **Next**.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
- 13 If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.


Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- a Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - b Select a service from the **Services** list.
 - c Clear the **Auto-start service when OS starts** check box.
- 2 Repeat [step b](#) and [step c](#) for all other SQL Server services that are running on the server.

Preparing to install SQL Server on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2000 service”](#) on page 292
- [“Deporting the cluster disk group”](#) on page 292
- [“Importing the cluster disk group”](#) on page 293
- [“Adding drive letters to mount the volumes”](#) on page 293
- [“Renaming shared SQL Server 2000 files”](#) on page 295


Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2000 service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.

- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) on the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.

In the **Import Dynamic Disk Group** dialog box, click **OK**.


Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2000 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 295
- [“Removing shared SQL Server files”](#) on page 298

Installing SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 293 and [“Adding drive letters to mount the volumes”](#) on page 293.


Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.


Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000


- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
- b In the Services box, select the **SQL Server** option.
- c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
- d If you chose to run the service in the context of a domain user, specify the information for the user.

- e** Clear the **Auto Start Service** option.
 - f** Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g** Click **Next**.
- 11** Follow the wizard instructions to complete the installation.
 - 12** When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
 - 13** If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server on the second node](#)” on page 291 and “[Installing SQL Server 2000 on the second node](#)” on page 295 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1** On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2** Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.


Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.


To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example *SYSTEM2 \ INST1*.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text `"sp_helpserver"`
 - Press F5.
 - Make note of the name listed in the lower pane, for example `SYSTEM2\INST1`. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
`"sp_dropserver 'System_Name\Instance_Name.'"`
where *System_Name\Instance_Name* is the name noted in step 3.
For example, for named instance:
`"sp_dropserver 'SYSTEM2\INST1.'"`
For example, for a default instance:
`"sp_dropserver 'SYSTEM1.'"`
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
- In the top pane, enter
`"sp_addserver 'Virtual_Server_Name\Instance_Name', local"`
For example `'INST1-VS\INST1'`, `local` for a named instance, or `'INST1-VS'`, `local` for a default instance.

- Press F5.

Configuring the VCS SQL Server service group

The SQL Server Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.


Complete the following tasks before configuring the service group:

- ✓ Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2000, is installed on all cluster nodes. See “[Installing Veritas Storage Foundation for Windows](#)” on page 258.
- ✓ Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See “[Configuring the cluster](#)” on page 265.
- ✓ Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- ✓ Verify that the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- ✓ Verify that the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- ✓ Assign a unique virtual server name to SQL Server 2000. While it must be unique within the cluster, it must be the same as the virtual name of the SQL Server in the primary zone.
- ✓ Assign a unique virtual IP address to the SQL Server 2000 instance.

To create a SQL Server service group on the cluster


- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.

5 Specify the service group name and system list:




- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
 - 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 Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment. Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 302 are available for selection.
 - Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
 - If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.

- 11
- In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12
- Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
- Click **Next**.

- 13 In the Service Group Summary, review 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.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.


The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

Creating the primary system zone

In the service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 Select the SQL Server service group (INST1_SG) in the left pane and the Properties tab in the right pane.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.




- 7 Click **OK**.

Registering the virtual server in the SQL Server

In the SQL Server Enterprise Manager, register the VCS virtual server (INST1-VS) and delete the physical server name.


To register the VCS virtual server

- 1 Click **Start > All Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 In the right pane, expand the Microsoft SQL Servers.
- 3 Right-click the **SQL Server Group** and select **New SQL Server Registration**.
- 4 On the Welcome screen, click **Next**.
- 5 In the **Select a SQL Server** window, select the virtual server\SQL instance (for example INST1-VS\INST1).



- 6 Click the **Add** button.
- 7 Click **Next**.
- 8 In the **Select an Authentication Mode** window, select an authentication mode and click **Next**.

- 9 In the **Select SQL Server Group** window, select Create a new top-level SQL Server group, or leave it in the existing group. If a new SQL Server group is created, enter the **Group name**.



- 10 Click **Next**.
- 11 In the **Completing the Register the SQL Server Wizard** window, click **Finish**.
- 12 Wait for the **Register SQL Server Messages** window that shows the server registration completed successfully and click **Close**.

To delete the physical server name

- 1 In the Microsoft SQL Server **Enterprise Manager**, expand the Microsoft SQL Servers and the **SQL Server Group**.
- 2 Right-click the physical name of the SQL Server and instance (for example SYSTEM1\INST1) and select **Delete SQL Server Registration**.
- 3 Click **Yes** in the warning message window.
- 4 Verify the SQL servers in the Microsoft SQL Server **Enterprise Manager** window. In the left pane, the **Microsoft SQL Servers** should show the virtual server under **SFWHA Clustered SQL Servers** and no items under the **SQL Server Group**.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- ✓ Create volumes for a user-defined SQL Server database and its transaction log.
- ✓ Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- ✓ Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 282 for information on how to use the VEA console to create a volume.

Note: Best practice is to create a separate disk group with separate volumes for the user-defined database files.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.

- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.

Adding VMDg and MountV resources

Before running the SQL Server Agent wizard to add the VMDg and MountV resources:

- ✓ Make sure the SQL Server resources are online.
- ✓ Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

- Click **Next** to configure another SQL service group or an MSDTC service group.

Verifying the installation in the primary zone

Failover simulation is an important part of configuration testing.

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.

Creating a parallel environment in the secondary zone

After setting up a SFW HA environment in the primary zone, use the guidelines in the following list to complete the same tasks in the secondary zone.

Note: Before you begin to configure the secondary zone, offline the SQL Server resource, the SQL virtual server name resource, and the SQL virtual IP resource in the primary zone. The remaining resources should be online, including the VMDg resource. In VEA, make sure to remove all the drive letters from the configured volumes, to avoid conflicts when configuring the zones.

- ✓ [“Reviewing the prerequisites”](#) on page 250
- ✓ [“Reviewing the configuration”](#) on page 255
- ✓ [“Configuring the storage hardware and network”](#) on page 256
- ✓ [“Adding the IP addresses for the primary and secondary zones”](#) on page 314
- ✓ [“Adding the systems in the secondary zone to the cluster”](#) on page 314
- ✓ [“Configuring cluster disk groups and volumes”](#) on page 280

During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as the cluster at the primary zone:

 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- ✓ [“Installing and configuring SQL Server 2000 on the first node”](#) on page 288

Select the same options in the secondary zone as you did for the primary zone.
- ✓ [“Preparing to install SQL Server on the second node”](#) on page 291

The instance name must be the same in the primary zone and secondary zone.
- ✓ [“Installing SQL Server 2000 on the second node”](#) on page 295

Note: After you install SQL Server 2000 on the nodes in the secondary zone, make sure to use VEA to remove all the drive letters from the configured volumes, to avoid conflicts during the configuration of the zones.

Adding the IP addresses for the primary and secondary zones

Create IP addresses for both the primary and secondary zones. Connect to the primary system in each zone and add the IP address for the zone.

To create an IP address

- 1 On the primary system in the primary zone, open the Network Connections (Click **Start > Settings > Network Connections**).
- 2 Right-click the public network connection to be used and click **Properties**.
- 3 Select **Internet Protocol (TCP/IP)** and click **Properties**.
- 4 Click **Advanced**.
- 5 In the **IP Settings** tab, click **Add** to add a virtual IP address and subnet mask.
- 6 In the **TCP/IP Address** dialog box, enter a new IP address and a subnet mask and click **OK**, several times, until the **Network Connections** dialog box is closed.
- 7 Repeat these steps on the primary system in the secondary zone.


Adding the systems in the secondary zone to the cluster

Add the nodes in the secondary zone to the existing cluster with the following procedure.

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 panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

- 4 In 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 317.

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 316. 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 317.


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, type 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 the VCS private network communication, on 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.

Return to the task list [“Creating a parallel environment in the secondary zone”](#) on page 313.


Setting up the Replicated Data Sets (RDS)

Set up the Replicated Data Sets (RDS) in the primary zone and secondary zone. You can configure an RDS using the Create RDS wizard for both zones.

- ✓ Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- ✓ Verify that the cluster disk group is imported in the primary and secondary zone

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 (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name	Enter the name for the volume in the Name field.
Size	Enter a size for the volume in the Size field.
Layout	Select the desired volume layout.
Disk Selection	■ Choose Select disks automatically if you want VVR to select the disks for the Replicator Log. ■ Choose Select disks manually to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select Add to move the disks into the selected disks pane.

- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 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.

Add Secondary Wizard
Edit replication settings
 Edit replication settings or click next.

Primary side IP: 10.182.145.127
 Secondary side IP: 10.182.146.116
 Replication Mode: Synchronous Override
 Replicator Log Protection: AutoDCM

Advanced

DHCP addresses are not supported by VVR.

< Back Next > Cancel Help


- 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.

AutoDCM	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.
DCM	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.
Off	Disables Replicator Log overflow protection.
Override	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. 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.
Fail	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.

If the SQL Server user-defined database and files are in a separate disk group from the SQL Server system files, repeat the procedure “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 321 for the disk group (INST1_DB1_DG) that contains the SQL Server user-defined database and files. Provide unique

names for the Replicated Data Set name, and the Replicated Volume Group name. See “[Sample configuration](#)” on page 255 for a list of example names.

Configuring a hybrid RVG service group for replication

Create and configure a hybrid Replicated Volume Group (RVG) service group for replication.

The RVG service group is hybrid because it behaves as a failover service group within a zone and as a parallel service group between zones.

For additional information about service group types, see the *Veritas Cluster Server 5.0 Administrator’s Guide*.

Configure the RVG service group’s resources manually by copying and modifying components of the SQL Server service group. Then create new RVG resources and bring them online.

The RVG service group for RDC contains the following resources:

Table 9-3 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg for the first disk group	Volume Manager disk group with SQL system files
VvrRvg for the first disk group	Replicated volume group with SQL system files
VMDg for the second disk group	Volume Manager disk group with SQL user-defined files
VvrRvg for the second disk group	Replicated volume group with SQL user-defined files


Creating the RVG service group

Create a hybrid replicated volume (RVG) service group, to contain the resources for replication.

To create a hybrid RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.

- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.
- 3 In the **Add Service Group** window:



- a Enter a name for the service group, for example INST1_RVG_SG.
- b Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- c Select **Hybrid**.
- d Click **OK**.

Configuring the RVG service group for RDC replication

Configure the RVG service group's resources manually for RVG by completing the following tasks:

- [“Configuring the IP and NIC resources”](#) on page 331
Copy IP and NIC resources of the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
- [“Configuring the VMDg resources for the first disk group”](#) on page 332 and [“Configuring the VMDg resources for the second disk group”](#) on page 334
Copy the VMDg resources for all disk groups in the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
- [“Adding the VVR RVG resources for the disk groups”](#) on page 335
Create the VVR RVG resources for all the disk groups and enter the attributes for each of the disk groups and the replication IP address.
- [“Linking the VVR RVG resources to establish dependencies”](#) on page 338
Link the VVR RVG resources to establish the dependencies between the VMDg resources, the IP resource for replication, and the VVR RVG

resources for the disk groups. Configure the RVG service group’s VMDg resources to point to the disk groups that contain the RVGs.

- “[Deleting the VMDg resource from the SQL Server service group](#)” on page 339
Delete the VMDg resources from the SQL Server service group, because they depend on the replication and were configured in the RVG service group.

Configuring the IP and NIC resources


Configure the following resources and attributes for the IP and NIC:

Table 9-4 IP and NIC resources

Resource	Attributes to Modify
IP	Address
NIC	(none)

To create the IP resource and NIC resource

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the IP resource (INST1_SG-IP), and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the names of the IP and NIC resources for the RVG service group and click **OK**.



To modify the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.
- 2 In the **Properties View** window, for the **Address** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, enter the VVR IP address for the Primary Zone as the scalar value.
- 4 Close the **Properties View** window.

To enable the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (INST1_RVG_SG-IP) and select **Enabled**.
- 2 In the **Resources** tab display area, right-click the NIC resource (INST1_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg resources for the first disk group

Create the VMDg resource in the SQL Server service group, and clear the DGGuid attribute for the new VMDg.

Configure the following attributes in the SQL Server service group for the MountV resource:

Table 9-5 MountV resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
MountV (for the SQL Server system volume)	VMDg Resource Name Volume Name
MountV (for the registry volume)	VMDg Resource Name Volume Name
Resources for the disk group for the SQL user-defined database files:	
MountV (for the SQL Server user-defined database log)	VMDg Resource Name Volume Name
MountV (for the SQL Server user-defined database)	VMDg Resource Name Volume Name

To create the VMDg resource for the first disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the first disk group, with the SQL system files (INST1_SG-VMDg), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg.
- 6 Click **OK**.

To clear the DGGuid attribute for the new VMDg

- 1 In the **Resources** tab display area, right-click the new VMDg resource.
- 2 In the same **Properties View** window, for the **DGGuid** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, clear the scalar value for the **DGGuid** attribute.
- 4 Close the **Properties View** window.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server system data files (INST1_SG-MountV) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server system data files (INST1_DATA_FILES).
- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the registry volume (INST1_SG-MountV-1) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the **Volume Name** attribute is the registry volume (INST1_REGREP_VOL).

- 9 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11 Close the **Properties View** window.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg) and select **Enabled**.

Configuring the VMDg resources for the second disk group

Repeat the VMDg and MountV configuration for the second disk group.

To create the VMDg resource for the second disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the second disk group, with SQL user-defined files (INST1_SG-VMDg-1), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg-1.
- 6 Click **OK**.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined log (INST1_SG-MountV-2) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined log (INST1_DB1_LOG).

- 4
- In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5
- In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg-1.
- 6
- Close the **Properties View** window.
- 7
- In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined database (INST1_SG-MountV-3) and select **View > Properties View**.
- 8
- In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined database (INST1_DB1_VOL).
- 9
- In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 10
- In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11
- Close the **Properties View** window.

To enable the VMDg resource

- 1
- In the left pane, select the RVG service group (INST1_RVG_SG).
- 2
- In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg-1) and select **Enabled**.

Adding the VVR RVG resources for the disk groups


Add VVR RVG resources for replication of the disk groups.
Configure the following attributes in the RVG service group for the VvrRvg resource:

Table 9-6 VvrRvg resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
VvrRvg	VMDgResName IPResName
Resources for the disk group for the SQL user-defined database files:	
VvrRvg	VMDgResName IPResName


To create the VVR RVG resource for the first disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



- a Enter the **Resource Name** for the VVR RVG resource.
- b Select the **Resource Type** of VvrRvg.
- 3 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the RVG group that is being managed, for example INST1_RVG.
- 5 Click **OK**.
- 6 In the **Add Resource** window, for the **VMDGResName** attribute, click **Edit**.
- 7 In the **Edit Attribute** window, enter the name of the disk group containing the RVG, for example INST1_RVG_SG-VMDg.
- 8 Click **OK**.
- 9 In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, enter the name of the IP resource managing the IP address for replication, for example INST1_RVG_SG-IP.
- 11 Click **OK**.

- 12 In the **Add Resource** window, verify that the attributes have been modified:



Add Resource

Group name: INST1_RVG_SG

Resource name: INST1_VvrRvg

Resource Type: VvrRvg

Attribute name	Type	Dimension	Value	Edit
RVG	String	Scalar	INST1_RVG	[Edit]
VMDgResName	String	Scalar	INST1_RVG_SG-...	[Edit]
IPResName	String	Scalar	INST1_RVG_SG-IP	[Edit]
SRL	String	Scalar		[Edit]
RLinks	String	Vector		[Edit]


☒ Critical ☒ Enabled

Show Command OK Cancel

- 13 Click **OK**.

To create the VVR RVG resource for the second disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



Add Resource

Group name: INST1_RVG_SG

Resource name: INST1_DB1_VvrRvg

Resource Type: [Dropdown menu open showing: SQLAgService2005, SQLOlapService2005, SQLServer2000, SQLServer2005, ServiceMonitor, VMDg, VRTSWebApp, VvrRvg]


Attribute name	Type	Dimension	Value	Edit

☒ Critical ☐ Enabled

Show Command OK Cancel

- a Enter the **Resource Name** for the VVR RVG resource for the second disk group.
- b Select the **Resource Type** of VvrRvg.
- 3 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the RVG group that is being managed, for example INST1_DB1_RVG.
- 5 Click **OK**.
- 6 In the **Add Resource** window, for the **VMDGResName** attribute, click **Edit**.

- 7
- In the **Edit Attribute** window, enter the name of disk group containing the RVG, for example INST1_RVG_SG-VMDg-1.
- 8
- Click **OK**.
- 9
- In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10
- In the **Edit Attribute** window, enter the name IP resource managing the IP address for replication, for example INST1_RVG_SG-IP. In this example both disk groups are using the same IP resource for replication.
- 11
- Click **OK**.
- 12
- In the **Add Resource** window, verify that the attributes have been modified:



- 13
- Click **OK**.

Linking the VVR RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the VVR RVG service group to establish the dependencies between the resources. Start from the top parent and link the following resources:

Table 9-7 Dependencies for VVR RVG resources for RDC


Parent	Child
Resources for the disk group for the SQL system files:	
INST1_ VvrRvg	The IP for replication, for example INST1_RVG_SG-IP
INST1_ VvrRvg	The VMDg for the SQL system files, for example INST1_RVG_SG-VMDg
Resources for the disk group for the SQL user-defined database files:	
INST1_DB1_ VvrRvg	The IP for replication, for example INST1_RVG_SG-IP

Table 9-7 Dependencies for VVR RVG resources for RDC

Parent	Child
INST1_DB1_VvrRvg	The VMDg for the SQL user-defined database files, for example INST1_RVG_SG-VMDg-1

To link the VVR RVG resources

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example INST1_DB1_VvrRvg.
- 4 Click the child resource, for example INST1_RVG_SG-IP.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG resources:



Deleting the VMDg resource from the SQL Server service group

The VMDg resources must now be manually deleted from the SQL Server service group, because they depend on replication and were configured in the RVG service group.

To delete the VMDg Resources from the SQL Server service group

- 1
- In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) from the left pane.
- 2
- In the **Resources** tab display area, right-click the VMDg resource for the first disk group (INST1_SG-VMDg) and select **Delete**.
- 3
- Click **Yes** to confirm that you want to delete it (even if it is online).
- 4
- In the **Resources** tab display area, right-click the VMDg resource for the second disk group (INST1_SG-VMDg-1) and select **Delete**.
- 5
- Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

Add resources of type RVG Primary to the SQL Server service group for each of the SQL Server disk groups (system and user-defined) and configure the attributes.

Set the value of the **RvgResourceName** attribute to the name of the RVG resource for the RVGPrimary agent.

Configure the following attributes in the SQL Server service group for the RVG Primary resources:

Table 9-8 RVG Primary resources


Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
RVGPrimary	RvgResourceName
Resources for the disk group for the SQL user-defined database files:	
RVGPrimary	RvgResourceName

Creating the RVG Primary resources

For all disk groups, create an RVG Primary Resource for replication.

To create the RVG Primary resource for the SQL Server system disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:



- a Enter the **Resource Name** for the RVG Primary resource for the SQL Server system files disk group, for example INST1_RvgPrimary.
 - b Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Veritas Cluster Server 5.0 Administrator's Guide* for more information about the RVG Primary agent.
- 6 Verify that **Critical** and **Enabled** are both checked.
- 7 Click **OK**.

To create the RVG Primary resource for the SQL Server user-defined database disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:

- a Enter the **Resource Name** for the RVG Primary resource for the SQL Server user-defined database disk group, for example INST1_DB1_RvgPrimary.
 - b Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_DB1_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults.
- 6 Verify that **Critical** and **Enabled** are both checked.
- 7 Click **OK**.

Linking the RVG Primary resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the SQL Server service group (INST1_SG) to establish the dependencies between the resources for replication.

Start from the top parent and link the following resources:

Table 9-9 Dependencies for the RVG Primary resources for RDC

Parent	Child
INST1_ SG-MountV	INST1_RvgPrimary
INST1_ SG-MountV-1	INST1_RvgPrimary
INST1_ SG-MountV-2	INST1_DB1_RvgPrimary
INST1_ SG-MountV-3	INST1_DB1_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example INST1_ SG-MountV.
- 4 Click the child resource, for example INST1_RvgPrimary.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the SQL Server service group (INST1_SG) online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane on the **Resources** tab, right-click the first RVG Primary resource (INST1_RvgPrimary) and select **Online > SYSTEM1**.
- 3 In the right pane on the **Resources** tab, right click the second RVG Primary resource (INST1_DB1_RvgPrimary) and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG

In the RVG service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone for the RVG Service Group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.
- 7 Click **OK**.


Setting a dependency between the service groups

To ensure that the SQL Server service group and the RVG service group fail over and switch together, set up an online local hard dependency from the RVG service group to the SQL Server service group.

The SQL service group (for example, INST1_SG) is dependent on the replication service group (for example, INST1_RVG_GRP).

To set up an online local hard dependency

- 1 From VCS Cluster Explorer, in the left pane, select the cluster (MYCLUSTER).
- 2 In the right pane, select the **Service Groups** tab.
- 3 Click the **Link** button to create a dependency between service groups.
- 4 Click the SQL Server service group (the parent service group), for example INST1_SG.
- 5 Click the RVG service group (the child resource), for example INST1_RVG_SG.
- 6 In the **Link Service Groups** window:



- a Select the **Relationship** of **online local**.
- b Select the **Dependency Type** of **hard**.
- c Click **OK**.

Adding the nodes from the secondary zone to the RDC


Configuration of the systems in the Primary Zone (zone 0) is complete. The nodes in the Secondary Zone (zone 1) can now be added to the RDC configuration.

Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

To add the nodes from the secondary zone to the RVG

- 1 From the active node of the cluster in the primary zone, click **Start > All Programs > Veritas > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the **Welcome page**, and click **Next**.
- 3 In the **Wizard Options** dialog box:
 - a Click **Modify an existing replication service group**. The existing replication service group is selected, by default.
 - b Click **Next**.
- 4 If a VCS notice message appears, asking if you want to continue, click **Yes**.
- 5 Specify the system priority list:



- a In the **Available Cluster Systems** box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group's system list.
 To remove a node from the service group's system list, click the node in the **Systems in Priority Order** box, and click the left arrow icon.
- b To change the priority of a node in the system list, click the node in the **Systems in Priority Order** box, then click the up and down arrow icons. The node at the top of the list has the highest priority.
- c Click **Next**.

- 6

If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
- 7


Review the Disk Group and Replicated Volume Group Configuration and click **Next**.
- 8

In the IP Resource Options dialog box, select **Modify IP resource** and click **Next**.
- 9

If a VCS error appears, click **OK**.
- 10

In the Network Configuration dialog box, verify that the selected adapters are correct and click **Next**.
- 11

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.

- 12

Click **Next** to modify the replication service group.
- 13

When prompted, click **Yes** to modify the service group.
- 14

Click **Finish**.

Configuring secondary zone nodes in the RVG service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Configuring the IP resources for failover

Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.


Also modify the TCP/IP addresses on all the nodes in both zones, to include the IP address for replication.

In the event of a system or SQL 2000 failure, VCS attempts to fail over the SQL Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone.

To modify the IP resources in the RVG service group


- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Resources** tab.
- 3 Right-click the RVG IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.

4 In the Edit Attributes window:



- a Select **Per System**.
 - b Select the first node in the primary zone and enter the virtual IP address for the primary zone.
 - c Select the second node in the primary zone and enter the virtual IP address for the primary zone (the same IP address as the first node).
 - d Repeat for all nodes in the primary zone.
 - e Select the first node in the secondary zone (SYSTEM3) and enter the virtual IP address for the secondary zone.
 - f Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - g Repeat for all nodes in the secondary zone.
 - h Click **OK**.
- 5 In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone

have the same IP address, that is different from the IP address for the primary zone.



6 Close the Properties View window.

To modify the TCP/IP addresses on the primary nodes to include the replication IP

- 1 On the primary node in the primary zone, click **Start > Control Panel > Network Connections > Public**.
- 2 In the Public Status window, click the **Properties** button.
- 3 In the Public Properties window, select **Internet Protocol (TCP/IP)**.
- 4 Click the **Properties** button.
- 5 In the Internet Protocol (TCP/IP) Properties window, click the **Advanced** button.
- 6 In the Advanced TCP/IP Settings window, in the IP addresses area, click the **Add** button.
- 7 In the TCP/IP Address window, enter the IP address of the IP for replication, and the subnet mask.
- 8 Click **Add**.
- 9 In the Advanced TCP/IP Settings window, click **OK**.
- 10 In the Internet Protocol (TCP/IP) Properties window, click **OK**.
- 11 In the Public Properties window, click **OK**.
- 12 In the Public Status window, click **Close**.
- 13 Repeat these steps for the primary node in the secondary zone.

Adding the nodes from the secondary zone to the SQL Server service group

Use the SQL Server Agent Configuration Wizard to add the nodes from the secondary zone to the SQL Server service group.

To add the nodes from the secondary zone to the SQL Server service group

- 1 Select **Start > All Programs > Veritas > Veritas Cluster Server > SQL Server Enterprise Agent > SQL Configuration Wizard**.
- 2 In the Select Configuration Option dialog box, select **MS SQL Server - Service Group Configuration**. Select **Edit**, and click **Next**.
- 3 Verify that you have met the prerequisites listed and click **Next**.
- 4 Select the service group to be modified (INST1_SG) and click **Next**.
- 5 If a VCS notice message appears indicating that resources are online, click **Yes** to continue.
- 6 On the Service Group Configuration page, select the nodes in the secondary zone, use the arrow button to move them from **Available Cluster Systems** to **System in Priority Order**.

To change the priority of a system in the **Systems in Priority Order** list, select the system and click the up and down arrow icons. Arrange the systems in priority order in as failover targets for the group. The server that needs to come online first must be at the top of the list followed by the next one that will be brought online.

This set of nodes selected for the SQL Server service group must be the same as the nodes selected for the RVG service group. Ensure that the nodes are also in the same priority order.

- 7 Click **Next**.
- 8 On the SQL Server Instance Selection page, click **Next**.
- 9 The User Databases List page summarizes the databases for this instance of SQL. Click **Next**.
- 10 On the Detail Monitoring Configuration page, clear the box in the SQL Instance List to disable monitoring, as required. Detailed monitoring is not necessary. Click **Next**.
- 11 On the Registry Replication Path page, click **Next**.
- 12 On the Virtual Server Configuration page, verify that the public adapter is used on each system and click **Next**.
- 13 In the Service Group Summary, review the service group configuration and click **Next**.

- 14 A message appears if the configuration is currently in the Read Only mode. Click **Yes** to make the configuration read and write enabled. The wizard validates the configuration and modifies it.
- 15 Click **Finish**.

Configuring the zones in the SQL Server service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the SQL Server service group

- 1 From VCS Cluster Explorer, in the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Verifying the RDC configuration

Failover simulation is an important part of configuration testing. Test the failover by switching online nodes.

Note: This should never be tested on systems with live data. A reliable and tested backup should be available. A tested backup means that it has been tested successfully by a restore.

Switching online nodes

Switch nodes using Veritas Cluster Manager (Java Console). When you complete the procedure, you will see the online system role shift from one system to another.

To switch online nodes

- 1 Open the Veritas Cluster Manager (Java Console) (**Start > All Programs > Veritas > Veritas Cluster Manager (Java Console)**).
- 2 Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, in the File menu, click **New Cluster**. In the **New Cluster - Connectivity Configuration** window, enter the computer name in the **Host name** field and click **OK**.
- 3 In the **Machinename - Login window**, enter your user name and password in the respective fields and click **OK**.
- 4 Right-click the service group in the left pane, and select an alternate system name from the **Switch To** entry.
- 5 In the **Question** dialog box, click **Yes** to confirm you do want to switch the service group to the other node.

Configuring Replicated Data Clusters for SQL 2005

This chapter includes the following topics:

- [“Tasks for configuring Replicated Data Clusters for SQL 2005”](#) on page 354
- [“Reviewing the prerequisites”](#) on page 358
- [“Reviewing the configuration”](#) on page 363
- [“Configuring the storage hardware and network”](#) on page 364
- [“Installing Veritas Storage Foundation for Windows”](#) on page 366
- [“Configuring the cluster”](#) on page 373
- [“Configuring cluster disk groups and volumes”](#) on page 388
- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 396
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 400
- [“Installing SQL Server 2005 on the second node”](#) on page 403
- [“Setting the internal name of the clustered instance”](#) on page 407
- [“Configuring the VCS SQL Server service group”](#) on page 409
- [“Creating the primary system zone”](#) on page 414
- [“Registering the virtual server in the SQL Server”](#) on page 415
- [“Creating a SQL Server user-defined database”](#) on page 417
- [“Verifying the installation in the primary zone”](#) on page 420
- [“Creating a parallel environment in the secondary zone”](#) on page 421

- [“Adding the IP addresses for the primary and secondary zones”](#) on page 422
- [“Adding the systems in the secondary zone to the cluster”](#) on page 422
- [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 429
- [“Configuring a hybrid RVG service group for replication”](#) on page 437
- [“Setting a dependency between the service groups”](#) on page 454
- [“Adding the nodes from the secondary zone to the RDC”](#) on page 454
- [“Verifying the RDC configuration”](#) on page 462

Tasks for configuring Replicated Data Clusters for SQL 2005

Configure the high availability and SQL 2005 components on the primary and secondary zones. Then complete the Replicated Data Set solution by configuring the components for both zones.

Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details on VVR.

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

Table 10-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Reviewing the prerequisites” on page 358	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 363	✓ Understanding active/passive configuration and zone failover in a RDC environment
	✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 364	✓ Setting up the storage hardware for a cluster environment
	✓ Verifying the DNS entries for the systems on which SQL will be installed

Table 10-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Installing Veritas Storage Foundation for Windows” on page 366	✓ Verifying the driver signing option for the system ✓ Installing Veritas Storage Foundation for Windows HA (automatic installation) ✓ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ✓ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server ✓ Configuring VxSAS
“Configuring the cluster” on page 373	✓ Verifying static IP addresses and name resolution configured for each node ✓ Configuring cluster components using the Veritas Cluster Server Configuration Wizard ✓ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 388	✓ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ✓ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2005 on the first node” on page 396	✓ Installing and configuring SQL Server 2005 ✓ Setting SQL Server services to manual start ✓ Configuring SQL services

Table 10-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Preparing to install SQL Server 2005 on the second node” on page 400	✓ Stopping the SQL Service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Adding drive letters ✓ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2005 on the second node” on page 403	✓ Installing and configuring SQL Server 2005
“Setting the internal name of the clustered instance” on page 407	✓ Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 409	✓ Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating the primary system zone” on page 414	✓ Creating the primary system zone ✓ Adding the nodes to the primary zone
“Registering the virtual server in the SQL Server” on page 415	✓ Registering the virtual server in the SQL Server
“Creating a SQL Server user-defined database” on page 417	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server ✓ Adding resources for a user-defined database in VCS

Table 10-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Verifying the installation in the primary zone” on page 420	✓ Simulating failover ✓ Switching online nodes
“Creating a parallel environment in the secondary zone” on page 421	✓ Reviewing the prerequisites ✓ Reviewing the configuration ✓ Configuring the network and storage ✓ Installing SFW HA ✓ Configuring disk groups and volumes for SQL
“Setting up the Replicated Data Sets (RDS)” on page 429	✓ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones
“Configuring a hybrid RVG service group for replication” on page 437	✓ Creating a hybrid Replicated Volume Group (RVG) service group ✓ Configuring the hybrid RVG service group
“Setting a dependency between the service groups” on page 454	✓ Setting up a dependency from the VVR RVG Service Group to the SQL Server Service Group
“Adding the nodes from the secondary zone to the RDC” on page 454	✓ Adding the nodes from the secondary zone to the RVG service group ✓ Configuring the IP resources for failover ✓ Adding the nodes from the secondary zone to the SQL Server service group
“Verifying the RDC configuration” on page 462	✓ Verifying that failover occurs first within zones and then from the primary to the secondary zone

Reviewing the prerequisites

Verify that the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary zone and a secondary zone.

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) with the Veritas Cluster Server Database Agent for Microsoft SQL
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SP1 required)
with
 - Windows Server 2003 (32-bit) Enterprise Edition or Datacenter Edition (SP1 required)
or
 - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
or
 - Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SP1 required)
with
 - Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required)
or
 - Windows Server 2003 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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 361.
- 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 SQL 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.

- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

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.

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

Reviewing the configuration

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

- ✓ SQL virtual server: the IP address should be the same on all nodes at the primary and secondary zones
- ✓ Replication IP address for the primary zone
- ✓ Replication IP address for the secondary zone

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

Sample configuration

The sample setup has four servers, two for the primary zone and two for the secondary zone. The nodes will form two separate clusters, one at the primary zone and one at the secondary zone.

The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary zone

SYSTEM1 & SYSTEM2	First and second nodes of the primary zone
INST1_SG	Microsoft SQL Server 2005 service group
INST1-VS	Virtual SQL Server cluster
INST1	SQL Instance Name
INST1_DG	Cluster disk group for SQL system database and files
INST1_DATA_FILES	Volume for Microsoft SQL Server system data files
INST1_REGREP_VOL	Volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	Replicator log volume required by VVR
INST1_DB1_DG	Cluster disk group for SQL Server user-defined database and files
INST1_DB1_VOL	Volume for storing a Microsoft SQL Server user-defined database

INST1_DB1_LOG	Volume for storing a Microsoft SQL Server user-defined database log file
INST1_DB1_REPLOG	Replicator log volume required by VVR for SQL user-defined database

Secondary zone

SYSTEM3 & SYSTEM4 First and second nodes of the secondary zone

All the other parameters are the same as on the primary zone.

RDS and VVR
Components

INST1_RDS	RDS name for SQL system database and files
INST1_RVG	RVG name for SQL system database and files
INST1_RVG_SG	Replication service group for SQL system database and files
INST1_DB1_RDS	RDS name for SQL Server user-defined database and files
INST1_DB1_RVG	RVG name for SQL Server user-defined database and files
INST1_DB1_RVG_SG	Replication service group for SQL Server user-defined database and files

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 Veritas Storage Foundation 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. You must select the options to install VVR and the Veritas Cluster Server Database Agent for SQL. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation.

Setting Windows driver signing options

When installing on remote or local systems running Windows Server 2003, you must set the Windows driver signing options to allow the installation.

To change the driver signing options on each local system

- 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, note the current setting, and select **Ignore** or one of the other options from the table, to allow installation to proceed.
- 4 Click **OK**.
- 5 Repeat for each computer.

If you do not change these options, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing options to their previous states.

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**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Data- base Agent for SQL	Required to configure high availability for SQL Server.
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability.
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 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 `vxsasscfg.exe` from the command prompt of the required machine.

The Welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
Password	Specify a password in the Password field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same 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

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 that name resolution is configured for each node.
- ✓ Set the required privileges:
 - 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.

Note: Add only systems in the primary zone (zone 0) to the cluster at this time.

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 375.

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 375. 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 375.


- 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 376, proceed to the next step. Otherwise, proceed to [step 11](#) on page 378.

- 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 374, 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 384.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 385.

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 385. 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.

Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the disk group are shared and are available from all nodes in each zone. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

- INST1_REPLOG contains the VVR Replicator Log.
You can create this volume later while setting up the Replicated Data Sets.
See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 429.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and files.

The following disk group and volumes may be created now or later in the configuration process.

- INST1_DB1_DG is the disk group for the SQL Server user-defined database and files
- INST1_DB1_VOL contains the user database files
- INST1_DB1_LOG contains the user database log files
- INST1_DB1_REPLOG contains the VVR Storage Replicator Log.

Warning: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.


Creating a cluster disk group

Create a cluster disk group on the first node of the cluster. You can repeat the process to create the disk group for the SQL Server user-defined database and files at this time or later.

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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes in the cluster disk group (INST1_DG) for the system files on the first node of the cluster:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)

- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_REPLOG** contains the VVR Storage Replicator Log.
 You can create the INST1_REPLOG volume at this time or during the process of creating Replicated Data Sets.
 See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 429.

You can create the INST1_DB1_VOL and INST1_DB1_LOG volumes at this time or during the process of “[Creating a SQL Server user-defined database](#)” on page 417.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and the log files. The following volumes in the disk group (INST1_DB1_DG) for the SQL Server user-defined database and files may be created now or later in the configuration process:

- **INST1_DB1_VOL** contains the user database files
- **INST1_DB1_LOG** contains the user database log files
- **INST1_DB1_REPLOG** contains the VVR Storage Replicator Log.

Warning: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.


Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.

You can right-click the disk group you have just created, for example INST1_DG.

- 5
- At the New Volume wizard opening screen, click **Next**.
- 6
- Select the disks for the volume.




-
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
-
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
-
- Click **Next**.

7 Specify the parameters of the volume.

- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the **Mirror Info** area, select the appropriate mirroring options.
 - Verify that **Enable Logging** is not selected.
 - Click **Next**.
- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on


all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
 - For the Replicator Log volume only: Clear the Format this volume check box.
 - Click **NTFS**.
 - Select an allocation size or accept the Default.
 - The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 11 Click **Finish** to create the new volume.
- 12 Repeat these steps to create additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Installing and configuring SQL Server 2005 on the first node

Before installing Microsoft SQL Server 2005, ensure the cluster disk group is imported to the first node and the volumes are mounted.

Complete the following procedures to install and configure Microsoft SQL Server 2005:

- ✓ [Installing Microsoft SQL Server](#)
- ✓ [Setting SQL Server 2005 services to manual start](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters). See “[Importing the cluster disk group](#)” on page 401 and “[Adding drive letters to mount the volumes](#)” on page 402.

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005


- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.

- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.

If you install optional components on one node, install the same components in the same order on other nodes.


- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.


Setting SQL Server 2005 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.

- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2005 Service”](#) on page 401
- [“Deporting the cluster disk group”](#) on page 401
- [“Importing the cluster disk group”](#) on page 401
- [“Adding drive letters to mount the volumes”](#) on page 402
- [“Renaming shared SQL Server 2005 files”](#) on page 403

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2005 Service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

In order to install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) on the next node in the cluster (SYSTEM2).

To import a cluster disk group


- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
- *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2005 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 403
- [“Removing shared SQL Server files”](#) on page 406

Installing SQL Server

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 401 and [“Adding drive letters to mount the volumes”](#) on page 402.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.


To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
 - 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.

Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.

- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.

Repeat the procedures described in [“Preparing to install SQL Server 2005 on the second node”](#) on page 400 and [“Installing SQL Server 2005 on the second node”](#) on page 403 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.

- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, INST1-VS\INST1).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.


To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).


2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
- Enter the name in the format *System_Name\Instance_Name*.
- Select the appropriate authentication method.
- Enter valid user credentials if using Domain authentication and click **Connect**.

3 Find the SQL Server name as follows:

- Right-click the instance in the Object Explorer and click **New Query**.
- In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver
- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, *SYSTEM1\INST1*. (For a default instance, you see only *System_Name*.)

4 Delete the contents in the upper pane.

- 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
sp_dropserver "System_Name\Instance_Name"
where **System_Name\Instance_Name** is the name noted in [step 3](#).
For example, for a named instance:
`sp_dropserver "SYSTEM1\INST1"`
For example, for a default instance:
`sp_dropserver "SYSTEM1"`
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.

Configuring the VCS SQL Server service group


The Enterprise Agent Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

- ✓ Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2005, is installed on all cluster nodes. See "[Configuring the cluster](#)" on page 373.
- ✓ Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See "[Configuring the cluster](#)" on page 373.
- ✓ Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- ✓ Verify that the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.

- ✓ Verify that the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- ✓ Assign a unique virtual server name to SQL Server 2005. While it must be unique within the cluster, it must be the same as the virtual name of the SQL Server in the primary zone.
- ✓ Assign a unique virtual IP address to the SQL Server 2005 instance.


To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- 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 Systems in Priority Order list.

- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
 For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment. Click **Next**.
 - 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
 - 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 411 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script

- on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) for all the instances for which detail monitoring will be configured.
- 10 Click **Next**.
- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:

The screenshot shows the 'Virtual Server Configuration' window of the SQL Server Agent Configuration Wizard. The window title is 'SQL Server Agent Configuration Wizard'. The subtitle is 'Virtual Server Configuration' with the instruction 'Enter a virtual server name for the application and specify the virtual IP information.' The left sidebar contains a list of steps: Prerequisites, System List, Instance Selection, Databases, Detail Monitoring, Registry Replication, and Virtual Computer. The 'Virtual Computer' section is expanded, showing 'Summary' and 'Configuration Status'. The main area contains fields for 'Virtual Server Name', 'Virtual IP Address', and 'Subnet Mask', each with a text input box and a default value of '0 . 0 . 0 . 0'. Below these is a table titled 'Specify the adapter to be used on each system.' with columns 'System Name' and 'Adapter Display Name'. The table has two rows: 'PC177VM-3' and 'PC177VM-4', both with 'Public' as the adapter display name. At the bottom right is an 'Advanced Settings...' button. At the bottom are '< Back', 'Next >', and 'Cancel' buttons.

System Name	Adapter Display Name
PC177VM-3	Public
PC177VM-4	Public

- Enter the virtual name for the server, for example `INST1-VS`. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.


- Enter the subnet mask to which the virtual IP address belongs.
 - For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.
 The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review 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.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
- To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
 You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.
- The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

Creating the primary system zone

In the service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 Select the SQL Server service group (INST1_SG) in the left pane and the Properties tab in the right pane.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.




- 7 Click **OK**.

Registering the virtual server in the SQL Server

In the SQL Server Enterprise Manager, register the VCS virtual server (INST1-VS) and delete the physical server name.


To register the VCS virtual server

- 1 Click **Start > All Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 In the right pane, expand the Microsoft SQL Servers.
- 3 Right-click the **SQL Server Group** and select **New SQL Server Registration**.
- 4 On the Welcome screen, click **Next**.
- 5 In the **Select a SQL Server** window, select the virtual server\SQL instance (for example INST1-VS\INST1).



- 6 Click the **Add** button.
- 7 Click **Next**.
- 8 In the **Select an Authentication Mode** window, select an authentication mode and click **Next**.

- 9 In the **Select SQL Server Group** window, select Create a new top-level SQL Server group, or leave it in the existing group. If a new SQL Server group is created, enter the **Group name**.



- 10 Click **Next**.
- 11 In the **Completing the Register the SQL Server Wizard** window, click **Finish**.
- 12 Wait for the **Register SQL Server Messages** window that shows the server registration completed successfully and click **Close**.

To delete the physical server name

- 1 In the Microsoft SQL Server **Enterprise Manager**, expand the Microsoft SQL Servers and the **SQL Server Group**.
- 2 Right-click the physical name of the SQL Server and instance (for example SYSTEM1\INST1) and select **Delete SQL Server Registration**.
- 3 Click **Yes** in the warning message window.
- 4 Verify the SQL servers in the Microsoft SQL Server **Enterprise Manager** window. In the left pane, the **Microsoft SQL Servers** should show the virtual server under **SFWHA Clustered SQL Servers** and no items under the **SQL Server Group**.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- ✓ Create volumes for a user-defined SQL Server database and its transaction log.
- ✓ Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- ✓ Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 390 for information on how to use the VEA console to create a volume.

Note: Best practice is to create a separate disk group with separate volumes for the user-defined database files.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2005 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.

- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.

Adding VMDg and MountV resources

Before running the SQL Server Agent wizard to add the VMDg and MountV resources:

- ✓ Make sure the SQL Server resources are online.
- ✓ Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

- Click **Next** to configure another SQL service group or an MSDTC service group.

Verifying the installation in the primary zone

Failover simulation is an important part of configuration testing.

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.

Creating a parallel environment in the secondary zone

After setting up a SFW HA environment in the primary zone, use the guidelines in the following list to complete the same tasks in the secondary zone.

Note: Before you begin to configure the secondary zone, offline the SQL Server resource, the SQL virtual server name resource, and the SQL virtual IP resource in the primary zone. The remaining resources should be online, including the VMDg resource. In VEA, make sure to remove all the drive letters from the configured volumes, to avoid conflicts when configuring the zones.

- ✓ [“Reviewing the prerequisites”](#) on page 358
- ✓ [“Reviewing the configuration”](#) on page 363
- ✓ [“Configuring the storage hardware and network”](#) on page 364
- ✓ [“Adding the IP addresses for the primary and secondary zones”](#) on page 422
- ✓ [“Adding the systems in the secondary zone to the cluster”](#) on page 422
- ✓ [“Configuring cluster disk groups and volumes”](#) on page 388

During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as the cluster at the primary zone:

 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- ✓ [“Installing and configuring SQL Server 2005 on the first node”](#) on page 396

Select the same options in the secondary zone as you did for the primary zone.
- ✓ [“Preparing to install SQL Server 2005 on the second node”](#) on page 400

The instance name must be the same in the primary zone and secondary zone.
- ✓ [“Installing SQL Server 2005 on the second node”](#) on page 403

Note: After you install SQL Server 2005 on the nodes in the secondary zone, make sure to use VEA to remove all the drive letters from the configured volumes, to avoid conflicts during the configuration of the zones.

Adding the IP addresses for the primary and secondary zones

Create IP addresses for both the primary and secondary zones. Connect to the primary system in each zone and add the IP address for the zone.

To create an IP address

- 1 On the primary system in the primary zone, open the Network Connections (Click **Start > Settings > Network Connections**).
- 2 Right-click the public network connection to be used and click **Properties**.
- 3 Select **Internet Protocol (TCP/IP)** and click **Properties**.
- 4 Click **Advanced**.
- 5 In the **IP Settings** tab, click **Add** to add a virtual IP address and subnet mask.
- 6 In the **TCP/IP Address** dialog box, enter a new IP address and a subnet mask and click **OK** several times, until the **Network Connections** dialog box is closed.
- 7 Repeat these steps on the primary system in the secondary zone.


Adding the systems in the secondary zone to the cluster

Add the nodes in the secondary zone to the existing cluster with the following procedure.

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 panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

- 4 In 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 425.

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 424. 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 425.


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, type 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 the VCS private network communication, on 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.

Return to the task list [“Creating a parallel environment in the secondary zone”](#) on page 421.


Setting up the Replicated Data Sets (RDS)

Set up the Replicated Data Sets (RDS) in the primary zone and secondary zone. You can configure an RDS using the Create RDS wizard for both zones.

- ✓ Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- ✓ Verify that the cluster disk group is imported in the primary and secondary zone

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 (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume


- Click **Create Volume** and enter the following information in the dialog box that displays.

Name	Enter the name for the volume in the Name field.
Size	Enter a size for the volume in the Size field.
Layout	Select the desired volume layout.
Disk Selection	■ Choose Select disks automatically if you want VVR to select the disks for the Replicator Log. ■ Choose Select disks manually to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select Add to move the disks into the selected disks pane.

- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 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.

AutoDCM	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.
DCM	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.
Off	Disables Replicator Log overflow protection.
Override	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. 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.
Fail	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.

If the SQL Server user-defined database and files are in a separate disk group from the SQL Server system files, repeat the procedure “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 429 for the disk group (INST1_DB1_DG) that contains the SQL Server user-defined database and files. Provide unique

names for the Replicated Data Set name, and the Replicated Volume Group name. See “[Sample configuration](#)” on page 363 for a list of example names.

Configuring a hybrid RVG service group for replication

Create and configure a hybrid Replicated Volume Group (RVG) service group for replication.

The RVG service group is hybrid because it behaves as a failover service group within a zone and as a parallel service group between zones.

For additional information about service group types, see the *Veritas Cluster Server 5.0 Administrator’s Guide*.

Configure the RVG service group’s resources manually by copying and modifying components of the SQL Server service group. Then create new RVG resources and bring them online.

The RVG service group for RDC contains the following resources:

Table 10-3 Replication service group resources


Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg for the first disk group	Volume Manager disk group with SQL system files
VvrRvg for the first disk group	Replicated volume group with SQL system files
VMDg for the second disk group	Volume Manager disk group with SQL user-defined files
VvrRvg for the second disk group	Replicated volume group with SQL user-defined files

Creating the RVG service group

Create a hybrid replicated volume (RVG) service group, to contain the resources for replication.

To create a hybrid RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.
- 3 In the **Add Service Group** window:



- a Enter a name for the service group, for example INST1_RVG_SG.
- b Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- c Select **Hybrid**.
- d Click **OK**.

Configuring the RVG service group for RDC replication

Configure the RVG service group’s resources manually for RVG by completing the following tasks:

- **“Configuring the IP and NIC resources”**
Copy IP and NIC resources of the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
- **“Configuring the VMDg resources for the first disk group” and “Configuring the VMDg resources for the second disk group”**
Copy the VMDg resources for all disk groups in the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).

- [“Adding the VVR RVG resources for the disk groups”](#)
Create the VVR RVG resources for all the disk groups and enter the attributes for each of the disk groups and the replication IP address.
- [“Linking the VVR RVG resources to establish dependencies”](#)
Link the VVR RVG resources to establish the dependencies between the VMDg resources, the IP resource for replication, and the VVR RVG resources for the disk groups. Configure the RVG service group’s VMDg resources to point to the disk groups that contain the RVGs.
- [“Deleting the VMDg resource from the SQL Server service group”](#)
Delete the VMDg resources from the SQL Server service group, because they depend on the replication and were configured in the RVG service group.

Configuring the IP and NIC resources

Configure the following resources and attributes for the IP and NIC:


Table 10-4 IP and NIC resources

Resource	Attributes to Modify
IP	Address
NIC	(none)

To create the IP resource and NIC resource

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the IP resource (INST1_SG-IP), and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.

- 5 In the **Name Clashes** window, change the names of the IP and NIC resources for the RVG service group.



- 6 Click **OK**.

To modify the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.
- 2 In the **Properties View** window, for the **Address** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, enter the VVR IP address for the Primary Zone as the scalar value.
- 4 Close the **Properties View** window.

To enable the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (INST1_RVG_SG-IP) and select **Enabled**.
- 2 In the **Resources** tab display area, right-click the NIC resource (INST1_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg resources for the first disk group

Create the VMDg resource in the SQL Server service group, and clear the DGGuid attribute for the new VMDg.

Configure the following attributes in the SQL Server service group for the MountV resource:

Table 10-5 MountV resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
MountV (for the SQL Server system volume)	VMDg Resource Name Volume Name
MountV (for the registry volume)	VMDg Resource Name Volume Name
Resources for the disk group for the SQL user-defined database files:	
MountV (for the SQL Server user-defined database log)	VMDg Resource Name Volume Name
MountV (for the SQL Server user-defined database)	VMDg Resource Name Volume Name

To create the VMDg resource for the first disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the first disk group, with the SQL system files (INST1_SG-VMDg), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg.
- 6 Click **OK**.

To clear the DGGuid attribute for the new VMDg

- 1 In the **Resources** tab display area, right-click the new VMDg resource.
- 2 In the same **Properties View** window, for the **DGGuid** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, clear the scalar value for the **DGGuid** attribute.
- 4 Close the **Properties View** window.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server system data files (INST1_SG-MountV) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server system data files (INST1_DATA_FILES).
- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the registry volume (INST1_SG-MountV-1) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the **Volume Name** attribute is the registry volume (INST1_REGREP_VOL).
- 9 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11 Close the **Properties View** window.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg) and select **Enabled**.

Configuring the VMDg resources for the second disk group

Repeat the VMDg and MountV configuration for the second disk group.

To create the VMDg resource for the second disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the second disk group, with SQL user-defined files (INST1_SG-VMDg-1), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg-1.
- 6 Click **OK**.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined log (INST1_SG-MountV-2) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined log (INST1_DB1_LOG).
- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg-1.
- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined database (INST1_SG-MountV-3) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined database (INST1_DB1_VOL).
- 9 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.

- 10
- In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11
- Close the **Properties View** window.

To enable the VMDg resource

- 1
- In the left pane, select the RVG service group (INST1_RVG_SG).
- 2
- In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg-1) and select **Enabled**.

Adding the VVR RVG resources for the disk groups

Add VVR RVG resources for replication of the disk groups.


Configure the following attributes in the RVG service group for the VvrRvg resource:

Table 10-6 VvrRvg resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
VvrRvg	VMDgResName IPResName
Resources for the disk group for the SQL user-defined database files:	
VvrRvg	VMDgResName IPResName


To create the VVR RVG resource for the first disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



- a Enter the **Resource Name** for the VVR RVG resource.
- b Select the **Resource Type** of VvrRvg.
- 3 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the RVG group that is being managed, for example INST1_RVG.
- 5 Click **OK**.
- 6 In the **Add Resource** window, for the **VMDGResName** attribute, click **Edit**.
- 7 In the **Edit Attribute** window, enter the name of the disk group containing the RVG, for example INST1_RVG_SG-VMDg.
- 8 Click **OK**.
- 9 In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, enter the name of the IP resource managing the IP address for replication, for example INST1_RVG_SG-IP.
- 11 Click **OK**.


12 In the **Add Resource** window, verify that the attributes have been modified:



13 Click **OK**.


To create the VVR RVG resource for the second disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



- a Enter the **Resource Name** for the VVR RVG resource for the second disk group.
 - b Select the **Resource Type** of VvrRvg.
- 3 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
 - 4 In the **Edit Attribute** window, enter the name of the RVG group that is being managed, for example INST1_DB1_RVG.
 - 5 Click **OK**.
 - 6 In the **Add Resource** window, for the **VMDGResName** attribute, click **Edit**.

- 7
- In the **Edit Attribute** window, enter the name of disk group containing the RVG, for example INST1_RVG_SG-VMDg-1.
- 8
- Click **OK**.
- 9
- In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10
- In the **Edit Attribute** window, enter the name IP resource managing the IP address for replication, for example INST1_RVG_SG-IP. In this example both disk groups are using the same IP resource for replication.
- 11
- Click **OK**.
- 12
- In the **Add Resource** window, verify that the attributes have been modified:



- 13
- Click **OK**.

Linking the VVR RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the VVR RVG service group to establish the dependencies between the resources. Start from the top parent and link the following resources:


Table 10-7 Dependencies for VVR RVG resources for RDC

Parent	Child
Resources for the disk group for the SQL system files:	
INST1_ VvrRvg	The IP for replication, for example INST1_RVG_SG-IP
INST1_ VvrRvg	The VMDg for the SQL system files, for example INST1_RVG_SG-VMDg
Resources for the disk group for the SQL user-defined database files:	
INST1_DB1_ VvrRvg	The IP for replication, for example INST1_RVG_SG-IP
INST1_DB1_ VvrRvg	The VMDg for the SQL user-defined database files, for example INST1_RVG_SG-VMDg-1

To link the VVR RVG resources

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example INST1_DB1_VvrRvg.
- 4 Click the child resource, for example INST1_RVG_SG-IP.
- 5 When prompted to confirm, click **OK**.

6 Repeat these steps to link all the RVG resources:



Deleting the VMDg resource from the SQL Server service group

The VMDg resources must now be manually deleted from the SQL Server service group, because they depend on replication and were configured in the RVG service group.

To delete the VMDg Resources from the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) from the left pane.
- 2 In the **Resources** tab display area, right-click the VMDg resource for the first disk group (INST1_SG-VMDg) and select **Delete**.
- 3 Click **Yes** to confirm that you want to delete it (even if it is online).
- 4 In the **Resources** tab display area, right-click the VMDg resource for the second disk group (INST1_SG-VMDg-1) and select **Delete**.
- 5 Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

Add resources of type RVG Primary to the SQL Server service group for each of the SQL Server disk groups (system and user-defined) and configure the attributes.

Set the value of the **RvgResourceName** attribute to the name of the RVG resource for the RVGPrimary agent.

Configure the following attributes in the SQL Server service group for the RVG Primary resources:

Table 10-8 RVG Primary resources


Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
RVGPrimary	RvgResourceName
Resources for the disk group for the SQL user-defined database files:	
RVGPrimary	RvgResourceName

Creating the RVG Primary resources

For all disk groups, create an RVG Primary Resource for replication.

To create the RVG Primary resource for the SQL Server system disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:



- a Enter the **Resource Name** for the RVG Primary resource for the SQL Server system files disk group, for example INST1_RvgPrimary.
 - b Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Veritas Cluster Server 5.0 Administrator's Guide* for more information about the RVG Primary agent.
- 6 Verify that **Critical** and **Enabled** are both checked.
- 7 Click **OK**.

To create the RVG Primary resource for the SQL Server user-defined database disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:
 - a Enter the **Resource Name** for the RVG Primary resource for the SQL Server user-defined database disk group, for example INST1_DB1_RvgPrimary.
 - b Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_DB1_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults.
- 6 Verify that **Critical** and **Enabled** are both checked.
- 7 Click **OK**.

Linking the RVG Primary resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the SQL Server service group (INST1_SG) to establish the dependencies between the resources for replication.

Start from the top parent and link the following resources:

Table 10-9 Dependencies for the RVG Primary resources for RDC

Parent	Child
INST1_ SG-MountV	INST1_RvgPrimary
INST1_ SG-MountV-1	INST1_RvgPrimary
INST1_ SG-MountV-2	INST1_DB1_RvgPrimary
INST1_ SG-MountV-3	INST1_DB1_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 Click the **Link** button in the right pane.

- 3 Click the parent resource, for example INST1_SG-MountV.
- 4 Click the child resource, for example INST1_RvgPrimary.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the SQL Server service group (INST1_SG) online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane on the **Resources** tab, right-click the first RVG Primary resource (INST1_RvgPrimary) and select **Online > SYSTEM1**.
- 3 In the right pane on the **Resources** tab, right click the second RVG Primary resource (INST1_DB1_RvgPrimary) and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG

In the RVG service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone for the RVG Service Group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.
- 7 Click **OK**.


Setting a dependency between the service groups

To ensure that the SQL Server service group and the RVG service group fail over and switch together, set up an online local hard dependency from the RVG service group to the SQL Server service group.

The SQL service group (for example, INST1_SG) is dependent on the replication service group (for example, INST1_RVG_GRP).

To set up an online local hard dependency

- 1 From VCS Cluster Explorer, in the left pane, select the cluster (MYCLUSTER).
- 2 In the right pane, select the **Service Groups** tab.
- 3 Click the **Link** button to create a dependency between service groups.
- 4 Click the SQL Server service group (the parent service group), for example INST1_SG.
- 5 Click the RVG service group (the child resource), for example INST1_RVG_SG.
- 6 In the **Link Service Groups** window:



- a Select the **Relationship** of **online local**.
- b Select the **Dependency Type** of **hard**.
- c Click **OK**.

Adding the nodes from the secondary zone to the RDC


Configuration of the systems in the Primary Zone (zone 0) is complete. The nodes in the Secondary Zone (zone 1) can now be added to the RDC configuration.

Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

To add the nodes from the secondary zone to the RVG

- 1 From the active node of the cluster in the primary zone, click **Start > All Programs > Veritas > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the **Welcome page**, and click **Next**.
- 3 In the **Wizard Options** dialog box:
 - a Click **Modify an existing replication service group**. The existing replication service group is selected, by default.
 - b Click **Next**.
- 4 If a VCS notice message appears, asking if you want to continue, click **Yes**.
- 5 Specify the system priority list:



- a In the **Available Cluster Systems** box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group's system list.
 To remove a node from the service group's system list, click the node in the **Systems in Priority Order** box, and click the left arrow icon.

- b

To change the priority of a node in the system list, click the node in the **Systems in Priority Order** box, then click the up and down arrow icons. The node at the top of the list has the highest priority.
- c

Click **Next**.
- 6

If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
- 7


Review the Disk Group and Replicated Volume Group Configuration and click **Next**.
- 8

In the IP Resource Options dialog box, select **Modify IP resource** and click **Next**.
- 9

If a VCS error appears, click **OK**.
- 10

In the Network Configuration dialog box, verify that the selected adapters are correct and click **Next**.
- 11

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.
- d

Click **Next** to modify the replication service group.
- 12

When prompted, click **Yes** to modify the service group.
- 13

Click **Finish**.

Configuring secondary zone nodes in the RVG service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Configuring the IP resources for failover

Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.


Also modify the TCP/IP addresses on all the nodes in both zones, to include the IP address for replication.

In the event of a system or SQL 2005 failure, VCS attempts to fail over the SQL Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone.

To modify the IP resources in the RVG service group


- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Resources** tab.
- 3 Right-click the RVG IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.

4 In the Edit Attributes window:



- a Select **Per System**.
 - b Select the first node in the primary zone and enter the virtual IP address for the primary zone.
 - c Select the second node in the primary zone and enter the virtual IP address for the primary zone (the same IP address as the first node).
 - d Repeat for all nodes in the primary zone.
 - e Select the first node in the secondary zone (SYSTEM3) and enter the virtual IP address for the secondary zone.
 - f Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - g Repeat for all nodes in the secondary zone.
 - h Click **OK**.
- 5 In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone

have the same IP address, that is different from the IP address for the primary zone.



6 Close the Properties View window.

To modify the TCP/IP addresses on the primary nodes to include the replication IP

- 1 On the primary node in the primary zone, click **Start > Control Panel > Network Connections > Public**.
- 2 In the Public Status window, click the **Properties** button.
- 3 In the Public Properties window, select **Internet Protocol (TCP/IP)**.
- 4 Click the **Properties** button.
- 5 In the Internet Protocol (TCP/IP) Properties window, click the **Advanced** button.
- 6 In the Advanced TCP/IP Settings window, in the IP addresses area, click the **Add** button.
- 7 In the TCP/IP Address window, enter the IP address of the IP for replication, and the subnet mask.
- 8 Click **Add**.
- 9 In the Advanced TCP/IP Settings window, click **OK**.
- 10 In the Internet Protocol (TCP/IP) Properties window, click **OK**.
- 11 In the Public Properties window, click **OK**.
- 12 In the Public Status window, click **Close**.
- 13 Repeat these steps for the primary node in the secondary zone.

Adding the nodes from the secondary zone to the SQL Server service group

Use the SQL Server Agent Configuration Wizard to add the nodes from the secondary zone to the SQL Server service group.

To add the nodes from the secondary zone to the SQL Server service group

- 1 Select **Start > All Programs > Veritas > Veritas Cluster Server > SQL Server Enterprise Agent > SQL Configuration Wizard**.
- 2 In the Select Configuration Option dialog box, select **MS SQL Server - Service Group Configuration**. Select **Edit**, and click **Next**.
- 3 Verify that you have met the prerequisites listed and click **Next**.
- 4 Select the service group to be modified (INST1_SG) and click **Next**.
- 5 If a VCS notice message appears indicating that resources are online, click **Yes** to continue.
- 6 On the Service Group Configuration page, select the nodes in the secondary zone, use the arrow button to move them from **Available Cluster Systems** to **System in Priority Order**.

To change the priority of a system in the **Systems in Priority Order** list, select the system and click the up and down arrow icons. Arrange the systems in priority order in as failover targets for the group. The server that needs to come online first must be at the top of the list followed by the next one that will be brought online.

This set of nodes selected for the SQL Server service group must be the same as the nodes selected for the RVG service group. Ensure that the nodes are also in the same priority order.

- 7 Click **Next**.
- 8 On the SQL Server Instance Selection page, click **Next**.
- 9 The User Databases List page summarizes the databases for this instance of SQL. Click **Next**.
- 10 On the Detail Monitoring Configuration page, clear the box in the SQL Instance List to disable monitoring, as required. Detailed monitoring is not necessary. Click **Next**.
- 11 On the Registry Replication Path page, click **Next**.
- 12 On the Virtual Server Configuration page, verify that the public adapter is used on each system and click **Next**.
- 13 In the Service Group Summary, review the service group configuration and click **Next**.

- 14 A message appears if the configuration is currently in the Read Only mode. Click **Yes** to make the configuration read and write enabled. The wizard validates the configuration and modifies it.
- 15 Click **Finish**.

Configuring the zones in the SQL Server service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the SQL Server service group

- 1 From VCS Cluster Explorer, in the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **System Zones** attribute.
- 5 Click the **Edit** icon for the **System Zones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Verifying the RDC configuration

Failover simulation is an important part of configuration testing. Test the failover by switching online nodes.

Note: This should never be tested on systems with live data. A reliable and tested backup should be available. A tested backup means that it has been tested successfully by a restore.

Switching online nodes

Switch nodes using Veritas Cluster Manager (Java Console). When you complete the procedure, you will see the online system role shift from one system to another.

To switch online nodes

- 1 Open the Veritas Cluster Manager (Java Console) (**Start > All Programs > Veritas > Veritas Cluster Manager (Java Console)**).
- 2 Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, in the File menu, click **New Cluster**. In the **New Cluster - Connectivity Configuration** window, enter the computer name in the **Host name** field and click **OK**.
- 3 In the **Machinename - Login window**, enter your user name and password in the respective fields and click **OK**.
- 4 Right-click the service group in the left pane, and select an alternate system name from the **Switch To** entry.
- 5 In the **Question** dialog box, click **Yes** to confirm you do want to switch the service group to the other node.

Disaster Recovery

This section includes the following chapters:

- [Disaster recovery for SQL: Overview](#)
- [Deploying disaster recovery: New SQL Server 2000 installation](#)
- [Deploying disaster recovery: New SQL Server 2005 installation](#)
- [Configuring an MSDTC service group for disaster recovery](#)
- [Testing fault readiness by running a fire drill](#)

Disaster recovery for SQL: Overview

This chapter includes the following topics:

- [“What is a disaster recovery solution?”](#) on page 465
- [“What needs to be protected in a SQL Server environment?”](#) on page 466
- [“Typical disaster recovery configuration”](#) on page 467

What is a disaster recovery solution?

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.

A disaster recovery (DR) solution is a series of procedures which you can use to safely and efficiently restore application user data and services in the event of a catastrophic failure. A typical DR solution requires that you have a source or *primary site* and a destination or *secondary site*. The user application data on the primary site is replicated to the secondary site, using VVR. The cluster on the primary site provides data and services during normal operations; in the event of a disaster at the primary site and failure of the cluster, the secondary site provides the data and services.

This disaster recovery solution section includes procedures for installing and configuring clustered and replicated Microsoft SQL Server environments. The final chapter describes the fire drill procedure used to test the fault readiness of an established DR environment.

Why implement a disaster recovery solution

A DR solution is vital for businesses that rely on the availability of data. A well-designed DR solution prepares a business for unexpected disasters and provides the following benefits in a DR situation:

- Minimizes economic loss due to the unavailability or loss of data.
- Provides a plan for the safe and orderly recovery of data in the event of a disaster.
- Ensures safe and efficient recovery of data and services.
- Minimizes any decision making during DR.
- Reduces the reliance on key individuals.

Strategically planning a DR solution provides businesses with affordable ways to meet their service level agreements, comply with government regulations, and minimize their business risks.

Note: A DR solution requires a well-defined backup strategy. Refer to Veritas NetBackup product documentation for information on configuring backup.

Understanding replication

The term replication refers to the use of a tool or service to automate the process of maintaining a consistent copy of data from a designated source (primary site) on one or more remote locations (secondary sites).

In the event that the primary site data center is destroyed, the application data is readily available at the remote site, and the application can be restarted at the remote site. For more information on VVR refer to the *Veritas Volume Replicator Administrator's Guide*.

What needs to be protected in a SQL Server environment?

The following components of a SQL Server environment must be protected in the event of a disaster:

- **User Databases:** The most critical component in any SQL Server implementation is the user data that is stored in user-defined databases.
- **Logins:** Logins allow clients to connect to SQL Server and execute queries on user data. Logins are stored in the master database and each of the user-defined databases.


- **Jobs:** Jobs are a set of scheduled tasks that maintain SQL Server databases. The job configuration is stored in the msdb system database.
- **Alerts:** Alerts are actions that are taken when a specific event occurs. They are used to respond to and correct errors that occur in SQL Server. The alert configuration is stored in the msdb system database.
- **Operators:** Operators are contacts that address problems occurring in SQL Server. They are notified in the event of errors. The operator configuration is stored in the msdb system database.
- **Extended Stored Procedures:** Extended stored procedures are external routines that are called from within SQL Server. They are typically stored in DLL files on the file system.
- **Other Server Extensions:** SQL Server is a very flexible database engine and it is possible to extend its functionality in several ways. These extensions are also important to the operation of the SQL Server.

Typical disaster recovery configuration

The illustration below shows a typical configuration used to protect SQL Server in a clustered disaster recovery environment.


The VVR primary site is replicating the SQL Server application data to the secondary site. The setup contains clusters at the primary and secondary sites to make both SQL Server and VVR highly available.

Figure 11-1 Disaster recovery environment



When a disaster such as an earthquake causes a failure at the primary site, the secondary site takes over the primary role and restores the application services and data to users.

Figure 11-2 Secondary site online after a disaster



Deploying disaster recovery: New SQL Server 2000 installation

This chapter covers the following topics:

- [Tasks for a new disaster recovery installation of Microsoft SQL](#)
- [Before you begin](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Setting up the secondary site: Configuring SFW HA and setting up a cluster](#)
- [Verifying your primary site configuration](#)
- [Setting up security for VVR](#)
- [Configuring disaster recovery](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [Cloning the storage on the secondary site using the DR wizard](#)
- [Installing and configuring SQL Server 2000 on the first node \(secondary site\)](#)
- [Preparing to install SQL Server on the second node \(secondary site\)](#)
- [Installing SQL Server 2000 on the second node \(secondary site\)](#)
- [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 \(optional\)](#)
- [Recovery procedures for service group dependencies](#)

Tasks for a new disaster recovery installation of Microsoft SQL

After setting up an SFW HA environment for SQL on a primary site, you can create a secondary or “failover” site for disaster recovery.

Before setting up disaster recovery at the secondary site, you must complete the high availability configuration on the primary site.

See [Chapter 4, “Deploying SFW HA for high availability: New SQL Server 2000 installation”](#) on page 35.

You can then install and configure the high availability and SQL components on the secondary site, with the intent of creating a parallel setup for the SQL service group on both sites.

The parallel 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 a primary site to a secondary site.

Note: If you want to create the parallel configuration manually, without cloning the storage configuration or the service group, see [Appendix A, “Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation”](#).

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.

To configure MSDTC service groups, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 787.

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

Table 12-1 Tasks for deploying disaster recovery for SQL Server 2000	
Objective	Tasks
“Before you begin” on page 473	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 477	✓ Understanding site failover in a DR environment
	✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 480	✓ Setting up the storage hardware for a cluster environment
	✓ Verifying the DNS entries for the systems on which SQL will be installed
“Setting up the secondary site: Configuring SFW HA and setting up a cluster” on page 482	✓ Installing SFW HA
	✓ Configuring the cluster using the Veritas Cluster Server Configuration Wizard
“Verifying your primary site configuration” on page 503	✓ Verifying that SQL has been configured for high availability at the primary site and that the service groups are online
“Assigning user privileges (secure clusters only)” on page 507	✓ For a secure cluster only, assigning user privileges
“Cloning the storage on the secondary site using the DR wizard” on page 508	✓ Cloning the storage configuration on the secondary site using the DR wizard

Table 12-1 Tasks for deploying disaster recovery for SQL Server 2000

Objective	Tasks
“Installing and configuring SQL Server 2000 on the first node (secondary site)” on page 512	✓ Installing and configuring SQL Server 2000 ✓ Configuring SQL services
“Preparing to install SQL Server on the second node (secondary site)” on page 515	✓ Stopping the SQL Service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Adding drive letters ✓ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2000 on the second node (secondary site)” on page 519	✓ Installing SQL Server 2000
“Setting the internal name of the clustered instance” on page 523	✓ Setting the internal name of the clustered instance
“Cloning the service group configuration from the primary to the secondary site” on page 525	✓ Cloning the service group configuration from the primary to the secondary site using the DR wizard
“Configuring replication and global clustering” on page 527	✓ Configuring VVR components and global clustering using the DR wizard
“Verifying the disaster recovery configuration” on page 534	✓ Verifying that the secondary site has been fully configured for disaster recovery

Table 12-1 Tasks for deploying disaster recovery for SQL Server 2000

Objective	Tasks
“Establishing secure communication within the global cluster (optional)” on page 535	✓ Adding secure communication between local clusters within the global cluster (optional task)
“Recovery procedures for service group dependencies” on page 537	✓ Bringing the service groups online after failover to the secondary site

Before you begin

This DR solution requires a primary site and secondary site.

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>

Supported Software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the following options:
Veritas Volume Replicator and the Global Clustering Option
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with
Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 required for all editions)
or
Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition
or
Windows 2000 Server, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server (all require SP4 with Update Rollup 1)
or - Microsoft SQL Server 2000 (64-bit) Enterprise Edition
with
Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required for both)
or - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with
Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T):
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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 477.
- 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 SQL 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 *VCS 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.
- 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

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 doing the following tasks:

- 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

During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under [“Before you begin”](#) on page 473.


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

The following figure illustrates a typical clustered VVR configuration for an active/passive configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2000 application data is stored on the volumes that are under the control of the RVG.

You can also configure disaster recovery for an active/active configuration.

See [“Active/active configuration”](#) on page 48 in [Chapter 4, “Deploying SFW HA for high availability: New SQL Server 2000 installation”](#).

Figure 12-1 Typical VVR configuration



If the Microsoft SQL Server 2000 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2000 service group
SQL_CLUS1	virtual SQL Server cluster
INST1-VS	virtual server name
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

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**.

Setting up the secondary site: Configuring SFW HA and setting up a cluster

Install SFW HA and configure the cluster at the secondary site.

Begin with verifying that the requirements are met on the secondary site:

- See “[Before you begin](#)” on page 473.
- See “[Reviewing the configuration](#)” on page 477.
- See “[Configuring the storage hardware and network](#)” on page 480.

Then continue with the following topics:

- “[Installing SFW HA](#)” on page 482
- “[Configuring the cluster](#)” on page 488

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. You must select the option to install the Veritas Cluster Server Database Agent for SQL. 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
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 12-3 Installation behavior with unsigned drivers (continued)

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
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**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Data- base Agent for SQL	Required to configure high availability for SQL Server.
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also 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 that name resolution is configured for each node.
- ✓ Set the required privileges:
 - 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 490.

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 490. 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 490.


- 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 491, proceed to the next step. Otherwise, proceed to [step 11](#) on page 493.

- 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 489, 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 499.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 500.

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 500. 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**.

Veritas Cluster Server Configuration Wizard
Notifier SMTP Configuration
Specify information about SMTP recipients

✓ Domain Selection
✓ VCW Options
✓ Create Cluster
✓ Select Components
▶ **Configure**
Summary
Finish

SMTP Server Name / IP:

Enter SMTP recipients and select a severity level for each recipient.

Recipients	Severity
admin@example.com	Information


Click '+' to add a recipient.
Click '-' to remove a recipient.

Click 'Next' to continue.

Back Next Cancel

- 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 your primary site configuration

Before you begin configuring disaster recovery, make sure that SQL Server 2000 has been configured for high availability at the primary site. If you have not yet configured SQL for high availability at the primary site, go to High Availability (HA) Configuration in the Solutions Configuration Center and follow the steps in the order shown.

See [Chapter 4, “Deploying SFW HA for high availability: New SQL Server 2000 installation”](#) on page 35.

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.

Setting up security for VVR

The following procedure applies to a disaster recovery configuration using Veritas Volume Replicator (VVR).

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 `vxscfg.exe` from the command prompt of the required machine.

The Welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
Password	Specify a password in the Password field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same 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.

Configuring disaster recovery

The Disaster Recovery (DR) wizard clones the storage configuration and SQL 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. For the cloning to work effectively, the disk group and the volumes at the primary site must be created according to the prerequisites described in the section [“Prerequisites for configuring cluster disk groups and volumes”](#) on page 66.

Although all the tasks can be performed using this single wizard, you will need to exit the wizard after cloning the storage to install SQL. 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, as you page through the wizard and enter required information, the wizard will recognize if a task has been completed and proceed to the next task without requiring you to re-enter the information.

See the following procedures:

- [“Assigning user privileges \(secure clusters only\)”](#) on page 507
- [“Cloning the storage on the secondary site using the DR wizard”](#) on page 508
- [“Installing and configuring SQL Server 2000 on the first node \(secondary site\)”](#) on page 512
- [“Preparing to install SQL Server on the second node \(secondary site\)”](#) on page 515
- [“Installing SQL Server 2000 on the second node \(secondary site\)”](#) on page 519
- [“Setting the internal name of the clustered instance”](#) on page 523
- [“Cloning the service group configuration from the primary to the secondary site”](#) on page 525
- [“Configuring replication and global clustering”](#) on page 527
- [“Verifying the disaster recovery configuration”](#) on page 534

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 SQL Server 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 SQL Server 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 the 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 Microsoft SQL Server 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 SQL instance is online. If you have launched the wizard on the system where the instance is online at the primary site, you can also specify localhost 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. ■ If the volume does not exist then a new volume will be created. ■ 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. ■ If the volume is of a greater size than that on the primary then the volume will be recreated using the appropriate size. ■ 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.

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 Edit 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 SQL Server Installation panel, do one of the following:
 - Click **Finish** to exit the wizard and proceed with installing the SQL 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 SQL 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 SQL 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 SQL Server 2000 on the first node (secondary site)

In preparation for installing Microsoft SQL Server 2000, ensure the cluster disk group is imported to the first node and the volumes are mounted. Complete the following procedures to install and configure Microsoft SQL Server 2000:

- ✓ [Installing Microsoft SQL Server](#)
- ✓ [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters) See “[Importing the cluster disk group](#)” on page 517 and “[Adding drive letters to mount the volumes](#)” on page 517.

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.


To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.

- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
 Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.




- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
 - b In the Services box, select the **SQL Server** option.
 - c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
 - d If you chose to run the service in the context of a domain user, specify the information for the user.
 - e Clear the **Auto Start Service** option.
 - f Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g Click **Next**.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
- 13 If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.


Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- a Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - b Select a service from the **Services** list.
 - c Clear the **Auto-start service when OS starts** check box.
- 2 Repeat [step b](#) and [step c](#) for all other SQL Server services that are running on the server.

Preparing to install SQL Server on the second node (secondary site)

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2000 Service”](#) on page 516
- [“Deporting the cluster disk group”](#) on page 516
- [“Importing the cluster disk group”](#) on page 517
- [“Adding drive letters to mount the volumes”](#) on page 517
- [“Renaming shared SQL Server 2000 files”](#) on page 519


Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2000 Service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM1).
- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.

- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the Actions menu, select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - a In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - b Click **OK** in the Import Dynamic Disk Group dialog box.


Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2000 on the second node (secondary site)

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 519
- [“Removing shared SQL Server files”](#) on page 522

Installing SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters) See [“Importing the cluster disk group”](#) on page 517 and [“Adding drive letters to mount the volumes”](#) on page 517.


Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.


Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.




- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

- 10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
- b In the Services box, select the **SQL Server** option.
- c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
- d If you chose to run the service in the context of a domain user, specify the information for the user.

- e** Clear the **Auto Start Service** option.
 - f** Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g** Click **Next**.
- 11** Follow the wizard instructions to complete the installation.
 - 12** When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
 - 13** If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server on the second node \(secondary site\)](#)” on page 515 and “[Installing SQL Server 2000 on the second node \(secondary site\)](#)” on page 519 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1** On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2** Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.


Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.


To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example *SYSTEM2 \ INST1*.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text "sp_helpserver"
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
"sp_dropserver '*System_Name\Instance_Name*.'"
where *System_Name\Instance_Name* is the name noted in step 3.
For example, for named instance:
"sp_dropserver 'SYSTEM2\INST1.' "
For example, for a default instance:
"sp_dropserver 'SYSTEM1.' "
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
- In the top pane, enter
"sp_addserver '*Virtual_Server_Name\Instance_Name*',
local"
For example 'INST1-VS\INST1', local for a named instance, or
'INST1-VS', local for a default instance.

- Press F5.

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 site and created the required service groups. You will also need to ensure that you have installed the application on the secondary site. After verifying launch the DR wizard again.

The DR Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

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 Solutions for Microsoft SQL Server 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.
Local attributes on Primary Cluster	Displays the service group attributes for the cluster at the Primary. These include: ■ IP Resource: consists of the IP address and the subnet mask ■ NIC Resource: is the MAC address

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. Note: 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.
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.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 787.

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 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 Solutions for Microsoft SQL Server 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. Symantec recommends excluding tempdb from replication. If you earlier moved tempdb to a separate volume in the same disk group as the system database volumes, you can exclude tempdb from replication by removing the tempdb volume from the Selected RVG Volumes pane.
Primary SRL	Select the appropriate primary Replicator Log volume from the drop-down menu. Size: Enter an appropriate log size value in the corresponding Size field. If you did not create the Replicator Log volume earlier, click Create New on the drop-down menu.
Secondary SRL	Select the appropriate secondary Replicator Log volume from the drop-down menu. Size: Enter an appropriate size value in the corresponding Size field. If you did not create the Replicator Log volume earlier, click Create New on the drop-down menu.
Start Replication after the wizard completes	Select this check box to start replication automatically after the wizard completes the necessary configurations.

- Click **Advanced Settings** to specify some additional replication properties. The options on the dialog box are described column-wise, from left to right; refer to the *Veritas Volume Replicator Administrator's Guide* for additional information on VVR replication options:

Replication Mode	Select the required mode of replication; Synchronous , Asynchronous , or Synchronous Override . The default is synchronous override.
Log Protection	Select the appropriate log protection from the list. The Off 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 minimum 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 Off. When this option is selected the High Mark Value and the Low Mark Value are disabled. Select the Fail or Override option to enable Latency protection. This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail 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 checkpoint

To 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.

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 479.

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	■ The parent remains online on the primary site.	1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the child service group only.	2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.	
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains 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, firm dependency link.

Table 12-5 Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent goes offline on the primary site.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ An alert notification at the secondary site occurs for the child service group only.	Leave the RVG group online at the primary site.
	■ The RVG group remains online.	

Table 12-5 Online, local, firm dependency link

Failure condition	Results	Action required
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains 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, hard dependency link.

Table 12-6 Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent goes offline on the primary site.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ An alert notification at the secondary site occurs for the child service group only.	Do not take the RVG group offline at the primary site.
	■ The RVG group remains online.	
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.	

Deploying disaster recovery: New SQL Server 2005 installation

This chapter covers the following topics:

- [Tasks for a new disaster recovery installation of Microsoft SQL](#)
- [Before you begin](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Setting up the secondary site: Configuring SFW HA and setting up a cluster](#)
- [Verifying your primary site configuration](#)
- [Setting up security for VVR](#)
- [Configuring disaster recovery](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [Cloning the storage on the secondary site using the DR wizard](#)
- [Installing and configuring SQL Server 2005 on the first node \(secondary site\)](#)
- [Preparing to install SQL Server 2005 on the second node \(secondary site\)](#)
- [Installing SQL Server 2005 on the second node \(secondary site\)](#)
- [Setting the internal name of the clustered instance](#)

- [Cloning the service group configuration from the primary to the secondary site](#)
- [Configuring replication and global clustering](#)
- [Establishing secure communication within the global cluster \(optional\)](#)
- [Recovery procedures for service group dependencies](#)

Tasks for a new disaster recovery installation of Microsoft SQL

After setting up an SFW HA environment for SQL on a primary site, you can create a secondary or “failover” site for disaster recovery.

Before setting up disaster recovery at the secondary site, you must complete the high availability configuration on the primary site.

See [Chapter 6, “Deploying SFW HA for high availability: New SQL Server 2005 installation”](#) on page 135.

You can then install and configure the high availability and SQL components on the secondary site, with the intent of creating a parallel setup for the SQL service group on both sites.

The parallel 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.

Note: If you want to create the parallel configuration manually, without cloning the storage configuration or the service group, see [“Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation”](#) on page 697.

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.

To configure MSDTC service groups, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 787.

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

Table 13-1 Tasks for deploying disaster recovery for SQL Server 2005

Objective	Tasks
“ Before you begin ” on page 546	✓ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 550	✓ Understanding site failover in a DR environment ✓ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 553	✓ Setting up the storage hardware for a cluster environment ✓ Verifying the DNS entries for the systems on which SQL will be installed
“ Setting up the secondary site: Configuring SFW HA and setting up a cluster ” on page 555	✓ Installing SFW HA ✓ Configuring the cluster using the Veritas Cluster Server Configuration Wizard
“ Verifying your primary site configuration ” on page 576	✓ Verifying that SQL has been configured for high availability at the primary site and that the service groups are online
Assigning user privileges (secure clusters only)	✓ For a secure cluster only, assigning user privileges
“ Cloning the storage on the secondary site using the DR wizard ” on page 580	✓ Cloning the storage configuration on the secondary site using the DR wizard

Table 13-1 Tasks for deploying disaster recovery for SQL Server 2005

Objective	Tasks
“Installing and configuring SQL Server 2005 on the first node (secondary site)” on page 585	✓ Installing and configuring SQL Server 2005 ✓ Configuring SQL services
“Preparing to install SQL Server 2005 on the second node (secondary site)” on page 589	✓ Stopping the SQL Service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Adding drive letters ✓ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2005 on the second node (secondary site)” on page 592	✓ Installing SQL Server 2005
“Setting the internal name of the clustered instance” on page 596	✓ Setting the internal name of the clustered instance
“Cloning the service group configuration from the primary to the secondary site” on page 598	✓ Cloning the service group configuration from the primary to the secondary site using the DR wizard
“Configuring replication and global clustering” on page 601	✓ Configuring VVR components and global clustering using the DR wizard
“Verifying the disaster recovery configuration” on page 607	✓ Verifying that the secondary site has been fully configured for disaster recovery

Table 13-1 Tasks for deploying disaster recovery for SQL Server 2005

Objective	Tasks
“Establishing secure communication within the global cluster (optional)” on page 608	✓ Adding secure communication between local clusters within the global cluster (optional task)
“Recovery procedures for service group dependencies” on page 610	✓ Bringing the service groups online after failover to the secondary site

Before you begin

This DR solution requires a primary site and secondary site.

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 13-2 estimates disk space requirements for SFW HA.

Table 13-2 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) with the Veritas Cluster Server Database Agent for Microsoft SQL
- Microsoft SQL servers and their operating systems:

- Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SP1 required)
with
- Windows Server 2003 (32-bit) Enterprise Edition or Datacenter Edition (SP1 required)
or
- Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
or
- Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SP1 required)
with
- Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required)
or
- Windows Server 2003 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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 549.
- 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 SQL 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.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

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

During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under “[Before you begin](#)” on page 546.


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

[Figure 13-1](#) illustrates a typical clustered VVR configuration for an active/passive configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2005 application data is stored on the volumes that are under the control of the RVG.

You can also configure disaster recovery for an active/active configuration.

See “[Active/active configuration](#)” on page 148 in [Chapter 6, “Deploying SFW HA for high availability: New SQL Server 2005 installation”](#).

Figure 13-1 Typical VVR configuration



If the Microsoft SQL Server 2005 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2005 service group
SQL_CLUS1	virtual SQL Server cluster
INST1-VS	virtual server name
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

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**.

Setting up the secondary site: Configuring SFW HA and setting up a cluster

- Install SFW HA and configure the cluster at the secondary site.
- Begin with verifying that the requirements are met on the secondary site:
- See “[Before you begin](#)” on page 546.
 - See “[Reviewing the configuration](#)” on page 550,
 - See “[Configuring the storage hardware and network](#)” on page 553.
- Then continue with the following topics:
- “[Installing SFW HA](#)” on page 555
 - “[Configuring the cluster](#)” on page 561

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. You must select the option to install the Veritas Cluster Server Database Agent for SQL. 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 13-3](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 13-3 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 13-3 Installation behavior with unsigned drivers (continued)

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
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**. Be sure to select the following as appropriate for your installation:




Veritas Cluster Server Data- Required to configure high availability for SQL Server.
base Agent for SQL

Client Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also 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 that name resolution is configured for each node.
- ✓ Set the required privileges:
 - 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 563.

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 563. 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 563.


- 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 564, proceed to the next step. Otherwise, proceed to [step 11](#) on page 566.

- 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 562, 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:

The screenshot shows the 'Veritas Cluster Server Configuration Wizard' window. The title bar reads 'Veritas Cluster Server Configuration Wizard'. The main window has a left-hand navigation pane with a tree view containing the following items: 'Domain Selection', 'VCW Options', 'Cluster Details', 'Cluster Selection', 'Validate Systems', 'Edit Options', 'NIC Selection', 'Service Account', 'Security', 'Summary', and 'Finish'. The 'Security' item is currently selected and highlighted in yellow. The main content area is titled 'Configure Security Service Option' with the subtitle 'Select security options for the cluster'. It contains two radio button options. The first option is 'Use Single Sign-on' with the description 'Configure the Symantec Product Authentication Service for secure communication in the cluster.' The second option is 'Use VCS User Privileges' with the description '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.' The 'Use VCS User Privileges' option is selected. Below this description are three text input fields: 'User Name:' (containing 'admin'), 'Password:' (containing 'password'), and 'Re-enter Password:' (containing 'password'). At the bottom right of the window are three buttons: 'Back', 'Next', and '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 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 572.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 573.

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 573. 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 your primary site configuration

Before you begin configuring disaster recovery, make sure that SQL Server 2005 has been configured for high availability at the primary site. If you have not yet configured SQL for high availability at the primary site, go to High Availability (HA) Configuration in the Solutions Configuration Center and follow the steps in the order shown.

See [Chapter 6, “Deploying SFW HA for high availability: New SQL Server 2005 installation”](#) on page 135.

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.

Setting up security for VVR

The following procedure applies to a disaster recovery configuration using Veritas Volume Replicator (VVR).

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 `vxscfg.exe` from the command prompt of the required machine.

The Welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
Password	Specify a password in the Password field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same 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.

Configuring disaster recovery

The Disaster Recovery (DR) wizard clones the storage configuration and SQL service group configuration for the SQL instance from the primary site to the secondary site. It also configures VVR replication settings and connects the clusters into a global cluster. For the cloning to work effectively, the disk group and the volumes at the primary site must be created according to the prerequisites described in the section [“Configuring cluster disk groups and volumes for SQL Server 2005”](#) on page 164.

Although all the tasks can be performed using this single wizard, you will need to exit the wizard after cloning the storage to install SQL. 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, as you page through the wizard and enter required information, the wizard will recognize if a task has been completed and proceed to the next task without requiring you to re-enter the information.

Configuring disaster recovery includes the following tasks:

- [“Assigning user privileges \(secure clusters only\)”](#) on page 580
- [“Cloning the storage on the secondary site using the DR wizard”](#) on page 580
- [“Installing and configuring SQL Server 2005 on the first node \(secondary site\)”](#) on page 585
- [“Preparing to install SQL Server 2005 on the second node \(secondary site\)”](#) on page 589
- [“Installing SQL Server 2005 on the second node \(secondary site\)”](#) on page 592
- [“Setting the internal name of the clustered instance”](#) on page 596
- [“Cloning the service group configuration from the primary to the secondary site”](#) on page 598
- [“Configuring replication and global clustering”](#) on page 601
- [“Verifying the disaster recovery configuration”](#) on page 607

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 SQL Server 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 SQL Server 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 for a particular service group present at the primary site on to the secondary site. To do this successfully, the systems at the 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 Microsoft SQL Server 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 SQL instance 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. ■ If the volume does not exist then a new volume will be created. ■ 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. ■ If the volume is of a greater size than that on the primary then the volume will be recreated using the appropriate size. ■ 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.

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 Edit 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 (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 11 On the SQL Server Installation panel, do one of the following:
 - Click **Finish** to exit the wizard and proceed with installing the SQL 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 SQL 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 SQL 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 SQL Server 2005 on the first node (secondary site)

Complete the following procedures to install and configure Microsoft SQL Server 2005:

- ✓ [Installing SQL Server 2005 on the first node](#)
- ✓ [Setting SQL Server 2005 services to manual start](#)

Installing SQL Server 2005 on the first node

Install Microsoft SQL Server 2005 on the first node for the SQL Server instance using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.


To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.

- 5
- Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
- Expand **Database Services**, click **Data Files**, and click **Browse**.




-
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



-
- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6
- In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.

- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
 Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.


Setting SQL Server 2005 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.

- 3
- In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4
- In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5
- Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node (secondary site)

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2005 Service”](#) on page 589
- [“Deporting the cluster disk group”](#) on page 590
- [“Importing the cluster disk group”](#) on page 590
- [“Adding drive letters to mount the volumes”](#) on page 590
- [“Renaming shared SQL Server 2005 files”](#) on page 591

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2005 Service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

In order to install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM3) and then import it to the desired node (SYSTEM4).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM3).
- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM4).

To import a cluster disk group


- 1 From the **Actions** menu select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - a In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - b Click **OK** in the **Import Dynamic Disk Group** dialog box.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2005 on the second node (secondary site)

Follow the procedures provided in this section to install and configure SQL Server on additional nodes for this SQL instance:

- [“Installing SQL Server on the second node”](#) on page 592
- [“Removing shared SQL Server files”](#) on page 595

Installing SQL Server on the second node

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters). See [“Importing the cluster disk group”](#) on page 590 and [“Adding drive letters to mount the volumes”](#) on page 590.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.


Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.


To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
 Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
 If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user, password, and domain.
- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.

9 Follow the wizard instructions to complete the installation. Reboot if prompted.

10 Install any SQL service packs or hotfixes if required.

11 Set SQL services to manual start.

Refer to [“Setting SQL Server 2005 services to manual start”](#) on page 587 for the procedure.

Repeat the procedures described in [“Preparing to install SQL Server 2005 on the second node \(secondary site\)”](#) on page 589 and [“Installing SQL Server 2005 on the second node \(secondary site\)”](#) on page 592 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, INST1-VS\INST1).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.


To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.


To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).

- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver
 - Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, *SYSTEM1 \ INST1*. (For a default instance, you see only *System_Name*.)
- 4 Delete the contents in the upper pane.

- 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
sp_dropserver "System_Name\Instance_Name"
where **System_Name\Instance_Name** is the name noted in [step 3](#).
For example, for a named instance:
`sp_dropserver "SYSTEM1\INST1"`
For example, for a default instance:
`sp_dropserver "SYSTEM1"`
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.
- 9 Stop the SQL instance on the node.
See "[Stopping the SQL Server 2005 Service](#)" on page 589.

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 site on all nodes for this SQL instance and created the required service groups. You will also need to ensure that you have installed the application on the secondary site on all nodes for this SQL instance. After verifying launch the DR wizard again.

The DR Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes for this SQL instance within the cluster, simultaneously.

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 Solutions for Microsoft SQL Server 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.
Local attributes on Primary Cluster	Displays the service group attributes for the cluster at the Primary. These include: ■ IP Resource: consists of the IP address and the subnet mask ■ NIC Resource: is the MAC address
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. Note: 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.
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.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 787.

Configuring replication and global clustering

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 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 Solutions for Microsoft SQL Server 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. Symantec recommends excluding tempdb from replication. If you earlier moved tempdb to a separate volume in the same disk group as the system database volumes, you can exclude tempdb from replication by removing the tempdb volume from the Selected RVG Volumes pane.
Primary SRL	Select the appropriate primary Replicator Log volume from the drop-down menu. Size: Enter an appropriate log size value in the corresponding Size field. If you did not create the Replicator Log volume earlier, click Create New on the drop-down menu.
Secondary SRL	Select the appropriate secondary Replicator Log volume from the drop-down menu. Size: Enter an appropriate size value in the corresponding Size field. If you did not create the Replicator Log volume earlier, click Create New on the drop-down menu.
Start Replication after the wizard completes	Select this check box to start replication automatically after the wizard completes the necessary configurations.
■ Click Advanced Settings 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; Synchronous, Asynchronous, or Synchronous Override. The default is synchronous override.
Log Protection	Select the appropriate log protection from the list. The Off 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 minimum 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 Off. When this option is selected the High Mark Value and the Low Mark Value are disabled. Select the Fail or Override option to enable Latency protection. This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail 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.

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 553.

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 13-4 Online, local, soft dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent remains online on the primary site.	1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the child service group only.	2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.	
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains 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, firm dependency link.

Table 13-5 Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent goes offline on the primary site.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ An alert notification at the secondary site occurs for the child service group only.	Leave the RVG group online at the primary site.
	■ The RVG group remains online.	
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains 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, hard dependency link.

Table 13-6 Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent goes offline on the primary site.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ An alert notification at the secondary site occurs for the child service group only.	Do not take the RVG group offline at the primary site.
	■ The RVG group remains online.	

Table 13-6 Online, local, hard dependency link (continued)

Failure condition	Results	Action required
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.	

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 *FDnn*. The wizard renames the fire drill service group resources by adding the prefix *FDnn* 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 *FDnn*.

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 14-1 outlines the high-level objectives and the tasks to complete each objective.

Table 14-1 Tasks for configuring and running fire drills

Objective	Tasks
“Prerequisites for a fire drill” on page 617	✓ Verifying hardware and software prerequisites
“Preparing the fire drill configuration” on page 619	✓ Using the wizard to prepare the initial fire drill configuration
“Running a fire drill” on page 622	✓ Using the wizard to run the fire drill ✓ Performing your own tests of the application to confirm that it is operational
“Restoring the fire drill system to a prepared state” on page 623	✓ 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
“Deleting the fire drill configuration” on page 624	✓ Using the wizard to remove the fire drill configuration

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	Creates the configuration required to run a fire drill. This step takes some time as the wizard prepares the mirrors for the snapshots. If this option is unavailable, the fire drill configuration already exists on the specified system. See “Preparing the fire drill configuration” on page 619.
Run Fire Drill	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. 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. See “Running a fire drill” on page 622.
Restore to Prepared State	Restores the fire drill configuration for another fire drill or to prepare the fire drill configuration for deletion. This option becomes available once a fire drill has been run. The wizard snaps back the snapshot mirrors to reattach to the original volumes and takes the fire drill service group offline. See “Restoring the fire drill system to a prepared state” on page 623.
Delete Fire Drill Configuration	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. 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. See “Deleting the fire drill configuration” on page 624.

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 617.

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 Microsoft SQL**, 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 622.

- 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 623.
- 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 Microsoft SQL**, 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 623.
- 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 623.

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 Microsoft SQL**, 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 623.

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 Microsoft SQL**, 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**.

Appendices

This section includes the following appendices:

- [Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation](#)
- [Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation](#)
- [Configuring the DR components \(VVR and GCO\) without using the DR wizard](#)
- [Configuring an MSDTC service group for disaster recovery](#)

Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation

Note: This chapter covers an earlier method of deploying disaster recovery. The 5.0 release provides a new method, which uses the Solutions Configuration Center and the new Disaster Recovery (DR) wizard to clone storage configuration and service groups. See “[Deploying disaster recovery: New SQL Server 2000 installation](#)” on page 469.

After setting up a SFW or SFW HA environment for Microsoft SQL Server on a primary site, you must create a secondary or “failover” site for disaster recovery. This chapter provides information on how to install and configure the high availability and SQL components on the primary and secondary sites, with the intent of creating a parallel setup for the SQL service group on both sites. The configuration process is the same for both sites.

To configure MSDTC service groups, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 787.

To complete the process of creating a DR solution, you must proceed to “[Configuring the DR components \(VVR and GCO\) without using the DR wizard](#)” on page 765 after performing the tasks outlined in this chapter.

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

Table A-1 Tasks for manually implementing disaster recovery for a new SQL Server 2000 installation

Objective	Tasks
“Reviewing the prerequisites” on page 633	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 637	✓ Understanding active/passive configuration and site failover in a DR environment
	✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 640	✓ Setting up the storage hardware for a cluster environment
	✓ Verifying the DNS entries for the systems on which SQL will be installed
“Installing Veritas Storage Foundation for Windows” on page 642	✓ Verifying the driver signing option for the system
	✓ Installing Veritas Storage Foundation for Windows HA (automatic installation)
	✓ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR
	✓ Selecting the Global Cluster Option for VCS to enable wide-area failover
	✓ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server
	✓ Configuring VxSAS

Table A-1 Tasks for manually implementing disaster recovery for a new SQL Server 2000 installation

Objective	Tasks
“Configuring the cluster” on page 649	✓ Verifying static IP addresses and name resolution configured for each node ✓ Configuring cluster components using the Veritas Cluster Server Configuration Wizard ✓ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 665	✓ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ✓ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2000 on the first node” on page 673	✓ Installing and configuring SQL Server 2000 ✓ Configuring SQL services
“Preparing to install SQL Server on the second node” on page 676	✓ Stopping the SQL Service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Adding drive letters ✓ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2000 on the second node” on page 680	✓ Installing SQL Server 2000

Table A-1 Tasks for manually implementing disaster recovery for a new SQL Server 2000 installation

Objective	Tasks
“Setting the internal name of the clustered instance” on page 684	✓ Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 686	✓ Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating a SQL Server user-defined database” on page 690	✓ Creating volumes for a user-defined database and transaction log ✓ Creating a new user-defined database in SQL Server ✓ Adding resources for a user-defined database in VCS
“Verifying the cluster configuration” on page 693	✓ Simulating failover ✓ Switching online nodes
“Creating a parallel environment on the secondary site” on page 694	✓ Reviewing the prerequisites ✓ Reviewing the configuration ✓ Configuring the network and storage ✓ Installing SFW HA ✓ Configuring the cluster ✓ Configuring disk groups and volumes for SQL
“Installing DR components on primary and secondary sites” on page 695	✓ Completing the tasks outlined in “Configuring the DR components (VVR and GCO) without using the DR wizard” on page 765

Reviewing the prerequisites

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

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 Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the following options:
Veritas Volume Replicator and the Global Clustering Option
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with
Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 required for all editions)
or
Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition
or
Windows 2000 Server, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server (all require SP4 with Update Rollup 1)
or
 - Microsoft SQL Server 2000 (64-bit) Enterprise Edition
with
Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required for both)
or
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)

with

Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T):
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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 636.
- 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 SQL 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 *VCS 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.
- 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

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 doing the following tasks:

- 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.

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

Reviewing the configuration


During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under [“Reviewing the prerequisites”](#) on page 633.

The following figure illustrates a typical clustered VVR configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2000 application data is stored on the volumes that are under the control of the RVG.

Figure A-1 Typical VVR configuration



If the Microsoft SQL Server 2000 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2000 service group
SQL_CLUS1	virtual SQL Server cluster
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

Configuring the storage hardware and network

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent 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 Veritas Storage Foundation 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. You must select the option to install the Veritas Cluster Server Database Agent for SQL. 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**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL	Required to configure high availability for SQL Server.
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also 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 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 `vxsasscfg.exe` from the command prompt of the required machine.

The Welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
Password	Specify a password in the Password field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same 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

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 that name resolution is configured for each node.
- ✓ Set the required privileges:
 - 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 651.

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 651. 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 651.


- 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 652, proceed to the next step. Otherwise, proceed to [step 11](#) on page 654.

- 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 650, 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 660.
If you chose to configure global cluster components, proceed to “[Configuring the wide-area connector process for global clusters](#)” on page 664.
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**.

Veritas Cluster Server Configuration Wizard

Notifier SNMP Configuration
Specify information about SNMP console

✓ Domain Selection
✓ VCW Options
✓ Create Cluster
✓ Select Components
▶ **Configure**
Summary
Finish

Enter name or IP of the SNMP console and severity level for each

SNMP Console	Severity
snmpserv	Information
snmpserv1	SevereError

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.

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**.

Veritas Cluster Server Configuration Wizard

Notifier SMTP Configuration
Specify information about SMTP recipients

✓ Domain Selection
✓ VCW Options
✓ Create Cluster
✓ Select Components
► Configure
Summary
Finish

SMTP Server Name / IP: SMTPServer

Enter SMTP recipients and select a severity level for each recipient.

Recipients	Severity
admin@example.com	Information


Click '+' to add a recipient.
Click '-' to remove a recipient.

Click 'Next' to continue.

Back Next Cancel

- 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 664. 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 cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server use database files and log files on separate volumes.
- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files
- INST1_REPLOG: Contains the VVR Storage Replicator Log.

Caution: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Use the following procedures to create the appropriate disk groups and volumes. Note that instructions for INST1_REPLOG are provided in “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 767; you can wait until that chapter to configure VVR and its related volumes.

The first time you open the Veritas Enterprise Administrator (VEA), it will automatically connect to the localhost. It will also display the names of other hosts that are connected to VEA as nodes in the left side tree display.


Creating a cluster disk group

Create a cluster disk group on the first node of the cluster.

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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_DB1_VOL:** contains the user database files
- **INST1_DB1_LOG:** contains the user database log files
- **INST1_REPLOG:** Contains the VVR Storage Replicator Log.

Caution: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.


Refer to the steps below for the **INST1_DATA_FILES** and **INST1_REGREP_VOL** volumes. You can create the **INST1_DB1_VOL** and **INST1_DB1_LOG** volumes at this time or during the process of “[Creating a SQL Server user-defined database](#)” on page 690. You can create the **INST1_REPLOG** volume at this time or during the process of “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 767.

Note: To ensure that the drive letters you assign to the new volumes will always be available on all nodes, assign drive letters starting in the middle of the alphabet. This way when drive letters are assigned as additional internal devices are added to a node there will not be a conflict with the volume drive letters.

To create dynamic volumes


- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.




- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the parameters of the volume.



- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the **Mirror Info** area, select the appropriate mirroring options.
 - Verify that **Enable Logging** is not selected.
 - Click **Next**.
- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on


all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
 - For the Replicator Log volume only: Clear the Format this volume check box.
 - Click **NTFS**.
 - Select an allocation size or accept the Default.
 - The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 11 Click **Finish** to create the new volume.
- 12 Repeat these steps to create additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Installing and configuring SQL Server 2000 on the first node

In preparation for installing Microsoft SQL Server 2000, ensure the cluster disk group is imported to the first node and the volumes are mounted. Complete the following procedures to install and configure Microsoft SQL Server 2000:

- ✓ [Installing Microsoft SQL Server](#)
- ✓ [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters) See “[Importing the cluster disk group](#)” on page 678 and “[Adding drive letters to mount the volumes](#)” on page 678.

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.


To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.

- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.




- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

- 10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
 - b In the Services box, select the **SQL Server** option.
 - c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
 - d If you chose to run the service in the context of a domain user, specify the information for the user.
 - e Clear the **Auto Start Service** option.
 - f Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g Click **Next**.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
- 13 If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.


Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- a Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - b Select a service from the **Services** list.
 - c Clear the **Auto-start service when OS starts** check box.
- 2 Repeat [step b](#) and [step c](#) for all other SQL Server services that are running on the server.

Preparing to install SQL Server on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2000 Service”](#) on page 677
- [“Deporting the cluster disk group”](#) on page 677
- [“Importing the cluster disk group”](#) on page 678
- [“Adding drive letters to mount the volumes”](#) on page 678
- [“Renaming shared SQL Server 2000 files”](#) on page 680


Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2000 Service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM1).
- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.

- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the Actions menu, select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - a In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - b Click **OK** in the Import Dynamic Disk Group dialog box.


Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2000 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 680
- [“Removing shared SQL Server files”](#) on page 683

Installing SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the additional node and the volumes are mounted (are assigned drive letters). See [“Importing the cluster disk group”](#) on page 678 and [“Adding drive letters to mount the volumes”](#) on page 678.


Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.


Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000


- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
- b In the Services box, select the **SQL Server** option.
- c In the Service Settings box, select the type of user account. Select **Use the Local System account** if you do not want to use SQL log shipping for SQL log replication. Otherwise, select **Use a Domain User account**.
- d If you chose to run the service in the context of a domain user, specify the information for the user.

- e Clear the **Auto Start Service** option.
 - f Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g Click **Next**.
- 11 Follow the wizard instructions to complete the installation.
 - 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
 - 13 If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server on the second node](#)” on page 676 and “[Installing SQL Server 2000 on the second node](#)” on page 680 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.


Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.


To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example `SYSTEM2\INST1`.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
`“sp_dropserver ‘System_Name\Instance_Name.’”`
 where *System_Name\Instance_Name* is the name noted in step 3.
 For example, for named instance:
`“sp_dropserver ‘SYSTEM2\INST1.’”`
 For example, for a default instance:
`“sp_dropserver ‘SYSTEM1.’”`
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
- In the top pane, enter
`“sp_addserver ‘Virtual_Server_Name\Instance_Name’,
 local”`
 For example ‘INST1-VS\INST1’, local for a named instance, or
 ‘INST1-VS’, local for a default instance.

- Press F5.

Configuring the VCS SQL Server service group


The Enterprise Agent Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

- ✓ Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2000, is installed on all cluster nodes. See [“Configuring the cluster”](#) on page 649.
- ✓ Verify you have configured a VCS cluster using VCS Configuration Wizard (VCW). See [“Configuring the cluster”](#) on page 649.
- ✓ Verify you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- ✓ Verify the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- ✓ Verify the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- ✓ Stop the SQL 2000 Server service for the SQL instance. See [“Stopping the SQL Server 2000 Service”](#) on page 677.
- ✓ Assign a unique virtual server name to SQL Server 2000. While it must be unique within the cluster, it must be the same as the virtual name of the SQL Server in the primary site.
- ✓ Assign a unique virtual IP address to the SQL Server 2000 instance.

To create a SQL Server service group on the cluster


- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.

5 Specify the service group name and system list:




- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
 - 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 Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment. Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

- 8
- In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 687 are available for selection.
 - Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
 - If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9
- If you want to configure detail monitoring for additional instances, repeat [step 8](#) for all the instances for which detail monitoring will be configured.
- 10
- Click **Next**.

- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
- Click **Next**.

- 13 In the Service Group Summary, review 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.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 787.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- ✓ Create volumes for a user-defined SQL Server database and its transaction log.
- ✓ Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- ✓ Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 668 for information on how to use the VEA console to create a volume.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.
- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.

Adding VMDg and MountV resources

Before running the SQL Server Agent wizard to add the VMDg and MountV resources:

- ✓ Make sure the SQL Server resources are online.
- ✓ Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas**

Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard.

- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.

The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 787.

Verifying the cluster configuration

Failover simulation is an important part of configuration testing.

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.

Creating a parallel environment on the secondary site

After setting up a SFW HA environment on the primary site, use the guidelines in this chapter to complete the same tasks on the secondary site:

- ✓ [“Reviewing the prerequisites”](#) on page 633
- ✓ [“Reviewing the configuration”](#) on page 637
- ✓ [“Configuring the storage hardware and network”](#) on page 640
- ✓ [“Configuring the cluster”](#) on page 649
- ✓ [“Configuring the cluster”](#) on page 649
- ✓ [“Configuring cluster disk groups and volumes”](#) on page 665

During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:

- Cluster disk group name
- Volume sizes
- Volume names
- Drive letters
- ✓ [“Installing and configuring SQL Server 2000 on the first node”](#) on page 673
Select the same options at the secondary site as you did at the primary site.
- ✓ [“Preparing to install SQL Server on the second node”](#) on page 676
The instance name must be the same on the primary site and secondary site.
- ✓ [“Installing SQL Server 2000 on the second node”](#) on page 680
- ✓ [“Configuring the VCS SQL Server service group”](#) on page 686
The service group name, virtual computer name, and IP address must be the same on both the primary site and secondary site.

Note: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, the MSSearch resource (if present), and the SQL virtual IP resource on the primary site. The remaining resources should be online, including the VMDg resource.

Installing DR components on primary and secondary sites

To complete the process of creating a DR solution, proceed to “[Configuring the DR components \(VVR and GCO\) without using the DR wizard](#)” on page 765 after performing the tasks outlined in this chapter.

Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation

Note: This chapter covers an earlier method of deploying disaster recovery. The 5.0 release provides a new method, which uses the Solutions Configuration Center and the new Disaster Recovery (DR) wizard to clone storage configuration and service groups. See “[Deploying disaster recovery: New SQL Server 2005 installation](#)” on page 541.

After setting up a SFW or SFW HA environment for Microsoft SQL Server on a primary site, you must create a secondary or “failover” site for disaster recovery. This chapter provides information on installing and configuring the high availability and SQL components on the primary and secondary sites, with the intent of creating a parallel setup for the SQL service group on both sites. The configuration process is the same for both sites.

To configure MSDTC service groups, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 787. To complete the process of creating a DR solution, you must proceed to “[Configuring the DR components \(VVR and GCO\) without using the DR wizard](#)” on page 765 after performing the tasks outlined in this chapter.

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

Table B-1 Tasks for manually implementing disaster recovery for a new SQL Server 2005 installation

Objective	Tasks
“Reviewing the prerequisites” on page 701	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 705	✓ Understanding active/passive configuration and site failover in a DR environment ✓ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 708	✓ Setting up the storage hardware and network for a cluster environment ✓ Verifying the DNS entries for the systems on which SQL will be installed
“Installing Veritas Storage Foundation for Windows” on page 710	✓ Verifying the driver signing option for the system ✓ Installing Veritas Storage Foundation for Windows HA (automatic installation) ✓ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ✓ Selecting the Global Cluster Option for VCS to enable wide-area failover ✓ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server ✓ Configuring VxSAS

Table B-1 Tasks for manually implementing disaster recovery for a new SQL Server 2005 installation

Objective	Tasks
“Configuring the cluster” on page 717	✓ Verifying static IP addresses and name resolution configured for each node ✓ Configuring cluster components using the Veritas Cluster Server Configuration Wizard ✓ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 732	✓ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ✓ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2005 on the first node” on page 741	✓ Installing and configuring SQL Server 2005 ✓ Configuring SQL services
“Preparing to install SQL Server 2005 on the second node” on page 745	✓ Stopping the SQL Service ✓ Deporting the cluster disk group from the first node ✓ Importing the cluster disk group on an additional node ✓ Adding drive letters ✓ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2005 on the second node” on page 748	✓ Installing SQL Server 2005

Table B-1 Tasks for manually implementing disaster recovery for a new SQL Server 2005 installation

Objective	Tasks
“Setting the internal name of the clustered instance” on page 752	✓ Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 754	✓ Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating a SQL Server user-defined database” on page 759	✓ Creating volumes for a user-defined database and transaction log
	✓ Creating a new user-defined database in SQL Server
	✓ Adding resources for a user-defined database in VCS
“Verifying the cluster configuration” on page 762	✓ Simulating failover
	✓ Switching online nodes
“Creating a parallel environment on the secondary site” on page 763	✓ Reviewing the prerequisites
	✓ Reviewing the configuration
	✓ Configuring the network and storage
	✓ Installing SFW HA
	✓ Configuring the cluster
	✓ Configuring disk groups and volumes for SQL
“Installing DR components on the primary and secondary sites” on page 764	✓ Completing the tasks outlined in “Configuring the DR components (VVR and GCO) without using the DR wizard” on page 765

Reviewing the prerequisites

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

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) with the Veritas Cluster Server Database Agent for Microsoft SQL
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition (SP1 required)
with
 - Windows Server 2003 (32-bit) Enterprise Edition or Datacenter Edition (SP1 required)
or
 - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
or
 - Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition (SP1 required)
with
 - Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP1 required)
or
 - Windows Server 2003 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

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

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 704.
- 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 SQL 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.

- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

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.

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 B-2 estimates disk space requirements for SFW HA.

Table B-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


Reviewing the configuration

During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under [“Reviewing the prerequisites”](#) on page 701.

The following figure illustrates a typical clustered VVR configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2005 application data is stored on the volumes that are under the control of the RVG.

Figure B-1 Typical VVR configuration

If the Microsoft SQL Server 2005 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2005 service group
SQL_CLUS1	virtual SQL Server cluster
INST1-VS	virtual server name
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

Configuring the storage hardware and network

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent 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 Veritas Storage Foundation 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. You must select the option to install the Veritas Cluster Server Database Agent for SQL. 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 B-3 describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table B-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**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL	Required to configure high availability for SQL Server.
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also 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 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 `vxsasscfg.exe` from the command prompt of the required machine.

The Welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
Password	Specify a password in the Password field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same 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

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 that name resolution is configured for each node.
- ✓ Set the required privileges:
 - 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 719.

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 719. 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 719.


- 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 720, proceed to the next step. Otherwise, proceed to [step 11](#) on page 722.

- 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 718, 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 728.
If you chose to configure global cluster components, proceed to “[Configuring the wide-area connector process for global clusters](#)” on page 732.
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**.

Veritas Cluster Server Configuration Wizard

Notifier SNMP Configuration
Specify information about SNMP console

Enter name or IP of the SNMP console and severity level for each

SNMP Console	Severity
snmpserv	Information
snmpserv1	SevereError

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.

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**.

Veritas Cluster Server Configuration Wizard

Notifier SMTP Configuration
Specify information about SMTP recipients

✓ Domain Selection
✓ VCSW Options
✓ Create Cluster
✓ Select Components
► Configure
Summary
Finish

SMTP Server Name / IP:

Enter SMTP recipients and select a severity level for each recipient.

Recipients	Severity
admin@example.com	Information


Click '+' to add a recipient.
Click '-' to remove a recipient.

Click 'Next' to continue.

Back Next Cancel

- 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 732. 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 cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server use database files and log files on separate volumes.
- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files
- INST1_REPLOG: Contains the VVR Storage Replicator Log.

Caution: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Use the following procedures to create the appropriate disk groups and volumes. Note that instructions for INST1_REPLOG are provided in “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 767; you can wait until that chapter to configure VVR and its related volumes.

The first time you open the Veritas Enterprise Administrator (VEA), it will automatically connect to the localhost. It will also display the names of other hosts that are connected to VEA as nodes in the left side tree display.


Creating a cluster disk group

Create a cluster disk group on the first node of the cluster.

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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_DB1_VOL:** contains the user database files
- **INST1_DB1_LOG:** contains the user database log files
- **INST1_REPLOG:** Contains the VVR Storage Replicator Log.

Caution: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.


Refer to the steps below for the **INST1_DATA_FILES** and **INST1_REGREP_VOL** volumes. You can create the **INST1_DB1_VOL** and **INST1_DB1_LOG** volumes at this time or during the process of “[Creating a SQL Server user-defined database](#)” on page 759. You can create the **INST1_REPLOG** volume at this time or during the process of “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 767.

Note: To ensure that the drive letters you assign to the new volumes will always be available on all nodes, assign drive letters starting in the middle of the alphabet. This way when drive letters are assigned as additional internal devices are added to a node there will not be a conflict with the volume drive letters.

To create dynamic volumes


- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.




- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the parameters of the volume.



- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the **Mirror Info** area, select the appropriate mirroring options.
 - Verify that **Enable Logging** is not selected.
 - Click **Next**.
- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on


all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
 - For the Replicator Log volume only: Clear the Format this volume check box.
 - Click **NTFS**.
 - Select an allocation size or accept the Default.
 - The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 11 Click **Finish** to create the new volume.
- 12 Repeat these steps to create additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Installing and configuring SQL Server 2005 on the first node

In preparation for installing Microsoft SQL Server 2005, ensure the cluster disk group is imported to the first node and the volumes are mounted. Complete the following procedures to install and configure Microsoft SQL Server 2005:

- ✓ [Installing SQL Server 2005 on the first node](#)
- ✓ [Setting SQL Server 2005 services to manual start](#)

Installing SQL Server 2005 on the first node

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.


To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.

- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.

- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.


Setting SQL Server 2005 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.

- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2005 Service”](#) on page 745
- [“Deporting the cluster disk group”](#) on page 746
- [“Importing the cluster disk group”](#) on page 746
- [“Adding drive letters to mount the volumes”](#) on page 746
- [“Renaming shared SQL Server 2005 files”](#) on page 747

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2005 Service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

In order to install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM1).
- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group


- 1 From the **Actions** menu select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - a In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - b Click **OK** in the **Import Dynamic Disk Group** dialog box.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. Verify that all volumes are mounted.

To add a drive letter to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2005 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server on the second node”](#) on page 748
- [“Removing shared SQL Server files”](#) on page 751

Installing SQL Server on the second node

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters) See [“Importing the cluster disk group”](#) on page 746 and [“Adding drive letters to mount the volumes”](#) on page 746.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.


Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005


- 1 Navigate to the installation directory and launch **splash.hta**.

- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.
Refer to [“Setting SQL Server 2005 services to manual start”](#) on page 743 for the procedure.

Repeat the procedures described in [“Preparing to install SQL Server 2005 on the second node”](#) on page 745 and [“Installing SQL Server 2005 on the second node”](#) on page 748 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, INST1-VS\INST1).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.


To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.


To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).

- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
- Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver
 - Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, *SYSTEM1 \ INST1*. (For a default instance, you see only *System_Name*.)
- 4 Delete the contents in the upper pane.

- 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
`sp_dropserver "System_Name\Instance_Name"`
 where **System_Name\Instance_Name** is the name noted in [step 3](#).
 For example, for a named instance:
`sp_dropserver "SYSTEM1\INST1"`
 For example, for a default instance:
`sp_dropserver "SYSTEM1"`
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
`sp_addserver "Virtual_Server_Name\Instance_Name", local`
 For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
 For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.
- 9 Stop the SQL instance on the node.
 See ["Stopping the SQL Server 2005 Service"](#) on page 745.

Configuring the VCS SQL Server service group


The Database Agent Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

- ✓ Verify that SFW HA, along with the VCS database agent for SQL Server, is installed on all cluster nodes. See ["Installing Veritas Storage Foundation for Windows"](#) on page 710.
- ✓ Verify you have configured a VCS cluster using VCS Configuration Wizard (VCW). See ["Configuring the cluster"](#) on page 717.
- ✓ Verify you have VCS Administrator privileges. This user classification is required to create and configure a service group.


- ✓ Verify the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- ✓ Verify the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- ✓ Stop the SQL 2005 Server service for the SQL instance. See “[Stopping the SQL Server 2005 Service](#)” on page 745.
- ✓ Assign a unique virtual server name to SQL Server 2005. While it must be unique within the cluster, it must be the same as the virtual name of the SQL Server in the primary site.
- ✓ Assign a unique virtual IP address to the SQL Server 2005 instance.

To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:




- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
 - 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 Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment. Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 756 are available for selection.
 - Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
 - If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) for all the instances for which detail monitoring will be configured.
- 10 Click **Next**.
- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.

12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example `INST1-VS`. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
 - Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
 - Enter the subnet mask to which the virtual IP address belongs.
 - For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review 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.

- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 787.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- ✓ Create volumes for a user-defined SQL Server database and its transaction log.
- ✓ Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- ✓ Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 736 for information on how to use the VEA console to create a volume.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2005 database

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

Adding VMDg and MountV resources

Before running the SQL Server Agent wizard to add the VMDg and MountV resources:

- ✓ Make sure the SQL Server resources are online.
- ✓ Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.

- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.

The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 787.

Verifying the cluster configuration

Failover simulation is an important part of configuration testing.

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.

Creating a parallel environment on the secondary site

After setting up a SFW HA environment on the primary site, use the guidelines in this chapter to complete the same tasks on the secondary site:

Note: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, the MSSearch resource (if present), and the SQL virtual IP resource on the primary site. The remaining resources should be online, including the VMDg resource.

- ✓ [“Reviewing the prerequisites”](#) on page 701
 - ✓ [“Reviewing the configuration”](#) on page 705
 - ✓ [“Configuring the storage hardware and network”](#) on page 708
 - ✓ [“Configuring the cluster”](#) on page 717
 - ✓ [“Configuring cluster disk groups and volumes”](#) on page 732
During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
 - ✓ [“Installing and configuring SQL Server 2005 on the first node”](#) on page 741
Select the same options at the secondary site as you did at the primary site.
 - ✓ [“Preparing to install SQL Server 2005 on the second node”](#) on page 745
The instance name must be the same on the primary site and secondary site.
 - ✓ [“Installing SQL Server 2005 on the second node”](#) on page 748
 - ✓ [“Configuring the VCS SQL Server service group”](#) on page 754
-

Caution: Before you begin creating the SQL Server service group for the cluster at the secondary site, make sure that the SQL Server service group at the primary site is offline.

The service group name and virtual computer name must be the same on both the primary site and secondary site.

Installing DR components on the primary and secondary sites

To complete the process of creating a DR solution, proceed to “[Configuring the DR components \(VVR and GCO\) without using the DR wizard](#)” on page 765 after performing the tasks outlined in this chapter.

Configuring the DR components (VVR and GCO) without using the DR wizard

This chapter provides information on configuring the following disaster recovery components: VVR, the Veritas Volume Replicator Agent for VCS, and the Global Cluster Option. Refer to the *Veritas Volume Replicator Administrator's Guide* for additional details on VVR.

You configure these disaster recovery components for the primary and secondary sites after configuring the high availability and SQL components on both sites.

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

Table C-1 Tasks for configuring the DR components

Objective	Tasks
“Reviewing the prerequisites” on page 767	✓ Verifying HA prerequisites for DR components
“Setting up the Replicated Data Sets (RDS)” on page 767	✓ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites

Table C-1 Tasks for configuring the DR components

Objective	Tasks
“Creating the VVR RVG service group” on page 776	✓ Using the VVR Configuration Wizard to create a replication service group for the replicated volume group.
“Configuring the global cluster option for wide-area failover” on page 780	✓ Linking clusters (adding a remote cluster to a local cluster) ✓ Converting the application service group that is common to all the clusters to a global service group
“Establishing secure communication within the global cluster (optional)” on page 785	✓ Adding secure communication between local clusters within the global cluster (optional task)

Reviewing the prerequisites

The procedures assume that are deploying disaster recovery without using the DR wizard. All tasks in “[Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation](#)” on page 629, or “[Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation](#)” on page 697 must be completed prior to starting this part of the DR solution.

Setting up the Replicated Data Sets (RDS)


Set up the Replicated Data Sets (RDS) on the primary site and secondary site. You can configure an RDS using the Create RDS wizard for both sites.

- ✓ Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- ✓ Verify that the cluster disk group is imported on the primary and secondary site

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 (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name	Enter the name for the volume in the Name field.
Size	Enter a size for the volume in the Size field.
Layout	Select the desired volume layout.
Disk Selection	■ Choose Select disks automatically if you want VVR to select the disks for the Replicator Log. ■ Choose Select disks manually to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select Add to move the disks into the selected disks pane.


- Click **OK** to create the Replicator Log volume.

- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.
- 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.

AutoDCM	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.
DCM	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.
Off	Disables Replicator Log overflow protection.
Override	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. 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.
Fail	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


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 (INST1_RVG_GRP).
- b In the **Available Cluster Systems** box, click the nodes on which to configure the service group, and click the right-arrow icon to move the

nodes to the service group's system list. Make sure that the set of nodes selected for the replication service group is the same or a superset of nodes selected for the SQL Server service group. Ensure that the nodes are in the same priority order.

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

- c To change the priority of a node in the system list, click the node in the **Systems in Priority Order** box, then click the up and down arrow icons. The node at the top of the list has the highest priority.
 - d Click **Next**.
- 5 A message appears, indicating that the configuration will be changed from Read Only to Read/Write. Click **Yes** to continue.
 - 6 In the Disk Group and Replicated Volume Group Configuration dialog box:



- a Select **Configure RVGPrimary resource for selected RVG**.
 - b Select the replicated volume group for which you want to configure the RVG primary resource.
 - c Click **Next**.
- 7 In the IP Resource Options dialog box, select **Create a new IP resource** and click **Next**.

8 Enter the network information:

VVR Agent Configuration Wizard for VCS

Network Configuration
Configure your IP and NIC resources.

✓ Prerequisites
✓ Group Selection
✓ System Selection
✓ DG RVG Selection
✓ IP Options
▶ IP NIC Selection

Enter the attributes for new IP and NIC resources

Virtual IP Address: 10 . 182 . 146 . 125

Subnet Mask: 255 . 255 . 252 . 0


Specify the adapter to be used on each system.

System Name	Adapter Display Name
SYSTEM1	PXE (00-0F-1F-64-5E-3A)
SYSTEM2	PXE (00-0F-1F-66-C1-1F)

< Back Next > Cancel

- 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.

Note: The SQL service group (for example, INST1_SG) is dependent on the replication service group (for example, INST1_RVG_GRP).

Configuring the global cluster option for wide-area failover

The Global Cluster option is required to manage global clustering for wide-area disaster recovery. Creating a global cluster environment involves:

- ✓ 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 the Cluster Management Console (Single Cluster Mode) also referred to as Web Console, to perform global cluster operations; this guide only provides procedures 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.

Prerequisites

Creating a global cluster environment requires:

- All service groups properly configured and able to come online.
- The service group serving as the global group has the same unique name across all applicable clusters.
- The clusters use the same version of VCS.
- The clusters use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.
- The names of the clusters at the primary and secondary sites and the virtual IP addresses associated with them are registered in the DNS with reverse lookup.

Linking clusters: Adding a remote cluster to a local 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 service 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.

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.

Bringing a global service group online

After converting the local service group that is common to the global clusters to a global group, use the Cluster Explorer to bring the global service group online.

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**.

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.

Configuring an MSDTC service group for disaster recovery

This chapter includes the following topics:

- [“Tasks for configuring MSDTC for disaster recovery”](#) on page 787
- [“Reviewing the prerequisites”](#) on page 790
- [“Reviewing the configuration”](#) on page 790
- [“Configuring cluster disk groups and volumes”](#) on page 794
- [“Mounting drives used by the MSDTC service group”](#) on page 798
- [“Creating an MSDTC service group”](#) on page 799
- [“Creating an MSDTC client”](#) on page 800
- [“Setting up the secondary site: Creating a parallel environment”](#) on page 802
- [“Installing DR components on the primary and secondary sites”](#) on page 802

Tasks for configuring MSDTC for disaster recovery

The VCS Database Agent for SQL, which you install during SFW HA installation for Microsoft SQL high availability, includes an MSDTC agent. The MSDTC agent can provide high availability for the Microsoft Data Transaction Coordinator (MSDTC) service. The MSDTC agent comprises two parts: MSDTC client and MSDTC server.

You can configure high availability for MSDTC either before or after configuring high availability for Microsoft SQL. To configure high availability for MSDTC, you use the SQL Server Configuration Wizard to configure an MSDTC service group for the MSDTC server and to configure the MSDTC client.

To modify an existing MSDTC service group, see the *Veritas Cluster Server Database Agent for Microsoft SQL Configuration Guide*.

After configuring high availability for MDSTC on the primary site, you can configure MSDTC on the secondary site for disaster recovery.

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

Table D-1 Tasks for configuring MSDTC for disaster recovery

Objective	Tasks
“ Reviewing the prerequisites ” on page 790	✓ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 790	✓ Reviewing MSDTC service group configuration ✓ Reviewing the sample configuration
“ Mounting drives used by the MSDTC service group ” on page 798	✓ Mounting the drives used by the service group
“ Configuring cluster disk groups and volumes ” on page 794	✓ Configuring disk groups for an MSDTC service group ✓ Configuring volumes for an MSDTC service group
“ Creating an MSDTC service group ” on page 799	✓ Creating an MSDTC service group
“ Creating an MSDTC client ” on page 800	✓ Creating an MSDTC client

Table D-1 Tasks for configuring MSDTC for disaster recovery

Objective	Tasks
“Setting up the secondary site: Creating a parallel environment” on page 802	✓ Reviewing the prerequisites
	✓ Reviewing the MSDTC service group configuration
	✓ Mounting the drives used by the service group
	✓ Configuring disk groups for an MSDTC service group
	✓ Configuring volumes for an MSDTC service group
	✓ Creating an MSDTC service group
	✓ Creating an MSDTC client

Reviewing the prerequisites

You must meet the following prerequisites before creating and configuring the MSDTC service group:

- ✓ You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- ✓ You must be a Local Administrator on the node where you run the wizard.
- ✓ Verify that the VCS Database Agent for SQL is installed on all cluster nodes.
- ✓ Verify that the VCS cluster is configured using the VCS Configuration Wizard (VCW).
- ✓ Verify that the drives containing the MSDTC logs and registry replication directory are mounted on one node (the node on which you are configuring the service group) and unmounted on all other nodes.
- ✓ Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC service group.
- ✓ Assign a unique virtual server name and virtual IP address for the MSDTC server.
- ✓ Verify that the Distributed Transaction Coordinator service is stopped.

Reviewing the configuration

MSDTC servers can coexist with SQL servers on the same cluster nodes. If the MSDTC server and the SQL server are running on the same node, the MSDTC client is left in the default configuration. If the MSDTC server is not configured on the same node as the SQL server, then the MSDTC client must be configured on that node using the SQL Configuration Wizard.

For instance, a SQL Server configuration in a VCS cluster might span four nodes and two sets of shared storage. The shared storage can be managed using Veritas Storage Foundation for Windows (SFW). Two configurations are possible:

- SQL server is configured on different nodes than the MSDTC server
- SQL server is configured on the same node as the MSDTC server

Figure D-1

MSDTC server configured on different nodes than SQL server (primary site)






Figure D-2

MSDTC server configured on the same node as SQL server (primary site)



For replication, the secondary site must be configured exactly the same as the primary site.

Figure D-3 Replication configuration for MSDTC service group



Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your MSDTC service group.

A disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the cluster disk group and volumes:

- ✓ Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place MSDTC files and log files on separate volumes.
- ✓ Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- ✓ Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- ✓ Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (MSDTC_DG) on shared disks and then create the following volumes:

- MSDTC_LOG: contains the MSDTC log files.
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

Creating a cluster disk group

Create a cluster disk group on the first node of the cluster.

To create a 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 On the **Welcome** page of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group you want to create.
 - Enter the name of the disk group (for example, MSDTC_DG).
 - Click **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and click the **Add** button to move them to the **Selected disks** list.
 - Click **Next**.
- 7 Review the selected disks and click **Next**.
- 8 Review the summary information and click **Finish**.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- MSDTC_LOG: contains the MSDTC log files
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the MSDTC service group.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create a volume

- 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 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 the disk group you created for the MSDTC volumes and select **New Volume**.
- 5 On the **Welcome** page of the New Volume Wizard, click **Next**.
- 6 On the **Assign Disks for Volume** page, make sure the name of the disk group that you created for MSDTC appears in the Group name field.
- 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. After selecting the necessary disks, click **Next**.
- 8 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.
 - 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.
- 9 In the **Add Drive Letter and Path** page, 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.
- 10 Specify information about the attributes for the file system that will be created:
- 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**.
- 11 Review the Summary dialog box and click **Finish**.
- 12 Repeat the procedure as necessary to create the additional volumes needed for your deployment. Create the volumes on the first node of the cluster only.

Mounting drives used by the MSDTC service group

Ensure that the shared volumes created for the MSDTC log files and registry replication information are mounted on the node where you run the SQL Server Configuration Wizard to configure the service group. The volumes must be unmounted on all other nodes in the cluster.

If you created the MSDTC disk group and volumes on another node, mounting a volume involves deporting the disk group from the other node and then importing the disk group to the node where you are configuring the service group.

Occasionally, when a disk group is imported, a drive letter may not be associated with an existing volume. If this occurs, use VEA to add a drive letter or folder path and mount the volume so that it can be seen by the operating system.

To mount volumes

- 1 If the disk group is not imported, deport it from the other node and then import it on the current node:
 - Stop all processes accessing the volumes in the disk group, and from the VEA console's tree view, right-click the disk group and select **Deport Dynamic Disk Group**. Then select **Yes**.
 - In the VEA, connect to the node where you want to import the cluster disk group. In the tree, right-click **Storage Agent**, and click **Rescan** to update the disk information on the node. Then right-click the disk group and select **Import Dynamic Group** and click **OK**.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.
- 3 In the Drive Letter and Paths dialog box, select **Add**.
- 4 Select one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder. Assign the same drive letter or mount path that was assigned when the volume was created
 - *To assign a drive letter*
In the Assign Drive Letter dialog box, select the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder option** and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.
- 6 Repeat [step 2](#) through [step 5](#) for all the volumes to be mounted.
You can now create the MSDTC service group.

Creating an MSDTC service group

MSDTC is a global resource and is accessed by more than one SQL Server service group. Symantec recommends you to configure one MSDTC service group in a VCS cluster.

To configure an MSDTC service group

- 1 If you have just configured a SQL service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Select **MSDTC Server - Service Group Configuration** and **Create**. Click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 Specify the service group name and system list.
 - Enter a name for MSDTC 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. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
 - To change a system's priority in the **Systems in Priority Order** list, select the system and click the up and down arrows. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
 - Click **Next**. If the configuration is in read-only mode, the wizard prompts you before changing it to read-write mode. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 Specify the information related to the virtual server.
 - Enter a virtual server name for the node on which the DTC service is running. Ensure that the virtual server name you enter is unique in the cluster.
 - Enter a unique virtual IP address for the MSDTC server.
 - Enter the subnet mask to which the virtual IP address belongs.

- For each system in the cluster, select the public network adapter name. Click the **Adapter Display Name** field to view the adapters associated with a system.
The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop down list.
 - Click **Next**.
- 6 Specify the MSDTC log and replication directory and click **Next**. Symantec recommends using different paths for these directories. If the directory does not exist, the wizard creates it.
 - 7 Review the service group configuration, change the resource names if desired and click **Next** to create the service group.
 - 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 the resources, if desired. 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**.
 - 8 A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes** to create the service group.
Various messages indicate the status of these commands.
 - 9 In the Configuration Complete panel, check **Bring the service group online** to bring the configured service group online. To bring the service group online later, uncheck the option.
 - 10 Click **Next** to create an MSDTC client, or click **Finish** to exit the wizard.

Creating an MSDTC client

Set the MSDTC client to run on nodes where:

- A SQL instance is configured to run
- The MSDTC server is not configured to run.

Before configuring the MSDTC configuration wizard to configure an MSDTC client:

- ✓ Verify the MSDTC service group is online in the cluster.
- ✓ Verify the node on which you run the wizard is not a part of an MSDTC service group system list.

To configure an MSDTC client

- 1 If you have just configured the MSDTC service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 In the Select Configuration Option panel, select **MSDTC Client Configuration** and click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 In the System Selection panel, specify the nodes on which the MSDTC client will be configured and click **Next**:
 - Select the nodes in the **Available Cluster Systems** list. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
 - Click the right arrow to add them to the **Selected systems** list.
 - The **Unavailable Cluster Systems** lists the nodes that have an MSDTC service group configured and are therefore not available for setting up an MSDTC client.
- 5 If the MSDTC service group is not online in the cluster, an informational message appears informing you that the wizard will bring the MSDTC service group online. Click **Yes**.
- 6 On the Specify DTC Node panel, specify the virtual DTC server name and click **Next**.
- 7 On the Configuration Complete panel, click **Finish** to exit the wizard.

Setting up the secondary site: Creating a parallel environment

After setting up a SFW HA environment on the primary site, use the guidelines in this chapter to complete the same tasks on the secondary site:

- ✓ [“Reviewing the prerequisites”](#) on page 790
- ✓ [“Reviewing the configuration”](#) on page 790
- ✓ [“Configuring cluster disk groups and volumes”](#) on page 794
During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- ✓ [“Creating an MSDTC service group”](#) on page 799
Select the same options at the secondary site as you did at the primary site.
- ✓ [“Creating an MSDTC client”](#) on page 800
The instance name must be the same on the primary site and secondary site.

Caution: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, and the SQL virtual IP resource on the primary site. The remaining resources should be online, including the VMDg resource.

The service group name, virtual computer name, and IP address must be the same on both the primary site and secondary site.

Installing DR components on the primary and secondary sites

To complete the process of creating a DR solution, proceed to [“Configuring the DR components \(VVR and GCO\) without using the DR wizard”](#) on page 765 after performing the tasks outlined in this chapter. You can configure SQL Server first if you have not already configured it, and then proceed to the DR configuration tasks.

Index

A

- active/active configuration
 - illustration 49, 149
 - IP addresses 50, 150
 - key information 54, 154
 - overview 48
 - service group requirements 104, 205
- active/passive configuration
 - illustration 47, 147
 - IP addresses 48
 - overview 47, 146

C

- cloning
 - secondary storage 508, 580
 - SQL Server 2000 service group 525
 - SQL Server 2005 service group 598
- Cluster Management Console 111
- clusters
 - assigning user privileges 507, 580
 - configuring (DR) 265, 373, 488, 561, 649, 717
 - configuring (HA) 74, 173
 - configuring disaster recovery 506, 579
 - creating a disk group (DR) 281, 389, 666, 734
 - creating a disk group (HA) 68, 167
 - deporting the cluster disk group 93, 194, 292, 401, 516, 590, 677, 746
 - importing the cluster disk group 94, 194, 293, 401, 517, 590, 678, 746
 - preconditions for a cluster disk group (DR) 280, 388, 665, 732
 - preconditions for a cluster disk group (HA) 65, 66, 164
 - preconditions for a cluster disk group (MSDTC) 224, 794
 - setting the internal name of the clustered instance (2000) 100, 299, 407, 523, 684
 - setting the internal name of the clustered instance (2005) 200, 596, 752
 - switching online nodes 351, 462

- clusters (continued)
 - verifying HA configuration at primary site 503, 576
 - verifying the SQL Server 2000 configuration 110
 - verifying the SQL Server 2005 configuration 211
- configuring
 - cluster (DR) 265, 373, 488, 561, 649, 717
 - cluster (HA) 74, 173
 - DR components (VCS) 765
 - network and storage 57, 156, 553, 708
 - network and storage for SQL Server 2000 DR 256, 364, 480, 640
 - SQL Server 2000 DR overview 255, 363, 477, 637
 - SQL Server 2000 HA overview 47
 - SQL Server 2005 DR overview 550, 705
 - SQL Server 2005 HA overview 146
 - SQL Server service group 2000 103, 301, 409, 525, 686
 - SQL Server service group 2005 205, 598, 754
- converting
 - moving data files and databases (HA) 132
 - standalone SQL Server to HA 119

D

- database
 - new database for SQL Server 2000 102, 309, 417, 690, 691
 - new database for SQL Server 2005 203, 759, 760
- disaster recovery (DR)
 - cloning secondary storage 508, 580
 - cloning SQL Server 2000 service group 525
 - cloning SQL Server 2005 service group 598
 - configuration example 52, 151
 - configuring 506, 579
 - configuring DR components 765
 - configuring GCO with DR wizard 527, 601
 - configuring replication with DR wizard 527

disaster recovery (DR) (continued)

- defined 465
- deploying for SQL Server 2000 469, 629
- deploying for SQL Server 2005 541, 697
- illustrated 467
- IP addresses 53, 153
- overview 465
- protection in a SQL Server environment 466
- rationale 466
- setting up a secondary SQL Server 2000
 - site 313, 421, 694
- setting up a secondary SQL Server 2005
 - site 579, 763
- typical configuration 467

disk groups

- cloning secondary storage 508, 580
- creating a cluster disk group (DR) 281, 389, 666, 734
- creating a cluster disk group (HA) 68, 167
- creating a disk group (MSDTC) 225, 795
- deporting the cluster disk group 93, 194, 292, 401, 516, 590, 677, 746
- importing the cluster disk group 94, 194, 293, 401, 517, 590, 678, 746
- multiple instances 67, 165
- overview 65, 164
- preconditions for a cluster disk group (DR) 280, 388, 665, 732
- preconditions for a cluster disk group (HA) 66, 164
- preconditions for a cluster disk group (MSDTC) 224, 794
- sample configuration 66, 165
- where to create 66, 164

DR wizard

- cloning secondary storage 508, 580
- cloning SQL Server 2000 service group 525
- cloning SQL Server 2005 service group 598
- configuring replication and GCO 527, 601

drive letters

- adding drive letters to mount volumes 94, 195, 293, 402, 517, 590, 678, 746

driver signing options

- resetting 65, 163, 265, 373, 488, 561, 649, 717

F**failover considerations for SQL Server 2000 48****Fire Drill Wizard**

- actions 618
- deleting the configuration 624
- overview 615
- preparing the configuration 619
- prerequisites for a fire drill 617
- restoring the prepared configuration 623
- running a fire drill 622

G**Global Cluster Option (GCO)**

- configuration tasks 765
- configuring with the DR wizard 527, 601
- prerequisites 780
- secure configuration 535, 608, 785

H**hardware, configuring 57, 156, 256, 364, 480, 553, 640, 708****high availability (HA)**

- adding HA to standalone SQL Servers 119
- defined 33
- deploying for SQL Server 2000 35
- deploying for SQL Server 2005 135
- overview for SQL Server 33
- rationale 34
- VCS SQL Server agent 34

I**installing**

- before installing SQL Server 2000 on a second node 92, 291, 400, 515, 676
- before installing SQL Server 2005 on a second node 193, 589, 745
- multiple instances 88, 187
- SFW HA (DR) 258, 366, 482, 555, 642, 710
- SFW HA (HA) 59, 158
- SQL Server 2000 on a second node (VCS) 96, 295, 403, 519, 680
- SQL Server 2000 on the first node 89, 288, 512, 673
- SQL Server 2000 tasks (DR) 469, 629
- SQL Server 2000 tasks (HA) 35
- SQL Server 2000, secondary site tasks 313, 421, 694

installing (continued)

- SQL Server 2005 on a second node (VCS) 196, 592, 748
- SQL Server 2005 on the first node 188, 396, 585, 741
- SQL Server 2005 tasks (DR) 541, 697
- SQL Server 2005 tasks (HA) 135
- SQL Server 2005, secondary site tasks 585, 763
- verifying the installation 110, 312, 420, 693, 762

internal name

- setting for the clustered instance (SQL 2000) 100, 299, 407, 523, 684
- setting for the clustered instance (SQL 2005) 200, 596, 752

IP addresses

- active/active configuration 50, 150
- active/passive configuration 48, 148
- disaster recovery configuration 53, 153

M**MSDTC**

- configuration tasks for DR (VCS) 787
- configuration tasks for HA (VCS) 219
- configuring the client (VCS) 233, 800
- creating an MSDTC service group 231, 799
- DR configuration overview 790
- HA configuration overview 221
- mounting drives 230, 798
- service group configuration (VCS) 221, 790

multiple instances

- assigning the port 103, 204
- disk groups 67, 165
- failover considerations for SQL Server 2000 48
- installation order 88, 187
- instance name 89, 188
- service groups 104, 205
- virtual server name 89, 188

N

- network, configuring 57, 156, 256, 364, 480, 553, 640, 708

O**options**

- driver signing 65, 163, 265, 373, 488, 561, 649, 717

P**port**

- assigning for multiple instances 103, 204

prerequisites

- MSDTC service group (DR) 790
- MSDTC service group (HA) 221
- SQL Server 2000 (DR) 250, 358, 473, 633
- SQL Server 2000 (HA) 43
- SQL Server 2000 service group 104
- SQL Server 2005 (DR) 546, 701
- SQL Server 2005 (HA) 142
- SQL Server 2005 service group 206
- SQL Server standalone conversion (HA) 122

primary host, defined 465**primary site**

- verifying the cluster configuration 503, 576

Q

- Query Analyzer 100, 299, 523, 684

R**replicated data clusters**

- setting up 240

replication

- configuring with the DR wizard 527, 601
- defined 466
- setting up a Replicated Data Set (RDS) 321, 429, 767

resetting

- driver signing options 65, 163, 265, 373, 488, 561, 649, 717

resources

- adding to SQL 2000 service group 116
- adding to SQL 2005 service group 217

S**sample configurations**

- SQL Server (DR) 255, 363, 478, 551, 639, 707
- SQL Server (HA) active/active 49
- SQL Server (HA) active/passive 47, 147
- standalone SQL Server conversion (HA) 127

secondary host, defined 465**secure clusters**

- assigning user privileges 507, 580

- secure GCO, establishing 535, 608, 785
- Security Services
 - configuring 80, 179, 271, 379, 494, 567, 655, 723
- service groups
 - configuring for SQL Server 2000 103
 - configuring for SQL Server 2005 205
 - dependencies 479, 537, 553, 610
 - modifying for SQL 2000 116
 - modifying for SQL 2005 217
 - prerequisites for configuring for SQL Server 2000 104
 - prerequisites for configuring for SQL Server 2005 206
 - priority order 104, 205
 - requirements for active/active
 - configurations 104, 205
 - requirements for multiple instances 104, 205
 - VCS SQL Configuration wizard 105, 207
- setting bandwidth
 - using RDS wizard 530, 604
- SFW for HA
 - installing (DR) 258, 366, 482, 555, 642, 710
 - installing (HA) 59, 158
- Solutions Configuration Center
 - context sensitivity 24
 - overview 23
 - running wizards remotely 28
 - starting 24
 - wizard descriptions 28
 - workflow for active/active configuration 55, 155
- SQL Server 2000
 - before installing on a second node 92, 291, 400, 515, 676
 - configuring a service group (VCS) 103, 301, 409, 525, 686
 - configuring network and storage 57, 256, 364, 480, 640
 - creating a new database 102, 309, 417, 691
 - DR configuration overview 255, 363, 477, 637
 - DR sample configuration 255, 363, 478, 639
 - HA configuration overview 47
 - HA sample configuration 47
 - installing on a second node (DR) 295, 403, 519, 680
 - installing on a second node (HA) 96
 - installing on the first node (DR) 288, 512, 673
 - installing on the first node (HA) 89
- SQL Server 2000 (continued)
 - new database 102
 - new installation tasks (DR) 469, 629
 - new installation tasks (HA) 35
 - overview of user-defined database 309, 417, 690
 - removing shared files 96, 295, 403, 519, 680
 - removing shared files on the second node 99, 199
 - setting services to manual start 92, 291, 399, 515, 676
 - stopping the SQL Server service 93, 292, 401, 516, 677
 - tasks for setting up a secondary site 313, 421, 694
 - user-defined database 102
- SQL Server 2005
 - before installing on a second node 193, 589, 745
 - configuring a service group 598
 - configuring a service group (VCS) 205
 - configuring network and storage 156, 553, 708
 - configuring the SQL Server service group 754
 - creating a new database 203, 760
 - DR configuration overview 550, 705
 - DR sample configuration 551, 707
 - HA configuration overview 146
 - HA sample configuration 49, 147, 149
 - installing on a second node 196, 592, 748
 - installing on the first node 188, 396, 585, 741
 - Management Studio 200, 407, 596, 752
 - new database 203
 - new installation tasks (DR) 541, 697
 - new installation tasks (HA) 135
 - overview of user-defined database 759
 - removing shared files 196, 591, 747
 - setting services to manual start 191, 587, 743
 - start a SQL Server service 200, 407, 596, 752
 - stopping the SQL Server service 193, 589, 745
 - tasks for setting up a secondary site 763
 - user-defined database 203
- SQL Server Configuration Wizard 105, 207
- SQL Server service group
 - configuring for DR 525, 598
 - configuring for HA 103, 205
- standalone SQL Server
 - configuring network and storage 128
 - converting to clustered for HA 119

- standalone SQL Server (continued)
 - installing and configuring VSFW (HA) 131
 - overview of HA configuration 126
 - preparing for conversion (HA) 130
 - sample HA configuration 127
- storage configuration
 - SQL Server 2000 DR 256, 364, 480, 640
 - SQL Server 2005 DR 553, 708
- switching online nodes 351, 462

T

- tempdb database 103, 204

U

- user privileges
 - assigning 507, 580

V

- VCS
 - configuring the cluster 173
 - configuring the cluster for DR 265, 373, 488, 561, 649, 717
 - configuring the cluster for HA 74
 - configuring the SQL Server service group 103, 205, 301, 409, 525, 686, 754
 - switching online nodes 351, 462
- VCS Configuration Wizard 74, 173
- VCS SQL Configuration Wizard 105, 207
- verifying HA configuration at primary site 503, 576
- virtual server name 89, 188
- volumes
 - adding drive letters 94, 195, 293, 402, 517, 590, 678, 746
 - considerations for VVR and disaster recovery 68, 166
 - creating on a cluster disk group 69, 168, 282, 390, 668, 736
 - multiple instances 165
 - preconditions on a cluster disk group 66, 164, 280, 388, 665, 732
 - preconditions on a cluster disk group (MSDTC) 224, 794
 - sample configuration 66, 165
- VVR
 - configuration diagram (DR) 52, 152, 478, 551, 638, 706
 - configuration diagram (MSDTC DR) 793

- VVR (continued)
 - configuration tasks 765
 - configuring replication with DR wizard 527, 601
 - moving tempdb to separate volume 103, 204
 - setting up RDS (VCS) 321, 429, 767

