

Veritas Storage Foundation[™] and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange

Windows 2000, Windows Server 2003

5.0

Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange

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 Exchange

About the solutions guides	19
About high availability	20
About disaster recovery	20
How this guide is organized	20

Chapter 2 Using the Solutions Configuration Center

About the Solutions Configuration Center	21
Starting the Configuration Center	22
Available options from the Configuration Center	22
About running the Configuration Center wizards	26
Following the workflow in the Configuration Center	27

Section 2 High Availability

Chapter 3 High availability for Exchange: Overview

What is high availability?	31
Why implement a high availability solution?	32
How the VCS application agent makes Microsoft Exchange highly available	32
Typical HA configurations for Exchange	32

Chapter 4 Deploying SFW HA for high availability: New installation

Tasks for a new HA installation of Microsoft Exchange	34
Reviewing the requirements	36
Disk space requirements	36
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	37

Reviewing the configuration	40
IP addresses required during configuration	42
Sample configuration	43
Configuring the storage hardware and network	43
Preparing the forest and domain	45
Configuring SFW HA: Prior to installing Exchange	45
Installing Veritas Storage Foundation HA for Windows	45
Setting Windows driver signing options	45
Installing Storage Foundation HA for Windows	46
Configuring VxSAS	51
Resetting the driver signing options	52
Configuring disk groups and volumes	52
Creating a disk group	54
Creating volumes	55
Managing disk groups and volumes	60
Importing a disk group and mounting a volume	60
Unmounting a volume and deporting a disk group	61
Configuring the cluster	61
Configuring Web console	72
Configuring notification	73
Installing Exchange on the first node	76
Exchange pre-installation: First node	77
Exchange installation: First node	80
Exchange post-installation: First node	80
Moving Exchange databases to shared storage	81
Installing Exchange on additional nodes	84
Exchange pre-installation: additional nodes	85
Exchange installation: additional nodes	86
Exchange post-installation: additional nodes	88
Configuring the Exchange service group for VCS	89
Prerequisites	90
Verifying the cluster configuration	96
Configuring the Cluster Management Console connection	97

Chapter 5

Deploying SFW HA for high availability: Standalone Exchange servers

Reviewing the requirements	106
Disk space requirements	106
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	106
Reviewing the configuration	110

- Configuring the network and storage 113
- Installing Veritas Storage Foundation HA for Windows 115
 - Setting Windows driver signing options 115
 - Installing Storage Foundation HA for Windows 116
 - Resetting the driver signing options 121
- Configuring disk groups and volumes 121
 - Creating a disk group 123
 - Creating volumes 124
- Managing disk groups and volumes 128
 - Importing a disk group and mounting a volume 129
 - Unmounting a volume and deporting a disk group 129
- Converting the standalone Exchange server into a “Clustered”
 - Exchange server 130
- Adding the standalone Exchange server to a cluster 132
 - Prerequisites for a new cluster 133
 - Creating a new cluster and adding nodes 134
 - Configuring Web console 145
 - Configuring notification 146
 - Prerequisites for adding nodes to an existing cluster 149
 - Adding nodes to an existing cluster 150
 - Modifying values for ClusterService group attributes 157
- Moving Exchange databases to shared storage 160
- Installing Exchange on additional nodes 163
 - Exchange pre-installation: additional nodes 165
 - Exchange installation: additional nodes 166
 - Exchange post-installation: additional nodes 168
- Configuring the Exchange service group for VCS 169
 - Prerequisites 169
- Verifying the cluster configuration 176

Chapter 6

Deploying SFW HA for high availability: Configuring a new any-to-any failover

- Reviewing the configuration 180
 - Any-to-any configuration 181
 - Sample configuration 182
- Reviewing the requirements 183
 - Disk space requirements 183
 - Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA) 183

Configuring the storage hardware and network	187
Preparing the forest and domain	188
Installing Veritas Storage Foundation HA for Windows	189
Setting Windows driver signing options	189
Installing Storage Foundation HA for Windows	190
Resetting the driver signing options	195
Configuring the cluster	195
Configuring Web console	206
Configuring notification	207
Configuring the first Exchange Virtual Server	210
Configuring disk groups and volumes	211
Managing disk groups and volumes	217
Importing a disk group and mounting a volume	217
Unmounting a volume and deporting a disk group	218
Installing Exchange on the first node	219
Moving Exchange databases to shared storage	223
Installing Exchange on additional nodes	226
Configuring the Exchange service group for VCS	231
Verifying the cluster configuration	238
Configuring another Exchange virtual server for an any-to-any failover	239
Configuring disk groups and volumes	239
Managing disk groups and volumes	241
Importing a disk group and mounting a volume	241
Unmounting a volume and deporting a disk group	242
Installing Exchange on the first node of an additional Exchange Virtual Server	242
Moving Exchange databases to shared storage	246
Specifying a common node for failover	250
Configuring the Exchange service group for an additional Exchange Virtual Server	252
Verifying the cluster configuration	258

Chapter 7

Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover

Reviewing the requirements	262
Disk space requirements	262
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	263
Reviewing the configuration	266
Any-to-any configuration	267
Sample configuration	268

- Configuring new nodes: Prior to creating additional Exchange virtual server268
 - Configuring the network and storage269
 - Installing Veritas Storage Foundation HA for Windows270
 - Setting Windows driver signing options270
 - Installing Storage Foundation HA for Windows271
 - Resetting the driver signing options276
 - Configuring the cluster276
 - Adding a node to a cluster276
 - Modifying values for ClusterService group attributes283
 - Configuring disk groups and volumes286
 - Creating a disk group288
 - Creating volumes289
 - Managing disk groups and volumes293
 - Importing a disk group and mounting a volume294
 - Unmounting a volume and deporting a disk group294
 - Installing Exchange on the new nodes295
 - Moving Exchange databases to shared storage (EVS2)299
 - Installing Exchange on additional nodes303
 - Exchange pre-installation: Additional nodes303
 - Exchange installation: Additional nodes305
 - Exchange post-installation: Additional nodes306
- Specifying a common node for failover308
 - Preparing the cluster with the any-to-any option308
 - Configuring the Exchange service group for VCS310
- Verifying the cluster configuration316

Section 3 Campus Cluster

Chapter 8 Campus cluster for Exchange: Overview

- What is a campus cluster?321
- Why implement a campus cluster?321
- What is high availability?322
- Why implement a high availability solution?322
- How the VCS application agent makes Microsoft Exchange highly available322
- Campus cluster failover using the ForceImport attribute323

Chapter 9 Deploying SFW HA for Campus Cluster: New Installation

Reviewing the requirements	327
Disk space requirements	327
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	328
Reviewing the configuration	331
Configuring the network and storage	332
Installing Veritas Storage Foundation HA for Windows	334
Setting Windows driver signing options	334
Installing Storage Foundation HA for Windows	335
Resetting the driver signing options	340
Configuring the cluster	340
Configuring Web console	351
Configuring notification	352
Configuring disk groups and volumes	355
Configuring the disks and volumes	357
Creating a dynamic (cluster) disk group	358
Creating a volume	360
Managing disk groups and volumes	364
Importing a disk group and mounting a volume	364
Unmounting a volume and deporting a disk group	365
Preparing the forest and domain	366
Installing Exchange on the first node	366
Exchange pre-installation: First node	367
Exchange installation: First node	369
Exchange post-installation: First node	370
Moving Exchange databases to shared storage	370
Installing Exchange on additional nodes	374
Exchange pre-installation: Additional nodes	374
Exchange installation: Additional nodes	376
Exchange post-installation: Additional nodes	377
Configuring the Exchange service group for VCS	379
Prerequisites	379
Modifying the IP resource in the Exchange service group	385
Verifying the campus cluster: Switching the service group	387
Possible tasks after creating the campus cluster	387
Setting the ForceImport attribute to 1 after a site failure	387

Section 4 Disaster Recovery

Chapter 10 Disaster Recovery for Exchange: Overview

What is a disaster recovery solution?	391
Why implement a DR solution?	391
Typical DR configurations for Exchange	392

Chapter 11 Deploying Disaster Recovery using the DR wizard for cloning: New Exchange Server installation

Tasks for a new disaster recovery installation of Microsoft Exchange	395
Before you begin	398
Disk space requirements	398
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	399
Reviewing the configuration	402
Supported disaster recovery configurations for service group dependencies	404
Configuring the storage hardware and network	405
Managing disk groups and volumes	406
Importing a disk group and mounting a volume	406
Unmounting a volume and deporting a disk group	407
Preparing the forest and domain	407
Setting up the secondary site: Installing SFW HA and configuring a cluster	408
Installing SFW HA	408
Setting Windows driver signing options	408
Installing Storage Foundation HA for Windows	409
Resetting the driver signing options	414
Configuring the cluster	414
Configuring Web console	425
Configuring notification	426
Verifying your primary site configuration	429
Setting up security for VVR	430
Configuring disaster recovery	432
Assigning user privileges (secure clusters only)	433
Cloning the storage on the secondary site using the DR wizard	433
Installing Exchange on the first node with DR option (secondary site)	438
Prerequisites for installing Exchange Server	438
Exchange pre-installation on first node (secondary site)	440
Exchange installation on first node (secondary site)	442
Exchange post-installation on first node (secondary site)	443

- Installing Exchange on additional nodes (secondary site) 444
 - Exchange pre-installation: Additional nodes 445
 - Exchange installation: Additional nodes 446
 - Exchange post-installation: Additional nodes 447
- Cloning the service group configuration on to the secondary site
 - using the DR wizard 449
- Configuring replication and global clustering 451
- Verifying the disaster recovery configuration 457
- Establishing secure communication within the global cluster
 - (optional) 458
- Recovery procedures for service group dependencies 460
- Possible task after creating the DR Environment: Adding a new
 - failover node 464
 - Preparing the new node 464
 - Preparing the existing DR environment 464
 - Installing Exchange on the new node 465
 - Modifying the replication and Exchange service groups 465
 - Reversing replication direction 466

Chapter 12 Deploying SFW HA for Disaster Recovery: Configuring any-to-any failover

- Tasks for deploying a new disaster recovery any-to-any configuration
 - 468
- Reviewing the configuration 470
 - Disaster recovery configuration 470
 - Any-to-any configuration 472
 - Sample any-to-any configuration for disaster recovery 473
- Configuring disaster recovery for the first Exchange virtual server 474
- Verifying your primary site configuration for an additional Exchange
 - virtual server 475
- Adding the user to the service group (secure clusters only) 475
- Configuring disaster recovery for the second Exchange virtual server ... 476
- Cloning the storage on the secondary site using the DR wizard 476
- Installing Exchange on the first node of an additional EVS (secondary
 - site) 480
 - Exchange pre-installation on first node of an additional EVS
 - (secondary site) 482
 - Exchange installation on first node of an additional EVS
 - (secondary site) 484
 - Exchange post-installation on first node of an additional EVS
 - (secondary site) 485
- Specifying a common node for failover 486

Chapter 12 Deploying SFW HA for Disaster Recovery: Configuring any-to-any failover (*continued*)

- Cloning the service group configuration on to the secondary site using the DR wizard488
- Configuring replication and global clustering491
- Verifying the disaster recovery configuration497
- Establishing secure communication within the global cluster (optional)498
- Possible tasks after creating the DR environment500

Chapter 13 Testing fault readiness by running a fire drill

- About disaster recovery fire drills501
- About the Fire Drill Wizard501
- Tasks for configuring and running fire drills503
- Prerequisites for a fire drill503
- Fire Drill Wizard actions504
- Preparing the fire drill configuration505
- Running a fire drill508
- Restoring the fire drill system to a prepared state509
- Deleting the fire drill configuration510

Section 5 Appendices

Appendix A Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation

- Reviewing the requirements519
 - Disk space requirements519
 - Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)520
- Reviewing the configuration523
- Configuring the network and storage525
- Configuring SFW HA: Prior to installing Exchange526
- Installing Veritas Storage Foundation HA for Windows527
 - Setting Windows driver signing options527
 - Installing Storage Foundation HA for Windows528
 - Configuring VxSAS533
 - Resetting the driver signing options535
 - Configuring the cluster535
- Configuring disk groups and volumes for disaster recovery551
 - Creating a disk group553
 - Creating volumes554

Managing disk groups and volumes	559
Importing a disk group and mounting a volume	559
Unmounting a volume and deporting a disk group	560
Preparing the forest and domain (Primary site)	560
Installing Exchange on the first node (Primary site)	561
Exchange pre-installation: First node	562
Exchange installation: First node	564
Exchange post-installation: First node	565
Moving Exchange databases (Primary site)	565
Installing Exchange on additional nodes (Primary site)	568
Exchange pre-installation: Additional nodes	569
Exchange installation: Additional nodes	570
Exchange post-installation: Additional nodes	572
Configuring the Exchange service group for VCS (Primary site)	574
Prerequisites	574
Setting up the secondary Site: Configuring SFW HA prior to installing Exchange	580
Installing Exchange on the first node and Additional nodes (Secondary site)	580
Installing Exchange on the first node with DR Option (Secondary site)	581
Exchange pre-installation on first node (Secondary site)	582
Exchange installation on first node (Secondary site)	585
Exchange post-installation on first node (Secondary site)	586
Installing Exchange on additional nodes (secondary site)	587
Exchange pre-installation: Additional nodes	588
Exchange installation: Additional nodes	589
Exchange post-installation: Additional nodes	590
Configuring SFW HA: After installing Exchange on secondary site	592
Copying the .CRK file to the primary site	592
Backing up and restoring the Exchange disk group	593
Configuring the Exchange service group for VCS (secondary site)	593
Prerequisites	593
Verifying the cluster configuration	599
Configuring DR components on primary and Secondary sites	600
Possible task after creating the DR Environment: Adding a new failover node	600
Preparing the new node	600
Preparing the existing DR environment	601
Installing Exchange on the new node	601
Modifying the replication and Exchange service groups	602
Reversing replication direction	602

Appendix B	Configuring the DR components (VVR and GCO)	
	Reviewing the prerequisites	607
	Setting up the replicated data sets (RDS) for VVR	607
	Creating the VVR RVG service group	616
	Configuring the global cluster option for wide-area failover	621
	Prerequisites	621
	Linking clusters: Adding a remote cluster to a local cluster	621
	Converting a local Exchange service group to a global service group	623
	Bringing a global service group online	625
	Administering global service groups	625
	Deleting a remote cluster	627
	Establishing secure communication within the global cluster (optional)	631
Index		635

Introduction

This section introduces Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange and contains information on using the Solutions Configuration Center. This section contains the following chapters:

- [Chapter 1, “Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange”](#) on page 19
- [Chapter 2, “Using the Solutions Configuration Center”](#) on page 21

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange

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 Exchange* 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 Exchange*.

Separate guides are available for Microsoft SQL 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 Microsoft Exchange includes procedures for installing and configuring clustered Microsoft Exchange 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 Microsoft Exchange includes procedures for installing, configuring, and testing clustered and replicated Microsoft Exchange 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 Exchange* 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

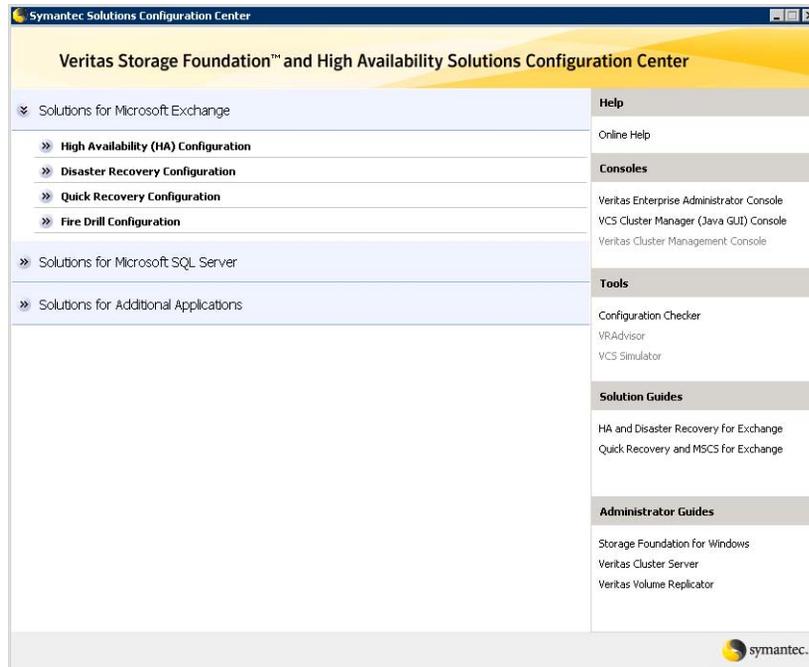


Figure 2-2 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 2-2 Solutions Configuration Center for Microsoft SQL Server

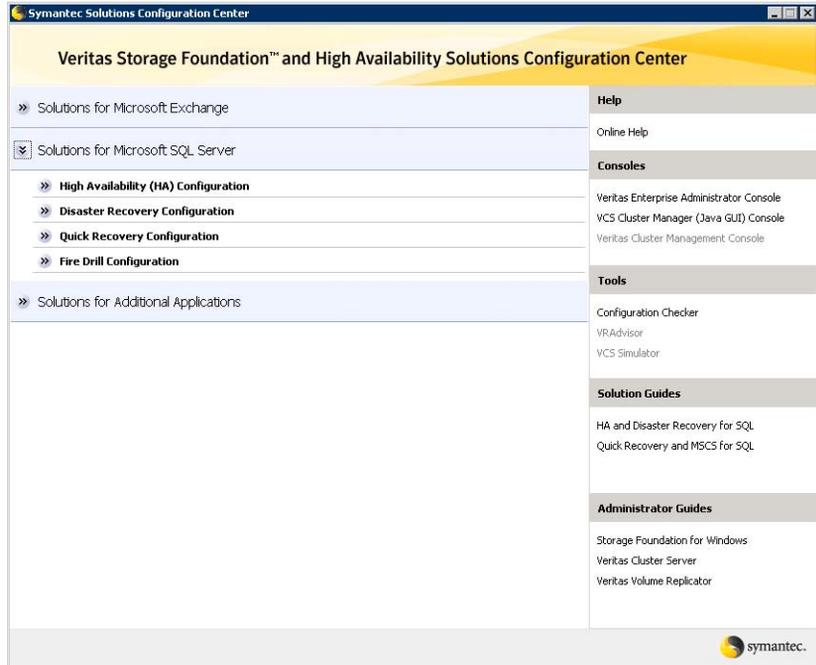
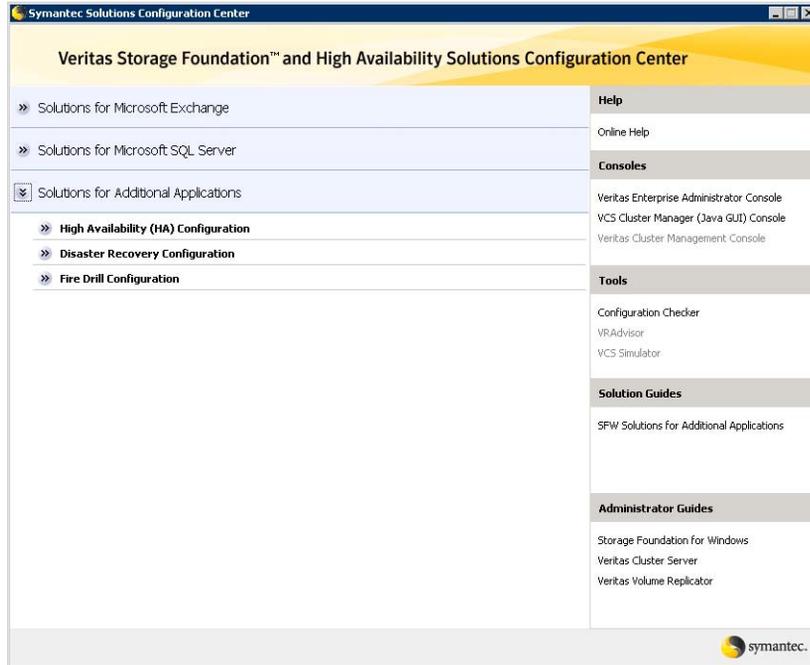


Figure 2-3 shows the choices available when you click Solutions for Additional Applications.

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

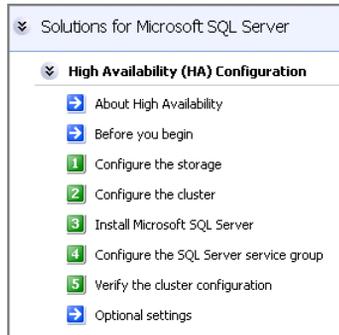
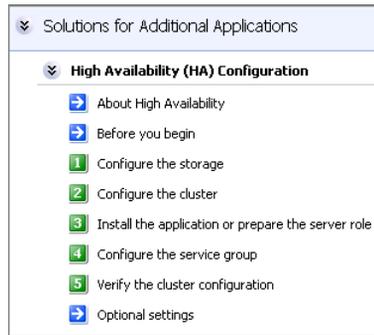


Figure 2-9 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

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



High Availability

Local clustering provides high availability (HA) through database and application failover. Use local clusters to recover data in the event of application, operating system, or hardware failure, and to minimize planned and unplanned downtime.

Refer to the following chapters to install and configure a clustered Exchange environment using Veritas Storage Foundation HA for Windows:

- [Chapter 3, “High availability for Exchange: Overview” on page 31](#)
- [Chapter 4, “Deploying SFW HA for high availability: New installation” on page 33](#)
- [Chapter 6, “Deploying SFW HA for high availability: Configuring a new any-to-any failover” on page 177](#)

High availability for Exchange: Overview

This chapter covers the following topics:

- [What is high availability?](#)
- [Why implement a high availability solution?](#)
- [How the VCS application agent makes Microsoft Exchange highly available](#)
- [Typical HA configurations for Exchange](#)

What is high availability?

High Availability (HA) is a state where data and applications are highly available because software or hardware maintain the continued functioning in the event of computer failure. HA can refer to any software or hardware that provides fault tolerance, but generally the term is associated with clustering. This section focuses on configurations that use Veritas Storage Foundation HA for Windows (SFW HA).

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, and shares a common namespace. It is 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 goal for critical applications. Clustered systems have several advantages, including fault tolerance, high availability, scalability, simplified management, and support for rolling upgrades.

Using VCS as a local high availability solution prepares 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; accommodates flexible failover policies, any-to-any configurations, and shared standby servers for Exchange

How the VCS application agent makes Microsoft Exchange highly available

If a configured Exchange service is not running or if a configured virtual server is not available, the VCS application agent for Microsoft Exchange Server detects an application failure. When this occurs, the Exchange service group is failed over to the next available system in the service group's system list. The configured Exchange services and virtual servers are started on the new system. This ensures continuous availability for Exchange data and configured mailboxes.

Typical HA configurations for Exchange

Typical HA configurations for Exchange are as follows:

- Active/passive failover configuration
- Any-to-any failover configuration

Deploying SFW HA for high availability: New installation

This chapter covers the following topics:

- [Tasks for a new HA installation of Microsoft Exchange](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Preparing the forest and domain](#)
- [Configuring SFW HA: Prior to installing Exchange](#)
- [Configuring disk groups and volumes](#)
- [Configuring the cluster](#)
- [Managing disk groups and volumes](#)
- [Installing Exchange on the first node](#)
- [Moving Exchange databases to shared storage](#)
- [Installing Exchange on additional nodes](#)
- [Configuring the Exchange service group for VCS](#)
- [Verifying the cluster configuration](#)
- [Configuring the Cluster Management Console connection](#)

Tasks for a new HA installation of Microsoft Exchange

This chapter provides information on how to install and configure a new Veritas Storage Foundation High Availability or Disaster Recovery environment for Exchange. This environment involves an active/passive configuration with one-to-one failover capabilities.

See [“Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover”](#) on page 259.

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.

After completing the high availability installation of Exchange, see [“Deploying Disaster Recovery using the DR wizard for cloning: New Exchange Server installation”](#) on page 393.

[Table 4-1](#) outlines the high-level objectives and the tasks to complete each objective:

Table 4-1 Task List

Objective	Tasks
“Reviewing the requirements” on page 36	<ul style="list-style-type: none">■ Verify hardware and software prerequisites
“Reviewing the configuration” on page 40	<ul style="list-style-type: none">■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster
“Configuring the storage hardware and network” on page 43	<ul style="list-style-type: none">■ Set up the network and storage for a cluster environment■ Verify the DNS entries for the systems on which Exchange will be installed

Table 4-1 Task List

Objective	Tasks
“Configuring SFW HA: Prior to installing Exchange” on page 45	<ul style="list-style-type: none"> ■ Verify the driver signing options for Windows 2003 systems ■ Install SFW, VCS, and the Veritas Cluster Server Application Agent for Microsoft Exchange ■ Restore driver signing options for Windows 2003 systems
“Configuring disk groups and volumes” on page 52	<ul style="list-style-type: none"> ■ Use the VEA console to create disk groups ■ Use the VEA console to create the data, log, RegRep, and Shared volumes ■ Manage disk groups and volumes, with instructions for mounting and unmounting volumes
“Configuring the cluster” on page 61	<ul style="list-style-type: none"> ■ Verify static IP addresses and name resolution configured for each node ■ Configure cluster components using the Veritas Cluster Server Configuration Wizard
“Preparing the forest and domain” on page 45	<ul style="list-style-type: none"> ■ Set up the forest and domain prior to the Exchange installation
“Installing Exchange on the first node” on page 76	<ul style="list-style-type: none"> ■ Review the prerequisite checklist ■ Run the Exchange Setup Wizard for Veritas Cluster Server and Microsoft Exchange Server installation
“Moving Exchange databases to shared storage” on page 81	<ul style="list-style-type: none"> ■ Move databases on the first node from the local drive to the shared drive using the Exchange Setup Wizard for Veritas Cluster Server

Table 4-1 Task List

Objective	Tasks
“Installing Exchange on additional nodes” on page 84	<ul style="list-style-type: none"> Run the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes
“Configuring the Exchange service group for VCS” on page 89	<ul style="list-style-type: none"> Create the Exchange service group using the VCS Exchange Configuration Wizard.
“Verifying the cluster configuration” on page 96	<ul style="list-style-type: none"> Verify the cluster configuration by switching service groups and shutting down an active cluster node

Reviewing the requirements

Before installation, review these product installation requirements for your systems. 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 4-2](#) estimates disk space requirements for SFW HA.

Table 4-2 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 Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:
 - Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition (SP2 required)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
or
Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP1 required for all editions)
or
Windows Server 2003 R2 (Standard Edition, Enterprise Edition, or Datacenter Edition)
or
 - Microsoft Exchange 2000 Server or Microsoft Exchange 2000 Enterprise Server (SP3 with August 2004 rollup patch required for both editions)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)

Note: Microsoft support for Exchange Server 2003 is limited to 32-bit versions of the Windows 2003 operating system.

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 40.
- 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 Exchange Virtual Server (EVS)
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

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

Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

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

- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Domain Servers group.
- 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.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- You must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.
- 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.

- You must install IIS, SMTP, NNTP, ASP.NET, and WWW services before you install Microsoft Exchange Server.

Best practices

Symantec recommends that you perform the following tasks:

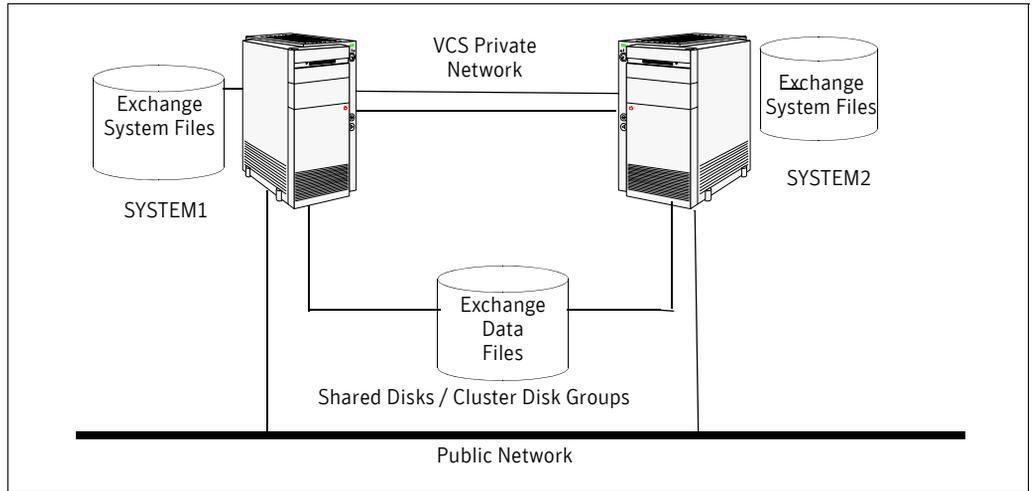
- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the 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

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 if the active node fails.

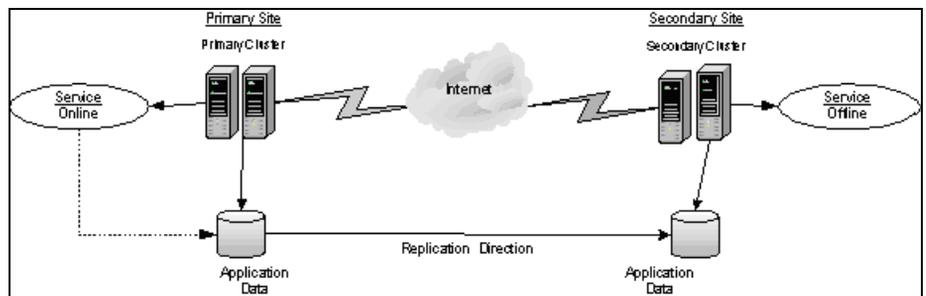
In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. In this case, EVS1 can fail over from SYSTEM1 or SYSTEM2. [Figure 4-1](#) illustrates an active/passive failover configuration with an Exchange virtual server.

Figure 4-1 Active/Passive failover configuration



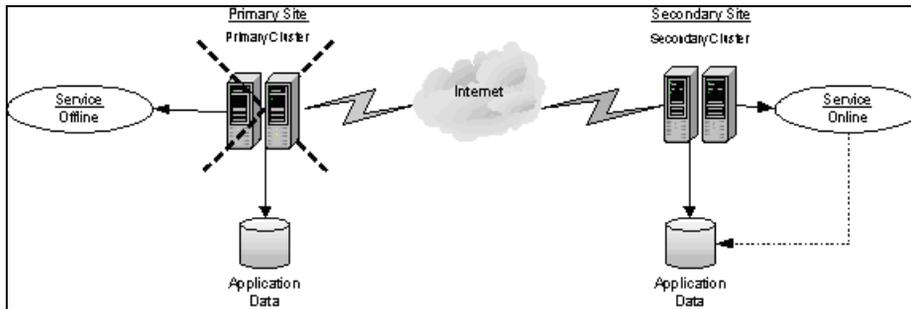
In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. [Figure 4-2](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 4-2 Disaster Recovery environment



When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients. [Figure 4-3](#) illustrates this type of failure:

Figure 4-3 Failure in a disaster recovery environment



IP addresses required during configuration

You should have the following IP addresses available before you start the configuration process:

Exchange virtual server	The virtual IP address for the Exchange server. For a disaster recovery configuration, the virtual IP address for the Exchange server at the primary and disaster recovery site can only 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	Used by Veritas Cluster Management Console (Single Cluster Mode), also referred to as Web Console. Used by VCS notifier. For a disaster recovery configuration, used by the Global Cluster Option. For a disaster recovery configuration, a separate IP address is required for the secondary site.
Replication IP address (disaster recovery configuration only)	For a disaster recovery configuration, an IP address is required for each Replicated Data Set (RDS) one for the primary site and one for the secondary site. Two IP addresses are required per Replicated Volume Group (RVP).

Sample configuration

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

Table 4-3 Sample configuration

Name	Object
SYSTEM1, SYSTEM2	Physical node names
EVS1	Microsoft Exchange Virtual Server
EVS1_SG1	Microsoft Exchange service group
EVS1_SG1_DG	Cluster disk group name
EVS1_SG1_DB1	Volume for storing the Microsoft Exchange Server database
EVS1_SG1_LOG	Volume for storing a Microsoft Exchange Server database log file
EVS1_SG1_REGREP	Volume that contain the list of registry keys that must be replicated among cluster systems for the Exchange server
EVS1_SG1_SHARED	Volume for storing Microsoft Exchange Server MTA database, SMTP, and message tracking

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**.

Preparing the forest and domain

Microsoft requires the preparation of the forest and domain prior to an Exchange installation. See Microsoft documentation for instructions for preparing the forest and domain for an Exchange installation. Do not repeat this process for additional Exchange installations.

Configuring SFW HA: Prior to installing Exchange

Before installing Exchange on the primary site, complete the following procedures:

- Install the SFW HA software.
See [“Installing Veritas Storage Foundation HA for Windows”](#) on page 45.
- Set up a VCS environment.
See [“Configuring the cluster”](#) on page 61
- Create the required disk groups and volumes.
See [“Configuring disk groups and volumes”](#) on page 52
See [“Managing disk groups and volumes”](#) on page 60.

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 Application Agent for Exchange. 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

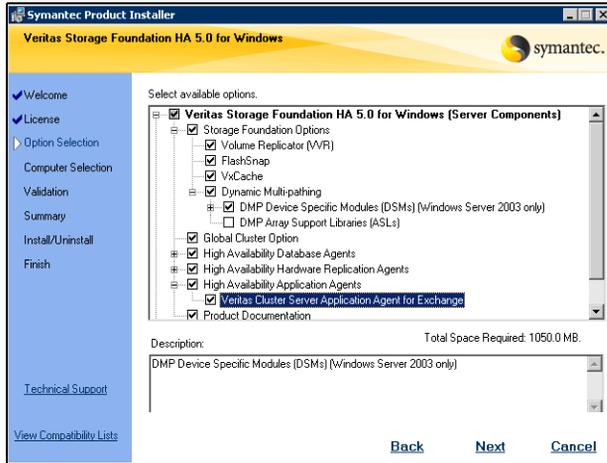
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.

- 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.

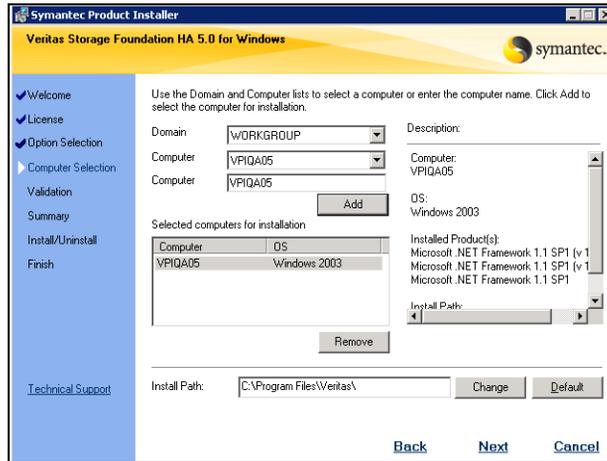
9 Select the appropriate SFW product options for your installation. Click **Next**.



The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

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

10 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

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

Computer

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

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

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

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

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
 - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
 - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

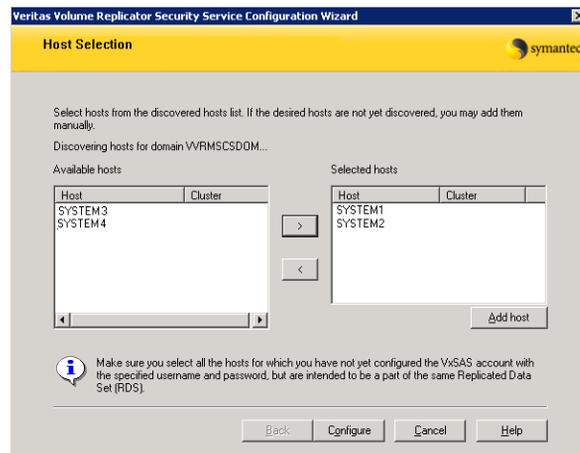
Configuring VxSAS

In DR installations the wizard for Veritas Volume Replicator Security Service (VxSAS) launches automatically after reboot. Proceed with the wizard to configure this service.

See the *Veritas Storage Foundation Volume Replicator, Administrator's Guide*.

To configure the VxSAS service

- 1 After reboot, review the Welcome message, and click **Next**.
- 2 Enter account information:
 - Specify the administrative account name in the Account name field.
 - Specify a password in the Password field.
 - Click **Next**.
Specify the same user name and password on all nodes that will be part of the VVR configuration.
- 3 Select the domain where the systems that will be part of the RDS reside and click **Next**.
- 4 Select the hosts that will be part of the RDS:



- Select the required hosts from the computers selection pane. If you see all required hosts, Click **Configure**.
- If you do not see the required host:
 - Click **Add Host**.
 - Specify the name or the IP of the new host.

- Click **Add**.
 - Click **Configure**.
- 5 Review the configuration results summary for the selected hosts. Click **Back** to select or add hosts as required.
 - 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 disk groups and volumes

Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to grow existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

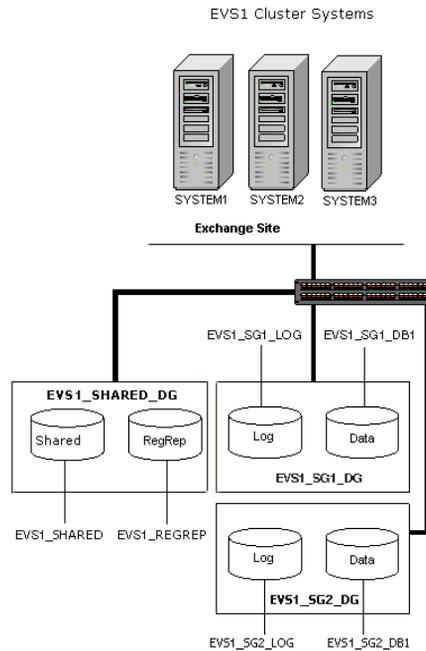
Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs.

Typically, a SFW disk group corresponds to an Exchange storage group.

[Figure 4-4](#) displays a detailed view of the disk groups and volumes in an HA environment.

Figure 4-4 Disk groups and volumes for Exchange virtual server EVS1 in HA setup



Use the following procedures to create the appropriate disk groups and volumes. The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange disk group EVS1_SG1_DG contains two volumes:

- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume. This contains the EVS1_SG1_LOG volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

Exchange storage group EVS1_SHARED_DG create contains two volumes:

- EVS1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.
- EVS1_SHARED: Contains the MTA database, SMTP, and message tracking.

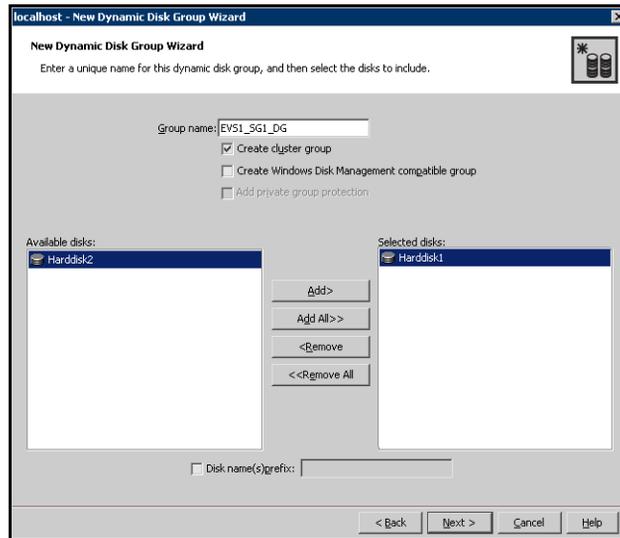
Note: For additional Exchange storage groups, place the disks associated with the additional storage group's volumes in their own disk group.

Creating a disk group

To create a dynamic (cluster) disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_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

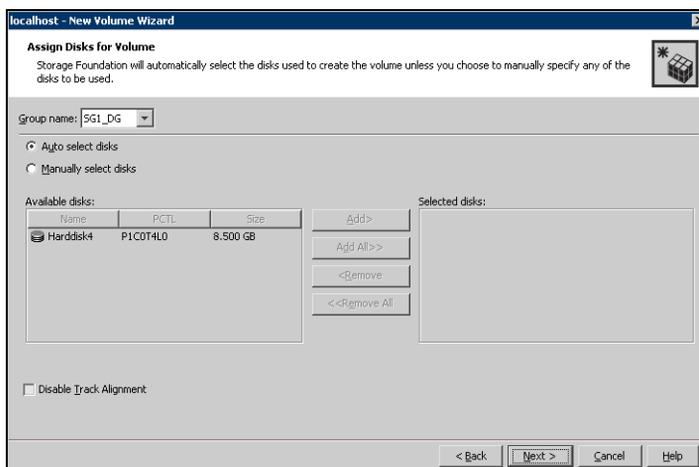
This procedure assumes you are starting with the EVS1_SG1_DB1 volume. Refer to the steps below for the Data, Log, RegRep, and Shared volumes.

For Disaster Recovery configuration only, the Disaster Recovery Configuration wizard can create the Storage Replicator Log volumes for you.

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

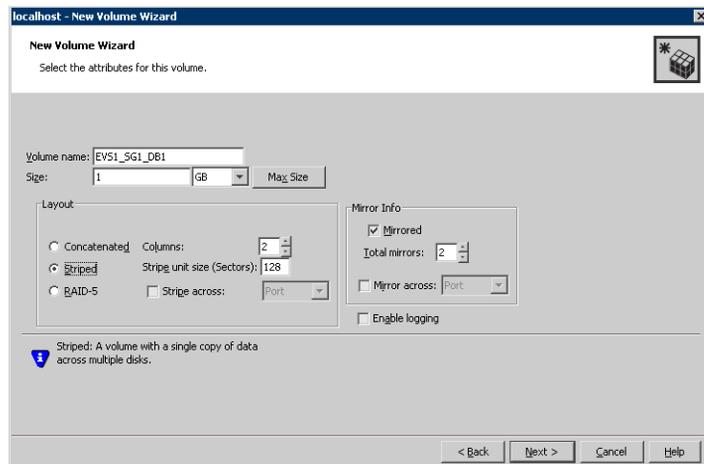
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu 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 EVS1_SG1_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**.
- Specify the volume attributes.

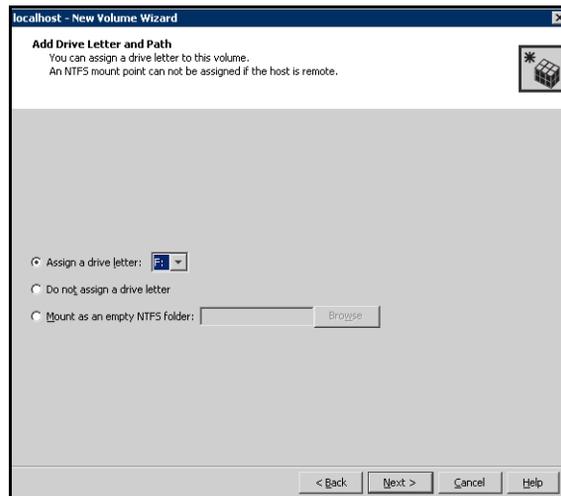


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 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.

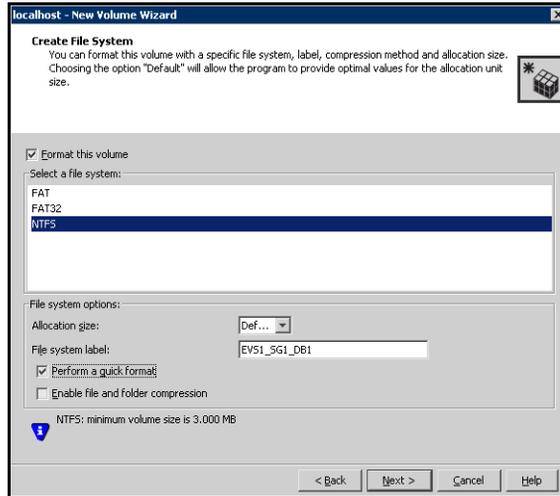
Do not assign the M: drive letter if you are using Exchange 2000, as it is reserved for internal use.

- 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 the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) Also, repeat these steps to create both the RegRep volume (EVS1_REGREP) and the EVS1_SHARED volumes in the EVS1_SHARED_DG disk group.

15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3). Create the cluster disk group and volumes on the first node of the cluster only.

If you are configuring an any-to-any environment, you can also create similar disk groups and volumes for the other Exchange servers. For

example, create disk group (EVS2_SG1_DG) and volumes (EVS2_SG1_DB1, EVS2_REGREP, EVS2_SG1_LOG, and EVS2_SHARED).

Managing disk groups and volumes

This section includes the following procedures:

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

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*

Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.

- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Configuring the cluster

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

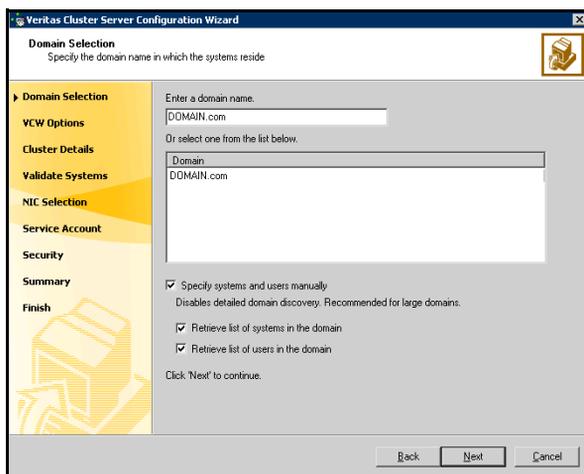
Complete the following tasks before creating a cluster:

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

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

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**.

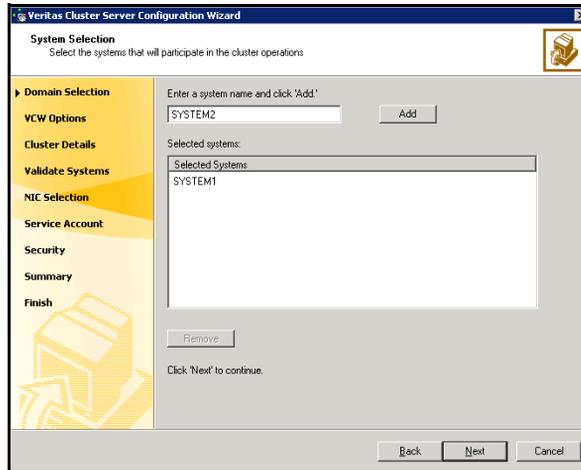
Proceed to [step 7](#) on page 64.

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**.

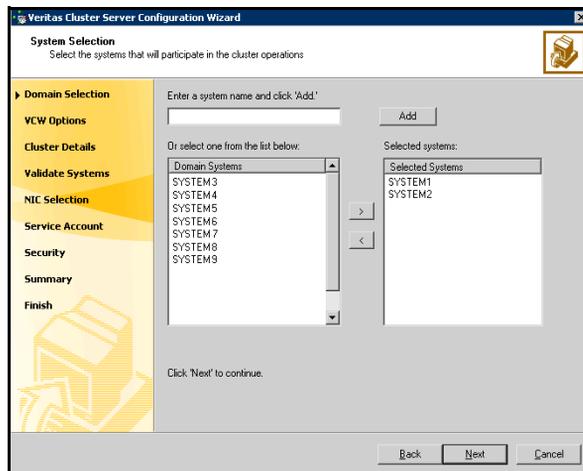
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 63. Otherwise, proceed to the next step.

- 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.



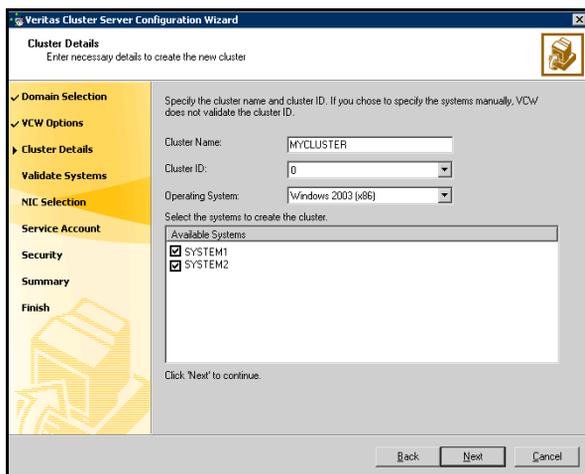
Proceed to [step 7](#) on page 64.

- 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.

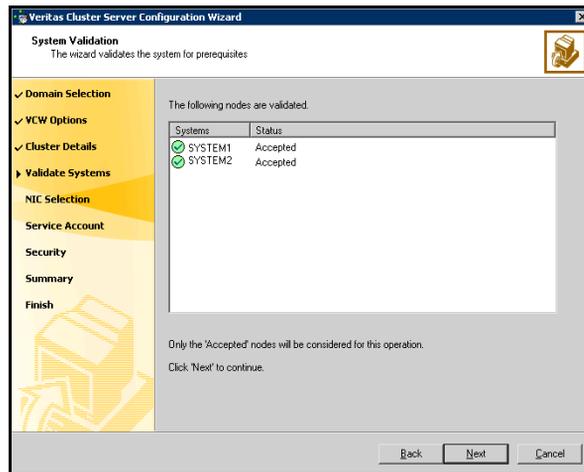
Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

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.

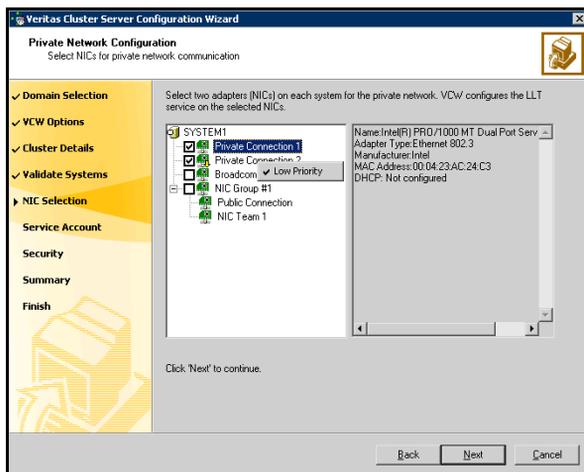
- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



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 64, proceed to the next step. Otherwise, proceed to [step 11](#) on page 66.

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

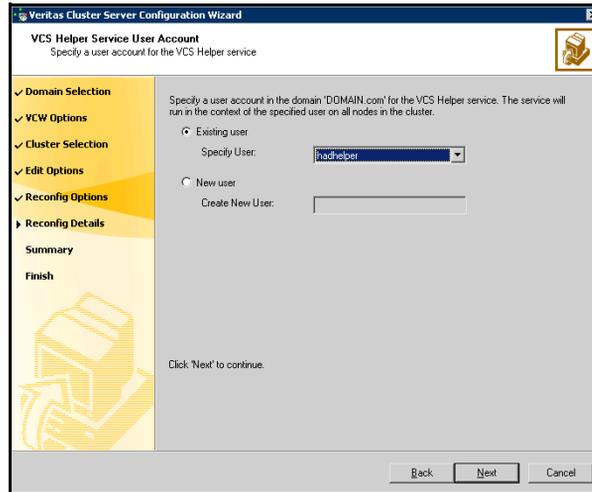


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs 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 62, 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**.

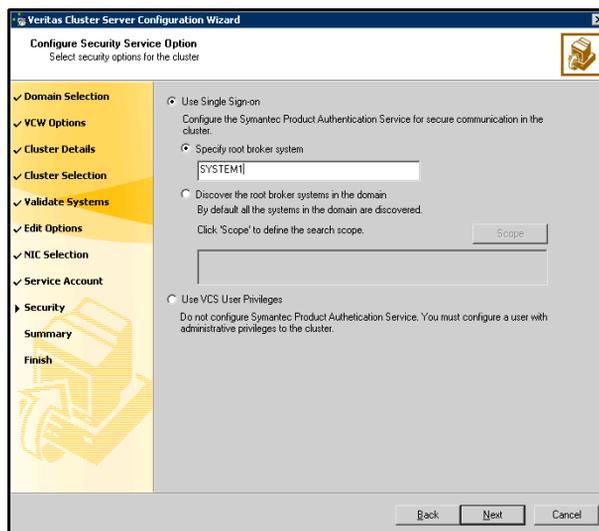
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- 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.

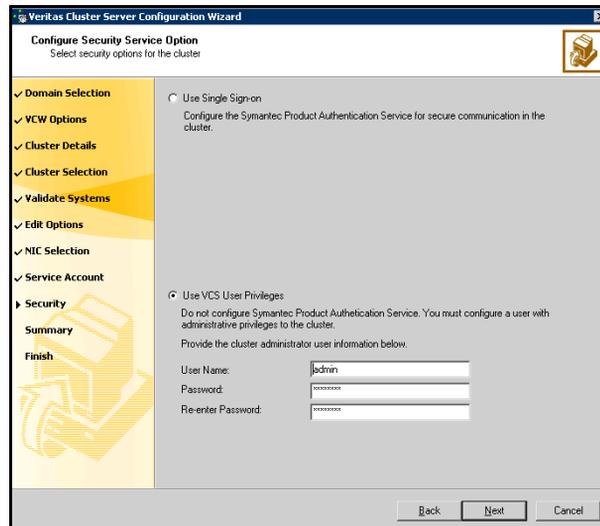
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- 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 'Configure Security Service Option' step of the Veritas Cluster Server Configuration Wizard. The window title is 'Veritas Cluster Server Configuration Wizard' and the subtitle is 'Configure Security Service Option'. The main instruction is 'Select security options for the cluster'. On the left, a navigation pane shows steps: Domain Selection, VCS Options, Cluster Details, Cluster Selection, Validate Systems, Edit Options, NIC Selection, Service Account, Security, Summary, and Finish. The 'Security' step is selected. The main area has two radio button options: 'Use Single Sign-on' (unselected) and 'Use VCS User Privileges' (selected). The 'Use VCS User Privileges' option includes the text: '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.' Below this text are three input fields: 'User Name:' with 'admin' entered, 'Password:' with 'password' entered, and 'Re-enter Password:' with 'password' entered. At the bottom right, there are 'Back', 'Next', and 'Cancel' buttons.

- 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.

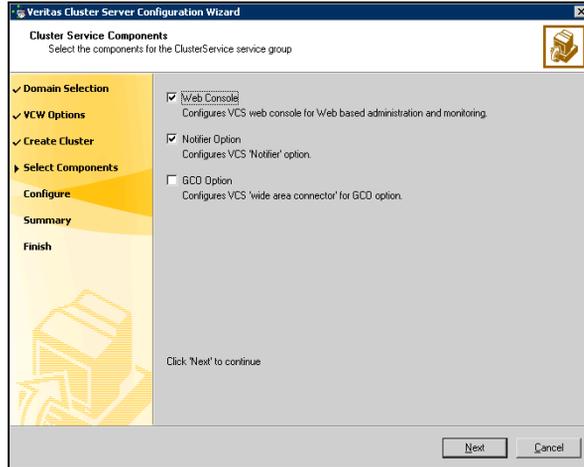
■ Click **Next**.

- 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**.



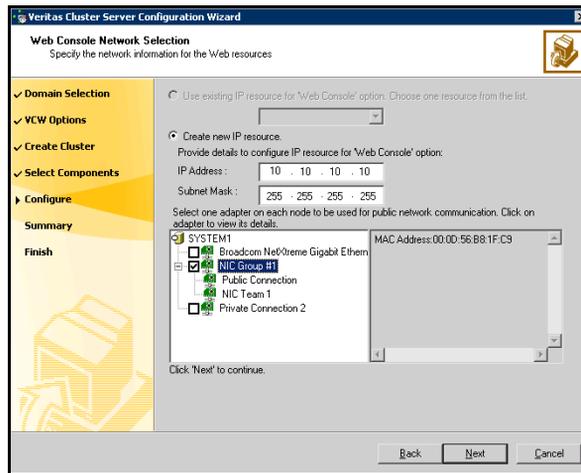
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 72.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 73.

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**.



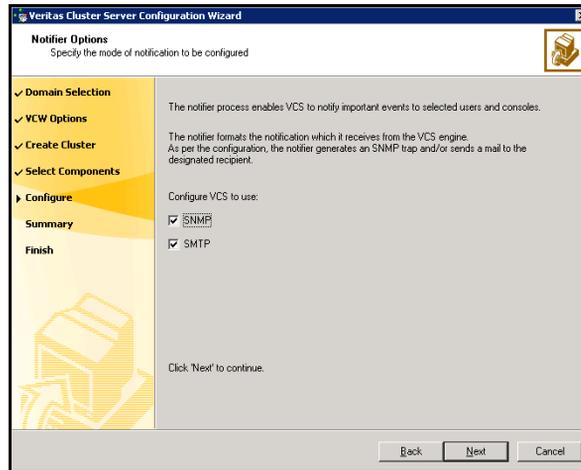
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 73. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

To configure notification

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



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

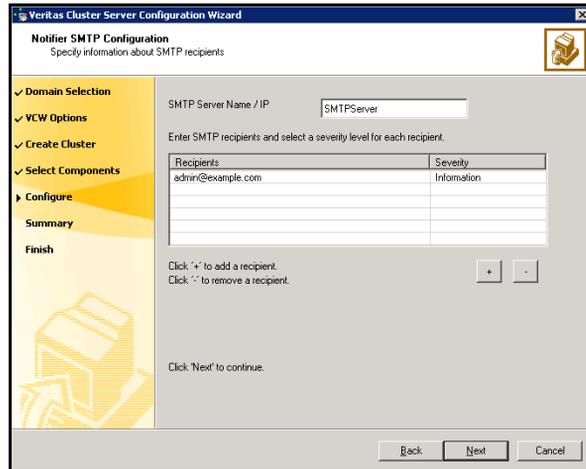
- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

The screenshot shows the 'Notifier SNMP Configuration' window in the Veritas Cluster Server Configuration Wizard. The window title is 'Veritas Cluster Server Configuration Wizard' and the subtitle is 'Notifier SNMP Configuration'. Below the subtitle is the instruction 'Specify information about SNMP console'. On the left side, there is a navigation pane with the following options: 'Domain Selection' (checked), 'VCW Options' (checked), 'Create Cluster' (checked), 'Select Components' (checked), 'Configure' (expanded), 'Summary', and 'Finish'. The main area contains a table for configuring SNMP consoles. The table has two columns: 'SNMP Console' and 'Severity'. The first row contains 'snmpserv' and 'Information'. The second row contains 'snmpserv1' and 'SevereError'. Below the table are two buttons: '+' and '-'. Below the buttons is a text input field for 'Enter SNMP Trap Port' with the value '162'. Below the input field is a note: 'Note: SNMP console must be MIB 2.0 compliant'. Below the note is the instruction 'Click "Next" to continue.'. At the bottom of the window are three buttons: 'Back', 'Next', and 'Cancel'.

SNMP Console	Severity
snmpserv	Information
snmpserv1	SevereError

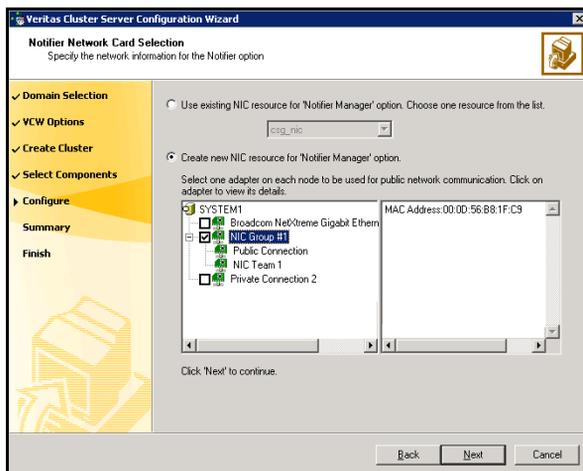
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.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

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



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

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



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

Installing Exchange on the first node

Installing Exchange on the first node is described in three stages that involve preinstallation, installation, and post-installation procedures. Complete the following tasks before installing Exchange Server:

- Prepare the forest and domain.
See “[Preparing the forest and domain](#)” on page 45.
- Verify the disk group is imported on the first node of the cluster.
See “[Importing a disk group and mounting a shared volume](#)” on page 60.

- Mount the volume containing the information for registry replication (EVS1_REGREP).
See “[Importing a disk group and mounting a shared volume](#)” on page 60.
- Verify that all systems on which Exchange Server will be installed have IIS installed. You must install SMTP, NNTP, and WWW services on all systems. If you install Exchange on Windows 2003, make sure to install ASP.NET service.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. The VCS requires the Exchange installation must place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.
 - Make sure the HAD Helper domain user account has the “Add workstations to domain” privilege enabled in the Active Directory. To verify this, click Start > Administrative Tools > Local Security Policy on the domain controller to launch the security policy display. Click Local Policies > User Rights Management and make sure the user account has this privilege.
 - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
 - The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

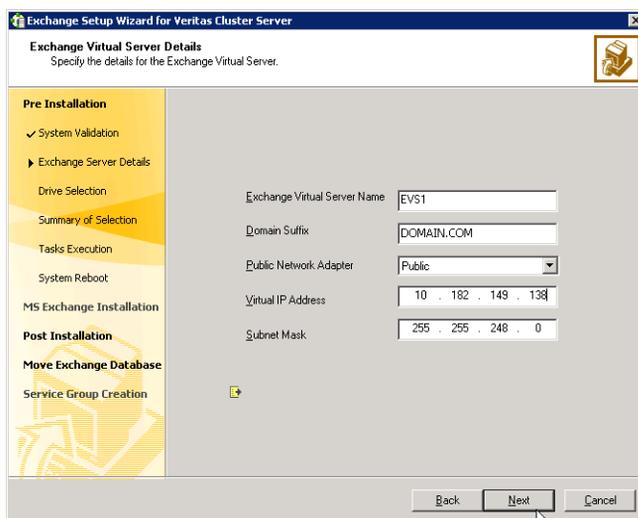
Exchange pre-installation: First node

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. You must install Exchange on a virtual node to facilitate high availability. After you have run the wizard, you will be requested to restart the

node. So, close all open applications and save your data before running the wizard.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a new Exchange Virtual Server** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.
- 7 Specify information related to the network.



- Enter a unique virtual name for the Exchange server. Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.

- Enter a domain suffix for the virtual server.
- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
- Enter a unique virtual IP address for the Exchange virtual server.
- Enter the subnet mask for the virtual IP address.
- Click **Next**.

The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met. Click **Next**.

- 8 Select a drive where the registry replication data will be stored and click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: First node

Install Exchange on the node selected in, “[Exchange pre-installation: First node](#)” on page 77.

Exchange 2000 requires service pack 3 with the August 2004 rollup patch. Exchange 2003 requires service pack 2. The procedure below is based on Exchange 2003 SP2.

This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Install Exchange Server using the Microsoft Exchange installation program. See the Microsoft Exchange documentation for instructions.
- 2 Reboot the node if prompted to do so.
- 3 For Exchange 2000 or Exchange 2003, install the required service pack.

Exchange post-installation: First node

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.

- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

Moving Exchange databases to shared storage

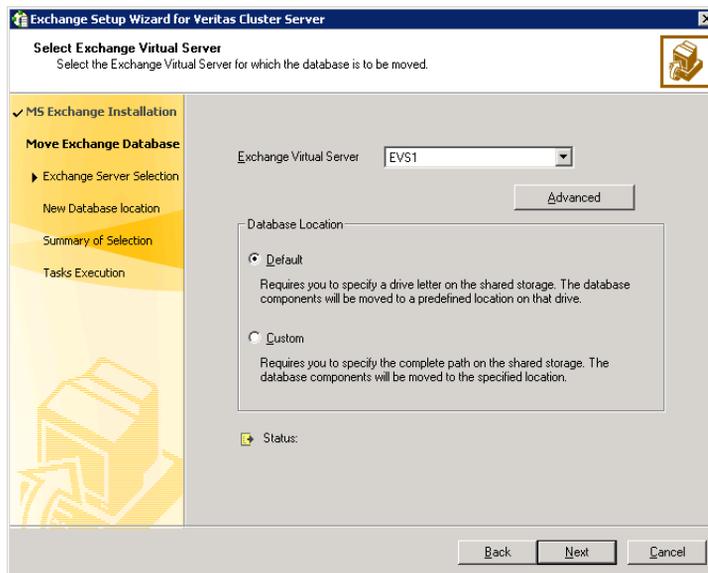
After completing the Exchange installation on the first node, move the Exchange databases on the first node from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

- Make sure the data queue is empty on the SMTP server.
- Make sure to import the disk group and mount the volumes for the Exchange database, MTA data, and transaction logs. See “[Managing disk groups and volumes](#)” on page 60.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.

5 In the Select Exchange Virtual Server dialog box:



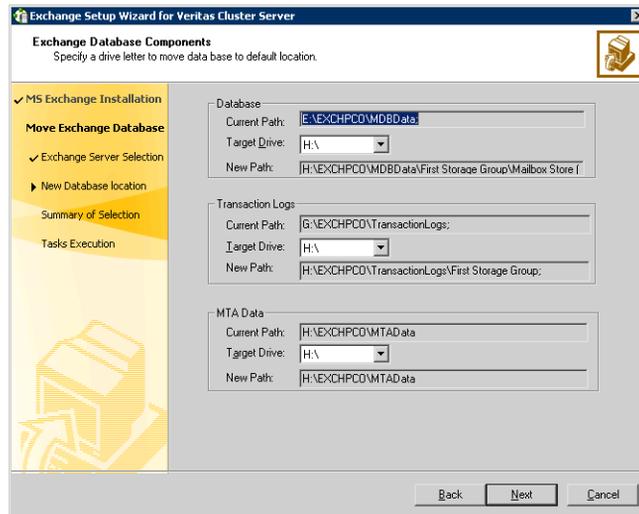
- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

Warning: The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Click **OK**.
 - Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
 - Click **Next**.
- 6 Do one of the following:
- If you would like to move the first mailbox store, public store, and MTA data to the generated default paths on the volumes for these components, proceed to the next step.
 - Otherwise, proceed to [step 8](#) on page 84 to specify the path location on the volumes that you will designate for these components.

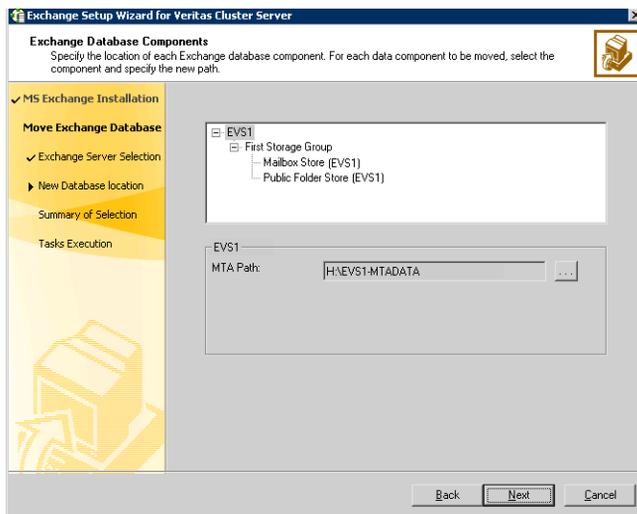
Warning: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

- 7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components will be moved to a default location on that drive. In the Exchange Database Components dialog box:



- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Specify the drive where the Exchange MTA Data will be moved.
- Click **Next** and proceed to [step 9](#) on page 84.

- 8 For the option of a custom database location, specify the location for specific Microsoft Exchange data components. In the Exchange Database Components dialog box:



For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the path for the Exchange database components contains only ANSI characters.

- 9 Review the summary of your selections and click **Next**.
- 10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.
- 11 Click **Finish** to exit the wizard.

Installing Exchange on additional nodes

After moving the Exchange databases to shared storage, install Exchange on additional nodes in the cluster for the same Exchange virtual server (EVS1). You must run pre-installation, installation, and post-installation procedures for each additional node. Make sure to review the prerequisites for permissions.

See “[Installing Exchange on the first node](#)” on page 76.

Exchange pre-installation: additional nodes

Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes.

See “[Unmounting a volume and deporting a disk group](#)” on page 60.

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server.
The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically

generated when setting up an additional failover mode for the virtual server in the same cluster.

- Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
 - 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.
 - 11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
 - 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: additional nodes

- Install Exchange on the node selected in “[Installing Exchange on additional nodes](#)” on page 84.
- Install any required service packs.
- Install the same Exchange version and components on all nodes.

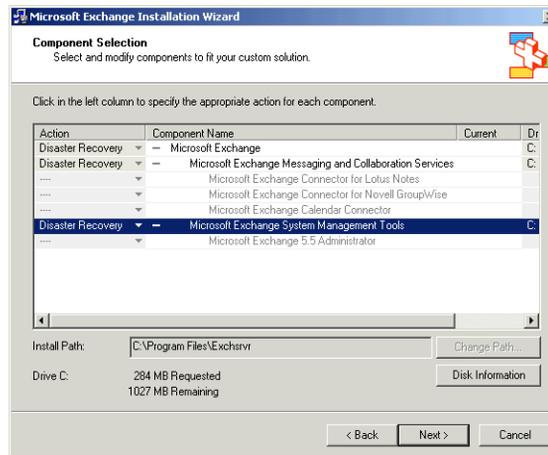
The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe /disasterrecovery
```

 where <drive letter> is the location where the Exchange software is located.
- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

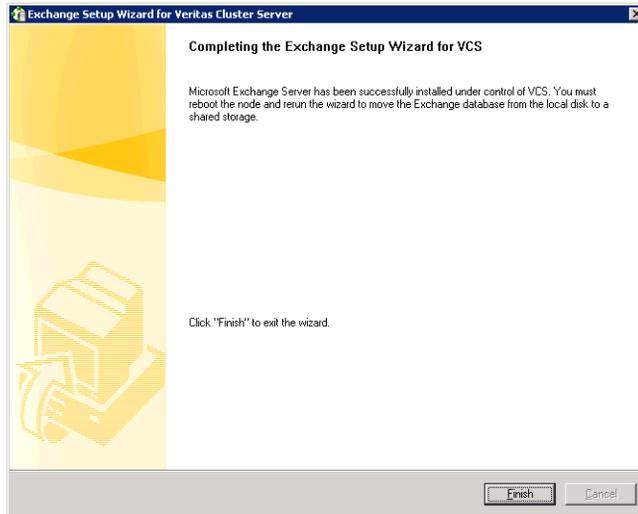
Exchange post-installation: additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.

- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

To add the nodes later, use the Exchange service group configuration wizard. See “[Configuring the Exchange service group for VCS](#)” on page 89.

Configuring the Exchange service group for VCS

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

Prerequisites

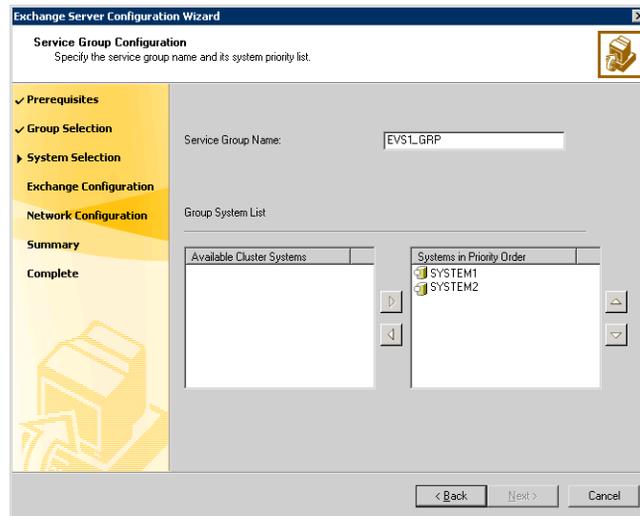
- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node on which you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk group and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - Registry changes related to Exchange
 - Transaction logs for the first storage group
 - MTA databaseSee “[Importing a disk group and mounting a shared volume](#)” on page 60 for instructions on mounting and “[Unmounting a volume and deporting a disk group](#)” on page 60 for instructions on unmounting.
- Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to an already configured service group.

To configure the Exchange service group

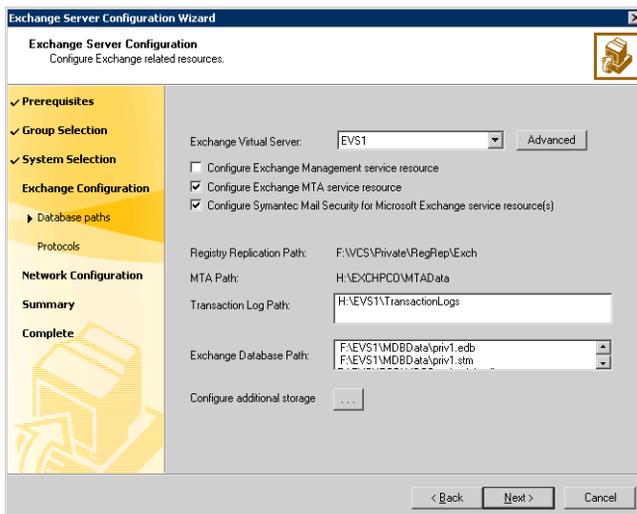
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.

- 3 In the Wizard Options dialog box, choose the **Create service group** option and click **Next**.
- 4 Specify the service group name and system list:



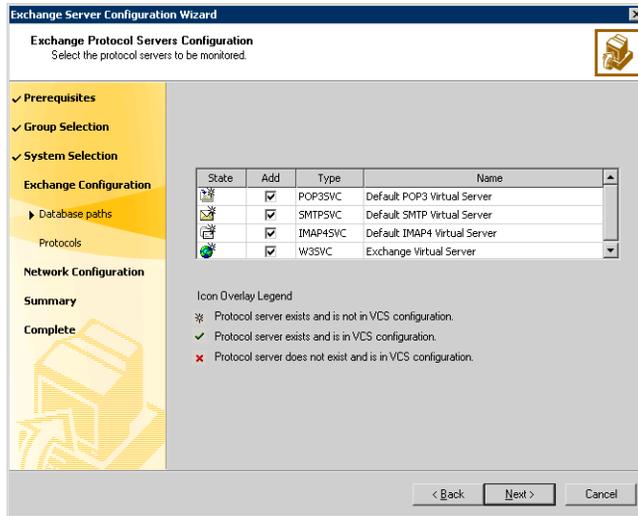
- Enter a name for the Exchange service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group’s system list. Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- To remove a system from the service group’s system list, select the system in the Systems in Priority Order list and click the left arrow.
- To change a node’s priority in the service group’s system list, select the node in the Systems in Priority Order list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:

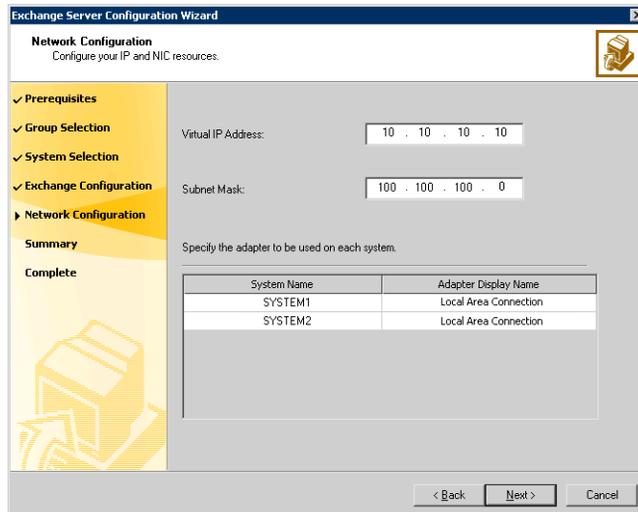


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.
- Specify whether you want to configure the Exchange Management service.
 - Specify whether you want to configure the Exchange MTA service resource. The Exchange Message Transfer Agent is specific to legacy Exchange message routing based on X.400. There are specific circumstances where it may or may not be required for your Exchange server. Please refer to Microsoft documentation on Exchange message routing and the use of MTA for further clarification and applicability to its use within your Exchange environment.

- If you are utilizing Symantec Mail Security for Exchange be sure to indicate that you would like to configure this resource.
 - Click **Next**.
- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



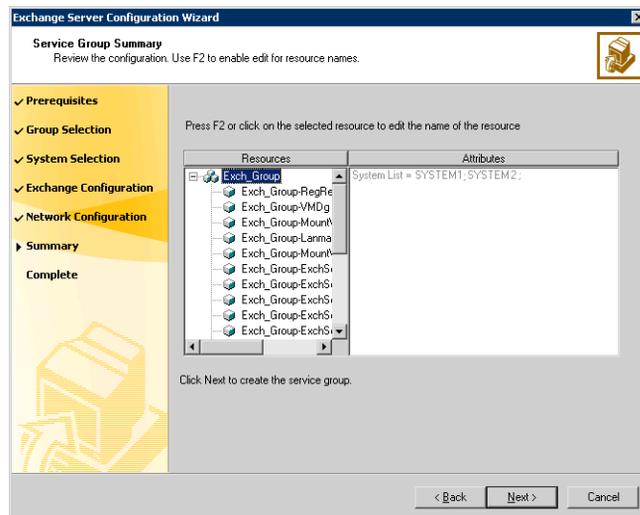
- 7 Specify information related to the network:



- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

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

- Click **Next**.
- 8 Review the service group configuration and change the resource names, if desired:



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 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.

- Click **Next**.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 9 In the Completing the Exchange Configuration dialog box, select the **Bring the service group online** check box to bring the service group online on the local system.
 - 10 Click **Finish**. After bringing the service group online, you must run the Exchange System Manager so that all the stores are automatically mounted on start-up.

To reconfigure mounting of stores at start-up

- 1 Start Exchange System Manager.
- 2 In the left pane, navigate to your storage group.
If administrative groups are not configured, Servers > Exchange Server > Storage Group.
If more than one administrative group is configured, expand Administrative Groups > Your Administrative Group > Servers > Exchange Server > Storage Group.
- 3 Right-click the Exchange database and choose **Properties** from the pop-up menu.
- 4 Click the Database tab.
- 5 Clear the **Do not mount this store at start-up** check box.
- 6 Click **OK**.

Repeat these steps for all the Exchange databases that were previously mounted.

If you need to configure additional storage groups or mailbox stores on the shared storage you should do that now. Import disk groups and mount volumes that have been created for the additional storage groups or mailbox stores, and create the new storage groups and mailbox stores in Exchange System Manager, and then rerun the Exchange Configuration Wizard to bring them under VCS control. If you already designated the additional mounted volumes and disk groups when you ran the configuration wizard the first time then you can just create the storage groups and mailbox stores in Exchange System Manager.

Verifying the cluster configuration

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

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

To switch service groups

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

To shut down an active cluster node

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

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.

Deploying SFW HA for high availability: Standalone Exchange servers

This chapter covers the following topics:

- [Reviewing the requirements](#)
- [Configuring the network and storage](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Configuring disk groups and volumes](#)
- [Managing disk groups and volumes](#)
- [Converting the standalone Exchange server into a “Clustered” Exchange server](#)
- [Adding the standalone Exchange server to a cluster](#)
- [Moving Exchange databases to shared storage](#)
- [Installing Exchange on additional nodes](#)
- [Configuring the Exchange service group for VCS](#)
- [Verifying the cluster configuration](#)

This chapter provides information on how to convert a standalone Exchange server into a “clustered” Exchange server in a new Veritas Storage Foundation HA environment. This environment involves an active/passive configuration with one to one failover capabilities.

See [Chapter 7, “Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover”](#) on page 259.

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

Table 5-1 Task List

Objective	Tasks
“ Reviewing the requirements ” on page 106	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 110	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster
“ Configuring the network and storage ” on page 113	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which Exchange will be installed
“ Installing Veritas Storage Foundation HA for Windows ” on page 115	<ul style="list-style-type: none"> ■ Checking the prerequisites ■ Verifying the driver signing options for Windows 2003 systems ■ Installing SFW, VCS, and the Veritas Cluster Server Application Agent for Microsoft Exchange ■ Restoring driver signing options for Windows 2003 systems
“ Configuring disk groups and volumes ” on page 121	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data, log, RegRep, and MTA volumes

Table 5-1 Task List

Objective	Tasks
“ Managing disk groups and volumes ” on page 128	<ul style="list-style-type: none"> ■ Managing disk group and volume operations, with instructions for mounting and unmounting volumes
“ Converting the standalone Exchange server into a “Clustered” Exchange server ” on page 130	<ul style="list-style-type: none"> ■ Converting the standalone Exchange server into a cluster node using the Exchange Setup Wizard for Veritas Cluster Server
“ Adding the standalone Exchange server to a cluster ” on page 132	<ul style="list-style-type: none"> ■ Configuring the cluster ■ For a new cluster, creating the cluster, “Creating a new cluster and adding nodes” on page 134 ■ For an existing cluster, adding the new nodes to the cluster, “Adding nodes to an existing cluster” on page 150
“ Moving Exchange databases to shared storage ” on page 160	<ul style="list-style-type: none"> ■ Moving databases on the first node from the local drive to the shared drive using the Exchange Setup Wizard for Veritas Cluster Server
“ Installing Exchange on additional nodes ” on page 163	<ul style="list-style-type: none"> ■ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes
“ Configuring the Exchange service group for VCS ” on page 169	<ul style="list-style-type: none"> ■ Creating the Exchange service group using the VCS Exchange Configuration Wizard
“ Verifying the cluster configuration ” on page 176	<ul style="list-style-type: none"> ■ Verifying the cluster configuration by switching service groups and shutting down an active cluster node

Reviewing the requirements

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

This HA solution requires a standalone Microsoft Exchange 2000 or 2003 server.

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

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 Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:

- Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition (SP2 required)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
or
Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP1 required for all editions)
or
Windows Server 2003 R2 (Standard Edition, Enterprise Edition, or Datacenter Edition)
or
- Microsoft Exchange 2000 Server or Microsoft Exchange 2000 Enterprise Server (SP3 with August 2004 rollup patch required for both editions)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)

Note: Microsoft support for Exchange Server 2003 is limited to 32-bit versions of the Windows 2003 operating system.

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 109.
- 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 Exchange Virtual Server (EVS)
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

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

See the *Veritas Cluster Server Bundled Agents Reference Guide*.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Domain Servers group.

- 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.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- You must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.
- 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.
- You must install IIS, SMTP, NNTP, ASP.NET, and WWW services before you install Microsoft Exchange Server.

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

Complete the tasks in this chapter to create an active/passive configuration for Exchange with one to one failover capabilities, starting from a single standalone Exchange server.

In Scenario I, you start with a standalone Exchange server and a new node.

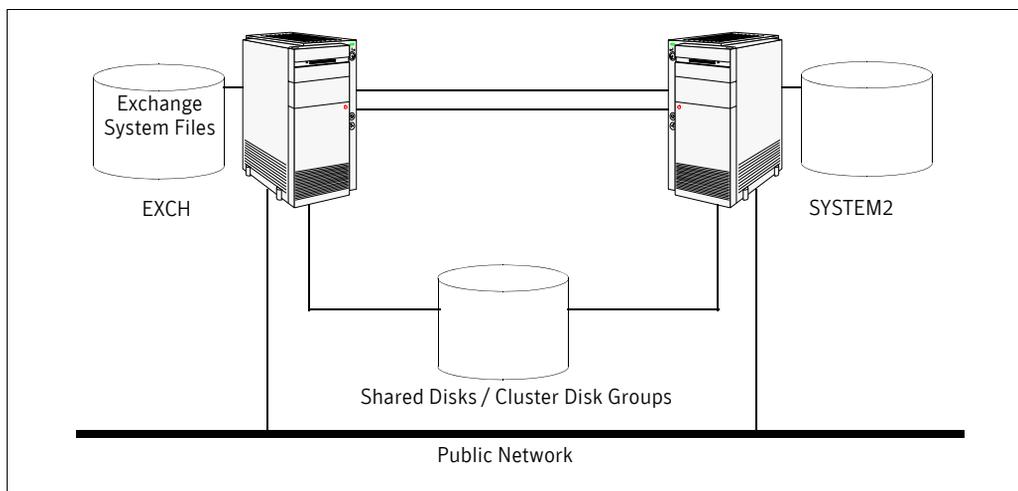
In Scenario II, you start with a standalone Exchange server and a cluster which may be running other applications.

Scenario I

In Scenario I, start with two nodes:

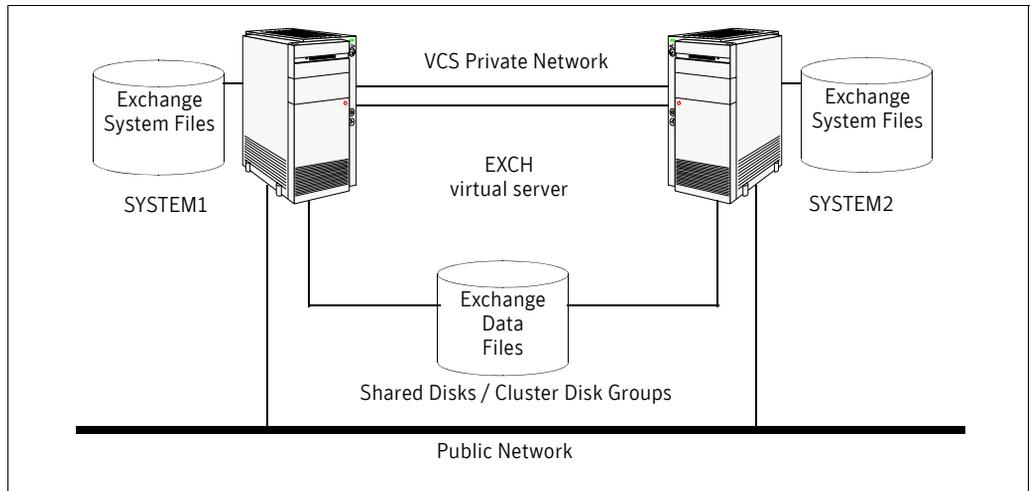
- EXCH which is a standalone Exchange server
- SYSTEM2, a new node which will join the standalone Exchange server to form a cluster

Figure 5-1 Standalone initial configuration



During the following procedures, the initial standalone Exchange server will become part of a new cluster which includes SYSTEM2, be renamed, and become an Exchange virtual server, allowing failover capabilities.

Figure 5-2 Standalone to active / passive completed configuration



In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. In this case, the Exchange virtual server can fail over from SYSTEM1 to SYSTEM2 and vice versa.

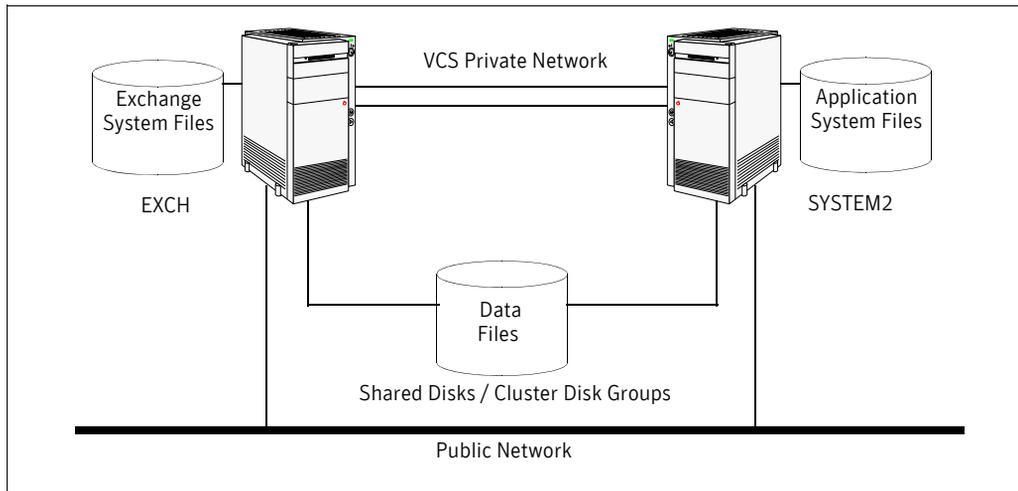
During the conversion of a standalone Exchange server into a clustered server, the existing node name of the standalone Exchange server becomes the name of the Exchange virtual server. For example, if the name of your Exchange server is EXCH, EXCH becomes the name of the Exchange virtual server.

Scenario II

In scenario II, start with a cluster:

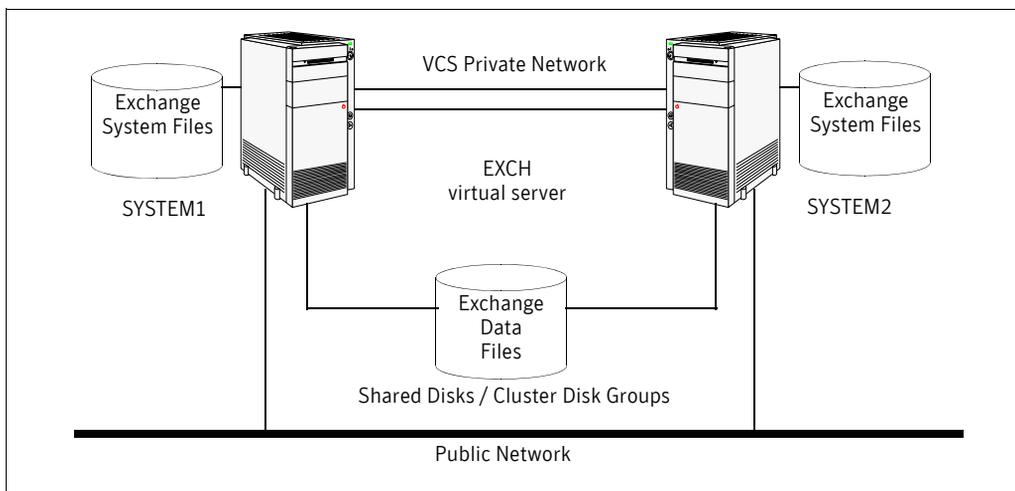
- EXCH which is a standalone Exchange server
- SYSTEM2, a node which not running as an Exchange server, but is part of a cluster

Figure 5-3 Standalone initial configuration with a cluster



During the following procedures, the initial standalone Exchange server will receive a new physical node name and the original physical node name becomes the name of the Exchange virtual server, allowing failover capabilities within the existing cluster.

Figure 5-4 Standalone to active / passive completed configuration



In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. In this case, the Exchange virtual server can fail over from SYSTEM1 to SYSTEM2 and vice versa.

Sample configuration

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

Table 5-3 Sample configuration

Name	Object
(EXCH) SYSTEM1, SYSTEM2	Physical node names; SYSTEM1 was EXCH standalone.
EVS1 (EXCH)	Microsoft Exchange Virtual Server
EVS1_GRP	Microsoft Exchange service group
EVS1_SG1_DG	Cluster disk group name
EVS1_SG1_DB1	volume for storing the Microsoft Exchange Server database
EVS1_SG1_LOG	Volume for storing a Microsoft Exchange Server database log file
EVS1_REGREP	Volume that contain the list of registry keys that must be replicated among cluster systems for the Exchange server
EVS1_SHARED	Volume for storing Microsoft Exchange Server MTA database for the Exchange server

Configuring the network and storage

Use the following procedures to configure the hardware and verify DNS settings, if this has not already been completed for all the nodes an existing cluster.

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.

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.
- 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.
- 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 that will run Exchange

- 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:
 - 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

Make sure to review the prerequisites for permissions in [“Reviewing the requirements”](#) on page 106.

When you specify the domain and the computers for the installation, specify the current physical names of your systems. Initially, the physical node names in the configuration example are EXCH (the existing standalone Exchange server), and SYSTEM2 (the new node). However, in the example below, the names used are SYSTEM1 and SYSTEM2

In the following examples, EVS1 is the name of the first Exchange virtual server. During the conversion of a standalone Exchange server into a clustered server, the existing node name of the standalone Exchange server will become the name of the Exchange virtual server. For example, if the name of your Exchange server is EXCH, then EXCH will become the name of the Exchange virtual server.

Install SFW HA on all the nodes where it is not currently installed. For a standalone Exchange server plus a new node see [“Scenario I”](#) on page 110, SFW HA must be installed on both the standalone Exchange server and the node that will serve as the failover node.

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 Application Agent for Exchange. 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 5-4 describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 5-4 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

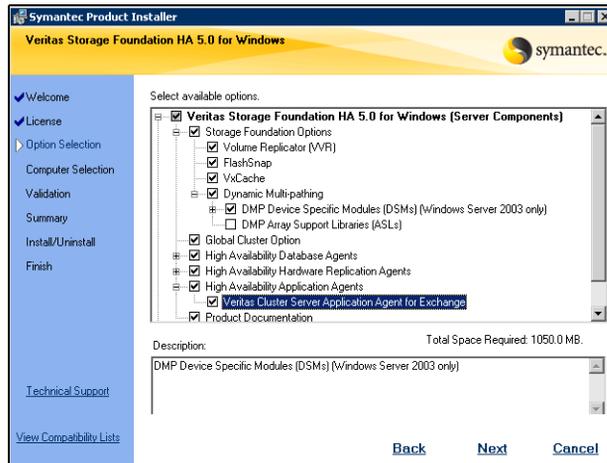
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.

- 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.

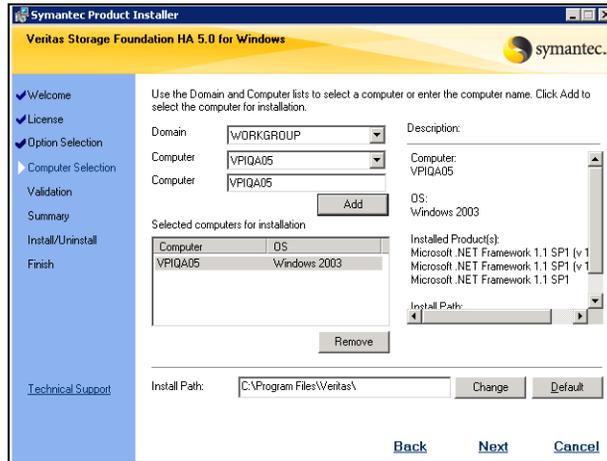
9 Select the appropriate SFW product options for your installation. Click **Next**.



The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

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

10 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

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

Computer

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

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

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

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

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
 - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
 - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring disk groups and volumes

Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to increase existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

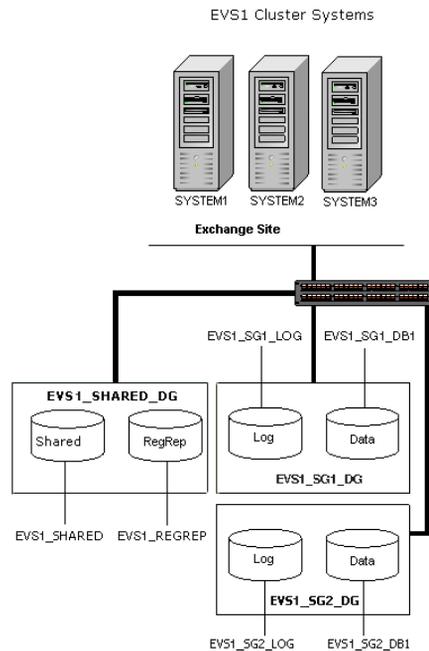
Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.

Typically, a SFW disk group corresponds to an Exchange storage group.

[Figure 5-5](#) a detailed view of the disk groups and volumes in an HA environment.

Figure 5-5 Disk groups and volumes for Exchange virtual server EVS1 in HA setup



Use the following procedures to create the appropriate disk groups and volumes. The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange storage group EVS1_SG1_DG create contains two volumes:

- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume. This will contain the EVS1_SG1_LOG volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

Exchange storage group EVS1_SHARED_DG create contains two volumes:

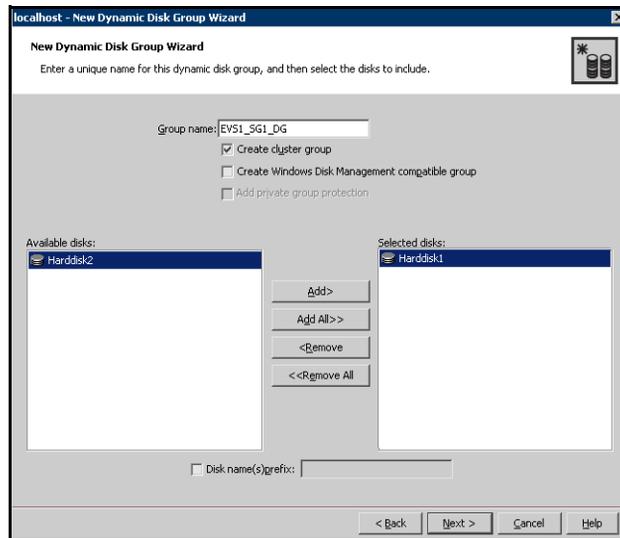
- EVS1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.
- EVS1_SHARED: Contains the MTA database, SMTP, and message tracking.

Additional storage groups (for example, EVS1_SG2_DG) only contain the data, and log volumes. The RegRep and SHARED volumes are included in the first storage group.

Creating a disk group

To create a dynamic (cluster) disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_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.

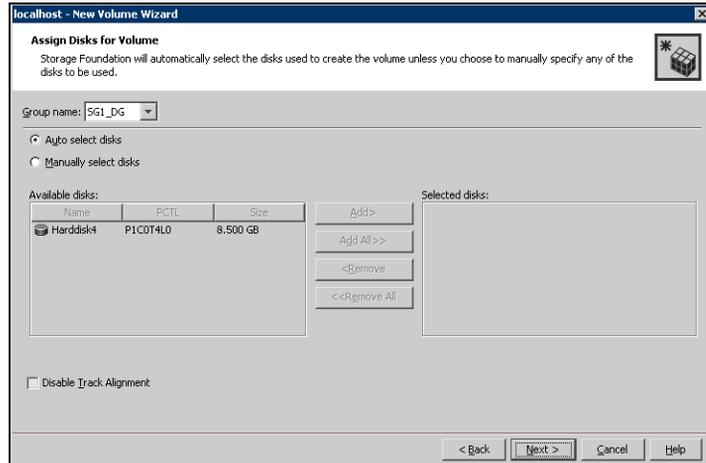
Creating volumes

This procedure assumes you are starting with the EXCH_SG1_DB1 volume.

To create dynamic volumes

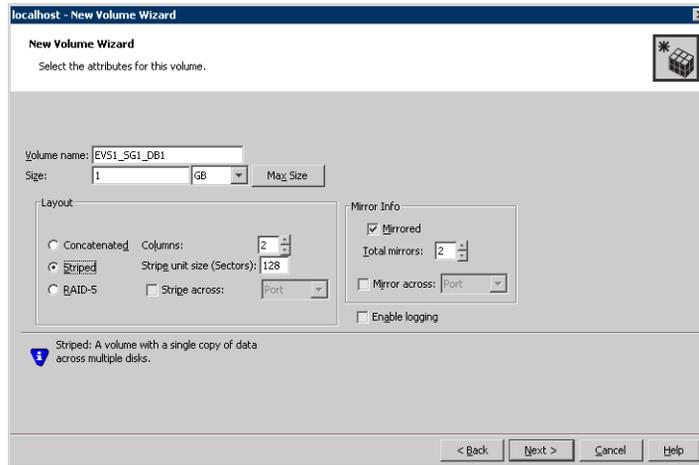
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu 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 EVS1_SG1_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.



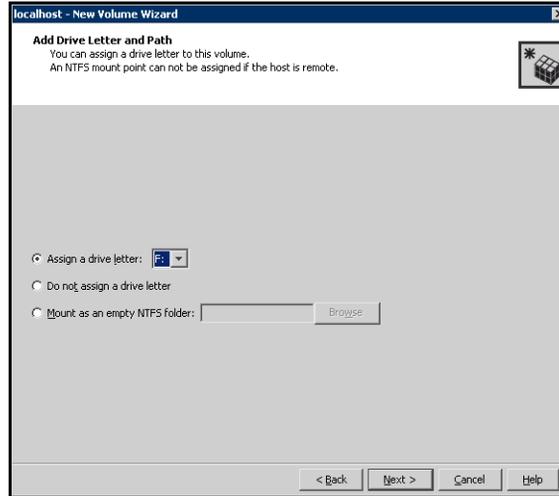
- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
 You may also check **Disable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

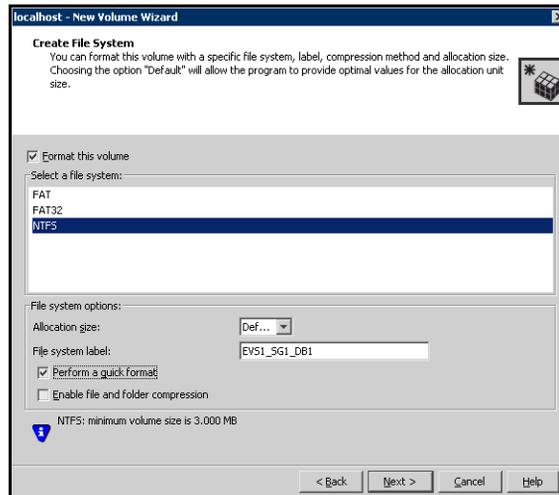


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- Do not assign the M: drive letter if you are using Exchange 2000, as it is reserved for internal use.
- 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 the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) Also, repeat these steps to create both the RegRep volume (EVS1_REGREP) and the EVS1_SHARED volumes in the EVS1_SHARED_DG disk group.
- 15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3). Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk groups and volumes

This section includes the following procedures:

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

Note: Verify the volume created to store registry replication information is mounted on this node and unmounted from other nodes in the cluster.

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Converting the standalone Exchange server into a “Clustered” Exchange server

Use the Exchange Setup Wizard to convert a standalone Exchange Server into a “clustered” Exchange server.

In this wizard, the node name of the standalone Exchange Server becomes the name of the Exchange virtual server and the existing node is given a new physical node name.

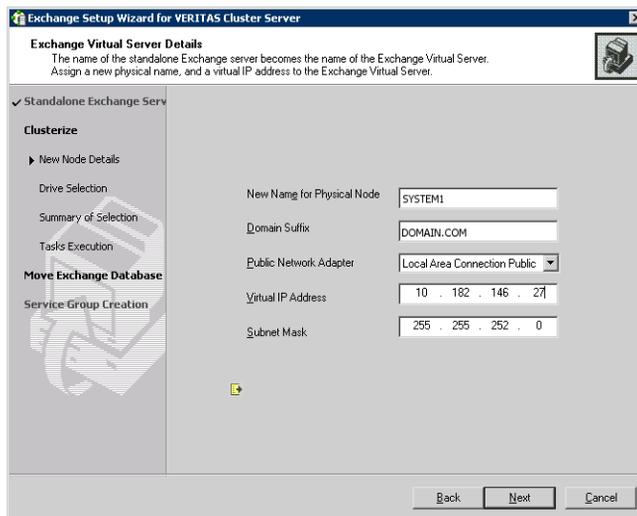
Renaming the existing standalone Exchange server allows Active Directory entries to remain valid. For example, if your existing standalone Exchange server is called EXCH, the name of the Exchange virtual server will become EXCH and the existing node is given a new physical node name, for example, SYSTEM1.

Note: Make sure the node hosting the Exchange virtual server, which will become highly available, is not configured as a root broker for a cluster.

To convert a standalone Exchange server into a “clustered” Exchange server

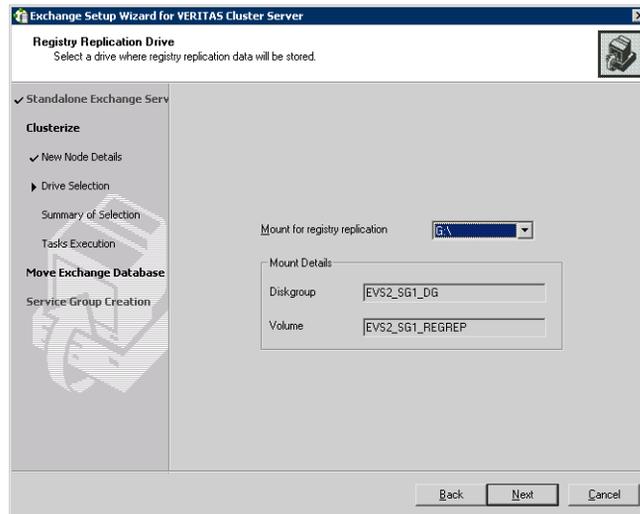
- 1 Start the Exchange Setup Wizard for VCS from the node having the standalone Exchange server installed.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard**.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the option **Make a standalone Exchange Server highly available** and click **Next**.

- 4 Specify information related to your network. Make sure to store the virtual name and IP address for future use.



- Enter a name for the node, for example SYSTEM1. This name for the node becomes the new name of the physical system after the process is completed. The original name of the system, for example, EXCH, is returned as the name of the Exchange virtual server so that the Active Directory entries remain valid.
- Enter the domain suffix.
- Select the appropriate public network adapter from the drop-down list. The installer displays all low priority TCP/IP enabled adapters on a system, including the private network adapters. Make sure that you select the adapters for the public network, and not those assigned to the private network.
- Enter a unique virtual IP address for the Exchange virtual server. If you plan to use the IP address of the node as the virtual IP address, you must assign a new static IP address to the node.
- Enter the subnet mask for the virtual IP address.
- Click **Next**.

5 Specify the information for registry replication:



- Select the drive letter (or directory in the case of folder mounts) for registry replication. Select a shared drive to allow failover to occur.
 - Click **Next**.
- 6 Review the summary information. Click **Next** to continue or **Back** to make changes.
 - 7 After reviewing the warning message about the renaming and rebooting of the system, click **Yes** to continue.
 - 8 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
 - 9 Click **Finish**.
 - 10 The wizard prompts you to restart the system. Click **Yes** to restart the system. Click **No** to restart the system later.
You must restart the system before continuing with the next step.

Adding the standalone Exchange server to a cluster

After converting the standalone Exchange server into a virtual server, create a cluster, if one does not already exist, and add all the nodes to the cluster.

Standalone Exchange server, plus a new node

If no cluster exists, check the prerequisites in “[Prerequisites for a new cluster](#)” on page 133 and then use “[Creating a new cluster and adding nodes](#)” on page 134 to create a new cluster and add all the nodes.

Standalone Exchange server and a cluster of nodes that may be running other applications

If a cluster already exists, check the prerequisites in “[Prerequisites for adding nodes to an existing cluster](#)” on page 149 and then continue with the procedure “[Adding nodes to an existing cluster](#)” on page 150 to add any new nodes, including the standalone Exchange server, to the cluster.

Prerequisites for a new cluster

The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures the Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

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

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

If no cluster exists, continue with “[Creating a new cluster and adding nodes](#)” on page 134 to create a new cluster and add the nodes.

If a cluster already exists (Scenario II), check the prerequisites in “[Prerequisites for adding nodes to an existing cluster](#)” on page 149 and then continue with the

procedure “[Adding nodes to an existing cluster](#)” on page 150 to add any new nodes, including the standalone Exchange server, to the cluster.

Creating a new cluster and adding nodes

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 the 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.

In the examples, below the system names are SYSTEM1 and SYSTEM2. Remember that the original name of your existing standalone Exchange server, EXCH in the example, is now the name of the virtual server, and the physical node has a new name, in the example, SYSTEM1.

Complete the following tasks before creating a cluster:

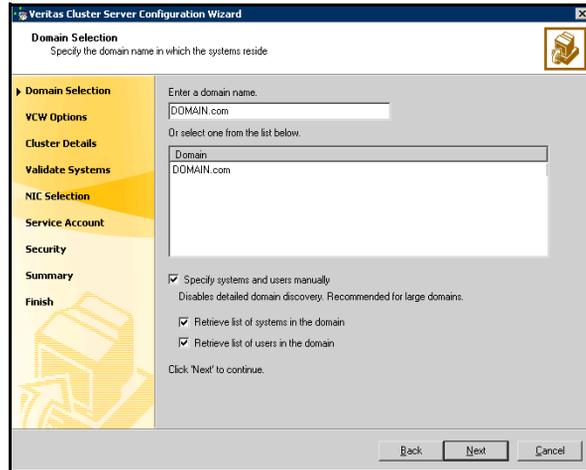
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative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

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

To configure a VCS cluster

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

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



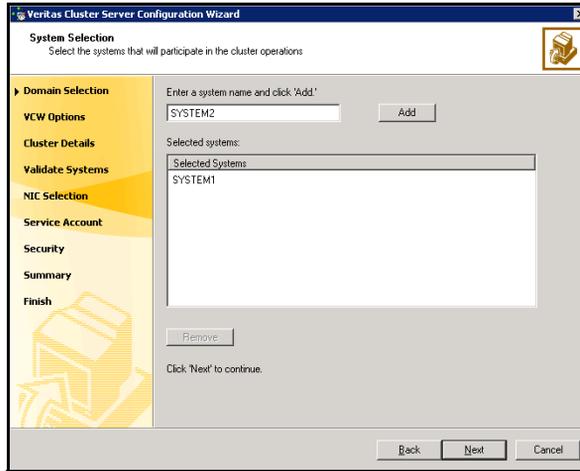
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
Proceed to [step 7](#) on page 137.

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

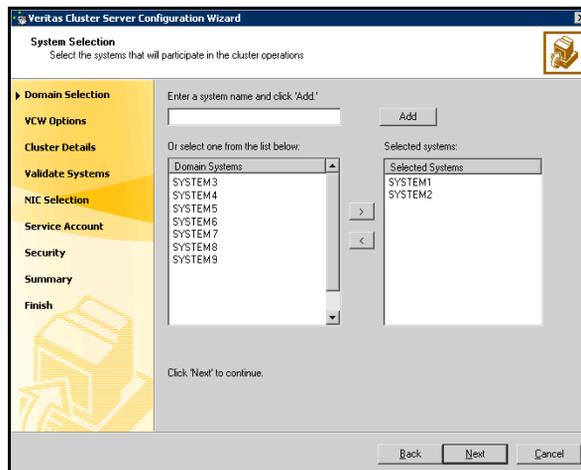
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 136. Otherwise, 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.



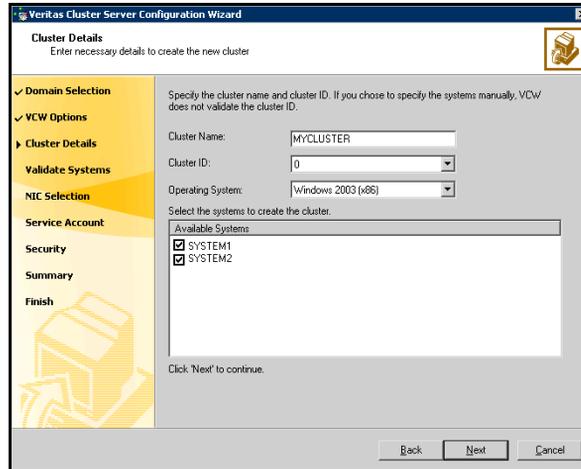
Proceed to [step 7](#) on page 137.

- 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**.

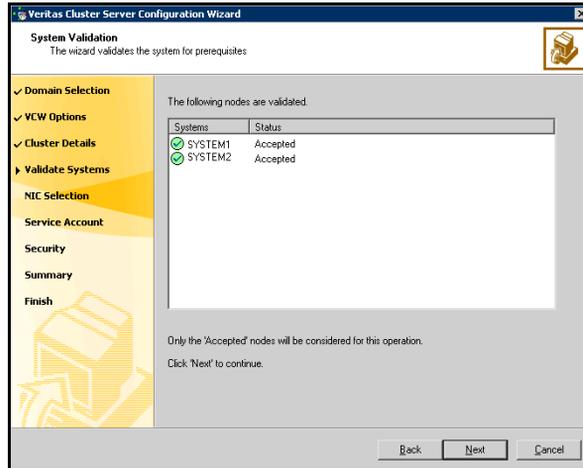


Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

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.

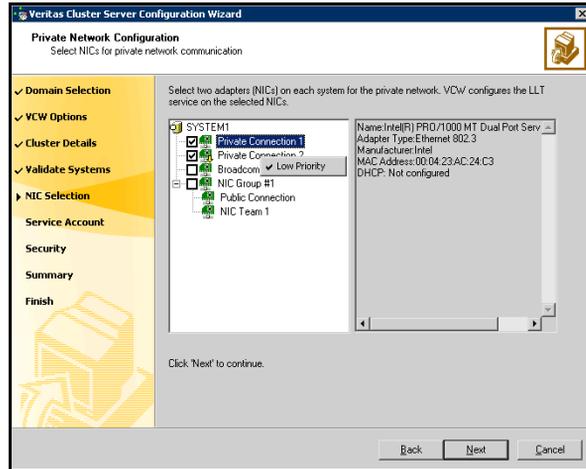
- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



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 137, proceed to the next step. Otherwise, proceed to [step 11](#) on page 139.

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

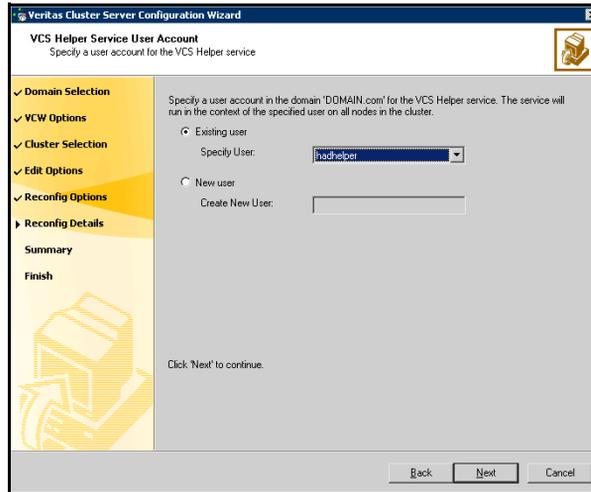


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs 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 135, 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**.

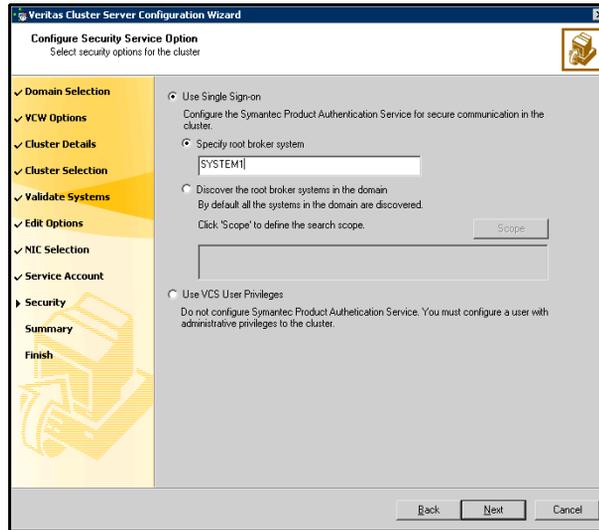
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- 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.

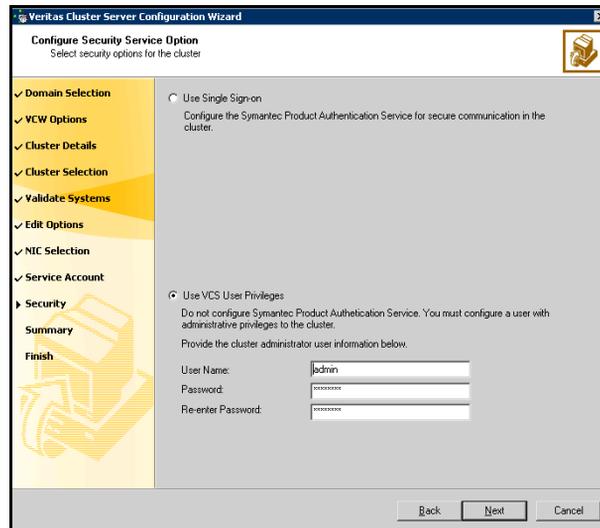
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- 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.

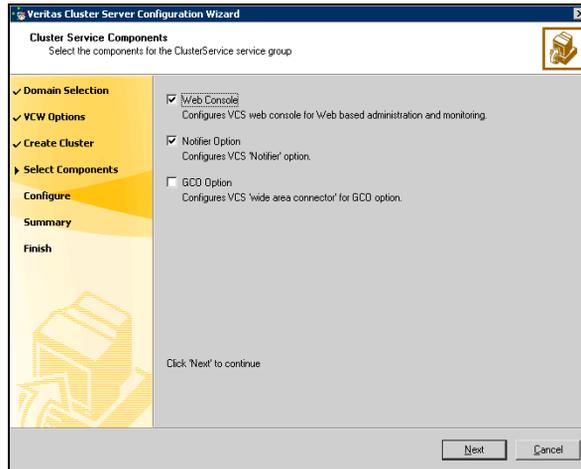
- Click **Next**.

- 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**.



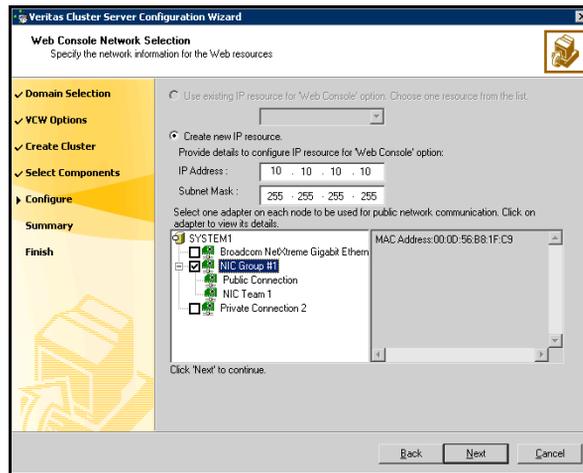
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 145.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 146.

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**.



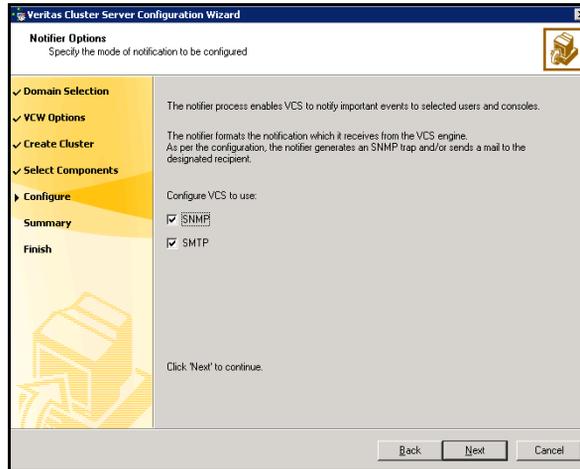
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 146. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

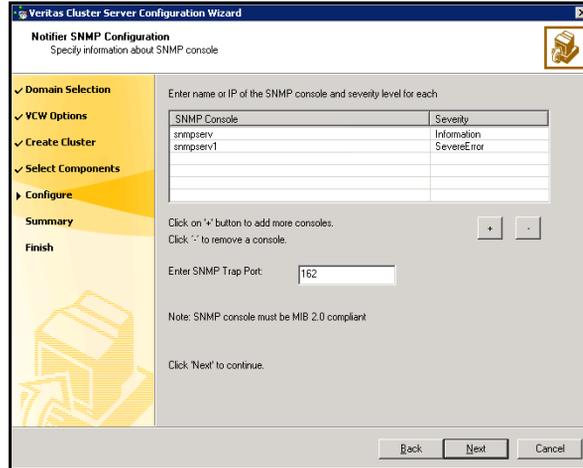
To configure notification

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



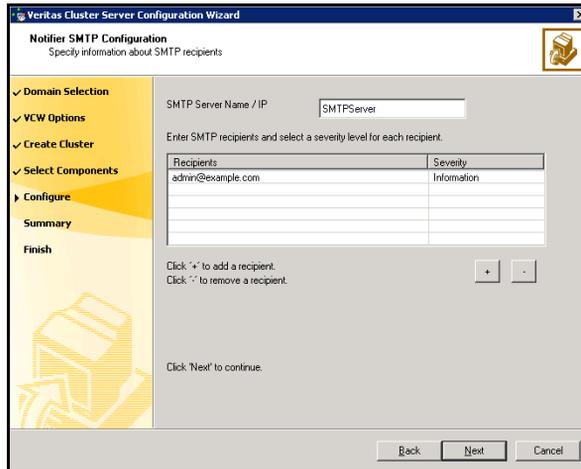
You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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



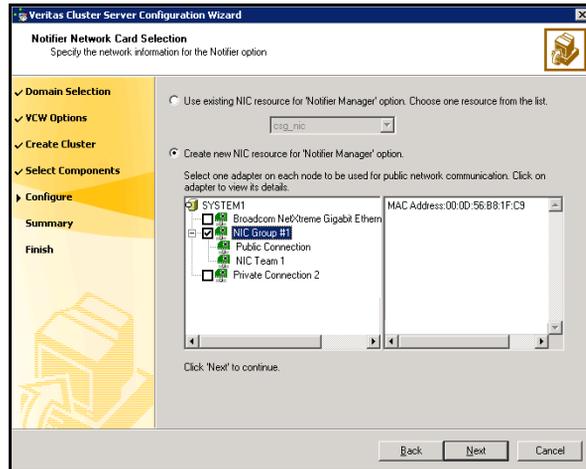
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.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

If you have completed the previous procedure, skip to [“Moving Exchange databases to shared storage”](#) on page 160.

Prerequisites for adding nodes to an existing cluster

This is scenario II, a standalone Exchange server and a cluster of nodes that may be running other applications. The standalone Exchange server and any new nodes must be added to the existing cluster.

In the examples, below the system names are SYSTEM1 and SYSTEM2. Remember that the original name of your existing standalone Exchange server, EXCH in the example, is now the name of the virtual server, and the physical node has a new name, in the example, SYSTEM1.

Check this list of prerequisites before beginning the procedure to add the nodes to the existing cluster:

- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk group and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - Registry changes related to Exchange
 - Transaction logs for the first storage group
 - MTA databaseSee “[Importing a disk group and mounting a shared volume](#)” on page 128 for instructions on mounting and “[Unmounting a volume and deporting a disk group](#)” on page 128 for instructions on unmounting.
- Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator’s Guide* to add additional resources to the EVS1_SG1_DG disk group.

Adding nodes to an existing cluster

This procedure applies only to an existing cluster running other applications, and you want to bring your standalone Exchange server into the cluster.

In the examples below the system names are SYSTEM1 and SYSTEM2.

Remember that the original name of your existing standalone Exchange server,

EXCH in the example, is now the name of the virtual server, and the physical node has a new name, in the example, SYSTEM1.

This section includes optional instructions to configure the ClusterService group for the VCS Cluster Management Console (Single Cluster Mode) also referred to as Web Console or notification after adding a node to the cluster.

To add a node to a cluster

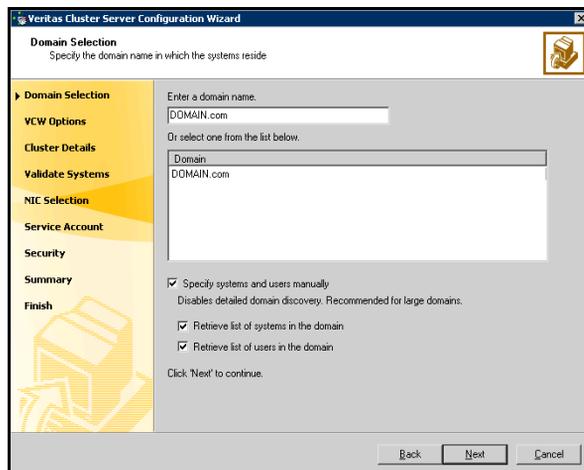
Note: Run the VCS Configuration Wizard from the standalone node or a node in the cluster.

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 154.

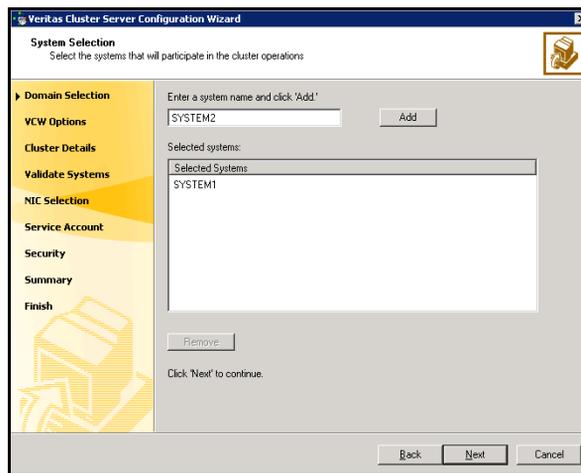
To specify systems and user names manually (recommended for large domains):

- 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 153. Otherwise proceed to the next step.

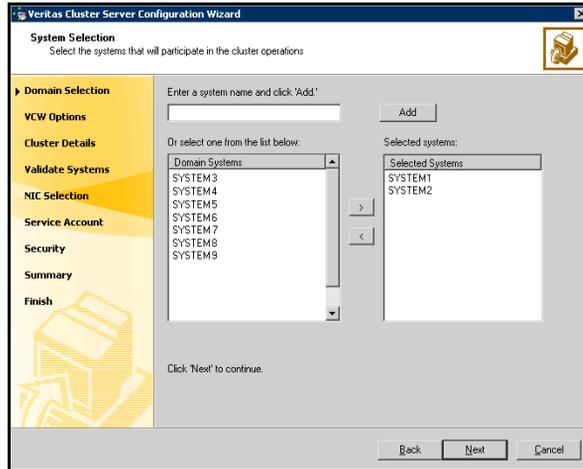
- 5 On the System Selection panel, complete the following and click **Next**.



- 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 154.

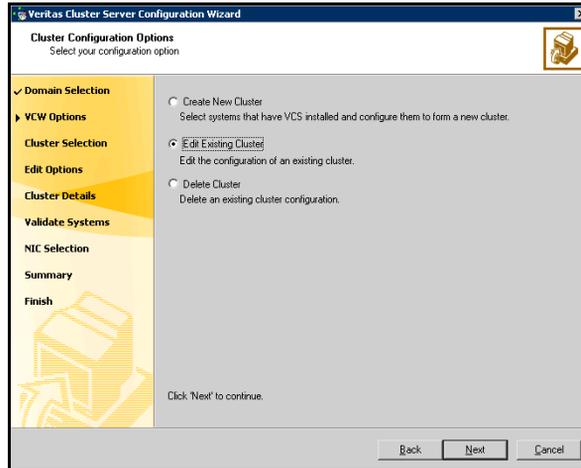
- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems.



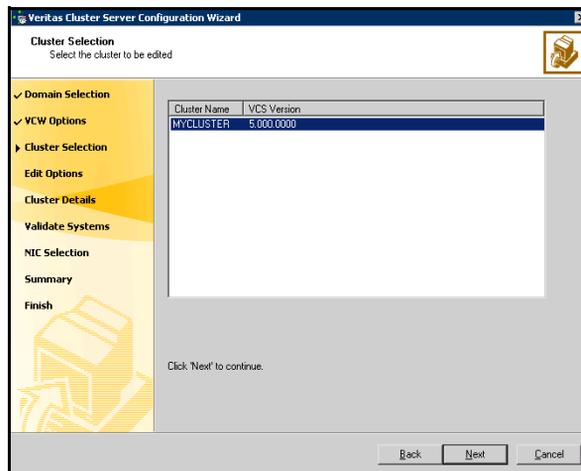
Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

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**.

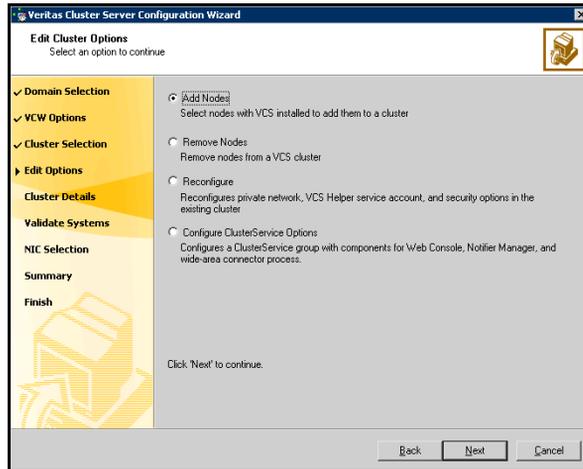


- 8 On the Cluster Selection panel, select the cluster to be edited and click **Next**.



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

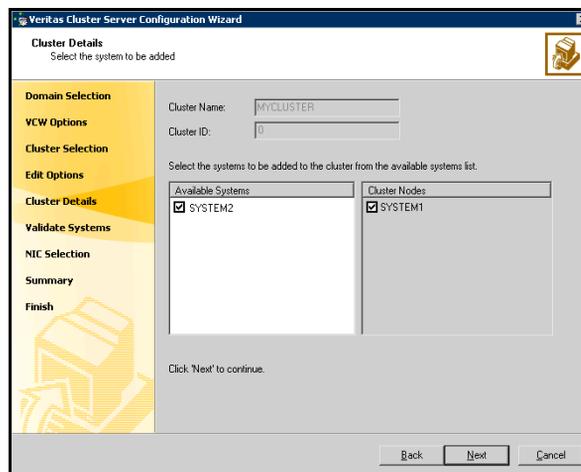
- 9 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**.



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**.

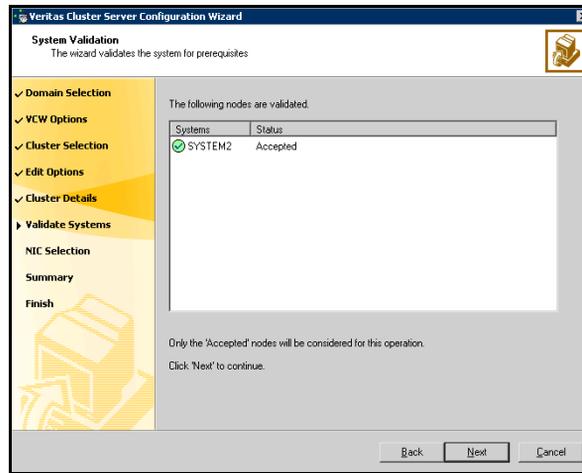
The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges (a cluster that is not a secure cluster).

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



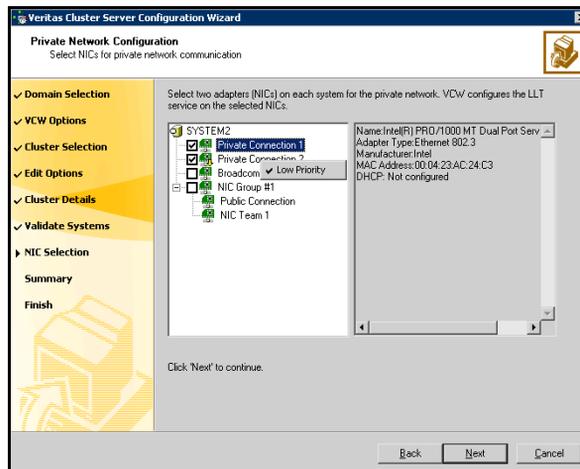
The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

- 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.

If you need the VCS Web Console or notification for the cluster, proceed to the next procedure, “[Modifying the ClusterService group for VCS](#)” on page 158.

If you do not need to configure the VCS Web Console and notification, skip to the next task list in “[Moving Exchange databases to shared storage](#)” on page 160.

If a new ClusterService group needs created, be sure to complete the procedure, “[Configuring the Exchange service group for VCS](#)” on page 169 when this procedure appears in the sequence.

Modifying the ClusterService group for VCS

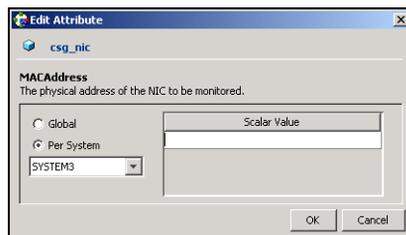
If you plan on setting up the VCS Cluster Management Console (Single Cluster Mode) also referred to as Web Console, or notification for the cluster, you must manually alter the resources for the ClusterService service group.

- Use the VCS Java Console to configure the NIC, IP, and VCSweb resources in the ClusterService group for the VCS Web Console.
- Use the VCS Java Console to configure the NIC resource in the ClusterService group for notification.

Note: Refer to the *Veritas Cluster Server Administrator’s Guide* for complete details on using the VCS Java Console and configuring the VCS Web Console and Notifier resource.

To configure the ClusterService group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 From the Service Groups tab of the Cluster Explorer configuration tree, expand the NIC resource type and select the **csg_nic** resource.
- 3 In the Properties tab of the view panel, click the **Edit** icon for the **MACAddress** attribute.
- 4 In the Edit Attribute dialog box:

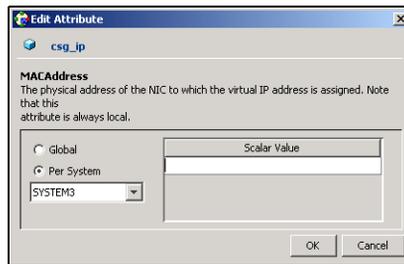


- Select the **per system** option, and select the newly added system.
- Enter the scalar value. To obtain the MAC address, run the `ipconfig /all` command from the command prompt on that system.
- Click **OK**.

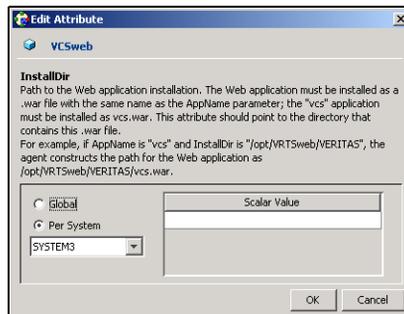
If you are only configuring notification, skip to the tasks “[Moving Exchange databases to shared storage](#)” on page 160.

If you are configuring the Web Console, proceed to step 5.

- 5 From the Service Groups tab of the Cluster Explorer configuration tree, expand the IP resource type and click the **csg_ip** resource.
- 6 In the Properties tab of the view panel, click the **Edit** icon for the **MACAddress** attribute.
- 7 In the Edit Attribute dialog box:



- Select the **per system** option, and select the newly added system.
 - Enter the scalar value. To obtain the MAC address, run the `ipconfig /all` command from the command prompt on that system.
 - Click **OK**.
- 8 From the **Service Groups** tab of the Cluster Explorer configuration tree, expand the VRTSWebApp resource type and select the **VCSweb** resource.
 - 9 In the **Properties** tab of the view panel, click the **Edit** icon for the **InstallDir** attribute.
 - 10 In the Edit Attribute dialog box:



- Select the **per system** option, and select the system.

- Enter the scalar value. From the command prompt on that system, type the following command to obtain the value:
`C: \>set VCS_ROOT`
Attach "\\VRTSweb\Veritas" to the end of the generated value to determine the scalar value.
- Click **OK**.

11 On the File menu of Cluster Explorer, click **Save Configuration**.

Moving Exchange databases to shared storage

Move the Exchange databases on the existing standalone node, which will belong to the new Exchange virtual server, from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

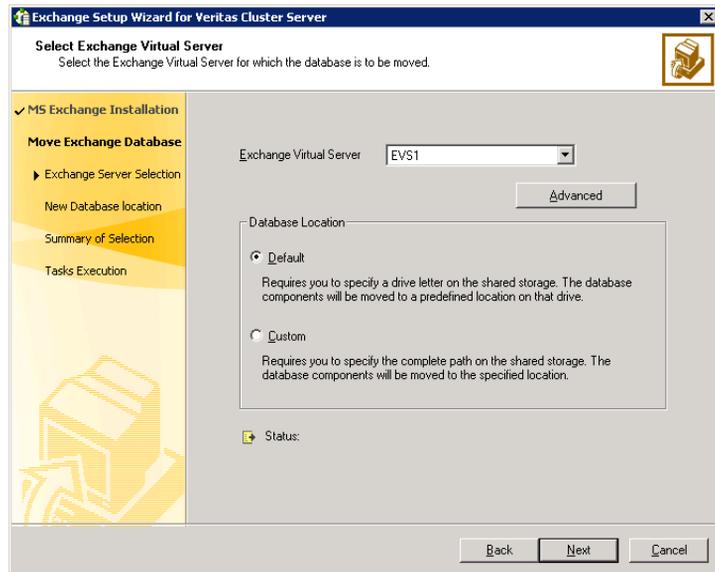
- Make sure the data queue is empty on the SMTP server.
- Make sure to import the disk group and mount the volumes for the Exchange database, MTA data, and transaction logs.
See "[Managing disk groups and volumes](#)" on page 128.

In the following example, your former standalone Exchange server is called EVS1, for the first Exchange Virtual server. Remember that your standalone Exchange server was renamed to the Exchange virtual server, to preserve Active Directory entries.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.

5 In the Select Exchange Virtual Server dialog box:



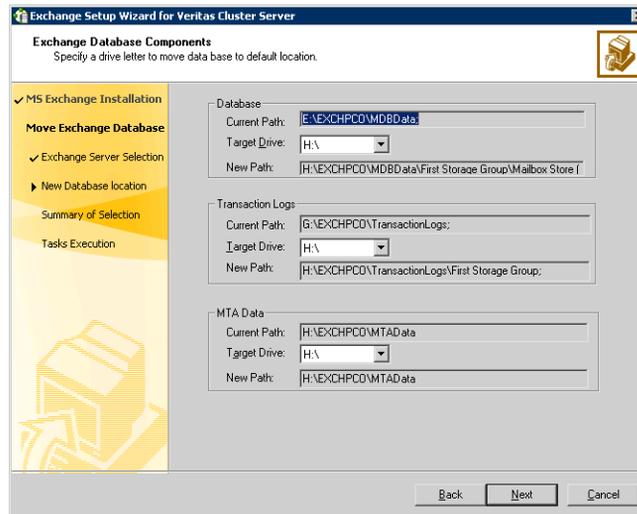
- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

Warning: The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Click **OK**.
 - Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
 - Click **Next**.
- 6 Do one of the following:
- If you would like to move the first mailbox store, public store, and MTA data to the generated default paths on the volumes for these components, proceed to the next step.
 - Otherwise, proceed to [step 8](#) on page 163 to specify the path location on the volumes that you will designate for these components.

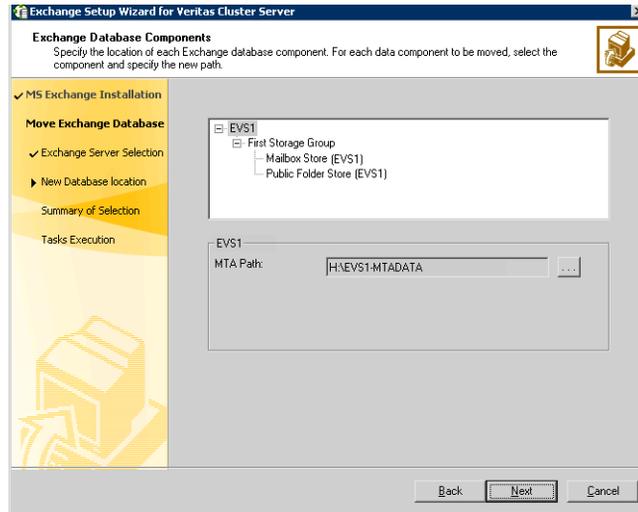
Warning: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

- 7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components will be moved to a default location on that drive. In the Exchange Database Components dialog box:



- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Specify the drive where the Exchange MTA Data will be moved.
- Click **Next** and proceed to [step 9](#) on page 163.

- 8 For the option of a custom database location, specify the location for specific Microsoft Exchange data components. In the Exchange Database Components dialog box:



For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the path for the Exchange database components contains only ANSI characters.

- 9 Review the summary of your selections and click **Next**.
- 10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.
- 11 Click **Finish** to exit the wizard.

Installing Exchange on additional nodes

After moving the Exchange databases to shared storage, install Exchange on additional nodes in the cluster for the same Exchange virtual server. You must run preinstallation, installation, and post-installation procedures for each additional node.

Installing Exchange on additional nodes is described in three stages that involve preinstallation, installation, and post-installation procedures. Complete the following tasks before installing Exchange Server:

- Verify the disk group is imported on the existing Exchange node of the cluster. Refer to “[Importing a disk group and mounting a shared volume](#)” on page 128 for instructions.
- Mount the volume containing the information for registry replication (EVS1_SG1_REGREP). Refer to “[Importing a disk group and mounting a shared volume](#)” on page 128 for instructions.
- Verify that all systems on which Exchange Server will be installed have IIS installed. You must install SMTP, NNTP, and WWW services on all systems. If you install Exchange on Windows 2003, make sure to install ASP.NET service.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.
 - Make sure the HAD Helper domain user account has “Add workstations to domain” privilege enabled in the Active Directory. To verify this, click **Start > Administrative Tools > Local Security Policy** on the domain controller to launch the security policy display. Click **Local Policies > User Rights Management** and make sure the user account has this privilege.
 - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.

The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

Exchange pre-installation: additional nodes

Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes.

See “[Unmounting a volume and deporting a disk group](#)” on page 128.

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

Remember that the Exchange virtual server name was formerly the name of your standalone Exchange server. In the example below, the name EVS1 is the example virtual server name.

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server.
The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when

the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically generated when setting up an additional failover mode for the virtual server in the same cluster.

- Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
 - 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.
 - 11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
 - 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: additional nodes

Install Exchange on the same node selected in “[Installing Exchange on additional nodes](#)” on page 163.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

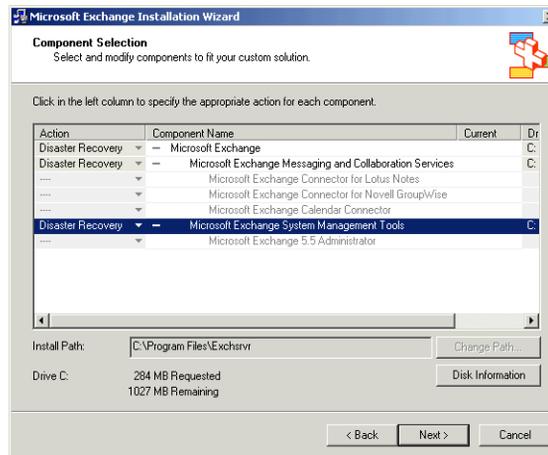
The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe /disasterrecovery
```

 where <drive letter> is the location where the Exchange software is located.
- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

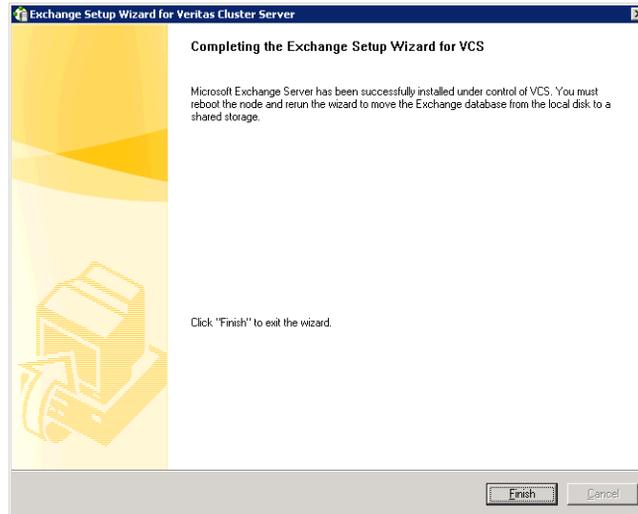
Exchange post-installation: additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.

- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

Configuring the Exchange service group for VCS

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

Prerequisites

- You must be a Cluster Administrator.

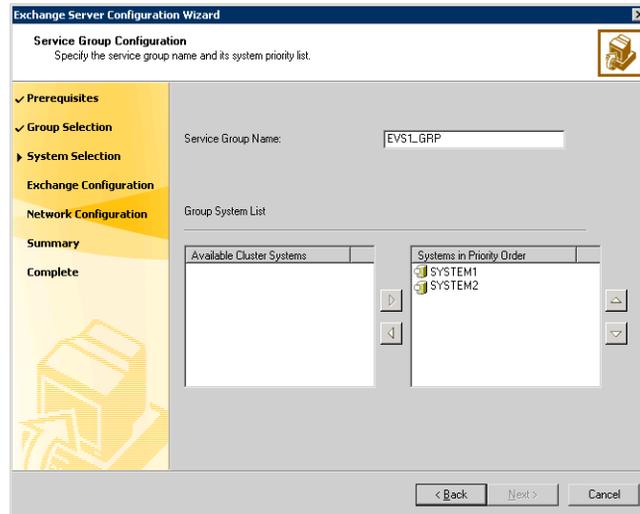
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk group and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - Registry changes related to Exchange
 - Transaction logs for the first storage group
 - MTA databaseSee “[Importing a disk group and mounting a shared volume](#)” on page 128 for instructions on mounting and “[Unmounting a volume and deporting a disk group](#)” on page 128 for instructions on unmounting.
- Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to the EVS1_SG1_DG disk group.

To configure the Exchange service group

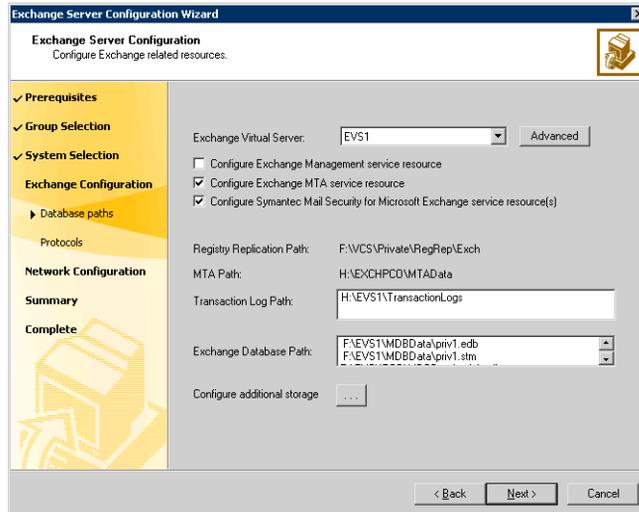
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options dialog box, choose the **Create service group** option and click **Next**.

4 Specify the service group name and system list:



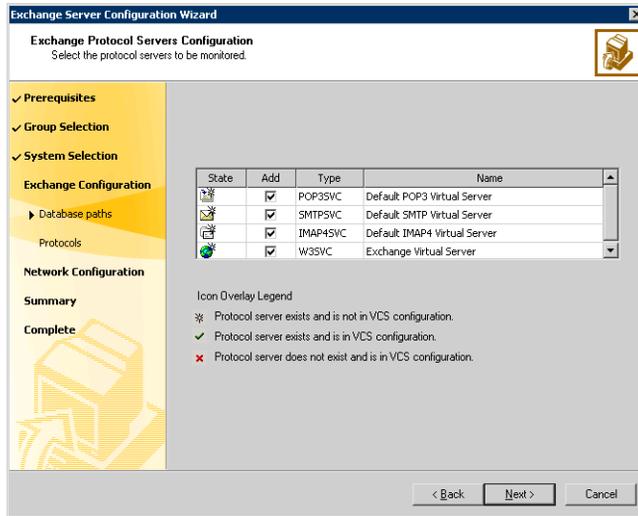
- Enter a name for the Exchange service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group’s system list. Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- To remove a system from the service group’s system list, select the system in the Systems in Priority Order list and click the left arrow.
- To change a node’s priority in the service group’s system list, select the node in the Systems in Priority Order list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:

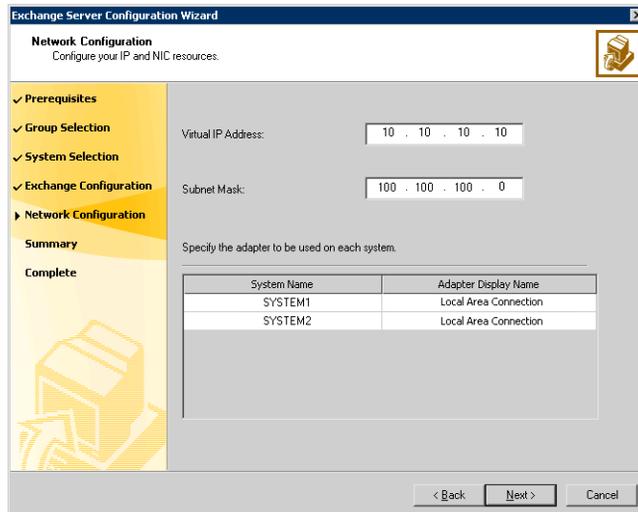


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.
- Specify whether you want to configure the Exchange Management service.
 - Specify whether you want to configure the Exchange MTA service resource. The Exchange Message Transfer Agent is specific to legacy Exchange message routing based on X.400. There are specific circumstances where it may or may not be required for your Exchange server. Please refer to Microsoft documentation on Exchange message routing and the use of MTA for further clarification and applicability to its use within your Exchange environment.

- If you are utilizing Symantec Mail Security for Exchange be sure to indicate that you would like to configure this resource.
 - Click **Next**.
- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



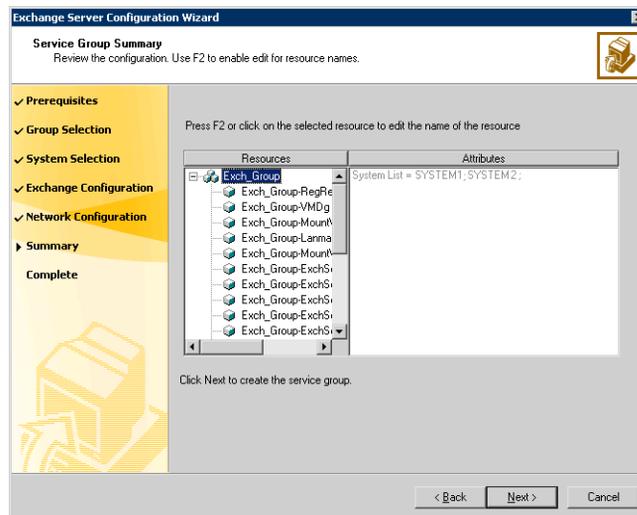
- 7 Specify information related to the network:



- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

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

- Click **Next**.
- 8 Review the service group configuration and change the resource names, if desired:



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 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.

- Click **Next**.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 9 In the Completing the Exchange Configuration dialog box, select the **Bring the service group online** check box to bring the service group online on the local system.
 - 10 Click **Finish**. After bringing the service group online, you must run the Exchange System Manager so that all the stores are automatically mounted on start-up.

To reconfigure mounting of stores at start-up

- 1 Start Exchange System Manager.
- 2 In the left pane, navigate to your storage group.
If administrative groups are not configured, Servers > Exchange Server > Storage Group.
If more than one administrative group is configured, expand Administrative Groups > Your Administrative Group > Servers > Exchange Server > Storage Group.
- 3 Right-click the Exchange database and choose **Properties** from the pop-up menu.
- 4 Click the Database tab.
- 5 Clear the **Do not mount this store at start-up** check box.
- 6 Click **OK**.

Repeat these steps for all the Exchange databases that were previously mounted.

If you need to configure additional storage groups or mailbox stores on the shared storage you should do that now. Import disk groups and mount volumes that have been created for the additional storage groups or mailbox stores, and create the new storage groups and mailbox stores in Exchange System Manager, and then rerun the Exchange Configuration Wizard to bring them under VCS control. If you already designated the additional mounted volumes and disk groups when you ran the configuration wizard the first time then you can just create the storage groups and mailbox stores in Exchange System Manager.

Your SFW HA environment is now complete.

Verifying the cluster configuration

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

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

To switch service groups

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

To shut down an active cluster node

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

Deploying SFW HA for high availability: Configuring a new any-to-any failover

This chapter covers the following topics:

- [Reviewing the configuration](#)
- [Reviewing the requirements](#)
- [Configuring the storage hardware and network](#)
- [Preparing the forest and domain](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Configuring the cluster](#)
- [Configuring the first Exchange Virtual Server](#)
- [Configuring another Exchange virtual server for an any-to-any failover](#)

You can either install and configure a new “any-to-any” SFW HA environment for Exchange to provide a production node with multiple failover nodes or, you can transform an existing active/passive SFW HA environment for Exchange into an any-to-any environment.

See [Chapter 7, “Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover”](#) on page 259.

[Table 6-1](#) outlines the high-level objectives to create a new any-to-any environment and the tasks to complete each objective:

Table 6-1 Task list

Objective	Tasks
“Reviewing the configuration” on page 180	<ul style="list-style-type: none"> ■ Understanding a basic any-to-any Exchange configuration
“Reviewing the requirements” on page 183	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Configuring the storage hardware and network” on page 187	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which Exchange will be installed
“Preparing the forest and domain” on page 188	Setting up the forest and domain prior to the Exchange installation
“Installing Veritas Storage Foundation HA for Windows” on page 189	<ul style="list-style-type: none"> ■ Verifying the driver signing options for Windows 2003 systems ■ Installing SFW, VCS, and the Veritas Cluster Server Application Agent for Microsoft Exchange ■ Restoring driver signing options for Windows 2003 systems
“Configuring the cluster” on page 195	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the Veritas Cluster Server Configuration Wizard

Table 6-1 Task list (continued)

Objective	Tasks
“Managing disk groups and volumes” on page 217	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the Data, Log, RegRep, and SHARED volumes ■ Managing disk groups and volumes, with instructions for mounting and unmounting volumes
“Installing Exchange on the first node” on page 219	<ul style="list-style-type: none"> ■ Reviewing the prerequisite checklist ■ Running the Exchange Setup Wizard for Veritas Cluster Server and Microsoft Exchange Server installation ■ Performing this “First Node” installation on each of the active Exchange nodes in the final configuration. ■ After this task is complete, two or more Exchange Virtual Servers will exist, one for each of the active Exchange servers in the final configuration.
“Moving Exchange databases to shared storage” on page 223	<ul style="list-style-type: none"> ■ Moving databases on the first node from the local drive to the shared drive using the Exchange Setup Wizard for Veritas Cluster Server ■ Repeating this task for each of the active Exchange nodes in the final configuration, making sure that each of the active Exchange servers has a separate area for its databases. Do not share databases between separate Exchange servers.

Table 6-1 Task list (continued)

Objective	Tasks
“Installing Exchange on additional nodes” on page 226	<ul style="list-style-type: none"> ■ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes ■ Perform this task for all of the failover systems.
“Configuring the Exchange service group for VCS” on page 231	<ul style="list-style-type: none"> ■ Preparing the cluster for any-to-any failover using the Exchange Setup Wizard. This step must be completed on each of the Exchange Virtual Servers. ■ Configuring the Exchange service group for the second Exchange Virtual Server. If necessary, you can later add common failover nodes to the Exchange service group’s system list.
“Verifying the cluster configuration” on page 238	Verifying the cluster configuration by switching service groups and shutting down an active cluster node.

Reviewing the configuration

Configure an any-to-any configuration with new nodes transformed into an any-to-any configuration as in [Table 6-2](#):

Table 6-2 New nodes to any-to-any cluster

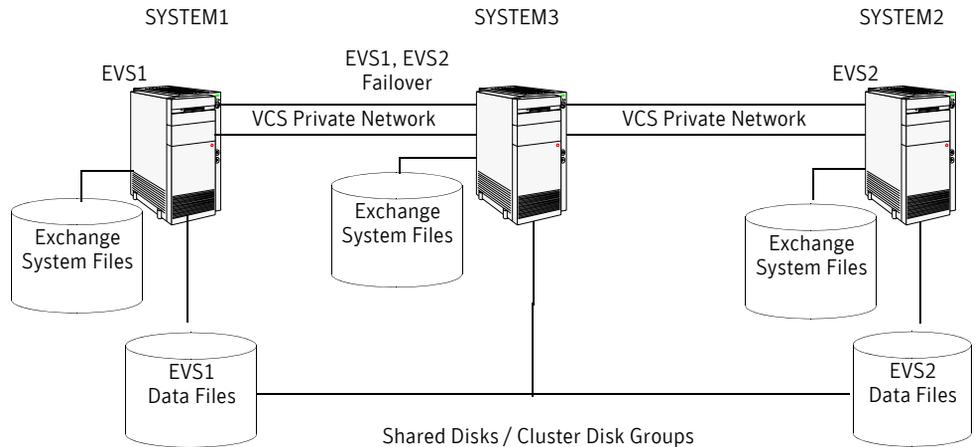
Exchange virtual server	Nodes	Any-to-any common failover node
EVS1	SYSTEM1	SYSTEM3
EVS2	SYSTEM2	SYSTEM3

With individual nodes, no failover capability exists. In an any-to-any configuration, the active Exchange nodes can share failover nodes. Additional failover nodes can also exist in an any-to-any configuration.

Any-to-any configuration

In an any-to-any configuration, each Exchange virtual server in the cluster is configured in a separate service group. Each service group can fail over to any configured node in the cluster, provided that no other Exchange virtual server is online on that node. You must ensure that an Exchange service group does not fail over to a node on which another Exchange service group is online. [Figure 6-1](#) shows an example of a three-node cluster.

Figure 6-1 Three-node cluster in an any-to-any configuration



For example, consider a three-node cluster hosting two Exchange Virtual Servers, EVS1 and EVS2. The virtual servers are configured in VCS in two service groups such that SYSTEM1 has first priority for the EVS1 service group and SYSTEM2 has first priority for the EVS2 service group, while SYSTEM3 is shared as a common failover node between the 2 virtual servers. If SYSTEM1 fails, the service group containing the EVS1 resources is failed over to SYSTEM3. If SYSTEM2 fails, the service group containing the EVS2 resources fails over to SYSTEM3.

Note: EVS1 and EVS2 cannot be online at the same time on SYSTEM3.

Sample configuration

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

Table 6-3 Sample configuration

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3	Physical node names
EVS1, EVS2	Microsoft Exchange Virtual Servers
EVS1_GRP, EVS2_GRP	Microsoft Exchange service groups
EVS1_SG1_DG, EVS2_SG1_DG	Cluster disk group names
EVS1_SG1_DB1, EVS2_SG1_DB1	Volumes for storing the Microsoft Exchange Server database
EVS1_SG1_LOG, EVS2_SG1_LOG	Volumes for storing a Microsoft Exchange Server database log file
EVS1_REGREP, EVS2_REGREP	Volumes that contain the list of registry keys that must be replicated among cluster systems for the Exchange server
EVS1_SHARED, EVS2_SHARED	Volumes for storing Microsoft Exchange Server MTA database, SMTP and message tracking for Exchange server

Reviewing the requirements

Two or more Exchange virtual servers can exist in an any-to-any configuration. Refer to “[Reviewing the configuration](#)” on page 180.

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

Disk space requirements

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

[Table 6-4](#) estimates disk space requirements for SFW HA.

Table 6-4 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 Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:

- Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition (SP2 required)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
or
Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP1 required for all editions)
or
Windows Server 2003 R2 (Standard Edition, Enterprise Edition, or Datacenter Edition)
or
- Microsoft Exchange 2000 Server or Microsoft Exchange 2000 Enterprise Server (SP3 with August 2004 rollup patch required for both editions)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)

Note: Microsoft support for Exchange Server 2003 is limited to 32-bit versions of the Windows 2003 operating system.

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 186.
- 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 Exchange Virtual Server (EVS)
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

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

See the *Veritas Cluster Server Bundled Agents Reference Guide*.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Domain Servers group.

- 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.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- You must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.
- 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.
- You must install IIS, SMTP, NNTP, ASP.NET, and WWW services before you install Microsoft Exchange Server.

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.

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**.

Preparing the forest and domain

Microsoft requires the preparation of the forest and domain prior to an Exchange installation. See Microsoft documentation for instructions for preparing the forest and domain for an Exchange installation. Do not repeat this process for additional Exchange installations.

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 Application Agent for Exchange. 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 6-5](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 6-5 Installation behavior with unsigned drivers

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

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

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.

- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

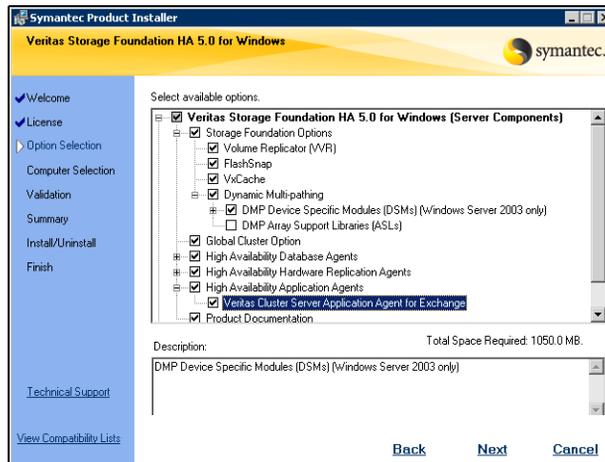
To install the product

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



- 4 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 5 Review the Welcome message and click **Next**.

- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
 If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 8 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.
- 9 Select the appropriate SFW product options for your installation. Click **Next**.

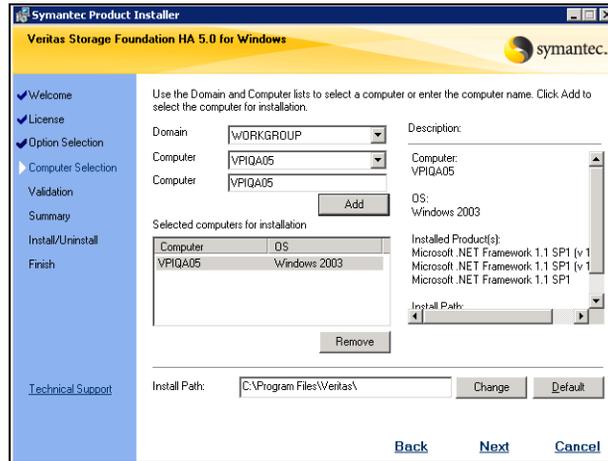


The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

Veritas Cluster Server Application Agent for Exchange	Required to configure high availability for Exchange Server.
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console. Required to install the Solutions Configuration Center which provides information and wizards to assist configuration.
Global Cluster Option	Required for a disaster recovery configuration only.

Veritas Volume Replicator If you plan to use VVR for replication, you must also select the option to install VVR.

10 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

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

Computer

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

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

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

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

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
 - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
 - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

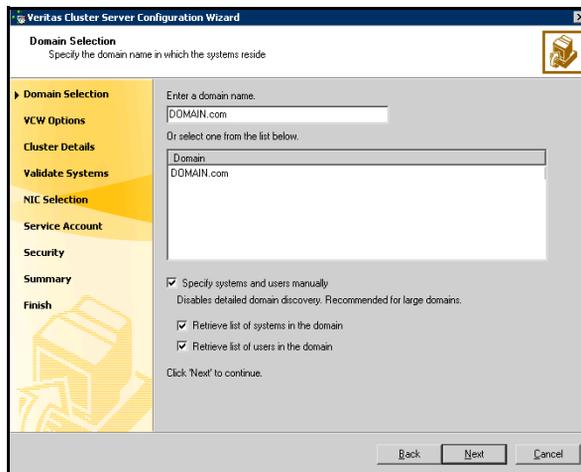
Complete the following tasks before creating a cluster:

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

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

To configure a VCS cluster

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



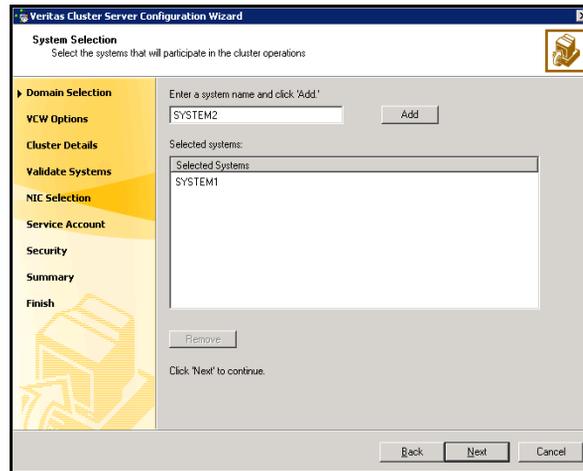
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
Proceed to [step 7](#) on page 198.

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

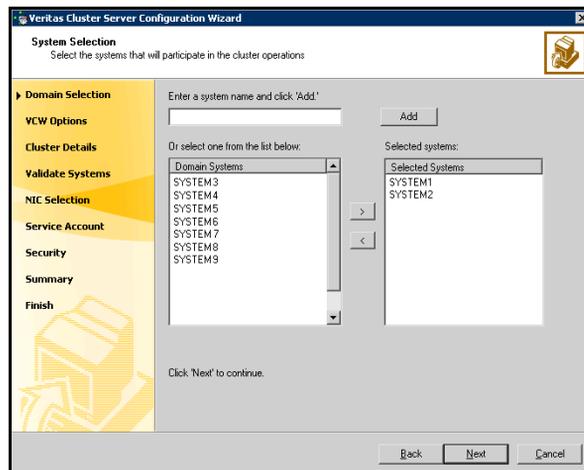
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 197. Otherwise, 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.



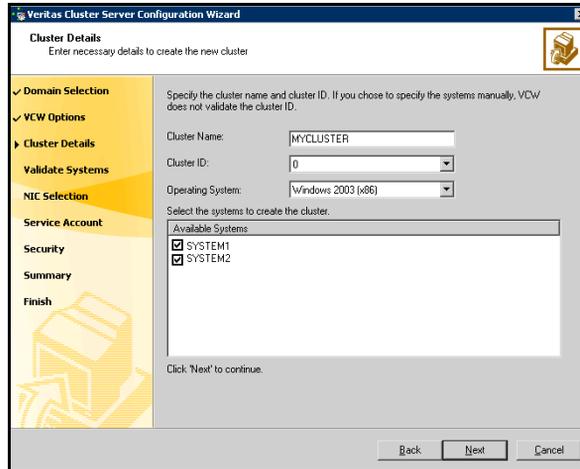
Proceed to [step 7](#) on page 198.

- 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.

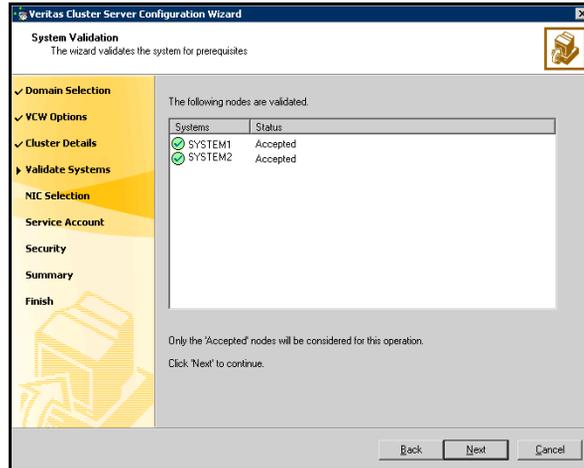
Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

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.

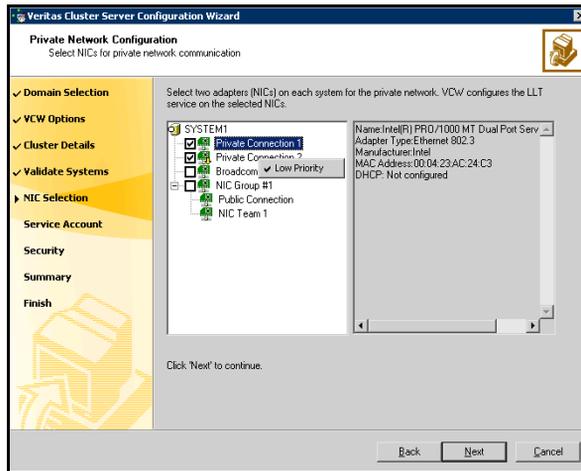
- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



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 198, proceed to the next step. Otherwise, proceed to [step 11](#) on page 200.

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

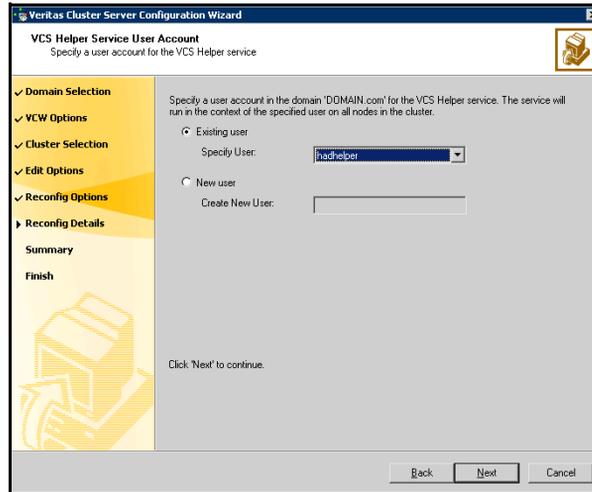


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs 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 196, 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**.

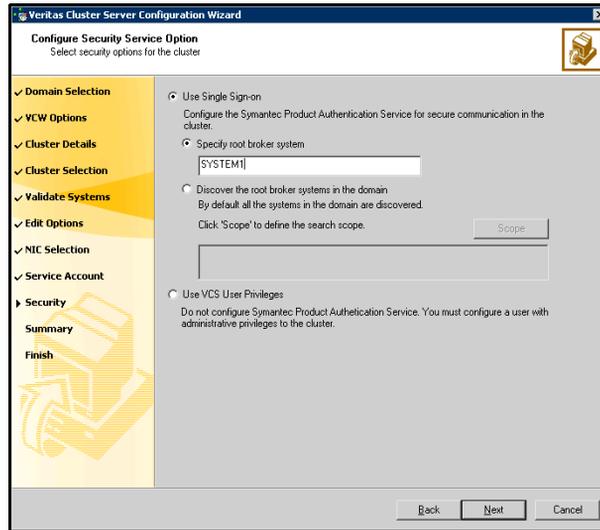
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- 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.

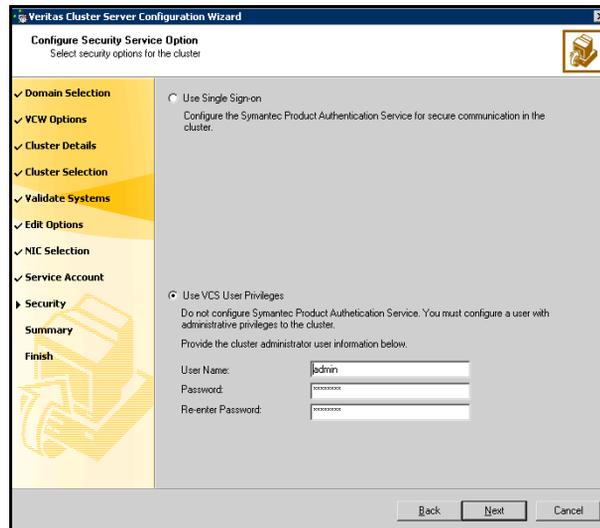
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- 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.

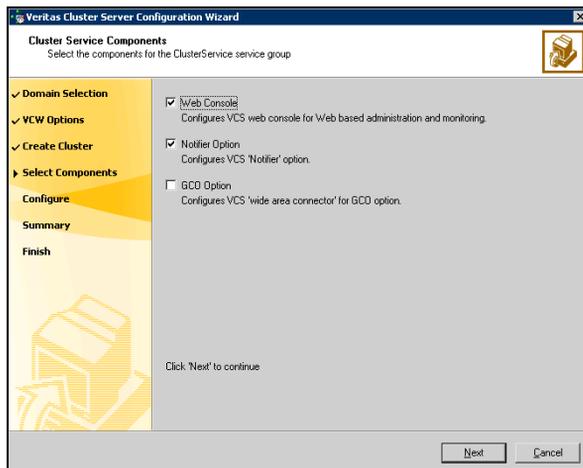
■ Click **Next**.

- 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**.



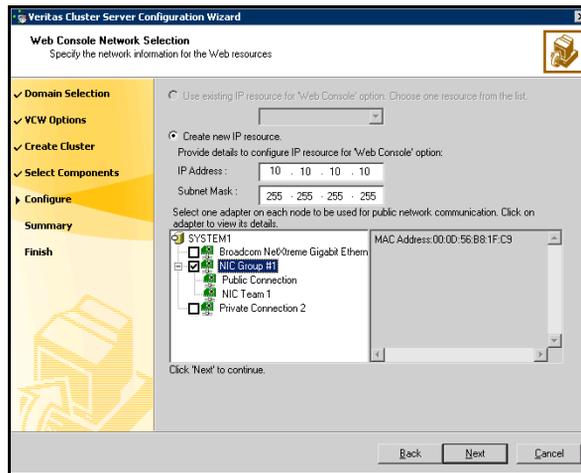
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 206.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 207.

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**.



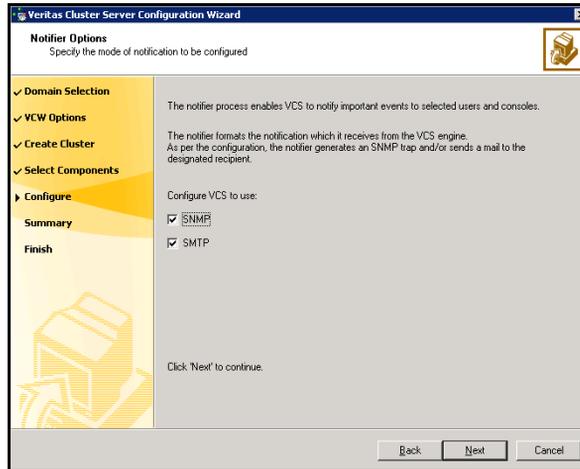
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 207. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

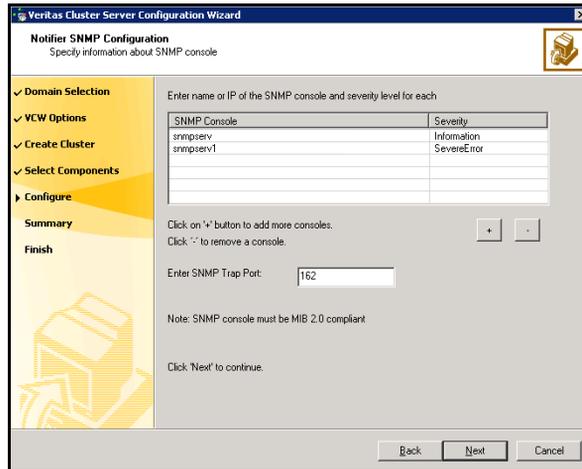
To configure notification

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



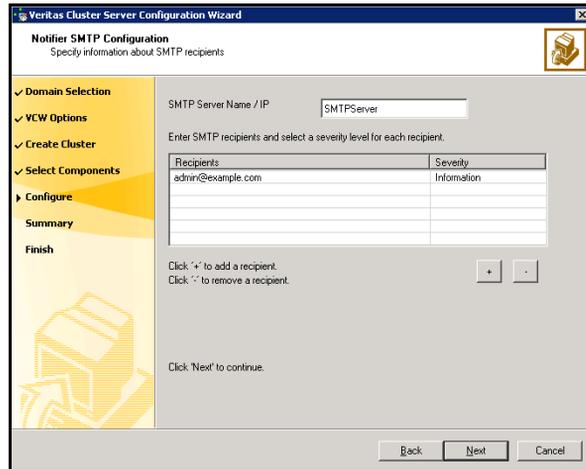
You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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



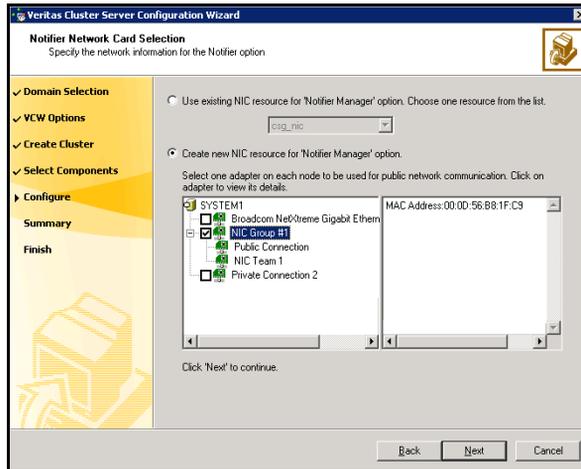
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- 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 the first Exchange Virtual Server

Use the procedures described in this section to install and configure a new Veritas Storage Foundation HA environment for Exchange on a new cluster with the any-to-any configuration.

See “[Reviewing the configuration](#)” on page 180.

All the “First Node” installation tasks need to be repeated on all of the active Exchange nodes in the any-to-any configuration.

Configuring disk groups and volumes

Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to grow existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

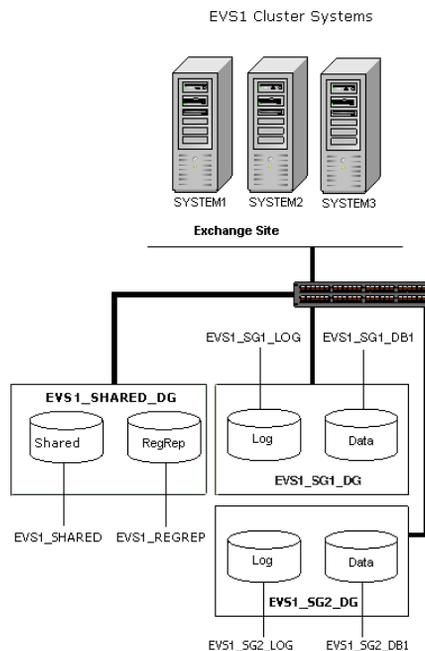
Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.

Typically, a SFW disk group corresponds to an Exchange storage group.

Figure 6-2 is a detailed view of the disk groups and volumes in an HA environment.

Figure 6-2 Disk groups and volumes for Exchange virtual server EVS1 in HA setup



Use the following procedures to create the appropriate disk groups and volumes. The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange disk group EVS1_SG1_DG contains two volumes:

- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

Exchange disk group EVS1_SHARED_DG create contains two volumes:

- EVS1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.
- EVS1_SHARED: Contains the MTA database, SMTP, and message tracking.

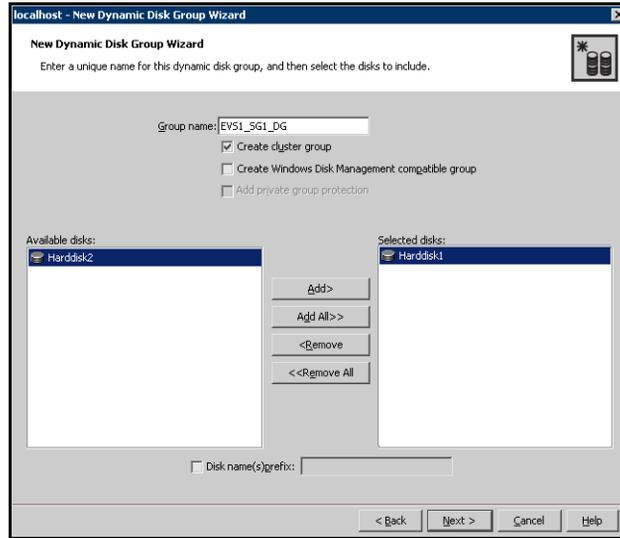
Note: Additional storage groups (for example, EVS1_SG2_DG) only contain the database, and log volumes; the RegRep and SHARED volumes are included in the EVS1_SHARED_DG disk group.

Creating a disk group

To create a dynamic (cluster) disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group:



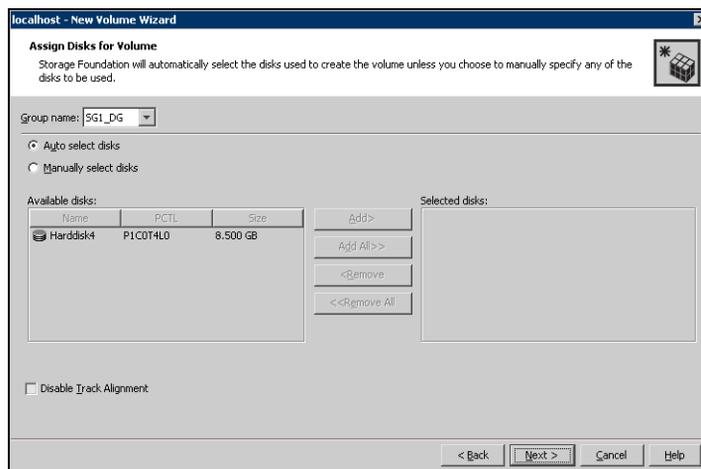
- Enter the name of the disk group (for example, EVS1_SG1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
 Optionally, check the **Disk names prefix** checkbox 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 assumes you are starting with the EVS1_SG1_DB1 volume.

To create dynamic volumes

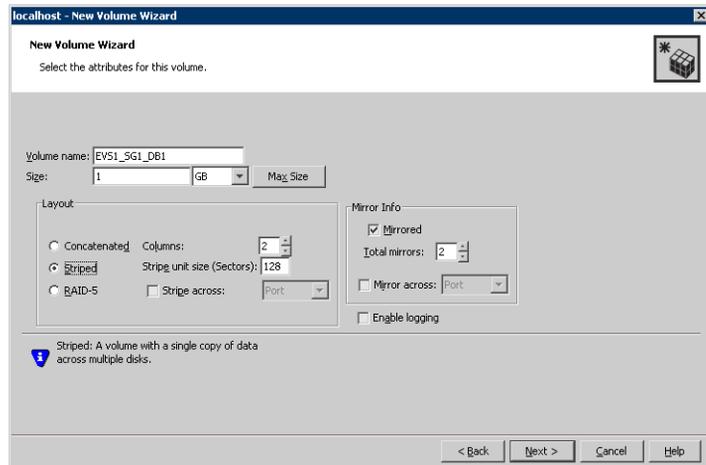
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu 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 EVS1_SG1_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.

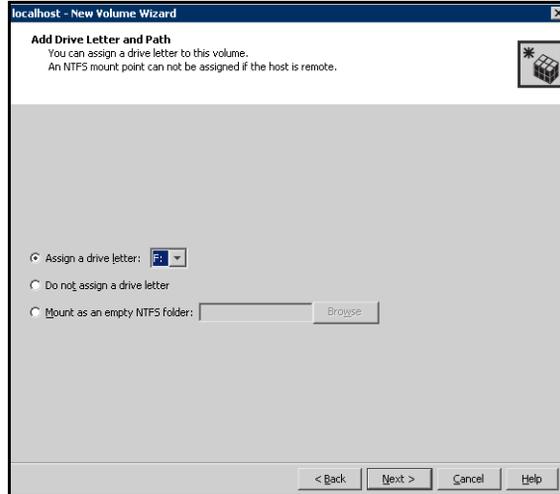
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling **Track Alignment** means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

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



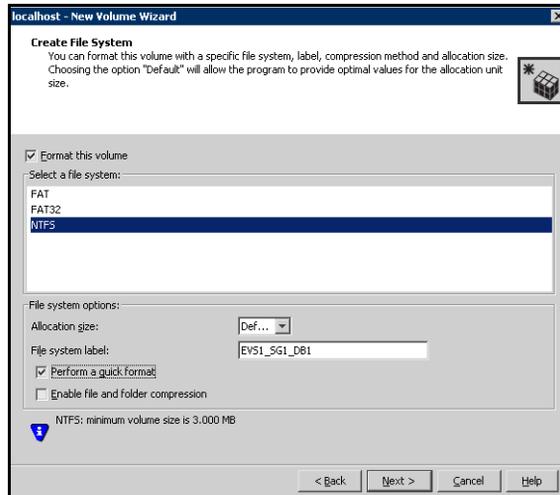
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
 Do not assign the M: drive letter if you are using Exchange 2000, as it is reserved for internal use.

- 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 the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) Also, repeat these steps to create both the RegRep volume (EVS1_REGREP) and the EVS1_SHARED volumes in the EVS1_SHARED_DG disk group.
- 15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3). Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk groups and volumes

This section includes the following procedures:

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

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Installing Exchange on the first node

Installing Exchange on the first node of EVS1 is described in three stages that involve preinstallation, installation, and post-installation procedures. Complete the following tasks before installing Exchange Server:

- Prepare the forest and domain. This must be done one time only, prior to the first time you install Exchange in your domain.
See “[Preparing the forest and domain](#)” on page 188.
- Verify the disk group is imported on the first node of the cluster.
See “[Managing disk groups and volumes](#)” on page 217.
- Mount the volume containing the information for registry replication (EVS1_REGREP).
See “[Managing disk groups and volumes](#)” on page 217.
- Verify that all systems on which Exchange Server will be installed have IIS installed; you must install SMTP, NNTP, and WWW services on all systems. If you install Exchange on Windows 2003, make sure to install ASP.NET service.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.
 - Make sure the HAD Helper domain user account has the “Add workstations to domain” privilege enabled in the Active Directory. To verify this, click **Start > Administrative Tools > Local Security Policy** on the domain controller to launch the security policy display. Click **Local Policies > User Rights Management** and make sure the user account has this privilege.

- If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

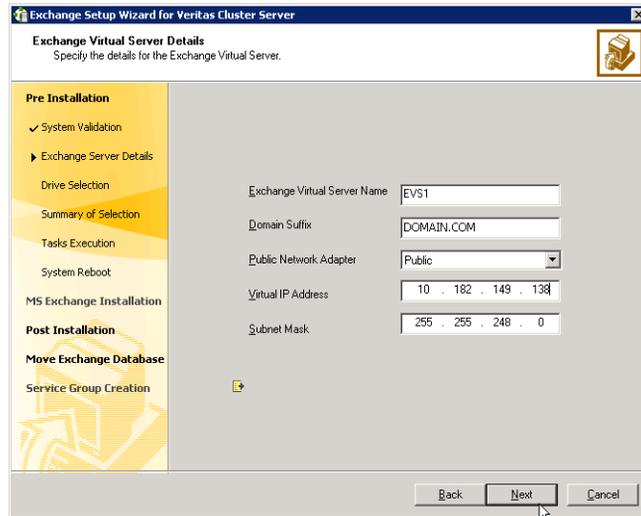
Exchange pre-installation: First node

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. You must install Exchange on a virtual node to facilitate high availability.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a new Exchange Virtual Server** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.

7 Specify information related to the network.



- Enter a unique virtual name for the Exchange server.
Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.
- Enter a domain suffix for the virtual server.
- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
- Enter a unique virtual IP address for the Exchange virtual server.
- Enter the subnet mask for the virtual IP address.
- Click **Next**.

The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met.

Click **Next**.

- 8 Select a drive where the registry replication data will be stored and click **Next**.

- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: First node

Install Exchange on the node selected in “[Exchange pre-installation: First node](#)” on page 220.

Exchange 2000 requires service pack 3 with the August 2004 rollup patch. Exchange 2003 requires service pack 2. The procedure below is based on Exchange 2003 SP2.

This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Install Exchange Server using the Microsoft Exchange installation program. See the Microsoft Exchange documentation for instructions.
- 2 Reboot the node if prompted to do so.
- 3 For Exchange 2000 or Exchange 2003, install the required service pack.

Exchange post-installation: First node

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

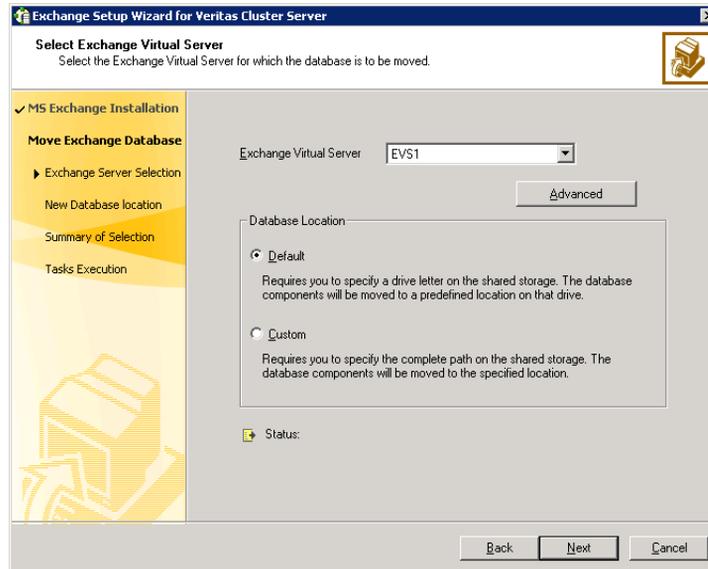
Moving Exchange databases to shared storage

After completing the Exchange installation on the first node, move the Exchange databases on the first node from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

- Make sure the data queue is empty on the SMTP server.
- Make sure to import the disk groups and mount the volumes for the Exchange database, MTA data, and transaction logs. Refer to [“Managing disk groups and volumes”](#) on page 217 for instructions.
- Start VEA and go to SYSTEM1. Select the storageagent and import the disk groups. Make sure the volumes have been assigned a drive letter.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.
- 5 In the Select Exchange Virtual Server dialog box:



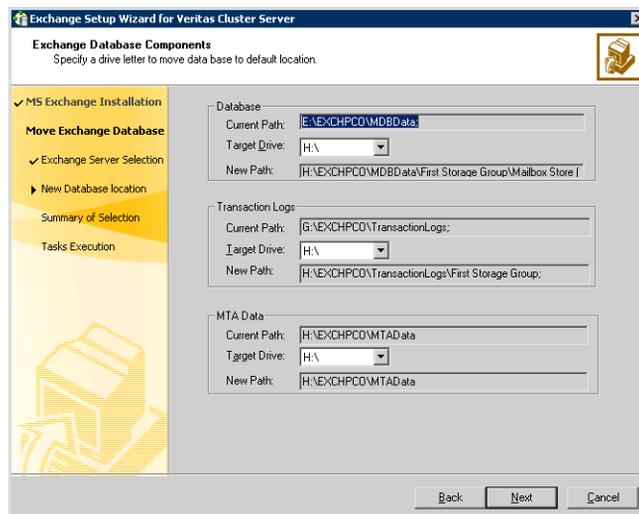
- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

Warning: The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Click **OK**.
 - Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
 - Click **Next**.
- 6 Do one of the following:
- If you would like to move the first mailbox store, public store, and MTA data to the generated default paths on the volumes for these components, proceed to the next step.
 - Otherwise, proceed to [step 8](#) on page 226 to specify the path location on the volumes that you will designate for these components.

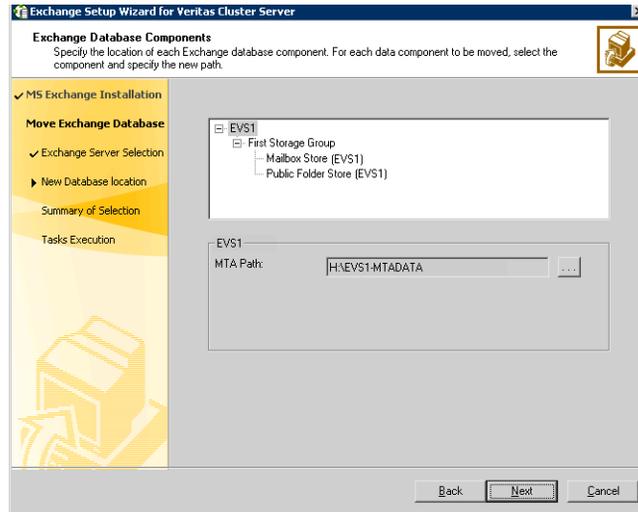
Warning: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

- 7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components will be moved to a default location on that drive. In the Exchange Database Components dialog box:



- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Specify the drive where the Exchange MTA Data will be moved.
- Click **Next** and proceed to [step 9](#) on page 226.

- 8 For the option of a custom database location, specify the location for specific Microsoft Exchange data components. In the Exchange Database Components dialog box:



For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the path for the Exchange database components contains only ANSI characters.

- 9 Review the summary of your selections and click **Next**.
- 10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.
- 11 Click **Finish** to exit the wizard.

Installing Exchange on additional nodes

After moving the Exchange databases to shared storage, install Exchange on all failover nodes in the cluster for the same Exchange virtual server (EVS1). You must run preinstallation, installation, and post-installation procedures for each failover node.

Note: Make sure to review the prerequisites for permissions in [“Installing Exchange on the first node”](#) on page 219.

Exchange pre-installation: additional nodes

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

Note: Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes. See [“Managing disk groups and volumes”](#) on page 217 for instructions.

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server. The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.

- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically generated when setting up an additional failover mode for the virtual server in the same cluster.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.
- 11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: additional nodes

Install Exchange on the node selected in “[Installing Exchange on additional nodes](#)” on page 226.

- Install any required service packs.

- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

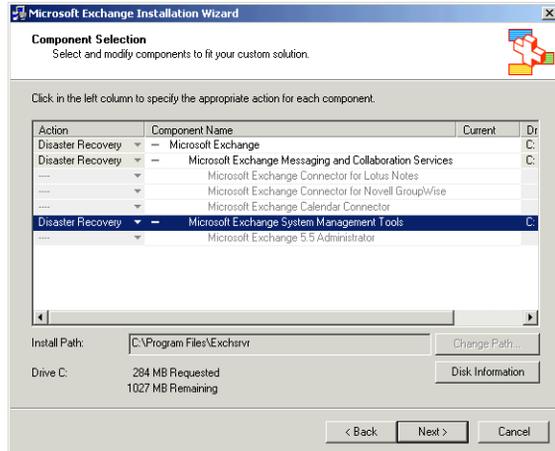
To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe
    /disasterrecovery
```

where <drive letter> is the location where the Exchange software is located.

- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

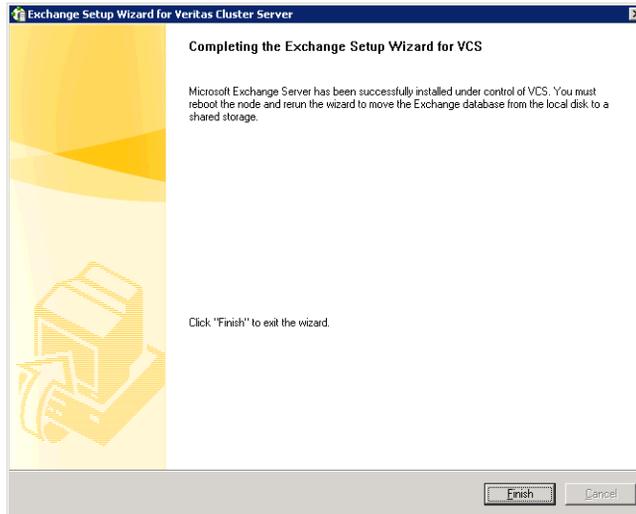
Exchange post-installation: additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.

- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

Configuring the Exchange service group for VCS

A new Exchange service group must be configured for the new Exchange virtual server, EVS1.

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

Prerequisites

- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk groups and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - Registry changes related to Exchange
 - Transaction logs for the first storage group
 - MTA databaseSee “[Managing disk groups and volumes](#)” on page 217 for instructions on mounting and unmounting.
- Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

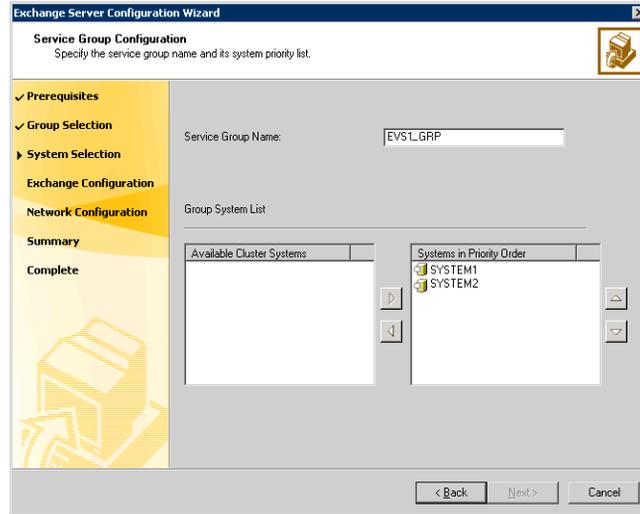
Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to an already configured service group.

To configure the Exchange service group

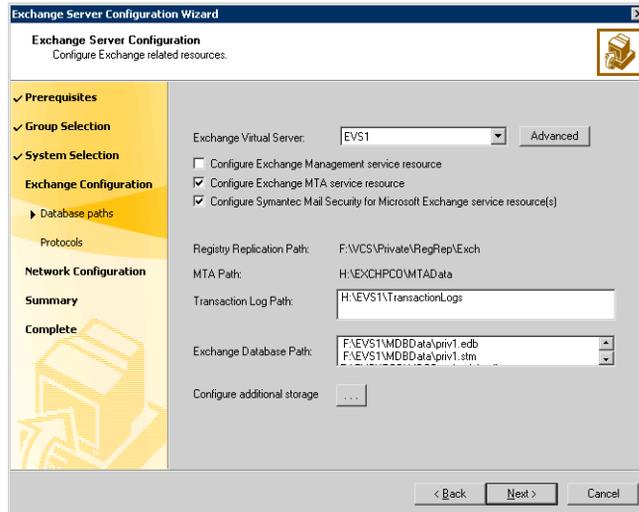
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options dialog box, choose the **Create service group** option and click **Next**.

4 Specify the service group name and system list:



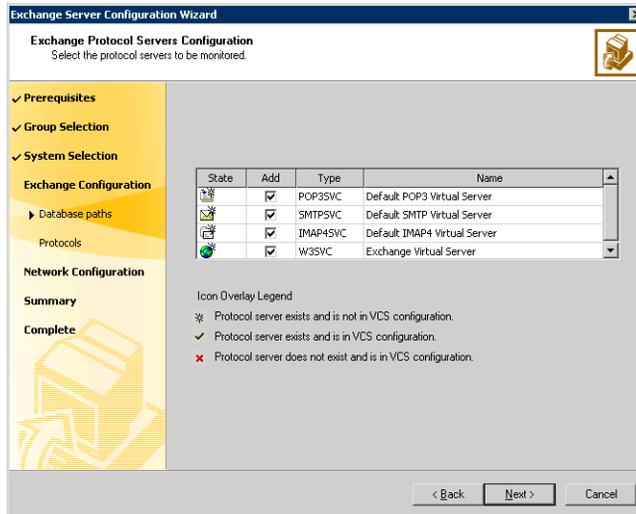
- Enter a name for the Exchange service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group’s system list. Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- To remove a system from the service group’s system list, select the system in the Systems in Priority Order list and click the left arrow.
- To change a node’s priority in the service group’s system list, select the node in the Systems in Priority Order list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:

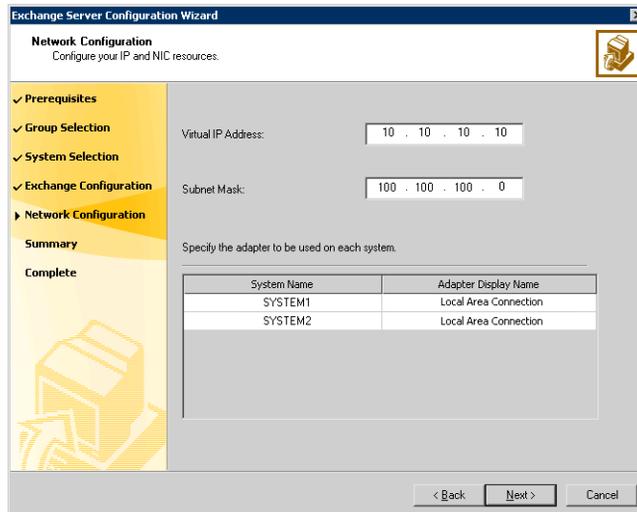


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.
- Specify whether you want to configure the Exchange Management service.
 - Specify whether you want to configure the Exchange MTA service resource. The Exchange Message Transfer Agent is specific to legacy Exchange message routing based on X.400. There are specific circumstances where it may or may not be required for your Exchange server. Please refer to Microsoft documentation on Exchange message routing and the use of MTA for further clarification and applicability to its use within your Exchange environment.

- If you are utilizing Symantec Mail Security for Exchange be sure to indicate that you would like to configure this resource.
 - Click **Next**.
- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



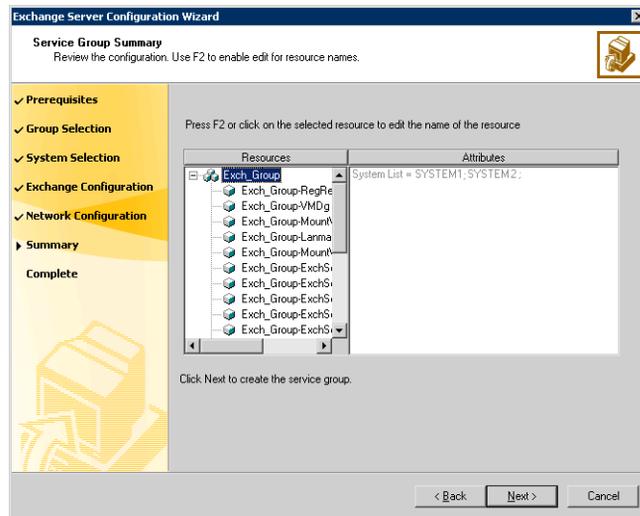
- 7 Specify information related to the network:



- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

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

- Click **Next**.
- 8 Review the service group configuration and change the resource names, if desired:



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 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.

- Click **Next**.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 9 In the Completing the Exchange Configuration dialog box, select the **Bring the service group online** check box to bring the service group online on the local system.
 - 10 Click **Finish**. After bringing the service group online, you must run the Exchange System Manager so that all the stores are automatically mounted on start-up.

To reconfigure mounting of stores at start-up

- 1 Start Exchange System Manager.
- 2 In the left pane, navigate to your storage group.
If administrative groups are not configured, Servers > Exchange Server > Storage Group.
If more than one administrative group is configured, expand Administrative Groups > Your Administrative Group > Servers > Exchange Server > Storage Group.
- 3 Right-click the Exchange database and choose **Properties** from the pop-up menu.
- 4 Click the Database tab.
- 5 Clear the **Do not mount this store at start-up** check box.
- 6 Click **OK**.

Repeat these steps for all the Exchange databases that were previously mounted.

If you need to configure additional storage groups or mailbox stores on the shared storage you should do that now. Import disk groups and mount volumes that have been created for the additional storage groups or mailbox stores, and create the new storage groups and mailbox stores in Exchange System Manager, and then rerun the Exchange Configuration Wizard to bring them under VCS control. If you already designated the additional mounted volumes and disk groups when you ran the configuration wizard the first time then you can just create the storage groups and mailbox stores in Exchange System Manager.

Verifying the cluster configuration

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

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

To switch service groups

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

To shut down an active cluster node

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

Configuring another Exchange virtual server for an any-to-any failover

Configure the next virtual server EVS2 on nodes 2 and 3.

See “[Reviewing the configuration](#)” on page 180.

Configuring disk groups and volumes

Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to grow existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

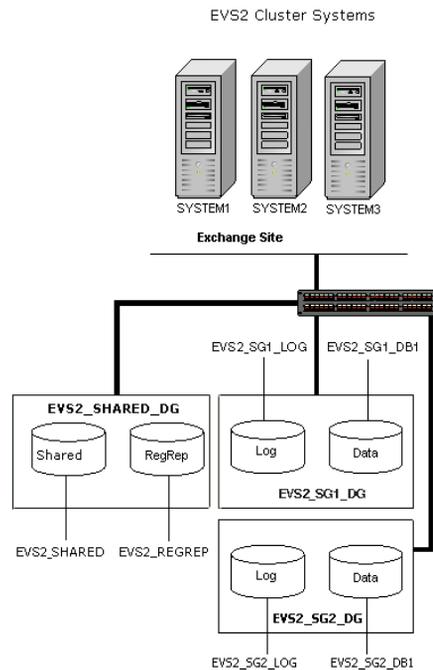
Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.

Typically, a SFW disk group corresponds to an Exchange storage group.

[Figure 6-3](#) shows a detailed view of the disk groups and volumes in an HA environment.

Figure 6-3 Disk groups and volumes for Exchange virtual server EVS2 in HA setup



Use the following procedures to create the appropriate disk groups and volumes. The general guidelines for disk group and volume setup for EVS2_SG2_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange disk group EVS2_SG2_DG create contains two volumes:

- EVS2_SG2_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.
- EVS2_SG2_LOG: Contains the transaction log for the storage group.

Exchange storage group EVS2_SHARED_DG create contains two volumes:

- EVS2_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.
- EVS2_SHARED: Contains the MTA database, SMTP, and message tracking.

Note: For additional Exchange storage groups, place the disks associated with the additional storage group's volumes in their own disk group.

For instructions on creating disk groups:

see “[Creating a disk group](#)” on page 212

For instructions on creating volumes:

see “[Creating volumes](#)” on page 213.

Managing disk groups and volumes

This section includes the following procedures:

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

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*

- Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Installing Exchange on the first node of an additional Exchange Virtual Server

Installing Exchange on the first node of EVS2 is described in three stages that involve preinstallation, installation, and post-installation procedures. Complete the following tasks before installing Exchange Server:

- Verify the disk group is imported on the first node of the cluster.
See “[Managing disk groups and volumes](#)” on page 217.
- Mount the volume containing the information for registry replication (EVS2_REGREP).
See “[Managing disk groups and volumes](#)” on page 217.
- Verify that all systems on which Exchange Server will be installed have IIS installed; you must install SMTP, NNTP, and WWW services on all systems. If you install Exchange on Windows 2003, make sure to install ASP.NET service.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.

- Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.
 - Make sure the HAD Helper domain user account has “Add workstations to domain” privilege enabled in the Active Directory. To verify this, click **Start > Administrative Tools > Local Security Policy** on the domain controller to launch the security policy display. Click **Local Policies > User Rights Management** and make sure the user account has this privilege.
 - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
 - The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

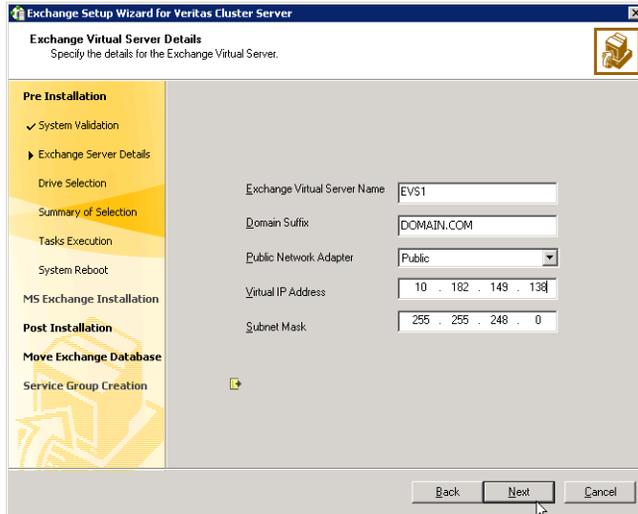
Exchange pre-installation: first node of an additional Exchange Virtual Server

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. You must install Exchange on a virtual node to facilitate high availability.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a new Exchange Virtual Server** option and click **Next**.

- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.
- 7 Specify information related to the network.



- Enter a unique virtual name for the Exchange server.
Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.
- Enter a domain suffix for the virtual server.
- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
- Enter a unique virtual IP address for the Exchange virtual server.
- Enter the subnet mask for the virtual IP address.
- Click **Next**.
The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met. Click **Next**.

- 8 Select a drive where the registry replication data will be stored and click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: first node of an additional Exchange Virtual Server

Install Exchange on the node selected in “[Exchange pre-installation: first node of an additional Exchange Virtual Server](#)” on page 243.

Exchange 2000 requires service pack 3 with the August 2004 rollup patch. Exchange 2003 requires service pack 2. The procedure below is based on Exchange 2003 SP2.

This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Install Exchange Server using the Microsoft Exchange installation program. See the Microsoft Exchange documentation for instructions.
- 2 Reboot the node if prompted to do so.

- 3 For Exchange 2000 or Exchange 2003, install the required service pack.

Exchange post-installation: first node of an additional Exchange Virtual Server

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

Moving Exchange databases to shared storage

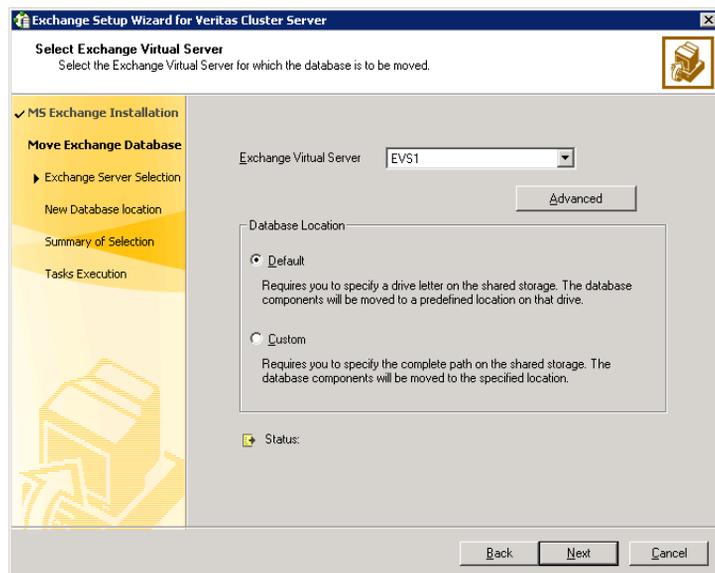
After completing the Exchange installation on the first node, move the Exchange databases on the first node from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

- Make sure the data queue is empty on the SMTP server.
- Make sure to import the disk group and mount the volumes for the Exchange database, MTA data, and transaction logs.
See "[Managing disk groups and volumes](#)" on page 217.

- Start VEA and go to SYSTEM1. Select the storageagent and import the disk groups. Make sure the volumes have been assigned a drive letter.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.
- 5 In the Select Exchange Virtual Server dialog box:



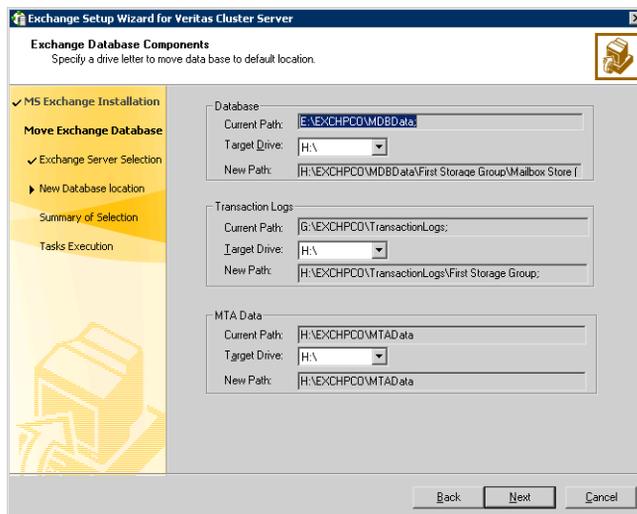
- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

Warning: The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

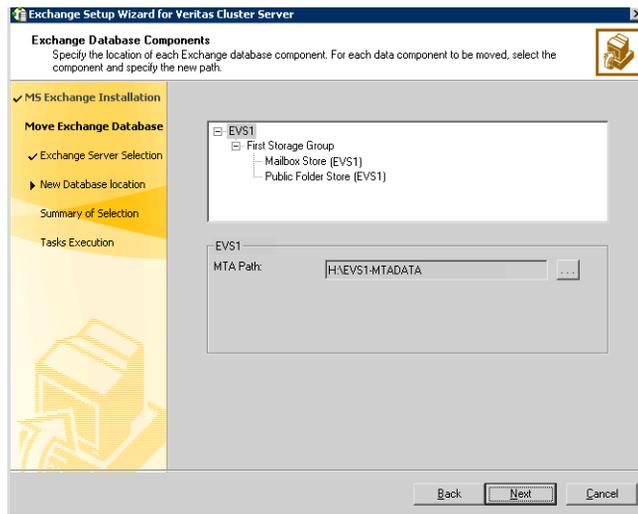
- Click **OK**.
 - Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
 - Click **Next**.
- 6 Do one of the following:
- If you would like to move the first mailbox store, public store, and MTA data to the generated default paths on the volumes for these components, proceed to the next step.
 - Otherwise, proceed to [step 8](#) on page 226 to specify the path location on the volumes that you will designate for these components.

Warning: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

- 7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components will be moved to a default location on that drive. In the Exchange Database Components dialog box:



- Specify the drive where the Exchange database will be moved.
 - Specify the drive where the Exchange Transaction Logs will be moved.
 - Specify the drive where the Exchange MTA Data will be moved.
 - Click **Next** and proceed to [step 9](#) on page 226.
- 8 For the option of a custom database location, specify the location for specific Microsoft Exchange data components. In the Exchange Database Components dialog box:



For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the path for the Exchange database components contains only ANSI characters.

- 9 Review the summary of your selections and click **Next**.
- 10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.
- 11 Click **Finish** to exit the wizard.

Specifying a common node for failover

Specifying a common node for failover involves preparing the cluster with the Exchange Setup Wizard for VCS.

The failover node for the first Exchange virtual server, EVS1, was specified when the EVS1 service group was created. After the designated Exchange virtual servers have been installed in the cluster, launch the Exchange Setup Wizard with the any-to-any option from any system in the cluster. Repeat the task below for each additional Exchange Virtual Server.

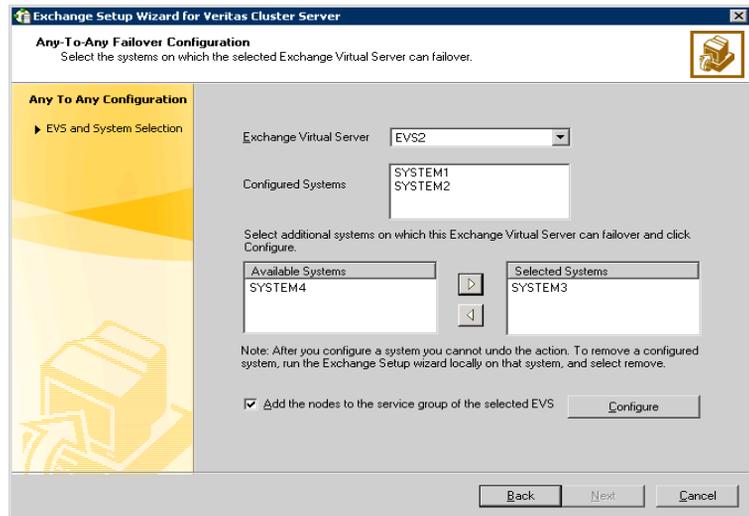
Note: The Exchange software was installed on the common failover node during the installation process for the first EVS. You do not install Exchange a second time on the common failover node.

To prepare the cluster with the any-to-any option

In our example EVS1 is already configured with SYSTEM3 as a failover node. Execute this wizard on EVS2 only

- 1 Start the Exchange Setup Wizard for VCS from any node configured to host an Exchange service group. Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard**.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Options dialog box, choose the **Configure any-to-any failover** option and click **Next**.

- 5 Select systems to be configured for any-to-any failover. The **Configured Systems** box lists the nodes on which the Exchange Server service group can fail over. Do the following in order:



- Select the Exchange virtual server to which you want to add the additional failover nodes.
 - The Configured Systems box displays the nodes on which the Exchange Server has been installed.
 - From the **Available Systems** box, select the systems to be configured for any-to-any failover.
 - The **Available Systems** box lists only those systems that have the same version and service pack level of Microsoft Exchange as the selected Exchange virtual server.
 - Click the right arrow to move the selected systems to the **Selected Systems** box. To remove a system from the box, select the system and click the left arrow.
 - Select Add the nodes to the service group of the selected EVS to add the selected systems to the SystemList of the service group for the selected Exchange virtual server.
 - Click **Configure**.
 - Click **Next**.
- 6 Click **Finish**.

Configuring the Exchange service group for an additional Exchange Virtual Server

A new Exchange service group must be configured for the new Exchange virtual server, EVS2.

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

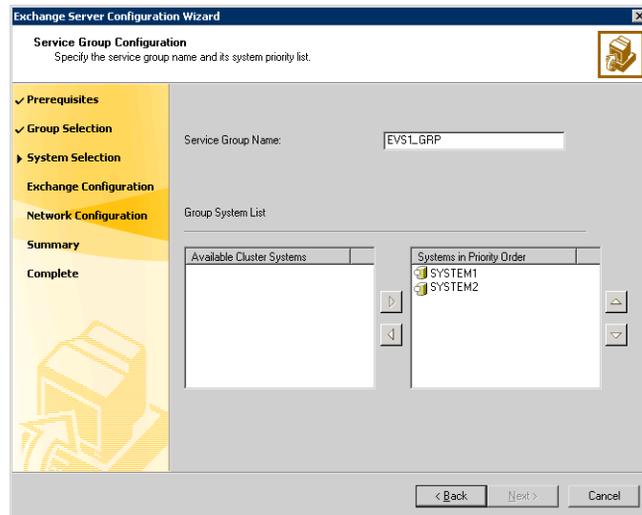
Prerequisites

- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk groups and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - registry changes related to Exchange
 - transaction logs for the first storage group
 - MTA databaseSee “[Managing disk groups and volumes](#)” on page 217.
- Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to an already configured service group.

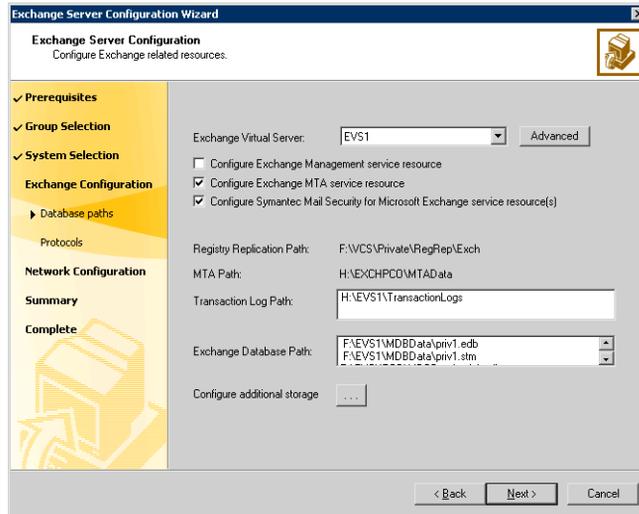
To configure the Exchange service group

- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options dialog box, choose the **Create service group** option and click **Next**.
- 4 Specify the service group name and system list:



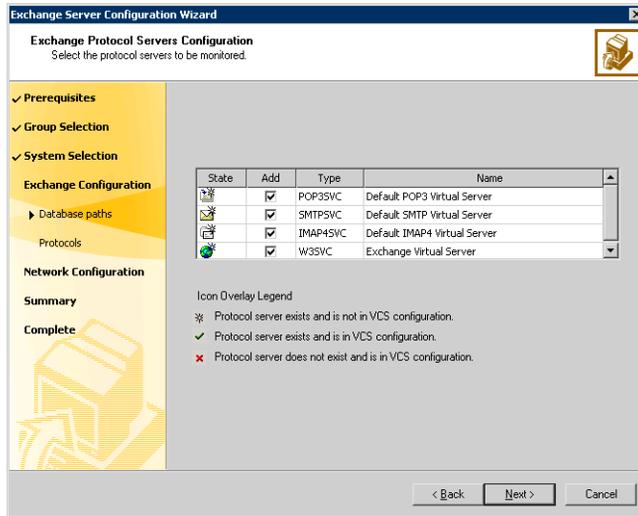
- Enter a name for the Exchange service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group’s system list. Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- To remove a system from the service group’s system list, select the system in the Systems in Priority Order list and click the left arrow.
- To change a node’s priority in the service group’s system list, select the node in the Systems in Priority Order list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:

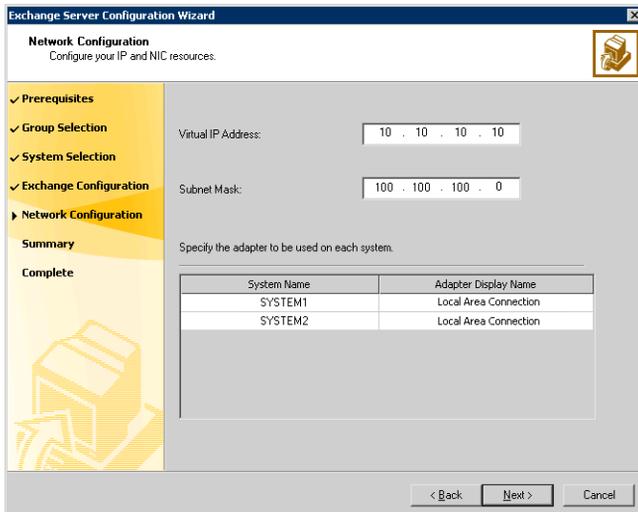


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.
- Specify whether you want to configure the Exchange Management service.
 - Specify whether you want to configure the Exchange MTA service resource. The Exchange Message Transfer Agent is specific to legacy Exchange message routing based on X.400. There are specific circumstances where it may or may not be required for your Exchange server. Please refer to Microsoft documentation on Exchange message routing and the use of MTA for further clarification and applicability to its use within your Exchange environment.

- If you are utilizing Symantec Mail Security for Exchange be sure to indicate that you would like to configure this resource.
 - Click **Next**.
- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



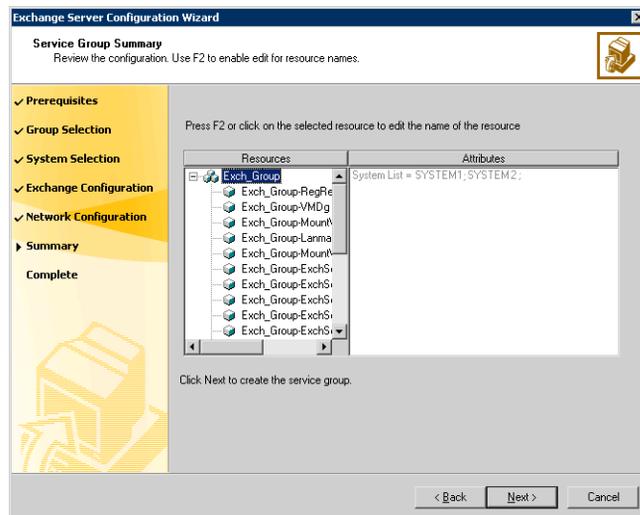
- 7 Specify information related to the network:



- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

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

- Click **Next**.
- 8 Review the service group configuration and change the resource names, if desired:



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 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.

- Click **Next**.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 9 In the Completing the Exchange Configuration dialog box, select the **Bring the service group online** check box to bring the service group online on the local system.
 - 10 Click **Finish**. After bringing the service group online, you must run the Exchange System Manager so that all the stores are automatically mounted on start-up.

To reconfigure mounting of stores at start-up

- 1 Start Exchange System Manager.
- 2 In the left pane, navigate to your storage group.
 If administrative groups are not configured, Servers > Exchange Server > Storage Group.
 If more than one administrative group is configured, expand Administrative Groups > Your Administrative Group > Servers > Exchange Server > Storage Group.
- 3 Right-click the Exchange database and choose **Properties** from the pop-up menu.
- 4 Click the Database tab.
- 5 Clear the **Do not mount this store at start-up** check box.
- 6 Click **OK**.

Repeat these steps for all the Exchange databases that were previously mounted.

If you need to configure additional storage groups or mailbox stores on the shared storage you should do that now. Import disk groups and mount volumes that have been created for the additional storage groups or mailbox stores, and create the new storage groups and mailbox stores in Exchange System Manager, and then rerun the Exchange Configuration Wizard to bring them under VCS control. If you already designated the additional mounted volumes and disk groups when you ran the configuration wizard the first time then you can just create the storage groups and mailbox stores in Exchange System Manager.

Verifying the cluster configuration

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

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

To switch service groups

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

To shut down an active cluster node

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

Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover

This chapter covers the following topics:

- [Reviewing the requirements](#)
- [Configuring new nodes: Prior to creating additional Exchange virtual server](#)
- [Specifying a common node for failover](#)
- [Verifying the cluster configuration](#)

You can install and configure an “any-to-any” SFW HA environment for Exchange to provide a production node with multiple failover nodes by transforming an active/passive SFW HA environment for Exchange, which involves one-to-one failover capabilities, into an any-to-any environment. The

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

Table 7-1 Task List

Objective	Tasks
“ Reviewing the requirements ” on page 262	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 266	<ul style="list-style-type: none"> ■ Understanding a basic any-to-any Exchange configuration, starting with an existing active / passive cluster and adding nodes.
“ Configuring the network and storage ” on page 269	<ul style="list-style-type: none"> ■ Configuring the network and storage for the new systems if this system is not already part of the existing active/passive cluster.
“ Installing Veritas Storage Foundation HA for Windows ” on page 270	<ul style="list-style-type: none"> ■ Running the installation only on the new systems, if these systems are not already part of the existing active/passive cluster.
“ Configuring the cluster ” on page 276	<ul style="list-style-type: none"> ■ Adding the new nodes to the cluster if there are not enough nodes in the existing active/passive cluster to facilitate an any-to-any configuration. ■ If new nodes need added, completing the steps in “Adding a node to a cluster” on page 276.
“ Configuring disk groups and volumes ” on page 286	<ul style="list-style-type: none"> ■ Creating the disk groups and volumes for any new nodes. These disk groups and volumes must be on a separate disk from the volumes for the existing active / passive cluster. ■ Using the VEA console to create disk groups ■ Using the VEA console to create the Data, Log, RegRep, and SHARED volumes ■ Managing disk groups and volumes, with instructions for mounting and unmounting volumes

Table 7-1 Task List

Objective	Tasks
“Installing Exchange on the new nodes” on page 295	<ul style="list-style-type: none"> ■ Installing Exchange on the new active Exchange node, a “first node.”
“Moving Exchange databases to shared storage (EVS2)” on page 299	<ul style="list-style-type: none"> ■ Moving the databases from the new active Exchange node to shared storage.
“Installing Exchange on additional nodes” on page 303	<ul style="list-style-type: none"> ■ Adding additional new failover nodes, if any.
“Specifying a common node for failover” on page 308	<ul style="list-style-type: none"> ■ Preparing the cluster for any-to-any failover using the Exchange Setup Wizard. This step must be completed on each of the Exchange Virtual Servers. ■ Configuring the Exchange service group for the second Exchange Virtual Server. If necessary, you can later add common failover nodes to the Exchange service group’s system list. ■ The nodes currently in your active / passive cluster and members of the first Exchange Virtual Server (EVS1 in the example) need only the final task “Specifying a common node for failover” on page 308.
“Verifying the cluster configuration” on page 316	<ul style="list-style-type: none"> ■ Verifying the cluster configuration by switching service groups and shutting down an active cluster node.

Reviewing the requirements

Refer to “[Reviewing the configuration](#)” on page 266 for an overview of an any-to-any configuration.

To create an any-to-any cluster, refer to the prerequisites below.

- To transform an active/passive cluster to an any-to-any cluster, you must already have one active/passive cluster.
 - For a new Exchange server, refer to the chapter “[Deploying SFW HA for high availability: New installation](#)” on page 33 to create one active / passive cluster.
 - For an existing standalone Exchange server, refer to the chapter “[Deploying SFW HA for high availability: Standalone Exchange servers](#)” on page 103 to create one active / passive cluster.
- Two or more Exchange virtual servers can exist in an any-to-any configuration.

Use the procedures in this chapter to create a second Exchange server, if two Exchange servers are not already present in the configuration, and to add a second virtual server to an existing active/passive cluster.

Review the following 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 7-2](#) estimates disk space requirements for SFW HA.

Table 7-2 Disk space requirements

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

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

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

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

Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA) with the Veritas Cluster Server Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:
 - Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition (SP2 required)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
or
Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP1 required for all editions)
or
Windows Server 2003 R2 (Standard Edition, Enterprise Edition, or Datacenter Edition)
or
 - Microsoft Exchange 2000 Server or Microsoft Exchange 2000 Enterprise Server (SP3 with August 2004 rollup patch required for both editions)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)

Note: Microsoft support for Exchange Server 2003 is limited to 32-bit versions of the Windows 2003 operating system.

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 266.
- 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 Exchange Virtual Server (EVS)
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

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

Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

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

- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Domain Servers group.
- 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.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- You must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.
- 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.

- You must install IIS, SMTP, NNTP, ASP.NET, and WWW services before you install Microsoft Exchange Server.

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

The following table details the systems in an active / passive configuration (SYSTEM1 and SYSTEM2) plus a new node (SYSTEM3) transformed into an any-to-any configuration:

Table 7-3 Existing Active / Passive Configuration to Any-to-Any Cluster

Exchange Virtual Server	Active/Passive	Any-to-Any Common Failover Node
EVS1	SYSTEM1, SYSTEM3	SYSTEM1, SYSTEM3
EVS2	SYSTEM2	SYSTEM2, SYSTEM3

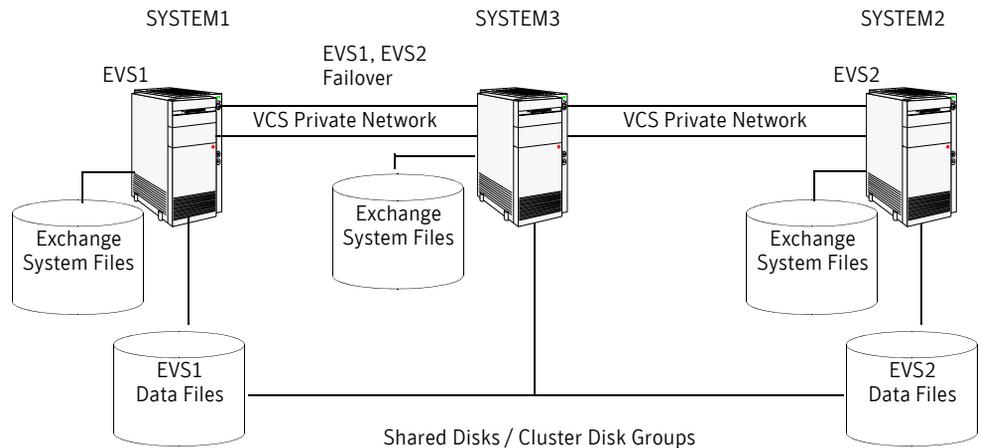
In an active / passive configuration, a separate failover system is required for each active Exchange node. In an any-to-any configuration, the active Exchange nodes can share failover nodes.

Additional failover nodes can also exist in an any-to-any configuration.

Any-to-any configuration

In an any-to-any configuration, each Exchange virtual server in the cluster is configured in a separate service group. Each service group can fail over to any configured node in the cluster, provided that no other Exchange virtual server is online on that node. You must ensure that an Exchange service group does not fail over to a node on which another Exchange service group is online. [Figure 7-1](#) shows an example of a three-node cluster.

Figure 7-1 Three-node cluster in an any-to-any configuration



For example, consider a three-node cluster hosting two Exchange Virtual Servers, EVS1 and EVS2. The virtual servers are configured in VCS in two service groups such that SYSTEM1 hosts the EVS1 service group and SYSTEM2 hosts the EVS2 service group. If SYSTEM1 fails, the service group containing the EVS1 resources is failed over to SYSTEM3. If SYSTEM2 fails, the service group EVS2 fails over to SYSTEM3.

Note: EVS1 and EVS2 cannot be online at the same time on SYSTEM2.

Sample configuration

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

Table 7-4 Sample Configuration

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3	physical node names
EVS1, EVS2	Microsoft Exchange Virtual Servers
EVS1_GRP, EVS2_GRP	Microsoft Exchange service groups
EVS1_SG1_DG, EVS2_SG1_DG	cluster disk group names
EVS1_SG1_DB1, EVS2_SG1_DB1	volumes for storing the Microsoft Exchange Server database
EVS1_SG1_LOG, EVS2_SG1_LOG	volumes for storing a Microsoft Exchange Server database log file
EVS1_REGREP, EVS2_REGREP	volumes that contain the list of registry keys that must be replicated among cluster systems for the Exchange server
EVS1_SHARED, EVS2_SHARED	volumes for storing Microsoft Exchange Server MTA database, SMTP and message tracking for Exchange server

Configuring new nodes: Prior to creating additional Exchange virtual server

In the example, SYSTEM3 is a new node to be added to an existing active / passive cluster consisting of SYSTEM1 and SYSTEM2. In addition, SYSTEM3 will become a second Exchange server in the configuration.

Your final configuration will consist of two independent Exchange Servers (called active or “first” nodes in the tasks below) and one or more failover nodes (called “additional nodes” in the tasks below).

Configuring the network and storage

Configure the network and storage for the new systems (SYSTEM3 in the example) if this system is not already part of the existing active/passive cluster (in the example, SYSTEM1 and SYSTEM2).

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

Run the installation only on the new systems (SYSTEM3 in the example) if this system is not already part of the existing active/passive cluster (SYSTEM1 and SYSTEM2). Also include any other new nodes in this installation.

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 Application Agent for Exchange. For a disaster recovery configuration, select the option to install GCO. If you plan to use VVR for replication, you must also select the option to install VVR.

Setting Windows driver signing options

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

Table 7-5 describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 7-5 Installation behavior with unsigned drivers

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

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

To change the driver signing options on each system

- 1 Log on locally to the system.
 - 2 Open the Control Panel and click **System**.
 - 3 Click the **Hardware** tab and click **Driver Signing**.
 - 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
 - 5 Click **OK**.
 - 6 Repeat for each computer.
- If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

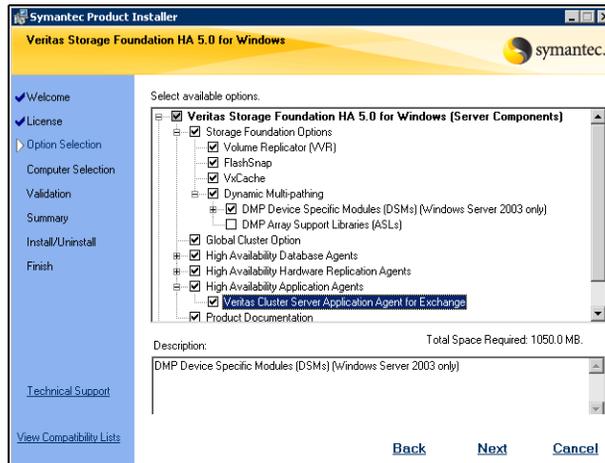
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.

- 2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.
- 3 Click **Storage Foundation HA 5.0 for Windows**.



- 4 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 8 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

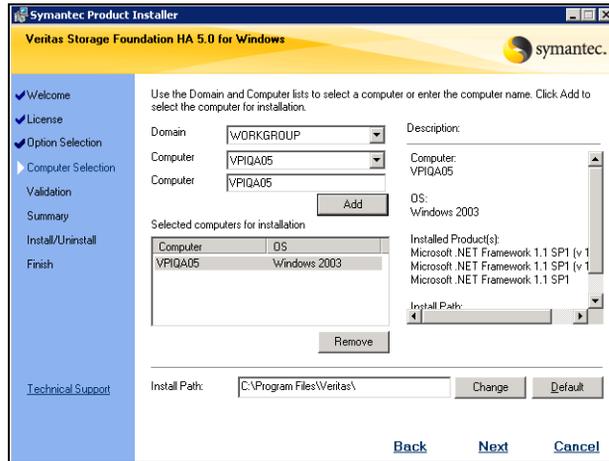
9 Select the appropriate SFW product options for your installation. Click **Next**.



The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

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

10 Select the domain and the computers for the installation and click **Next**.



- Domain** Select a domain from the list.
Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.
- Computer** To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.
To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.
Click a computer's name to see its description.
- Install Path** Optionally, change the installation path.
- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
 - To restore the default path, select a computer and click **Default**.
The default path is:
C:\Program Files\Veritas
For 64-bit installations, the default path is:
C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
 - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
 - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

Add the new nodes (SYSTEM3 in the example) to the cluster if you do not already have enough nodes in the existing active/passive cluster to facilitate an any-to-any configuration.

If you need to add nodes to the cluster, complete the steps in “[Adding a node to a cluster](#)” on page 276.

Adding a node to a cluster

This procedure only applies to any-to-any clusters transformed from active/passive clusters that do not already have all the nodes for the any-to-any configuration in the cluster.

This section includes optional instructions to configure the ClusterService group for the VCS Cluster Management Console (Single Cluster Mode) also referred to as Web Console, or notification after adding a node to the cluster.

To add a node to a cluster

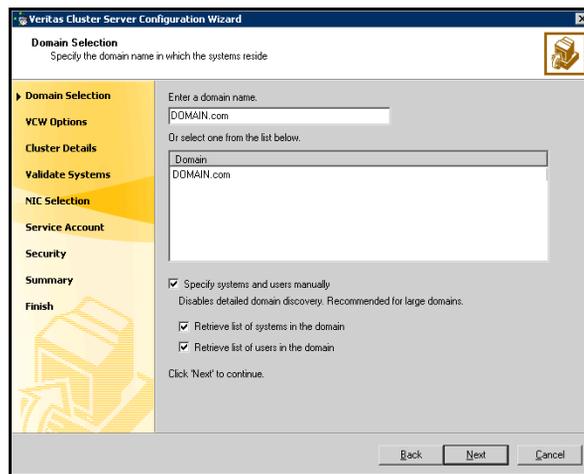
Note: Run the VCS Configuration Wizard from the standalone node or a node in the cluster.

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 280.

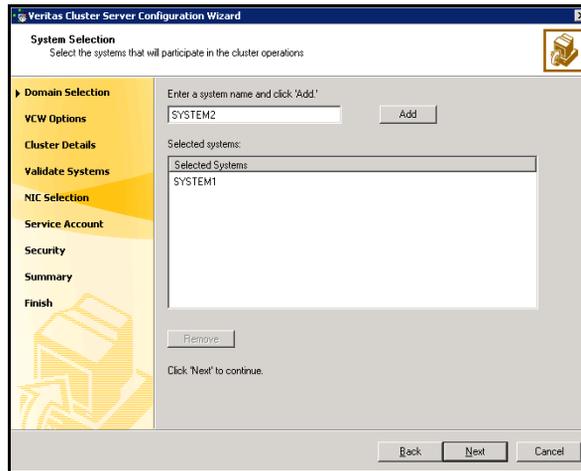
To specify systems and user names manually (recommended for large domains):

- 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**.

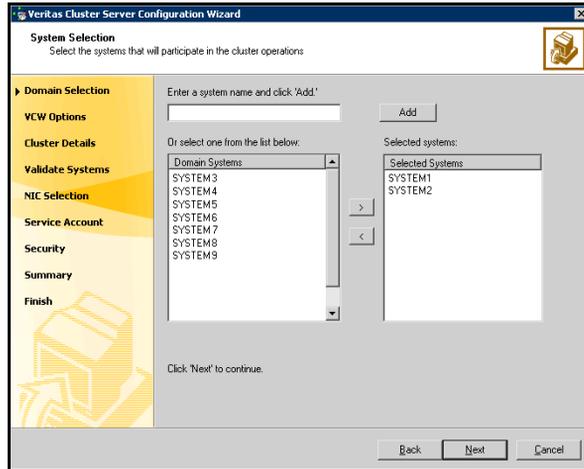
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 279. Otherwise proceed to the next step.

- 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 280.

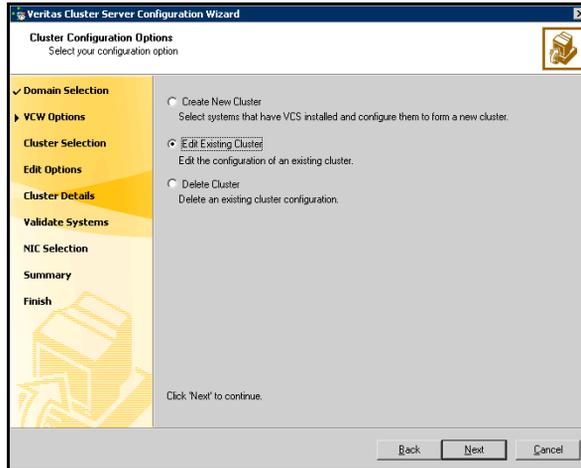
- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems.



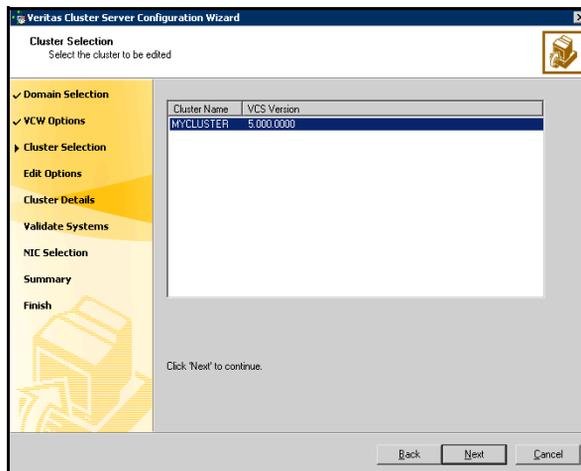
Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

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**.

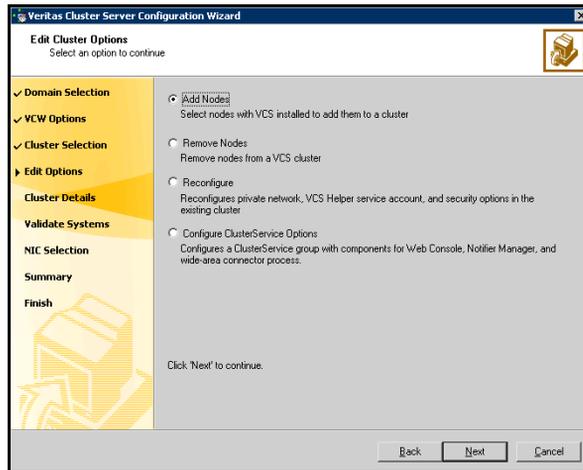


- 8 On the Cluster Selection panel, select the cluster to be edited and click **Next**.



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

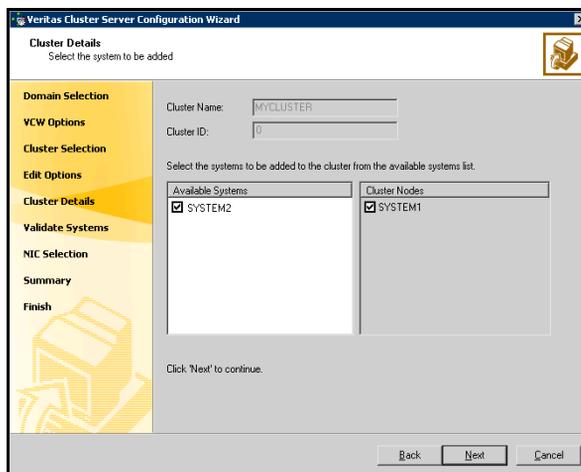
- 9 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**.



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**.

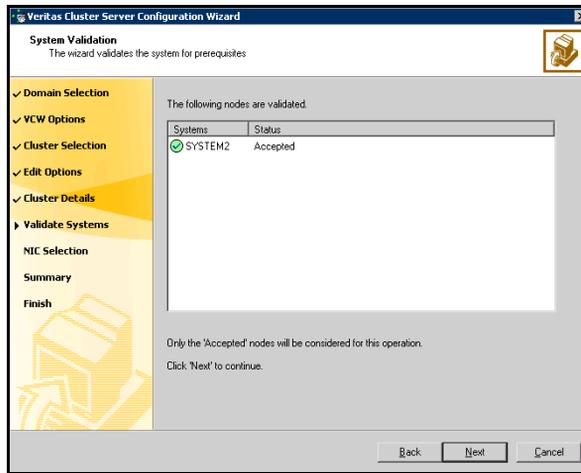
The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges (a cluster that is not a secure cluster).

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



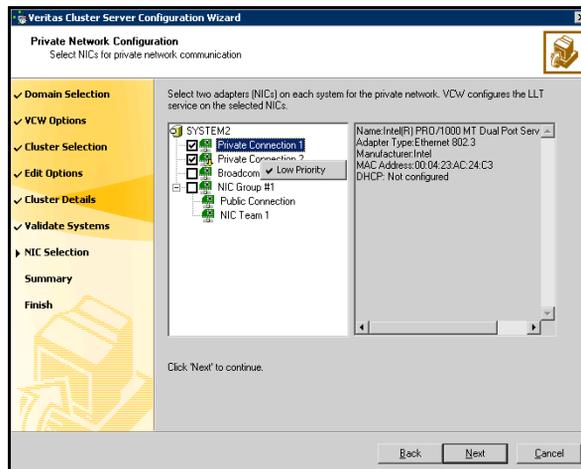
The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

- 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.

If you need the VCS Web Console or notification for the cluster, proceed to [“Configuring the ClusterService group for VCS”](#) on page 284.

If you do not need to configure the VCS Web Console and notification, return to the task list in [“Configuring the network and storage”](#) on page 269.

Configuring the ClusterService group for VCS

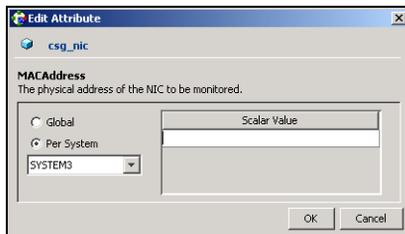
If you plan on setting up the VCS Cluster Management Console (Single Cluster Mode) also referred to as Web Console, or notification for the cluster, you must manually alter the resources for the ClusterService service group.

- Use the VCS Java Console to configure the NIC, IP, and VCSweb resources in the ClusterService group for the VCS Web Console.
- Use the VCS Java Console to configure the NIC resource in the ClusterService group for notification.

Note: Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on using the VCS Java Console and configuring the VCS Web Console and Notifier resource.

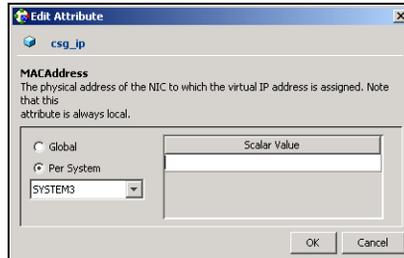
To configure the ClusterService group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 From the **Service Groups** tab of the Cluster Explorer configuration tree, expand the NIC resource type and select the **csg_nic** resource.
- 3 In the **Properties** tab of the view panel, click the **Edit** icon for the **MACAddress** attribute.
- 4 In the Edit Attribute dialog box:

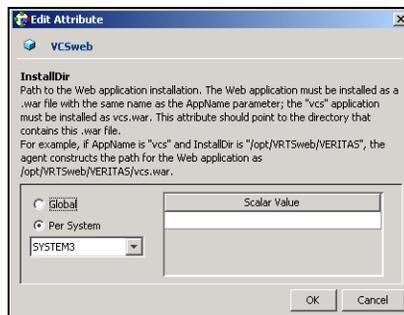


- Select the **per system** option, and select the newly added system.
 - Enter the scalar value. To obtain the MAC address, run the `ipconfig /all` command from the command prompt on that system.
 - Click **OK**.
- If you are only configuring notification, return to the task list in [“Configuring the network and storage”](#) on page 269. If you are configuring the Web Console, proceed to step 5.
- 5 From the **Service Groups** tab of the Cluster Explorer configuration tree, expand the IP resource type and click the **csg_ip** resource.

- 6 In the **Properties** tab of the view panel, click the **Edit** icon for the **MACAddress** attribute.
- 7 In the Edit Attribute dialog box:



- a Select the **per system** option, and select the newly added system.
 - b Enter the scalar value. To obtain the MAC address, run the `ipconfig / all` command from the command prompt on that system.
 - c Click **OK**.
- 8 From the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **VRTSWebApp** resource type and select the **VCSweb** resource.
 - 9 In the **Properties** tab of the view panel, click the **Edit** icon for the **InstallDir** attribute.
 - 10 In the Edit Attribute dialog box:



- a Select the **per system** option, and select the system.
- b Enter the scalar value. From the command prompt on that system, type the following command to obtain the value:
`C: \>set VCS_ROOT`
Attach "`\VRTSweb\Veritas`" to the end of the generated value to determine the scalar value.

c Click **OK**.

11 On the **File** menu of Cluster Explorer, click **Save Configuration**.

After configuring the ClusterService group, return to the task list in “[Configuring new nodes: Prior to creating additional Exchange virtual server](#)” on page 268.

Configuring disk groups and volumes

If you added a new node (SYSTEM3) to an existing active/passive cluster, create the disk groups and volumes for this node.

Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to grow existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

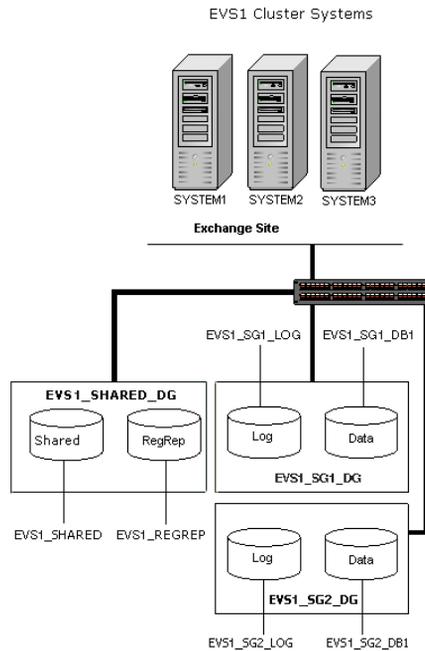
Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.

Typically, a SFW disk group corresponds to an Exchange storage group.

[Figure 7-2](#) shows a detailed view of the disk groups and volumes in an HA environment.

Figure 7-2 Disk groups and volumes for Exchange virtual server EVS1 in HA setup



Use the following procedures to create the appropriate disk groups and volumes. The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange storage group EVS1_SG1_DG create contains two volumes:

- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume. This contains the EVS1_SG1_LOG volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

Exchange storage group EVS1_SHARED_DG create contains two volumes:

- EVS1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.
- EVS1_SHARED: Contains the MTA database, SMTP, and message tracking.

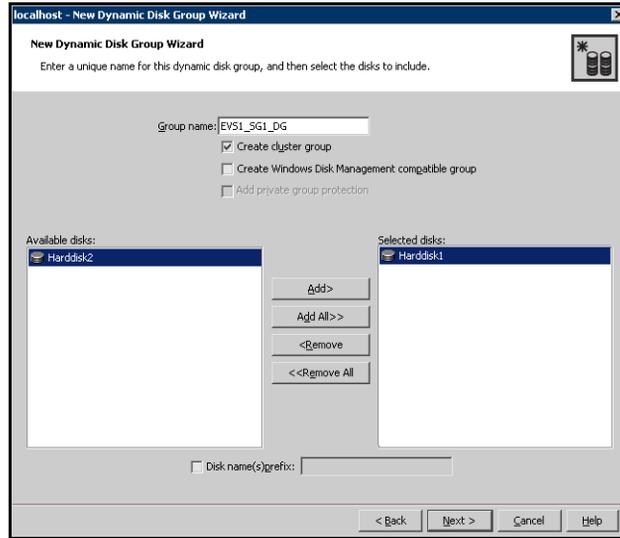
Note: For additional Exchange storage groups, place the disks associated with the additional storage group's volumes in their own disk group.

Creating a disk group

To create a dynamic (cluster) disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group:



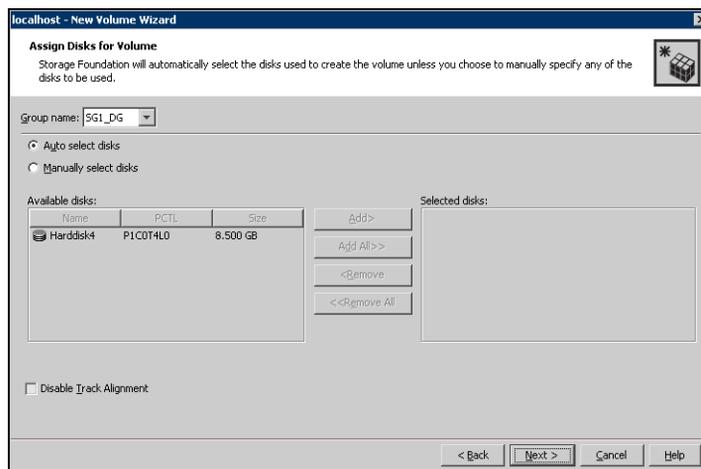
- Enter the name of the disk group (for example, EVS1_SG1_DG).
 - Click the checkbox for **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox 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 assumes you are starting with the EVS1_SG1_DB1 volume.

To create dynamic volumes

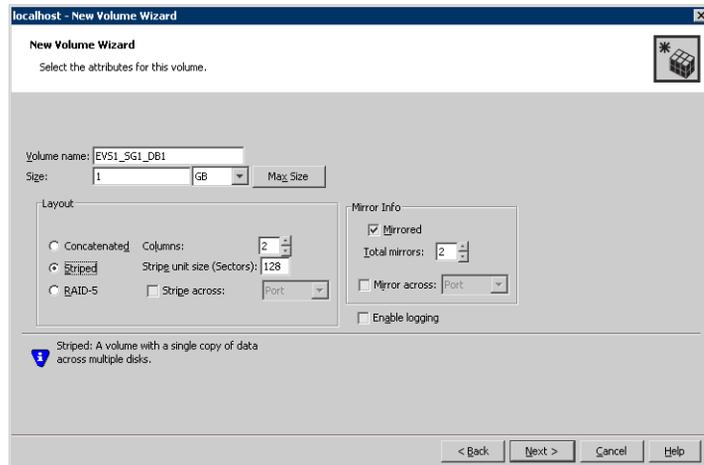
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu 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 EVS1_SG1_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.

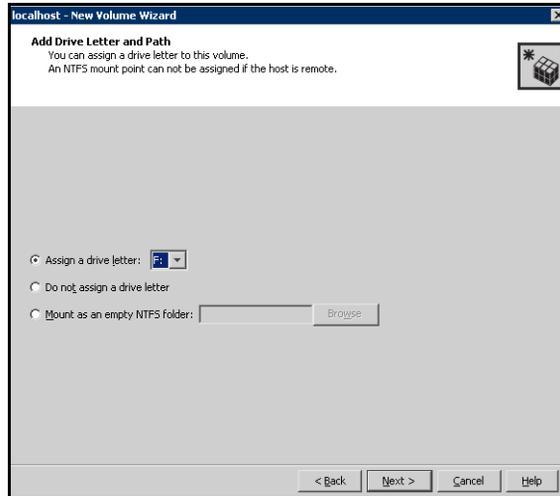
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling **Track Alignment** means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

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



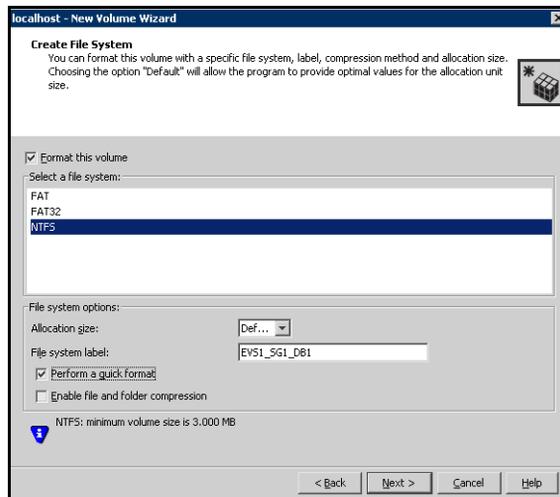
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
 Do not assign the M: drive letter if you are using Exchange 2000, as it is reserved for internal use.

- 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 the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) Also, repeat these steps to create both the RegRep volume (EVS1_REGREP) and the EVS1_SHARED volumes in the EVS1_SHARED_DG disk group.

15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3). Create the cluster disk group and volumes on the first node of the cluster only.

Create similar disk groups and volumes for other Exchange servers. For example:

- Create disk group (EVS2_SG1_DG).
- Create volumes (EVS2_SG1_DB1, EVS2_REGREP, EVS2_SG1_LOG, and EVS2_SHARED).

Managing disk groups and volumes

This section includes the following procedures:

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

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.

- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.

- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Installing Exchange on the new nodes

An any-to-any configuration requires a minimum of two Exchange virtual servers; you must create an Exchange virtual server (EVS2) for SYSTEM3 that will ultimately include the common failover node (SYSTEM2).

Install Exchange on the new active Exchange node, a “first node.” In the example the new active Exchange node for EVS2 is SYSTEM3.

For an existing active/passive cluster (SYSTEM1 and SYSTEM2) run the installation only on the new system (SYSTEM3). This system becomes an active Exchange server, or “first node” for the second Exchange Virtual Server (EVS2).

- Verify the disk group is imported on the first node of the cluster. Refer to [“Importing a disk group and mounting a shared volume”](#) on page 293 for instructions.
- Mount the volume containing the information for registry replication (EVS1_SG1_REGREP). Refer to [“Importing a disk group and mounting a shared volume”](#) on page 293 for instructions.
- Verify that all systems on which Exchange Server will be installed have IIS installed; you must install SMTP, NNTP, and WWW services on all systems. If you install Exchange on Windows 2003, make sure to install ASP.NET service.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.

- Make sure the HAD Helper domain user account has “Add workstations to domain” privilege enabled in the Active Directory. To verify this, click **Start > Administrative Tools > Local Security Policy** on the domain controller to launch the security policy display. Click **Local Policies > User Rights Management** and make sure the user account has this privilege.
- If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

Exchange pre-installation: First node

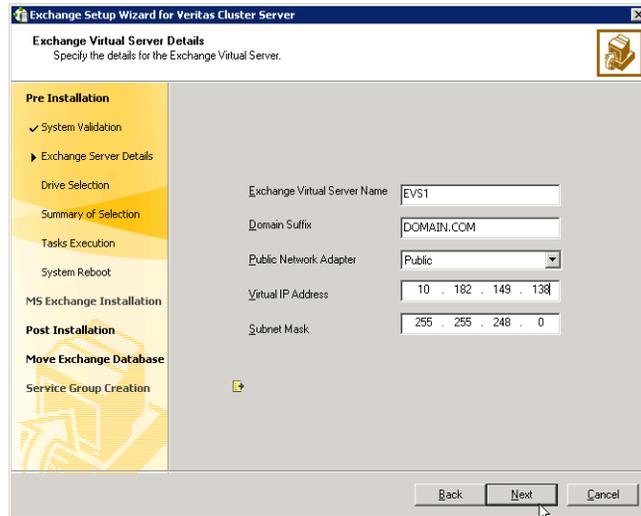
Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. You must install Exchange on a virtual node to facilitate high availability.

For this “first node” installation for the second Exchange virtual server, EVS2, the mount information for the registry replication information (in step 1 below) refers only to the second Exchange virtual server. The current active / passive cluster, for the first Exchange virtual server, EVS1, can continue to operate normally during this procedure.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a new Exchange Virtual Server** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.

7 Specify information related to the network.



- Enter a unique virtual name for the Exchange server.
Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.
 - Enter a domain suffix for the virtual server.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Enter a unique virtual IP address for the Exchange virtual server.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met. Click **Next**.
- 8 Select a drive where the registry replication data will be stored and click **Next**.

- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: First node

Install Exchange on the same node selected in “[Exchange pre-installation: First node](#)” on page 296.

Exchange 2000 requires service pack 3 with the August 2004 rollup patch. Exchange 2003 requires service pack 2. The procedure below is based on Exchange 2003 SP2.

This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Install Exchange Server using the Microsoft Exchange installation program. See the Microsoft Exchange documentation for instructions.
- 2 Reboot the node if prompted to do so.
- 3 For Exchange 2000 or Exchange 2003, install the required service pack.

Exchange post-installation: First node

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

Moving Exchange databases to shared storage (EVS2)

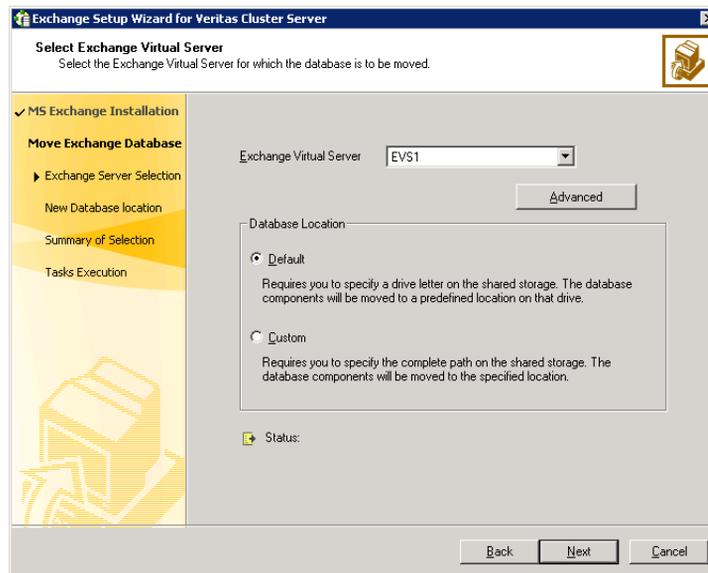
After completing the Exchange installation on the new first node (SYSTEM3), move the Exchange databases on that “first node” from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

- Complete the pre-installation, installation, and post-installation procedures to create the new Exchange virtual server on SYSTEM3. See “[Installing Exchange on the new nodes](#)” on page 295.
- The databases from the new active Exchange node (SYSTEM3) must be on separate disks from the shared storage for the existing active / passive first Exchange virtual server, EVS1. See “[Configuring disk groups and volumes](#)” on page 286.

- Make sure the data queue is empty on the SMTP server.
- Make sure to import the disk group and mount the volumes for the Exchange database, MTA data, and transaction logs.
See “[Managing disk groups and volumes](#)” on page 293.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.
- 5 In the Select Exchange Virtual Server dialog box:



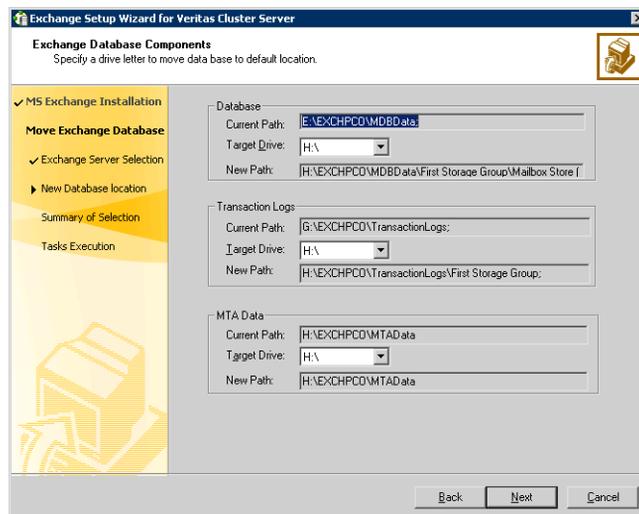
- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

Warning: The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

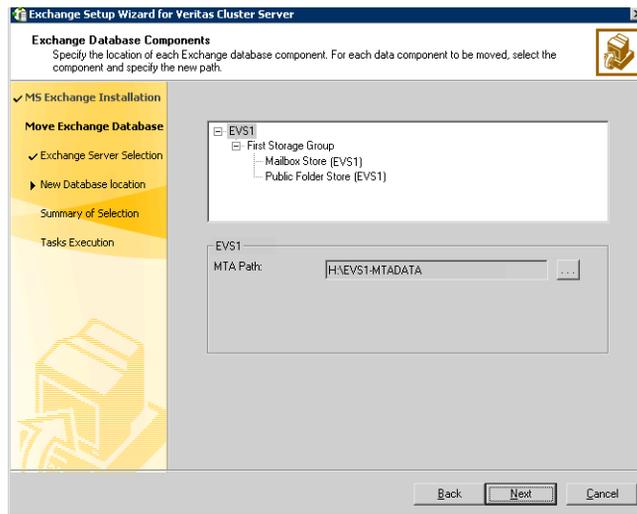
- Click **OK**.
 - Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
 - Click **Next**.
- 6 Do one of the following:
- If you would like to move the first mailbox store, public store, and MTA data to the generated default paths on the volumes for these components, proceed to the next step.
 - Otherwise, proceed to [step 8](#) on page 302 to specify the path location on the volumes that you will designate for these components.

Warning: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

- 7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components will be moved to a default location on that drive. In the Exchange Database Components dialog box:



- Specify the drive where the Exchange database will be moved.
 - Specify the drive where the Exchange Transaction Logs will be moved.
 - Specify the drive where the Exchange MTA Data will be moved.
 - Click **Next** and proceed to [step 9](#) on page 302.
- 8 For the option of a custom database location, specify the location for specific Microsoft Exchange data components. In the Exchange Database Components dialog box:



For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the path for the Exchange database components contains only ANSI characters.

- 9 Review the summary of your selections and click **Next**.
- 10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.
- 11 Click **Finish** to exit the wizard.

Installing Exchange on additional nodes

In the example with SYSTEM1 and SYSTEM2 in an existing active / passive cluster, no “additional nodes” exist. If other new failover nodes are in your configuration, or existing nodes do not yet have Exchange installed, complete the following tasks on these nodes.

After moving the Exchange databases to shared storage, install Exchange on additional nodes in the cluster for the same Exchange virtual server (EVS1 or EVS2). You must run preinstallation, installation, and post-installation procedures for each additional node.

For any-to-any add the additional nodes to the first Exchange virtual server, EVS1. In the later procedure “[Preparing the cluster with the any-to-any option](#)” on page 308, they will also be added to the second Exchange virtual server, EVS2, to configure the any-to-any failover.

Note: Make sure to review the prerequisites for permissions in “[Installing Exchange on the new nodes](#)” on page 295.

Exchange pre-installation: Additional nodes

Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes.

See “[Unmounting a volume and deporting a disk group](#)” on page 293.

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.

- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server.
The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically generated when setting up an additional failover mode for the virtual server in the same cluster.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.
- 11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: Additional nodes

Install Exchange on the same node selected in “[Installing Exchange on additional nodes](#)” on page 303.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

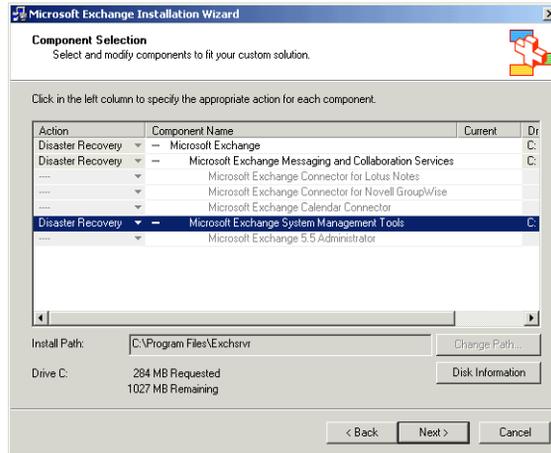
- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe  
/disasterrecovery
```

where <drive letter> is the location where the Exchange software is located.

- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools

components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

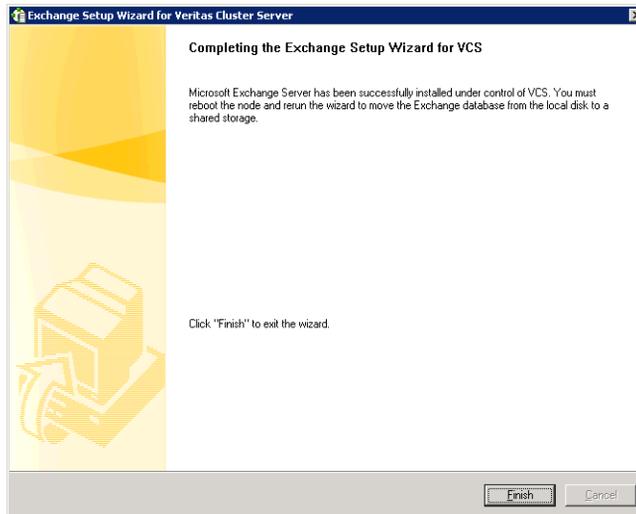
Exchange post-installation: Additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.
- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

Specifying a common node for failover

Specifying a common node for failover involves preparing the cluster with the Exchange Setup Wizard for VCS, and configuring the Exchange service group.

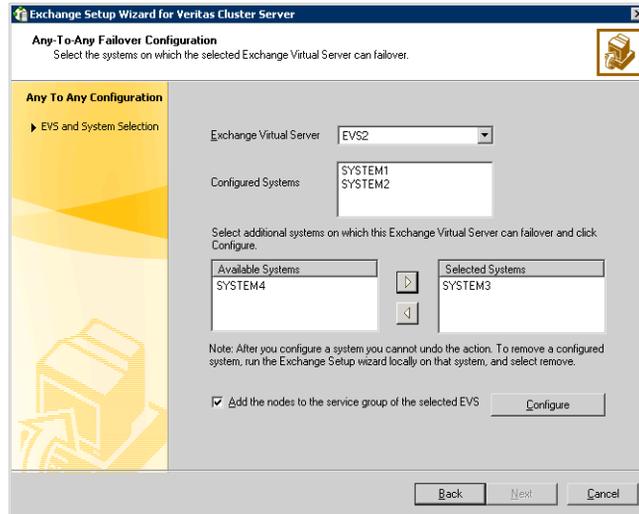
Preparing the cluster with the any-to-any option

Launch the Exchange Setup Wizard with the any-to-any option from any system in the cluster.

To prepare the cluster with the any-to-any option

- 1 Start the Exchange Setup Wizard for VCS from any node configured to host an Exchange service group. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard**)
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Options dialog box, choose the **Configure any-to-any failover** option and click **Next**.

- 5 Select systems to be configured for any-to-any failover. The **Configured Systems** box lists the nodes on which the Exchange Server service group can fail over.



- Select the Exchange virtual server to which you want to add the additional failover nodes.
 - From the **Available Systems** box, select the systems to be configured for any-to-any failover.
 Select any new nodes that were added as “additional nodes” on the first Exchange virtual server in [“Installing Exchange on additional nodes”](#) on page 303, as well as the existing nodes that will be any-to-any failover nodes.
 The **Available Systems** box lists only those systems that have the same version and service pack level of Microsoft Exchange as the selected Exchange virtual server.
 - Click the right arrow to move the selected systems to the **Selected Systems** box. To remove a system from the box, select the system and click the left arrow.
 - Specify whether you want to add the systems to the SystemList of the service group for the selected EVS.
 - Click **Configure**.
 - Click **Next**.
- 6 Click **Finish**.

The failover nodes for the first Exchange virtual server, EVS1, were already set in the existing active / passive cluster.

Configuring the Exchange service group for VCS

A new Exchange service group must be configured for the new Exchange virtual server, EVS2.

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

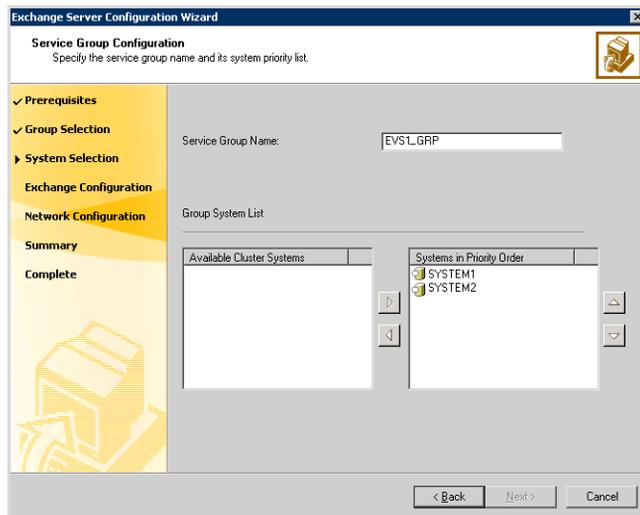
Prerequisites

- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk group and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - registry changes related to Exchange
 - transaction logs for the first storage group
 - MTA databaseFor instructions on mounting, see [“Importing a disk group and mounting a shared volume”](#) on page 293. For instructions on unmounting, see [“Unmounting a volume and deporting a disk group”](#) on page 293.
- Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to an already configured service group.

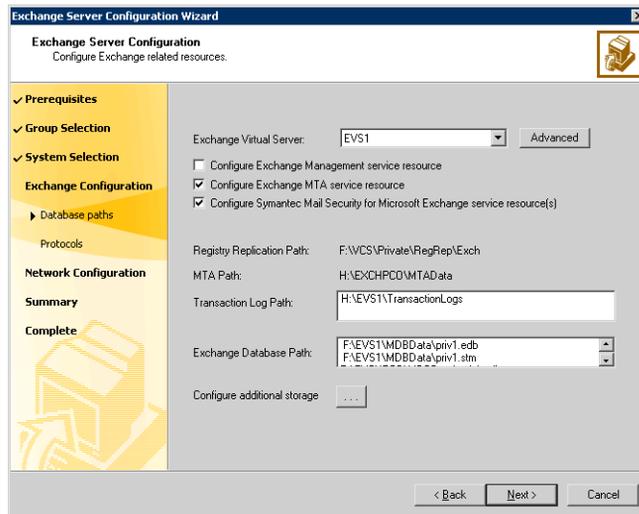
To configure the Exchange service group

- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options dialog box, choose the **Create service group** option and click **Next**.
- 4 Specify the service group name and system list:



- Enter a name for the Exchange service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list. Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.

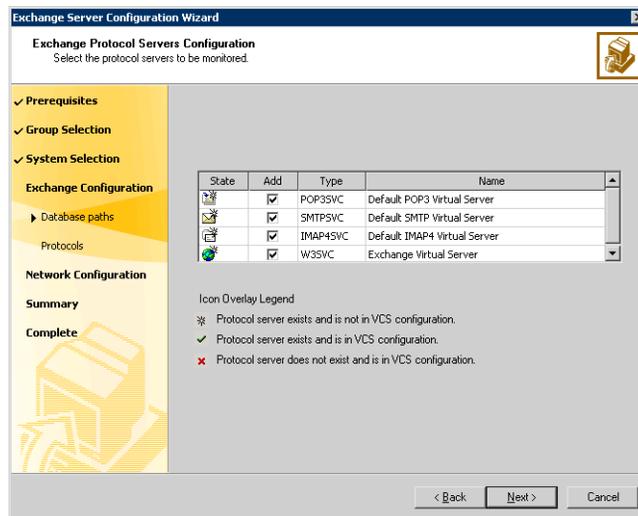
- To change a node's priority in the service group's system list, select the node in the Systems in Priority Order list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
 - Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:



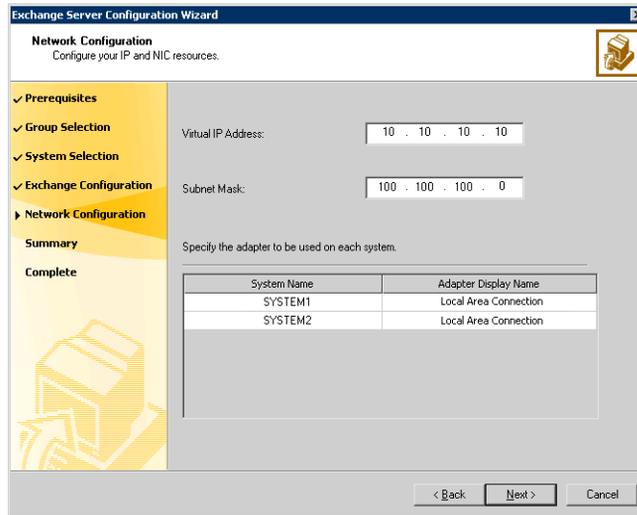
- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.
- Specify whether you want to configure the Exchange Management service.
 - Specify whether you want to configure the Exchange MTA service resource. The Exchange Message Transfer Agent is specific to

legacy Exchange message routing based on X.400. There are specific circumstances where it may or may not be required for your Exchange server. Please refer to Microsoft documentation on Exchange message routing and the use of MTA for further clarification and applicability to its use within your Exchange environment.

- If you are utilizing Symantec Mail Security for Exchange be sure to indicate that you would like to configure this resource.
 - Click **Next**.
- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



7 Specify information related to the network:

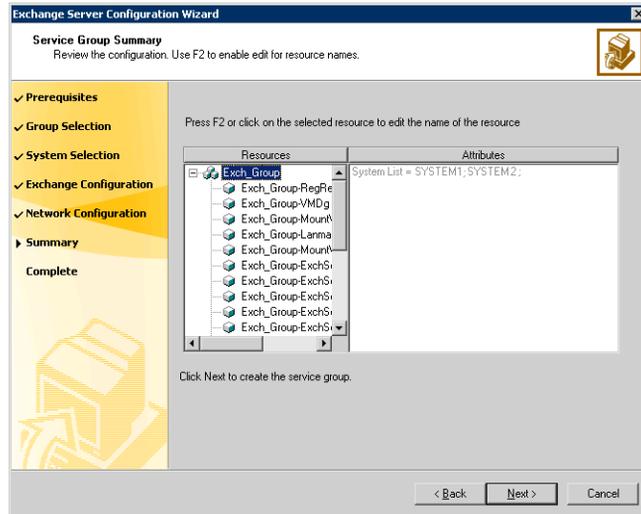


- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

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

- Click **Next**.

- 8 Review the service group configuration and change the resource names, if desired:



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 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.
 - Click **Next**.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 9 In the Completing the Exchange Configuration dialog box, select the **Bring the service group online** check box to bring the service group online on the local system.
 - 10 Click **Finish**. After bringing the service group online, you must run the Exchange System Manager so that all the stores are automatically mounted on start-up.

To reconfigure mounting of stores at start-up

- 1 Start Exchange System Manager.

- 2 In the left pane, navigate to your storage group.
If administrative groups are not configured, Servers > Exchange Server > Storage Group.
If more than one administrative group is configured, expand Administrative Groups > Your Administrative Group > Servers > Exchange Server > Storage Group.
- 3 Right-click the Exchange database and choose **Properties** from the pop-up menu.
- 4 Click the Database tab.
- 5 Clear the **Do not mount this store at start-up** check box.
- 6 Click **OK**.

Repeat these steps for all the Exchange databases that were previously mounted.

If you need to configure additional storage groups or mailbox stores on the shared storage you should do that now. Import disk groups and mount volumes that have been created for the additional storage groups or mailbox stores, and create the new storage groups and mailbox stores in Exchange System Manager, and then rerun the Exchange Configuration Wizard to bring them under VCS control. If you already designated the additional mounted volumes and disk groups when you ran the configuration wizard the first time then you can just create the storage groups and mailbox stores in Exchange System Manager.

Verifying the cluster configuration

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

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

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.

- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

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

Campus Cluster

Administrators can use campus clusters to protect data from natural disasters, such as floods and hurricanes, and unpredictable power outages. Campus clusters provide a layer of protection that extends beyond local high availability but is not as complex as disaster recovery with replication.

Refer to the following chapters to install and configure Exchange in an SFW HA campus cluster:

- [Chapter 8, “Campus cluster for Exchange: Overview” on page 321](#)
- [Chapter 9, “Deploying SFW HA for Campus Cluster: New Installation” on page 325](#)

Campus cluster for Exchange: Overview

What is a campus cluster?

Campus clusters are multiple-node clusters that provide protection against disasters. These clusters are in separate buildings (or sites) with mirrored SAN-attached storage located in each building. Typical campus clusters involve two sites; you can use more than two sites for additional redundancy.

In a typical configuration, each node has its own storage array with the same number of disks and contains mirrored data of the storage on the other array. Refer to [Chapter 9, “Deploying SFW HA for Campus Cluster: New Installation” on page 325](#), for details on a typical active/passive configuration for a campus cluster.

Why implement a campus cluster?

In the event of a site disaster, such as power failure in a building, campus clusters offer a level of high availability that surpasses mirroring or clustering at a single site by dispersing the clustered servers into different buildings or sites. This environment also provides a simpler solution for disaster recovery than a more elaborate Veritas DR environment with replication software; however, a campus cluster generally stretches a shorter distance than a replication-based solution depending on the hardware.

What is high availability?

High Availability (HA) refers to a state where data and applications are highly available because software or hardware maintains the continued functioning in the event of computer failure. HA can refer to any software or hardware that provides fault tolerance, but generally the term has become associated with clustering.

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 goal for critical applications. Clustered systems have several advantages, including fault tolerance, high availability, scalability, simplified management, and support for rolling upgrades.

Using VCS 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; accommodates flexible failover policies, any-to-any configurations, and shared standby servers for Exchange.

How the VCS application agent makes Microsoft Exchange highly available

The VCS application agent for Microsoft Exchange Server detects an application failure if a configured Exchange service is not running or if a configured virtual server is not available. When this occurs, the Exchange service group is failed over to the next available system in the service group's system list. The configured Exchange services and virtual servers are started on the new system.

This ensures continuous availability for Exchange data and configured mailboxes.

Campus cluster failover using the ForceImport attribute

To ensure proper failover in a VCS campus cluster, you must verify the value of the ForceImport attribute of the VMDg resource. The table below lists failure situations and the outcomes depending on the settings for the ForceImport attribute. You can set this attribute to 1 (forcing the import of the disk groups to the other node) or 0 (not forcing the import). Use the VCS Java Console or command line to modify the ForceImport attribute.

Table 8-6 Failure Situations

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
1) Application fault May mean the services stopped for an application, a NIC failed, or a database table went offline.	Application automatically moves to another node.	Service Group failover is automatic on the standby or preferred system or node.
2) Server failure May mean a power cord became unplugged or a failure caused the system to stop responding.	Application automatically moves to other node. 100% of the disks are still available.	Service Group failover is automatic on the standby or preferred system or node. 100% of the mirrored disks are still available.
3) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.	No interruption of service. Remaining disks in mirror are still accessible from the other node.	The Service Group does not failover. 50% of the mirrored disk is still available at remaining site.
4) Zone failure Complete Site failure, all accessibility to the servers and storage is lost.	Manual intervention required to online the Service Group at remaining site. Can not automatically import 50% of mirrored disk.	Automatic failover of Service Group to online site. Force Import must be set to True before site failure to ensure VCS can import 50% of mirrored disk.

Table 8-6 Failure Situations (continued)

Failure Situation	Forcelmport set to 0 (import not forced)	Forcelmport set to 1 (automatic force import)
<p>5) Split-brain (loss of both heartbeats)</p> <p>If the public network link serves as a low-priority heartbeat, the assumption is made that the link is also lost.</p>	<p>No interruption of service. Can't import disks because the original node still has the SCSI reservation.</p>	<p>No interruption of service. Failover does not occur due to Service Group resources remaining online on the original nodes. Example: Online node has SCSI reservation to own disk.</p>
<p>6) Storage interconnect lost</p> <p>Fibre interconnect severed.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>No interruption of service. Service Group resources remain online, but 50% of the mirror disk becomes detached.</p>
<p>7) Split-brain and storage interconnect lost</p> <p>If a single pipe is used between buildings for the Ethernet and storage, this situation can occur.</p>	<p>No interruption of service. Cannot import with only 50% of disks available. Disks on the same node are functioning. Mirroring is not working.</p>	<p>Automatically imports 50% of mirrored disk to the alternate node.</p> <p>Disks online for a short period in both locations but offlined again due to IP and other resources being online on original node. No interruption of service.</p>

Deploying SFW HA for Campus Cluster: New Installation

This chapter covers the following topics:

- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the network and storage](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Configuring the cluster](#)
- [Managing disk groups and volumes](#)
- [Preparing the forest and domain](#)
- [Installing Exchange on the first node](#)
- [Moving Exchange databases to shared storage](#)
- [Installing Exchange on additional nodes](#)
- [Configuring the Exchange service group for VCS](#)
- [Modifying the IP resource in the Exchange service group](#)
- [Verifying the campus cluster: Switching the service group](#)
- [Possible tasks after creating the campus cluster](#)

This chapter provides information on how to install and configure a new Veritas Storage Foundation HA environment for Exchange in a campus cluster. This environment provides high availability and disaster recovery that extends

beyond local clustering and mirroring at a single site, but is not as complex as Veritas DR solutions with replication.

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

Table 9-1 Task list

Objective	Tasks
“Reviewing the requirements” on page 327	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 331	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive Exchange configuration in a two-node campus cluster
“Configuring the network and storage” on page 332	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which Exchange will be installed
“Installing Veritas Storage Foundation HA for Windows” on page 334	<ul style="list-style-type: none"> ■ Verifying the driver signing options for Windows 2003 remote systems ■ Installing SFW and VCS (automatic installation) and installing Veritas Cluster Server Application Agent for Microsoft Exchange ■ Restoring driver signing options for the Windows 2003 remote systems
“Configuring the cluster” on page 340	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the Veritas Cluster Server Configuration Wizard
“Configuring disk groups and volumes” on page 355	<ul style="list-style-type: none"> ■ Creating disk groups ■ Creating the data, log, RegRep, and MTA volumes
“Preparing the forest and domain” on page 366	<ul style="list-style-type: none"> ■ Setting up the forest and domain prior to the Exchange installation
“Installing Exchange on the first node” on page 366	<ul style="list-style-type: none"> ■ Reviewing the prerequisite checklist ■ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server Installation Wizard
“Moving Exchange databases to shared storage” on page 370	<ul style="list-style-type: none"> ■ Moving databases on the first node from the local drive to the shared drive using the Exchange Setup Wizard for Veritas Cluster Server

Table 9-1 Task list (continued)

Objective	Tasks
“Installing Exchange on additional nodes” on page 374	<ul style="list-style-type: none"> Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server Installation Wizard
“Configuring the Exchange service group for VCS” on page 379	<ul style="list-style-type: none"> Creating the Exchange service group using the VCS Exchange Configuration Wizard.
“Modifying the IP resource in the Exchange service group” on page 385	<ul style="list-style-type: none"> Modifying the Address and SubNetMask attributes if the sites are in different subnets.
“Possible tasks after creating the campus cluster” on page 387	<ul style="list-style-type: none"> If a site failure occurs, setting the ForceImport attribute of the VMDg resource to 1 to ensure proper failover.

Reviewing the requirements

The campus cluster solution allows for clustered systems with mirrored or synchronously replicated storage arrays to be implemented in separate datacenters, located either within the same building or separate buildings. For example, datacenter A could be located in building A and datacenter B located in building B. This guide will refer to these different areas as Site A and Site B.

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

Disk space requirements

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

[Table 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

Table 9-2 Disk space requirements

Installation options	Installation on system drive	Installation on non-system drive
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 Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:
 - Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition (SP2 required)
with
 Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
or
 Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP1 required for all editions)
or
 Windows Server 2003 R2 (Standard Edition, Enterprise Edition, or Datacenter Edition)
or
 - Microsoft Exchange 2000 Server or Microsoft Exchange 2000 Enterprise Server (SP3 with August 2004 rollup patch required for both editions)
with

Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)

Note: Microsoft support for Exchange Server 2003 is limited to 32-bit versions of the Windows 2003 operating system.

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 331.
- 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 Exchange Virtual Server (EVS)
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.

- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
See the *Veritas Cluster Server Bundled Agents Reference Guide*.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Domain Servers group.
- 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.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- You must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.
- 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.
- You must install IIS, SMTP, NNTP, ASP.NET, and WWW services before you install Microsoft Exchange Server.

Best practices

Symantec recommends that you perform the following tasks:

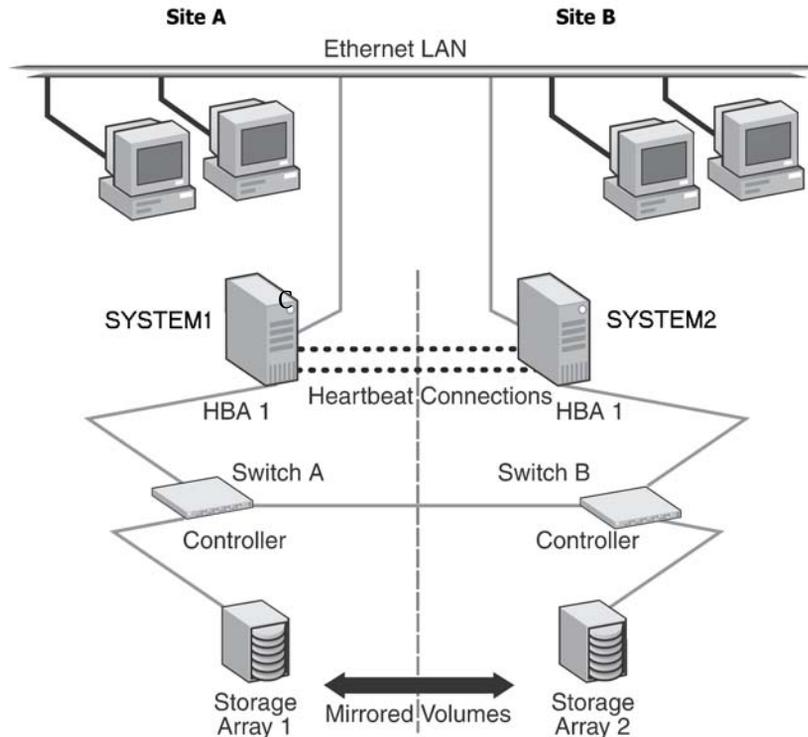
- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the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

Reviewing the configuration

This chapter uses the example of a two-node campus cluster with each node in a separate site (Site A or Site B). In this example, each node has its own storage array with the same number of disks and contains mirrored data of the storage on the other array.

The campus cluster involves an active/passive configuration for Exchange with one to one failover capabilities. In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed

by a service group configured with a set of nodes in the cluster. In this case, EVS1 can fail over from SYSTEM1 to SYSTEM2 and vice versa.



The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number of disks and contains mirrored data of the storage on the other array. The example describes a generic database application.

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

Configuring the network and storage

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

To configure the hardware

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

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
- 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.
- 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 that will run Exchange

- 1 Open the Control Panel (**Start > Settings > 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 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.
 - c 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:
 - a Select the **Internet Protocol (TCP/IP)** check box.
 - b 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 Application Agent for Exchange. 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 9-3](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 9-3 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

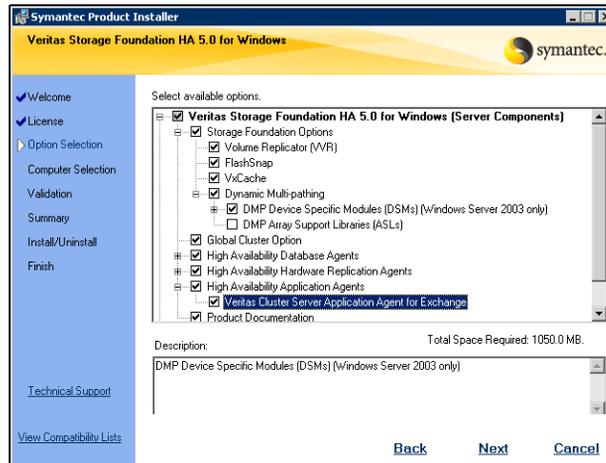
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.

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.

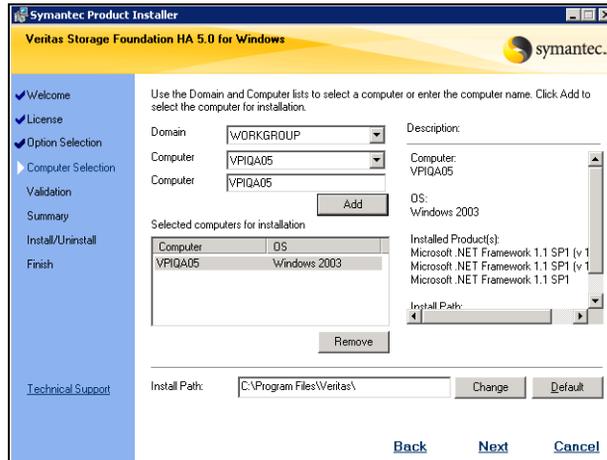
9 Select the appropriate SFW product options for your installation. Click **Next**.



The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

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

10 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

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

Computer

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

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

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

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

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
 - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
 - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

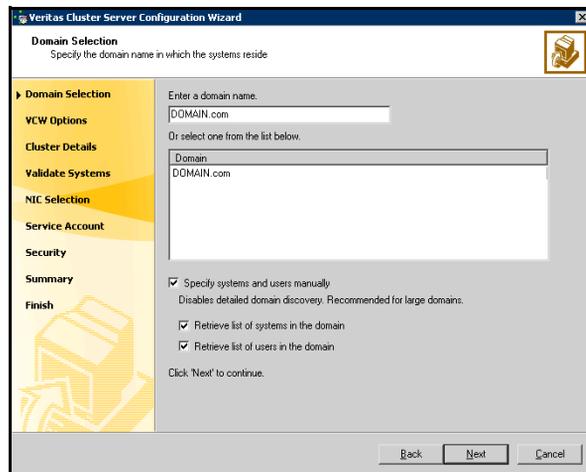
Complete the following tasks before creating a cluster:

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

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

To configure a VCS cluster

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



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

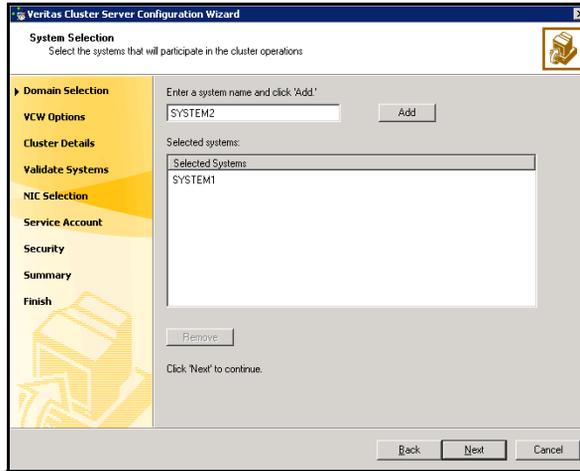
- Clear the **Specify systems and users manually** check box.
- Click **Next**.
 Proceed to [step 7](#) on page 343.

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**.

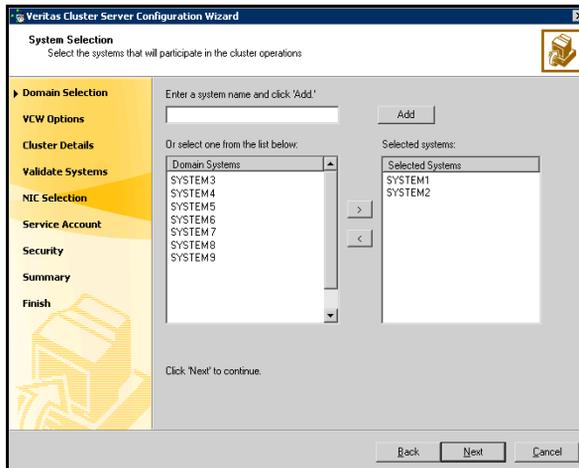
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 342. Otherwise, proceed to the next step.

- 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.



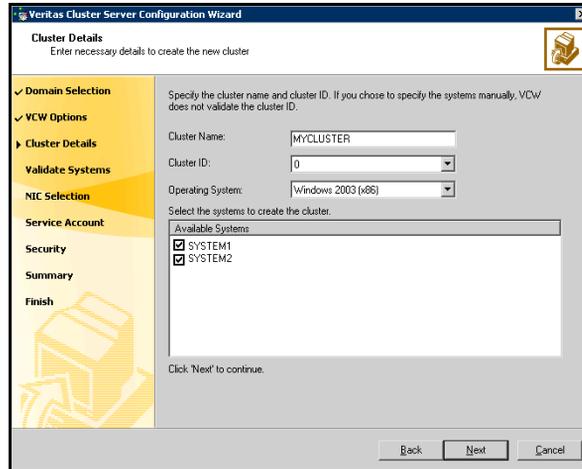
Proceed to [step 7](#) on page 343.

- 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.

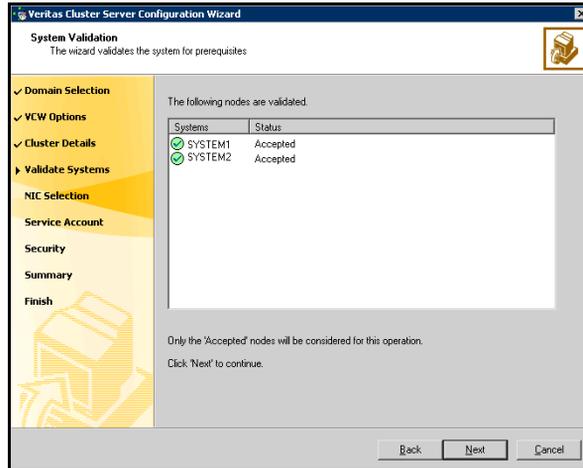
Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

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.

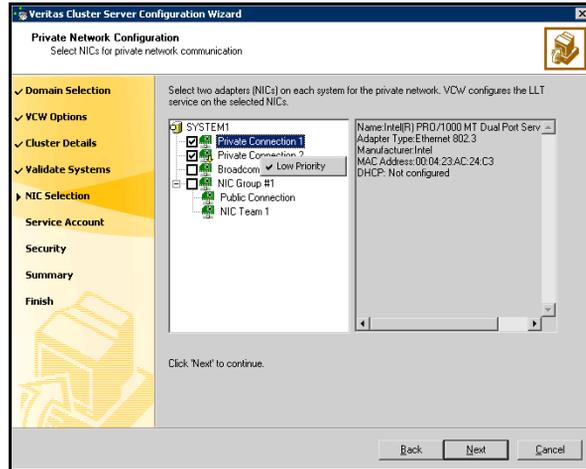
- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



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 343, proceed to the next step. Otherwise, proceed to [step 11](#) on page 345.

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

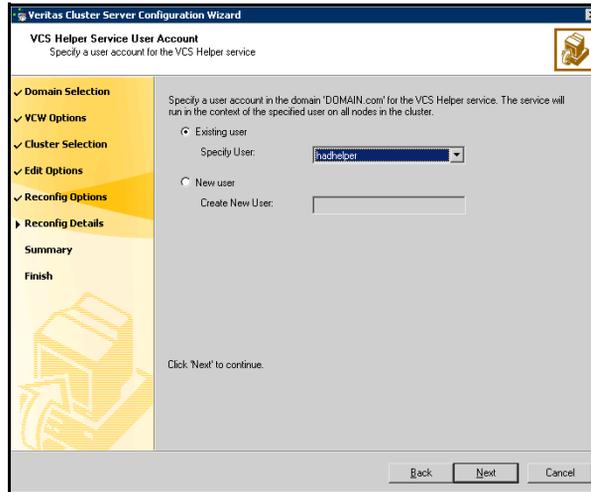


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs 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 341, 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**.

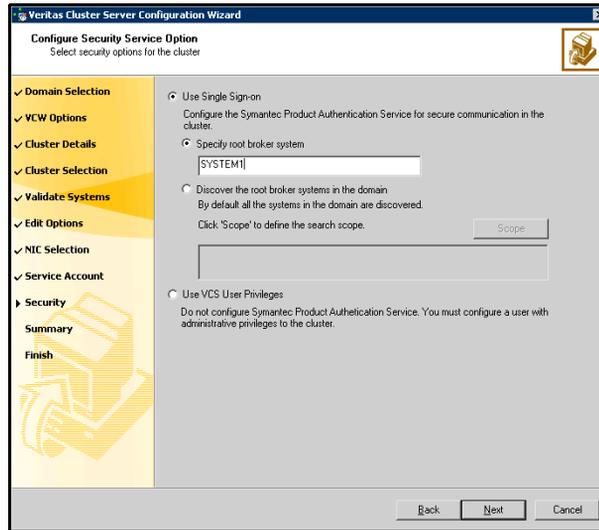
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- 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.

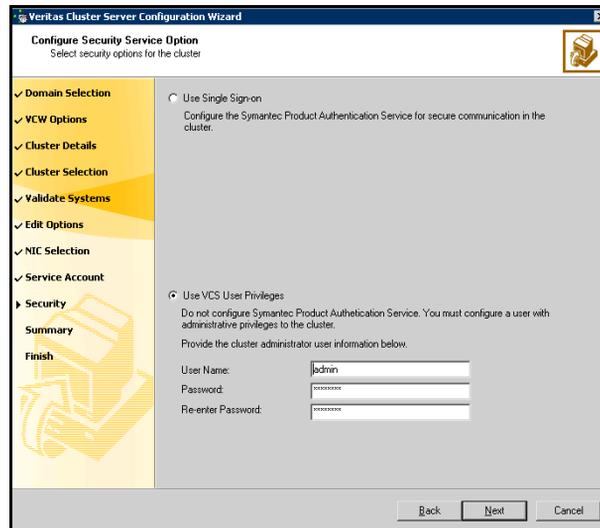
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- 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.

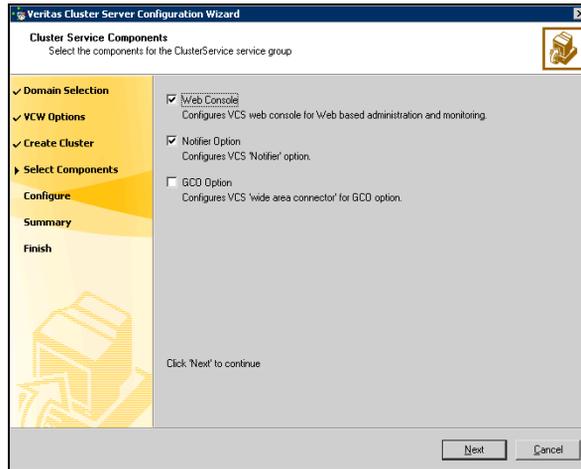
- Click **Next**.

- 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**.



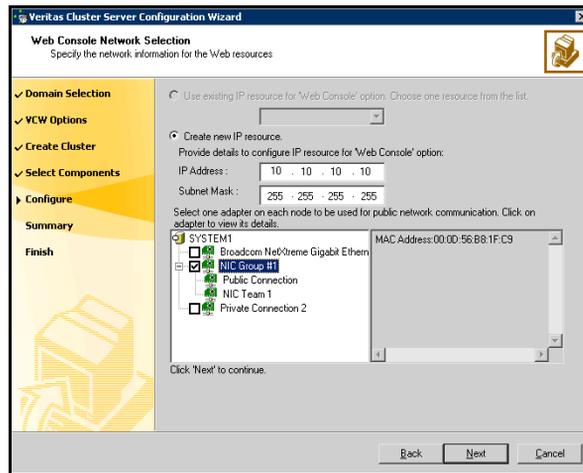
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 351.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 352.

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**.



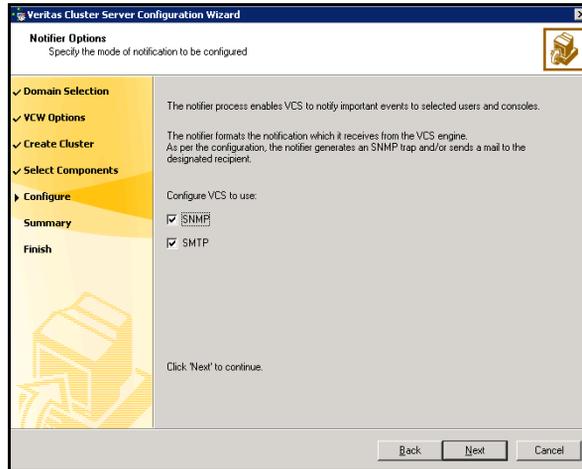
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 352. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

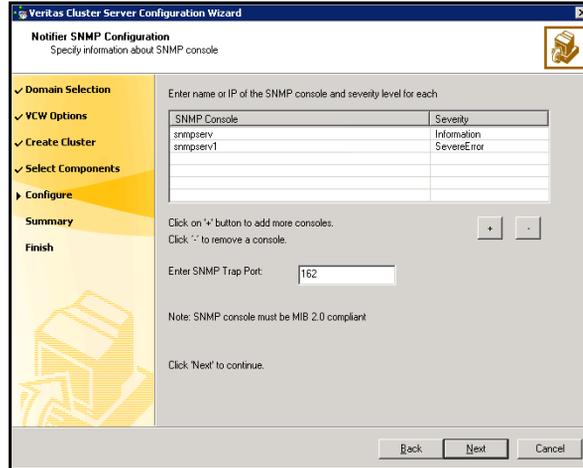
To configure notification

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



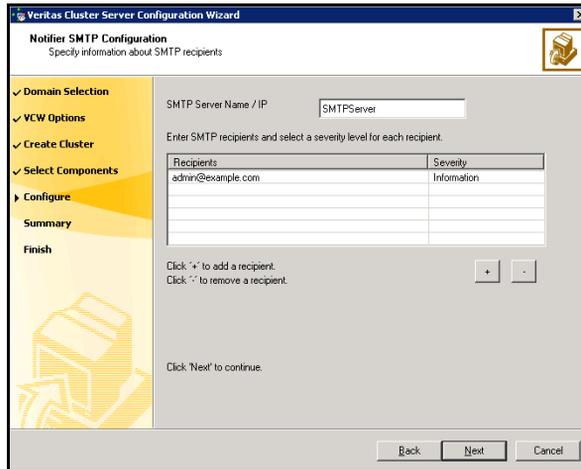
You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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



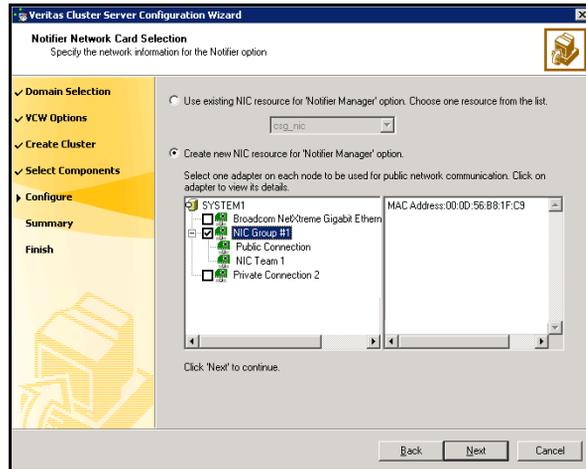
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- 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 disk groups and volumes

Before installing Exchange, you must create disk groups and mirrored volumes using the VEA console installed with SFW. This is also an opportunity to increase the size of existing volumes, add storage groups, and create volumes to support additional databases for storage groups.

Before you create a disk group, consider the following items:

- The type of volume configurations that are required
- The number of LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup

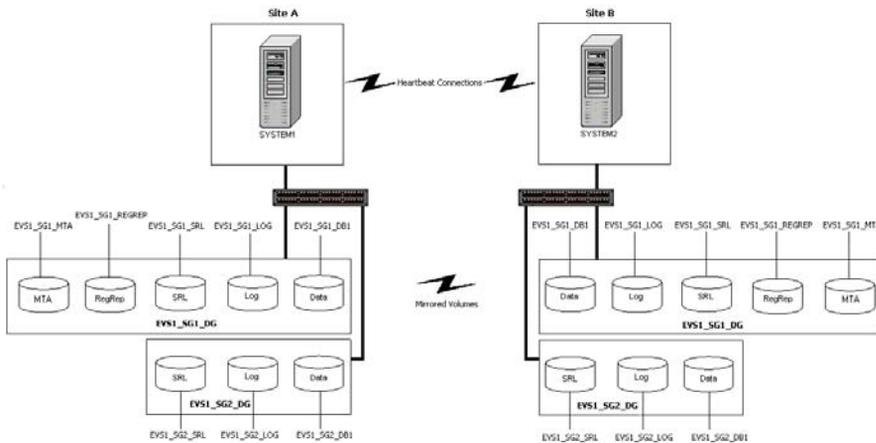
- The size of databases and logs that depend on the traffic load
- The disk groups and number of disks on each site
- Types of volumes required and location of the plex of each volume in the storage array

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

Note: Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

Typically, a SFW disk group corresponds to an Exchange storage group. [Figure 9-1](#) is a detailed view of the disk groups and volumes.

Figure 9-1 Disk groups and volumes for Exchange virtual server EVS1 in a campus cluster



Exchange storage group EVS1_SG1_DG contains the following volumes:

EVS1_SG1_DB1 Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.

EVS1_SG1_REGREP	Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.
EVS1_SG1_LOG	Contains the transaction log for the storage group.
EVS1_SG1_MTA	Contains the MTA database

Additional storage groups (for example, EVS1_SG2_DG) only contain the data, and log volumes; the RegRep and MTA volumes are included in the first storage group.

Use the following procedures to create disk groups and volumes. The guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Configuring the disks and volumes

Ensure that each disk group has the same number of disks on each site. Each volume must be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

See the following sections:

- [“Creating a dynamic \(cluster\) disk group”](#) on page 358
- [“Creating a volume”](#) on page 360

Considerations when creating new volumes

Consider the following when creating new volumes.

- For campus clusters, when creating a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.
- Symantec recommends using either simple mirrored (concatenated) or striped mirrored options for the new volumes. Striped mirrored gives you better performance compared to concatenated.

When selecting striped mirrored, select two columns in order to stripe one enclosure that is mirrored to the second enclosure.

- You cannot selecting RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.
The internal names for the disks that the current system can access for available storage are displayed, with names Harddisk1, Harddisk2, etc. The list includes both disks internal to the local system and any external storage that is available.

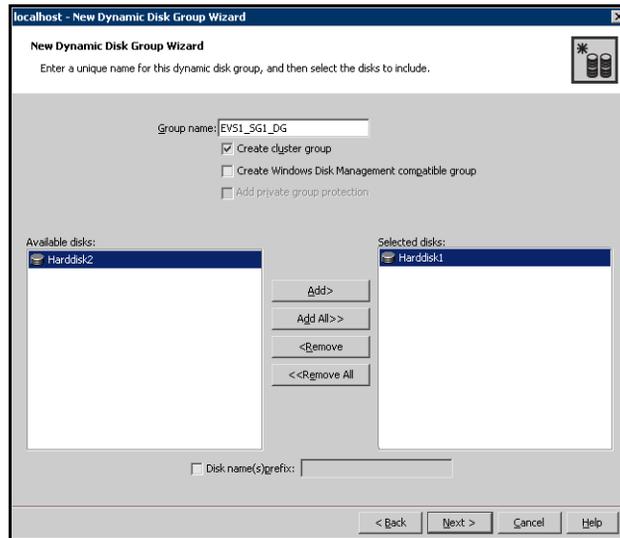
Creating a dynamic (cluster) disk group

Use the following procedure to create a dynamic cluster disk group.

To create a dynamic (cluster) disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_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.

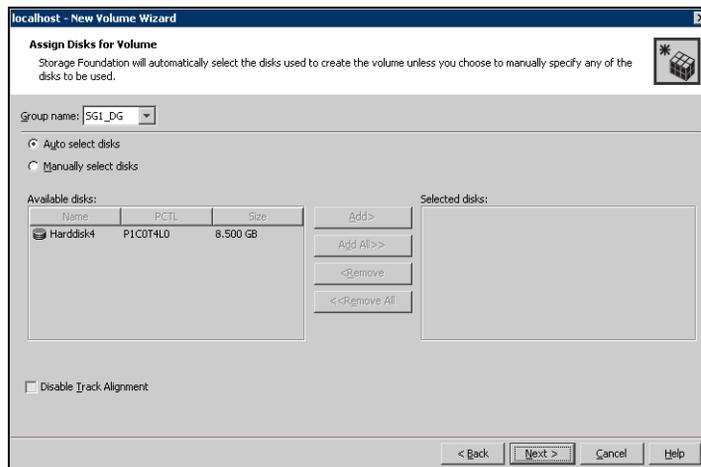
Proceed to create the appropriate volumes on each disk.

Creating a volume

This procedure assumes you are starting with the EVS1_SG1_DB1 volume.

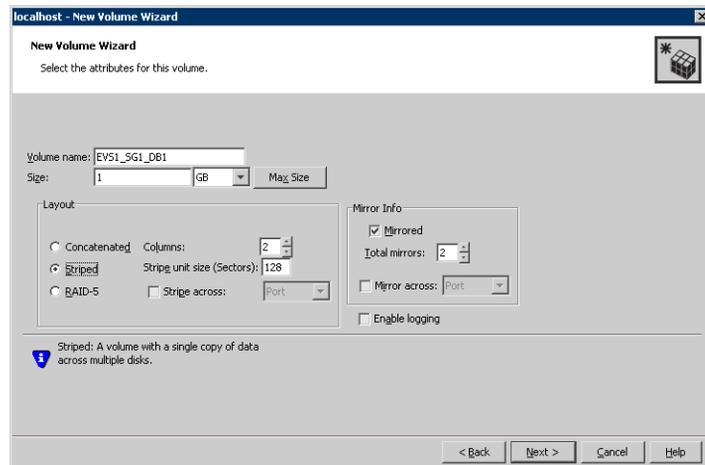
To create dynamic volumes

- 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 SG1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume; make sure the appropriate disk group name appears in the Group name drop-down list.



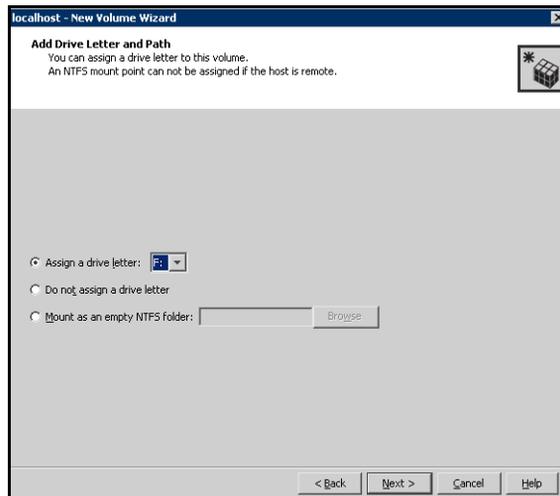
- 7 Select auto or manual disk selection and enable or disable track alignment.

- Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
 - To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.
- 9 Specify the volume attributes.



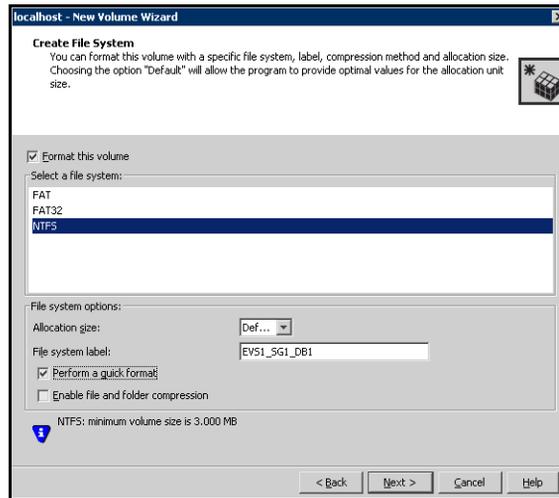
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume.
- If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.

- Select a volume layout type. For campus clusters, select either **Concatenated** or **Striped**. Since campus clusters are mirrored volumes, you also must select the **Mirrored** checkbox.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
 - In the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
 - Verify that **Enable logging** is not selected and click **Next**.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- Do not assign the M: drive letter if you are using Exchange 2000, as it is reserved for internal use.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
 - To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.

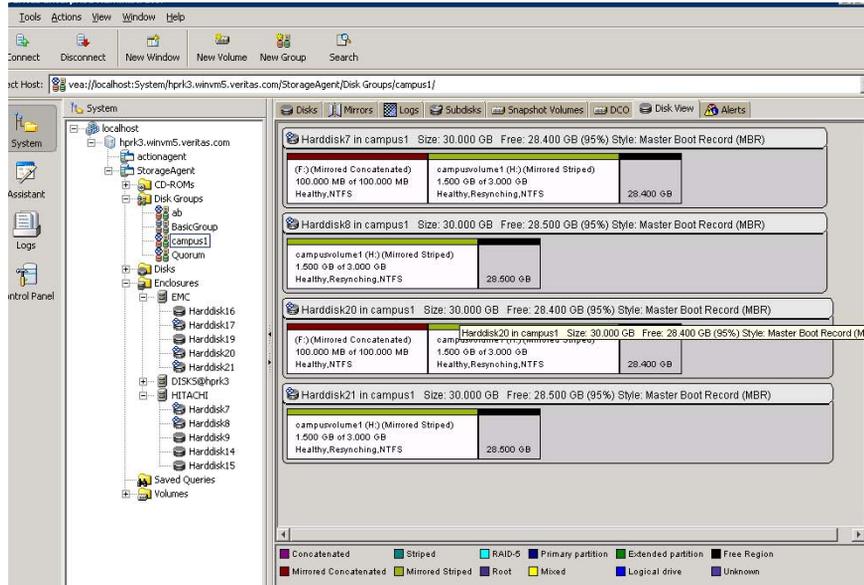
12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and select **NTFS**.
 - Select an allocation size or accept the Default.
 - The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance. Click **Next**.
- 13 Click **Finish** to create the new volume.
- 14 Repeat these steps to create the RegRep volume (EVS1_SG1_REGREP), the MTA volume (EVS1_SG1_MTA), and the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.)
- 15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3).

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

Figure 9-2 View of disks with volumes in VEA console



Managing disk groups and volumes

This section includes the following procedures:

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

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Preparing the forest and domain

Microsoft requires the preparation of the forest and domain prior to an Exchange installation. See Microsoft documentation for instructions for preparing the forest and domain for an Exchange installation. Do not repeat this process for additional Exchange installations.

Installing Exchange on the first node

Installing Exchange on the first node is described in three stages that involve preinstallation, installation, and post-installation procedures. Complete the following tasks before installing Exchange Server:

- ✓ Prepare the forest and domain. Refer to [“Preparing the forest and domain”](#) on page 366 for instructions.
- ✓ Verify the disk group is imported on the first node of the cluster. Refer to [“Importing a disk group and mounting a shared volume”](#) on page 364 for instructions.
- ✓ Mount the volume containing the information for registry replication (EVS1_SG1_REGREP). Refer to [“Unmounting a volume and deporting a disk group”](#) on page 364 for instructions.
- ✓ Verify that all systems on which Exchange Server will be installed have IIS installed; you must install SMTP, NNTP, and WWW services on all systems. If you install Exchange on Windows 2003, make sure to install ASP.NET service.
- ✓ Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- ✓ Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.

- If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

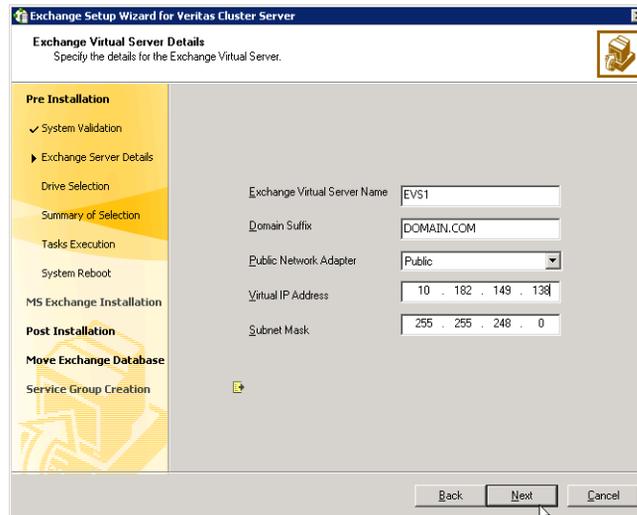
Exchange pre-installation: First node

You must install Exchange on a virtual node to facilitate high availability. Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a new Exchange Virtual Server** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.

7 Specify information related to the network.



- Enter a unique virtual name for the Exchange server.
Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.
 - Enter a domain suffix for the virtual server.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Enter a unique virtual IP address for the Exchange virtual server.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met. Click **Next**.
- 8 Select a drive where the registry replication data will be stored and click **Next**.

- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: First node

Install Exchange on the same node selected in “[Exchange pre-installation: First node](#)” on page 367.

Exchange 2000 requires service pack 3 with the August 2004 rollup patch. Exchange 2003 requires service pack 2. The procedure below is based on Exchange 2003 SP2.

This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Install Exchange Server using the Microsoft Exchange installation program. See the Microsoft Exchange documentation for instructions.
- 2 Reboot the node if prompted to do so.
- 3 For Exchange 2000 or Exchange 2003, install the required service pack.

Exchange post-installation: First node

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

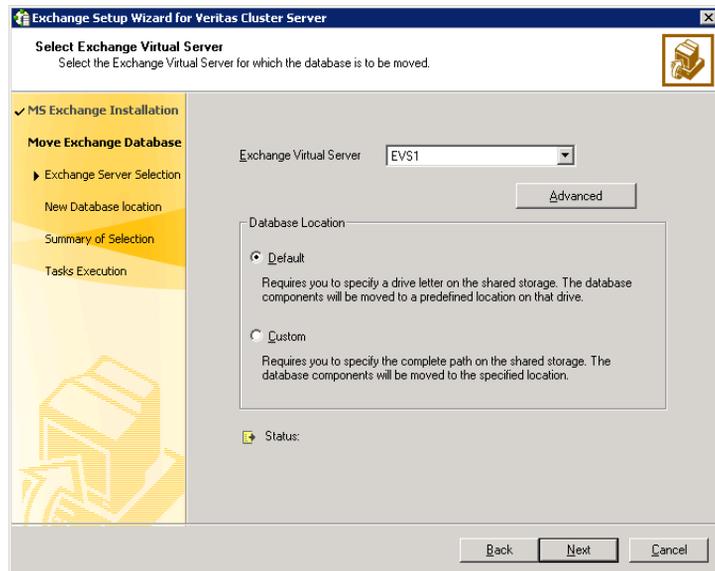
Moving Exchange databases to shared storage

After completing the Exchange installation on the first node, move the Exchange databases on the first node from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

- Make sure the data queue is empty on the SMTP server.
- Make sure to import the disk group and mount the volumes for the Exchange database, MTA data, and transaction logs. Refer to “[Managing disk groups and volumes](#)” on page 364 for instructions.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.
- 5 In the Select Exchange Virtual Server dialog box:



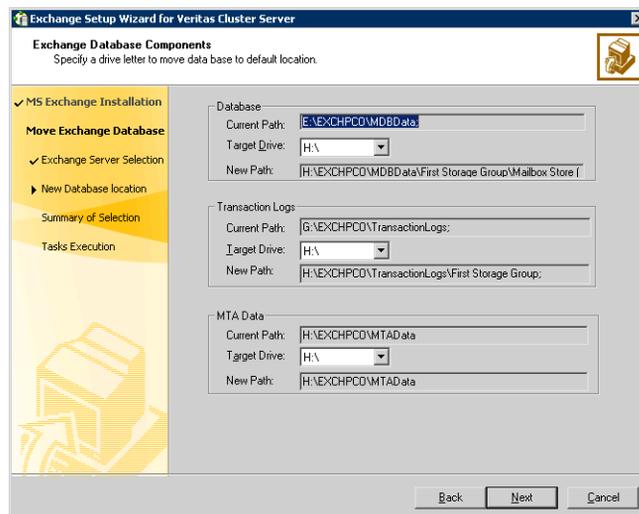
- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

Warning: The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Click **OK**.
 - Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
 - Click **Next**.
- 6 Do one of the following:
- If you would like to move the first mailbox store, public store, and MTA data to the generated default paths on the volumes for these components, proceed to the next step.
 - Otherwise, proceed to [step 8](#) on page 373 to specify the path location on the volumes that you will designate for these components.

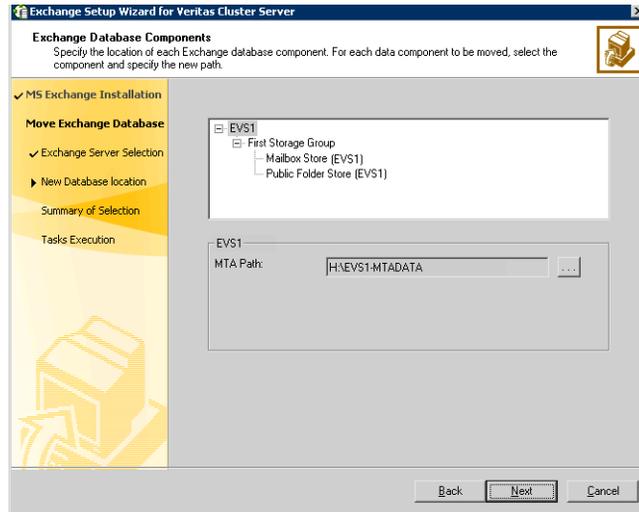
Warning: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

- 7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components will be moved to a default location on that drive. In the Exchange Database Components dialog box:



- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Specify the drive where the Exchange MTA Data will be moved.
- Click **Next** and proceed to [step 9](#) on page 373.

- 8 For the option of a custom database location, specify the location for specific Microsoft Exchange data components. In the Exchange Database Components dialog box:



For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the path for the Exchange database components contains only ANSI characters.

- 9 Review the summary of your selections and click **Next**.
- 10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.
- 11 Click **Finish** to exit the wizard.

Installing Exchange on additional nodes

After moving the Exchange databases to shared storage, install Exchange on additional nodes in the cluster for the same Exchange virtual server (EVS1). You must run preinstallation, installation, and post-installation procedures for each additional node.

Note: Make sure to review the prerequisites for permissions in “[Installing Exchange on the first node](#)” on page 366.

Exchange pre-installation: Additional nodes

Note: Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes.

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server.

The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.

- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically generated when setting up an additional failover mode for the virtual server in the same cluster.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
 - 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.
 - 11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
 - 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: Additional nodes

Install Exchange on the same node selected in “[Exchange pre-installation: Additional nodes](#)” on page 374.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

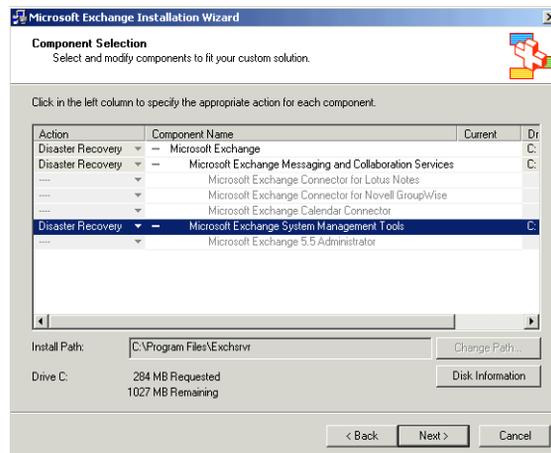
The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe /disasterrecovery
```

where **<drive letter>** is the location where the Exchange software is located.
- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

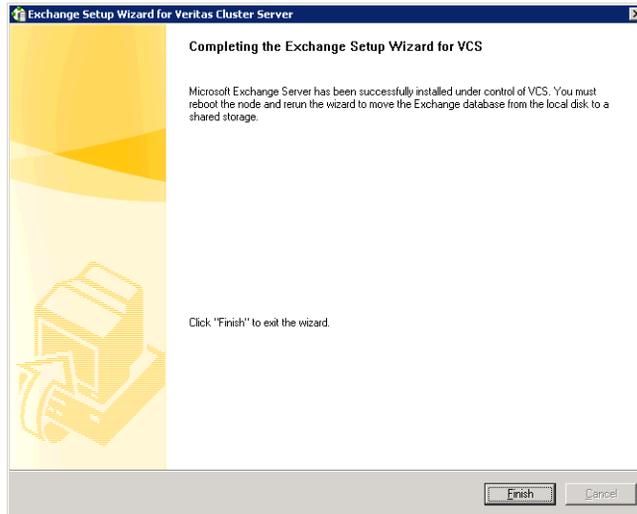
Exchange post-installation: Additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.

- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

Configuring the Exchange service group for VCS

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

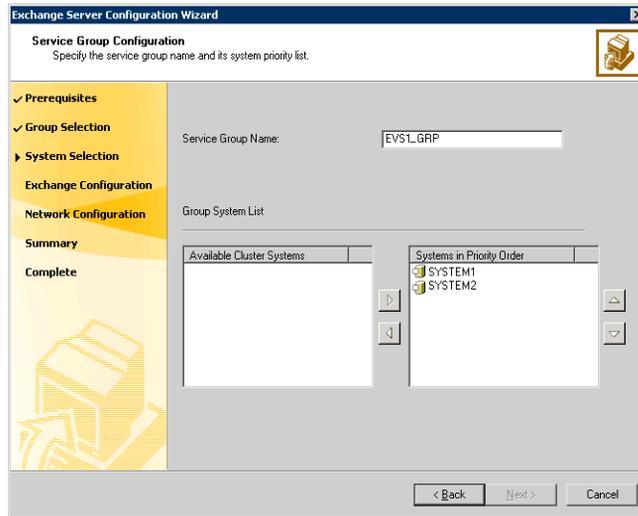
Prerequisites

- You must be a Cluster Administrator.
- Verify that Command Server is running on all nodes in the cluster.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard.
- Mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - registry changes related to Exchange
 - transaction logs for the first storage group
 - MTA database
- Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

To configure the Exchange service group

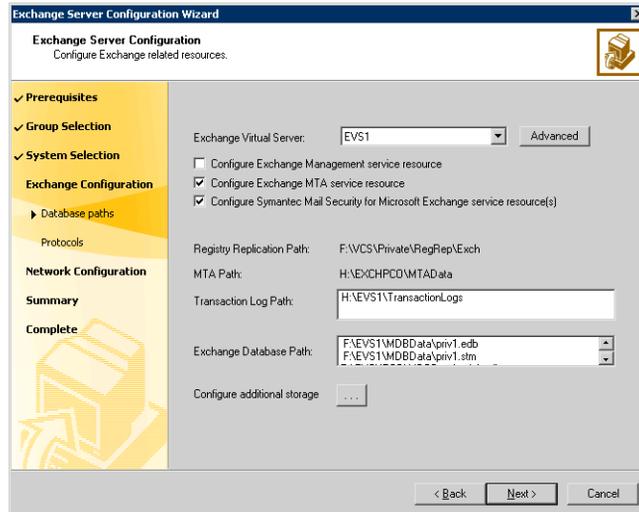
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options dialog box, choose the **Create service group** option and click **Next**.

4 Specify the service group name and system list:



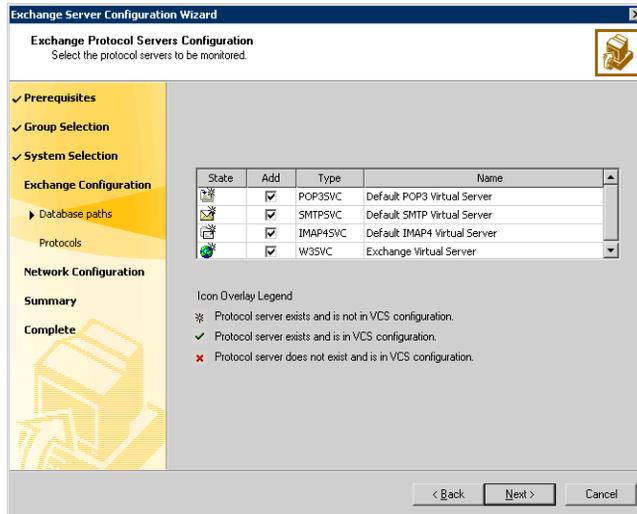
- Enter a name for the Exchange service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list. Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.
- To change a node's priority in the service group's system list, select the node in the Systems in Priority Order list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:

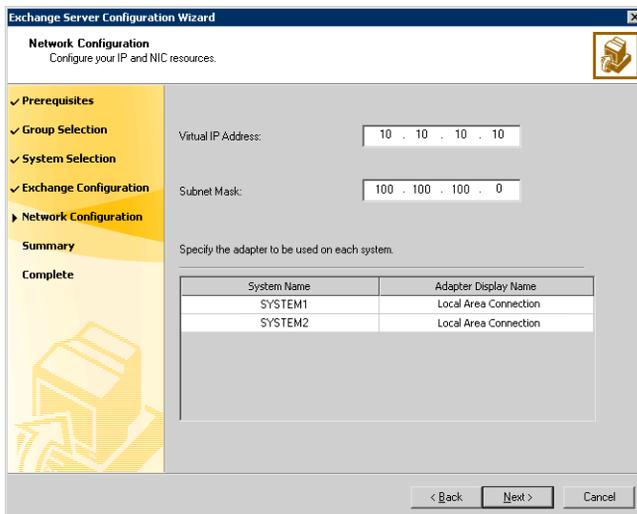


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.
- Specify whether you want to configure the Exchange Management service.
 - Specify whether you want to configure the Exchange MTA service resource. The Exchange Message Transfer Agent is specific to legacy Exchange message routing based on X.400. There are specific circumstances where it may or may not be required for your Exchange server. Please refer to Microsoft documentation on Exchange message routing and the use of MTA for further clarification and applicability to its use within your Exchange environment.

- If you are utilizing Symantec Mail Security for Exchange be sure to indicate that you would like to configure this resource.
 - Click **Next**.
- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



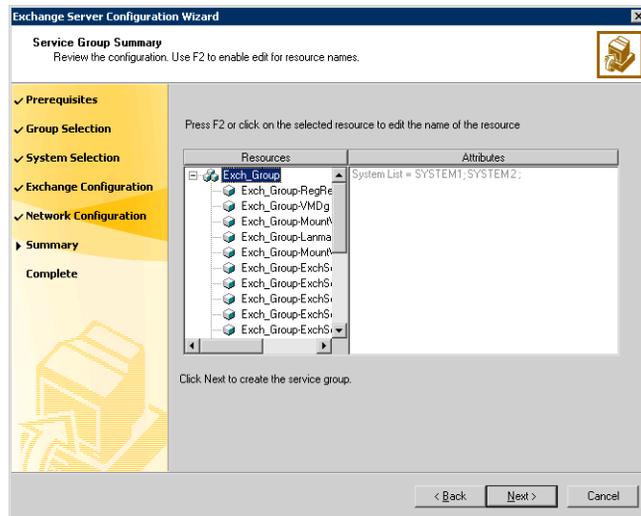
- 7 Specify information related to the network:



- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
 If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

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

- Click **Next**.
- 8 Review the service group configuration and change the resource names, if desired:



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 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.

- Click **Next**.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 9 In the Completing the Exchange Configuration dialog box, select the **Bring the service group online** check box to bring the service group online on the local system.
 - 10 Click **Finish**. After bringing the service group online, you must run the Exchange System Manager so that all the stores are automatically mounted on start-up.

To reconfigure mounting of stores at start-up

- 1 Start Exchange System Manager.
- 2 In the left pane, navigate to your storage group.
If administrative groups are not configured, Servers > Exchange Server > Storage Group.
If more than one administrative group is configured, expand Administrative Groups > Your Administrative Group > Servers > Exchange Server > Storage Group.
- 3 Right-click the Exchange database and choose **Properties** from the pop-up menu.
- 4 Click the Database tab.
- 5 Clear the **Do not mount this store at start-up** check box.
- 6 Click **OK**.

Repeat these steps for all the Exchange databases that were previously mounted.

If you need to configure additional storage groups or mailbox stores on the shared storage you should do that now. Import disk groups and mount volumes that have been created for the additional storage groups or mailbox stores, and create the new storage groups and mailbox stores in Exchange System Manager, and then rerun the Exchange Configuration Wizard to bring them under VCS control. If you already designated the additional mounted volumes and disk groups when you ran the configuration wizard the first time then you can just create the storage groups and mailbox stores in Exchange System Manager.

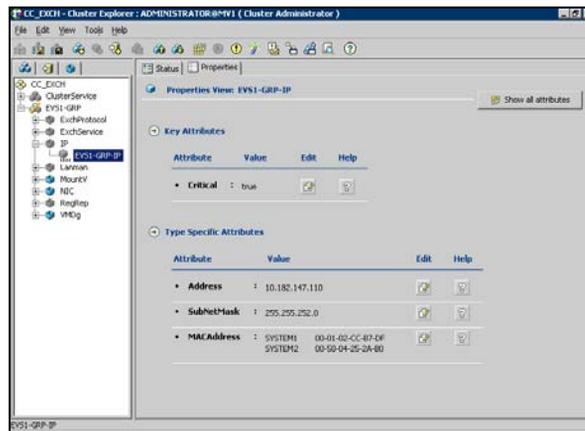
Modifying the IP resource in the Exchange service group

Note: This procedure is only applicable to a campus cluster with sites in different subnets.

Use the Java Console to modify the Address and SubNetMask attributes of the IP resource in the Exchange service group.

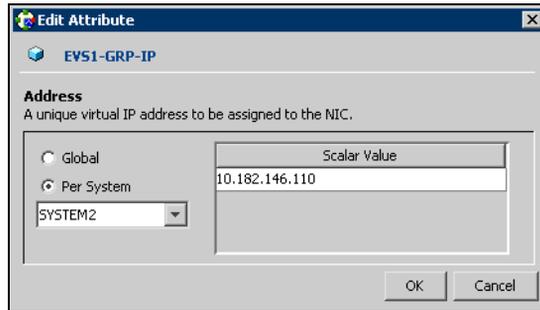
To modify the IP resource

- 1 From the Cluster Explorer configuration tree, select the IP resource (EVS1-GRP-IP) in the Exchange service group (EVS1-GRP).

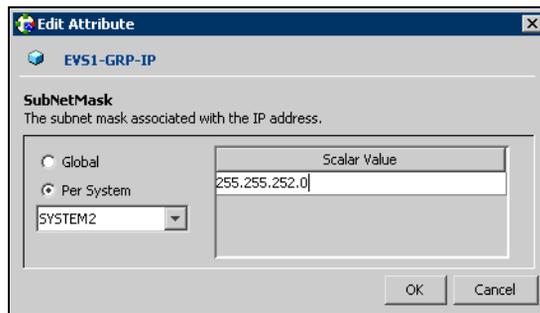


- 2 In the Properties View, click the **Edit** icon for the **Address** attribute.

- 3 In the Edit Attribute dialog box:



- a Select the **Per System** option.
 - b Select the system at Site B.
 - c Enter the virtual IP address at Site B.
 - d Click **OK**.
- 4 In the Properties View, click the **Edit** icon for the **SubNetMask** attribute.
 - 5 In the Edit Attribute dialog box:



- a Select the **Per System** option.
 - b Select the system at Site B.
 - c Enter the subnet mask at Site B.
 - d Click **OK**.
- 6 From the **File** menu of Cluster Explorer, click **Close Configuration**.

Verifying the campus cluster: Switching the service group

To verify the campus cluster is functioning properly

- 1 Bring the service group online on one node:
 - a In the Cluster Explorer configuration tree, right-click the service group.
 - b Click **Online**, and click the appropriate system from the menu.
- 2 Switch the service group to the other node:
 - a In the Cluster Explorer configuration tree, right-click the service group.
 - b Click **Switch To**, and click the appropriate system from the menu.

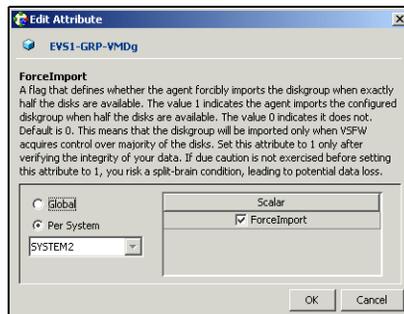
Possible tasks after creating the campus cluster

Setting the ForceImport attribute to 1 after a site failure

You must set the ForceImport attribute for the VMDg resource to 1 after a site failure to ensure proper failover. Refer to [Chapter 8, “Campus cluster for Exchange: Overview” on page 321](#), for a complete review of campus cluster failover using the ForceImport attribute.

To set the ForceImport attribute to 1

- 1 From the Cluster Explorer configuration tree, select the VMDg resource (EVS1-GRP-VMDg) in the Exchange service group (EVS1-GRP).
- 2 In the Properties View, click the **Edit** icon for the **ForceImport** attribute.
- 3 In the Edit Attribute dialog box:



- a Select the **Per System** option.
 - b Select the system in Site B.
 - c Select the **ForceImport** check box.
 - d Click **OK**.
- 4 From the **File** menu of Cluster Explorer, click **Close Configuration**.
- 5 After the failover takes place, revert the ForceImport attribute to its original value.

Disaster Recovery

This section includes the following chapters:

- [Chapter 10, “Disaster Recovery for Exchange: Overview” on page 391](#)
- [Chapter 11, “Deploying Disaster Recovery using the DR wizard for cloning: New Exchange Server installation” on page 393](#)
- [Chapter 12, “Deploying SFW HA for Disaster Recovery: Configuring any-to-any failover” on page 467](#)
- [Chapter 13, “Testing fault readiness by running a fire drill” on page 501](#)

Disaster Recovery for Exchange: Overview

This chapter includes the following topics:

- [What is a disaster recovery solution?](#)
- [Why implement a DR solution?](#)
- [Typical DR configurations for Exchange](#)

What is a disaster recovery solution?

A disaster recovery (DR) solution is a series of procedures you can use to safely and efficiently restore application data and services in the event of a catastrophic failure. A typical DR solution requires clusters on *primary* and *secondary* sites with replication between those sites. The cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails.

Why implement a DR 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 can fail over to a site hundreds or thousands of miles away.

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 or Backup Exec product documentation for information on configuring backup.

Typical DR configurations for Exchange

Note: Refer to the chapters in this section for details on active/passive and any-to-any configurations.

The DR chapters of this guide cover the following configurations:

- Using an active/passive configuration, create a new SFW HA environment with DR capabilities for Exchange on primary and secondary sites. This section provides guidelines on how to create a SFW HA environment with DR capabilities if SFW is already installed on the primary site.
- Using an active/passive configuration, integrate a standalone Exchange server into a new SFW HA environment with DR capabilities for Exchange on primary and secondary sites.
- Using an any-to-any configuration, create a new SFW HA environment with DR capabilities for Exchange, or transform an active/passive DR environment for Exchange into an any-to-any environment, on primary and secondary sites.
- Using an active/passive configuration, upgrade an existing SFW environment on a site to a new SFW HA environment with DR capabilities for Exchange. This section provides guidelines on how to upgrade your environment with or without a standalone Exchange server.

Deploying Disaster Recovery using the DR wizard for cloning: New Exchange Server installation

This chapter covers the following topics:

- [Tasks for a new disaster recovery installation of Microsoft Exchange](#)
- [Before you begin](#)
- [Setting up the secondary site: Installing SFW HA and configuring 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 Exchange on the first node with DR option \(secondary site\)](#)
- [Installing Exchange on additional nodes \(secondary site\)](#)
- [Cloning the service group configuration on to the secondary site using the DR wizard](#)
- [Configuring replication and global clustering](#)

- [Verifying the disaster recovery configuration](#)
- [Establishing secure communication within the global cluster \(optional\)](#)
- [Recovery procedures for service group dependencies](#)
- [Possible task after creating the DR Environment: Adding a new failover node](#)

Tasks for a new disaster recovery installation of Microsoft Exchange

After setting up an SFW HA environment for Exchange on a primary site, you can create a secondary or “failover” site for disaster recovery. This chapter provides information on how to install and configure the high availability and Exchange components on the secondary site, with the intent of creating an identical setup for the Exchange service group on both sites. This environment involves an active/passive configuration with one to one failover capabilities. The identical configuration can be achieved using the Disaster Recovery (DR) wizard. The DR wizard helps you to clone the storage configuration and the service group configuration from the primary site to the secondary site.

Note: If you want to create the identical configuration manually, without cloning the storage configuration or the service group, see [Appendix A, “Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation” on page 515](#).

You can either choose to configure replication using VVR or an agent-supported array-based hardware replication and then use the DR wizard to configure global clustering. If you plan to use array-based hardware replication, then you must first complete configuring global clustering using the DR wizard and then proceed with configuring replication. Irrespective of the method you choose for replication, you must set up Global Clustering to complete the Disaster Recovery configuration.

Caution: To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment, you must first run the wizard before configuring replication.

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

Table 11-1 Task List

Objective	Tasks
“Before you begin” on page 398	✓ Verifying hardware and software prerequisites

Table 11-1 Task List

Objective	Tasks
“Reviewing the configuration” on page 402	<ul style="list-style-type: none"> ✓ Understanding Active/Passive configuration and site failover in a DR environment
“Configuring the storage hardware and network” on page 405	<ul style="list-style-type: none"> ✓ Setting up the network and storage for a cluster environment ✓ Verifying the DNS entries for the systems on which Exchange will be installed
“Setting up the secondary site: Installing SFW HA and configuring a cluster” on page 408	<ul style="list-style-type: none"> ✓ Reviewing the prerequisites ✓ Reviewing the configuration ✓ Configuring the network and storage ✓ Installing SFW HA ✓ Configuring the cluster using the Veritas Cluster Server Configuration Wizard
“Configuring disaster recovery” on page 432	<ul style="list-style-type: none"> ✓ Verifying that Exchange has been configured for high availability at the Primary site
“Cloning the storage on the secondary site using the DR wizard” on page 433	<ul style="list-style-type: none"> ✓ Cloning the storage configuration on the secondary
“Installing Exchange on the first node with DR option (secondary site)” on page 438	<ul style="list-style-type: none"> ✓ Reviewing the prerequisite checklist ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation on the first node using the disaster recovery option ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes in the same virtual server

Table 11-1 Task List

Objective	Tasks
“Cloning the service group configuration on to the secondary site using the DR wizard” on page 449	✓ Cloning the service group configuration from the primary to the secondary site using the DR wizard
“Configuring replication and global clustering” on page 451	✓ Configuring VVR components and global clustering using the DR wizard
“Verifying the disaster recovery configuration” on page 457	✓ Verifying the disaster recovery configuration
“Establishing secure communication within the global cluster (optional)” on page 458	✓ Adding secure communication between local clusters within the global cluster (optional task)
“Recovery procedures for service group dependencies” on page 460	✓ Reviewing actions required for disaster recovery if there are service group dependencies
“Possible task after creating the DR Environment: Adding a new failover node” on page 464	✓ Completing required tasks when adding a new failover system to either the primary or 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 11-2](#) estimates disk space requirements for SFW HA.

Table 11-2 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 Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:
 - Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition (SP2 required)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
or
Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP1 required for all editions)
or
Windows Server 2003 R2 (Standard Edition, Enterprise Edition, or Datacenter Edition)
or
 - Microsoft Exchange 2000 Server or Microsoft Exchange 2000 Enterprise Server (SP3 with August 2004 rollup patch required for both editions)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)

Note: Microsoft support for Exchange Server 2003 is limited to 32-bit versions of the Windows 2003 operating system.

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 402.
- 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 Exchange Virtual Server (EVS)
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

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

Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

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

- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Domain Servers group.
- 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.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- You must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.
- 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.

- You must install IIS, SMTP, NNTP, ASP.NET, and WWW services before you install Microsoft Exchange Server.

Best practices

Symantec recommends that you perform the following tasks:

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

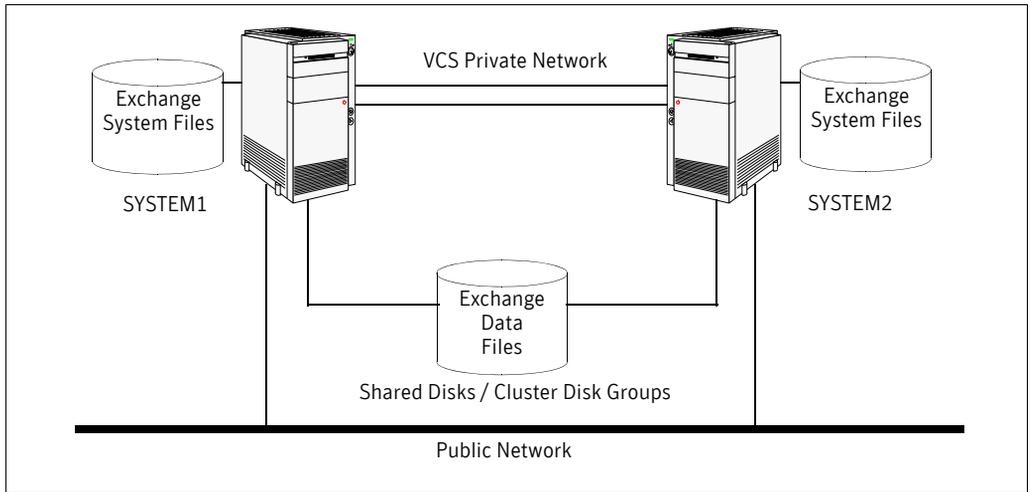
Reviewing the configuration

This overview highlights the high availability within a cluster, and the disaster recovery between two sites.

In an active/passive configuration with one to one failover capabilities, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. If you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM4 and SYSTEM5 on the secondary site), EVS1 can fail over from SYSTEM1 to SYSTEM2 or vice versa on the primary site, and SYSTEM4 to SYSTEM5 or vice versa on the secondary site.

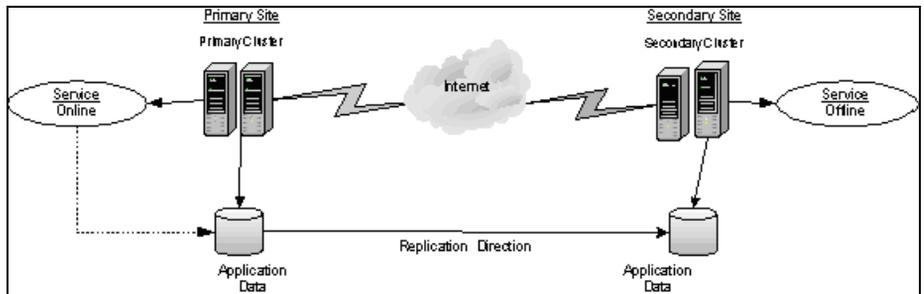
[Figure 11-1](#) provides a view of a cluster configuration on the primary site:

Figure 11-1 Cluster configuration on the primary site



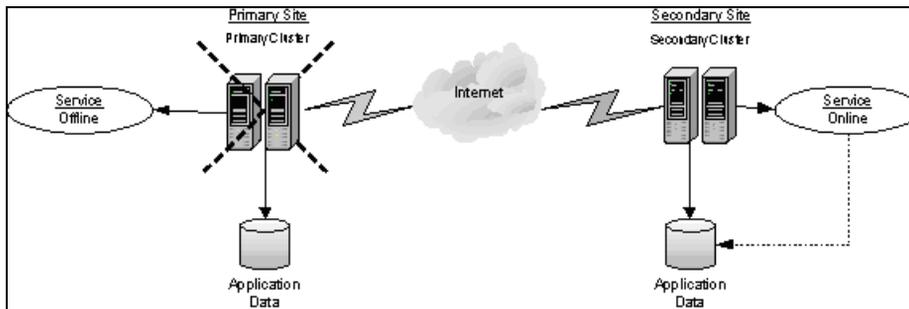
In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. [Figure 11-2](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 11-2 Disaster Recovery environment



When a failure occurs at the primary site, the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients. [Figure 11-3](#) illustrates this type of failure:

Figure 11-3 Application services restored after primary site failure



Supported disaster recovery configurations for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

Service group dependency configurations are described in detail in the VCS documentation.

See *Veritas Cluster Server Administrator's Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

The Disaster Recovery wizard supports only one level of dependency (one child). If you need to configure more levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

The wizard clones dependent service groups as global groups.

Configuring the storage hardware and network

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk.
Use Windows Disk Management (**Start > Control Panel > Administrative Tools > Computer Management**) or Veritas Enterprise Administrator (VEA) on each system to verify that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Right-click on Network Connections and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 From the status window, click **Properties**.

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

Managing disk groups and volumes

This section includes the following procedures:

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

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

See “[Installing Exchange on the first node with DR option \(secondary site\)](#)” on page 438.

Preparing the forest and domain

Microsoft requires the preparation of the forest and domain prior to an Exchange installation. See Microsoft documentation for instructions for preparing the forest and domain for an Exchange installation. Do not repeat this process for additional Exchange installations.

Setting up the secondary site: Installing SFW HA and configuring a cluster

After completing the configuration on the primary site, repeat the appropriate tasks to complete the SFW HA installation at the secondary site.

Because the Disaster Recovery wizard is capable of cloning the storage, you only need to complete configuring SFW HA at the secondary site. The storage configuration will be handled by the DR wizard. Begin with reviewing the requirements on the secondary site, similar to the primary site and continue with the procedures given below.

- ✓ Reviewing the requirements
See “[Before you begin](#)” on page 398.

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 Application Agent for Exchange. 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 11-3](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 11-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed

Table 11-3 Installation behavior with unsigned drivers (continued)

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

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

To change the driver signing options on each system

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

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

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

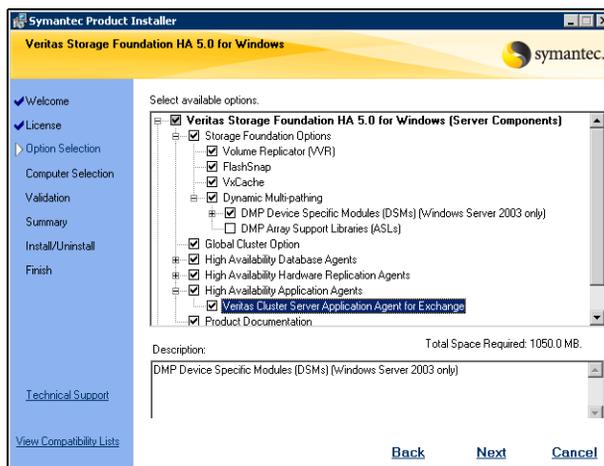
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.

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



- 4 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 8 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

9 Select the appropriate SFW product options for your installation. Click **Next**.



The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

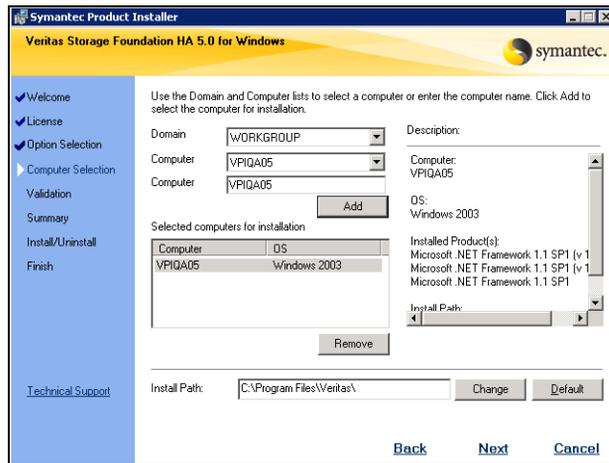
Veritas Cluster Server Application Agent for Exchange Required to configure high availability for Exchange Server.

Client Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console. Required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option Required for a disaster recovery configuration only.

Veritas Volume Replicator If you plan to use VVR for replication, you must also select the option to install VVR.

10 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

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

Computer

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

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

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

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

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
 - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
 - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

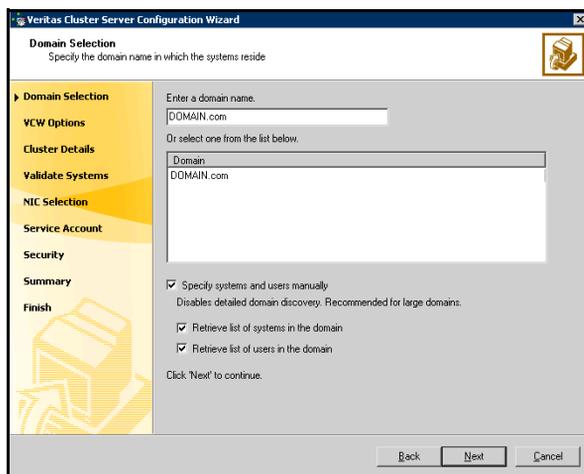
- ✓ Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- ✓ Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative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

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

Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.

To configure a VCS cluster

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



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

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

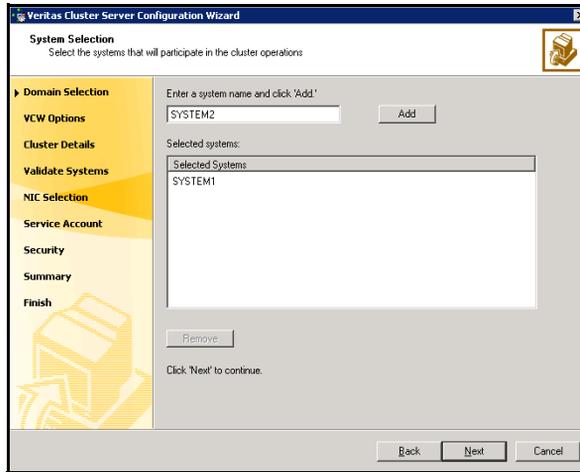
Proceed to [step 7](#) on page 417.

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**.

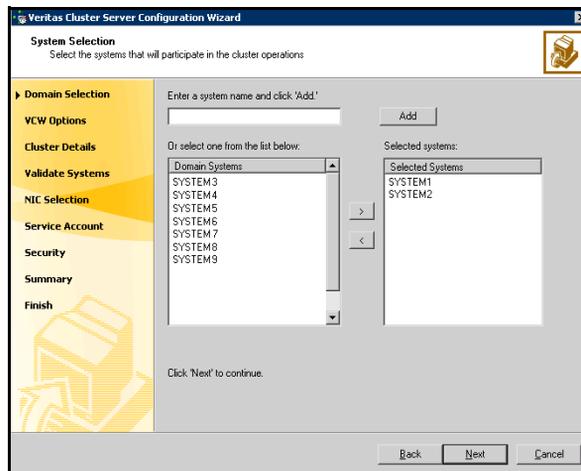
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 416. Otherwise, proceed to the next step.

- 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.



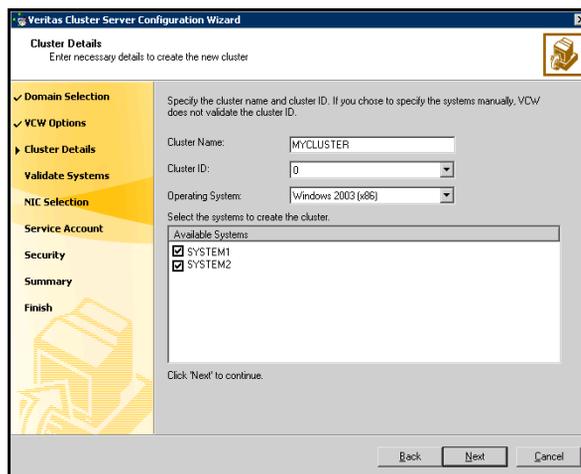
Proceed to [step 7](#) on page 417.

- 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.

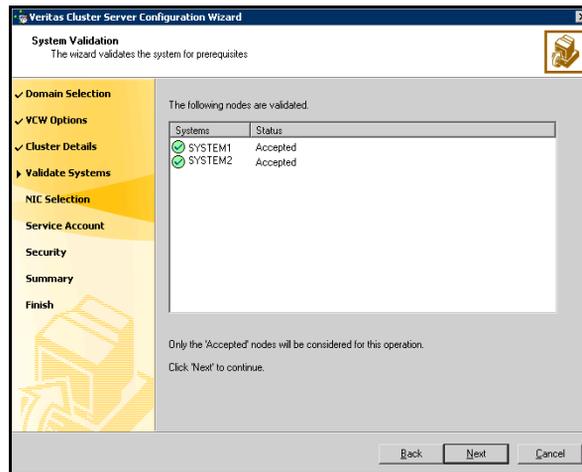
Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

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.

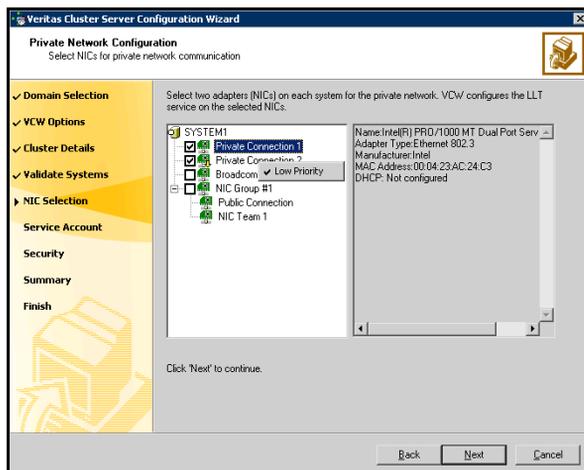
- 9 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.



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 417, proceed to the next step. Otherwise, proceed to [step 11](#) on page 419.

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

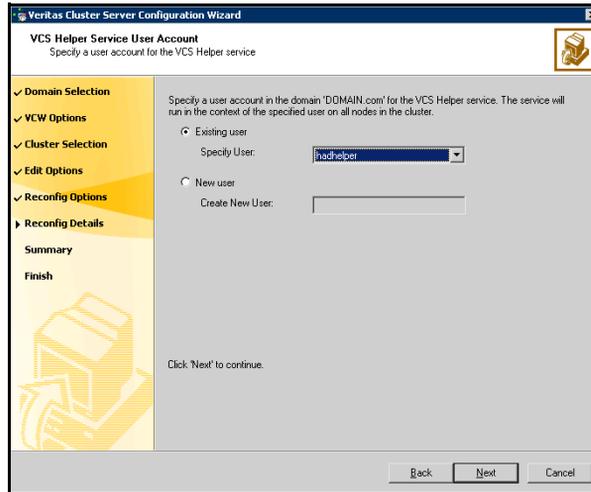


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs 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 415, 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**.

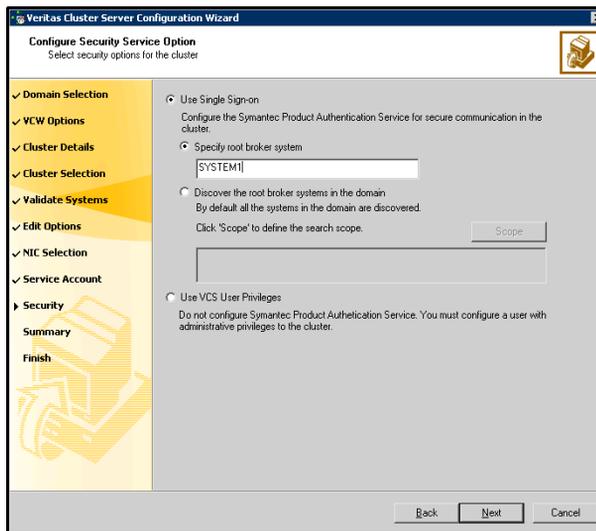
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- 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.

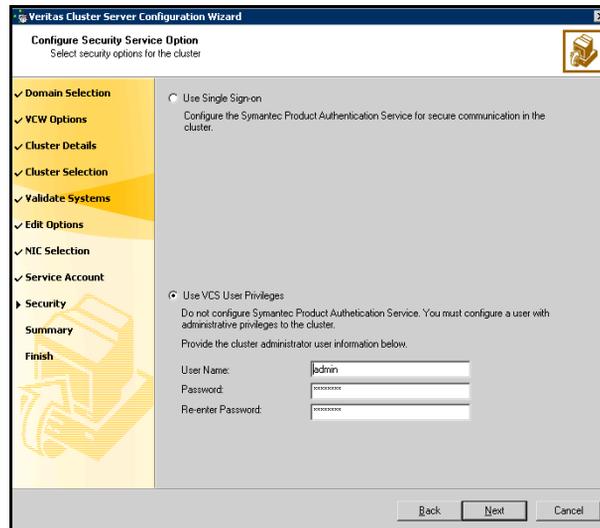
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- 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.

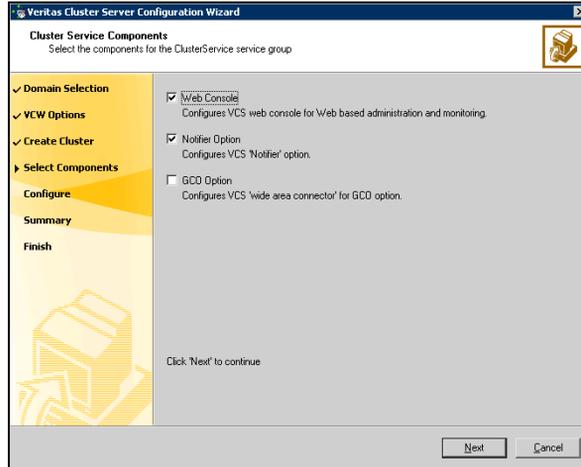
- Click **Next**.

- 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**.



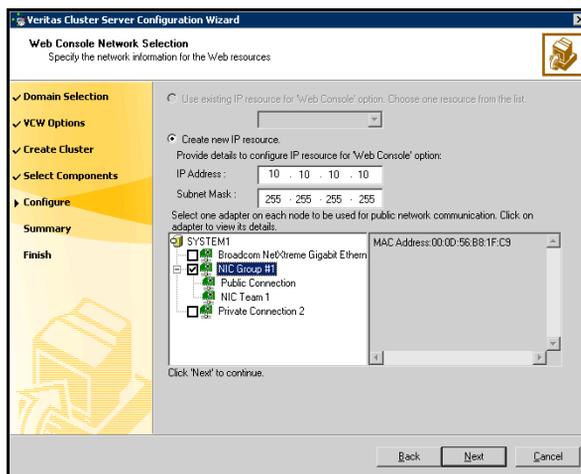
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 425.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 426.

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**.



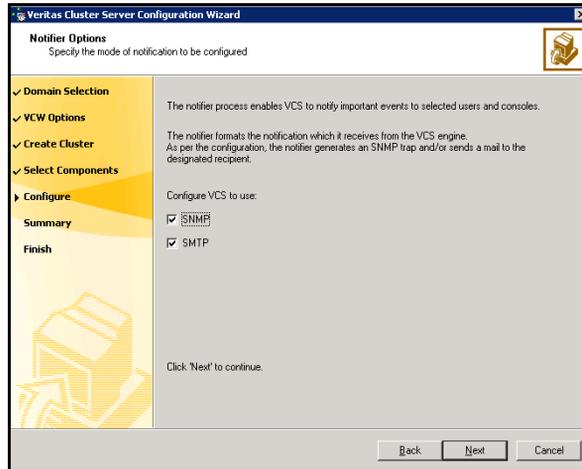
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 426. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

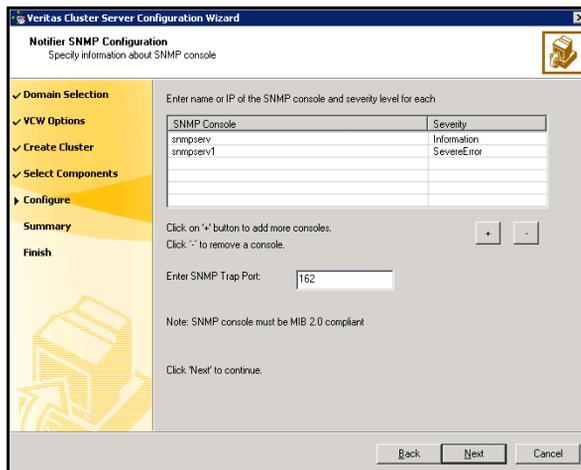
To configure notification

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



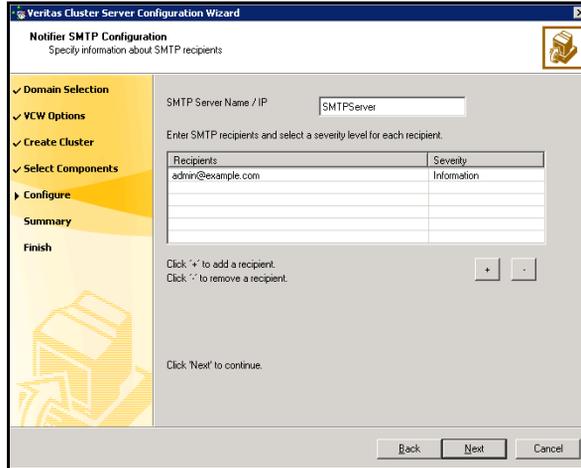
You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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



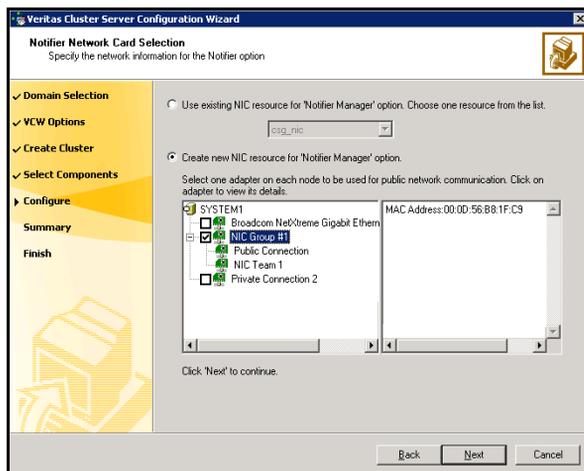
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.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

Make sure that Exchange has been configured for high availability at the primary site. If you have not yet configured Exchange for High Availability at the primary site, go to High Availability (HA) Configuration section and follow the steps in the order shown.

See “[Deploying SFW HA for high availability: New installation](#)”.

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

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server 2003 environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.

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

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
----------------------------------	--

Password	Specify a password in the 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 service group configuration from the primary site to the secondary site. It also configures VVR replication settings and connects the clusters into a global cluster. Although all the tasks can be performed using this single wizard, you will need to exit the wizard after cloning the storage to install Exchange. 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, clicking **Next** takes you to the start page of the process following the one that you had last completed.

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 Exchange 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 Exchange service group and any dependent service groups except for the RVG service group.
`hauser -add user [-priv <Administrator|Operator> [-group service_groups]]`
- 4 Reset the configuration to read-only:
`haconf -dump -makero`

To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Reset the configuration to read-only:
`haconf -dump -makero`

Cloning the storage on the secondary site using the DR wizard

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at secondary site must have adequate free storage. If you have created the

configuration but there is a mismatch in the volume sizes, then the wizard can correct this and then complete the configuration.

To clone the storage configuration from the primary site to the secondary site

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.

From the **Solutions Configurations Center** expand the Solutions for Microsoft Exchange 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 Exchange virtual Server 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	<p>Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary.</p> <ul style="list-style-type: none"> ■ 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 Exchange Installation panel, do one of the following:
 - Click **Finish** to exit the wizard and proceed with installing the Exchange 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 Exchange 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 Exchange 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 Exchange on the first node with DR option (secondary site)

Installing Exchange on the first node for the EVS includes procedures for three stages that involve pre-installation, installation, and post-installation procedures. In this procedure, virtual Exchange server EVS1 will fail over from SYSTEM4 to SYSTEM5.

See the following topics:

- [“Prerequisites for installing Exchange Server”](#) on page 438
- [“Exchange pre-installation on first node \(secondary site\)”](#) on page 440
- [“Exchange installation on first node \(secondary site\)”](#) on page 442
- [“Exchange post-installation on first node \(secondary site\)”](#) on page 443

Prerequisites for installing Exchange Server

Complete the following tasks before installing Exchange Server:

- ✓ Verify the disk group is imported on the first node of the cluster.
See [“Managing disk groups and volumes”](#) on page 406.
- ✓ Mount the volume containing the information for registry replication.
- ✓ Verify that all systems on which Exchange Server is to be installed have IIS installed; the SMTP, NNTP, and WWW services must be installed on all systems. For installing Exchange on Windows 2003, ASP.NET service must also be installed.
- ✓ Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange binaries to be installed at the same location on all the nodes. For example, if you install Exchange at `c:\Program Files\exchsrvr` on one node, then you must install Exchange at `c:\Program Files\exchsrvr` on all the other nodes.
- ✓ Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.

- If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- ✓ Make sure to use the same drive letters employed on the primary site.
- ✓ Make sure to take the Exchange service group offline on the primary site; otherwise, the installation on the secondary site will not function properly.

- ✓ Verify that you have completed the forest and domain preparation.
See “[Preparing the forest and domain](#)” on page 45

Exchange pre-installation on first node (secondary site)

Perform the Exchange pre-installation procedure.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for Exchange disaster recovery setup** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.
- 7 Specify information related to the network.

Exchange Setup Wizard for Veritas Cluster Server

Exchange Virtual Server Details
Specify the details for the Exchange Virtual Server.

Pre Installation

- ✓ System Validation
- Exchange Server Details
- Drive Selection
- Summary of Selection
- Tasks Execution
- System Reboot

MS Exchange Installation

Post Installation

- Move Exchange Database
- Service Group Creation

Exchange Virtual Server Name: EVS1

Domain Suffix: DOMAIN.COM

Public Network Adapter: Public

Virtual IP Address: 10 . 182 . 149 . 138

Subnet Mask: 255 . 255 . 248 . 0

Back Next Cancel

- Enter a unique virtual name for the Exchange server.

Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.

- Enter a domain suffix for the virtual server.
- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
- Enter a unique virtual IP address for the Exchange virtual server.
- Enter the subnet mask for the virtual IP address.
- Click **Next**.

The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met. Click **Next**.

- 8 Select a drive where the registry replication data will be stored and click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation on first node (secondary site)

Install Exchange on the same node selected in “[Exchange pre-installation on first node \(secondary site\)](#)” on page 440.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

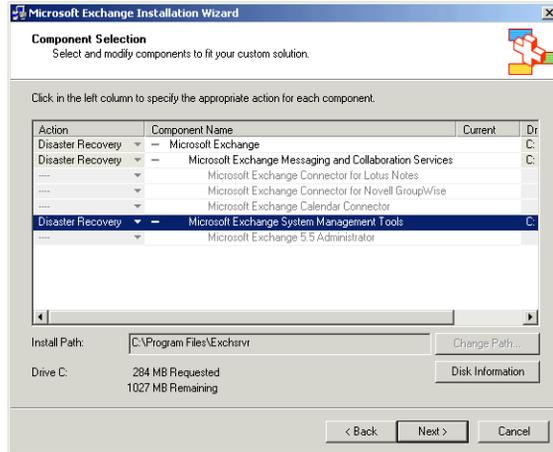
To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe  
/disasterrecovery
```

where **<drive letter>** is the location where the Exchange software is located.
- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools

components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

Exchange post-installation on first node (secondary site)

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

Installing Exchange on additional nodes (secondary site)

Install Exchange on additional nodes in the cluster for the same Exchange virtual server (EVS1). You must run pre-installation, installation, and post-installation procedures for each additional node.

Note: In an any-to-any configuration, the steps for installing Exchange on the additional nodes (failover nodes) can be completed for the first Exchange server, and do not need to be repeated for the common failover nodes for additional Exchange servers if the common failover nodes were already installed with Exchange under the context of the first Exchange server.

Make sure to complete the following tasks before the Exchange installation:

- ✓ Review the prerequisites for permissions.
See “[Installing Exchange on the first node with DR option \(secondary site\)](#)” on page 438.
- ✓ Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes.
See “[Managing disk groups and volumes](#)” on page 406.

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

Exchange pre-installation: Additional nodes

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server.
The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically generated when setting up an additional failover mode for the virtual server in the same cluster.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.

11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.

12 Click **Reboot**.

The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: Additional nodes

Install Exchange on the same node selected in “[Installing Exchange on additional nodes \(secondary site\)](#)” on page 444.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

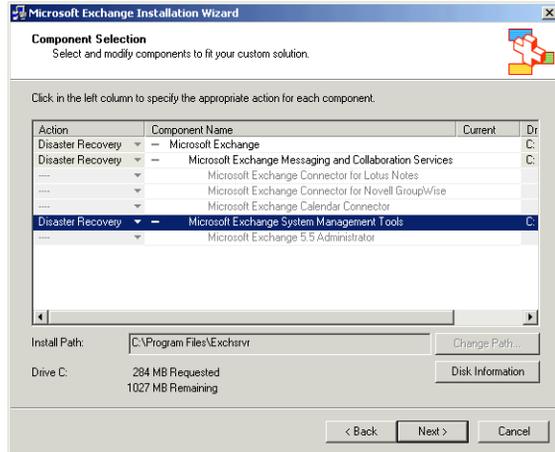
1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe  
/disasterrecovery
```

where <drive letter> is the location where the Exchange software is located.

2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools

components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

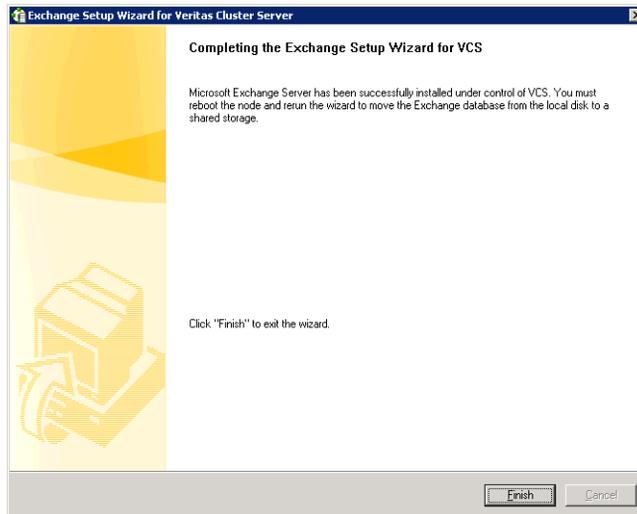
Exchange post-installation: Additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.
- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

If you wish to add the nodes later, use the Exchange service group configuration wizard. See “[Possible task after creating the DR Environment: Adding a new failover node](#)” on page 464 for instructions.

Cloning the service group configuration on to the secondary site using the DR wizard

Prior to cloning the service group on the secondary site verify that you have installed the application on the Primary and created the required service groups. You will also need to ensure that you have installed the application on the Secondary. After verifying launch the DR wizard.

To 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 Exchange 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: <ul style="list-style-type: none"> ■ 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.

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 Exchange 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	<p>Displays the list of available volumes that have not been selected to be a part of the RVG.</p> <p>Either double-click on the volume name or use the > option to move the volumes into the Selected RVG Volumes pane.</p>
Selected RVG Volumes	<p>Displays the list of volumes that have been selected to be a part of the RVG.</p> <p>To remove a selected volume, either double-click the volume name or use the < option to move the volumes into the Available Volumes pane.</p>
Primary SRL	<p>Select the appropriate primary Replicator Log volume from the drop-down menu.</p> <p>Size: Enter an appropriate log size value in the corresponding Size field.</p> <p>If you did not create the Replicator Log volume earlier, click Create New on the drop-down menu.</p>
Secondary SRL	<p>Select the appropriate secondary Replicator Log volume from the drop-down menu.</p> <p>Size: Enter an appropriate size value in the corresponding Size field.</p> <p>If you did not create the Replicator Log volume earlier, click Create New on the drop-down menu.</p>
Start Replication after the wizard completes	<p>Select this check box to start replication automatically after the wizard completes the necessary configurations.</p>
<p>■ 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:</p>	
Replication Mode	<p>Select the required mode of replication; Synchronous, Asynchronous, or Synchronous Override. The default is synchronous override.</p>
Log Protection	<p>Select the appropriate log protection from the list.</p> <p>The Off option disables Replicator Log Overflow protection.</p>

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, Replicator Log 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 VvrRvg 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 404.

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 11-4 Online, local, soft dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent remains online on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

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 11-5 Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	Secondary site: Bring the service groups online in the appropriate order (child first, then parent). Leave the RVG group online at the primary site.
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

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 11-6 Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	Secondary site: Bring the service groups online in the appropriate order (child first, then parent). Do not take the RVG group offline at the primary site.

Table 11-6 Online, local, hard dependency link (continued)

Failure condition	Results	Action required
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Possible task after creating the DR Environment: Adding a new failover node

The following procedures describe how to add an additional node to the cluster at either the primary or secondary site after your disaster recovery environment is in operation. The clusters at each site are not required to have the same number of nodes or the same failover configuration.

See the following topics:

- [“Preparing the new node”](#) on page 464
- [“Preparing the existing DR environment”](#) on page 464
- [“Installing Exchange on the new node”](#) on page 465
- [“Modifying the replication and Exchange service groups”](#) on page 465
- [“Reversing replication direction”](#) on page 466

Preparing the new node

Install SFW HA on the new system and then add the system to the cluster.

To install SFW HA and add the system to the cluster

- 1 Refer to [“Installing SFW HA”](#) on page 408 for installation instructions.
- 2 Use the **Cluster Operations** option of the VCS Configuration wizard (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**) to add the new system to the cluster. If necessary, refer to the *Veritas Cluster Server Administrator’s Guide* for information on this procedure.

Preparing the existing DR environment

If you plan to add a failover node to the secondary site, you must temporarily switch the roles of the Primary and Secondary sites so that the current site becomes the Primary. This action reverses the direction of replication.

To prepare the existing DR environment

- 1 If you adding the failover node to the cluster at the primary site, proceed directly to [step 2](#). If you are adding a failover node to the secondary site, you must switch the roles of the primary and secondary sites. This action reverses the direction of replication.
 - a In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.

- b Click **Switch To**, and click **Remote switch**.
- c In the **Switch global group** dialog box:
 - Click the cluster at the secondary site you want to switch the group to.
 - Click the specific system where you want to bring the global Exchange service group online.
 - Click **OK**.
- 2 Take the global Exchange service group offline at the current primary site.
- 3 Take the VVR replication service group offline.

Installing Exchange on the new node

Install Exchange on the new node, but do not add the node to the service group SystemList.

To prepare the node and install Exchange

- 1 Import the disk group on the new node. Follow the procedure described in [“Managing disk groups and volumes”](#) on page 406.
- 2 From the VEA navigation tree, right-click the RVG for the primary site, and click **Enable Data Access**.
- 3 Run the pre-installation, installation, and post-installation steps described in [“Installing Exchange on additional nodes \(secondary site\)”](#) on page 444; reboot when prompted in these procedures.

Note: During the last step of the post-installation wizard, do *not* check the check box to add the node to the SystemList.

Modifying the replication and Exchange service groups

Add the new failover node to the system lists in the Replication and Exchange service groups.

To add the failover node to the system lists

- 1 Bring the replication service group online on an existing cluster node of the current primary site.
- 2 Bring the MountV resources of the corresponding Exchange service group online on the same node.
- 3 Use the **Modify an existing replication service group** option of the Volume Replicator Agent Configuration Wizard (**Start > All Programs > Veritas > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard**) to

add the new node to the system list for the replication service group. If necessary, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for information on this procedure.

- 4 Use the **Modify service group** option of the Exchange Server Configuration Wizard (**Start > All Programs > Veritas > Veritas Cluster Server > Application Agent for Exchange > Configuration Wizard**) to add the new node to the system list for the Exchange service group. Check the check box to bring the service group online after the wizard completes. If necessary, refer to the *Veritas Cluster Server Administrator's Guide* for information on this procedure.
- 5 After bringing the Exchange service group online, you must use Exchange System Manager to configure all the database stores to automatically mount on start-up.

Reversing replication direction

If you added a failover node at the original secondary site and migrated the RVG in [“Preparing the existing DR environment”](#) on page 464, move the global Exchange service group back to the original primary site and reverse the direction of replication. These actions switch the Primary and Secondary sites back to their original roles.

To reverse the replication direction

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 In the **Switch global group** dialog box:
 - a Click the cluster to switch the group to.
 - b Click the specific system where you want to bring the global Exchange service group online.
 - c Click **OK**.

Deploying SFW HA for Disaster Recovery: Configuring any-to-any failover

This chapter covers the following topics:

- [“Tasks for deploying a new disaster recovery any-to-any configuration”](#) on page 468
- [“Reviewing the configuration”](#) on page 470
- [“Configuring disaster recovery for the first Exchange virtual server”](#) on page 474
- [“Verifying your primary site configuration for an additional Exchange virtual server”](#) on page 475
- [“Adding the user to the service group \(secure clusters only\)”](#) on page 475
- [“Configuring disaster recovery for the second Exchange virtual server”](#) on page 476
- [“Cloning the storage on the secondary site using the DR wizard”](#) on page 476
- [“Installing Exchange on the first node of an additional EVS \(secondary site\)”](#) on page 480
- [“Specifying a common node for failover”](#) on page 486
- [“Cloning the service group configuration on to the secondary site using the DR wizard”](#) on page 488

- [“Configuring replication and global clustering”](#) on page 491
- [“Verifying the disaster recovery configuration”](#) on page 497
- [“Establishing secure communication within the global cluster \(optional\)”](#) on page 498
- [“Possible tasks after creating the DR environment”](#) on page 500

Tasks for deploying a new disaster recovery any-to-any configuration

You can configure an any-to-any SFW HA environment for Exchange to provide a production node with multiple failover nodes.

See [Chapter 6, “Deploying SFW HA for high availability: Configuring a new any-to-any failover”](#) on page 177 .

See [Chapter 7, “Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover”](#) on page 259.

After setting up the any-to-any SFW HA environment for Exchange on a primary site, you can create a secondary or “failover” site for disaster recovery. You create an identical any-to-any configuration for the Exchange service groups on the secondary site.

The identical configuration can be achieved using the Disaster Recovery (DR) wizard. The DR wizard helps you to clone the storage configuration and the service group configuration from the primary site to the secondary site.

This chapter provides information on configuring the secondary site for disaster recovery for an any-to-any configuration using the Disaster Recovery wizard.

After creating the identical configuration on both sites, you must set up replication and Global Clustering. You can choose to configure replication using VVR or an agent-supported array-based hardware replication. If you plan to use array-based hardware replication, then you must first complete configuring global clustering using the DR wizard before you 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.

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

Table 12-1 Tasks for deploying a new any-to-any DR configuration

Objective	Tasks
“ Reviewing the configuration ” on page 470	<ul style="list-style-type: none"> ✓ Understanding any-to-any configuration and site failover in a DR environment
“ Configuring disaster recovery for the first Exchange virtual server ” on page 474	<ul style="list-style-type: none"> ✓ Completing all steps for deploying disaster recovery for the first Exchange virtual server
“ Verifying your primary site configuration for an additional Exchange virtual server ” on page 475	<ul style="list-style-type: none"> ✓ Verifying that the second Exchange virtual server has been configured for high availability at the primary site
“ Adding the user to the service group (secure clusters only) ” on page 475	<ul style="list-style-type: none"> ✓ For a secure cluster only, add the user to the second Exchange service group
“ Cloning the storage on the secondary site using the DR wizard ” on page 476	<ul style="list-style-type: none"> ✓ Cloning the storage configuration on the secondary site
“ Installing Exchange on the first node of an additional EVS (secondary site) ” on page 480	<ul style="list-style-type: none"> ✓ Reviewing the prerequisite checklist ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation on the first node using the disaster recovery option ✓ Specifying a common failover node
“ Cloning the service group configuration on to the secondary site using the DR wizard ” on page 488	<ul style="list-style-type: none"> ✓ Cloning the service group configuration from the primary to the secondary site using the DR wizard

Table 12-1 Tasks for deploying a new any-to-any DR configuration

Objective	Tasks
“Configuring replication and global clustering” on page 491	✓ Configuring VVR components and global clustering using the DR wizard
“Verifying the disaster recovery configuration” on page 497	✓ Verifying the disaster recovery configuration
“Establishing secure communication within the global cluster (optional)” on page 498	✓ Adding secure communication between local clusters within the global cluster (optional task)
“Possible tasks after creating the DR environment” on page 500	✓ Reviewing actions required for disaster recovery if there are service group dependencies
	✓ Adding a new failover node to a local cluster

Reviewing the configuration

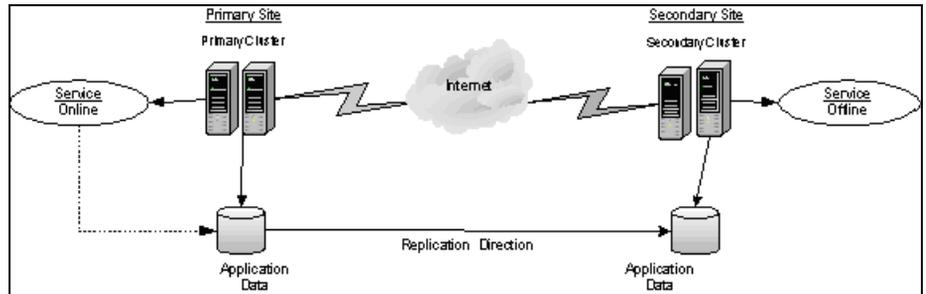
Before configuring disaster recovery for an any-to-any failover configuration, review the following topics:

- [“Disaster recovery configuration”](#) on page 470
- [“Any-to-any configuration”](#) on page 472
- [“Sample any-to-any configuration for disaster recovery”](#) on page 473

Disaster recovery configuration

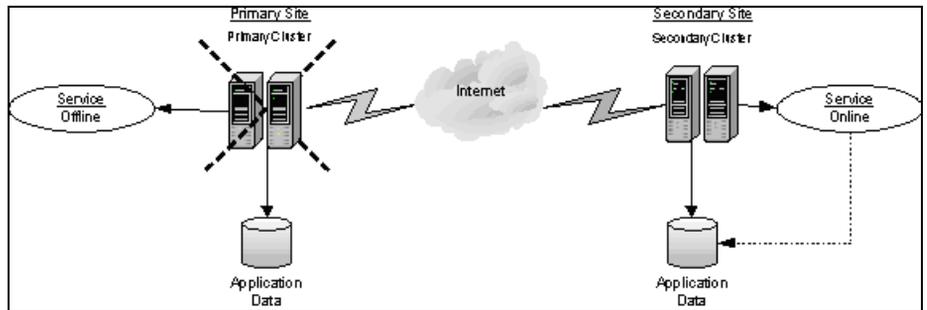
In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. [Figure 12-1](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 12-1 Disaster Recovery Environment



When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients. Figure 12-2 illustrates this type of failure:

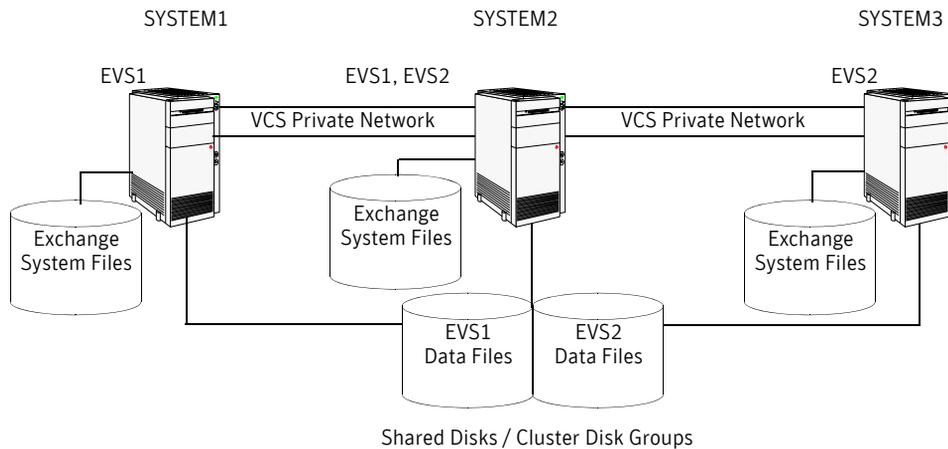
Figure 12-2 Application services restored after primary site failure



Any-to-any configuration

In an any-to-any configuration, each Exchange virtual server in the cluster is configured in a separate service group. Each service group can fail over to any configured node in the cluster, provided that no other Exchange virtual server is online on that node. The SFW HA software ensures that an Exchange service group does not fail over to a node on which another Exchange service group is online. [Figure 12-3](#) show an example of an any-to-any configuration.

Figure 12-3 Any-to-any configuration



For example, consider a three-node cluster hosting two Exchange Virtual Servers, EVS1 and EVS2. The virtual servers are configured in VCS in two service groups such that SYSTEM1 and SYSTEM2 host the EVS1 service group and SYSTEM3 and SYSTEM2 host the EVS2 service group. If SYSTEM1 fails, the service group containing the EVS1 resources is failed over to SYSTEM2. If SYSTEM3 fails, the service group containing the EVS2 resources is failed over to SYSTEM2.

Note: EVS1 and EVS2 cannot be online at the same time on SYSTEM2.

Likewise, on the secondary site, SYSTEM4 and SYSTEM5 host the EVS1 service group and SYSTEM6 and SYSTEM5 host the EVS2 service group. If SYSTEM4 fails, the service group containing the EVS1 resources is failed over to SYSTEM5. If SYSTEM6 fails, the service group containing the EVS2 resources is failed over to SYSTEM5.

Sample any-to-any configuration for disaster recovery

The following table shows the systems used in a three-node any-to-any disaster recovery configuration.

Table 12-2 Systems in an any-to-any DR configuration

Exchange Virtual Server	Any-to-Any Cluster
EVS1 (Primary Site)	SYSTEM1, SYSTEM2
EVS2 (Primary Site)	SYSTEM2, SYSTEM3
EVS1 (Secondary Site)	SYSTEM4, SYSTEM5
EVS2 (Secondary Site)	SYSTEM5, SYSTEM6

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

Table 12-3 Sample Configuration

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3 (Primary Site)	physical node names
SYSTEM4, SYSTEM5, SYSTEM6 (Secondary Site)	
EVS1, EVS2	Microsoft Exchange Virtual Servers
EVS1_GRP, EVS2_GRP	Microsoft Exchange service groups
EVS1_SG1_DG, EVS2_SG1_DG	cluster disk group names
EVS1_SG1_DB1, EVS2_SG1_DB1	volumes for storing the Microsoft Exchange Server database
EVS1_SG1_LOG, EVS2_SG1_LOG	volumes for storing a Microsoft Exchange Server database log file
EVS1_SG1_REGREP, EVS2_SG1_REGREP	volumes that contain the list of registry keys that must be replicated among cluster systems for the Exchange server

Table 12-3 Sample Configuration

Name	Object
EVS1_SG1_MTA,	volumes for storing Microsoft Exchange Server
EVS2_SG1_MTA	MTA database for the Exchange server

Configuring disaster recovery for the first Exchange virtual server

When you configure disaster recovery for the first Exchange virtual server (EVS1), you follow the same instructions as when you configure disaster recovery for an active/passive configuration with one EVS.

See [Chapter 11, “Deploying Disaster Recovery using the DR wizard for cloning: New Exchange Server installation”](#) on page 393.

However, note the following any-to-any considerations:

- You can prepare for SFW HA installation and do the SFW HA installation on all the nodes, not only the nodes to be used for EVS1. For example, if you plan to install EVS1 on SYSTEM4 and SYSTEM5 and install EVS2 on SYSTEM6, you can install SFW HA on all three systems: SYSTEM4, SYSTEM5, and SYSTEM6.
See [“Installing SFW HA”](#) on page 408.
- When configuring the cluster on the secondary site, include the nodes for both EVS1 and EVS2, for example: SYSTEM4, SYSTEM5, and SYSTEM6.
See [“Configuring the cluster”](#) on page 414.
- If you are using VVR as your replication solution, when configuring the VVR Security Service (VxSAS), include the nodes for both EVS1 and EVS2, for example: SYSTEM4, SYSTEM5, and SYSTEM6.
See [“Setting up security for VVR”](#) on page 430.

Once you have configured EVS1 on the secondary site, you can return to this chapter and continue with the steps for configuring EVS2.

Verifying your primary site configuration for an additional Exchange virtual server

Make sure that the additional Exchange virtual server has been configured for high availability at the primary site.

See “[Configuring another Exchange virtual server for an any-to-any failover](#)” on page 239 in [Chapter 6](#), “[Deploying SFW HA for high availability: Configuring a new any-to-any failover](#)”.

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.

For a secure cluster only, add the user to the service group on the primary site.

See “[Adding the user to the service group \(secure clusters only\)](#)” on page 475.

Adding the user to the service group (secure clusters only)

When you configured the first Exchange Virtual Server, you assigned user privileges to the cluster and modified the attribute of the service group to add the user. For the second Exchange Virtual Server, you only need to do the steps to modify the attribute of the service group to add the user. You do this task at the primary site.

To add the user to the service group at the primary site

- 1 Set the configuration to read/write mode:

```
haconf -makerw
```

- 2 Modify the attribute of the service group to add the user. Specify the Exchange service group and any dependent service groups except for the RVG service group.

```
hauser -add user [-priv <Administrator|Operator> [-group  
service_groups]]
```

- 3 Reset the configuration to read-only:

```
haconf -dump -makero
```

Configuring disaster recovery for the second Exchange virtual server

Configuring disaster recovery for the second Exchange virtual server (EVS2) is similar to configuring disaster recovery for EVS1. However setting up the Exchange failover node is different.

Configuring disaster recovery for EVS2 includes the following tasks:

- [Adding the user to the service group \(secure clusters only\)](#)
- [Cloning the storage on the secondary site using the DR wizard](#)
- [Installing Exchange on the first node of an additional EVS \(secondary site\)](#)
- [Specifying a common node for failover](#)
- [Cloning the service group configuration on to the secondary site using the DR wizard](#)
- [Configuring replication and global clustering](#)
- [Verifying the disaster recovery configuration](#)

Cloning the storage on the secondary site using the DR wizard

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, then the wizard can correct this and then complete the configuration.

To clone the storage configuration from the primary site to the secondary site

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start** > **All Programs** > **Symantec** > **Veritas Cluster Server** > **Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Microsoft Exchange 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 Exchange virtual Server 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	<p>Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary.</p> <ul style="list-style-type: none">■ 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 Exchange Installation panel, do one of the following:
 - Click **Finish** to exit the wizard and proceed with installing the Exchange 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 Exchange 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 Exchange 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 Exchange on the first node of an additional EVS (secondary site)

Installing Exchange on the first node for an additional EVS is described in three stages that involve pre-installation, installation, and post-installation procedures. In this example, the virtual Exchange server (EVS2) will fail over from SYSTEM6 to SYSTEM5. Complete the following tasks before installing Exchange Server:

- ✓ Verify the disk group is imported on the first node of the cluster.
See “[Managing disk groups and volumes](#)” on page 406.
- ✓ Mount the volume containing the information for registry replication.
- ✓ Verify that all systems on which Exchange Server is to be installed have IIS installed; the SMTP, NNTP, and WWW services must be installed on all systems. For installing Exchange on Windows 2003, ASP.NET service must also be installed.

- ✓ Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange binaries to be installed at the same location on all the nodes. For example, if you install Exchange at `c:\Program Files\exchsrvr` on one node, then you must install Exchange at `c:\Program Files\exchsrvr` on all the other nodes.
- ✓ Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.
 - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
 - The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- ✓ Make sure to use the same drive letters employed on the primary site.
- ✓ Make sure to take the Exchange service group offline on the primary site; otherwise, the installation on the secondary site will not function properly.

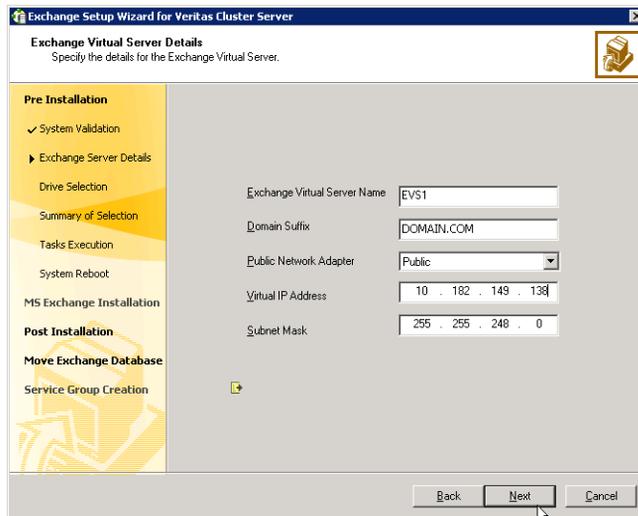
- ✓ Verify that you have completed the forest and domain preparation.
See “[Preparing the forest and domain](#)” on page 45

Exchange pre-installation on first node of an additional EVS (secondary site)

Perform the Exchange pre-installation procedure.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for Exchange disaster recovery setup** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.
- 7 Specify information related to the network.



The screenshot shows the 'Exchange Setup Wizard for Veritas Cluster Server' window. The title bar reads 'Exchange Setup Wizard for Veritas Cluster Server'. The main window has a title 'Exchange Virtual Server Details' and a subtitle 'Specify the details for the Exchange Virtual Server.' On the left, there is a navigation pane with sections: 'Pre Installation' (containing 'System Validation', 'Exchange Server Details', 'Drive Selection', 'Summary of Selection', 'Tasks Execution', 'System Reboot'), 'MS Exchange Installation', 'Post Installation', and 'Move Exchange Database' (containing 'Service Group Creation'). The 'Exchange Server Details' section is expanded. The main area contains the following fields: 'Exchange Virtual Server Name' (text box with 'EVS1'), 'Domain Suffix' (text box with 'DOMAIN.COM'), 'Public Network Adapter' (dropdown menu with 'Public'), 'Virtual IP Address' (text box with '10 . 182 . 149 . 138'), and 'Subnet Mask' (text box with '255 . 255 . 248 . 0'). At the bottom, there are 'Back', 'Next', and 'Cancel' buttons.

- Enter a unique virtual name for the Exchange server.

Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.

- Enter a domain suffix for the virtual server.
- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
- Enter a unique virtual IP address for the Exchange virtual server.
- Enter the subnet mask for the virtual IP address.
- Click **Next**.

The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met. Click **Next**.

- 8 Select a drive where the registry replication data will be stored and click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation on first node of an additional EVS (secondary site)

Install Exchange on the same node selected in “[Exchange pre-installation on first node of an additional EVS \(secondary site\)](#)” on page 482.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

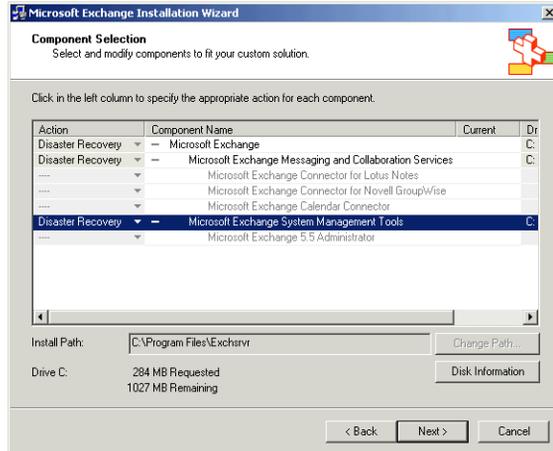
To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe  
/disasterrecovery
```

where **<drive letter>** is the location where the Exchange software is located.
- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools

components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

Exchange post-installation on first node of an additional EVS (secondary site)

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#).

If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

Specifying a common node for failover

Specifying a common node for failover involves preparing the cluster with the Exchange Setup Wizard for VCS.

The failover node for the first Exchange virtual server, EVS1, was already set when creating the “first node” Exchange virtual server in [“Configuring disaster recovery for the first Exchange virtual server”](#) on page 474. Repeat these tasks for each of the Exchange Virtual Servers.

Preparing the cluster with the any-to-any option

When the designated Exchange virtual servers have been installed on the cluster, launch the Exchange Setup Wizard with the any-to-any option to specify the common failover node.

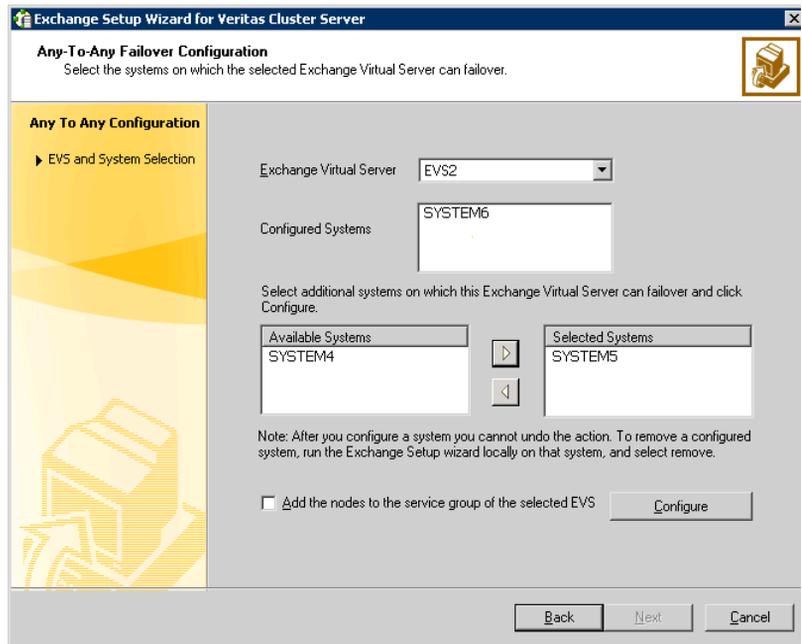
In the example, EVS1 is already configured with SYSTEM5 as a failover node, so you launch this wizard from EVS2 only.

Note: The Exchange software was installed on the common failover node during the installation process for the first EVS. You do not install Exchange a second time on the common failover node.

To prepare the cluster with the any-to-any option

- 1 Start the Exchange Setup Wizard for VCS from any node configured to host an Exchange service group. Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard**.

- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Options dialog box, choose the **Configure any-to-any failover** option and click **Next**.
- 5 Select systems to be configured for any-to-any failover. The **Configured Systems** box lists the nodes on which the Exchange Server service group can fail over. Do the following in order:



- Select the Exchange virtual server to which you want to add the additional failover nodes (in this case, EVS2).
- The Configured Systems box displays the nodes on which the selected Exchange virtual server has been installed.
- From the **Available Systems** box, select the systems to be configured for any-to-any failover.
- The **Available Systems** box lists only those systems that have the same version and service pack level of Microsoft Exchange as the selected Exchange virtual server.

- Click the right arrow to move the selected systems to the **Selected Systems** box. To remove a system from the box, select the system and click the left arrow.
 - Normally you run this wizard before you create the Exchange service group. If so, ensure that you clear **Add the nodes to the service group of the selected EVS**. If for some reason you already created the Exchange service group for this EVS, select this option to add the selected systems to the SystemList of the service group for the selected Exchange virtual server.
 - Click **Configure**.
 - Click **Next**.
- 6 Click **Finish**.

Cloning the service group configuration on to the secondary site using the DR wizard

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 on the primary site. You will also need to ensure that you have installed the application on the secondary site. After verifying launch the DR wizard.

To clone the service group configuration from the primary site to the secondary site

- 1 At the primary site, verify that you have brought the application service group online.
- 2 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Microsoft Exchange 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: <ul style="list-style-type: none"> ■ 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.

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 Exchange 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.
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, Replicator Log 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 VvrRvg 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.

Possible tasks after creating the DR environment

After you create the DR environment, you may need to perform the following additional tasks:

- Recovery procedures for service group dependencies
See [“Recovery procedures for service group dependencies”](#) on page 460.
- Adding a new failover node
See [“Possible task after creating the DR Environment: Adding a new failover node”](#) on page 464.

Testing fault readiness by running a fire drill

Topics in this chapter include:

- [About disaster recovery fire drills](#)
- [About the Fire Drill Wizard](#)
- [Tasks for configuring and running fire drills](#)
- [Prerequisites for a fire drill](#)
- [Fire Drill Wizard actions](#)
- [Preparing the fire drill configuration](#)
- [Running a fire drill](#)
- [Restoring the fire drill system to a prepared state](#)
- [Deleting the fire drill configuration](#)

About disaster recovery fire drills

A disaster recovery plan should include regular testing of an environment to ensure that a DR solution is effective and ready should disaster strike. This testing is called a fire drill. SFW HA provides a Fire Drill Wizard to help you set up and run a fire drill.

About the Fire Drill Wizard

The Fire Drill Wizard tests the fault readiness of a configuration by mimicking a failover from the primary site to the secondary site. The wizard does this without stopping the application at the primary site and disrupting user access.

The wizard prepares for the fire drill by completing the following steps:

- Creates a fire drill service group on the secondary site
The fire drill service group is a copy of the application service group, using the same service group name with the prefix *FD_{nn}*. The wizard renames the fire drill service group resources with a prefix *FD_{nn}* and changes attribute values as necessary to refer to the FD resources.
- Prepares a copy (mirror) of the production data on the secondary site
You assign one or more disks for the mirrored volumes while running the wizard. Mirror preparation can take some time, so you can exit the wizard once this step is started and let the preparation continue in the background.

Once these steps are complete, the wizard can run the fire drill. Running the fire drill detaches the mirrors from the original volumes to create point-in-time snapshots of the production data. It also brings the application online in the fire drill service group at the secondary site. Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to failover and come online at the secondary site should the need arise.

Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online on the primary site.

Running the fire drill creates a fire drill disk group on the secondary site with a snapshot of the application data to use for testing purposes. The wizard assigns the fire drill disk group name by prefixing the original disk group name with *FD_{nn}*.

After running the fire drill, you can choose to restore the fire drill configuration to a prepared state for use in regularly testing the disaster recovery solution, or you can delete the fire drill configuration and recreate it as needed using the wizard.

Tasks for configuring and running fire drills

The Fire Drill Wizard helps you configure and run a fire drill.

[Table 13-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 13-1 Tasks for configuring and running fire drills

Objective	Tasks
“Prerequisites for a fire drill” on page 503	<ul style="list-style-type: none"> ✓ Verifying hardware and software prerequisites
“Preparing the fire drill configuration” on page 505	<ul style="list-style-type: none"> ✓ Using the wizard to prepare the initial fire drill configuration
“Running a fire drill” on page 508	<ul style="list-style-type: none"> ✓ 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 509	<ul style="list-style-type: none"> ✓ 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 510	<ul style="list-style-type: none"> ✓ 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	<p>Creates the configuration required to run a fire drill. This step takes some time as the wizard prepares the mirrors for the snapshots.</p> <p>If this option is unavailable, the fire drill configuration already exists on the specified system.</p> <p>See “Preparing the fire drill configuration” on page 505.</p>
Run Fire Drill	<p>Runs the fire drill. The wizard creates the volume snapshots and brings the fire drill service group online. Optionally you can specify a script to be run once the fire drill is complete.</p> <p>If a fire drill has been run, you must restore the fire drill configuration to a prepared state before the wizard re-enables this option.</p> <p>See “Running a fire drill” on page 508.</p>
Restore to Prepared State	<p>Restores the fire drill configuration for another fire drill or to prepare the fire drill configuration for deletion.</p> <p>This option becomes available once a fire drill has been run.</p> <p>The wizard snaps back the snapshot mirrors to reattach to the original volumes and takes the fire drill service group offline.</p> <p>See “Restoring the fire drill system to a prepared state” on page 509.</p>
Delete Fire Drill Configuration	<p>Deletes the fire drill configuration to free up disk space. The wizard deletes the service group on the secondary site and performs a snap abort to delete the snapshot mirrors created on the secondary site for use in the fire drill.</p> <p>If a fire drill has been run, this option is disabled until you first restore the fire drill configuration to a prepared state. This ensures that mirrors are reattached and the fire drill service group is offline before the configuration is deleted.</p> <p>See “Deleting the fire drill configuration” on page 510.</p>

Preparing the fire drill configuration

Preparing the fire drill configuration creates a fire drill service group and snapshot mirrors of production data at the specified node on the secondary site. You specify the application service group and the secondary system to use. Only one service group can be prepared for a fire drill at one time.

Note: Preparing the snapshot mirrors takes some time to complete.

Before you prepare the fire drill configuration, you should verify that you meet the prerequisites.

See “[Prerequisites for a fire drill](#)” on page 503.

To prepare for the fire drill

- 1 Open the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**).
- 2 Start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, review the information and click **Next**.
- 4 In the System Selection panel, specify the primary site system on which the service group to be used for the fire drill is online and click **Next**.
All systems containing online global service groups are available to select. The default system is the node where you launched the wizard. When selecting a system you can specify either a fully qualified host name or IP address.
- 5 In the Service Group Selection panel, select the service group that you want to use for the fire drill and click **Next**. (You can select only one service group at a time for a fire drill.)
- 6 In the Secondary System Selection panel, select the cluster and the system to be used for the fire drill at the secondary site, and then click **Next**.
The selected system must have access to the replicated data and to disks for the snapshots that will be created for the fire drill.
- 7 In the Fire Drill Mode Selection panel, the available options depend on whether or not the fire drill service group already exists on this system and whether it is on or offline. Choose one of the following and click **Next**:

If the Prepare for Fire Drill option is available, a fire drill service group does not exist on this system. Click **Prepare for Fire Drill** and continue with the remaining steps in this procedure.

If the Run Fire Drill option is available, a fire drill service group has already been prepared. You can run the fire drill with no further preparation. Click **Run Fire Drill** and follow the procedure for running a fire drill.
See “[Running a fire drill](#)” on page 508.

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 509.

- 8 In the Fire Drill Service Group Settings panel, assign the virtual IP address and virtual name (Lanman name) to be used for the fire drill service group that will be created on the secondary site. These must be an address and name not currently in use.

If the service group contains more than one IP and Lanman resource, this panel does not display. After the fire drill service group is created, the wizard notifies you to manually update the IP and Lanman resources in the fire drill service group.

- 9 In the Disk Selection panel, review the information and make the selections as follows and click **Next**:

Volume	Select the volumes for the fire drill snapshots. By default all volumes associated with the service group are selected. If you deselect a volume that might result in the fire drill service group failing to come online, the wizard displays a warning message.
Disk Group	Shows the name of the disk group that contains the original volumes. This field is display only.
Fire Drill DG	Shows the name of the fire drill disk group that running the fire drill will create on the secondary system to contain the snapshots. This field is display only. For the fire drill disk group name, the wizard prefixes the original disk group name with <i>FDnn</i> .
Disk	Click the plus icon to the right of the Disk column and specify the disk to be used for the snapshot volume. Repeat for each row that contains a selected volume. You can store multiple snapshot volumes on the same disk, if the production volumes reside on disks in the same disk group.
Mount Details	Shows the mount details for the snapshot volumes on the secondary system, which match the mounts for the production volumes. This field is display only.

- 10 Wait while the wizard completes the preparation tasks. First the fire drill service group is created on the secondary site (but remains offline). Next the

snapshot mirrors for the volumes are prepared; this can take some time. You may want to minimize the wizard while the task runs in the background. You can also track the mirror preparation progress in the VEA. When done, the wizard displays a message that the fire drill preparation is complete.

- 11 To run the fire drill now, choose **Next**, or click **Finish** to exit the wizard. If you choose **Finish** the fire drill preparation remains in place. The next time you run the wizard, you choose the primary and secondary systems and service group and then can continue with running the fire drill.

Running a fire drill

After you complete the initial fire drill preparation step using the Fire Drill Wizard, you can run the fire drill immediately without exiting the wizard or run the wizard later to run the fire drill. Running the fire drill does the following:

- Creates the snapshots
 - Splits the fire drill disk group
 - Enables the firedrill resources
 - Brings the fire drill service group online
 - Optionally, executes a specified command to run a script
- For example, if you earlier created and populated a test table at the primary site, you could create a script to verify replication of the data. For the wizard to run the script, the script must exist on the secondary system that you are specifying for the fire drill.

To run a fire drill

- 1 If you completed the initial preparation and have not exited the wizard, go to [step 7](#). Otherwise, if you need to restart the wizard, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify the primary site system that contains the service group for which you want to run the fire drill and click **Next**.
- 5 In the Service Group Selection panel, select the service group and click **Next**.
- 6 In the Secondary System Selection panel, specify the system previously prepared for the fire drill at the secondary site.

- 7 In the Fire Drill Mode Selection panel, click **Run Fire Drill** and click **Next**. If a fire drill has been run previously, you must restore the fire drill configuration to a prepared state before the wizard enables the option to run another fire drill.
See “[Restoring the fire drill system to a prepared state](#)” on page 509.
- 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 509.

Restoring the fire drill system to a prepared state

After running a fire drill and verifying the results, use the Fire Drill Wizard to restore the fire drill system at the secondary site to a prepared state. When restoring the fire drill system to a prepared state, the wizard completes the following tasks:

- Takes the fire drill service group offline
- Disables the fire drill service group resources
- Imports the fire drill disk group
- Joins the fire drill disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

After running a fire drill, restoring the fire drill system to a prepared state is required to do any of the following:

- Run another fire drill.
- Restore the secondary system to a state where it can be used as failover for the application service group at the primary site.
- Delete the fire drill configuration.

To restore the fire drill system to a prepared state

- 1 If you completed running a fire drill and have not exited the wizard, go to [step 7](#). Otherwise, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify the system at the primary site that contains the service group on which the fire drill was run and click **Next**. The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 In the Fire Drill Mode Selection panel, click **Restore to Prepared State** and click **Next**.
If you have run a fire drill but not yet restored the configuration, this is the only option available. If the option is unavailable, the configuration has already been restored or the fire drill has not yet been run.
- 8 In the Restore Fire Drill screen, wait until the screen shows the restoration tasks are completed. Then click **Next** if you want to delete the fire drill configuration or click **Finish** to exit the wizard, leaving the fire drill configuration in a prepared state.

Deleting the fire drill configuration

If you no longer need a fire drill configuration you can delete it. Deleting a fire drill configuration deletes the fire drill service group on the secondary site and performs a snap abort of the snapshot mirrors created on the secondary site for use in the fire drill. It frees up the disk space used for the snapshot mirrors for other use.

If you have run a fire drill and want to delete the configuration, you must first restore the fire drill configuration to a prepared state before the wizard enables the option to delete the fire drill configuration.

See “[Restoring the fire drill system to a prepared state](#)” on page 509.

To delete a fire drill configuration

- 1 If you have just used the wizard to restore the fire drill configuration and have not exited the wizard, go to [step 7](#). Otherwise continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify the system at the primary site that contains the service group on which the fire drill was run and click **Next**. The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 In the Fire Drill Mode Selection panel, click **Delete Fire Drill Configuration** and click **Next**.
- 8 In the Delete Fire Drill Configuration panel, wait until screen shows the deletion is complete and then click **Next** and **Finish**.

512 | Testing fault readiness by running a fire drill
| **Deleting the fire drill configuration**

Appendices

- This section contains the following appendices:
- [Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation](#)
- [Configuring the DR components \(VVR and GCO\)](#)

Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation

This chapter covers the following topics:

- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the network and storage](#)
- [Configuring SFW HA: Prior to installing Exchange](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Installing Exchange on additional nodes \(secondary site\)](#)
- [Configuring SFW HA: After installing Exchange on secondary site](#)
- [Copying the .CRK file to the primary site](#)
- [Backing up and restoring the Exchange disk group](#)
- [Configuring the Exchange service group for VCS \(secondary site\)](#)
- [Verifying the cluster configuration](#)
- [Configuring DR components on primary and Secondary sites](#)

- [Possible task after creating the DR Environment: Adding a new failover node](#)

After setting up a SFW or SFW HA environment for Exchange on a primary site, you can create a secondary or “failover” site for disaster recovery. This chapter provides information on how to install and configure the high availability and Exchange components on the primary and secondary sites, with the intent of creating a parallel setup for the Exchange service group on both sites. This environment involves an active/passive configuration with one to one failover capabilities; refer to [Appendix C, “Deploying SFW HA for Disaster Recovery: Configuring any-to-any failover”](#) on page 635 for information on any-to-any configurations for Exchange in a DR environment.

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 [Chapter 11, “Deploying Disaster Recovery using the DR wizard for cloning: New Exchange Server installation”](#) on page 393.

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

Table A-1 Task List

Objective	Tasks
“Reviewing the requirements” on page 519	✓ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 523	✓ Understanding Active/Passive configuration and site failover in a DR environment
“Configuring the network and storage” on page 525	<ul style="list-style-type: none"> ✓ Setting up the network and storage for a cluster environment ✓ Verifying the DNS entries for the systems on which Exchange will be installed

Table A-1 Task List

Objective	Tasks
“Configuring SFW HA: Prior to installing Exchange” on page 526	<ul style="list-style-type: none"> ✓ Installing SFW HA ✓ Configuring the cluster using the Veritas Cluster Server Configuration Wizard ✓ Configuring the disk groups and volumes ✓ Setting up the forest and domain prior to the Exchange installation ✓ Managing disk group and volume operations, with instructions for mounting and unmounting volumes
“Installing Exchange on the first node (Primary site)” on page 561	<ul style="list-style-type: none"> ✓ Reviewing the prerequisite checklist ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation
“Moving Exchange databases (Primary site)” on page 565	<ul style="list-style-type: none"> ✓ Moving databases on the first node from the local drive to the shared drive using the Exchange Setup Wizard for Veritas Cluster Server
“Installing Exchange on additional nodes (Primary site)” on page 568	<ul style="list-style-type: none"> ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes
“Configuring the Exchange service group for VCS (Primary site)” on page 574	<ul style="list-style-type: none"> ✓ Creating the Exchange service group using the VCS Exchange Configuration Wizard.

Table A-1 Task List

Objective	Tasks
“Setting up the secondary Site: Configuring SFW HA prior to installing Exchange” on page 580	<ul style="list-style-type: none"> ✓ Reviewing the prerequisites ✓ Reviewing the configuration ✓ Configuring the network and storage ✓ Installing SFW HA ✓ Configuring the cluster using the Veritas Cluster Server Configuration Wizard ✓ Configuring the disk groups and volumes
“Installing Exchange on the first node and Additional nodes (Secondary site)” on page 580	<ul style="list-style-type: none"> ✓ Reviewing the prerequisite checklist ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation on the first node using the disaster recovery option ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes in the same virtual server
“Configuring SFW HA: After installing Exchange on secondary site” on page 592	<ul style="list-style-type: none"> ✓ Copying the public cryptographic key of the Exchange virtual server from the secondary site to the primary site ✓ Backing up the Exchange disk group on the primary site and restoring it on the secondary site ✓ Configuring the Exchange service group for VCS
“Verifying the cluster configuration” on page 599	<ul style="list-style-type: none"> ✓ Verifying the cluster configuration by switching service groups and shutting down an active cluster node

Table A-1 Task List

Objective	Tasks
“Configuring DR components on primary and Secondary sites” on page 600	✓ Completing the tasks outlined in “Configuring the DR components (VVR and GCO)” on page 605
“Possible task after creating the DR Environment: Adding a new failover node” on page 600	✓ Reviewing required tasks when adding a new failover system to either the primary or secondary site

Reviewing the requirements

This DR solution requires a primary site and secondary site.

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

Disk space requirements

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

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

Table A-2 Disk space requirements

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

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

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

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 Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:
 - Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition (SP2 required)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)
or
Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP1 required for all editions)
or
Windows Server 2003 R2 (Standard Edition, Enterprise Edition, or Datacenter Edition)
 - or
 - Microsoft Exchange 2000 Server or Microsoft Exchange 2000 Enterprise Server (SP3 with August 2004 rollup patch required for both editions)
with
Windows 2000 Server, Advanced Server, or Datacenter Server (all require Service Pack 4 with Update Rollup1)

Note: Microsoft support for Exchange Server 2003 is limited to 32-bit versions of the Windows 2003 operating system.

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 523.
- 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 Exchange Virtual Server (EVS)
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

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

Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

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

- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Domain Servers group.
- 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.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- You must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.
- 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.

- You must install IIS, SMTP, NNTP, ASP.NET, and WWW services before you install Microsoft Exchange Server.

Best practices

Symantec recommends that you perform the following tasks:

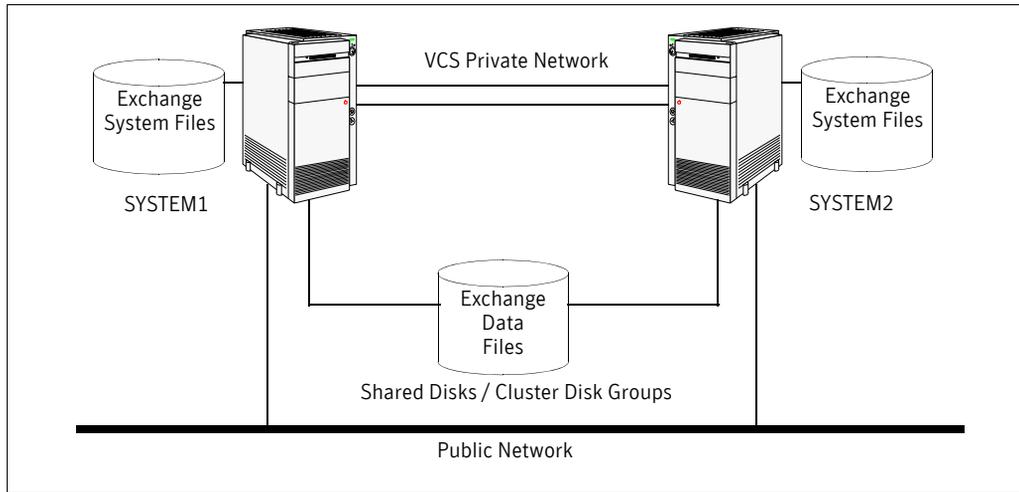
- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the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

Reviewing the configuration

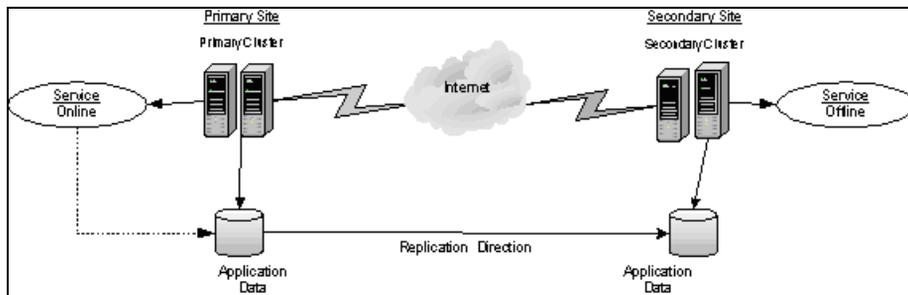
This overview highlights the high availability within a cluster, and the disaster recovery between two sites.

In an active/passive configuration with one to one failover capabilities, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. If you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM 5 and SYSTEM6 on the secondary site), EVS1 can fail over from SYSTEM1 to SYSTEM2 or vice versa on the primary site, and SYSTEM5 to SYSTEM6 or vice versa on the secondary site.

[Figure A-1](#) provides a view of a cluster configuration on the primary site:

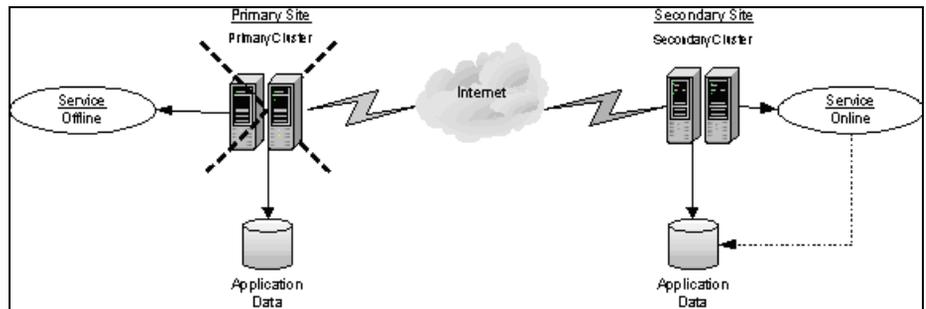
Figure A-1 Cluster configuration on the primary site

In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. [Figure A-2](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure A-2 Data replication in a disaster recovery environment

When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients. [Figure A-3](#) illustrates this type of failure:

Figure A-3 Application services restored after primary site failure



Configuring the network and storage

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent 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**.

Configuring SFW HA: Prior to installing Exchange

Before installing Exchange on the primary site, complete the following procedures:

- ✓ Install the SFW HA software.
See [“Installing Veritas Storage Foundation HA for Windows”](#) on page 527.
- ✓ Set up a VCS environment.
See [“Configuring the cluster”](#) on page 535
- ✓ Create the required disk groups and volumes.
See [“Configuring disk groups and volumes for disaster recovery”](#) on page 551.
See [“Managing disk groups and volumes”](#) on page 559.

- ✓ Prepare the Forest and Domain as required by Microsoft Exchange. See “[Preparing the forest and domain \(Primary site\)](#)” on page 560.

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 Application Agent for Exchange. 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.

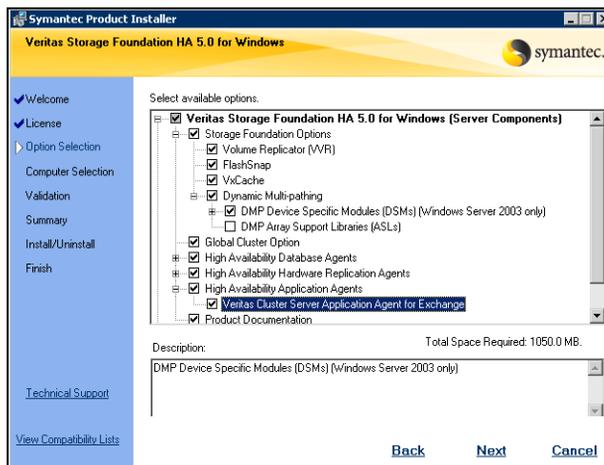
To install the product

- 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 for your installation. Click **Next**.



The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

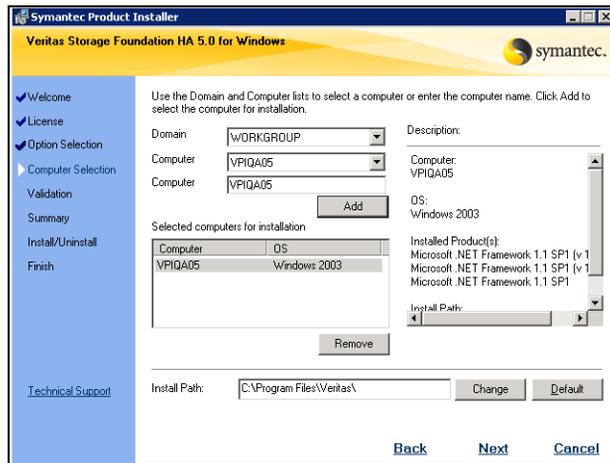
Veritas Cluster Server Application Agent for Exchange Required to configure high availability for Exchange Server.

Client

Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console.
Required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option	Required for a disaster recovery configuration only.
Veritas Volume Replicator	If you plan to use VVR for replication, you must also select the option to install VVR.

10 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

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

Computer

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

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

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

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

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
 - For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
 - For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Configuring VxSAS

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server 2003 environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.

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

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account)	Enter the administrative account name in the Account name field.
----------------------------------	--

Password	Specify a password in the 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 the Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

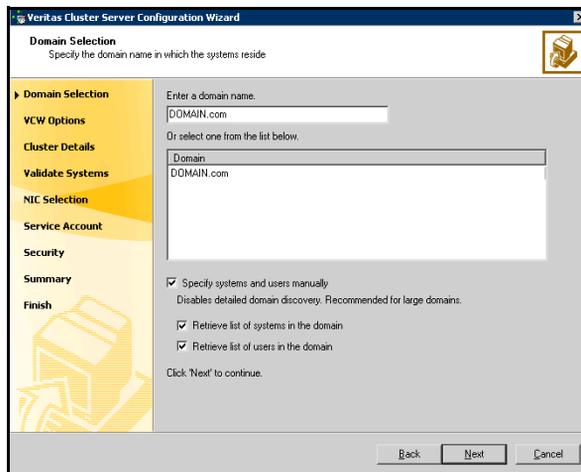
Complete the following tasks before creating a cluster:

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

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

To configure a VCS cluster

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



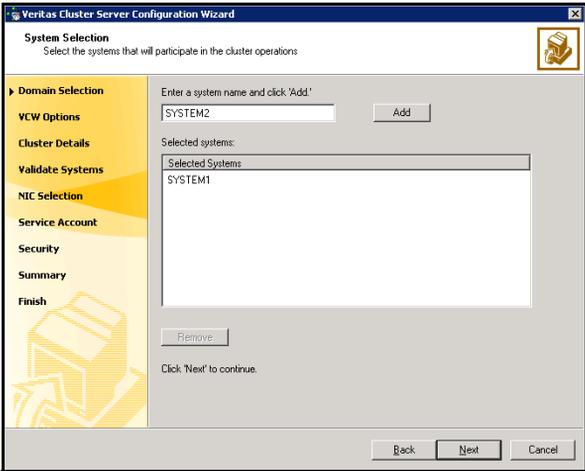
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
Proceed to [step 7](#) on page 538.

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

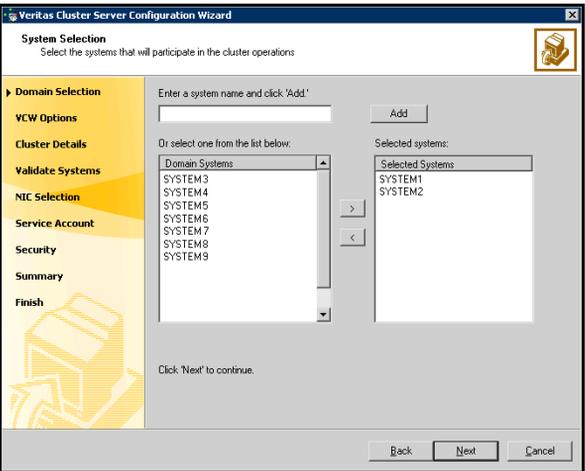
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 537. Otherwise 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.



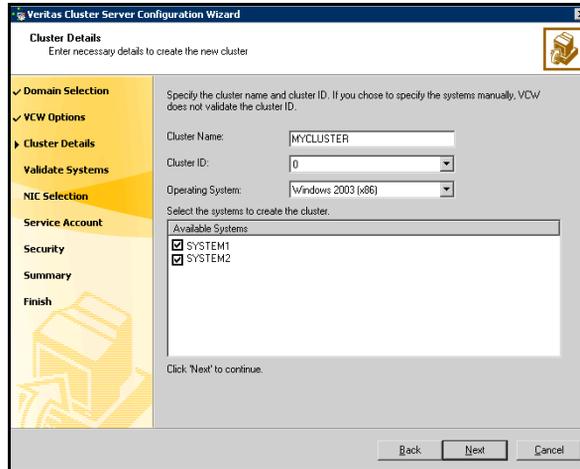
Proceed to [step 7](#) on page 538.

- 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**.

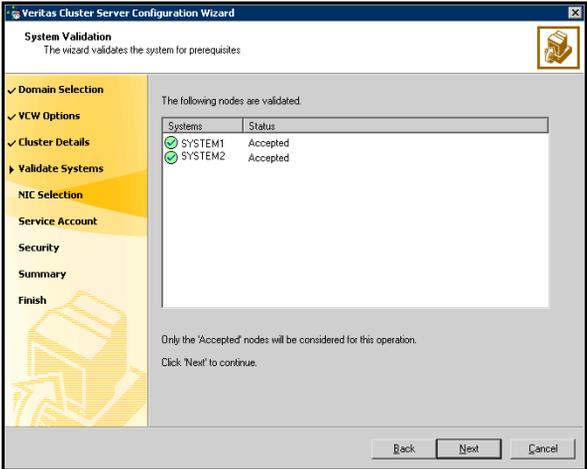


Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

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.

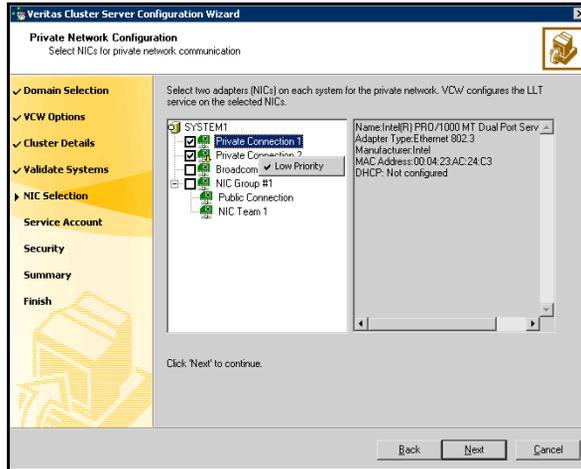
Operating System	From the drop-down list, select the operating system that the systems are running.
Available Systems	Select the systems that will be part of the cluster. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts 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**.



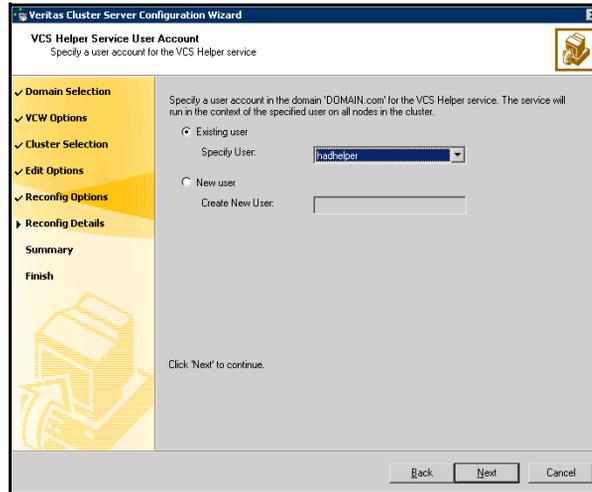
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem. If you chose to configure a private link heartbeat in [step 8](#) on page 538, proceed to the next step. Otherwise, proceed to [step 11](#) on page 540.

- 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 536, 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**.

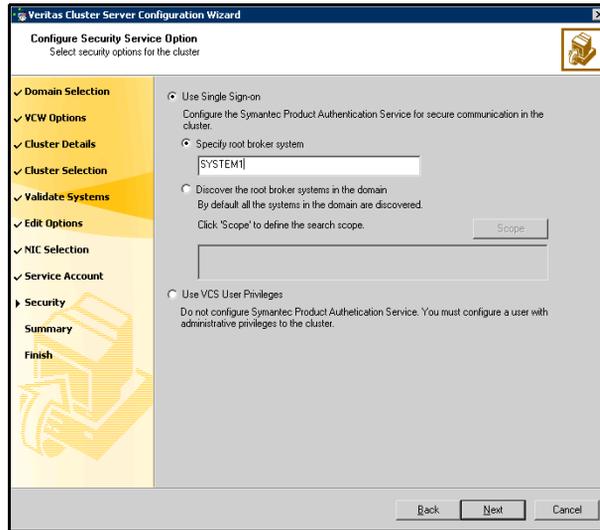
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- 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.

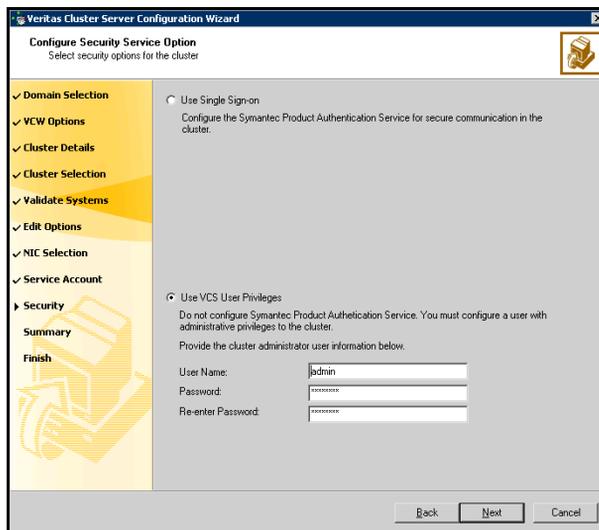
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- 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.

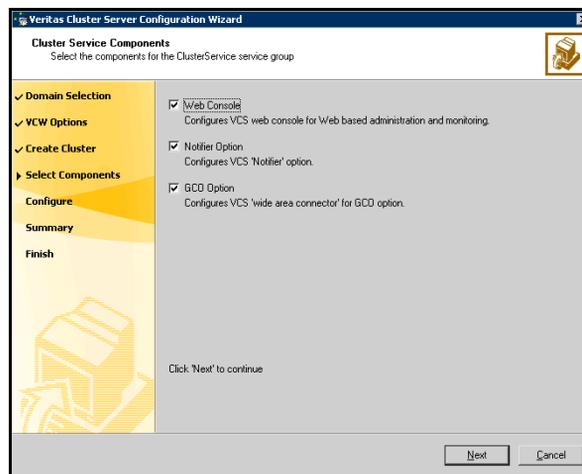
The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

- 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 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.

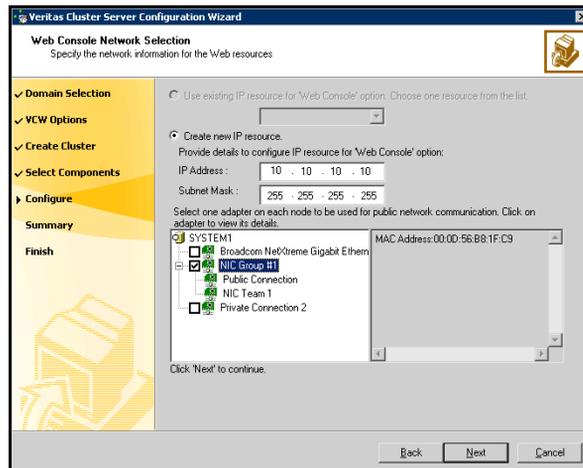
The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

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.

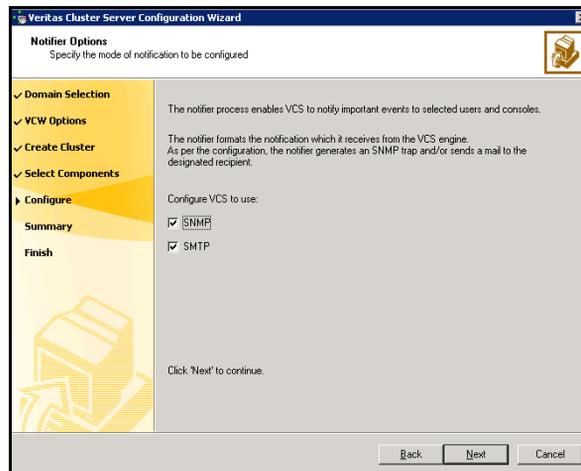
- Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to “[Configuring notification](#)” on page 546.
If you chose to configure global cluster components, proceed to “[Configuring the wide-area connector process for global clusters](#)” on page 550.
Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

To configure notification

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



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

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

The screenshot shows the 'Notifier SNMP Configuration' window. On the left is a navigation pane with options: Domain Selection, VCW Options, Create Cluster, Select Components, Configure (selected), Summary, and Finish. The main area contains a table for configuring SNMP consoles and a text input for the trap port.

SNMP Console	Severity
snmpserv	Information
snmpserv1	SevereError

Enter name or IP of the SNMP console and severity level for each

Click on "+" button to add more consoles.
Click "-" to remove a console.

Enter SNMP Trap Port:

Note: SNMP console must be MIB 2.0 compliant

Click 'Next' to continue.

Buttons: Back, Next, Cancel

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

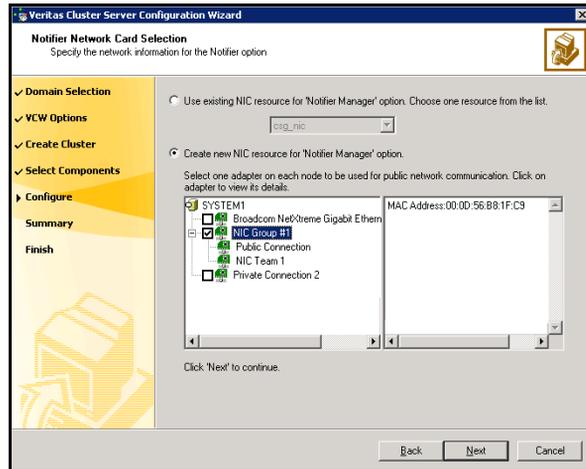
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.

The screenshot shows the 'Notifier SMTP Configuration' window of the Veritas Cluster Server Configuration Wizard. The window title is 'Veritas Cluster Server Configuration Wizard' and the subtitle is 'Notifier SMTP Configuration'. Below the subtitle, it says 'Specify information about SMTP recipients'. On the left side, there is a navigation pane with the following steps: 'Domain Selection', 'VCW Options', 'Create Cluster', 'Select Components', 'Configure', 'Summary', and 'Finish'. The 'Configure' step is currently selected. In the main area, there is a text box for 'SMTP Server Name / IP' containing 'SMTPServer'. Below this, it says 'Enter SMTP recipients and select a severity level for each recipient.' There is a table with two columns: 'Recipients' and 'Severity'. The first row contains 'admin@example.com' and 'Information'. Below the table, there are instructions: 'Click '+' to add a recipient.' and 'Click '-' to remove a recipient.' There are '+' and '-' buttons. At the bottom, it says 'Click 'Next' to continue.' and there are 'Back', 'Next', and 'Cancel' buttons.

Recipients	Severity
admin@example.com	Information

- 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**.



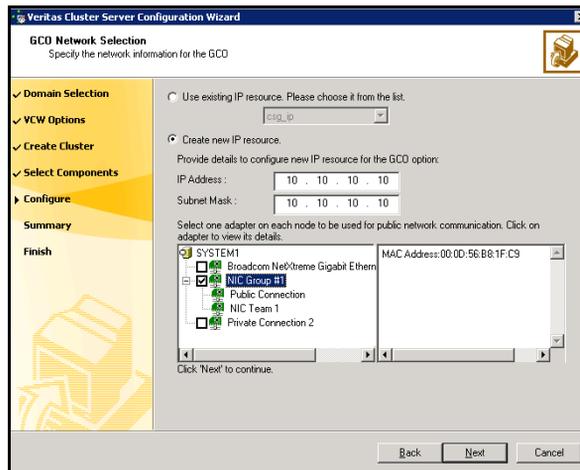
- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard 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 550. 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.

The wizard does not set up a global cluster environment; it configures a resource for the wide-area connector that is required for inter-cluster communication. Set up the global cluster after both sites are configured.

See “[Configuring the DR components \(VVR and GCO\)](#)” on page 605.

Configuring disk groups and volumes for disaster recovery

Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to grow existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

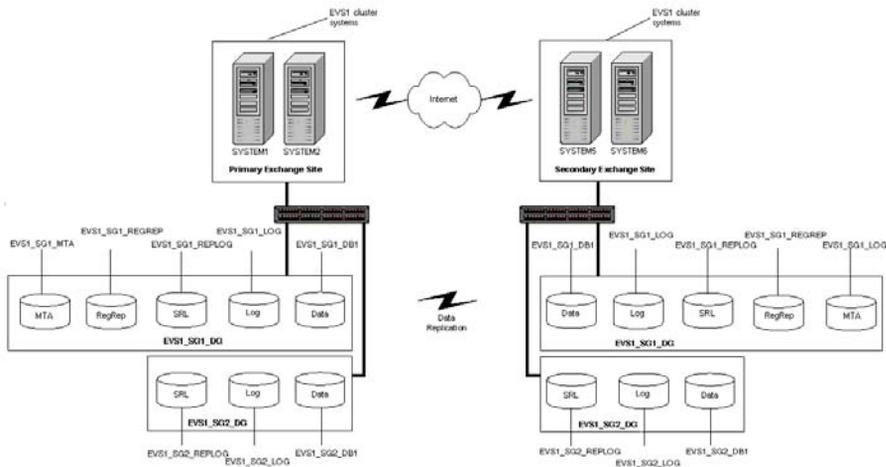
Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.

Note: While creating disk groups and volumes for the secondary site, make sure to use the same names of volumes as those on the primary site. The size of the volumes on the secondary site must be equal to or larger than the size of the volumes on the primary site.

Typically, a SFW disk group corresponds to an Exchange storage group. Below is a detailed view of the disk groups and volumes on both the primary and secondary sites of a disaster recovery environment:

Figure A-4 Disk groups and volumes for Exchange virtual server EVS1 in DR setup



Exchange storage group EVS1_SG1_DG contains five volumes in a disaster recovery environment:

- EVS_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.
- EVS1_SG1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.
- EVS1_SG1_MTA: Contains the MTA database.
- EVS1_SG1_REPLOG: Contains the VVR Storage Replicator Log. This volume is required in a DR solution; it is *not* required in an HA solution.

The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Additional storage groups (for example, EVS1_SG2_DG) only contain the data, log, and VVR Storage Replicator Log volumes; the RegRep and MTA volumes are included in the first storage group.

Caution: 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.

Setting up the Replicated Data Sets and the associated log (EVS1_SG1_REPLOG) is in a wizard. You can wait until that wizard is run in a later steps to configure VVR and its related volumes.

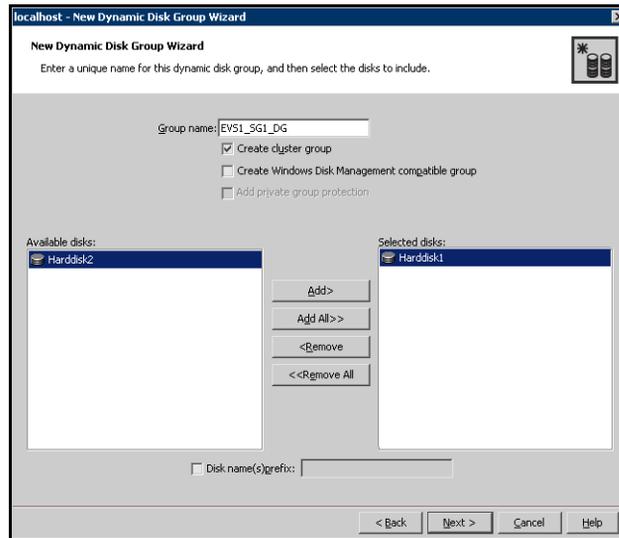
See “[Setting up the replicated data sets \(RDS\) for VVR](#)” on page 607.

Creating a disk group

To create a dynamic (cluster) disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.

6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_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

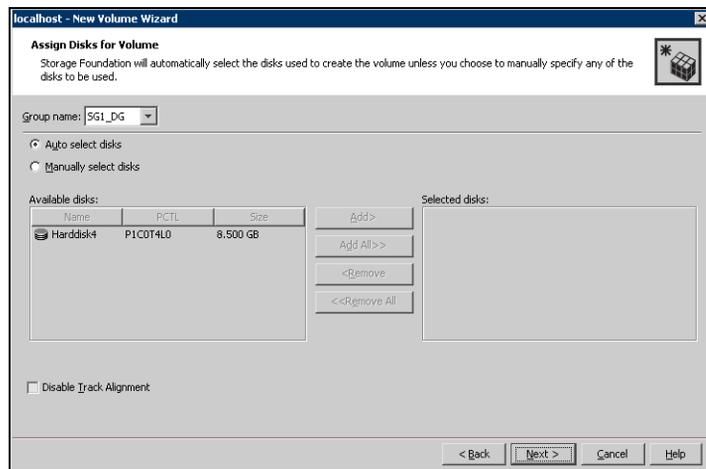
This procedure assumes you are starting with the EVS1_SG1_DB1 volume. Refer to the steps below for the data, log, RegRep, and MTA volumes.

You can create the Storage Replicator Log volumes during the Replicated Data Sets wizard at a later time.

See “[Setting up the replicated data sets \(RDS\) for VVR](#)” on page 607.

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 SG1_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.

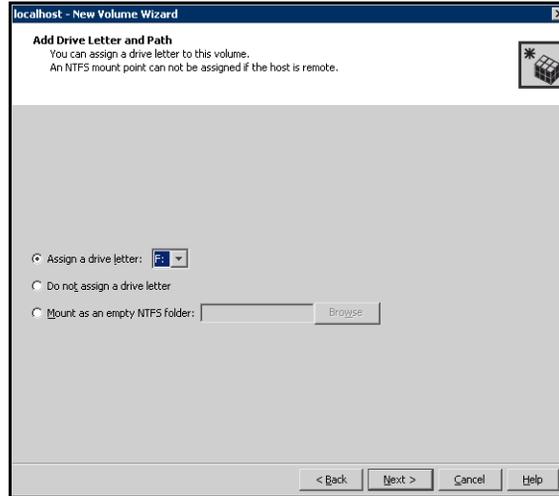
- Click **Next**.

7 Specify the 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.
- 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.

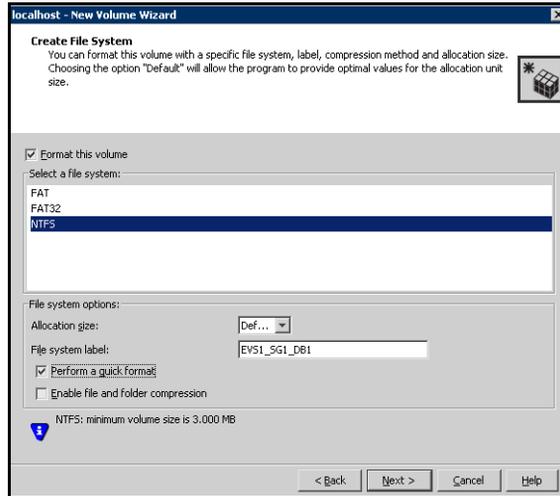


- To assign a drive letter:
Select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder:
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- For the Replicator Log volume only:
Select **Do not assign a drive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

Note: Do not assign the M: drive letter if you are using Exchange 2000, as it is reserved for internal use.

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 the RegRep volume (EVS1_SG1_REGREP), the MTA volume (EVS1_SG1_MTA), and the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.)
- 13 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3).

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

- ◆ If you are configuring an any-to-any environment, you can also create similar disk groups and volumes for the other Exchange servers. For example, create disk group (EVS2_SG1_DG) and volumes (EVS2_SG1_DB1, EVS2_SG1_REGREP, EVS2_SG1_LOG, and EVS2_SG1_MTA).
- ◆ If you are configuring the secondary site of a Disaster Recovery installation, continue with the installation of Exchange on the secondary site. See “[Installing Exchange on the first node and Additional nodes \(Secondary site\)](#)” on page 580.

Managing disk groups and volumes

This section includes the following procedures:

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

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.

- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

If you are configuring a secondary site, return to that procedure.

See “[Installing Exchange on the first node and Additional nodes \(Secondary site\)](#)” on page 580.

Preparing the forest and domain (Primary site)

Microsoft requires the preparation of the forest and domain prior to an Exchange installation. See Microsoft documentation for instructions for preparing the forest and domain for an Exchange installation. Do not repeat this process for additional Exchange installations.

Installing Exchange on the first node (Primary site)

Installing Exchange on the first node is described in three stages that involve pre-installation, installation, and post-installation procedures. Complete the following tasks before installing Exchange Server:

- ✓ Prepare the forest and domain.
See “[Preparing the forest and domain \(Primary site\)](#)” on page 560.
- ✓ Verify the disk group is imported on the first node of the cluster.
See “[Importing a disk group and mounting a shared volume](#)” on page 559.
- ✓ Mount the volume containing the information for registry replication.
See “[Importing a disk group and mounting a shared volume](#)” on page 559.
- ✓ Verify that all systems on which Exchange Server is to be installed have IIS installed; the SMTP, NNTP, and WWW services must be installed on all systems. For installing Exchange on Windows 2003, ASP.NET service must also be installed.
- ✓ Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- ✓ Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.
 - Make sure the HAD Helper domain user account has “Add workstations to domain” privilege enabled in the Active Directory. To verify this, click **Start > Administrative Tools > Local Security Policy** on the domain controller to launch the security policy display. Click **Local Policies > User Rights Management** and make sure the user account has this privilege.

- If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

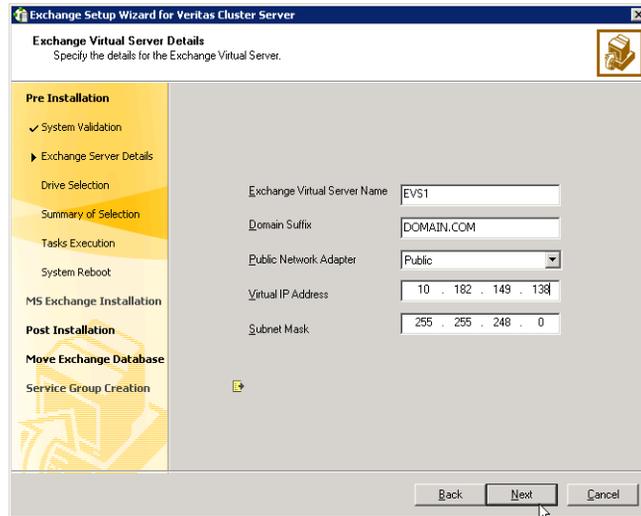
Exchange pre-installation: First node

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. You must install Exchange on a virtual node to facilitate high availability.

To perform Exchange pre-installation

- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
- 2 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a new Exchange Virtual Server** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.

7 Specify information related to the network.



- Enter a unique virtual name for the Exchange server.
Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.
- Enter a domain suffix for the virtual server.
- Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
- Enter a unique virtual IP address for the Exchange virtual server.
- Enter the subnet mask for the virtual IP address.
- Click **Next**.

The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server is not online on the network. If the Exchange virtual server is still online at the primary site you will be prompted to offline the group. Enter the VCS administrative user name and password for the primary cluster and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met.

Click **Next**.

- 8 Select a drive where the registry replication data will be stored and click **Next**.

- 9 Review the summary of your selections and click **Next**.
- 10 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 11 The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
- 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the name specified for the Exchange virtual server is temporarily assigned to the node. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Server Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

- 13 Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.
Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: First node

Install Exchange on the node selected in the Exchange pre-installation.

Exchange 2000 requires service pack 3 with the August 2004 rollup patch. Exchange 2003 requires service pack 2. The procedure below is based on Exchange 2003 SP2.

This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Install Exchange Server using the Microsoft Exchange installation program. See the Microsoft Exchange documentation for instructions.
- 2 Reboot the node if prompted to do so.
- 3 For Exchange 2000 or Exchange 2003, install the required service pack.

Exchange post-installation: First node

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

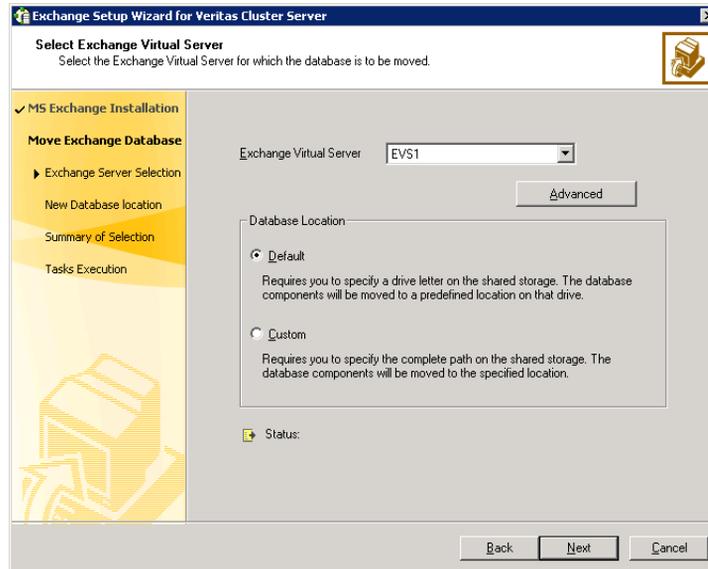
Moving Exchange databases (Primary site)

After completing the Exchange installation on the first node, move the Exchange databases on the first node from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

- ✓ Make sure the data queue is empty on the SMTP server.
- ✓ Make sure to import the disk group and mount the volumes for the Exchange database, MTA data, and transaction logs.
See "[Managing disk groups and volumes](#)" on page 559.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
- 4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.
- 5 In the Select Exchange Virtual Server dialog box:



- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

Warning: The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Click **OK**.
 - Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
 - Click **Next**.
- 6 Do one of the following:
- If you would like to move the first mailbox store, public store, and MTA data to the generated default paths on the volumes for these components, proceed to the next step.
 - Otherwise, proceed to [step 8](#) on page 568 to specify the path location on the volumes that you will designate for these components.

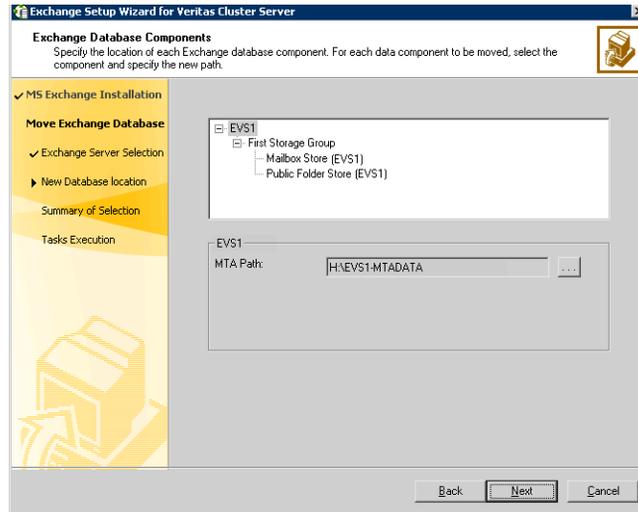
Warning: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

- 7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components will be moved to a default location on that drive. In the Exchange Database Components dialog box:

The screenshot shows the 'Exchange Setup Wizard for Veritas Cluster Server' dialog box, specifically the 'Exchange Database Components' step. The title bar reads 'Exchange Setup Wizard for Veritas Cluster Server'. Below the title bar, the text says 'Exchange Database Components' and 'Specify a drive letter to move data base to default location.' The dialog box is divided into a left sidebar and a main content area. The sidebar has a yellow background and contains the following items: 'MS Exchange Installation' (checked), 'Move Exchange Database' (checked), 'Exchange Server Selection' (checked), 'New Database location' (selected), 'Summary of Selection', and 'Tasks Execution'. The main content area has a grey background and contains three sections: 'Database', 'Transaction Logs', and 'MTA Data'. Each section has three input fields: 'Current Path', 'Target Drive', and 'New Path'. The 'Database' section shows 'Current Path: E:\EXCHPCO\MDBData', 'Target Drive: H:\', and 'New Path: H:\EXCHPCO\MDBData\First Storage Group\Mailbox Store 1'. The 'Transaction Logs' section shows 'Current Path: G:\EXCHPCO\TransactionLogs', 'Target Drive: H:\', and 'New Path: H:\EXCHPCO\TransactionLogs\First Storage Group'. The 'MTA Data' section shows 'Current Path: H:\EXCHPCO\MTAData', 'Target Drive: H:\', and 'New Path: H:\EXCHPCO\MTAData'. At the bottom right of the dialog box are three buttons: 'Back', 'Next', and 'Cancel'.

- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Specify the drive where the Exchange MTA Data will be moved.
- Click **Next** and proceed to [step 9](#) on page 568.

- 8 For the option of a custom database location, specify the location for specific Microsoft Exchange data components. In the Exchange Database Components dialog box:



For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the path for the Exchange database components contains only ANSI characters.

- 9 Review the summary of your selections and click **Next**.
- 10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.
- 11 Click **Finish** to exit the wizard.

Installing Exchange on additional nodes (Primary site)

After moving the Exchange databases to shared storage, install Exchange on additional nodes in the cluster for the same Exchange virtual server (EVS1). You must run pre-installation, installation, and post-installation procedures for each additional node.

Make sure to review the prerequisites for permissions. See “[Installing Exchange on the first node \(Primary site\)](#)” on page 561.

Exchange pre-installation: Additional nodes

Note: Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes. See “[Unmounting a volume and deporting a disk group](#)” on page 559 for instructions.

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server. The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.

- Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically generated when setting up an additional failover mode for the virtual server in the same cluster.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
 - 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.
 - 11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.
 - 12 Click **Reboot**.
The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: Additional nodes

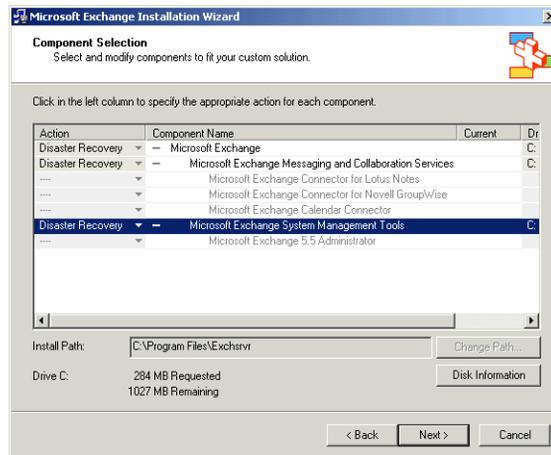
Install Exchange on the same node selected in the Exchange pre-installation.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:
`<drive letter>:\SETUP\I386\setup.exe /disasterrecovery`
where `<drive letter>` is the location where the Exchange software is located.
- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

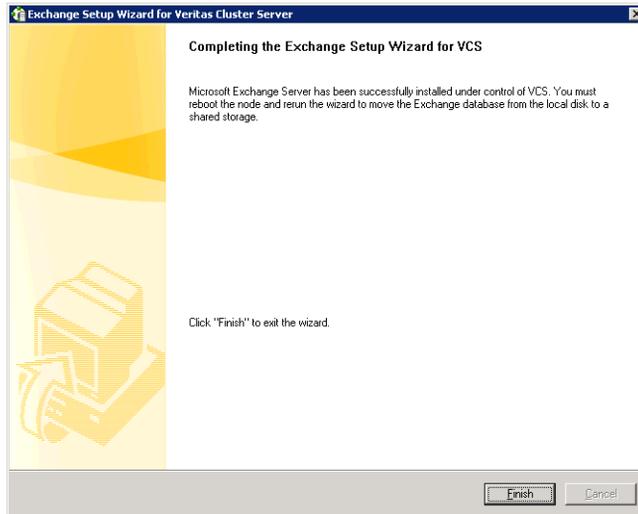
Exchange post-installation: Additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.

- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

Configuring the Exchange service group for VCS (Primary site)

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

Prerequisites

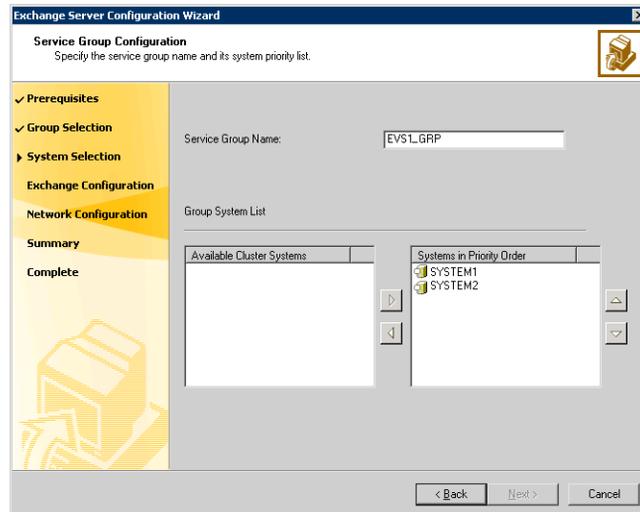
- ✓ You must be a Cluster Administrator.
- ✓ You must be a Local Administrator on the node where you run the wizard.
- ✓ Verify that Command Server is running on all nodes in the cluster.
- ✓ Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard.
- ✓ Import the disk group and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database
 - registry changes related to Exchange
 - transaction logs for the first storage group
 - MTA database
- ✓ Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- ✓ Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- ✓ Verify Microsoft Exchange is installed and configured identically on all nodes.

Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to the EVS1_SG1_DG disk group.

To configure the Exchange service group

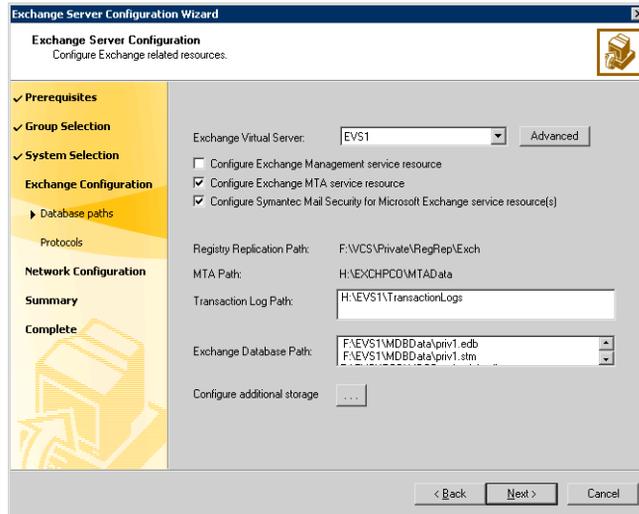
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Configuration Wizard**

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



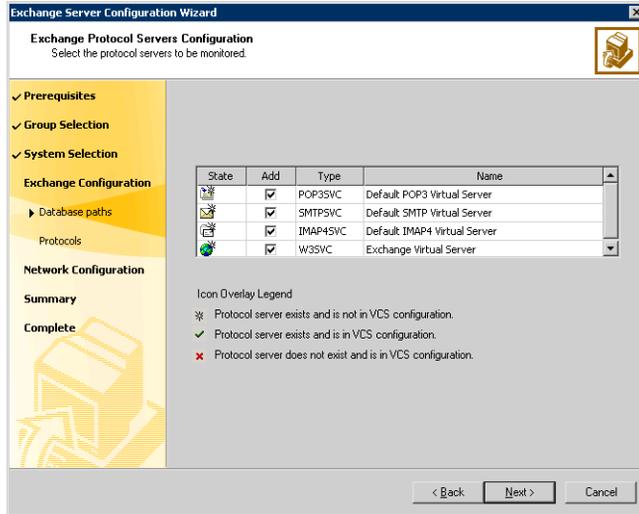
- Enter a name for the Exchange service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list. Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.
- To change a node's priority in the service group's system list, select the node in the Systems in Priority Order list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:

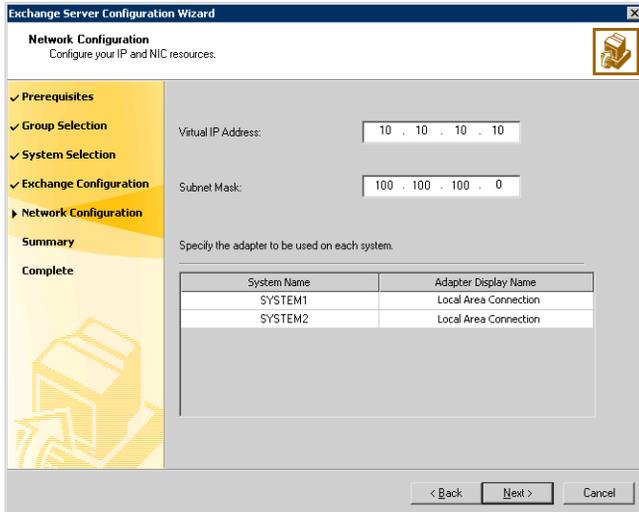


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.
- Specify whether you want to configure the Exchange Management service.
 - Specify whether you want to configure the Exchange MTA service resource. The Exchange Message Transfer Agent is specific to legacy Exchange message routing based on X.400. There are specific circumstances where it may or may not be required for your Exchange server. Please refer to Microsoft documentation on Exchange message routing and the use of MTA for further clarification and applicability to its use within your Exchange environment.

- If you are utilizing Symantec Mail Security for Exchange be sure to indicate that you would like to configure this resource.
 - Click **Next**.
- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



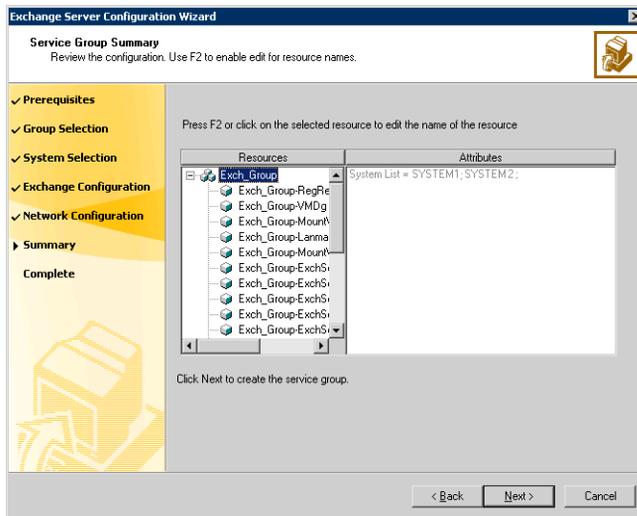
- 7 Specify information related to the network:



- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

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

- Click **Next**.
- 8 Review the service group configuration and change the resource names, if desired:



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 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.

- Click **Next**.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 9 In the Completing the Exchange Configuration dialog box, select the **Bring the service group online** check box to bring the service group online on the local system.
 - 10 Click **Finish**. After bringing the service group online, you must run the Exchange System Manager so that all the stores are automatically mounted on start-up.

To reconfigure mounting of stores at start-up

- 1 Start Exchange System Manager.
- 2 In the left pane, navigate to your storage group.
If administrative groups are not configured, Servers > Exchange Server > Storage Group.
If more than one administrative group is configured, expand Administrative Groups > Your Administrative Group > Servers > Exchange Server > Storage Group.
- 3 Right-click the Exchange database and choose **Properties** from the pop-up menu.
- 4 Click the Database tab.
- 5 Clear the **Do not mount this store at start-up** check box.
- 6 Click **OK**.

Repeat these steps for all the Exchange databases that were previously mounted.

If you need to configure additional storage groups or mailbox stores on the shared storage you should do that now. Import disk groups and mount volumes that have been created for the additional storage groups or mailbox stores, and create the new storage groups and mailbox stores in Exchange System Manager, and then rerun the Exchange Configuration Wizard to bring them under VCS control. If you already designated the additional mounted volumes and disk groups when you ran the configuration wizard the first time then you can just create the storage groups and mailbox stores in Exchange System Manager.

Setting up the secondary Site: Configuring SFW HA prior to installing Exchange

After setting up a SFW HA environment on the primary site, repeat the same tasks on the secondary site prior to the Exchange installation.

All the tasks for a Secondary Site apply either to a Disaster Recovery new installation or to an installation where the Primary Site has configured a standalone Exchange server into the cluster.

Begin with reviewing the requirements on the secondary site, similar to the primary site and continue with same the procedures as for the primary site.

See “[Reviewing the requirements](#)” on page 519

- ✓ Reviewing the requirements
- ✓ Reviewing the configuration
- ✓ Configuring the network and storage
- ✓ Installing SFW HA
- ✓ Configuring the cluster using the Veritas Cluster Server Configuration Wizard
- ✓ Configuring the disk groups and volumes

After completing these tasks, proceed to installing Exchange on the secondary site.

Installing Exchange on the first node and Additional nodes (Secondary site)

Complete the following tasks:

- ✓ Reviewing the prerequisite checklist
- ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation on the first node using the disaster recovery option
Make sure to perform the first node pre-installation, installation, and post-installation procedures.
- ✓ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes in the same virtual server

Make sure to perform the additional node pre-installation, installation, and post-installation procedures.

Installing Exchange on the first node with DR Option (Secondary site)

Installing Exchange on the first node is described in three stages that involve pre-installation, installation, and post-installation procedures. In this procedure, virtual Exchange server, EVS1, will fail over from SYSTEM5 to SYSTEM 6. Complete the following tasks before installing Exchange Server:

- ✓ Verify the disk group is imported on the first node of the cluster.
See [“Managing disk groups and volumes”](#) on page 559.
- ✓ Mount the volume containing the information for registry replication.
- ✓ Verify that all systems on which Exchange Server is to be installed have IIS installed; the SMTP, NNTP, and WWW services must be installed on all systems. For installing Exchange on Windows 2003, ASP.NET service must also be installed.
- ✓ Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- ✓ Set the required permissions:
 - You must be a domain user.
 - You must be an Exchange Full Administrator.
 - You must be a member of the Exchange Domain Servers group.
 - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
 - You must have write permissions on the DNS server to perform DNS updates.
 - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
 - The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- ✓ Make sure to use the same drive letters employed on the primary site.

- ✓ Make sure to take the Exchange service group offline on the primary site; otherwise, the wizard will be prompt you take the service group offline.

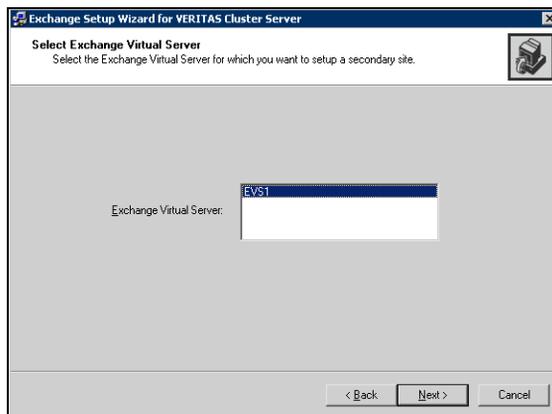
Exchange pre-installation on first node (Secondary site)

To perform Exchange pre-installation

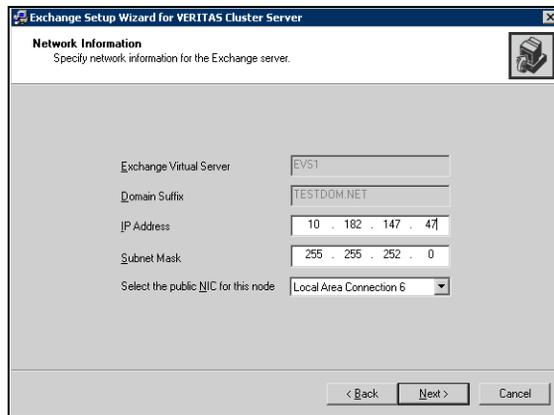
- 1 Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.

Note: After you have run the wizard, you will be requested to restart the node. So, close all open applications and save your data before running the wizard.

- 2 Click **Start > All Programs > Veritas > Veritas Cluster Server > Application Agent for Exchange > Exchange Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for Exchange disaster recovery setup** option and click **Next**.
- 6 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.
- 7 In the Select Exchange Virtual Server dialog box:



- a Select the Exchange virtual server for disaster recovery.
 - b Click **Next**.
- 8 The installer verifies that the selected node meets the Exchange requirements, click **Next**. If the service group on the primary node has not been taken offline the installer prompts you to do so without exiting the installer, or you can cancel the installation wizard and take the service group offline manually.
- 9 Enter the name of a failover node, and click **Next**.
- 10 Specify the information related to your network. Make sure to store the virtual name and IP address for future use.



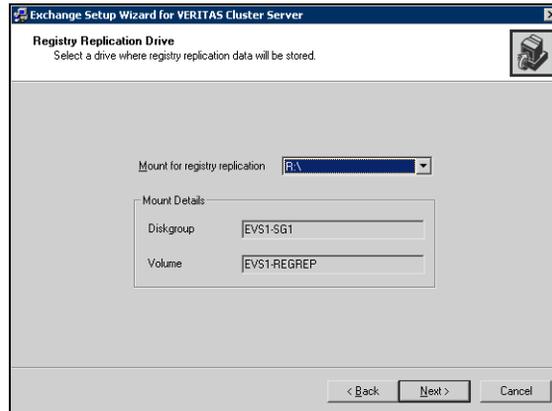
The screenshot shows the 'Exchange Setup Wizard for VERITAS Cluster Server' window. The title bar includes a help icon, the text 'Exchange Setup Wizard for VERITAS Cluster Server', and a close icon. The main window has a title 'Network Information' and a subtitle 'Specify network information for the Exchange server.' Below this, there are five input fields: 'Exchange Virtual Server' with the value 'EVS1', 'Domain Suffix' with 'TESTDOM.NET', 'IP Address' with '10 . 182 . 147 . 47', 'Subnet Mask' with '255 . 255 . 252 . 0', and 'Select the public NIC for this node' with a dropdown menu showing 'Local Area Connection 6'. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

- a Verify the virtual computer name for the server.
- b Verify the domain suffix.
- c Enter a unique virtual IP address for the virtual server, or use the same IP address as the virtual server on the primary site.
- d Enter the subnet to which the virtual IP address belongs.
- e Select the public NIC.

Caution: The installer displays all TCP/IP enabled adapters on a node, including the private network adapters. Make sure that you select the adapters for the public network, and not those assigned to the private network.

- f Click **Next**.

11 In the Registry Replication Drive dialog box:



- a Select the same drive letter (or directory in the case of folder mounts) as the one used on the primary site for registry replication.
 - b Click **Next**.
- 12 Review the summary of selections and click **Next**.
- 13 After reviewing the warning of the renaming and rebooting of the system, click **Yes**.
- 14 After the installer performs configuration tasks, click **Next**.
- 15 If the wizard could not locate a DNS entry for the specified Exchange server and IP address, click **OK** to create one.
- 16 Click **Reboot**.
- 17 Click **Yes** to reboot the node.
- 18 When prompted to install Microsoft Exchange with the /disasterrecovery option, click **OK**. If you need to undo all actions performed by the wizard during the pre-installation phase, click **Revert Changes**. Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.

Caution: After you reboot the node, the values specified for the Exchange virtual server are temporarily assigned to the node. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation on first node (Secondary site)

Install Exchange on the same node selected in “[Exchange pre-installation on first node \(Secondary site\)](#)” on page 582.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

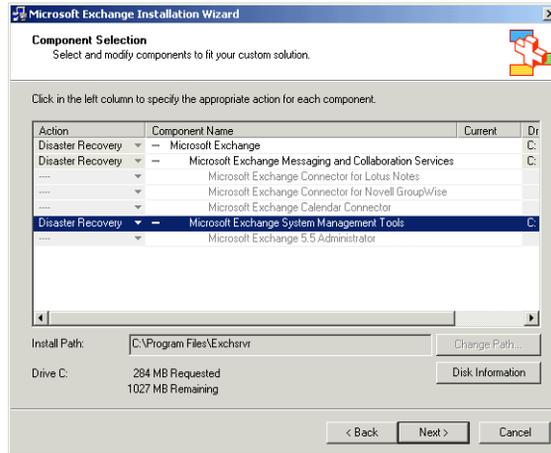
To install Exchange

- 1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe  
/disasterrecovery
```

where **<drive letter>** is the location where the Exchange software is located.
- 2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools

components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

Exchange post-installation on first node (Secondary site)

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name, and sets the Exchange services to manual so that the Exchange services can be controlled by VCS.

To perform Exchange post-installation

- 1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears indicating that the system will be renamed and restarted after you quit the wizard. This sets the node back to its physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.
- 6 Click **Finish**.
- 7 The wizard prompts you to reboot the node. Click **Yes** to reboot the node. Changes made during the post-installation steps do not take effect till you reboot the node.

Installing Exchange on additional nodes (secondary site)

Install Exchange on additional nodes in the cluster for the same Exchange virtual server (EVS1). You must run pre-installation, installation, and post-installation procedures for each additional node.

Note: In an any-to-any configuration, the steps for installing Exchanges on the additional nodes (failover nodes) can be completed for the first Exchange server, and do not need to be repeated for the common failover nodes for additional Exchange servers.

Make sure to complete the following tasks before the Exchange installation:

- ✓ Review the prerequisites for permissions.
See “[Installing Exchange on the first node with DR Option \(Secondary site\)](#)” on page 581.
- ✓ Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes.
See “[Managing disk groups and volumes](#)” on page 559.

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

Exchange pre-installation: Additional nodes

To perform Exchange pre-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 From the node to be added to an Exchange cluster, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 In the Available Option dialog box, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 In the Select Option dialog box, choose the **Create a failover node for existing Exchange Virtual Server** option and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 The wizard validates the system for the prerequisites. Various messages indicate the validation status. When the validations are done, click **Next**.
- 8 Specify network information for the Exchange virtual server.
The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.
 - Select the appropriate public NIC from the drop-down list. The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Optionally, enter a unique virtual IP address for the Exchange virtual server. By default, the text box displays the IP address assigned when the Exchange Virtual Server (EVS1) was created on the first node. You should not have to change the virtual IP address that is automatically generated when setting up an additional failover mode for the virtual server in the same cluster.
 - Enter the subnet mask for the virtual IP address.
 - Click **Next**.
- 9 Review the summary of your selections and click **Next**.
- 10 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.

11 The wizard runs commands to set up the VCS environment. Various messages indicate the status of each task. After all the commands are executed, click **Next**.

12 Click **Reboot**.

The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

Warning: After you reboot the node, the Exchange virtual server name is temporarily assigned to the node on which you run the wizard. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

On rebooting the node, the Exchange Setup wizard is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server.

Click **Revert** to undo all actions performed by the wizard during the pre-installation procedure.

Do not click **Continue** at this time. Wait until after the Exchange installation is complete.

Exchange installation: Additional nodes

Install Exchange on the same node selected in “[Installing Exchange on additional nodes \(secondary site\)](#)” on page 587.

- Install any required service packs.
- Install the same Exchange version and components on all nodes.

The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

To install Exchange

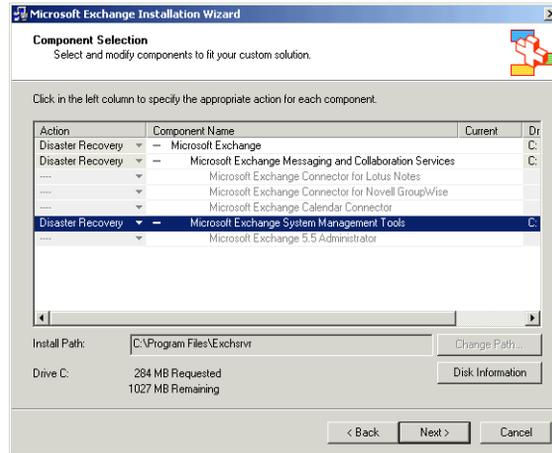
1 Begin the Exchange installation for disaster recovery at the command prompt using the /disasterrecovery option:

```
<drive letter>:\SETUP\I386\setup.exe  
/disasterrecovery
```

where <drive letter> is the location where the Exchange software is located.

2 During the wizard, verify or select **Disaster Recovery** in the **Action** column for the Microsoft Exchange, Microsoft Exchange Messaging and Collaboration services and Microsoft Exchange System Management Tools

components. Be sure to install the same components on all the nodes in the cluster.



- 3 When notified to restore databases from backup and reboot the node after completing the installation, click **OK** and complete the Microsoft Exchange wizard.
- 4 If prompted to reboot the node, click **Yes**.
- 5 For Exchange 2000 or Exchange 2003, install the service packs listed in the requirements. When installing service packs enter the following from the command line:

```
SETUP\I386\update.exe /disasterrecovery
```

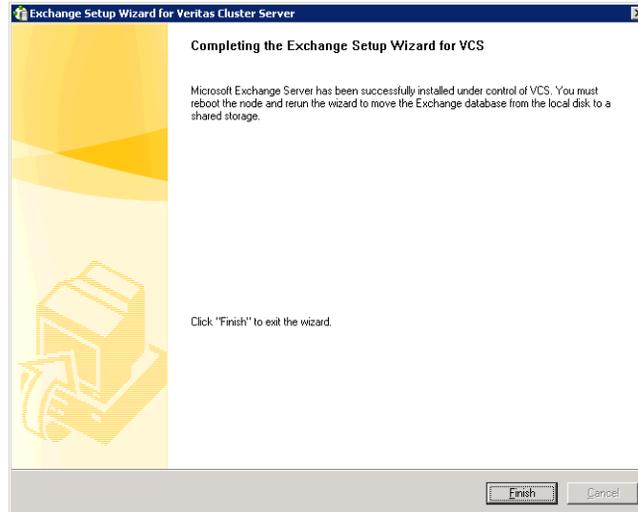
Exchange post-installation: Additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard for Veritas Cluster Server to complete the post-installation phase. This process reverts the node name to the physical name.

To perform Exchange post-installation

- 1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
- 2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 4](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

- 3 Review the information in the Welcome dialog box and click **Next**.
- 4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
- 5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Continue** or **Next**.
- 6 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step. You must do so only if service groups are already configured for the EVS.



If you wish to add the nodes later, you can do so by using the Exchange service group configuration wizard. For more information, see section “Modifying the replication and Exchange service groups”.

- 7 Click **Finish**.
- 8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.
- 9 Changes made during the post-installation steps do not take effect till you reboot the node.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to configure the Exchange service group.

If you wish to add the nodes later, use the Exchange service group configuration wizard. See “[Configuring the Exchange service group for VCS \(secondary site\)](#)” on page 593 for instructions.

Configuring SFW HA: After installing Exchange on secondary site

After installing Exchange on the first node and additional nodes, perform the following procedures on the secondary site:

- ✓ Copying the public cryptographic key of the Exchange virtual server from the secondary site to the primary site
- ✓ Backing up the Exchange disk group on the primary site and restoring it on the secondary site
- ✓ Configuring the Exchange service group for VCS

Copying the .CRK file to the primary site

The .CRK file is the public cryptographic key of the Exchange virtual server. This key is regenerated every time the virtual server is installed.

To copy the .CRK file from the secondary site to the primary site

- 1 On the desktop of any EVS1 system in the secondary site, click **Start > All Programs > Veritas > Veritas Enterprise Administrator**.
- 2 Connect to any system in EVS1 on the primary site.
- 3 Connect to any system in EVS1 on the secondary site.
- 4 From the VEA console, import the EVS1_SG1_DG disk group on the primary and secondary sites.
- 5 If a Service Group exists on the primary site, use Cluster Manager to online the regrep resource at the primary site.
- 6 Mount the EVS1_SG1_REGREP volume on the primary and secondary sites. For example, mount the volume on R:.
- 7 From the system on the secondary site connected to in step 3, run the following commands:

```
C:\>net use Z: \\<system name on primary site>\R$
C:\>copy /Y R:\VCS\Private\RegRep\Exch\EVS1.CRK
Z:\VCS\Private\RegRep\Exch\EVS1.CRK
C:\>net use Z: /d
```

In these commands, Z: is an example of the drive letter assigned to the mounted RegRep volume.

If the copy command asks whether to replace the file, reply yes.

- 8 Unmount the EVS1_SG1_REGREP volume on the primary and secondary sites. For the primary site, offline the resource in Cluster Manager.

Backing up and restoring the Exchange disk group

You must back up and restore the Exchange disk group; a DR installation of Microsoft Exchange does not create Exchange data files. Complete the following tasks before configuring the VCS Exchange service group on the secondary site:

- ✓ On the primary site, back up the Exchange disk group (EVS1_SG1_DG). This backup includes all four volumes.
- ✓ Restore the group in the corresponding location on the secondary site.
- If you navigated to the above procedure from configuring disaster recovery for the first Exchange Virtual Server in an any-to-any configuration, begin the procedure for creating the second Exchange Virtual Server.
- If you navigated to the above procedure from configuring disaster recovery for the second Exchange Virtual Server in an any-to-any configuration, begin the procedure for specifying the common nodes for failover.

Configuring the Exchange service group for VCS (secondary site)

Configuring the Exchange service group involves creating an Exchange service group and defining the attribute values for its resources.

Note: Do not bring the service group online if the service group on the primary site is offline.

Prerequisites

- ✓ You must be a Cluster Administrator.
- ✓ Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- ✓ Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- ✓ Mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
 - Exchange database

- registry changes related to Exchange
- transaction logs for the first storage group
- MTA database

See “[Managing disk groups and volumes](#)” on page 559.

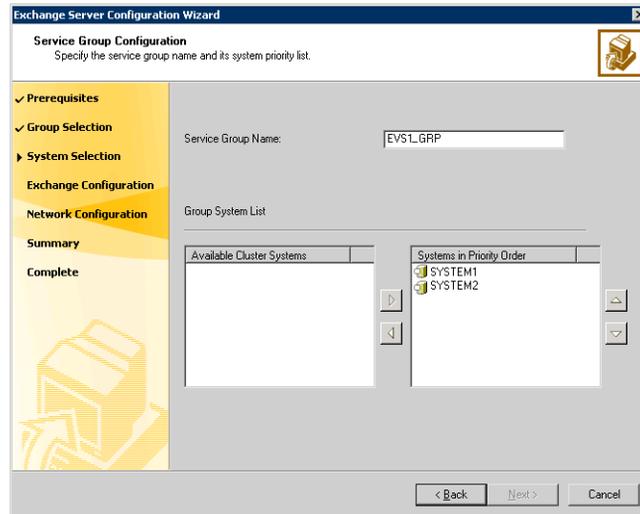
- ✓ Make sure to note the list of the Exchange services and virtual servers that the agent will monitor; the wizard prompts you for this information.
- ✓ Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- ✓ Verify Microsoft Exchange is installed and configured identically on all nodes.

Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to the EVS1_SG1_DG1 disk group

To configure the Exchange service group

- 1 Start the Exchange Server Configuration Wizard. (**Start > All Programs > Veritas > Veritas Cluster Server > Application Agent for Exchange > Configuration Wizard**)
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options dialog box, choose the **Create service group** option and click **Next**.

4 Specify the service group name and system list:

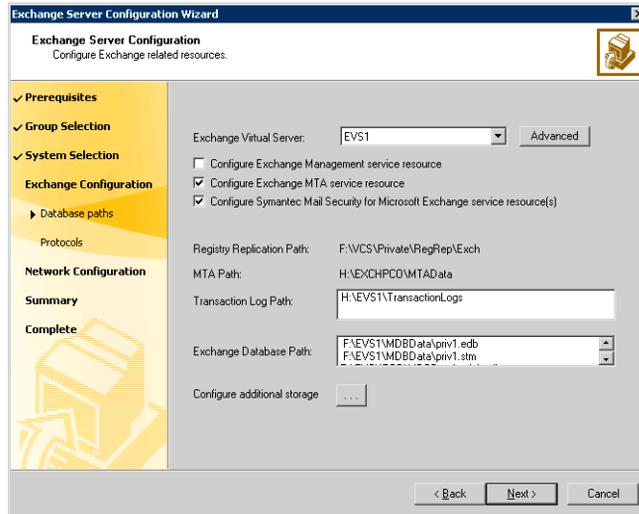


- a Enter a name for the Exchange service group.
- b In the **Available Cluster Systems** box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list.
To remove a system from the service group's system list, select the system in the **Systems in Priority Order** list and click the left arrow.

Note: Microsoft Exchange Server and Microsoft SQL Server 2000 can exist in the same cluster but cannot run on or fail over to the same system. If a SQL Server service group is configured in the cluster, make sure to select a distinct set of systems in the SystemList attribute for each application's service group.

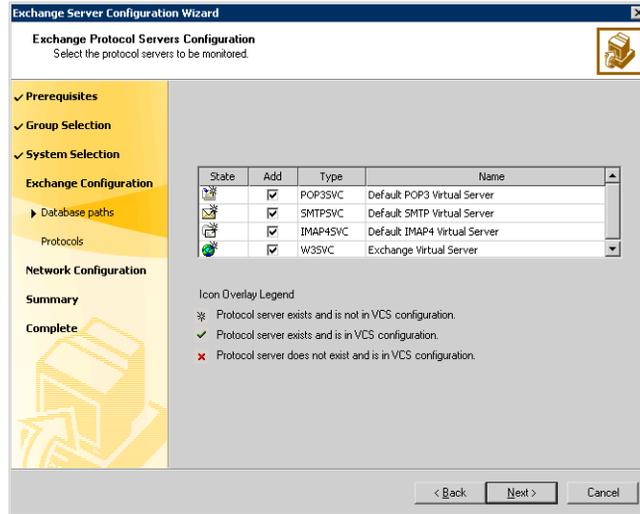
- c To change a node's priority in the service group's system list, select the node in the **Systems in Priority Order** list and click the up and down arrows. The node at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- d Click **Next**. If the configuration is in the read-only mode, the wizard prompts you before changing it to the read-write mode. The wizard starts validating your configuration. Various messages indicate the validation status.

- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:

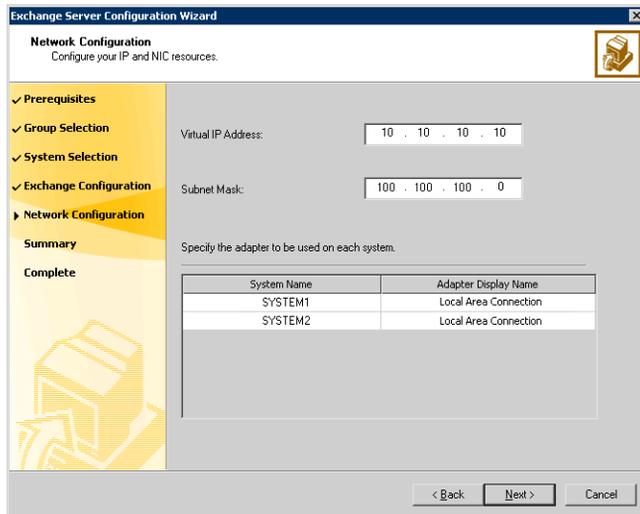


- a Specify the Exchange Virtual Server.
- b Verify the Exchange Database Path.
- c Verify the Transaction Log Path.
- d Specify whether you want to configure the Exchange Management service.
- e Click **Next**.

- 6 Select the check box next to the protocol servers to be monitored and click **Next**.



- 7 Specify information related to the network.



- a The **Virtual IP Address** and the **Subnet Mask** text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.

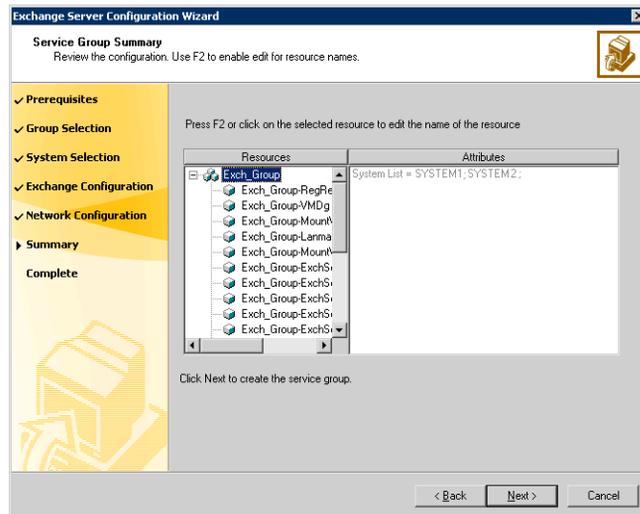
Note that if you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.

- b For each system in the cluster, select the public network adapter name. Select the **Adapter Name** field to view the adapters associated with a node.

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

- c Click **Next**.

- 8 Review the service group configuration and change the resource names, if desired.



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

- a The wizard assigns unique names to resources. Change names of 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.
- b Click **Next**.

- 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.

Configuring DR components on primary and Secondary sites

Proceed to “[Configuring the DR components \(VVR and GCO\)](#)” on page 605 to configure the GCO and VVR components of the DR solution.

Possible task after creating the DR Environment: Adding a new failover node

The following procedure describes how to add an additional node to the cluster at either the primary or secondary site after your disaster recovery environment is in operation. The clusters at each site are not required to have the same number of nodes or the same failover configuration.

Preparing the new node

Install SFW HA on the new system and then add the system to the cluster.

To install SFW HA and add the system to the cluster

- 1 Refer to “[Installing Veritas Storage Foundation HA for Windows](#)” on page 527 for installation instructions.
- 2 Use the **Cluster Operations** option of the VCS Configuration wizard (**Start > All Programs > Veritas > Veritas Cluster Server > VCS Configuration Wizard**) to add the new system to the cluster. If necessary, refer to the *Veritas Cluster Server Administrator's Guide* for information on this procedure.

Preparing the existing DR environment

If you plan to add a failover node to the secondary site, you must temporarily switch the roles of the Primary and Secondary sites so that the current site becomes Primary. This action reverses the direction of replication.

To prepare the existing DR environment

- 1 If you adding the failover node to the cluster at the primary site, proceed directly to [step 2](#). If you are adding a failover node to the secondary site, you must switch the roles of the Primary and Secondary sites. This action reverses the direction of replication.
 - a In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
 - b Click **Switch To**, and click **Remote switch**.
 - c In the **Switch global group** dialog box:
 - Click the cluster at the secondary site you want to switch the group to.
 - Click the specific system where you want to bring the global Exchange service group online.
 - Click **OK**.
- 2 Take the global Exchange service group offline at the current primary site.
- 3 Take the VVR replication service group offline.

Installing Exchange on the new node

Install Exchange on the new node, but do not add the node to the service group SystemList

To prepare the node and install Exchange

- 1 Import the disk group on the new node. Follow the procedure described in [“Managing disk groups and volumes”](#) on page 559.
- 2 From the VEA navigation tree, right-click the RVG for the primary site, and click **Enable Data Access**.
- 3 Run the pre-installation, installation, and post-installation steps described in [“Installing Exchange on additional nodes \(Primary site\)”](#) on page 568 or [“Installing Exchange on additional nodes \(secondary site\)”](#) on page 587; reboot when prompted in these procedures.

Note: During the last step of the post-installation wizard, do *not* check the check box to add the node to the SystemList

Modifying the replication and Exchange service groups

Add the new failover node to the system lists in the Replication and Exchange service groups.

To add the failover node to the system lists

- 1 Bring the replication service group online on an existing cluster node of the current primary site.
- 2 Bring the MountV resources of the corresponding Exchange service group online on the same node.
- 3 Use the **Modify an existing replication service group** option of the Volume Replicator Agent Configuration Wizard (**Start > All Programs > Veritas > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard**) to add the new node to the system list for the replication service group. If necessary, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for information on this procedure.
- 4 Use the **Modify service group** option of the Exchange Server Configuration Wizard (**Start > All Programs > Veritas > Veritas Cluster Server > Application Agent for Exchange > Configuration Wizard**) to add the new node to the system list for the Exchange service group. Check the check box to bring the service group online after the wizard completes. If necessary, refer to the *Veritas Cluster Server Administrator's Guide* for information on this procedure.
- 5 After bringing the Exchange service group online, you must use Exchange System Manager to configure all the database stores to automatically mount on start-up.

Reversing replication direction

If you added a failover node at the original secondary site and migrated the RVG in "[Preparing the existing DR environment](#)" on page 601, move the global Exchange service group back to the original primary site and reverse the direction of replication. These actions switch the Primary and Secondary sites back to their original roles.

To reverse the replication direction

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 In the **Switch global group** dialog box:

- a Click the cluster to switch the group to.
- b Click the specific system where you want to bring the global Exchange service group online.
- c Click **OK**.

Possible task after creating the DR Environment: Adding a new failover node

Configuring the DR components (VVR and GCO)

This chapter covers the following topics:

- [Reviewing the prerequisites](#)
- [Setting up the replicated data sets \(RDS\) for VVR](#)
- [Creating the VVR RVG service group](#)
- [Configuring the global cluster option for wide-area failover](#)
- [Administering global service groups](#)

Note: This chapter applies only if you are not using the Disaster Recovery wizard.

After configuring high availability and Exchange components on the primary and secondary sites, configure the DR components for both sites. This chapter provides information on configuring VVR, the Veritas Cluster Server Enterprise Agent for VVR, and the Global Cluster Option. Refer to the *Veritas Storage Foundation 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 B-1 Task List

Objective	Tasks
“ Reviewing the prerequisites ” on page 607	<ul style="list-style-type: none"> ✓ Verifying HA prerequisites for DR components
“ Setting up the replicated data sets (RDS) for VVR ” on page 607	<ul style="list-style-type: none"> ✓ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites ✓ Using the Setup Replicated Data Set Wizard to create Replicator Log volumes for the primary and secondary sites
“ Creating the VVR RVG service group ” on page 616	<ul style="list-style-type: none"> ✓ 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 621	<ul style="list-style-type: none"> ✓ 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 ✓ Converting the local service group to a global group ✓ Bringing the global service group online
“ Administering global service groups ” on page 625	<ul style="list-style-type: none"> ✓ Beginning VVR replication

Reviewing the prerequisites

- All tasks in “[Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation](#)” on page 515 or “[Deploying SFW HA for Disaster Recovery: Standalone Exchange servers](#)” on page 515 must be completed prior to starting this part of the DR solution.
- Verify that the Replicator log volume does not have a drive letter assigned.

Creating a global cluster environment requires the following conditions:

- All service groups are properly configured and able to come online.
- The clusters must use the same version of VCS.
- The clusters must use the same operating system.
- The names of the clusters at the primary and secondary sites and the virtual IPs associated with them must have been registered in the DNS with reverse lookup.

Setting up the replicated data sets (RDS) for VVR

Set up the Replicated Data Sets (RDS) on the primary site and secondary site. The wizard enables you to configure RDS for both sites.

- ✓ Verify that the data and replicator log 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 Replicator Log volume does not have a DCM.
- ✓ Verify that the cluster disk group is imported on the primary and secondary site

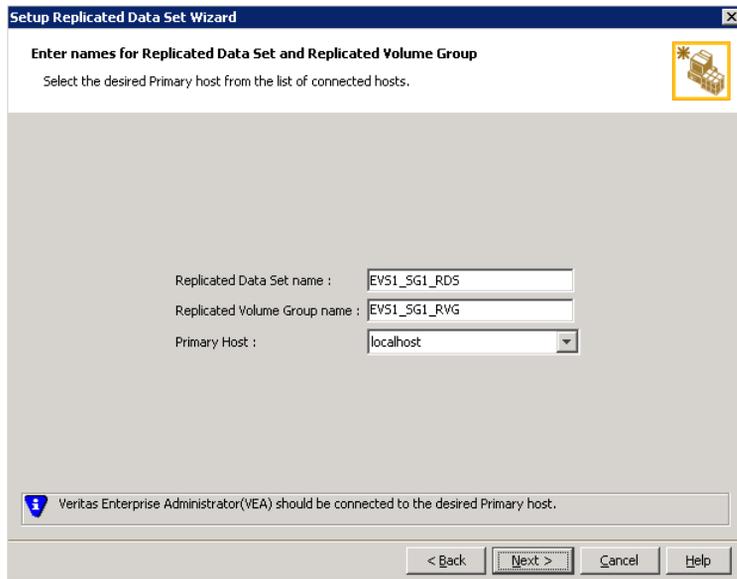
Configuring VVR involves setting up the Replicated Data Sets on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

To create the Replicated Data Set

- 1 From the cluster node on the primary site where the cluster disk group is imported, use the VEA console to launch the Setup Replicated Data Set

Wizard. Right-click **Replication Network** on the Management Host configuration tree, and click **Setup Replicated Data Set**.

- 2 Read the Welcome page and click **Next**.
- 3 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).



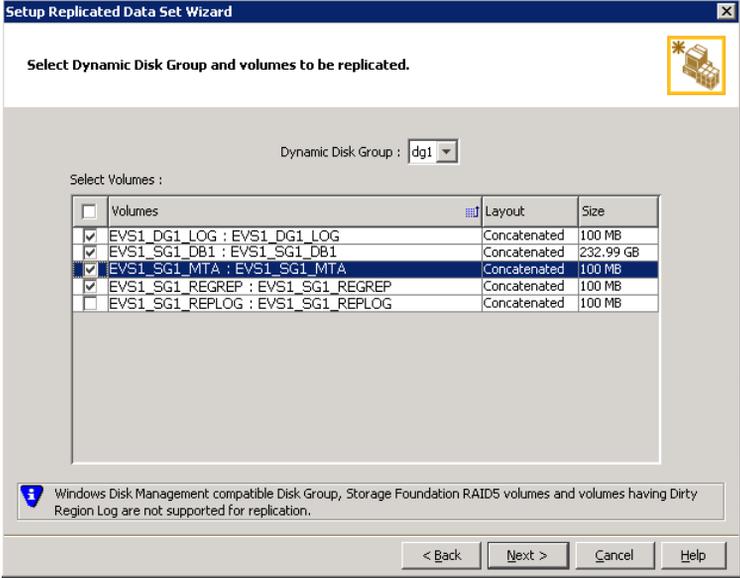
The screenshot shows a window titled "Setup Replicated Data Set Wizard". The main heading is "Enter names for Replicated Data Set and Replicated Volume Group". Below this, it says "Select the desired Primary host from the list of connected hosts." There is a small icon of a server rack in the top right corner. The form contains three input fields: "Replicated Data Set name" with the value "EV51_SG1_RDS", "Replicated Volume Group name" with the value "EV51_SG1_RVG", and "Primary Host" with a dropdown menu showing "localhost". At the bottom, there is a status bar with a blue shield icon and the text "Veritas Enterprise Administrator(VEA) should be connected to the desired Primary host." Below the status bar are four buttons: "< Back", "Next >", "Cancel", and "Help".

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.

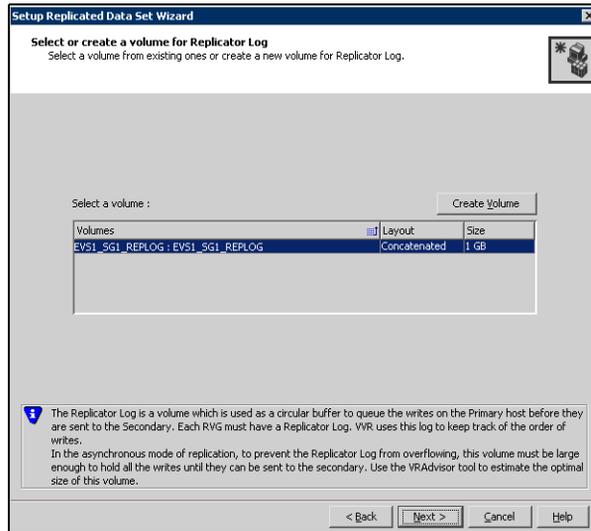


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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 (EVS1_SG1_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.

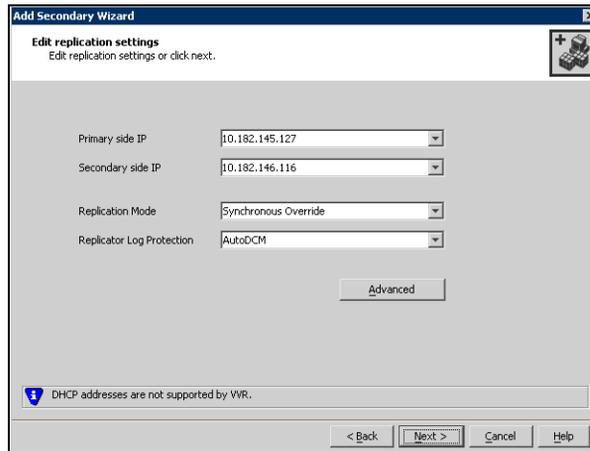
Name	Enter the name for the volume in the Name field.
Size	Enter a size for the volume in the Size field.
Layout	Select the desired volume layout.
Disk Selection	<ul style="list-style-type: none"> ■ Choose Select disks automatically if you want VVR to select the disks for the Replicator Log. ■ Choose 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.



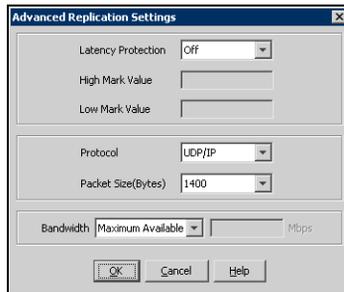
- Enter the virtual IP address for the Primary IP resource that will be used for replication.
- Select or specify an IP address for the Secondary IP resource.
- Specify the replication mode.

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	<p>Enables log protection. If the secondary site is still connected and the Replicator Log is about to overflow, VVR stalls write operations until five percent or 20 megabytes (whichever is smaller) of the space on the Replicator Log becomes available.</p> <p>If the secondary site becomes inactive because of a connection failure or administrative action, VVR disables Replicator Log protection and causes the Replicator Log to overflow.</p>
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

Run the wizard from the system that has the Exchange service group online. The procedure uses EVS1 as an example for all Exchange virtual servers.

Prerequisites:

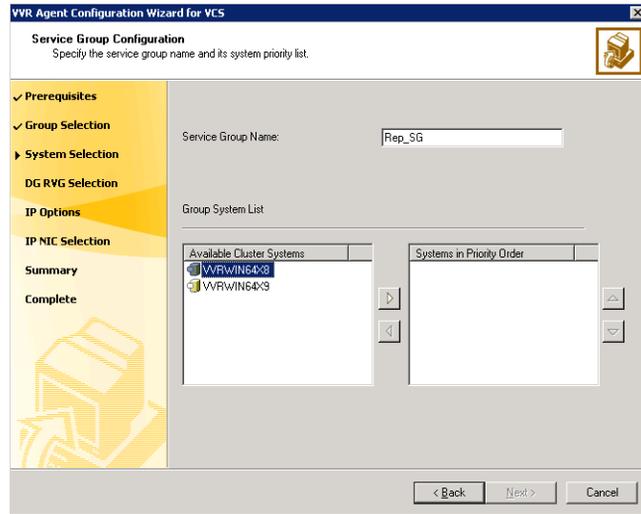
- ✓ Verify that the disk group is imported on the node on which you want to create the Replication Service Group.
- ✓ Verify VCS is running, by running the following command on the host on which the you intend to run the Volume Replicator Agent Configuration Wizard.

```
> hasys -state
```

To create a replication service group

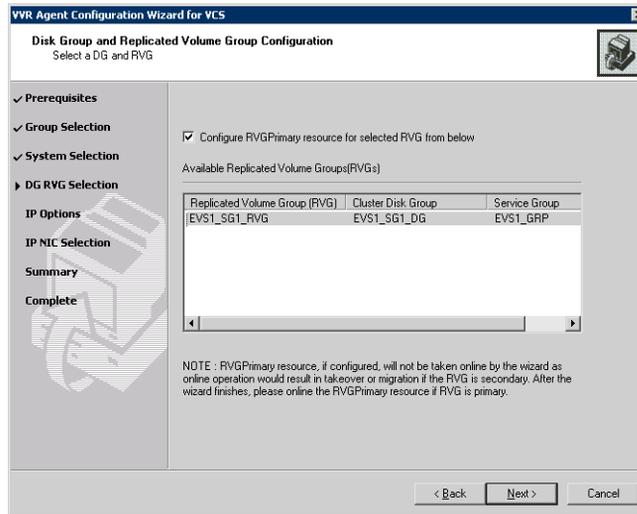
- 1 From the active node of the cluster at the primary site, click **Start > All Programs > Symantec > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the Welcome page, click **Next**.
- 3 In the **Wizard Options** dialog box:
 - a Click **Create a new replication service group**.
 - b Click **Next**.

4 Specify the service group name and system priority list:



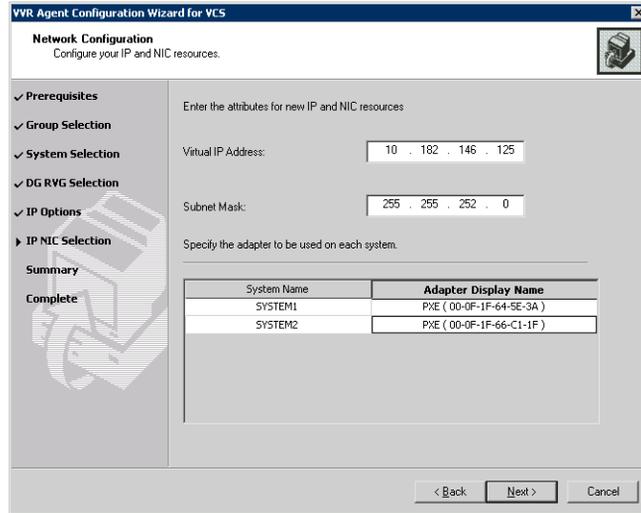
- a Enter the service group name (EVS1_RVG_GRP).
 - b In the **Available Cluster Systems** box, click the nodes on which to configure the service group, and click the right-arrow icon to move the nodes to the service group's system list. Make sure that the set of nodes selected for the replication service group is the same or a superset of nodes selected for the Exchange 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:



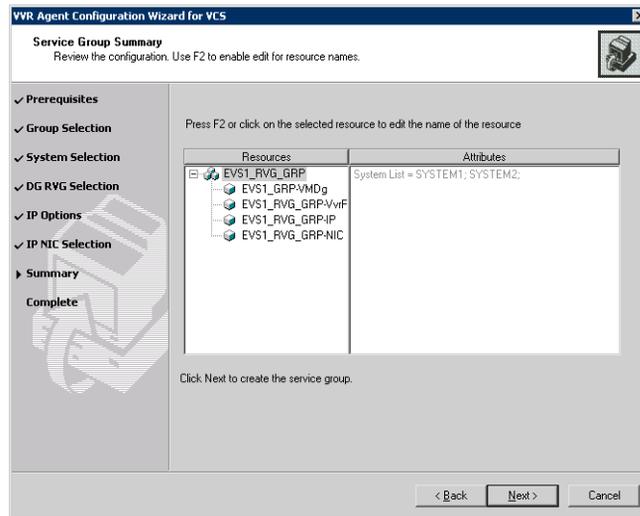
- a Select **Configure RVGPrimary resource for selected RVG**.
 - b Select the replicated volume group for which you want to configure the RVG primary resource.
 - c Click **Next**.
- 7 In the IP Resource Options dialog box, select **Create a new IP resource** and click **Next**.

8 Enter the network information:



- 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: When setting up replication for an application, EVS1-GRP of the Exchange application is dependent on EVS1-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. The process of creating a global cluster environment involves the following tasks:

- ✓ Connecting standalone clusters by adding a remote cluster to a local cluster.
- ✓ Converting the local service group that is common to all the clusters to a global service group.

Use the VCS Java Console or 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 the following conditions:

- All service groups are properly configured and able to come online.
- The service group that will serve as the global group has the same unique name across all applicable clusters.
- The clusters must use the same version of VCS.
- The clusters must use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment
- The names of the clusters at the primary and secondary sites and the virtual IPs associated with them must have been 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 Exchange service group to a global service group

After linking the clusters, use the Global Group Configuration wizard to convert a local Exchange 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**.

Administering global service groups

Administering global groups requires the following conditions:

- A group that will serve as the global group must have the same name across all applicable clusters.
- You must know the user name and password for the administrator to each cluster in the configuration.

Use the VCS Java Console or Web Console to bring a global group online, take a global group offline, or switch a global group on a remote cluster. The section below provides additional procedures for administering global groups from the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more

information on global cluster operations from the Java Console and Web Console.

Note: For remote cluster operations, the user must have the same name and privilege as the user logged on to the local cluster.

Taking a remote global service group offline

Use Cluster Explorer to take a remote global service group offline.

To take a remote global service group offline 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 **Offline**, and click **Remote offline**.
- 3 In the Offline global group dialog box:
 - a Click the remote cluster to take the group offline.
 - b Click the specific system, or click **All Systems**, to take the group offline.
 - c Click **OK**.

Switching a remote service group

Use Cluster Explorer to switch a remote service group.

To switch a remote service group 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 **Switch To**, and click **Remote switch**.
- 3 In the Switch global group dialog box:
 - a Click the cluster to switch the group.
 - b Click the specific system, or click **Any System**, to take the group offline.
 - c Click **OK**.

Deleting a remote cluster

If necessary, use the Remote Cluster Configuration wizard to delete a remote cluster. This operation involves the following tasks:

- Taking the wide area cluster (wac) resource in the ClusterService group offline on the cluster that will be removed from the global environment. For example, to delete cluster C2 from a global environment containing C1 and C2, log on to C2 and take the wac resource offline.
- Removing the name of the specified cluster (C2) from the cluster lists of the other global groups using the Global Group Configuration wizard. Note that the Remote Cluster Configuration wizard in Cluster Explorer automatically updates the cluster lists for heartbeats. Log on to the local cluster (C1) to complete this task before using the Global Group Configuration wizard.
- Deleting the cluster (C2) from the local cluster (C1) through the Remote Cluster Configuration wizard.

Note: You cannot delete a remote cluster if the cluster is part of a cluster list for global service groups or global heartbeats, or if the cluster is in the RUNNING, BUILD, INQUIRY, EXITING, or TRANSITIONING states.

Use Cluster Explorer to take the wide area cluster resource offline, remove a cluster from the cluster list for a global group, and delete a remote cluster from the local cluster.

To take the wide area cluster (wac) resource offline

- 1 From Cluster Monitor, log on to the cluster that will be deleted from the global cluster environment.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **wac** resource under the **Application** type in the **ClusterService** group.
or
Click a service group in the configuration tree, click the **Resources** tab, and right-click the **wac** resource in the view panel.
- 3 Click **Offline**, and click the appropriate system from the menu.

To remove a cluster from a cluster list for a global group

- 1 From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
- 2 Click **Next**.
- 3 Enter the details of the service group to modify:

- a Click the name of the service group.
 - b For global to local cluster conversion, click the left arrow to move the cluster name from the cluster list back to the **Available Clusters** box.
 - c 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:
If the cluster is not running 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**.If the cluster is running 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 using the connected cluster's credentials, or enter new credentials, including the user name, password, and domain.
 - d Click **OK**.
- 5 Click **Next**.
- 6 Click **Finish**.

To delete a remote cluster from the local cluster

- 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 Clusters**.
- 2 Review the required information for the **Remote Cluster Configuration** wizard and click **Next**.

- 3 On the **Wizard Options** dialog box:



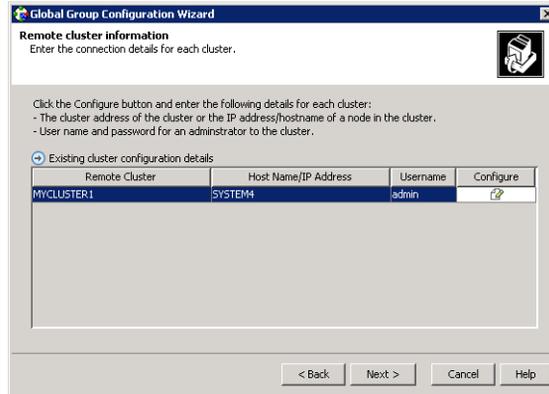
- a Click **Delete Cluster**.
- b Click **Next**.

- 4 In the Delete Cluster dialog box:



- a Click the name of the remote cluster to delete.
- b Click **Next**.

- 5 Review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:



If the cluster is not running 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**.

If the cluster is running 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 the domain.
If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
- d Click **OK**.

- 6 Click **Finish**.

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.
- ◆ If you navigated to the above procedure from [Appendix A, "Deploying Disaster Recovery: Manual implementation of a new Exchange Server"](#)

installation”, [“Configuring DR components on primary and Secondary sites”](#) on page 600, return to [“Possible task after creating the DR Environment: Adding a new failover node”](#) on page 600.

Index

A

- any-to-any HA
 - configuration 181, 187, 266, 267
 - configuring new nodes 268
 - creating a new cluster 187
 - disk space requirements 183, 262
 - new installation 177, 259
 - process overview 177, 259
 - sample configuration 182, 268
 - specifying a common node for failover 231, 308

C

- campus cluster
 - configuration 331
 - defined 321
 - disk space requirements 327
 - Exchange service group, modifying the IP resource 385
 - failover using the forceimport attribute 323
 - forceimport 387, 388
 - handling site failure 387
 - new installation 325, 327
 - overview 321
 - process overview 325
 - site failure 387
- cconfiguration
 - Exchange service group for VCS HA 89
- cluster
 - adding a node
 - any-to-any HA 276
 - configuring
 - any-to-any HA 276
 - DR primary site 535
 - HA 61
 - preparing
 - with the any-to-any option 308, 486

- cluster (continued)
 - verifying configuration
 - campus cluster 387
 - DR 599
 - HA 96, 176, 238, 258, 316
- clusters
 - assigning user privileges 433, 475
- clusterservice group for VCS
 - configuring for any-to-any 284
- configuration
 - any-to-any DR example 473
 - any-to-any HA 181, 187, 266, 267
 - any-to-any HA example 182, 268
 - campus cluster 331
 - DR 523
 - Exchange service group for VCS
 - any-to-any HA 231, 310
 - DR 574
 - DR secondary site 593
 - standalone to HA 169
 - HA 40
 - standalone Exchange server
 - to existing cluster 111
 - to HA 110, 113
 - to new cluster 110
- configuring
 - Exchange service group for VCS
 - campus cluster 379
- configuring VSFW HA
 - after installing Exchange
 - DR secondary site 592
 - prior to installing Exchange
 - DR primary site 526
 - DR secondary site 580
- csg_ip resource 158
- csg_nic resource 158

D

- disaster recovery
 - configuration 470
 - configurations for Exchange 392

- disaster recovery (continued)
 - defined 391
 - new installation 515
 - overview 391
 - see also DR
- disk groups
 - creating
 - any-to-any HA 288
 - campus cluster 358
 - DR primary site 553
 - HA 54, 212
 - standalone to HA 123
 - deporting
 - any-to-any HA 293
 - campus cluster 364
 - DR 559
 - HA 60, 188
 - importing
 - any-to-any HA 293
 - campus cluster 364
 - DR 559
 - HA 60, 188
 - overview
 - any-to-any HA 286
 - campus cluster 355
 - DR primary site 551
 - HA 52, 189
 - standalone to HA 121
- disk groups and volumes
 - configuring
 - any-to-any HA 286
 - campus cluster 355
 - campus cluster for site A and site B 357
 - DR primary site 551
 - HA 52, 189
 - standalone to HA 121
 - managing
 - any-to-any HA 293
 - campus cluster 364
 - DR 559
 - HA 60, 188
- disk space requirements
 - DR 398, 519
 - HA 36
 - standalone to HA 106
- DR
 - adding a new failover node
 - DR 464, 600

- DR (continued)
 - any-to-any
 - sample configuration 473
 - components
 - configuring on primary and secondary sites 464, 600
 - configuring VVR and GCO 605
 - copying the .crk file to the primary site 592
 - defined 391
 - disk space requirements 398, 519
 - new configuration 523
 - new installation 515
 - process overview 515
- driver signing options
 - any-to-any HA 276
 - resetting 52, 121, 195, 276, 340, 414, 535
 - HA 52

E

- Exchange
 - converting standalone servers to HA 103
 - disaster recovery overview 391
 - HA configurations 32
 - high availability overview 31
- Exchange databases
 - moving from standalone to shared storage
 - standalone to HA 160
 - moving to shared storage
 - any-to-any HA 299
 - campus cluster 370
 - DR primary site 565
 - HA 81, 223, 246
- Exchange disk group, backing up and restoring (DR) 593
- Exchange high availability, VCS application
 - agent 32, 322
- Exchange installation
 - additional nodes
 - any-to-any HA 303, 305
 - campus cluster 374, 376
 - DR 446, 570, 589
 - DR primary site 568
 - DR secondary site 444, 587
 - HA 84, 86, 226, 228
 - standalone to HA 163, 166
 - first node
 - any-to-any HA 298
 - campus cluster 366, 369

- Exchange installation (continued)
 - first node
 - DR 564
 - DR primary site 561
 - DR secondary site 442, 484, 581, 585
 - HA 76, 80, 219, 222, 245
 - first node and additional nodes
 - DR secondary site 580
 - new any-to-any nodes
 - any-to-any HA 295
- Exchange post-installation
 - additional nodes
 - any-to-any HA 306
 - campus cluster 377
 - DR 447, 572, 590
 - HA 88, 230
 - standalone to HA 168
 - first node
 - any-to-any HA 299
 - campus cluster 370
 - DR primary site 565
 - DR secondary site 443, 485, 586
 - HA 80, 223, 246
- Exchange pre-installation
 - additional nodes
 - any-to-any HA 303
 - campus cluster 374
 - DR 445, 569, 588
 - HA 85, 227
 - standalone to HA 165
 - first node
 - any-to-any HA 296
 - campus cluster 367
 - DR 562
 - DR secondary site 440, 482, 582
 - HA 77, 220, 243
- Exchange service group
 - configuring
 - any-to-any HA 231, 310
 - campus cluster 379
 - DR 593
 - DR primary site 574
 - HA 89
 - standalone to HA 169
 - configuring for an additional Exchange virtual server
 - HA 252

- Exchange service group (continued)
 - IP resource, modifying 385
 - prerequisites
 - any-to-any HA 232, 252, 310
 - campus cluster 379
 - DR 574, 593
 - HA 90
 - standalone to HA 169

F

- Fire Drill Wizard
 - actions 504
 - deleting the configuration 510
 - overview 501
 - preparing the configuration 505
 - prerequisites for a fire drill 503
 - restoring the prepared configuration 509
 - running a fire drill 508
- forceimport
 - attribute for campus cluster 323
 - defined 323
 - setting after a site failure 387, 388
- forest and domain, preparing for Exchange
 - campus cluster 366
 - DR primary site 560
 - HA 45, 188

G

- GCO
 - adding a remote cluster to a local cluster 621
 - administering global service groups 625
 - bringing a global service group online 625
 - configuring 605
 - configuring for wide-area failover 621
 - converting a local service group to a global service group 623
 - defined 621
 - prerequisites 607, 621
- Global Cluster Option
 - secure configuration 458, 498, 631
- global cluster option
 - overview 621
 - see also GCO
- global service group
 - defined 621

H**HA**

- defined 31
- disk space requirements 36
- installing
 - any-to-any 270
 - campus cluster 334
 - DR primary site 527
 - HA configuration 219
 - standalone Exchange server
 - conversion 115
- IPaddresses required 42
- new configuration 40
- new installation 33
- process overview 33
- Sample configuration 43
- HA configurations, Exchange 32
- high availability
 - defined 31
 - new installation 33
 - overview 31, 32, 322
 - see also HA

N

- network and storage, configuring
 - any-to-any HA 269
 - campus cluster 332
 - DR (primary site) 42, 408, 525
 - HA 187
 - standalone to HA 113

O

- options
 - driver signing 52, 121, 195, 276, 340, 414, 535

P

- prerequisites
 - Exchange service group
 - any-to-any HA 232, 252, 310
 - campus cluster 379
 - DR 574, 593
 - HA 90
 - standalone to HA 169

R

- requirements
 - any-to-any HA installation 183, 262
 - DR new installation 519
 - HA new installation 36
 - see also prerequisites
 - standalone Exchange server to HA 106
 - VSW HA standalone 106
- resetting
 - driver signing options 52, 121, 195, 276, 340, 414, 535

S

- secure clusters
 - assigning user privileges 433, 475
- secure GCO, establishing 458, 498, 631
- Security Services
 - configuring 67, 140, 201, 346, 420, 541
- setting bandwidth
 - using RDS wizard 454, 494
- site failure, forceimport attribute 387
- Solutions Configuration Center
 - context sensitivity 22
 - overview 21
 - running wizards remotely 26
 - starting 22
 - wizard descriptions 26
- standalone Exchange conversion
 - HA disk space requirements 106
- standalone Exchange server
 - adding a new node
 - HA 133
 - adding nodes
 - HA 134
 - adding nodes to an existing cluster
 - HA 150
 - adding to a cluster
 - HA 132
 - configuration
 - HA 110
 - converting
 - to HA 103
 - converting to a “clustered” Exchange server
 - HA 130
 - creating a new cluster
 - HA 134
 - modifying the clusterservice group for VCS
 - HA 158

standalone Exchange server (continued)

- prerequisites for a new cluster
 - HA 133
- prerequisites for adding nodes to an existing cluster
 - HA 149
- process overview
 - HA 103

U

- user privileges
 - assigning 433, 475

V

- VCS Application Agent 32, 322
- vcsweb resource 158
- verifying
 - cluster configuration for campus cluster 387
 - cluster configuration for DR 599
 - cluster configuration for HA 96, 176, 238, 258, 316
- volumes
 - creating
 - any-to-any HA 289
 - campus cluster 360
 - DR 554
 - HA 55, 213
 - standalone to HA 124
 - mounting
 - any-to-any HA 293
 - campus cluster 364
 - DR 559
 - HA 60, 188
 - overview
 - any-to-any HA 286
 - campus cluster 355
 - DR primary site 551
 - HA 52, 189
 - standalone to HA 121
 - unmounting
 - any-to-any HA 293
 - campus cluster 364
 - DR 559
 - HA 60, 188

VSW HA

- installing
 - any-to-any HA 270
 - campus cluster 334
 - DR primary site 527
 - HA 219
 - standalone to HA 115

VVR

- configuring 605
- creating replicator log volumes 607
- creating the VVR RVG service group 616
- prerequisites 607
- setting up the replicated data sets 607

VxSAS

- configuring 51

