

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

Windows Server 2003

5.1 Service Pack 1



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

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.1. Service Pack 1

Document version: 5.1.SP1.0

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice file accompanying this Symantec product for more information on the Third Party Programs.

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.

THE 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 as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

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

Contacting Technical Support

Customers with a current maintenance agreement may access Technical Support information at the following URL:

http://www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:

- Error messages and log files
- Troubleshooting that was performed before contacting Symantec
- Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

<http://customercare.symantec.com>

Customer service

Customer service information is available at the following URL:

<http://customercare.symantec.com>

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and maintenance contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

- Symantec Early Warning Solutions These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
- Managed Security Services These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
- Consulting Services Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
- Educational Services Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Section 1 Introduction

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

About the solutions guides	24
About high availability	24
About campus clusters	25
About replicated data clusters	25
About disaster recovery	25
How this guide is organized	26

Chapter 2 Using the Solutions Configuration Center

About the Solutions Configuration Center	27
Starting the Configuration Center	28
Available options from the Configuration Center	29
About running the Configuration Center wizards	35
Following the workflow in the Configuration Center	36
Solutions wizard logs	39

Section 2 High Availability

Chapter 3 High availability for Exchange: Overview

What is high availability?	43
Why implement a high availability solution?	44
How the VCS application agent makes Microsoft Exchange highly available	44
Typical HA configurations for Exchange	44

Chapter 4

Deploying SFW HA for high availability:
New installation

Tasks for a new HA installation of Microsoft Exchange	46
Reviewing the requirements	48
Disk space requirements	48
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	48
Supported Exchange 2003 versions	49
System requirements for SFW HA	50
Network requirements for SFW HA	50
Permission requirements for SFW HA	51
Additional requirements for SFW HA	52
Best practices for SFW HA	52
Reviewing the configuration	53
IP addresses required during configuration	55
Sample configuration	56
Configuring the storage hardware and network	56
Preparing the forest and domain	58
Configuring SFW HA: Prior to installing Exchange	58
Installing Veritas Storage Foundation HA for Windows	59
Setting Windows driver signing options	59
Installing Storage Foundation HA for Windows	61
Resetting the driver signing options	64
Configuring disk groups and volumes	65
Creating a disk group	67
Creating volumes	69
Managing disk groups and volumes	73
Importing a disk group and mounting a volume	73
Unmounting a volume and deporting a disk group	74
Configuring the cluster	75
Configuring Web console	88
Configuring notification	89
Installing Exchange on the first node	93
Exchange pre-installation: First node	94
Exchange installation: First node	97
Exchange post-installation: First node	97
Moving Exchange databases to shared storage	98
Installing Exchange on additional nodes	103
Exchange pre-installation: Additional nodes	103
Exchange installation: Additional nodes	105
Exchange post-installation: Additional nodes	106
Configuring the Exchange service group for VCS	108
Prerequisites	108

	Verifying the cluster configuration	115
	Configuring the Cluster Management Console connection	116
	Prerequisites for installing the cluster connector	116
	Installing the cluster connector on Windows clusters	117
	Configuring the cluster connector	118
Chapter 5	Deploying SFW HA for high availability: Standalone Exchange servers	
	Reviewing the requirements	123
	Disk space requirements	124
	Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	124
	Supported Exchange 2003 versions	125
	System requirements for SFW HA	126
	Network requirements for SFW HA	126
	Permission requirements for SFW HA	127
	Additional requirements for SFW HA	128
	Best practices for SFW HA	128
	Reviewing the configuration	129
	Configuring the network and storage	134
	Installing Veritas Storage Foundation HA for Windows	136
	Setting Windows driver signing options	136
	Installing Storage Foundation HA for Windows	138
	Resetting the driver signing options	142
	Configuring disk groups and volumes	142
	Considerations for converting existing shared storage to cluster disk groups and volumes	144
	Creating a disk group	145
	Creating volumes	147
	Managing disk groups and volumes	151
	Importing a disk group and mounting a volume	151
	Unmounting a volume and deporting a disk group	152
	Converting the standalone Exchange server into a “clustered” Exchange server	152
	Adding the standalone Exchange server to a cluster	155
	Prerequisites for a new cluster	155
	Creating a new cluster and adding nodes	156
	Configuring Web console	168
	Configuring notification	169
	Prerequisites for adding nodes to an existing cluster	172
	Adding nodes to an existing cluster	173
	Moving Exchange databases to shared storage	179
	Installing Exchange on additional nodes	184

Exchange pre-installation: Additional nodes	185
Exchange installation: Additional nodes	187
Exchange post-installation: Additional nodes	188
Configuring the Exchange service group for VCS	190
Prerequisites	190
Verifying the cluster configuration	197

Chapter 6

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

Reviewing the configuration	202
Any-to-any configuration	202
Configuring failover nodes for additional Exchange instances	203
Sample configuration	203
Reviewing the requirements	205
Disk space requirements	205
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	205
Supported Exchange 2003 versions	206
System requirements for SFW HA	207
Network requirements for SFW HA	207
Permission requirements for SFW HA	208
Additional requirements for SFW HA	209
Best practices for SFW HA	209
Configuring the storage hardware and network	210
Preparing the forest and domain	211
Installing Veritas Storage Foundation HA for Windows	212
Setting Windows driver signing options	212
Installing Storage Foundation HA for Windows	214
Resetting the driver signing options	217
Configuring the cluster	218
Configuring Web console	230
Configuring notification	231
Configuring the first Exchange virtual server	235
Configuring disk groups and volumes	235
Managing disk groups and volumes	242
Importing a disk group and mounting a volume	242
Unmounting a volume and deporting a disk group	243
Installing Exchange on the first node	244
Moving Exchange databases to shared storage	248
Installing Exchange on additional nodes	253
Configuring the Exchange service group for VCS	258
Verifying the cluster configuration	265

Configuring another Exchange virtual server for an any-to-any failover	266
Configuring disk groups and volumes	266
Managing disk groups and volumes	268
Importing a disk group and mounting a volume	268
Unmounting a volume and deporting a disk group	269
Installing Exchange on the first node of an additional Exchange virtual server	270
Moving Exchange databases to shared storage	274
Specifying a common node for failover	278
Configuring the Exchange service group for an additional Exchange virtual server	280
Verifying the cluster configuration	287

Chapter 7

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

Reviewing the requirements	292
Disk space requirements	292
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	293
Supported Exchange 2003 versions	293
System requirements for SFW HA	295
Network requirements for SFW HA	295
Permission requirements for SFW HA	296
Additional requirements for SFW HA	297
Best practices for SFW HA	297
Reviewing the configuration	298
Any-to-any configuration	298
Sample configuration	299
Configuring new nodes: prior to creating additional Exchange virtual server	300
Configuring the storage hardware and network	300
Preparing the forest and domain	302
Installing Veritas Storage Foundation HA for Windows	302
Setting Windows driver signing options	302
Installing Storage Foundation HA for Windows	304
Resetting the driver signing options	308
Configuring the cluster	308
Adding a node to a cluster	308
Configuring disk groups and volumes	314
Creating a disk group	316
Creating volumes	318
Managing disk groups and volumes	322

Importing a disk group and mounting a volume	322
Unmounting a volume and deporting a disk group	323
Installing Exchange on the new nodes	323
Moving Exchange databases to shared storage (EVS2)	328
Installing Exchange on additional nodes	332
Exchange pre-installation: Additional nodes	333
Exchange installation: Additional nodes	335
Exchange post-installation: Additional nodes	336
Specifying a common node for failover	338
Preparing the cluster with the any-to-any option	338
Configuring the Exchange service group for VCS	339
Verifying the cluster configuration	346

Section 3 Campus Cluster

Chapter 8 Campus cluster for Exchange: Overview

What is a campus cluster?	352
Why implement a campus cluster?	352
What is high availability?	352
Why implement a high availability solution?	353
How the VCS application agent makes Microsoft Exchange highly available	353

Chapter 9 Deploying SFW HA for Campus Cluster: New Installation

Reviewing the requirements	357
Disk space requirements	357
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	358
Supported Exchange 2003 versions	358
System requirements for SFW HA	360
Network requirements for SFW HA	360
Permission requirements for SFW HA	361
Additional requirements for SFW HA	362
Best practices for SFW HA	362
Reviewing the configuration	363
Campus cluster failover using the ForceImport attribute	365
Configuring the network and storage	367
Installing Veritas Storage Foundation HA for Windows	369
Setting Windows driver signing options	369
Installing Storage Foundation HA for Windows	371

Resetting the driver signing options	374
Configuring the cluster	375
Configuring Web console	387
Configuring notification	388
Configuring disk groups and volumes	392
Configuring the disks and volumes	394
Creating a dynamic (cluster) disk group	395
Creating a volume	397
Managing disk groups and volumes	401
Importing a disk group and mounting a volume	401
Unmounting a volume and deporting a disk group	402
Preparing the forest and domain	402
Installing Exchange on the first node	402
Exchange pre-installation: First node	404
Exchange installation: First node	406
Exchange post-installation: First node	406
Moving Exchange databases to shared storage	407
Installing Exchange on additional nodes	412
Exchange pre-installation: Additional nodes	412
Exchange installation: Additional nodes	414
Exchange post-installation: Additional nodes	415
Configuring the Exchange service group for VCS	417
Prerequisites	417
Modifying the IP resource in the Exchange service group	424
Verifying the campus cluster: switching the service group	426
Possible tasks after creating the campus cluster	427
Setting the ForceImport attribute to 1 after a site failure	427

Section 4 Replicated Data Clusters

Chapter 10 About Replicated Data Clusters

About Replicated Data Clusters	431
How VCS Replicated Data Clusters work	433
Setting up a Replicated Data Cluster configuration	434
Setting up replication	434
Configuring the service groups	435
Migrating the service group	436

Chapter 11 Configuring Replicated Data Clusters for Exchange

Tasks for configuring Replicated Data Clusters for Exchange Server	440
Reviewing the prerequisites	442
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	443
Supported Exchange 2003 versions	443
System requirements for SFW HA	445
Network requirements for SFW HA	445
Permission requirements for SFW HA	446
Additional requirements for SFW HA	447
Best practices for SFW HA	447
Reviewing the configuration	449
Sample configuration	449
Configuring the storage hardware and network	450
Preparing the forest and domain	451
Installing Veritas Storage Foundation HA for Windows	452
Setting Windows driver signing options	452
Installing Storage Foundation HA for Windows	453
Configuring VxSAS	456
Resetting the driver signing options	459
Configuring the cluster	459
Configuring Web console	471
Configuring notification	472
Configuring cluster disk groups and volumes	476
Creating a cluster disk group	477
Creating volumes	479
Managing disk groups and volumes	484
Importing a disk group and mounting a volume	484
Unmounting a volume and deporting a disk group	485
Installing Exchange on the first node	486
Exchange pre-installation: First node	487
Exchange installation: First node	489
Exchange post-installation: First node	489
Moving Exchange databases to shared storage	490
Installing Exchange on additional nodes	495
Exchange pre-installation: Additional nodes	495
Exchange installation: Additional nodes	497
Exchange post-installation: Additional nodes	498
Configuring the Exchange service group for VCS	500
Prerequisites	500
Creating the primary system zone	507
Verifying the installation in the primary zone	508

Creating a parallel environment in the secondary zone	509
Adding the systems in the secondary zone to the cluster	510
Setting up the Replicated Data Sets (RDS)	516
Configuring a hybrid RVG service group for replication	528
Creating the RVG service group	529
Configuring the RVG service group for RDC replication	530
Configuring the RVG Primary resources	535
Configuring the primary system zone for the RVG	537
Setting a dependency between the service groups	538
Adding the nodes from the secondary zone to the RDC	538
Adding the nodes from the secondary zone to the RVG service group	538
Configuring secondary zone nodes in the RVG service group	540
Configuring the IP resources for fail over	540
Adding the nodes from the secondary zone to the Exchange Server service group	542
Configuring the zones in the Exchange Server service group	543
Verifying the RDC configuration	544
Bringing the service group online	544
Switching online nodes	544
Additional instructions for GCO disaster recovery	545

Section 5 Disaster Recovery

Chapter 12 Disaster recovery for Exchange: Overview

What is a disaster recovery solution?	549
Why implement a DR solution?	549
Typical DR configurations for Exchange	550

Chapter 13 Deploying Disaster Recovery: New Exchange Server installation

Tasks for deploying a disaster recovery active-passive configuration of Microsoft Exchange	553
Reviewing the requirements	556
Disk space requirements	556
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	556
Supported Exchange 2003 versions	557
System requirements for SFW HA	559
Network requirements for SFW HA	559

Permission requirements for SFW HA	560
Additional requirements for SFW HA	561
Best practices for SFW HA	561
Reviewing the configuration	562
Supported disaster recovery configurations for service group dependencies	565
Configuring the storage hardware and network	565
Managing disk groups and volumes	567
Importing a disk group and mounting a volume	567
Unmounting a volume and deporting a disk group	568
Preparing the forest and domain	568
Setting up the secondary site: Installing SFW HA and configuring a cluster	569
Installing SFW HA	569
Setting Windows driver signing options	569
Installing Storage Foundation HA for Windows	571
Resetting the driver signing options	575
Configuring the cluster	575
Configuring Web console	588
Configuring notification	589
Verifying your primary site configuration	592
Setting up your replication environment	593
Setting up security for VVR	593
Requirements for EMC SRDF array-based hardware replication	596
Requirements for Hitachi TrueCopy array-based hardware replication	598
Assigning user privileges (secure clusters only)	600
Configuring disaster recovery with the DR wizard	601
Cloning the storage on the secondary site using the DR wizard (VVR replication option)	605
Creating temporary storage on the secondary site using the DR wizard (array-based replication)	610
Installing Exchange on the first node with DR option (secondary site) ...	613
Prerequisites for installing Exchange Server	614
Exchange pre-installation on first node (secondary site)	615
Exchange installation on first node (secondary site)	617
Exchange post-installation on first node (secondary site)	618
Installing Exchange on additional nodes (secondary site)	619
Exchange pre-installation: Additional nodes	619
Exchange installation: Additional nodes	621
Exchange post-installation: Additional nodes	622

Cloning the service group configuration on to the secondary site using the DR wizard	624
Configuring replication and global clustering	627
Configuring VVR replication and global clustering	627
Configuring EMC SRDF replication and global clustering	635
Configuring Hitachi TrueCopy replication and global clustering	638
Configuring global clustering only	641
Verifying the disaster recovery configuration	643
Establishing secure communication within the global cluster (optional)	645
Adding multiple DR sites (optional)	647
Recovery procedures for service group dependencies	647
Possible task after creating the DR environment:	
Adding a new failover node to a VVR environment	651
Preparing the new node	651
Preparing the existing DR environment	651
Installing Exchange on the new node	652
Modifying the replication and Exchange service groups	652
Reversing replication direction	653

Chapter 14 Deploying SFW HA for Disaster Recovery: Configuring any-to-any failover

Tasks for deploying a disaster recovery any-to-any configuration of Microsoft Exchange	656
Reviewing the configuration	658
Disaster recovery configuration	658
Any-to-any configuration	660
Sample any-to-any configuration for disaster recovery	661
Configuring disaster recovery for the first Exchange virtual server	662
Verifying your primary site configuration for an additional Exchange virtual server	663
Adding the user to the service group (secure clusters only)	663
Configuring disaster recovery for the second Exchange virtual server	664
Cloning the storage on the secondary site using the DR wizard	664
Installing Exchange on the first node of an additional EVS (secondary site)	665
Exchange pre-installation on first node of an additional EVS (secondary site)	666
Exchange installation on first node of an additional EVS (secondary site)	668
Exchange post-installation on first node of an additional EVS (secondary site)	669
Specifying a common node for failover	670

Cloning the service group configuration on to the secondary site using the DR wizard	672
Configuring replication and global clustering	672
Verifying the disaster recovery configuration	672
Establishing secure communication within the global cluster (optional)	674
Possible tasks after creating the DR environment	676

Chapter 15 Testing fault readiness by running a fire drill

About disaster recovery fire drills	677
About the Fire Drill Wizard	678
About Fire Drill Wizard general operations	678
About Fire Drill Wizard operations in a VVR environment	679
About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment	681
About post-fire drill scripts	683
Tasks for configuring and running fire drills	684
Prerequisites for a fire drill	686
Prerequisites for a fire drill in a VVR environment	686
Prerequisites for a fire drill in a Hitachi TrueCopy environment	687
Prerequisites for a fire drill in an EMC SRDF environment	688
Preparing the fire drill configuration	688
System Selection panel details	691
Service Group Selection panel details	691
Secondary System Selection panel details	691
Disk Selection panel details	691
Hitachi TrueCopy Path Information panel details	692
HTCSnap Resource Configuration panel details	693
SRDFSnap Resource Configuration panel details	693
Fire Drill Preparation panel details	694
Running a fire drill	694
Post fire drill operations panel details	696
Recreating a fire drill configuration that has changed	696
Restoring the fire drill system to a prepared state	699
Deleting the fire drill configuration	700
.....	701

Section 6 Appendices

Appendix A	Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation	
	Tasks for configuring disaster recovery	706
	Reviewing the configuration	709
	Verifying the primary site configuration	712
	Setting up the SFW HA environment (secondary site)	712
	Setting up security for VVR	713
	Installing Exchange on the first node and additional nodes (secondary site)	716
	Installing Exchange on the first node with DR Option (secondary site)	716
	Exchange pre-installation on first node (secondary site)	718
	Exchange installation on first node (secondary site)	719
	Exchange post-installation on first node (secondary site)	720
	Installing Exchange on additional nodes (secondary site)	721
	Exchange pre-installation: Additional nodes	722
	Exchange installation: additional nodes	724
	Exchange post-installation: additional nodes	725
	Copying the .CRK file to the primary site	727
	Backing up and restoring the Exchange disk group	727
	Configuring the Exchange service group for VCS (secondary site)	728
	Prerequisites	728
	Verifying the cluster configuration	734
	About configuring the DR components (VVR and GCO)	735
	Reviewing the prerequisites for configuring DR	736
	Setting up the replicated data sets (RDS) for VVR	736
	Creating the VVR RVG service group	748
	Configuring the global cluster option for wide-area failover	751
	Prerequisites	751
	Linking clusters:	
	Adding a remote cluster to a local cluster	752
	Converting a local Exchange service group to a global service group	753
	Bringing a global service group online	755
	Administering global service groups	757
	Deleting a remote cluster	758
	Establishing secure communication within the global cluster (optional)	762
	Possible task after creating the DR Environment:	

Adding a new failover node	764
Preparing the new node	764
Preparing the existing DR environment	764
Installing Exchange on the new node	765
Modifying the replication and Exchange service groups	765
Reversing replication direction	766

Index	767
-------------	-----

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 23
- [Chapter 2, “Using the Solutions Configuration Center”](#) on page 27

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange

This chapter contains the following topics:

- [“About the solutions guides”](#) on page 24
- [“About high availability”](#) on page 24
- [“About campus clusters”](#) on page 25
- [“About replicated data clusters”](#) on page 25
- [“About disaster recovery”](#) on page 25
- [“How this guide is organized”](#) on page 26

About the solutions guides

This guide describes the following solutions for Exchange:

- High availability (HA)
- Campus clusters
- Replicated data clusters
- Disaster recovery (DR)

Information on solutions for Quick Recovery and Microsoft Clustering are in *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering 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 campus clusters

Campus clusters are clusters 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 and contains mirrored data of the storage on the other array.

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

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

About replicated data clusters

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

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

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

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

This guide is organized to follow the workflow in the Solutions Configuration Center.

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

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

When setting up a site for disaster recovery, you first follow the instructions in the appropriate chapter in the high availability section and then continue with the 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](#)
- [Solutions wizard logs](#)

About the Solutions Configuration Center

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

- Microsoft Exchange Server 2003 and 2007
- Microsoft SQL Server 2000, 2005, and 2008
- Enterprise Vault Server (high availability new server and disaster recovery solutions)
- Additional applications

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

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

- Campus cluster disaster recovery, including the following:
 - Campus cluster using Veritas Cluster Server (SFW HA)
 - Campus cluster using Microsoft clustering
- Wide area disaster recovery involving multiple sites
- Quick Recovery for on-host recovery from logical errors in application data (available for Microsoft Exchange 2003 and 2007 and for Microsoft SQL Server 2005 and 2008)
- Fire drill to test the fault readiness of a disaster recovery environment

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

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

Starting the Configuration Center

You can start the Configuration Center in the following ways:

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

Available options from the Configuration Center

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

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

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

Figure 2-1 Solutions Configuration Center for Microsoft Exchange

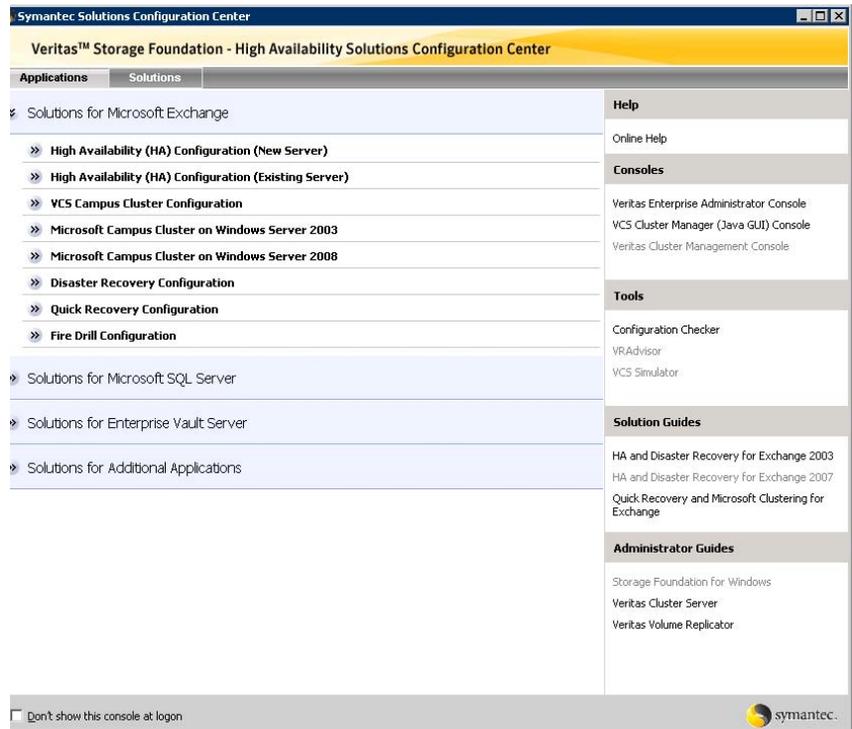


Figure 2-2 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 2-2 Solutions Configuration Center for Microsoft SQL Server

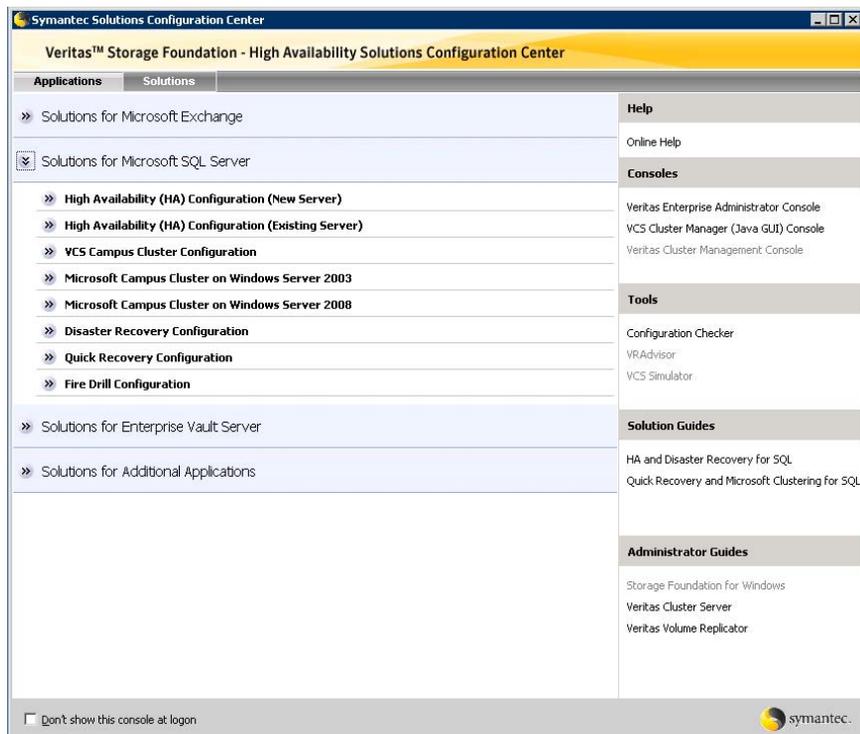


Figure 2-3 shows the choices available when you click Solutions for Enterprise Vault Server.

Figure 2-3 Solutions Configuration Center for Enterprise Vault Server

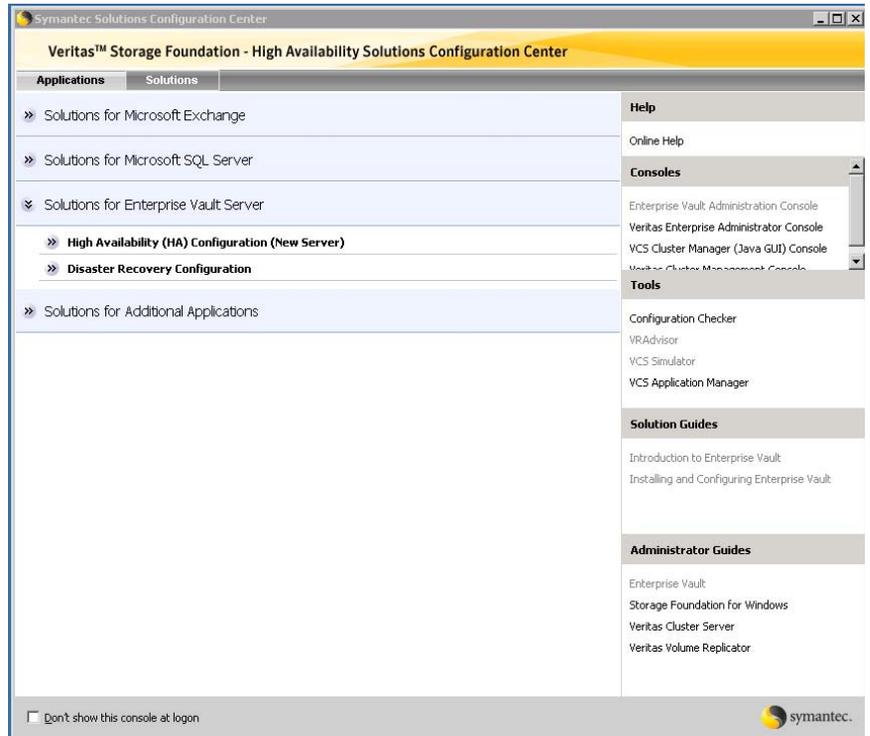
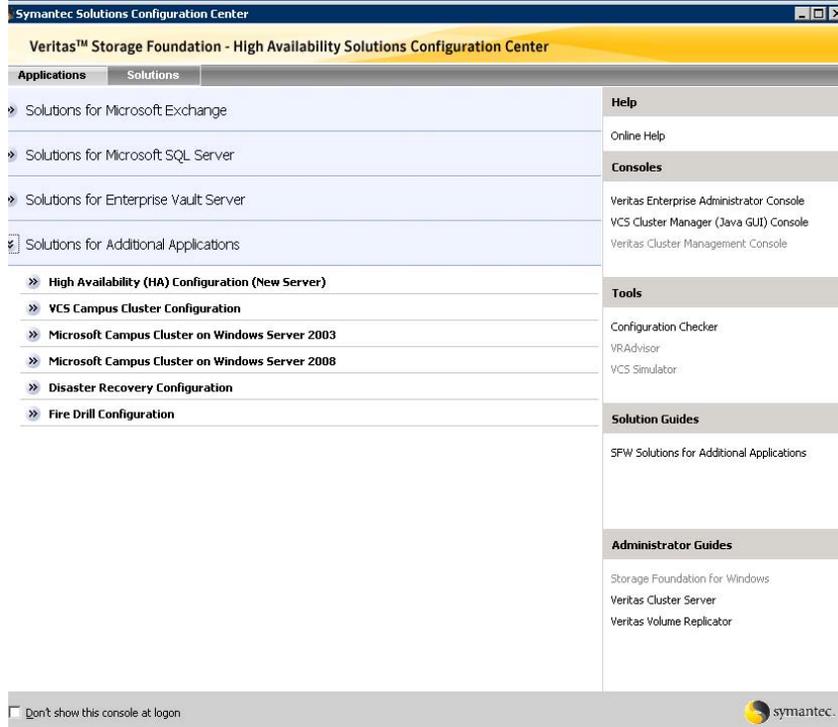


Figure 2-4 shows the choices available when you click Solutions for Additional Applications.

Figure 2-4 Solutions Configuration Center for additional applications



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

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

Figure 2-5 Context-sensitive step for Exchange



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

Figure 2-6 Context-sensitive step for SQL Server

4 Configure the SQL Server service group

The SQL Server Configuration Wizard helps you create a SQL Server service group to make the SQL Server instance highly available.

[Read how to create a new SQL Server service group](#)

Run this wizard to create a new SQL Server 2000 or SQL Server 2005 service group.

ⓘ Wizard must be run locally

SQL Server Configuration Wizard

Run this wizard to create a new SQL Server 2008 service group.

ⓘ Wizard must be run locally

SQL Server 2008 Configuration Wizard

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

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

4 Configure the Enterprise Vault service group

The Enterprise Vault Cluster Setup Wizard helps you create a service group to make the Enterprise Vault instance highly available.

Note: The wizard, by default, stores the Indexing and Shopping services data at the same location as that of the MSMQ data. You can modify this location after configuring the Enterprise Vault Server.

[Read how to create the Enterprise Vault service group](#)

ⓘ Wizard must be run locally

Enterprise Vault Cluster Setup Wizard

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

Figure 2-8 Context-sensitive step for additional applications

4 Configure the service group

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

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

Wizard must be run locally

File Share Configuration Wizard

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

Wizard must be run locally

Print Share Configuration Wizard

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

Wizard must be run locally

IIS Configuration Wizard

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

Wizard must be run locally

MSVirtual Machine Configuration Wizard

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

Wizard must be run locally

Oracle Agent Configuration Wizard

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

Wizard must be run locally

Application Configuration Wizard

About running the Configuration Center wizards

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

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

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

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

New Dynamic Disk Group Wizard	Launched from the Veritas Enterprise Administrator console
New Volume Wizard	Launched from the Veritas Enterprise Administrator console
Exchange Setup Wizard	Installs and configures Exchange for the high availability environment If Exchange is already installed, refer to the documentation for further instructions.
Exchange Configuration Wizard	Configures the service group for Exchange high availability

SQL Server Configuration Wizard	Configures the service group for SQL Server 2000 or SQL Server 2005 high availability You must first install SQL Server on each node according to the instructions in the documentation.
SQL Server 2008 Configuration Wizard	Configures the service group for SQL Server 2008 high availability You must first install SQL Server on each node according to the instructions in the documentation.
Enterprise Vault Cluster Setup Wizard	Configures the service group for Enterprise Vault Server high availability.
MSDTC Wizard	Configures an MSDTC Server service group for SQL Server 2000, 2005, or 2008 environments.
MSMQ Configuration Wizard	Configures a Microsoft Message Queuing (MSMQ) service group.

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

File Share Configuration Wizard	Configures FileShare for high availability.
Print Share Configuration Wizard	Configures PrintShare for high availability.
IIS Configuration Wizard	Configures IIS for high availability.
MSVirtual Machine Configuration Wizard	Configures MS Virtual Machine for high availability.
Oracle Agent Configuration Wizard	Configures Oracle for high availability
Application Configuration Wizard	Configures any other application service group for which application-specific wizards have not been provided.

Following the workflow in the Configuration Center

During the multi-step High Availability Configuration workflow, you may find it helpful to run an SFW HA client on another system and leave the Configuration Center open on that system. In this way, you can see what step comes next, drill

down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format. When setting up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration.

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

Figure 2-9 Workflow for configuring Exchange high availability

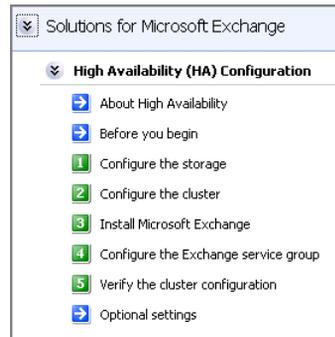


Figure 2-10 shows the high-level overview of the workflow steps for configuring high availability for SQL Server from the Solutions Configuration Center.

Figure 2-10 Workflow for configuring SQL Server high availability

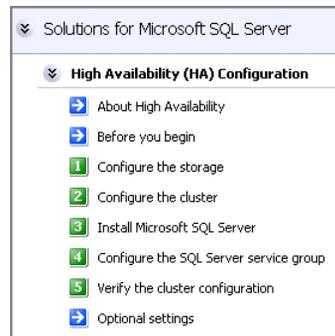


Figure 2-11 shows the high-level overview of the workflow steps for configuring high availability for Enterprise Vault Server from the Solutions Configuration Center.

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

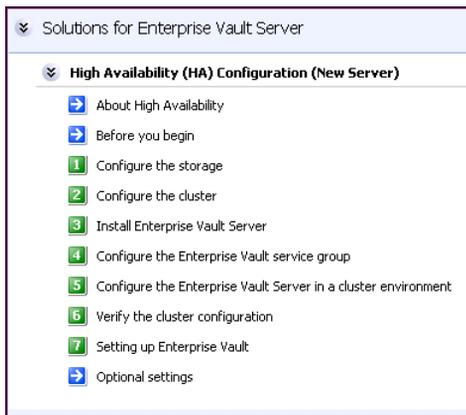
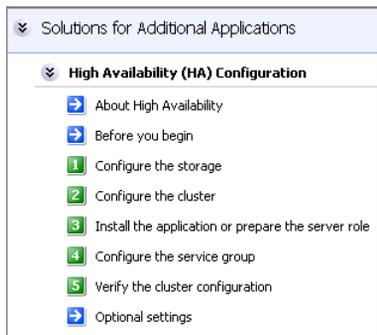


Figure 2-12 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

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



Solutions wizard logs

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

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

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

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

For Windows Server 2003:

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

For Windows Server 2008:

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


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 chapters in this section to install and configure a clustered Exchange environment using Veritas Storage Foundation HA for Windows.

This section contains the following chapters:

- [Chapter 3, “High availability for Exchange: Overview”](#) on page 43
- [Chapter 4, “Deploying SFW HA for high availability: New installation”](#) on page 45
- [Chapter 5, “Deploying SFW HA for high availability: Standalone Exchange servers”](#) on page 121
- [Chapter 6, “Deploying SFW HA for high availability: Configuring a new any-to-any failover”](#) on page 199
- [Chapter 7, “Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover”](#) on page 289

High availability for Exchange: Overview

This chapter contains the following topics:

- [“What is high availability?”](#) on page 43
- [“Why implement a high availability solution?”](#) on page 44
- [“How the VCS application agent makes Microsoft Exchange highly available”](#) on page 44
- [“Typical HA configurations for Exchange”](#) on page 44

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.

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 contains the following topics:

- [“Tasks for a new HA installation of Microsoft Exchange”](#) on page 46
- [“Reviewing the requirements”](#) on page 48
- [“Reviewing the configuration”](#) on page 53
- [“Configuring the storage hardware and network”](#) on page 56
- [“Preparing the forest and domain”](#) on page 58
- [“Configuring SFW HA: Prior to installing Exchange”](#) on page 58
- [“Configuring disk groups and volumes”](#) on page 65
- [“Configuring the cluster”](#) on page 75
- [“Managing disk groups and volumes”](#) on page 73
- [“Installing Exchange on the first node”](#) on page 93
- [“Moving Exchange databases to shared storage”](#) on page 98
- [“Installing Exchange on additional nodes”](#) on page 103
- [“Configuring the Exchange service group for VCS”](#) on page 108
- [“Verifying the cluster configuration”](#) on page 115
- [“Configuring the Cluster Management Console connection”](#) on page 116

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 environment for Exchange. This environment involves an active-passive configuration with one-to-one failover capabilities.

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. You must set up the secondary site only after completing configuring high availability at the primary site. See [“Deploying Disaster Recovery: New Exchange Server installation”](#) on page 551.

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

Table 4-1 Task list: Exchange active-passive HA configuration

Objective	Tasks
“Reviewing the requirements” on page 48	Verify hardware and software prerequisites
“Reviewing the configuration” on page 53	Understanding a typical active-passive Exchange configuration in a two-node cluster
“Configuring the storage hardware and network” on page 56	<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
“Configuring SFW HA: Prior to installing Exchange” on page 58	<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

Table 4-1 Task list: Exchange active-passive HA configuration (Continued)

Objective	Tasks
“Configuring disk groups and volumes” on page 65	<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 75	<ul style="list-style-type: none"> ■ Verify static IP addresses and name resolution configured for each node ■ Configure cluster components using the VCS Cluster Configuration Wizard (VCW)
“Preparing the forest and domain” on page 58	Set up the forest and domain prior to the Exchange installation
“Installing Exchange on the first node” on page 93	<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 98	Move 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 103	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 108	Create the Exchange service group using the Exchange Configuration Wizard for Veritas Cluster Server.
“Verifying the cluster configuration” on page 115	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.

Table 4-2 estimates disk space requirements for SFW HA.

Table 4-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

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

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/302144>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agent for Microsoft Exchange.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2003 versions

The following table lists the Microsoft Exchange Server 2003 versions supported with SFW HA 5.1 Service Pack 1.

Table 4-4 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Memory must be a minimum 256 MB of RAM per server for Exchange 2003; refer to the Microsoft documentation for more information.
- 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. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
See "[Best practices for SFW HA](#)" on page 52.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each Exchange Virtual Server (EVS).

- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone 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 Dynamic DNS (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 for SFW HA

The following permissions are required:

- 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 on 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 Exchange virtual server computer object 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, 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 for SFW HA

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 Server 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 2003.
- 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.

Best practices for SFW HA

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.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

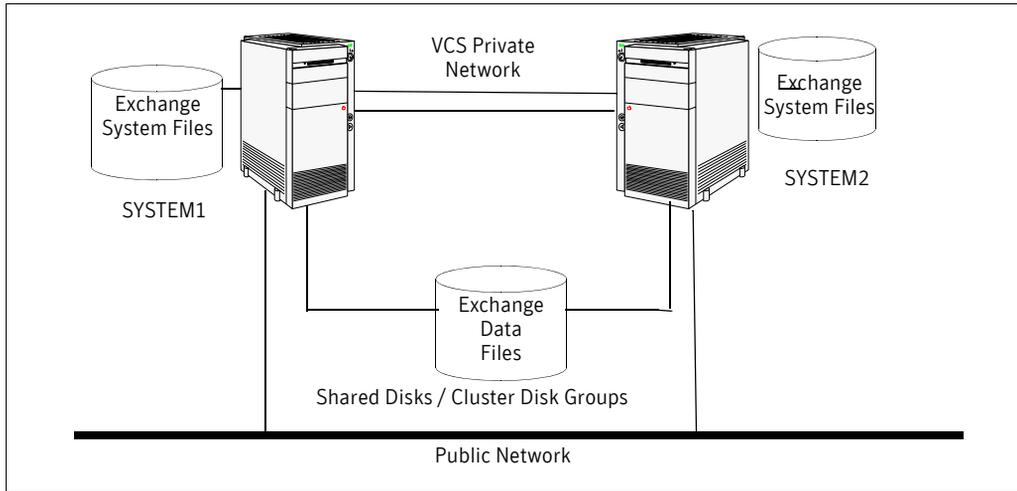
Reviewing the configuration

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.

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.

[Figure 4-1](#) illustrates an active-passive failover configuration with an Exchange virtual server. In this case, EVS1 can fail over from SYSTEM1 to SYSTEM2.

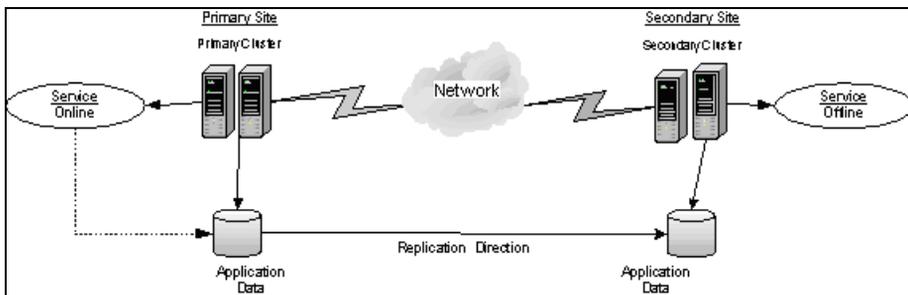
Figure 4-1 Active-passive failover configuration



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.

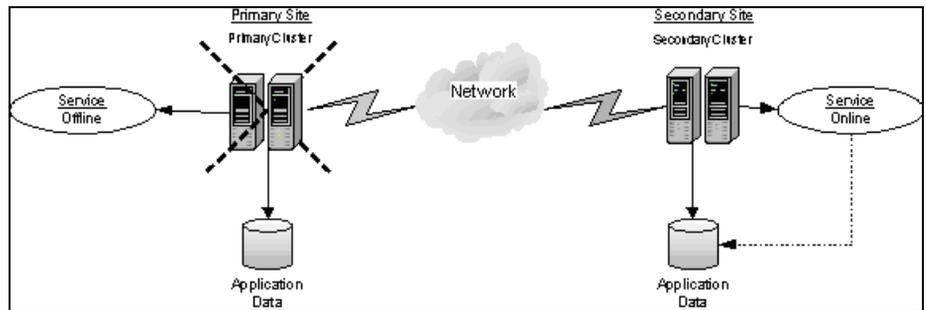
[Figure 4-2](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

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

Sample configuration

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

Table 4-6 Sample configuration

Name	Object
SYSTEM1, SYSTEM2	Physical node names
EVS1	Microsoft Exchange virtual server
EVS1_SG1	Microsoft Exchange service group
EVS1_SG1_DG EVS1_SHARED_DG	Cluster disk group names
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 contains 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 and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network 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 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, 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 59.
- Set up a VCS environment.
See “[Configuring the cluster](#)” on page 75
- Create the required disk groups and volumes.
See “[Configuring disk groups and volumes](#)” on page 65
See “[Managing disk groups and volumes](#)” on page 73.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 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, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

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

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 4-7 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 60.

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 Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The Select Product screen appears.
- 3 Review the links on the Select Product screen.
Links on this screen access Late Breaking News, the Configuration Checker, as well as begin the process to install Storage Foundation HA for Windows. Click on **Read Late Breaking News** for the latest information about updates, patches, and software issues regarding this release.
- 4 Click **Storage Foundation HA 5.1 SP1 for Windows**.
- 5 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 6 Review the Welcome message and the listed prerequisites. Ensure that any of the listed prerequisites applicable to your installation are met prior to proceeding.
Click **Next**.
- 7 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 AGREE TO the terms of the license agreement**, and then click **Next**.
- 8 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.
- 9 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.
- 10 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. Be sure to select the following as appropriate for your installation.

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, select the option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

11 Select the following 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. When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

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

- 12 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.
- 13 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.
- 14 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages.

If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning:

The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If applicable to your installation, perform the above procedure.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

If applicable to your installation, perform the above procedure.

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation 5.1 for Windows (Server Components) with the MSCS option, you may want MSCS Quorum Arbitration Time (Min. and Max) to be adjusted to ensure optimal functionality with Veritas dynamic volumes with MSCS.

For additional information, see the *Storage Foundation for Windows Administrator Guide* for details.

If applicable to your installation, perform the above procedure.

- 15 When finished reviewing the message or messages, click **OK**.
- 16 The Summary screen appears displaying an Install report. Review the information in the Install report. Click **Back** to make changes, if necessary. Click **Install** if information is validated.
- 17 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.
- 18 When the installation completes, review the summary screen and click **Next**.
- 19 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.
- 20 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 21 Review the log files and click **Finish**.
- 22 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 Veritas Enterprise Administrator (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.

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). Refer to the *Veritas Storage Foundation Administrator's Guide* for more information.

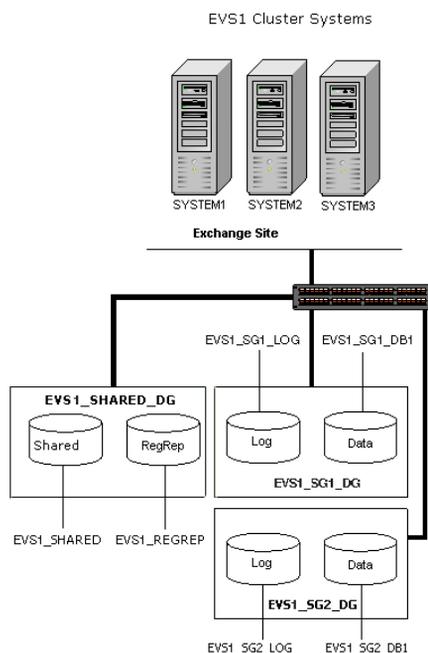
Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of volumes or 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.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

Exchange storage group EVS1_SHARED_DG contains the following 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: A disaster recovery configuration with VVR requires a Storage Replicator Log (SRL) volume (EVS1_REPLOG) for each disk group that contains volumes that are replicated. You can create the SRL volume now or you can create it later when you run the Disaster Recovery Wizard. If you create it later, ensure that you allow sufficient disk space for this volume. For more about VVR planning, see the *Veritas Volume Replicator, Administrator's Guide*.

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

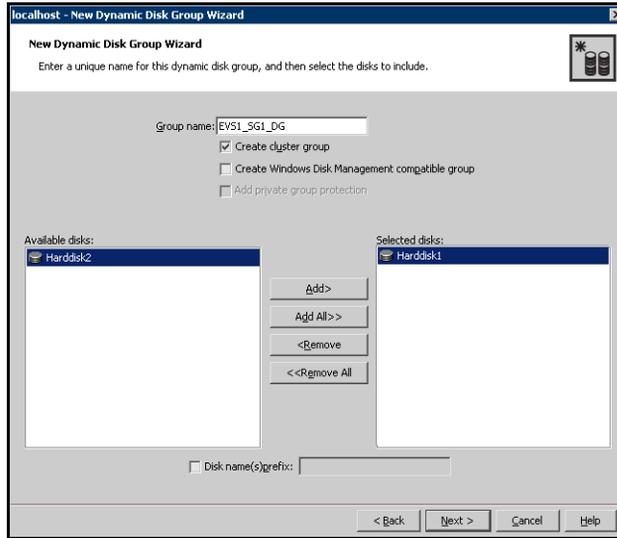
Creating a disk group

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

- 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.

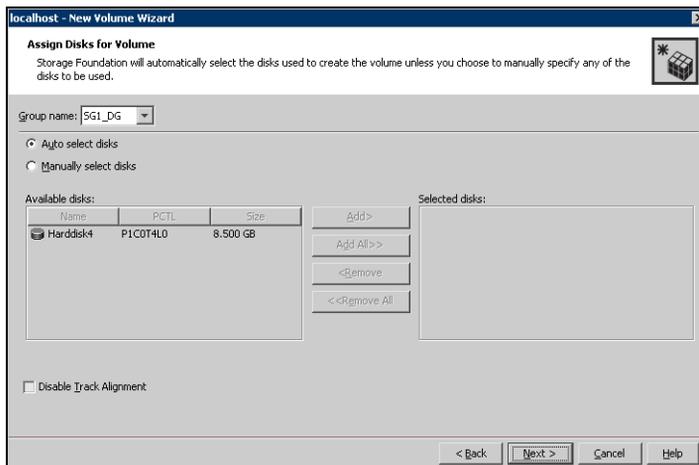
For Disaster Recovery configuration using VVR, the Disaster Recovery Configuration wizard can create the Storage Replicator Log volumes for you.

Note: Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

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.

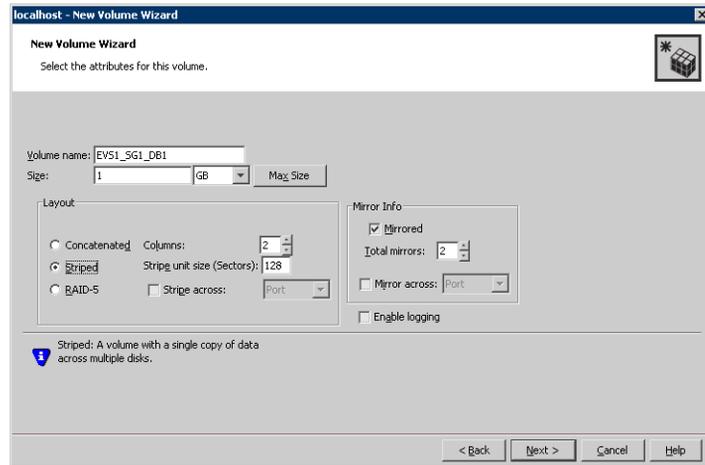


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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.

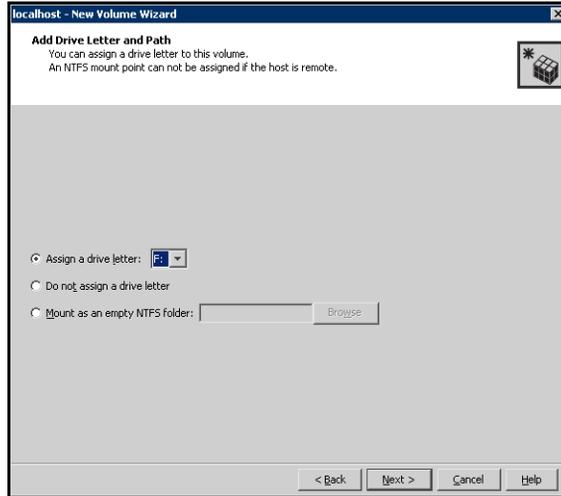
- 8 Click **Next**.

9 Specify the volume attributes.

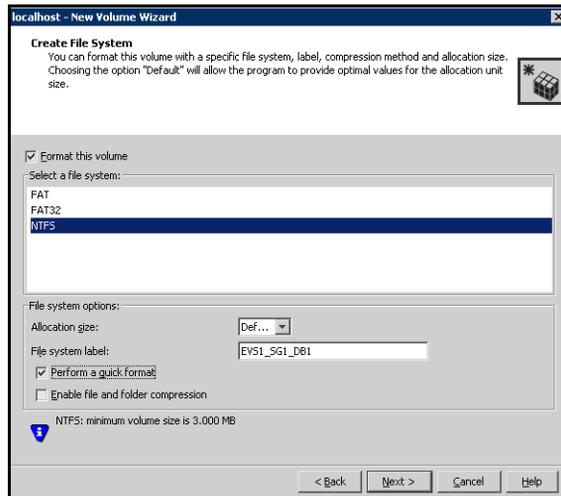


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** 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.
- 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 any additional volumes required. 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 Cluster Configuration Wizard (VCW) 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.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. 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. This account does not require domain admin privileges.
 - Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

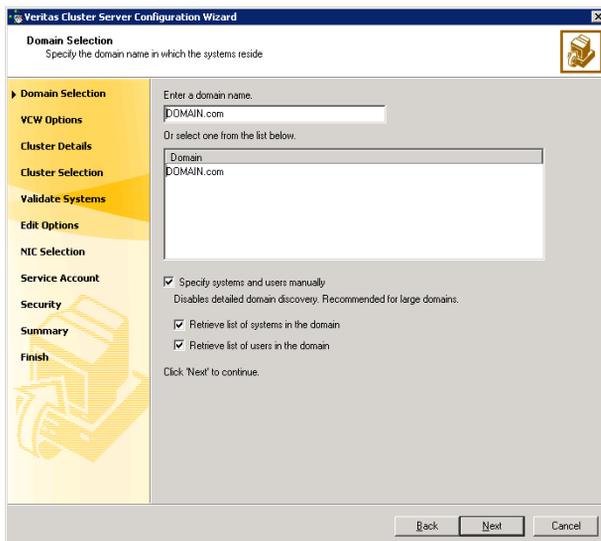
Note: If you are setting up a cluster with multiple Exchange instances (for example, in an any-to-any failover configuration), you may want to add all nodes that are going to host the Exchange instances, to the cluster. If you do that, you do not need to run this wizard again later to add those nodes.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.

Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.



Do one of the following:

- 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 8](#).

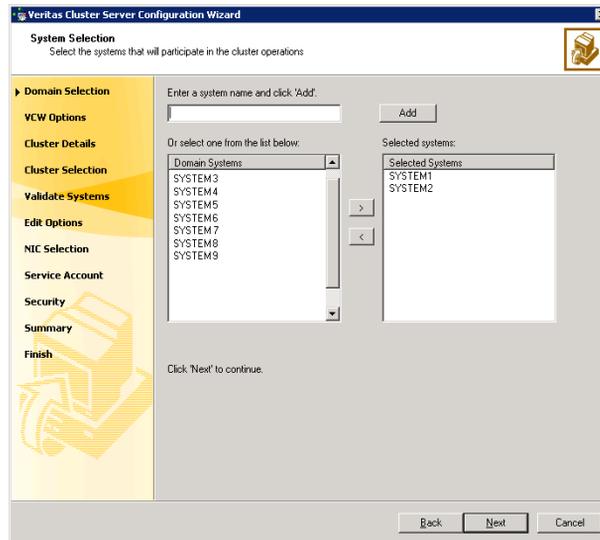
- To specify systems and user names manually (recommended for large domains):

- 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](#). 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 8](#).
- 6 On the System Selection panel, specify the systems for the 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.

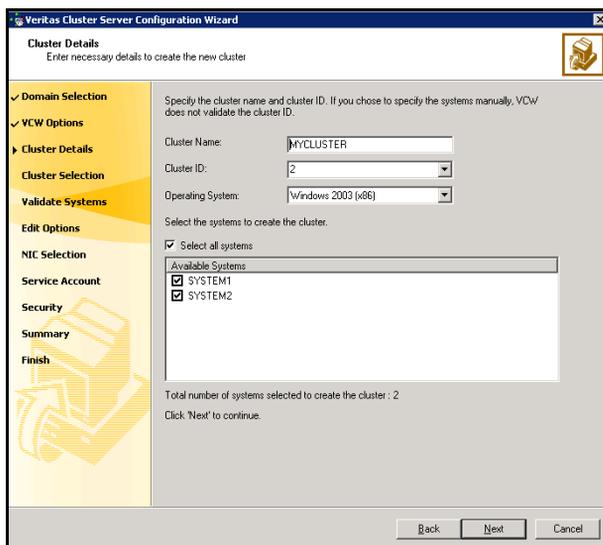
- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 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. The cluster ID can be any number from 0 to 255.

Caution: 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. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (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.

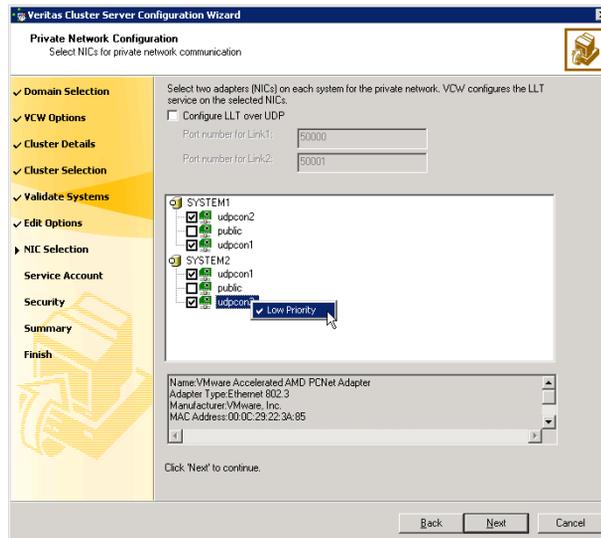
10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

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 the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#).

11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



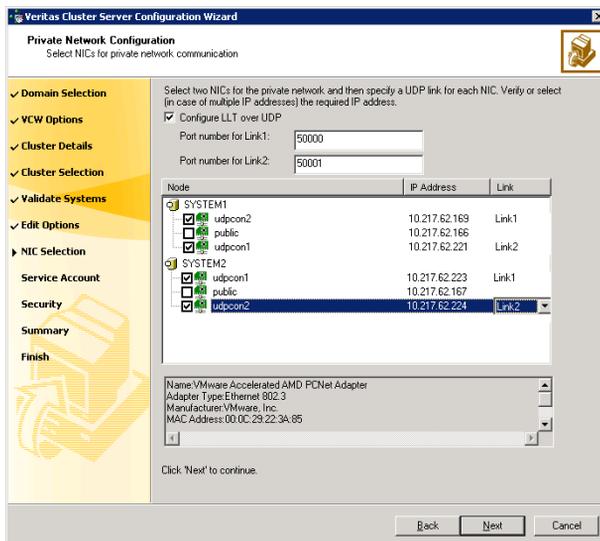
- Select the check boxes next to the two NICs to be assigned to the private network.

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 both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to

65535. The default ports numbers are 50000 and 50001 respectively.

- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

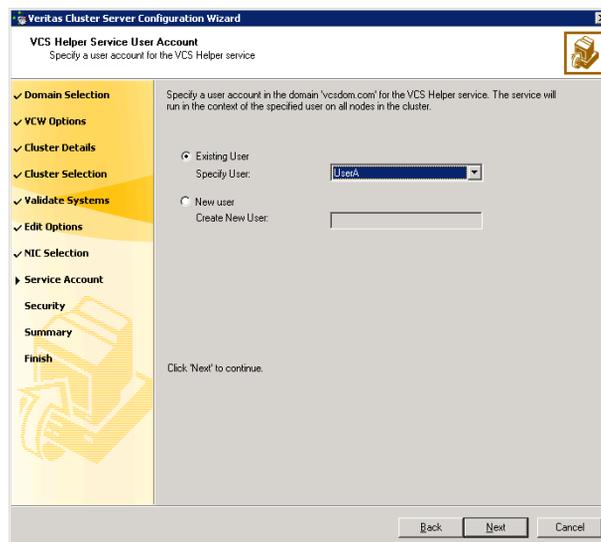
The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

This account does not require Domain Administrator privileges.



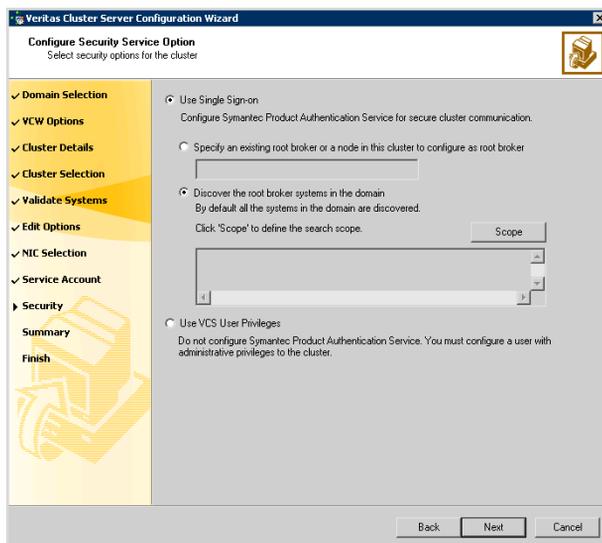
Specify a domain user as follows:

- 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](#), 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**.

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

Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability 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 an existing root broker or a node in this cluster to configure as root broker**, 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 then configures all nodes in the cluster as authentication brokers.

- If you want to search 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 a user *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**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 4-8 contains some more examples of search criteria.

Table 4-8 Search criteria examples

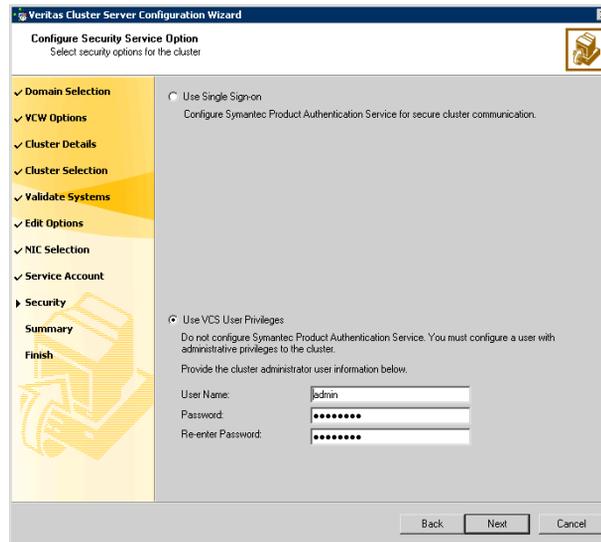
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- 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 a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS`Encrypt` utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password. After the cluster is configured, you can 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**.
- 14 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.

- 15 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**.

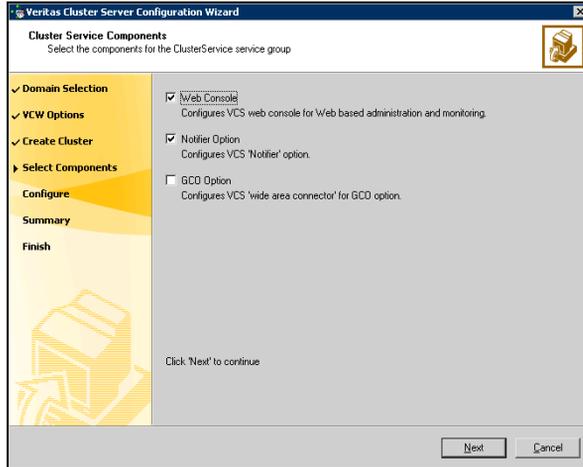
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

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.

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



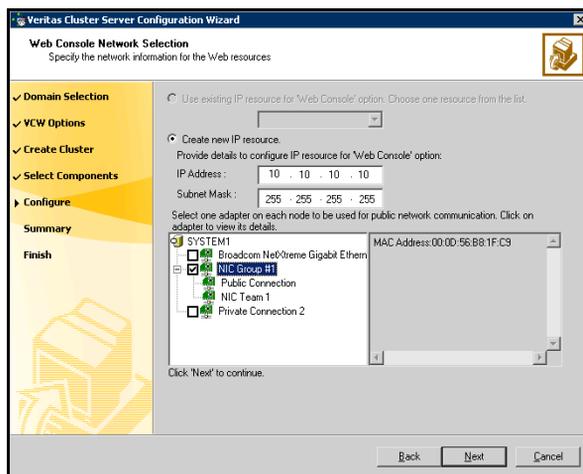
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 88.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 89.

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**.



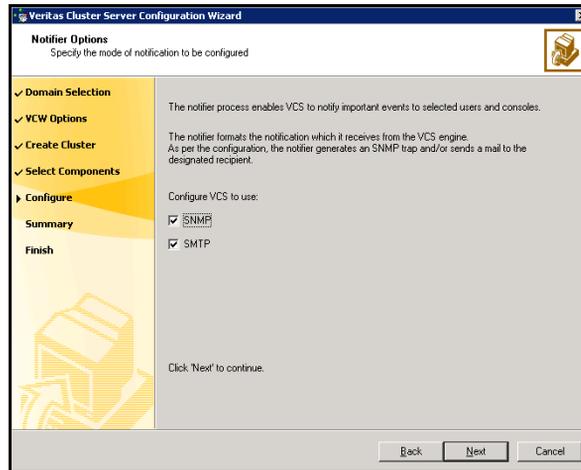
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 89. 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 send emails to designated recipients in response to certain events.

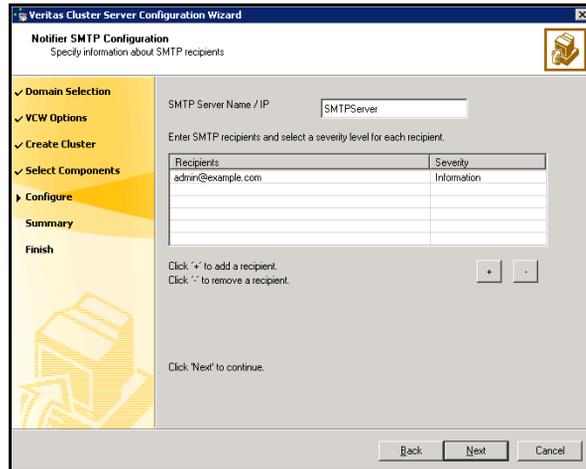
- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

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 items: '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 has 'snmpserv' in the first column and 'Information' in the second. The second row has 'snmpserv1' in the first column and 'SevereError' in the second. Below the table are two buttons: '+' and '-'. Below the buttons is a text input field labeled '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

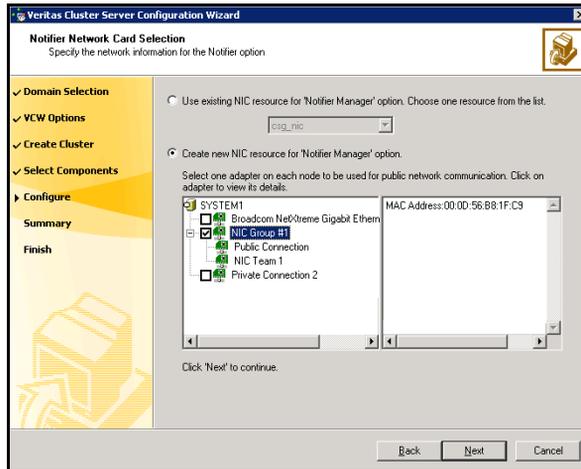
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.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

Ensure that you have completed the following tasks before installing Exchange:

- Prepare the forest and domain.
See [“Preparing the forest and domain”](#) on page 58.
- 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 73.
- Mount the volume containing the information for registry replication (EVS1_REGREP).
See [“Importing a disk group and mounting a shared volume”](#) on page 73.
- 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.
- Verify that the DNS and Active Directory Services are available. Make sure that a reverse lookup zone is created in the DNS.
Refer to Microsoft Exchange documentation for instructions on creating a reverse lookup zone.
- Verify the Dynamic Update option for the DNS server is set to "Secure Only."
- Set the required permissions:
 - You must be a domain user.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups

and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).

- You must be a member of the local Administrators group on all nodes where you install Exchange. You must have write permissions for objects corresponding to these nodes in the Active Directory.
- Either the logged-on user or the VCS Helper service user account must have write permissions on the DNS server to perform DNS updates.
- Make sure the VCS Helper service 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.

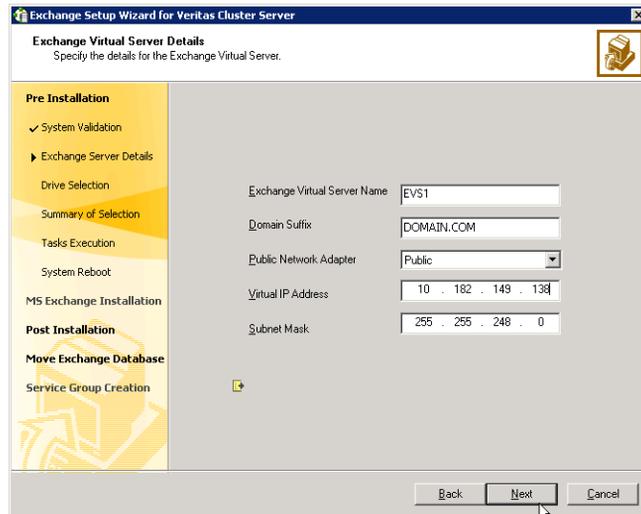
Exchange pre-installation: First node

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 Tools > 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.

Warning: 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 unique on the network.
- 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.

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

If you wish to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: First node

Install Exchange on the node on which you performed the pre-installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```


The state should display as **RUNNING**.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```



```
C:\>net start had
```
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.

- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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.
Once the node is rebooted, move the databases created during the Exchange installation, from the local drive to the shared storage.

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. This is necessary to ensure proper failover operations in the cluster.

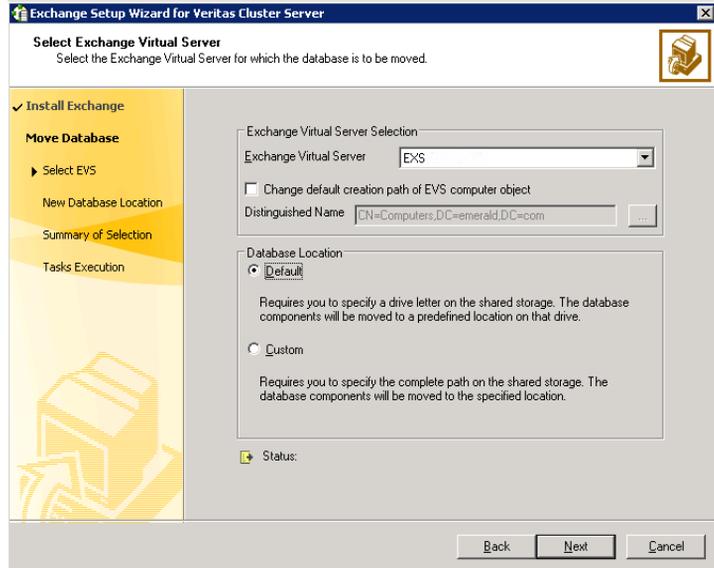
Complete the following tasks before moving the databases:

- 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 73.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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, click **Configure/Remove highly available Exchange Server** and then click **Next**.
- 4 In the Select Option dialog box, click **Move Exchange Databases** and then click **Next**.

- 5 In the Select Exchange Virtual Server dialog box, choose the Exchange virtual server and the database location option and then click **Next**.



Exchange Virtual Server

From the drop-down list, select the Exchange virtual server for which you want to move the databases.

Change default creation path of EVS computer object

Perform the following steps if you wish to change the default path for the Exchange virtual server object in Windows Active Directory:

- Check the **Change default creation path of EVS computer object** check box.
- Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. The Lanman agent performs Windows AD updates. These settings are applicable to the Lanman resource in the service group. By default, the Lanman resource adds the virtual server to the default container "Computers."

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

Default

Click **Default** if you wish to move the databases to predefined location on the shared storage. In the next step the wizard prompts you to specify the drive letter on the shared storage. The first mailbox store, public store, and MTA data are then moved to the generated default paths on the volumes that you specify.

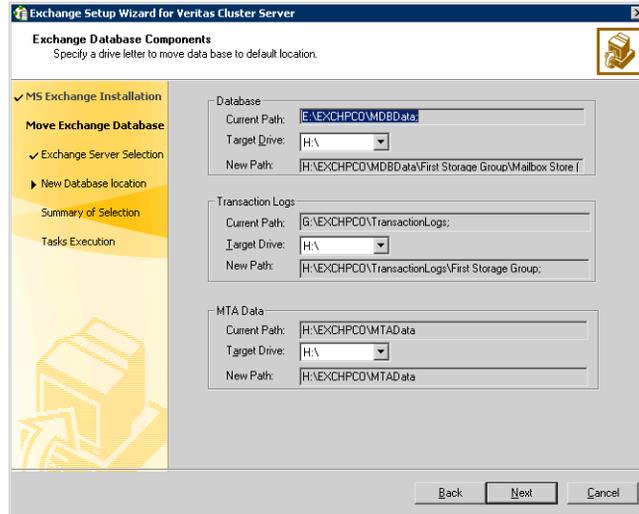
Caution: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

Custom

Click **Custom** if you wish to move the databases to a specific location on the shared storage. Choosing a custom location allows you to specify the Exchange database and streaming path. In the next step the wizard prompts you to specify the entire path of the location on the shared storage. The wizard then moves the databases to the specified directory.

If you chose the Default option, proceed to the next step. If you chose the Custom option, proceed to [step 7](#).

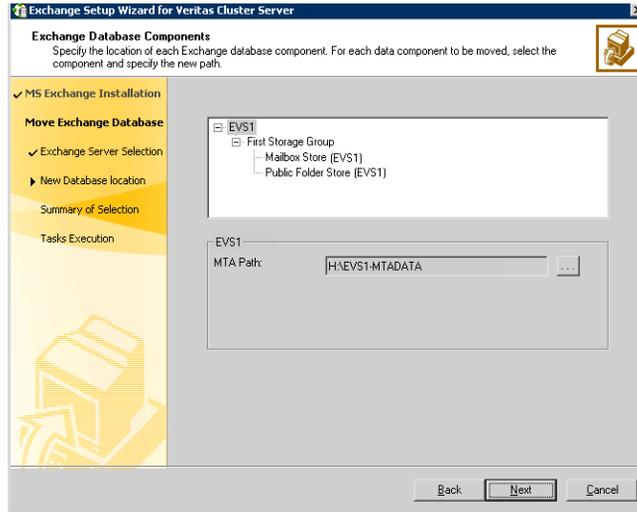
- 6 For the option of a default database location, specify the drives for moving the Exchange database components. The database components are then moved to a predefined location on that drive.



On the Exchange Database Components panel, complete the following steps:

- Specify a drive for moving the Exchange database.
- Specify a drive for moving the Exchange Transaction Logs.
- Specify a drive for moving the Exchange MTA Data.
- Click **Next** and proceed to [step 8](#).

- 7 For the option of a custom database location, specify the location for specific Microsoft Exchange data components and then click **Next**.



For each data component that you wish to move, select the component and then click the ellipsis (...) to browse for the folder where you want to move it.

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

- 8 Review the summary of your selections and then click **Next**.
The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task.
- 9 After all the tasks are completed successfully, click **Next**.
- 10 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 93.

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.

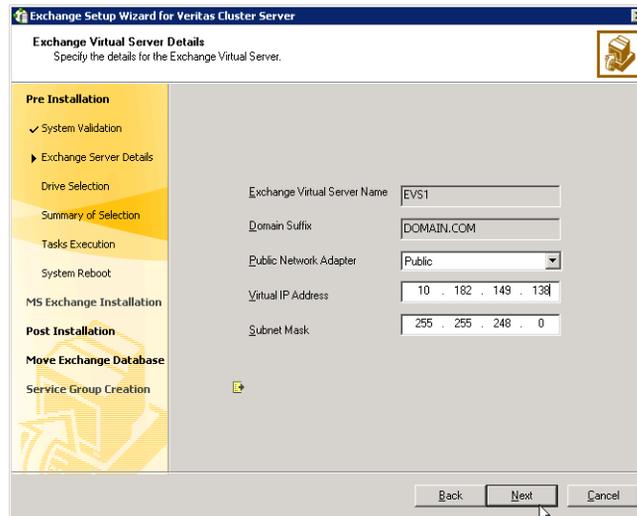
See .

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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: Additional nodes

- Install Exchange on the additional node on which you performed the 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.

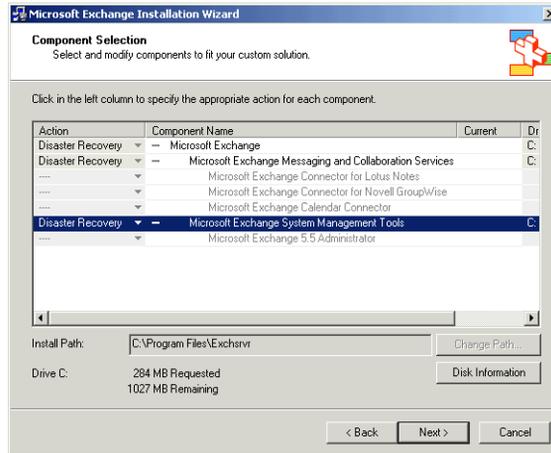
To install Exchange

- 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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:
`C:\>hasys -state`
The state should display as RUNNING.
If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`

- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

To add the nodes later, use the Exchange service group configuration wizard. See “[Configuring the Exchange service group for VCS](#)” on page 108.

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. Refer to the Exchange documentation for instructions.

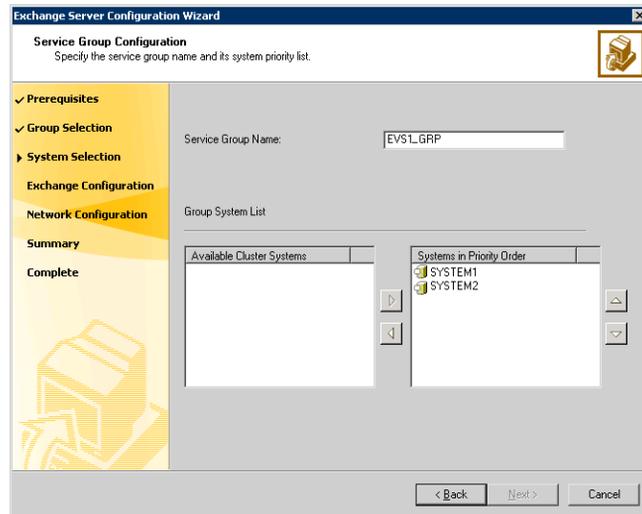
Prerequisites

- You must be a VCS Cluster Administrator. This privilege is required to configure service groups.
- 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 73 for instructions on mounting and “[Unmounting a volume and deporting a disk group](#)” on page 73 for instructions on unmounting.
- 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* for information on how to add additional resources to an already configured service group.

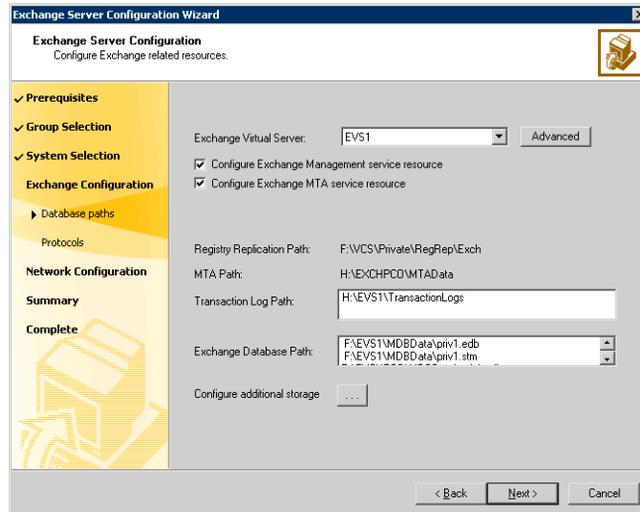
To configure the Exchange service group

- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**:
 The wizard starts validating your configuration. Various messages indicate the validation status.



- Enter a name for the Exchange service group.
 If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.
- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
- Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.
On the Lanman Advanced Configuration dialog box, complete the following:
 - In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
 - Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.

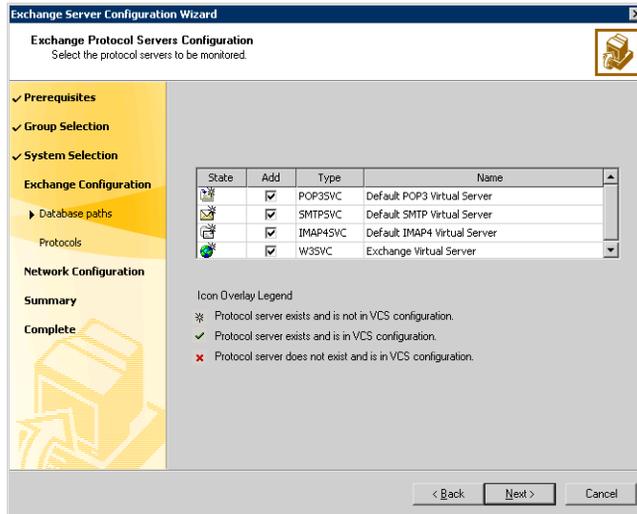
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.

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.

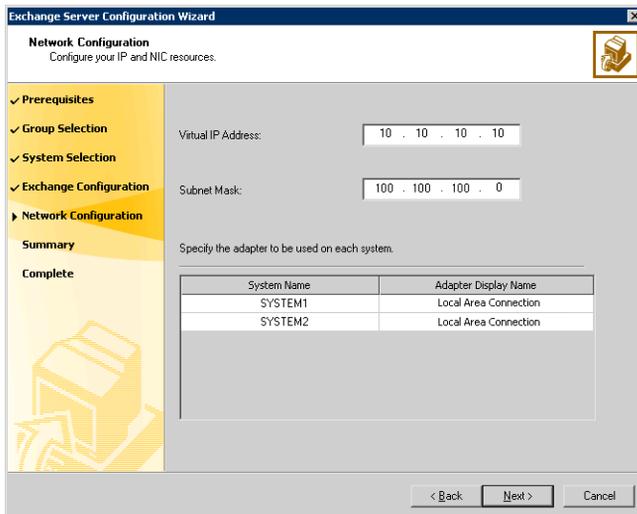
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:

 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- 6 On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.

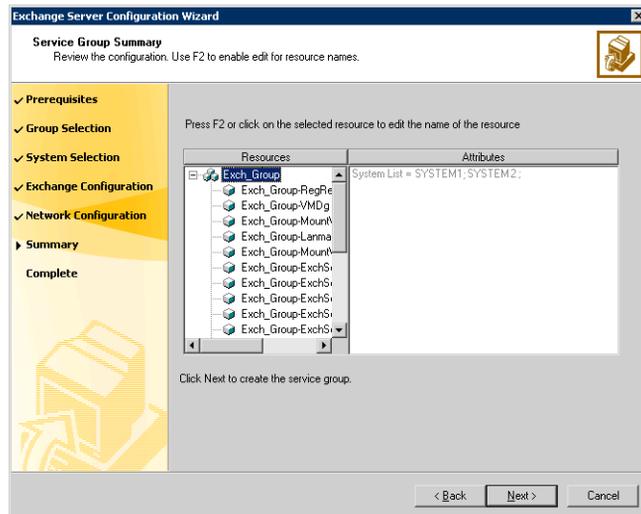


- 7 On the Network Configuration panel, specify information related to the network and then click **Next**:



- 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 Display Name** field to view the adapters associated with a node.
 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.

8 Review the service group configuration, change the resource names, if desired, and then click **Next**:



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 the **Esc** key.

9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status

of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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 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.
- After you install and configure the cluster connector, configure the CMC group on all the nodes in the cluster, and the state of the CMC group should ONLINE on one of the cluster nodes.

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 Start the Setup program to install the Cluster Connector for Windows.
- 2 In the Symantec Product Installer window, select **VCSMC Cluster Connector for Windows** to install the cluster connector.
- 3 In the Welcome dialog box, make sure all the prerequisites for installing the VCS MC Cluster Connector 5.1 for Windows are satisfied. Click **Next**.
- 4 In the VCS MC Cluster Connector 5.1 for Windows dialog box:
 - Select the domain name and the nodes on which the cluster connector will be installed. Click **Add**.
 - To change the install path, click **Change**.
 - Click **Next**.

- 5 The installer validates the selected nodes in the Validation dialog box. The installation proceeds only if all the nodes are accepted. Click **Next**.
- 6 The installer displays a summary of install options prior to the actual installation. Click **Next**.
- 7 The installation starts on all nodes simultaneously.
- 8 The installer displays the installation report after the installation is completed on all the nodes. Click **Next**.

Click **View Log Files** to see the log files of the installation process. You can check the ClusterConnector-0.log at the following path: C:\Program

Files\Symantec\VRTScmccc\log

Check the ClusterConnectorConfig-0.log in the same directory for the cluster connector configuration process.

Configuring the cluster connector

Perform the following steps to configure the cluster connector.

To configure the cluster connector

- 1 Install the management server and configure it. Refer to the *Veritas Cluster Management Console Implementation Guide*.
- 2 Install the cluster connector on a VCS cluster.
- 3 Run the cluster connector configuration utility, found in X:\Program Files\Symantec\VRTScmccc\bin\cc_configure.bat (where X is the driver letter on which the cluster connector is installed).
- 4 Enter the network IP address of the management server or the hostname.
- 5 Enter the certificate to add to the trusted keystore or enter 'q' to quit.
- 6 Enter an administrator user name: **root**
- 7 Enter the domain name. For example **vcs01.symantecexample.com**
- 8 Enter the domain type:
 - 1: Windows
 - 2: nis
 - 3: nisplus
 - 4: unixpwd
 - 5: ldap
 - 6: localhost
 - 0: QuitEnter the domain type [1]: 4

- 9 Enter the password.
- 10 Enter a unique identifier for the cluster:
Enter a unique identifier for the cluster: [43896e6c-0220-4832-9556-97082515c77b]/accept default:
This indicates the configuration is successful.
- 11 To verify that the CMC group and its resources are fully-functional i.e. they are online, can fail over, etc., check for the existence of the cluster on the management server.

Configuring the cluster connector using the management server console

This task enables you to configure an upgraded version of the cluster connector. Before you perform this task, you must first install an upgraded version of the cluster connector on the target clusters. This task configures only versions of the cluster connector that have already been installed on the target clusters.

To upgrade the cluster connector on discovered clusters

- 1 On the main tab bar, click **Administration**.
- 2 On the details tab bar, click **Configured Clusters**.
- 3 In the Configured Clusters table, do one of the following:
 - To select one or more clusters, check the check box next to each required cluster.
 - To select all clusters, check the check box at the top of the table.
- 4 On the Configuration task menu, select **Upgrade Cluster Connector**.
- 5 In the Upgrade Cluster Connector wizard, read the overview information and then click **Next**.
- 6 This launches the **Upgrade Cluster Connector** wizard to configure known (secure or non-secure clusters). Click **Next**.
- 7 In the Access Credentials for Target Clusters panel, specify the following options:
 - The type of security access that the cluster uses. The options are:
 - Classic VCS
This option enables only VCS users that are configured locally on this cluster to log in to the cluster.
 - VxAT
Otherwise known as Symantec Product Authentication Service, VxAT is the Symantec cross-product user authentication service. If you select VxAT, you must also specify the IP address of the Symantec Product authentication broker that you want to use.

- The cluster administrator user name, password, domain, and domain type required to establish a connection to the cluster. You must be a cluster-level administrator on each cluster that you want to add or discover. The **Domain** field requires a fully qualified domain name.
- 8 To configure clusters in the secure mode in the Discover Clusters dialog box:
 - Select **VxAT**.
 - Enter the access credentials (user name and password) of the target clusters.
 - Click **Next**.
 - 9 To configure clusters in the non-secure mode in the Discover Clusters dialog box:
 - Select **Classic VCS**.
 - Enter the access credentials (user name and password) of the target clusters.
 - Click **Next**.

If you have specified both VxAT security clusters and Classic VCS security clusters, this panel runs separately for each. The wizard enables you to select either the cluster's authentication broker or one of the predefined authentication brokers.
 - 10 In the Summary of Target Clusters panel, read the overview of your selections and then click **Finish**.

Deploying SFW HA for high availability: Standalone Exchange servers

This chapter contains the following topics:

- [“Reviewing the requirements”](#) on page 123
- [“Configuring the network and storage”](#) on page 134
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 136
- [“Configuring disk groups and volumes”](#) on page 142
- [“Managing disk groups and volumes”](#) on page 151
- [“Converting the standalone Exchange server into a “clustered” Exchange server”](#) on page 152
- [“Adding the standalone Exchange server to a cluster”](#) on page 155
- [“Moving Exchange databases to shared storage”](#) on page 179
- [“Installing Exchange on additional nodes”](#) on page 184
- [“Configuring the Exchange service group for VCS”](#) on page 190
- [“Verifying the cluster configuration”](#) on page 197

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.

Table 5-1 on page 122 outlines the high-level objectives and the tasks to complete each objective:

Table 5-1 Task list: Standalone Exchange HA configuration

Objective	Tasks
“Reviewing the requirements” on page 123	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 129	Understanding a typical active-passive Exchange configuration in a two-node cluster
“Configuring the network and storage” on page 134	<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 136	<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 142	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create volumes for data, log, RegRep, and MTA
“Managing disk groups and volumes” on page 151	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 152	Converting the standalone Exchange server into a cluster node using the Exchange Setup Wizard for Veritas Cluster Server

Table 5-1 Task list: Standalone Exchange HA configuration (Continued)

Objective	Tasks
“Adding the standalone Exchange server to a cluster” on page 155	<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 156 ■ For an existing cluster, adding the new nodes to the cluster, “Adding nodes to an existing cluster” on page 173
“Moving Exchange databases to shared storage” on page 179	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 184	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 190	Creating the Exchange service group using the VCS Exchange Configuration Wizard
“Verifying the cluster configuration” on page 197	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 Symantec recommended requirements may vary.

This HA solution is based on a standalone Microsoft Exchange server. See the requirements for the supported Exchange servers.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space.

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

Table 5-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

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

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/302144>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agent for Microsoft Exchange.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2003 versions

The following table lists the Microsoft Exchange Server 2003 versions supported with SFW HA 5.1 Service Pack 1.

Table 5-4 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Memory must be a minimum 256 MB of RAM per server for Exchange 2003; refer to the Microsoft documentation for more information.
- 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. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
See “[Best practices for SFW HA](#)” on page 128.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each Exchange Virtual Server (EVS).

- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone 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 Dynamic DNS (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 for SFW HA

The following permissions are required:

- 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 on 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 Exchange virtual server computer object 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, 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 for SFW HA

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 Server 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 2003.
- 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.

Best practices for SFW HA

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.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

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.

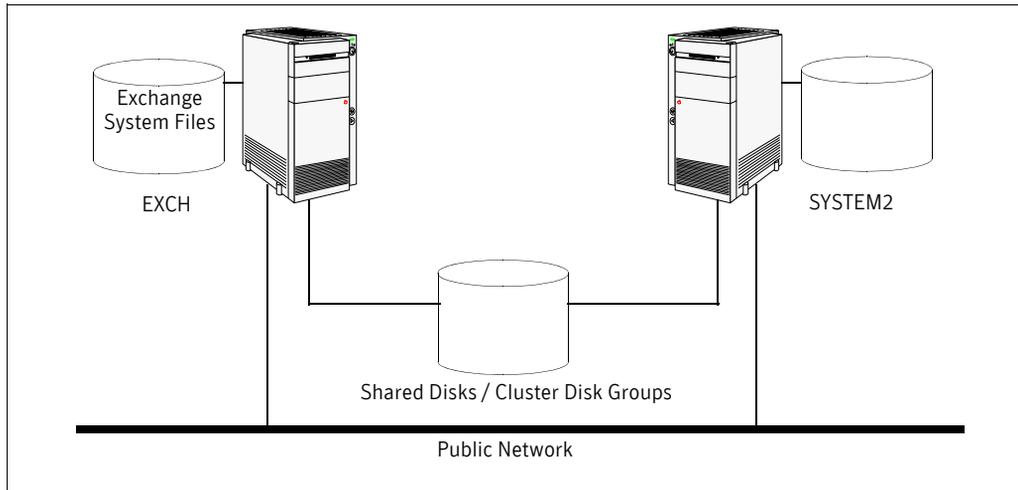
In Scenario II, you start with a standalone Exchange server and a cluster which may be running other applications.

Scenario I

In Scenario I, start with two nodes:

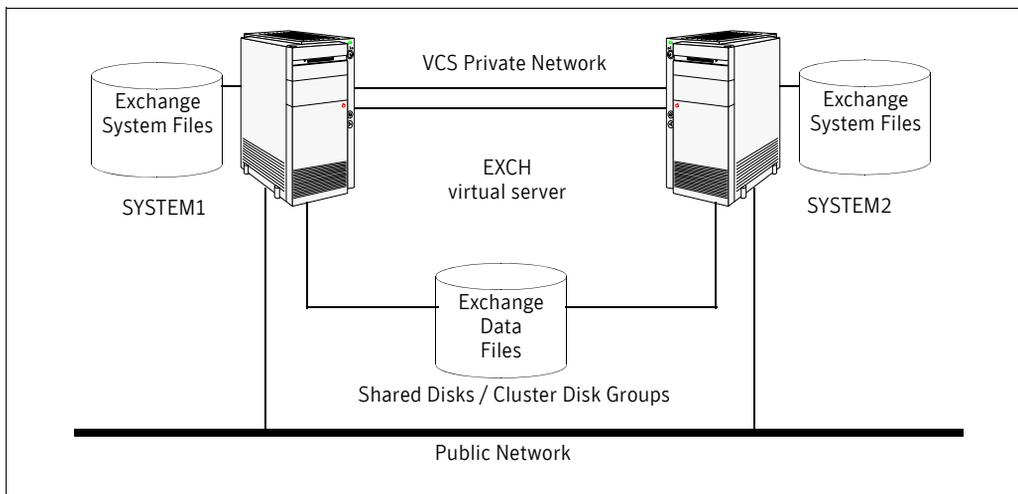
- EXCH which is a standalone Exchange server
- SYSTEM2, a new node which will join the standalone Exchange server to form a cluster

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.

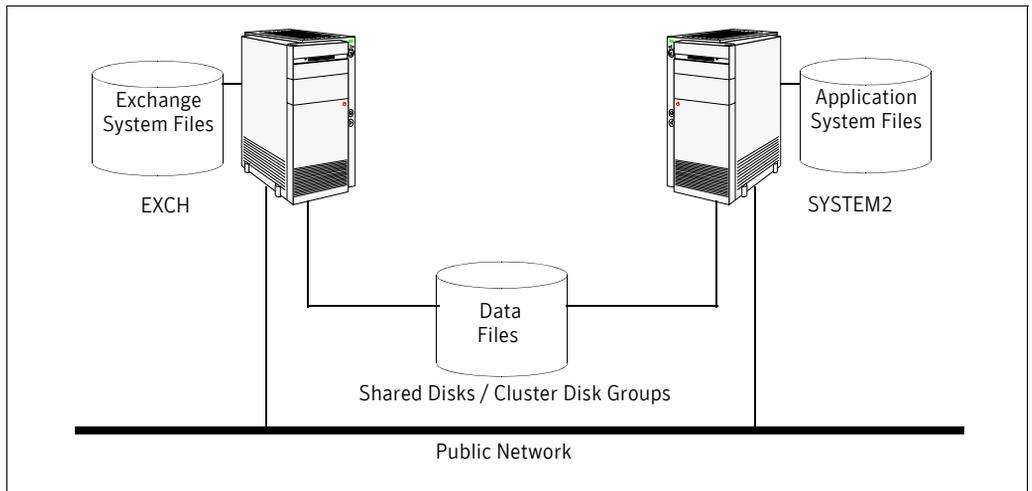
During the conversion of a standalone Exchange server into a clustered server, the existing node name of the standalone Exchange server becomes the name of the Exchange virtual server. For example, if the name of your Exchange server is EXCH, EXCH becomes the name of the Exchange virtual server.

Scenario II

In scenario II, start with a cluster:

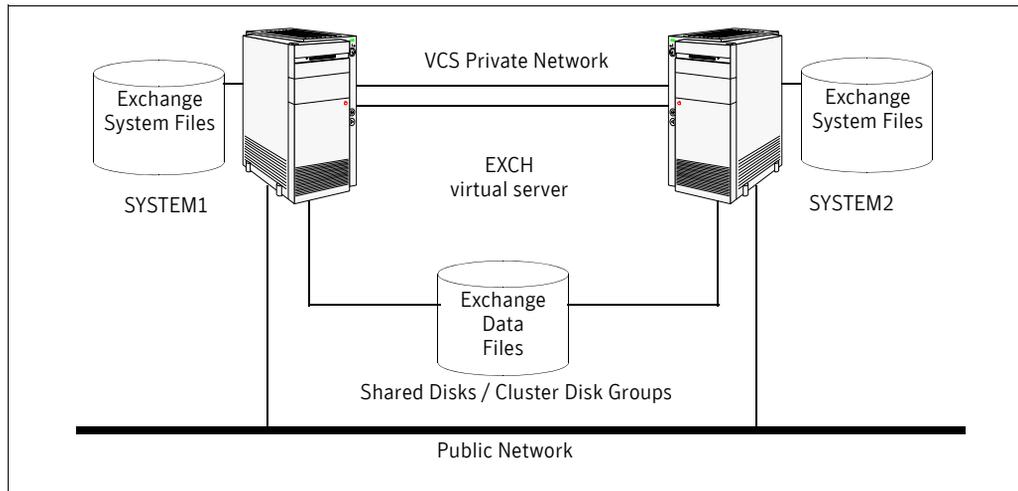
- EXCH which is a standalone Exchange server
- SYSTEM2, a node which not running as an Exchange server, but is part of a cluster

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 5-6 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

Table 5-6 Sample configuration

Name	Object
EVS1_REGREP	Volume that contains 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. 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 and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network 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 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, 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 123.

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 129, 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 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, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

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

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 5-7 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 138.

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 Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The Select Product screen appears.
- 3 Review the links on the Select Product screen.
Links on this screen access Late Breaking News, the Configuration Checker, as well as begin the process to install Storage Foundation HA for Windows. Click on **Read Late Breaking News** for the latest information about updates, patches, and software issues regarding this release.
- 4 Click **Storage Foundation HA 5.1 SP1 for Windows**.
- 5 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 6 Review the Welcome message and the listed prerequisites. Ensure that any of the listed prerequisites applicable to your installation are met prior to proceeding.

Click **Next**.

- 7 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 AGREE TO the terms of the license agreement**, and then click **Next**.
- 8 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.
- 9 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.
- 10 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. Be sure to select the following as appropriate for your installation.

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, select the option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

- 11 Select the following 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.

When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

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

- 12 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.
- 13 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.

14 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages.

If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning:

The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If applicable to your installation, perform the above procedure.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

If applicable to your installation, perform the above procedure.

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation 5.1 for Windows (Server Components) with the MSCS option, you may want MSCS Quorum Arbitration Time (Min. and Max) to be adjusted to ensure optimal functionality with Veritas dynamic volumes with MSCS.

For additional information, see the *Storage Foundation for Windows Administrator Guide* for details.

If applicable to your installation, perform the above procedure.

15 When finished reviewing the message or messages, click **OK**.

16 The Summary screen appears displaying an Install report. Review the information in the Install report. Click **Back** to make changes, if necessary. Click **Install** if information is validated.

17 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.

18 When the installation completes, review the summary screen and click **Next**.

- 19 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.
- 20 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 21 Review the log files and click **Finish**.
- 22 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 Veritas Enterprise Administrator (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.

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

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

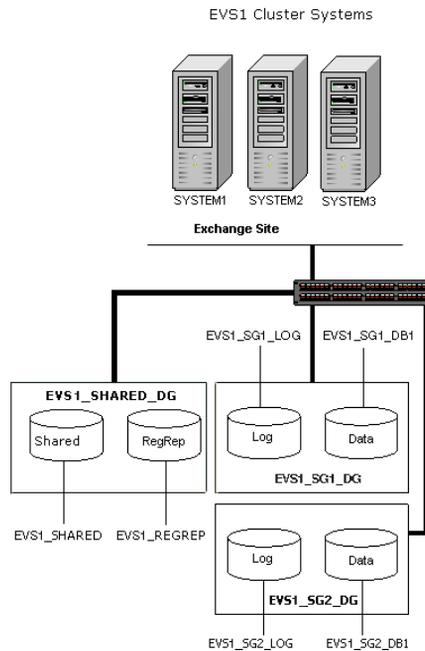
- The type of volume configurations that are required.
- The number of volumes or LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.

Typically, a SFW disk group corresponds to an Exchange storage group. You may be configuring new shared storage for the high availability environment, or the existing standalone Exchange stores and logs may already be on shared storage. If the existing stores and logs are already on shared storage, read the following topic:

[“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 144.

Figure 5-5 shows 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 contains the following 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 contains the following 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.

Note: If you are planning a disaster recovery configuration using Veritas Volume Replicator (VVR), you will need to allow additional disk space for a Storage Replicator Log volume (EVS1_REPLOG). This volume is created automatically when you run the wizard to set up replication. For more about VVR planning, see the *Veritas Volume Replicator, Administrator's Guide*.

Considerations for converting existing shared storage to cluster disk groups and volumes

The stores and logs for your existing standalone Exchange server may already be on shared storage. In this case, when you create cluster disk groups, you specify the disks that contain the existing stores and logs.

Creating a disk group converts the disks from basic disks to dynamic disks. Partitions on the disks are automatically converted to volumes on the dynamic disks.

Therefore, if your existing disk layout contains stores and logs in the same partition, they become part of the same volume in the cluster disk group. If the disk contains multiple partitions, each containing a storage group, each partition becomes a separate volume, but all will become part of the same cluster disk group. If this configuration does not meet your requirements, you may want to modify your disk layout before creating the cluster disk group.

For additional information on converting basic to dynamic disks, see *Veritas Storage Foundation Administrator's Guide*.

Symantec recommends creating a separate 100 MB RegRep volume that contains the list of registry keys that must be replicated among cluster systems for Exchange. However, if no additional disks are available on the shared

storage, you can specify an existing volume as the registry replication path during service group creation.

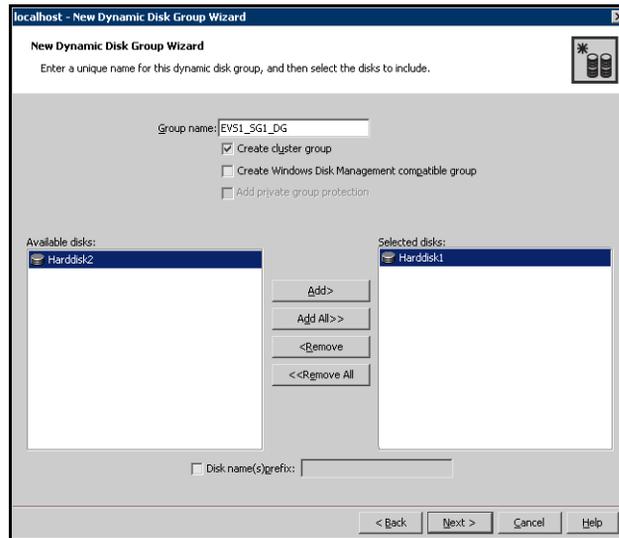
Creating a disk group

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

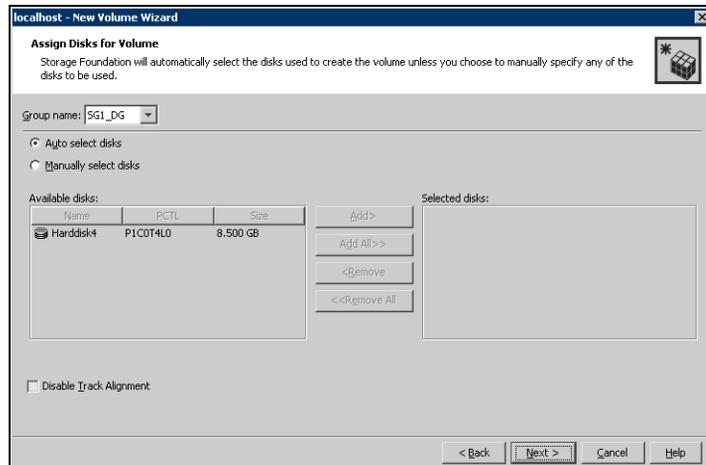
- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

Creating volumes

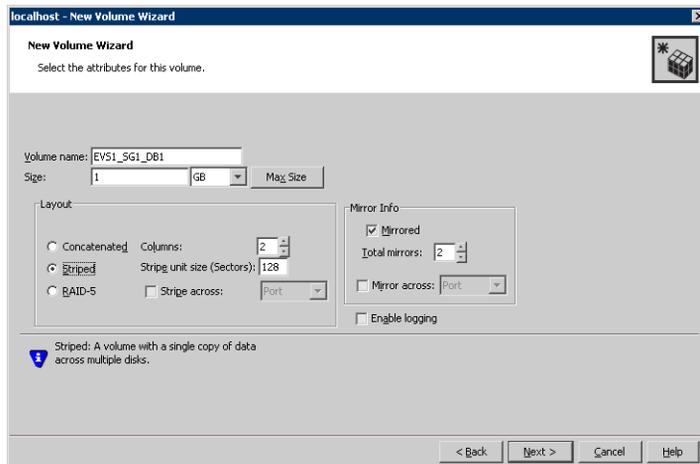
This procedure assumes you are starting with the EXCH_SG1_DB1 volume.

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.



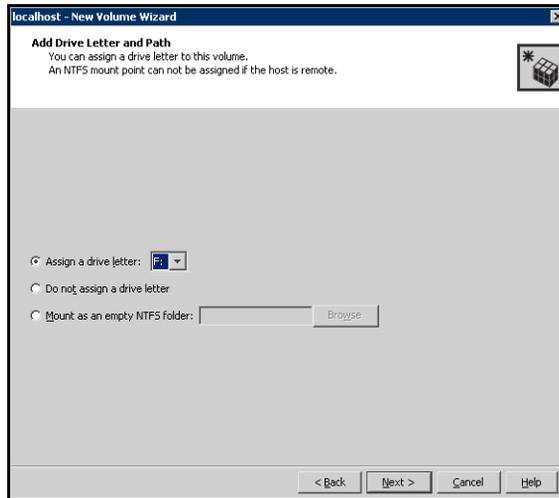
- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable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

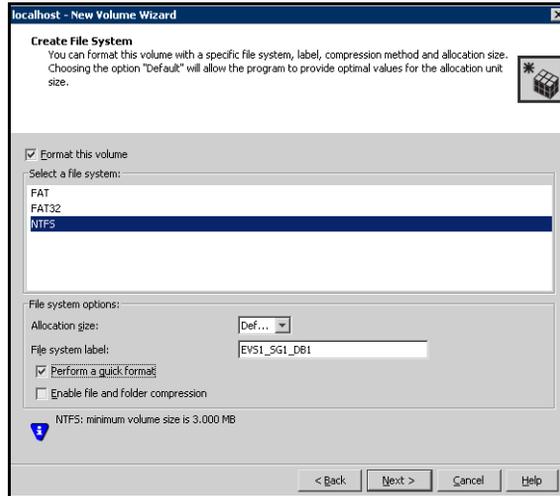
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

- 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 any additional volumes required. 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 for VCS to convert a standalone Exchange Server into a “clustered” Exchange server.

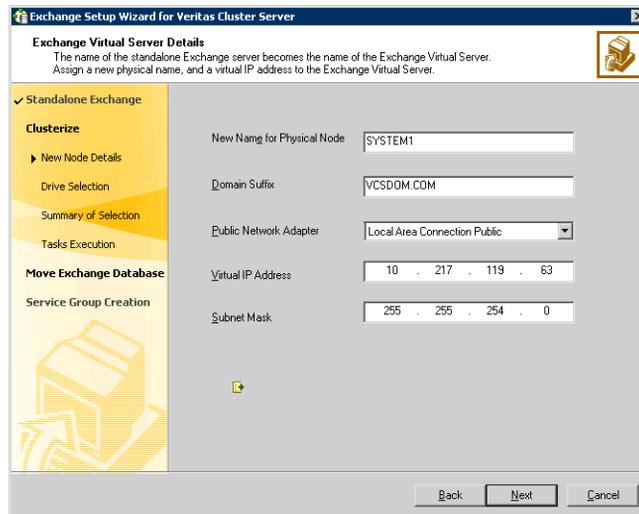
In this wizard, the node name of the standalone Exchange Server becomes the name of the Exchange virtual server and the existing node is given a new physical node name.

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

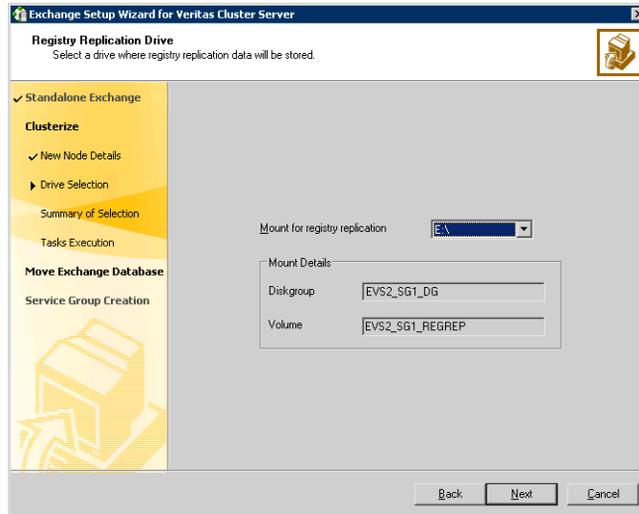
- 1 Start the Exchange Setup Wizard for VCS from the node having the standalone Exchange server installed.
 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Setup Wizard**.
- 2 Review the information in the Welcome dialog box and click **Next**.
- 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 155 and then use [“Creating a new cluster and adding nodes”](#) on page 156 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 172 and then continue with the procedure [“Adding nodes to an existing cluster”](#) on page 173 to add any new nodes, including the standalone Exchange server, to the cluster.

Prerequisites for a new cluster

The VCS Cluster Configuration Wizard (VCW) 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 service 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.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. 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. This account does not require domain admin privileges.

- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

If no cluster exists, continue with “[Creating a new cluster and adding nodes](#)” on page 156 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 172 and then continue with the procedure “[Adding nodes to an existing cluster](#)” on page 173 to add any new nodes, including the standalone Exchange server, to the cluster.

Creating a new cluster and adding nodes

After installing SFW HA, run VCW to configure and set up the components required to run a cluster.

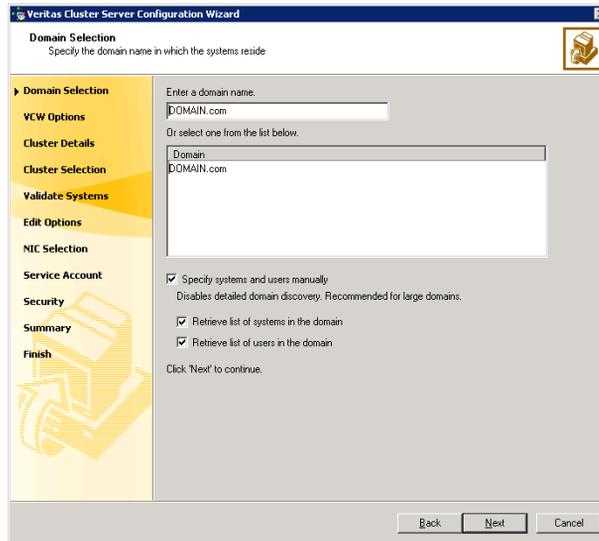
In the examples below, the system names are SYSTEM1 and SYSTEM2.

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.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.



Do one of the following:

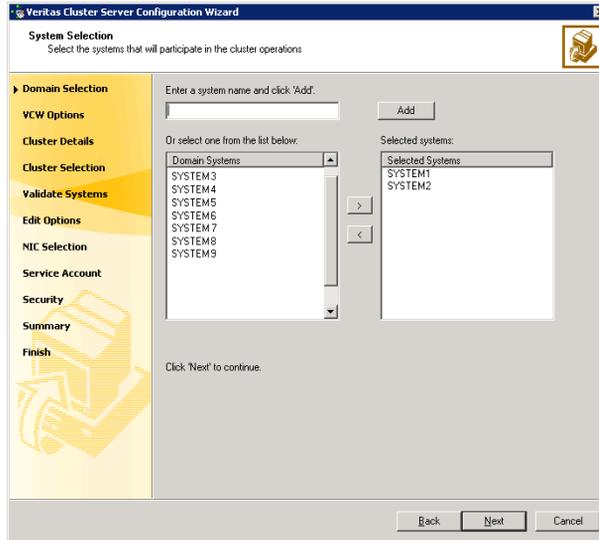
- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 158.
- 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 158. 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 8](#) on page 158.

- 6 On the System Selection panel, specify the systems for the 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 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

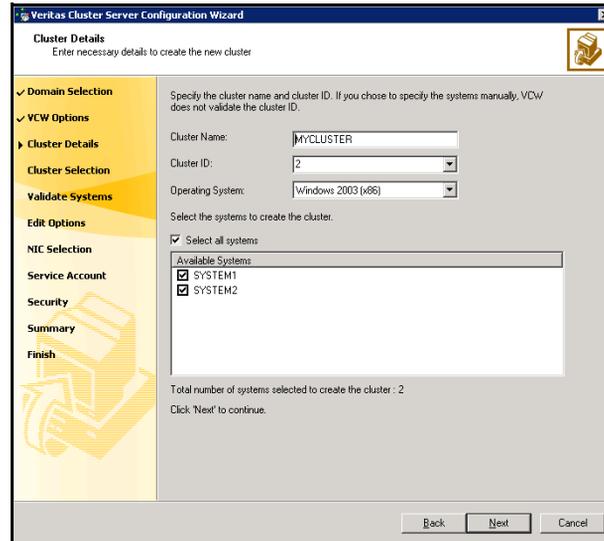
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.

- 9 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. The cluster ID can be any number from 0 to 255.

Caution: 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. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (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.

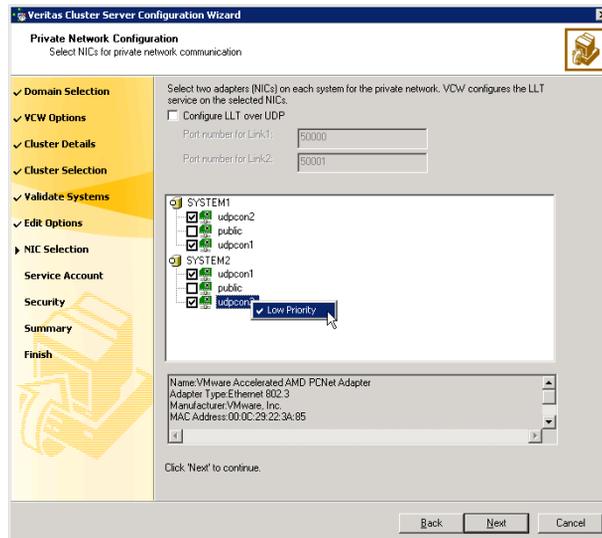
10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

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 the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 162.

11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



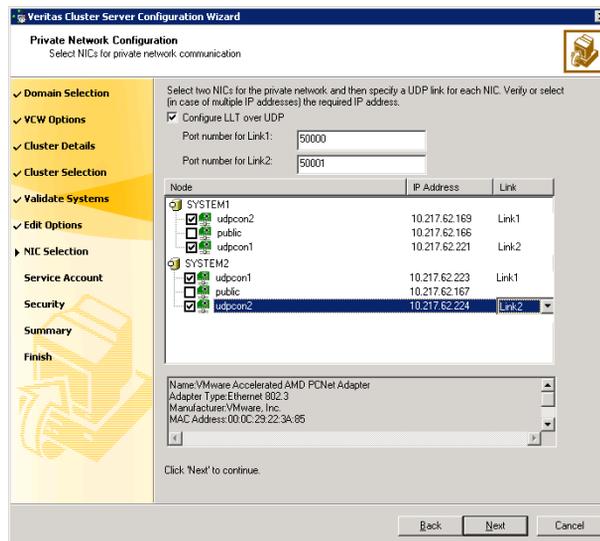
- Select the check boxes next to the two NICs to be assigned to the private network.

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 both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



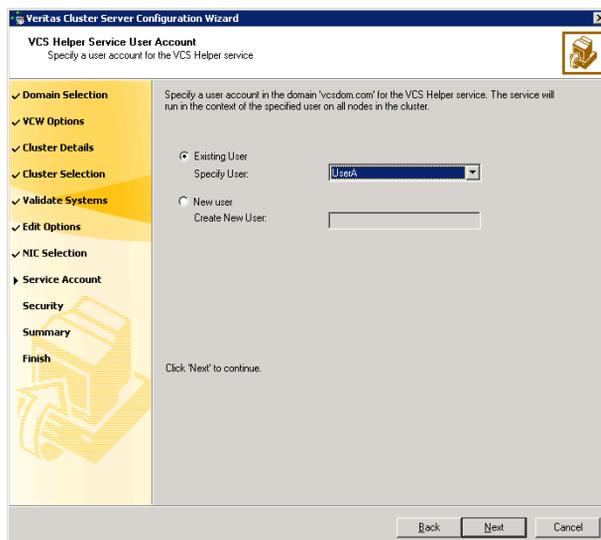
- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to

65535. The default ports numbers are 50000 and 50001 respectively.

- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.



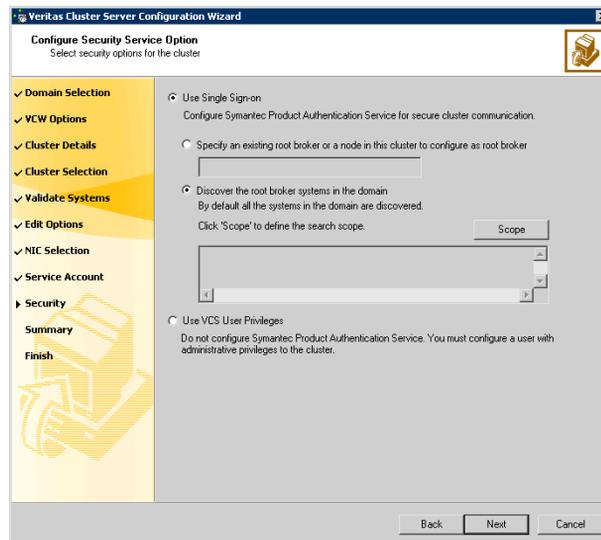
Specify a domain user as follows:

- 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 157, 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**.

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

Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability 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 an existing root broker or a node in this cluster to configure as root broker**, 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 then configures all nodes in the cluster as authentication brokers.

- If you want to search 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 a user *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**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 5-8 contains some more examples of search criteria.

Table 5-8 Search criteria examples

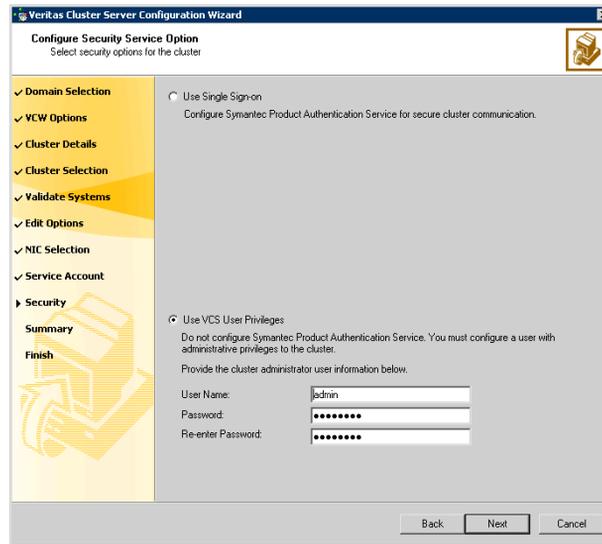
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- 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 a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS Encrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password. After the cluster is configured, you can 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**.

14 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.

- 15 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**.

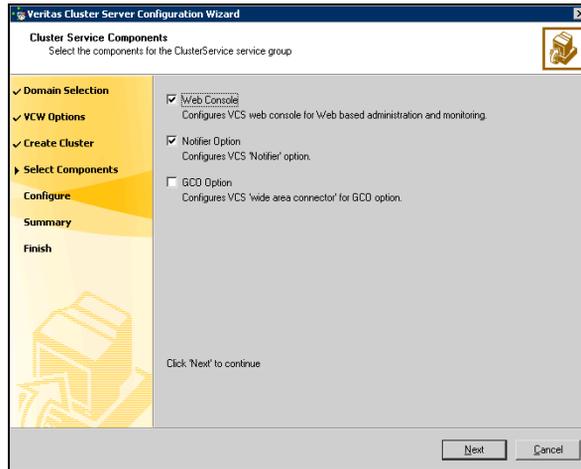
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

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.

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



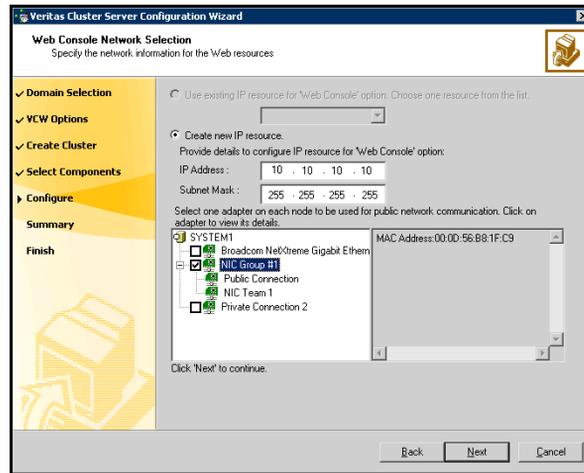
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 168.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 169.

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**.



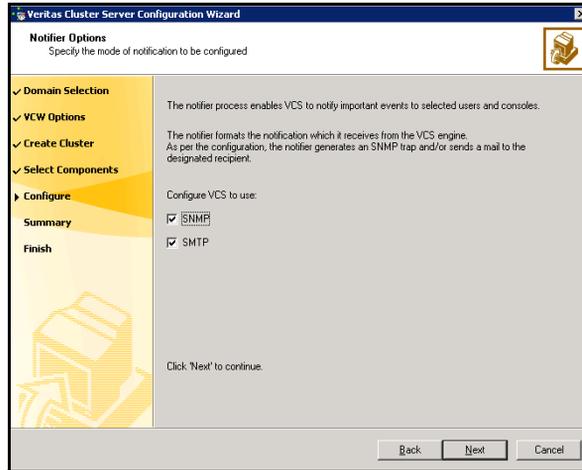
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 169. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

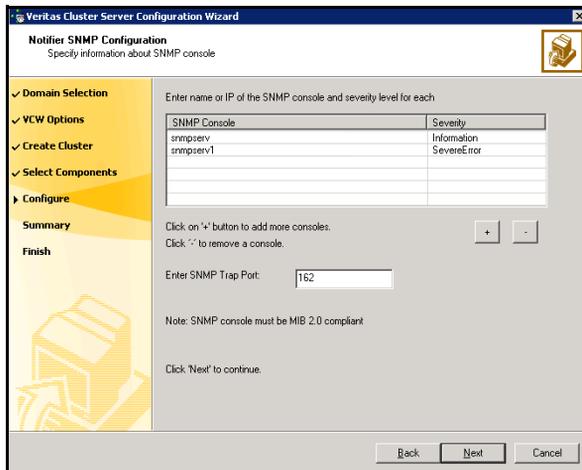
To configure notification

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

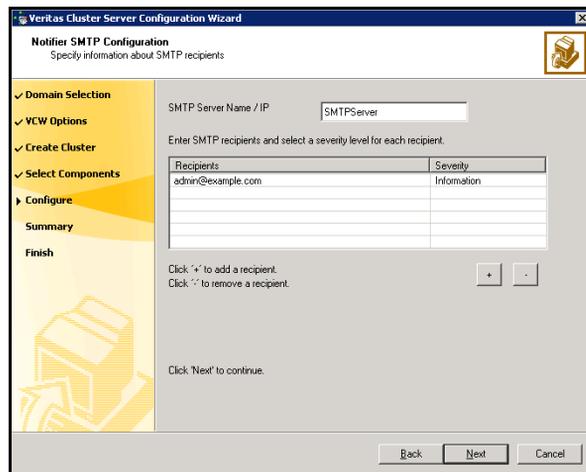


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

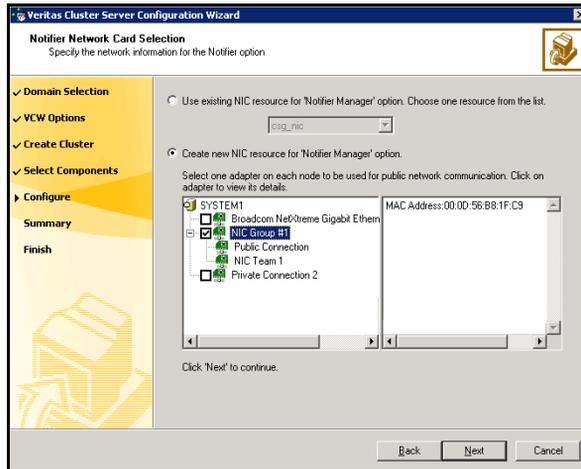


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.
 - 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 179.

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.

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 151 for instructions on mounting and “[Unmounting a volume and deporting a disk group](#)” on page 151 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 VCS Exchange agent resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on how 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.

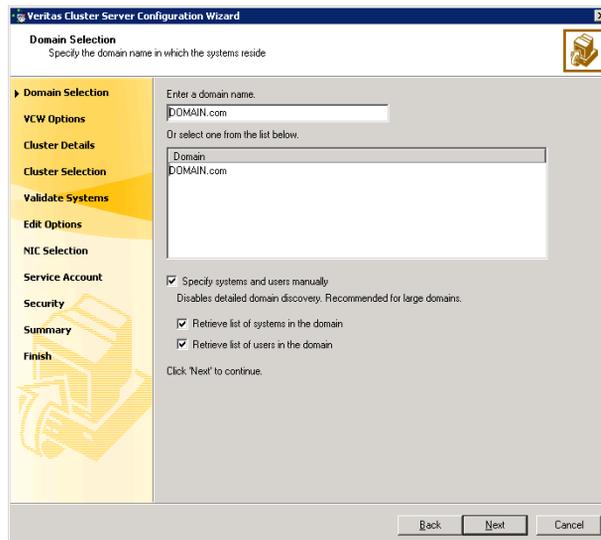
In the examples below the system names are SYSTEM1 and SYSTEM2. Remember that the original name of your existing standalone Exchange server, EXCH in the example, is now the name of the virtual server, and the physical node has a new name, in the example, SYSTEM1.

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.

Note: Run the VCS Cluster Configuration Wizard (VCW) from the standalone node or a node in the cluster.

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.



Do one of the following:

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

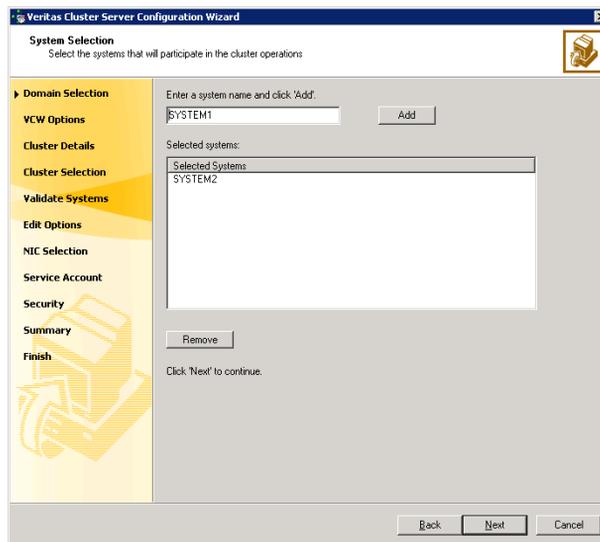
- Clear the **Specify systems and users manually** check box.
- Click **Next**.

Proceed to [step 8](#) on page 177.

- 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 176. Otherwise proceed to the next step.

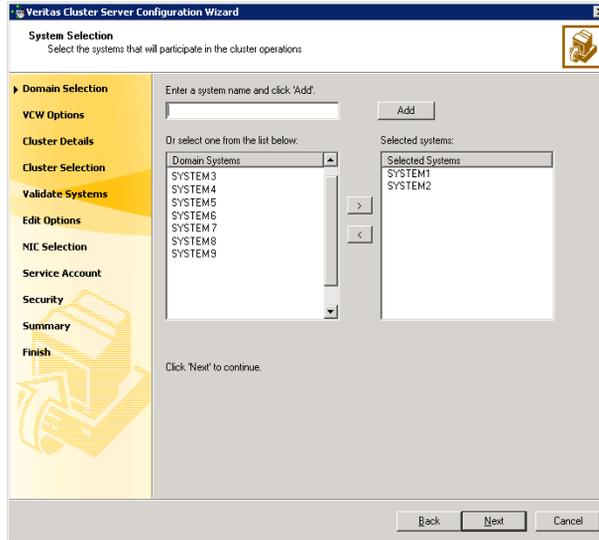
- 5 On the System Selection panel, complete the following and click **Next**.



- 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 8](#) on page 177.

- 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.

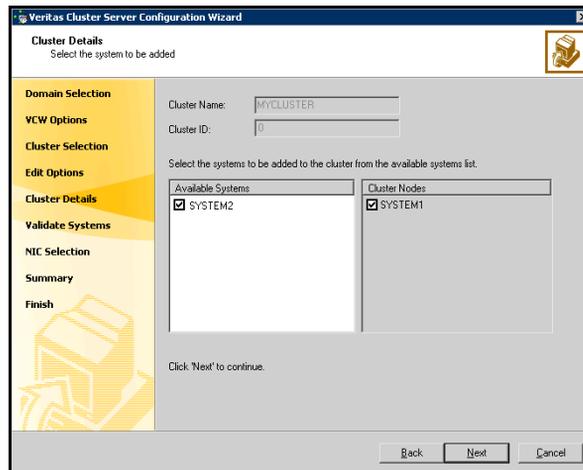
- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 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.
- 10 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, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.
- 11 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.

- 12 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.
- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**.
 How you configure the VCS private network communication depends on

how it is configured in the cluster. If LLT is configured over ethernet, you have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - 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 both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.
The wizard will configure the LLT service (over ethernet) on the selected network adapters.
- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Check the **Configure LLT over UDP** check box.
 - Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
 - Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard will configure the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 14 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.

- 15 Specify the password for the user in whose context the VCS Helper service runs.
- 16 Review the summary information and click **Add**.
- 17 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

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 179.

If a new ClusterService service group needs to be created, be sure to complete the procedure, [“Configuring the Exchange service group for VCS”](#) on page 190 when this procedure appears in the sequence.

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:

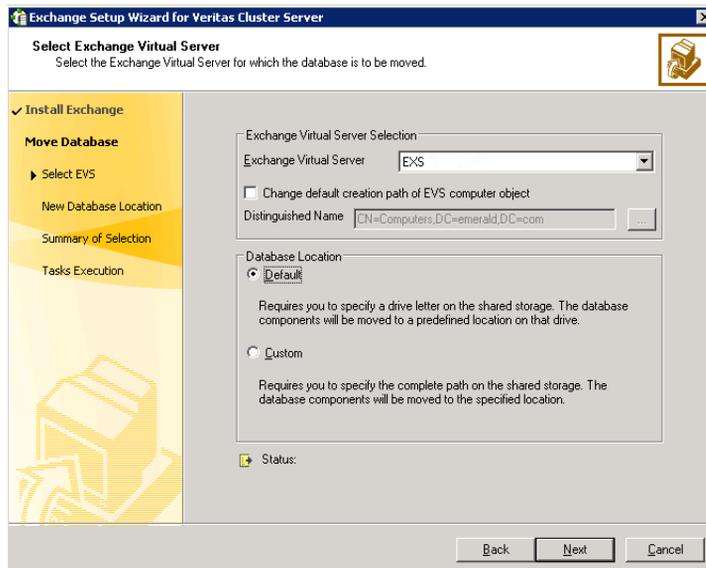
- Make sure the data queue is empty on the SMTP server.
- Make sure to import the disk group and mount the volumes for the Exchange database, MTA data, and transaction logs.
See [“Managing disk groups and volumes”](#) on page 151.

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 Tools > 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, click **Configure/Remove highly available Exchange Server** and then click **Next**.
- 4 In the Select Option dialog box, click **Move Exchange Databases** and then click **Next**.
- 5 In the Select Exchange Virtual Server dialog box, choose the Exchange virtual server and the database location option and then click **Next**.



Exchange Virtual Server

From the drop-down list, select the Exchange virtual server for which you want to move the databases.

Change default creation path of EVS computer object

Perform the following steps if you wish to change the default path for the Exchange virtual server object in Windows Active Directory:

- Check the **Change default creation path of EVS computer object** check box.
- Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**.
 To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
 The Lanman agent performs Windows AD updates. These settings are applicable to the Lanman resource in the service group.
 By default, the Lanman resource adds the virtual server to the default container "Computers."

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

Default

Click **Default** if you wish to move the databases to predefined location on the shared storage. In the next step the wizard prompts you to specify the drive letter on the shared storage. The first mailbox store, public store, and MTA data are then moved to the generated default paths on the volumes that you specify.

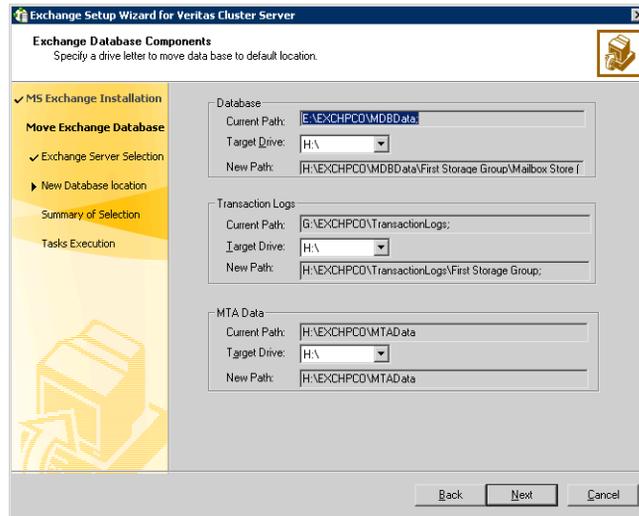
Caution: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

Custom

Click **Custom** if you wish to move the databases to a specific location on the shared storage. Choosing a custom location allows you to specify the Exchange database and streaming path. In the next step the wizard prompts you to specify the entire path of the location on the shared storage. The wizard then moves the databases to the specified directory.

If you chose the Default option, proceed to the next step. If you chose the Custom option, proceed to [step 7](#) on page 183.

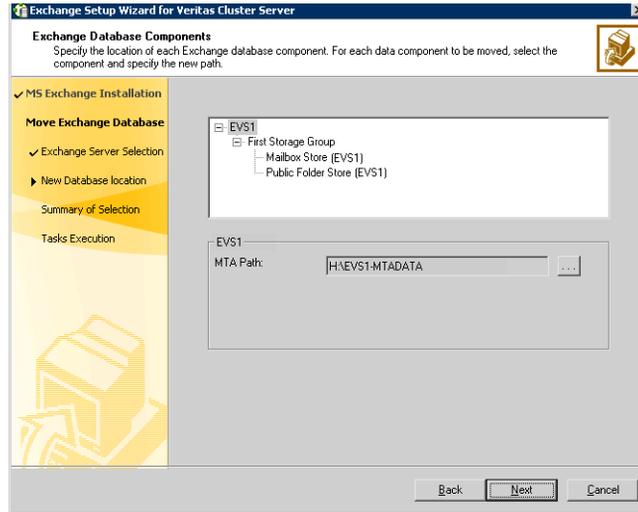
- 6 For the option of a default database location, specify the drives for moving the Exchange database components. The database components are then moved to a predefined location on that drive.



On the Exchange Database Components panel, complete the following steps:

- Specify a drive for moving the Exchange database.
- Specify a drive for moving the Exchange Transaction Logs.
- Specify a drive for moving the Exchange MTA Data.
- Click **Next** and proceed to [step 8](#) on page 183.

- 7 For the option of a custom database location, specify the location for specific Microsoft Exchange data components and then click **Next**.



For each data component that you wish to move, select the component and then click the ellipsis (...) to browse for the folder where you want to move it.

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

- 8 Review the summary of your selections and then click **Next**.
The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task.
- 9 After all the tasks are completed successfully, click **Next**.
- 10 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.
See “[Importing a disk group and mounting a shared volume](#)” on page 151 for instructions.
- Mount the volume containing the information for registry replication (EVS1_SG1_REGREP).
See “[Importing a disk group and mounting a shared volume](#)” on page 151 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).

- 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.
- Either the logged-on user or the VCS Helper service user account must have write permissions on the DNS server to perform DNS updates.
- Make sure the VCS Helper service 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: 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 151.

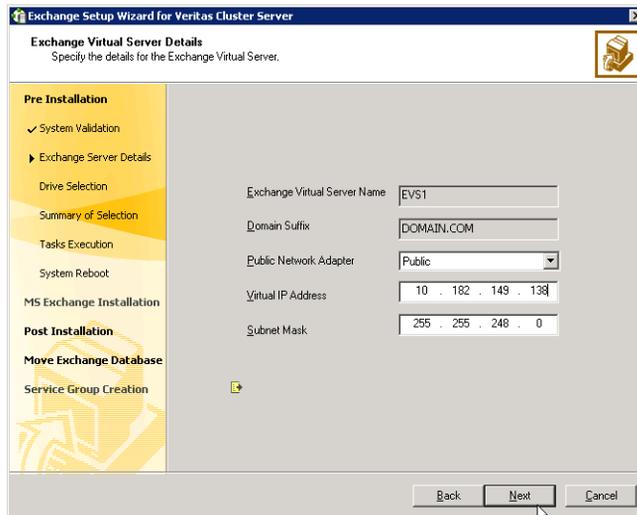
Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase on an additional node. This process changes the physical name of the node to a virtual name.

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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: Additional nodes

Install Exchange on the node on which you performed the 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.

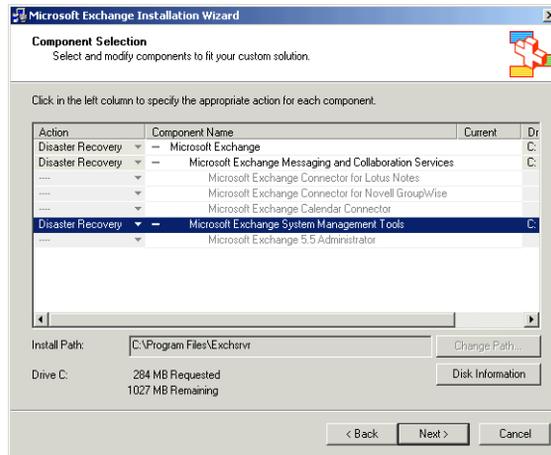
To install Exchange

- 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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:
`C:\>hasys -state`
The state should display as **RUNNING**.

If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```

```
C:\>net start had
```

- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

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.

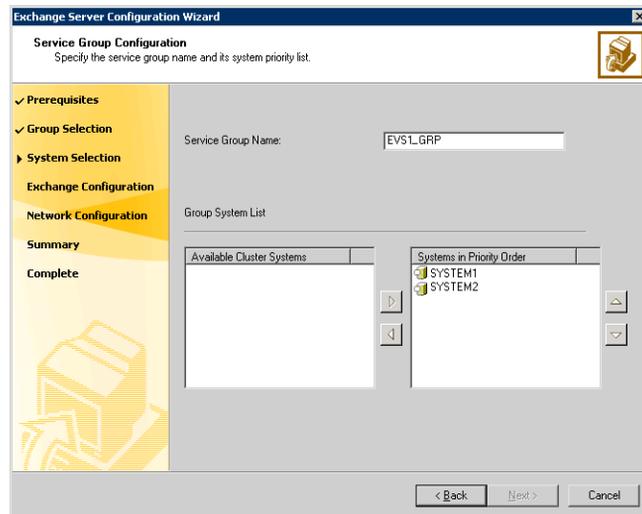
Prerequisites

- You must be a Cluster Administrator. This privilege is required to configure service groups.
- 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 151 for instructions on mounting and “[Unmounting a volume and deporting a disk group](#)” on page 151 for instructions on unmounting.
- 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 the VCS Exchange agent resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on how to add additional resources to the EVS1_SG1_DG disk group.

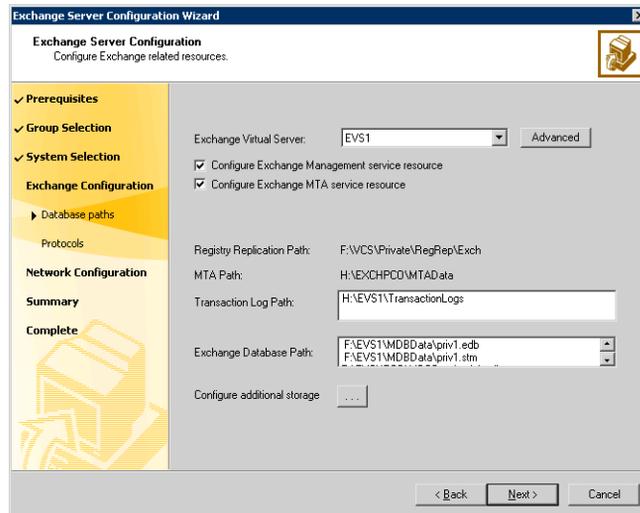
To configure the Exchange service group

- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**:
 The wizard starts validating your configuration. Various messages indicate the validation status.



- Enter a name for the Exchange service group.
 If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.
- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
- Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.

On the Lanman Advanced Configuration dialog box, complete the following:

- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
- Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.

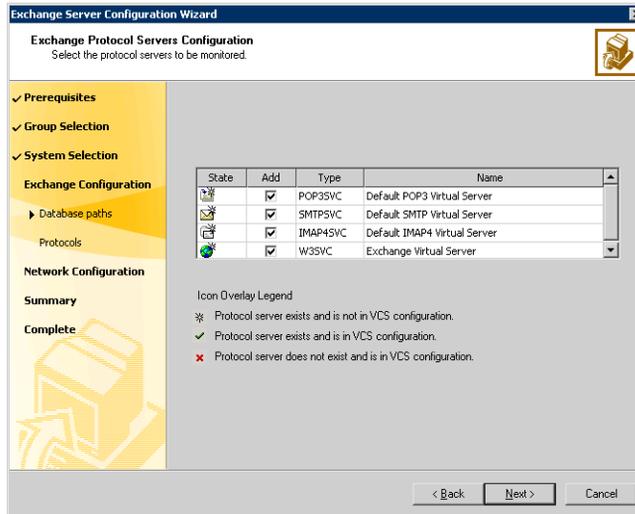
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.

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.

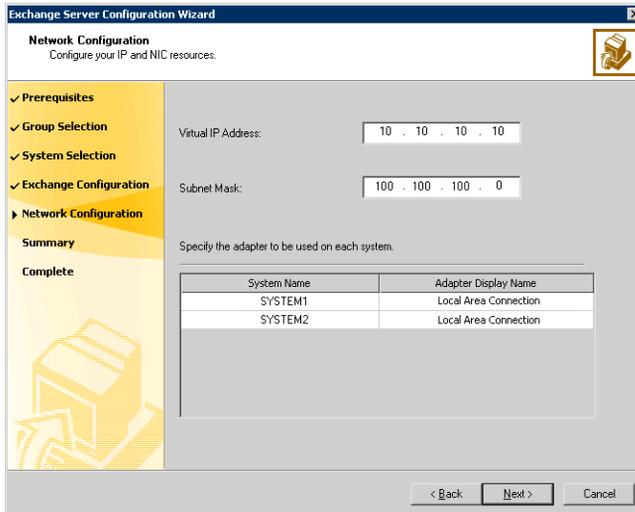
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:

 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- 6 On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.



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

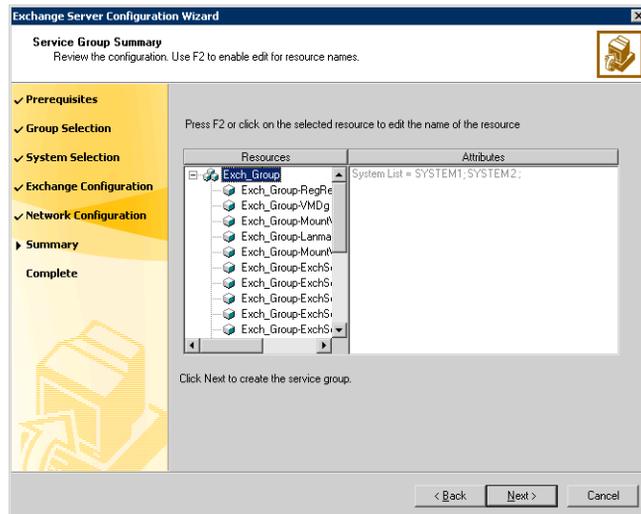


- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
 If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.

- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a node.

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.

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



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 the **Esc** key.

- 9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status

of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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 contains the following topics:

- [“Reviewing the configuration”](#) on page 202
- [“Reviewing the requirements”](#) on page 205
- [“Configuring the storage hardware and network”](#) on page 210
- [“Preparing the forest and domain”](#) on page 211
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 212
- [“Configuring the cluster”](#) on page 218
- [“Configuring the first Exchange virtual server”](#) on page 235
- [“Configuring another Exchange virtual server for an any-to-any failover”](#) on page 266

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 289.

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: New Exchange any-to-any failover configuration

Objective	Tasks
“Reviewing the configuration” on page 202	Understanding a basic any-to-any Exchange configuration
“Reviewing the requirements” on page 205	Verifying hardware and software prerequisites
“Configuring the storage hardware and network” on page 210	<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 211	Setting up the forest and domain prior to the Exchange installation
“Installing Veritas Storage Foundation HA for Windows” on page 212	<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 218	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard
“Managing disk groups and volumes” on page 242	<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

Table 6-1 Task list: New Exchange any-to-any failover configuration

Objective	Tasks
“Installing Exchange on the first node” on page 244	<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 248	<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.
“Installing Exchange on additional nodes” on page 253	<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 258	<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 265	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

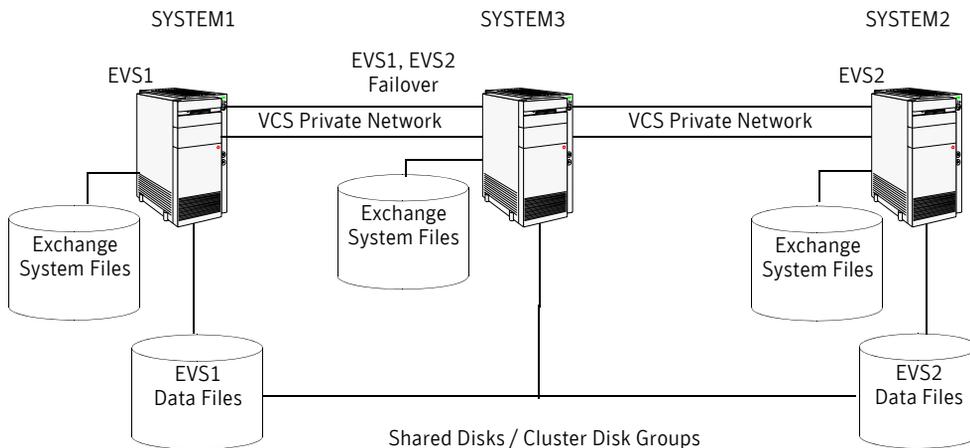
With individual nodes, no failover capability exists. In an any-to-any configuration, the active Exchange nodes can share failover nodes. Additional failover nodes can also exist in an any-to-any configuration.

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 in an any-to-any configuration.

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 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 two 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.

Configuring failover nodes for additional Exchange instances

How you configure failover nodes depends on if Exchange has already been installed on the target node.

In any-to-any configuration, the node you plan to use for failover may already have Exchange installed. For example, you configure an EVS1 cluster on SYSTEM1 and SYSTEM3. SYSTEM3 is the failover node for EVS1. Now you install EVS2 on SYSTEM2. You want to use SYSTEM3 as the failover node for EVS2. In this case, you do not install Exchange once again on SYSTEM3. Instead, you specify SYSTEM3 as a common node for failover.

See “[Specifying a common node for failover](#)” on page 278.

Sample configuration

[Table 6-3](#) on page 203 describes 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

Table 6-3 Sample configuration

Name	Object
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

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.

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

Table 6-4 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

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

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/302144>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agent for Microsoft Exchange.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2003 versions

The following table lists the Microsoft Exchange Server 2003 versions supported with SFW HA 5.1 Service Pack 1.

Table 6-6 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none">■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none">■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Memory must be a minimum 256 MB of RAM per server for Exchange 2003; refer to the Microsoft documentation for more information.
- 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. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
See "[Best practices for SFW HA](#)" on page 209.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each Exchange Virtual Server (EVS).

Reviewing the requirements

- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone 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 Dynamic DNS (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 for SFW HA

The following permissions are required:

- 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 on 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 Exchange virtual server computer object 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, 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 for SFW HA

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 Server 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 2003.
- 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.

Best practices for SFW HA

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.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

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 and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network 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 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, 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 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, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

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

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 6-8 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 213.

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 Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The Select Product screen appears.
- 3 Review the links on the Select Product screen.
Links on this screen access Late Breaking News, the Configuration Checker, as well as begin the process to install Storage Foundation HA for Windows. Click on **Read Late Breaking News** for the latest information about updates, patches, and software issues regarding this release.
- 4 Click **Storage Foundation HA 5.1 SP1 for Windows**.
- 5 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 6 Review the Welcome message and the listed prerequisites. Ensure that any of the listed prerequisites applicable to your installation are met prior to proceeding.
Click **Next**.
- 7 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 AGREE TO the terms of the license agreement**, and then click **Next**.
- 8 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.
- 9 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.
- 10 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. Be sure to select the following as appropriate for your installation.

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, select the option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

11 Select the following for the installation and click **Next**.

Domain	<p>Select a domain from the list.</p> <p>Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.</p>
Computer	<p>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.</p> <p>To remove a computer after adding it, click the name in the Selected computers for installation field and click Remove.</p> <p>Click a computer's name to see its description.</p> <p>When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.</p>

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

12 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.

13 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.

14 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages.

If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning:

The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If applicable to your installation, perform the above procedure.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

If applicable to your installation, perform the above procedure.

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation 5.1 for Windows (Server Components) with the MSCS option, you may want MSCS Quorum Arbitration Time (Min. and Max) to be adjusted to ensure optimal functionality with Veritas dynamic volumes with MSCS.

For additional information, see the *Storage Foundation for Windows Administrator Guide* for details.

If applicable to your installation, perform the above procedure.

- 15 When finished reviewing the message or messages, click **OK**.
- 16 The Summary screen appears displaying an Install report. Review the information in the Install report. Click **Back** to make changes, if necessary. Click **Install** if information is validated.
- 17 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.
- 18 When the installation completes, review the summary screen and click **Next**.
- 19 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.
- 20 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 21 Review the log files and click **Finish**.
- 22 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 Cluster Configuration Wizard (VCW) 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.

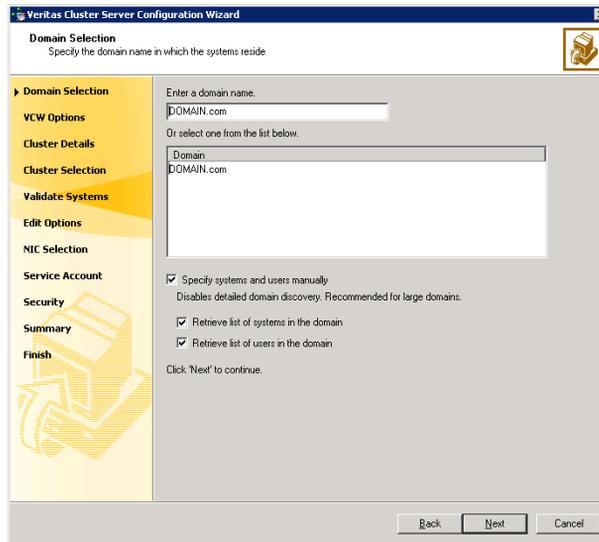
Complete the following tasks before creating a cluster:

- 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 HAD Helper service, you must know the password for the user account.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. 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. This account does not require domain admin privileges.
 - Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.



Do one of the following:

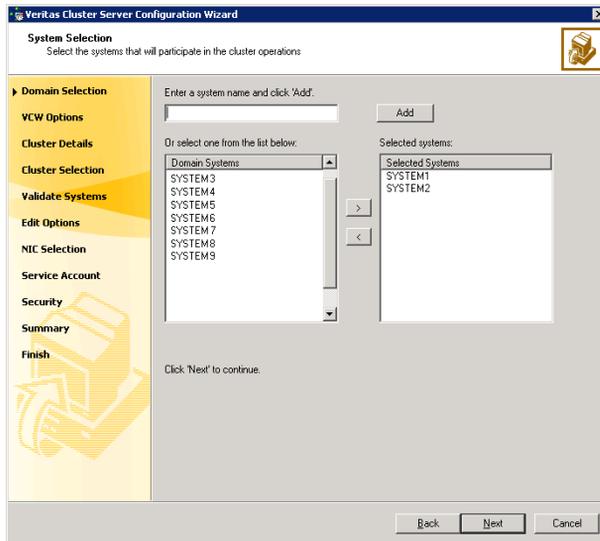
- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 221.
 - 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 220. 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 8](#) on page 221.

- 6 On the System Selection panel, specify the systems for the 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.

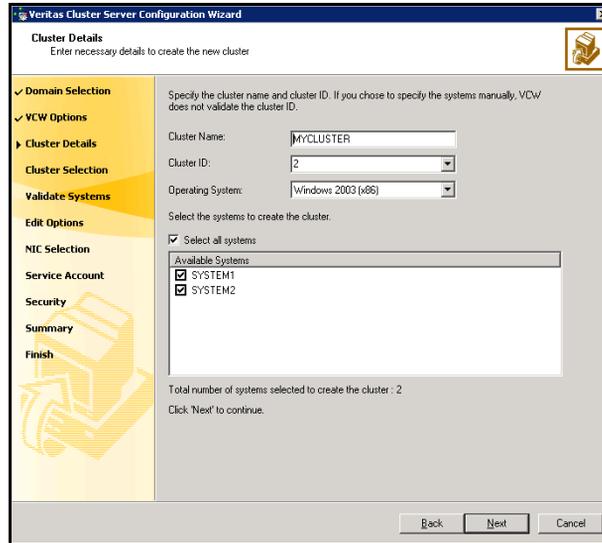
- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 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. The cluster ID can be any number from 0 to 255.

Caution: 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. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (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.

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

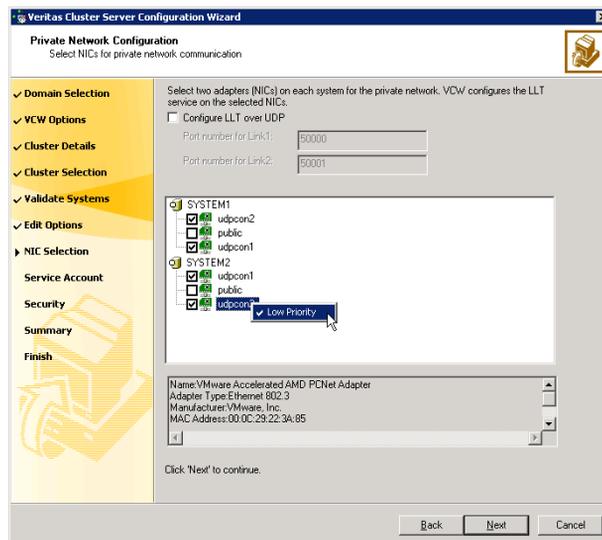
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 224.

11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer.

Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



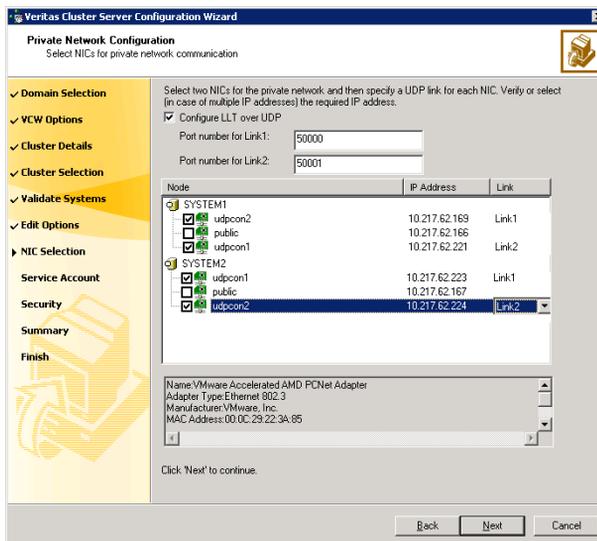
- Select the check boxes next to the two NICs to be assigned to the private network.

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 both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



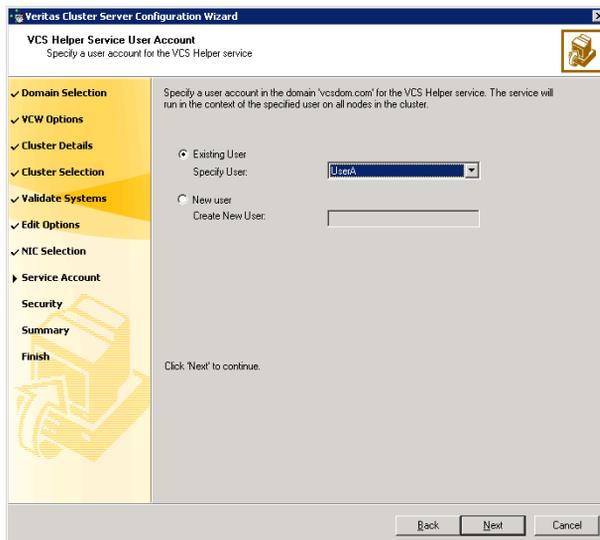
- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to

65535. The default ports numbers are 50000 and 50001 respectively.

- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.



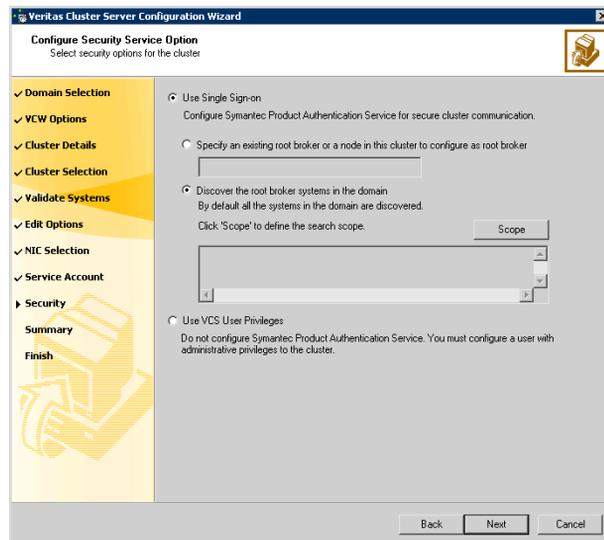
Specify a domain user as follows:

- 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 219, 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**.

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

Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability 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 an existing root broker or a node in this cluster to configure as root broker**, 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 then configures all nodes in the cluster as authentication brokers.

- If you want to search 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 a user *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**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 6-9 contains some more examples of search criteria.

Table 6-9 Search criteria examples

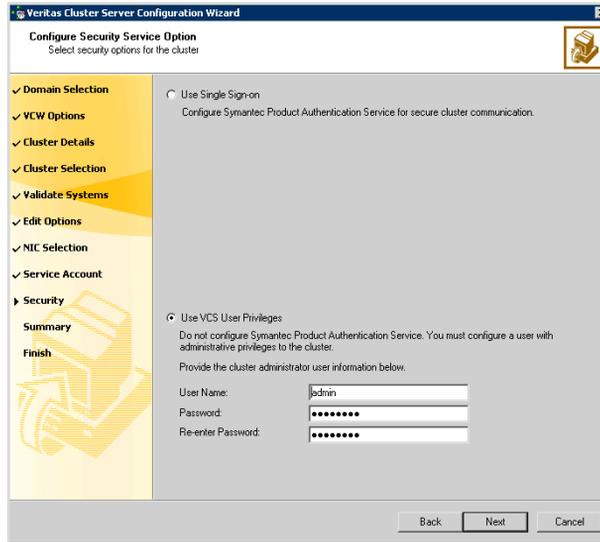
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- 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 a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS`Encrypt` utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password. After the cluster is configured, you can 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**.

14 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.

- 15 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**.

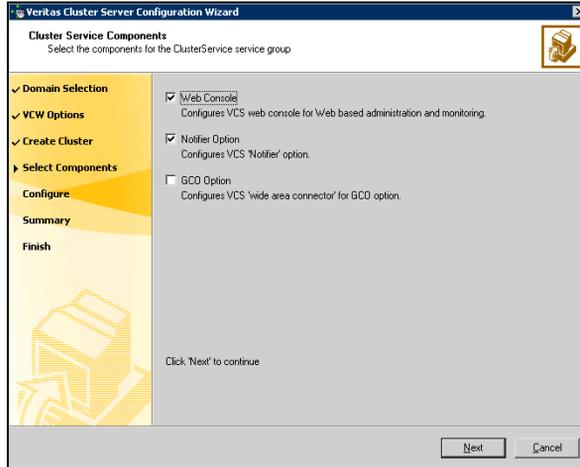
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

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.

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



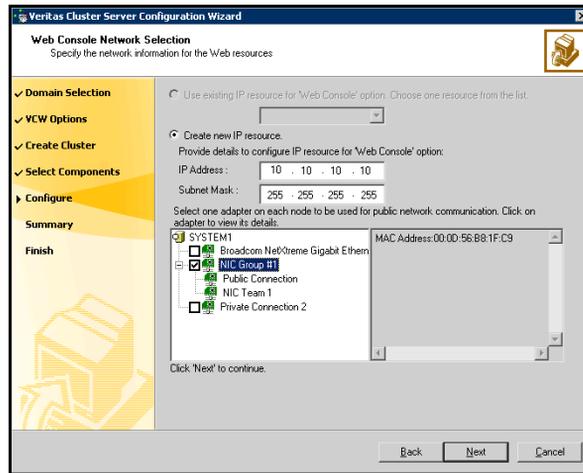
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 230.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 231.

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**.



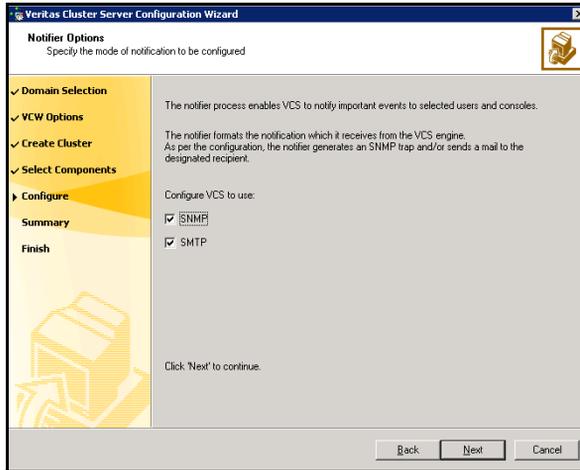
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 231. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

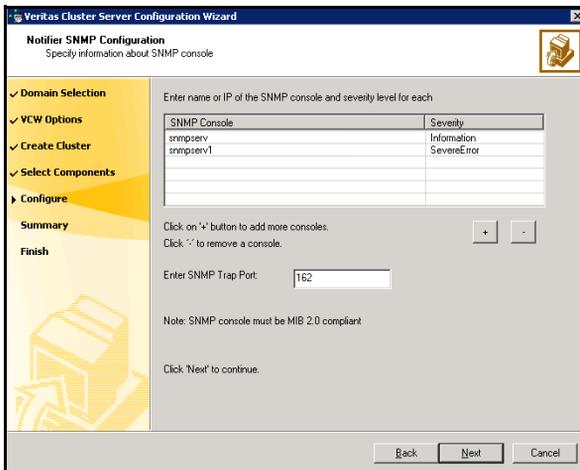
To configure notification

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

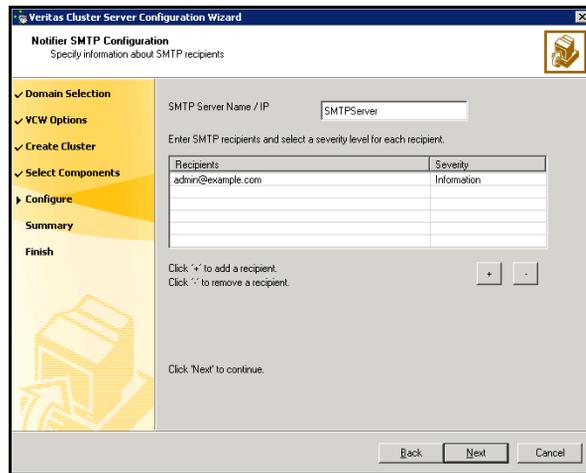


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

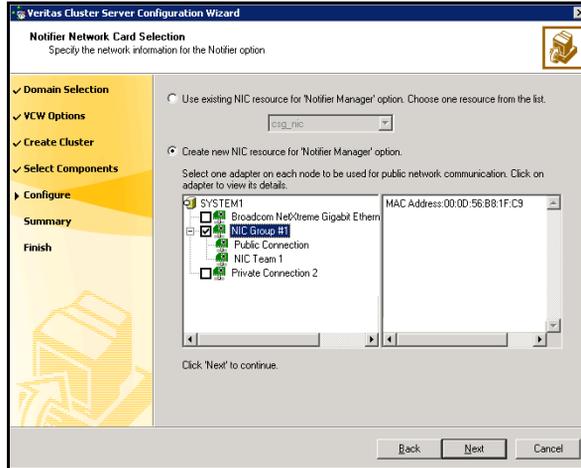


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.
 - 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 202.

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.

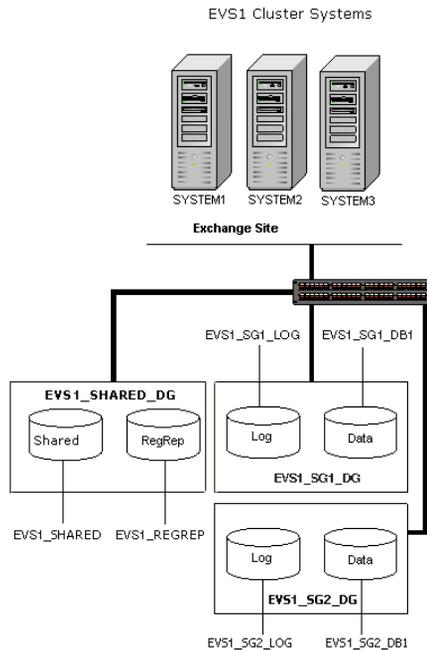
Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator’s Guide* for more information.

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 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_LOG: Contains the transaction log for the storage group.

Exchange disk group EVS1_SHARED_DG create contains the following 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: A disaster recovery configuration with VVR requires a Storage Replicator Log (SRL) volume (EVS1_REPLOG) for each disk group that contains volumes that are replicated. You can create the SRL volume now or you can create it later when you run the Disaster Recovery Wizard. If you create it later, ensure that you allow sufficient disk space for this volume. For more about VVR planning, see the *Veritas Volume Replicator, Administrator's Guide*.

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

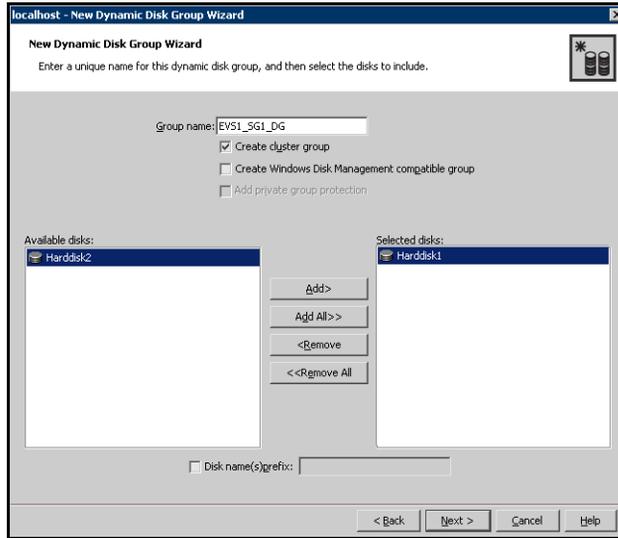
Steps to create a disk group are as follows.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

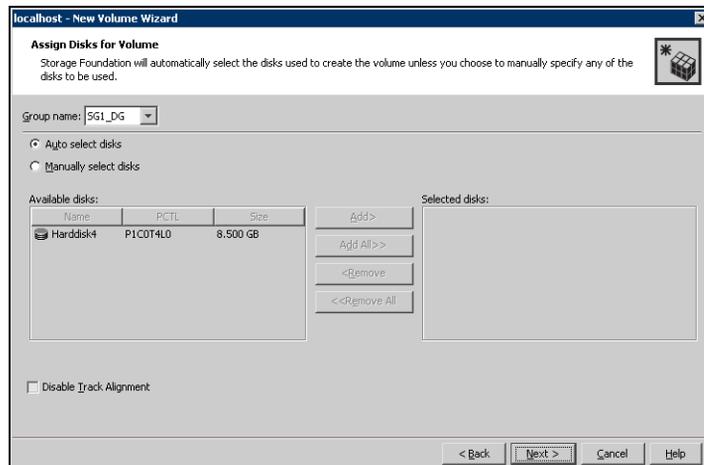
- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

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.

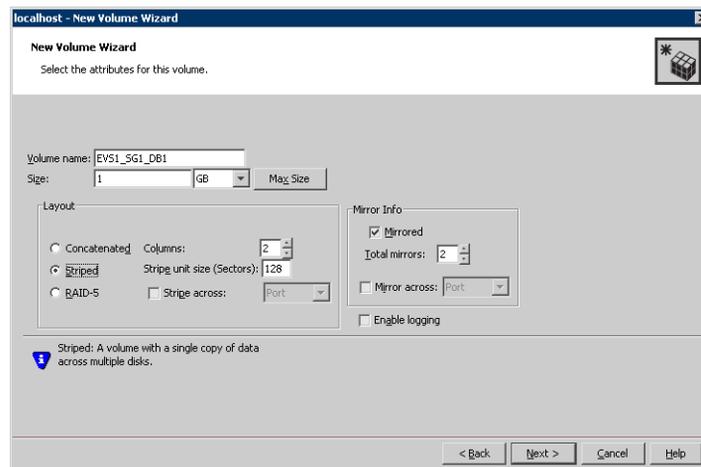


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

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

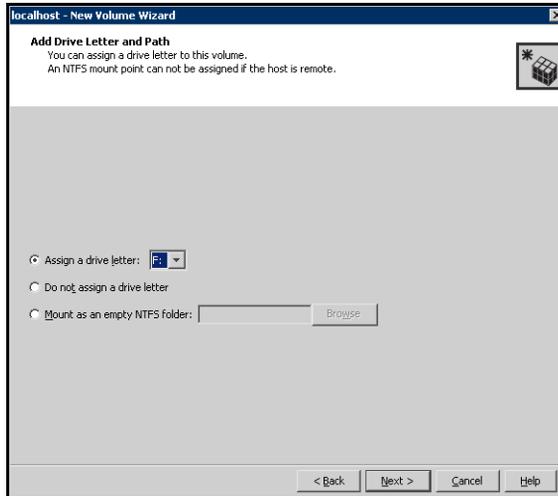
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling **Track Alignment** means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

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



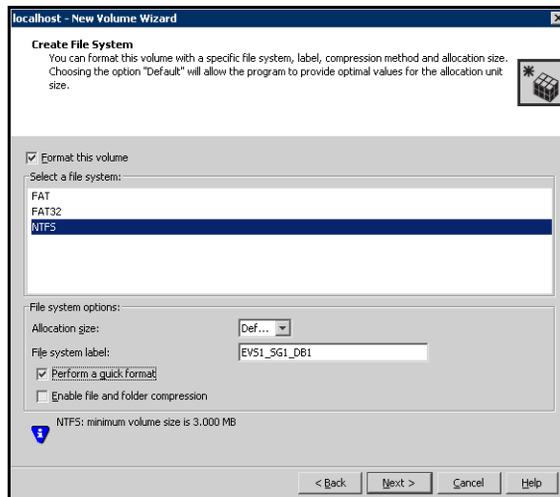
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns**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

- 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 any additional volumes required. Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk groups and volumes

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - *Control Panel - System Settings*). See the Veritas Storage Foundation Administrator's Guide for more information.

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 pre-installation, 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 211.
- Verify the disk group is imported on the first node of the cluster.
See “[Managing disk groups and volumes](#)” on page 242.
- Mount the volume containing the information for registry replication (EVS1_REGREP).
See “[Managing disk groups and volumes](#)” on page 242.
- 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - 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.

- Either the logged-on user or the VCS Helper service user account must have write permissions on the DNS server to perform DNS updates.
- Make sure the VCS Helper service 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.

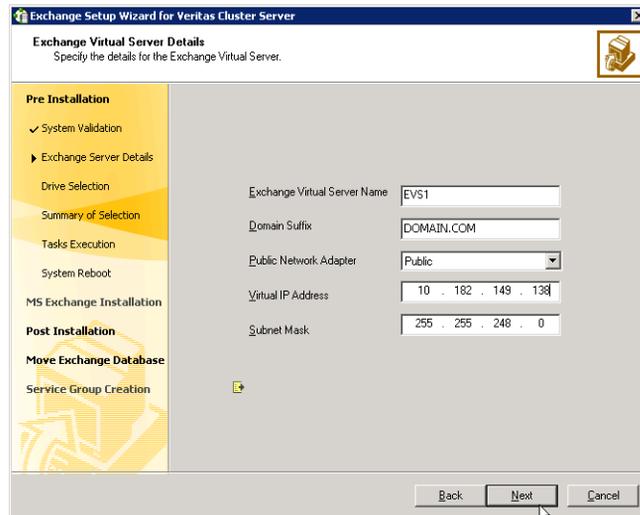
Exchange pre-installation: First node

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 Tools > 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.

Warning: 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 unique on the network.
- 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.

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

If you wish to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: First node

Install Exchange on the node on which you performed the pre-installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```

The state should display as **RUNNING**.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```

```
C:\>net start had
```
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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.
Once the node is rebooted, move the databases created during the Exchange installation, from the local drive to the shared storage.

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.

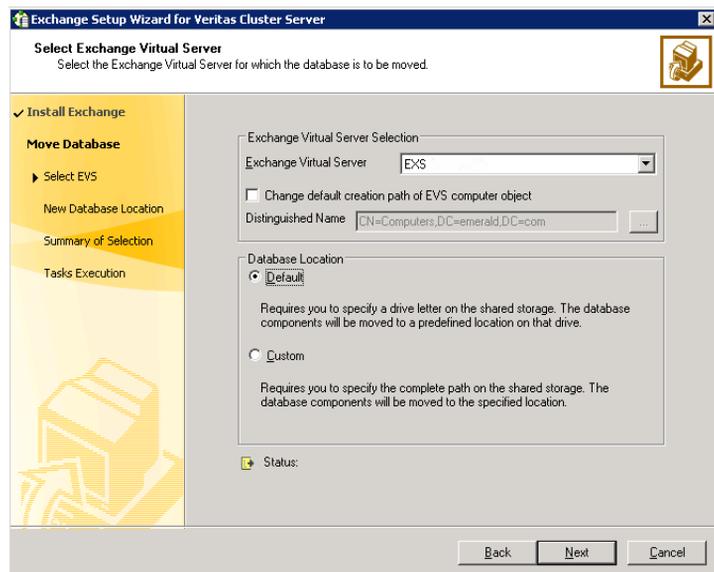
Complete the following tasks before moving the databases:

- 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.
See [“Managing disk groups and volumes”](#) on page 242 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 Tools > 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, click **Configure/Remove highly available Exchange Server** and then click **Next**.
- 4 In the Select Option dialog box, click **Move Exchange Databases** and then click **Next**.
- 5 In the Select Exchange Virtual Server dialog box, choose the Exchange virtual server and the database location option and then click **Next**.



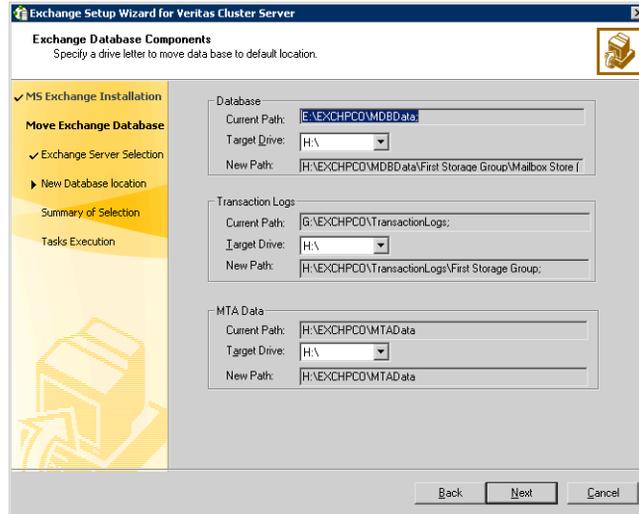
Exchange Virtual Server

From the drop-down list, select the Exchange virtual server for which you want to move the databases.

Change default creation path of EVS computer object	<p>Perform the following steps if you wish to change the default path for the Exchange virtual server object in Windows Active Directory:</p> <ul style="list-style-type: none">■ Check the Change default creation path of EVS computer object check box.■ Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format CN=containername,DC=domainname,DC=com. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. The Lanman agent performs Windows AD updates. These settings are applicable to the Lanman resource in the service group. By default, the Lanman resource adds the virtual server to the default container "Computers."
Default	<p>Click Default if you wish to move the databases to predefined location on the shared storage. In the next step the wizard prompts you to specify the drive letter on the shared storage. The first mailbox store, public store, and MTA data are then moved to the generated default paths on the volumes that you specify.</p> <p>Caution: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).</p>
Custom	<p>Click Custom if you wish to move the databases to a specific location on the shared storage. Choosing a custom location allows you to specify the Exchange database and streaming path. In the next step the wizard prompts you to specify the entire path of the location on the shared storage. The wizard then moves the databases to the specified directory.</p>

If you chose the Default option, proceed to the next step. If you chose the Custom option, proceed to [step 7](#) on page 252.

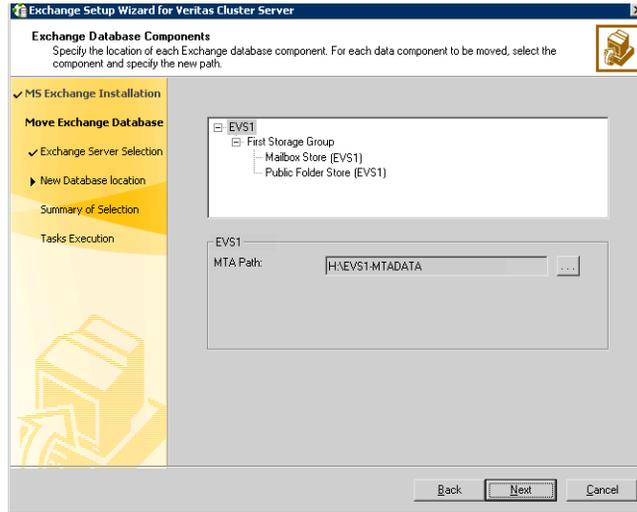
- 6 For the option of a default database location, specify the drives for moving the Exchange database components. The database components are then moved to a predefined location on that drive.



On the Exchange Database Components panel, complete the following steps:

- Specify a drive for moving the Exchange database.
- Specify a drive for moving the Exchange Transaction Logs.
- Specify a drive for moving the Exchange MTA Data.
- Click **Next** and proceed to [step 8](#) on page 252.

- 7 For the option of a custom database location, specify the location for specific Microsoft Exchange data components and then click **Next**.



For each data component that you wish to move, select the component and then click the ellipsis (...) to browse for the folder where you want to move it.

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

- 8 Review the summary of your selections and then click **Next**.
The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task.
- 9 After all the tasks are completed successfully, click **Next**.
- 10 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 pre-installation, 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 244.

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.

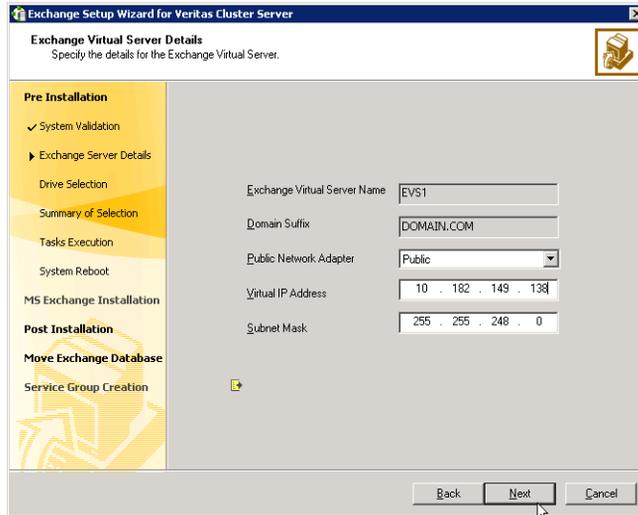
See [“Managing disk groups and volumes”](#) on page 242 for instructions.

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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: Additional nodes

Install Exchange on the additional node on which you performed the 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.

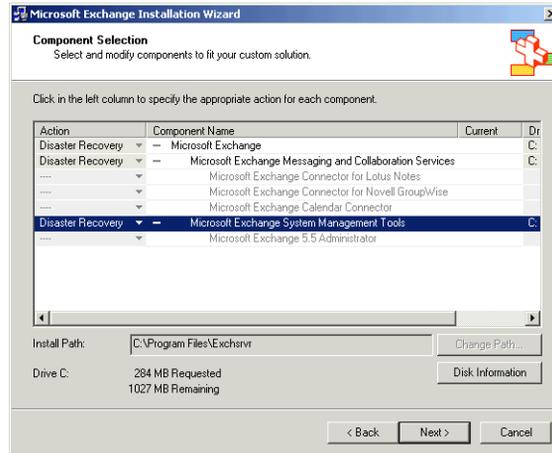
To install Exchange

- 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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:
`C:\>hasys -state`
The state should display as **RUNNING**.
If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`

- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

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

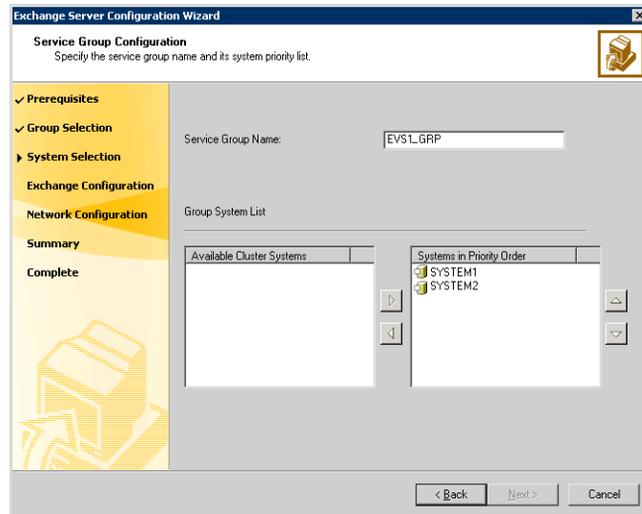
- You must be a Cluster Administrator.
- You must be a local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- 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 242 for instructions on mounting and unmounting.
- 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 the VCS Exchange agent resource types, attribute definitions, resource dependencies, and sample service group configurations.

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

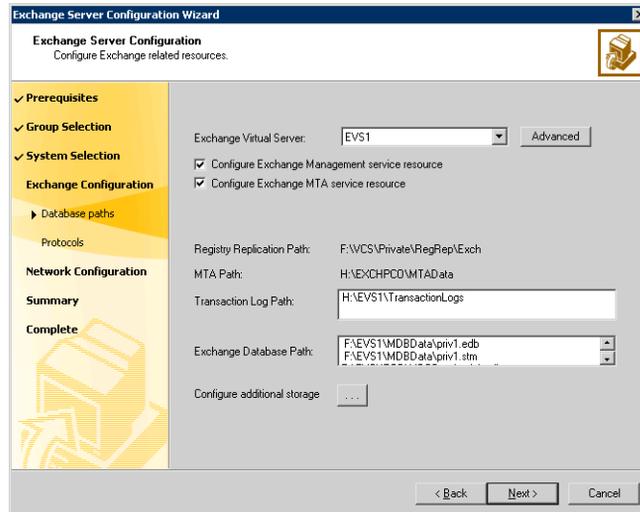
To configure the Exchange service group

- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**: The wizard starts validating your configuration. Various messages indicate the validation status.



- Enter a name for the Exchange service group.
If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.
- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
- Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.

On the Lanman Advanced Configuration dialog box, complete the following:

- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
- Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.

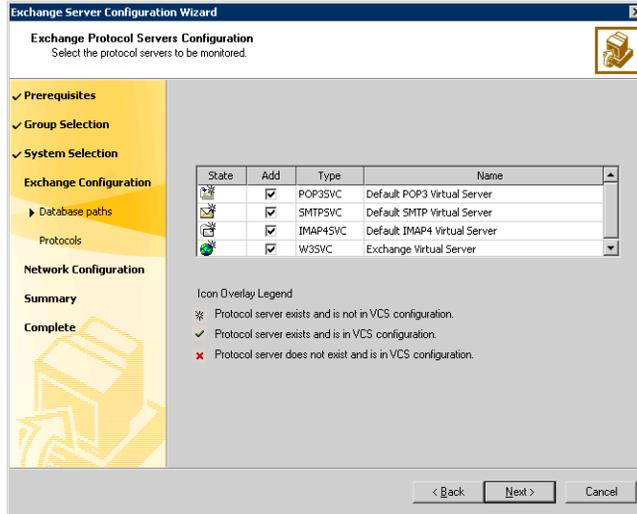
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.

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.

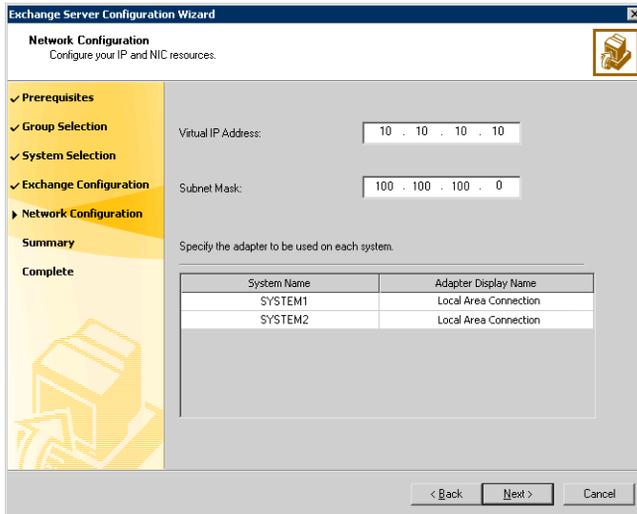
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:

 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- 6 On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.



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

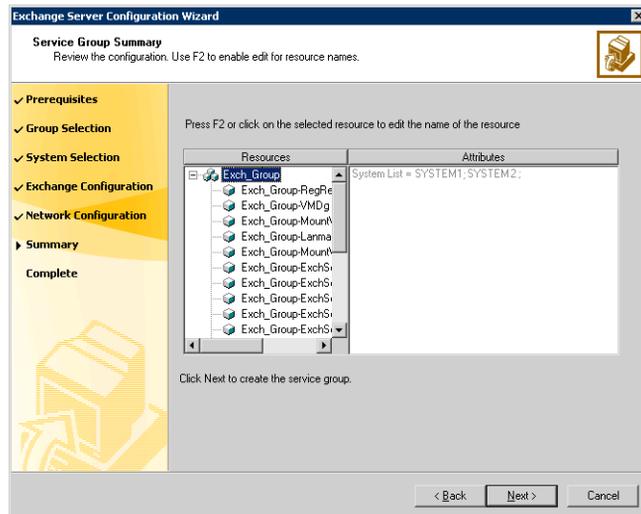


- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
 If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.

- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a node.

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.

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



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 the **Esc** key.

- 9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status

of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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 SYSTEM2 and SYSTEM3.

See “[Reviewing the configuration](#)” on page 202.

Configuring disk groups and volumes

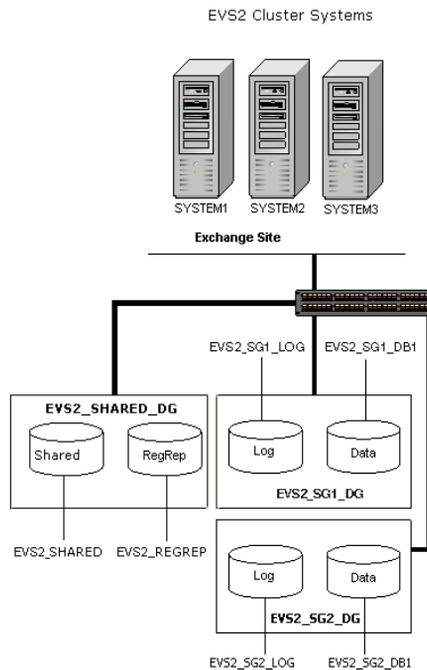
Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to grow existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

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 group 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 the following 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 the following 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 237

For instructions on creating volumes:

see “[Creating volumes](#)” on page 239.

Managing disk groups and volumes

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

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 pre-installation, 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 242.
- Mount the volume containing the information for registry replication (EVS2_REGREP).
See “[Managing disk groups and volumes](#)” on page 242.
- 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - 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.
 - Either the logged-on user or the VCS Helper service user account must have write permissions on the DNS server to perform DNS updates.

- Make sure the VCS Helper service 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.

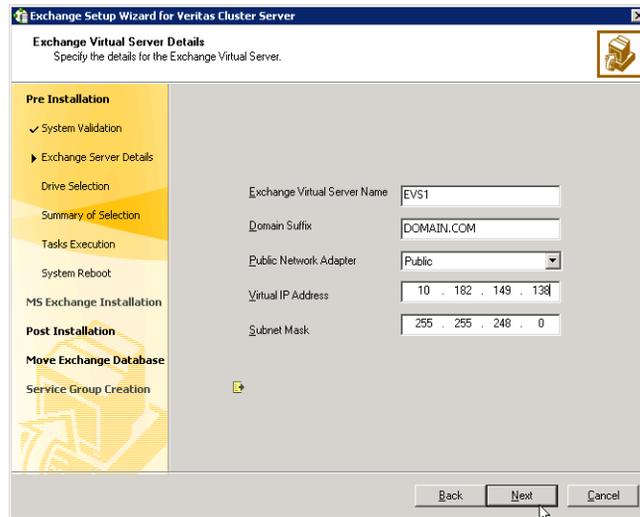
Exchange pre-installation: first node of an additional Exchange Virtual Server

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 Tools > 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.

Warning: 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 unique on the network.
- 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.

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

If you wish to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: first node of an additional Exchange virtual server

Install Exchange on the node on which you performed the pre-installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```


The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```



```
C:\>net start had
```
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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

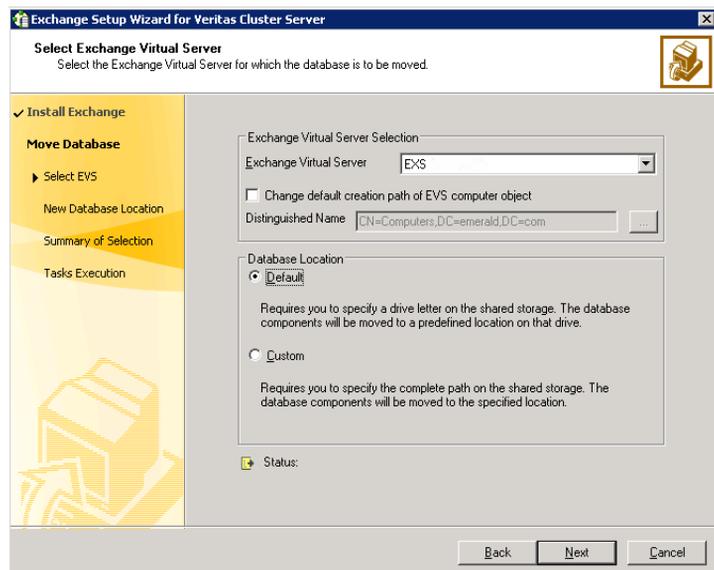
After completing the Exchange installation on the first node, move the Exchange databases on the first node from the local drive to the shared drive to ensure proper failover operations in the cluster. Complete the following tasks before moving the databases:

- 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, transaction logs, and MTA data.
See [“Managing disk groups and volumes”](#) on page 242.

- Start VEA and go to SYSTEM1, 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 Tools > 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, click **Configure/Remove highly available Exchange Server** and then click **Next**.
- 4 In the Select Option dialog box, click **Move Exchange Databases** and then click **Next**.
- 5 In the Select Exchange Virtual Server dialog box, choose the Exchange virtual server and the database location option and then click **Next**.



Exchange Virtual Server

From the drop-down list, select the Exchange virtual server for which you want to move the databases.

Change default creation path of EVS computer object

Perform the following steps if you wish to change the default path for the Exchange virtual server object in Windows Active Directory:

- Check the **Change default creation path of EVS computer object** check box.
- Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**.
To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
The Lanman agent performs Windows AD updates. These settings are applicable to the Lanman resource in the service group.
By default, the Lanman resource adds the virtual server to the default container "Computers."

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

Default

Click **Default** if you wish to move the databases to predefined location on the shared storage. In the next step the wizard prompts you to specify the drive letter on the shared storage. The first mailbox store, public store, and MTA data are then moved to the generated default paths on the volumes that you specify.

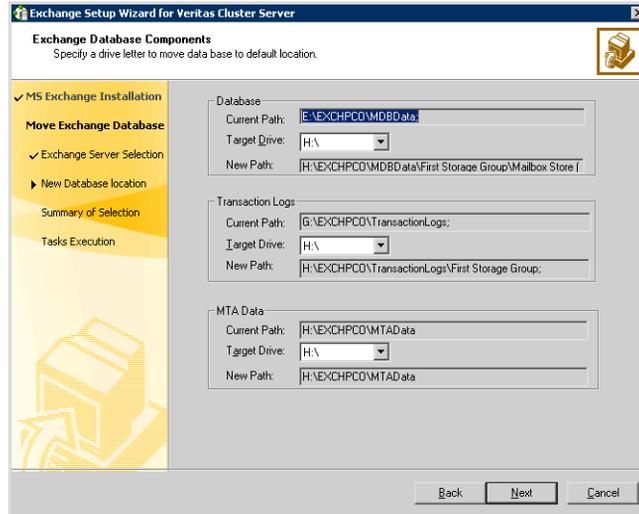
Caution: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

Custom

Click **Custom** if you wish to move the databases to a specific location on the shared storage. Choosing a custom location allows you to specify the Exchange database and streaming path. In the next step the wizard prompts you to specify the entire path of the location on the shared storage. The wizard then moves the databases to the specified directory.

If you chose the Default option, proceed to the next step. If you chose the Custom option, proceed to [step 7](#) on page 252.

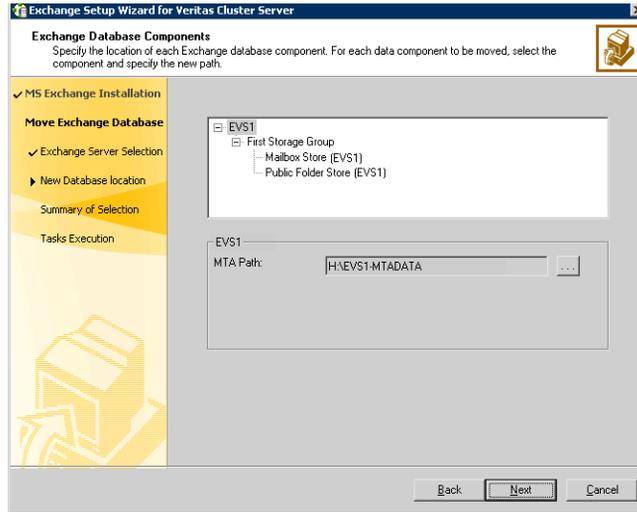
- 6 For the option of a default database location, specify the drives for moving the Exchange database components. The database components are then moved to a predefined location on that drive.



On the Exchange Database Components panel, complete the following steps:

- Specify a drive for moving the Exchange database.
- Specify a drive for moving the Exchange Transaction Logs.
- Specify a drive for moving the Exchange MTA Data.
- Click **Next** and proceed to [step 8](#) on page 252.

- 7 For the option of a custom database location, specify the location for specific Microsoft Exchange data components and then click **Next**.



For each data component that you wish to move, select the component and then click the ellipsis (...) to browse for the folder where you want to move it.

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

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

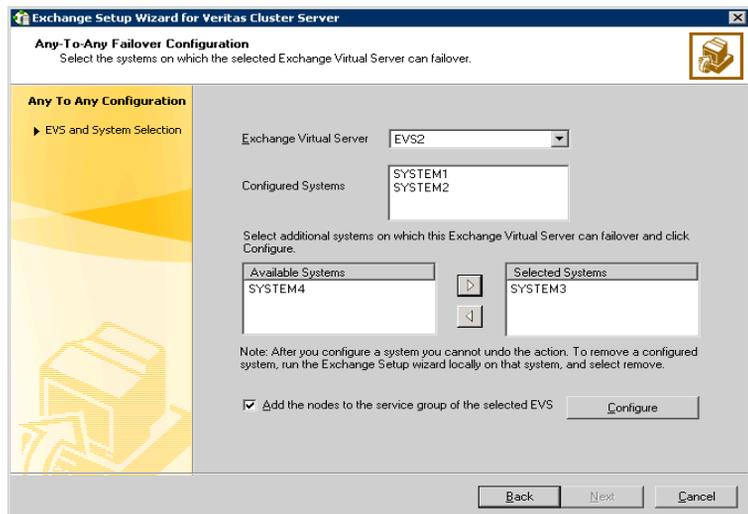
Repeat the task below for each additional Exchange Virtual Server.

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.

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

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 Tools > 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. This privilege is required to configure service groups.
- 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 database

See “[Managing disk groups and volumes](#)” on page 242.

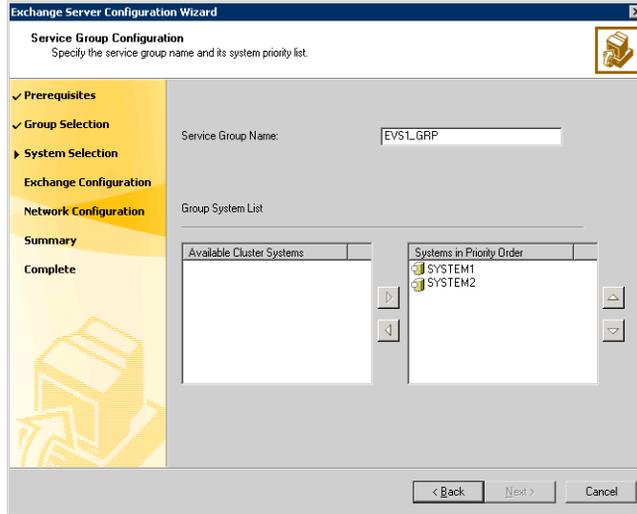
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name. Refer to the appropriate DNS documentation for further information.
- Verify Microsoft Exchange is installed and configured identically on all nodes.

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

To configure the Exchange service group

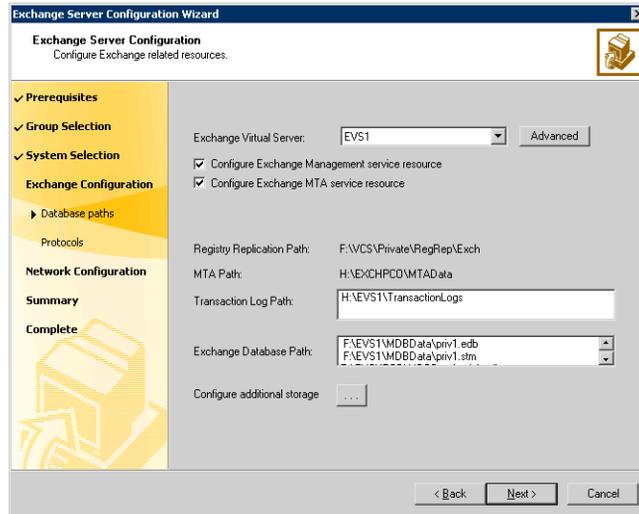
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**:

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



- Enter a name for the Exchange service group.
If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.

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



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
 - Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.
- On the Lanman Advanced Configuration dialog box, complete the following:
- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
 - Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.

If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

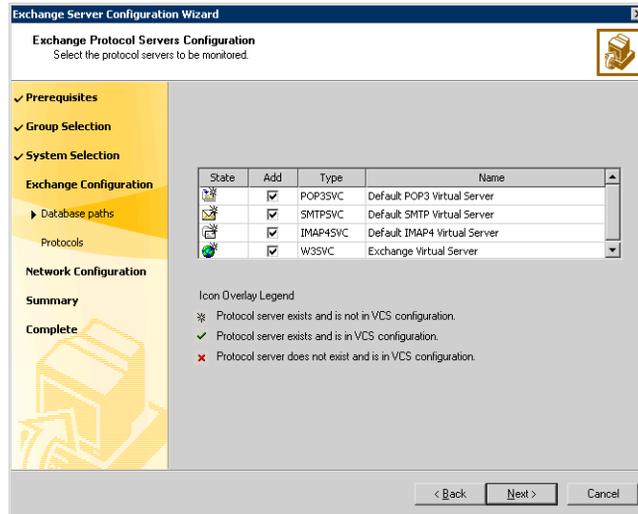
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.

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.

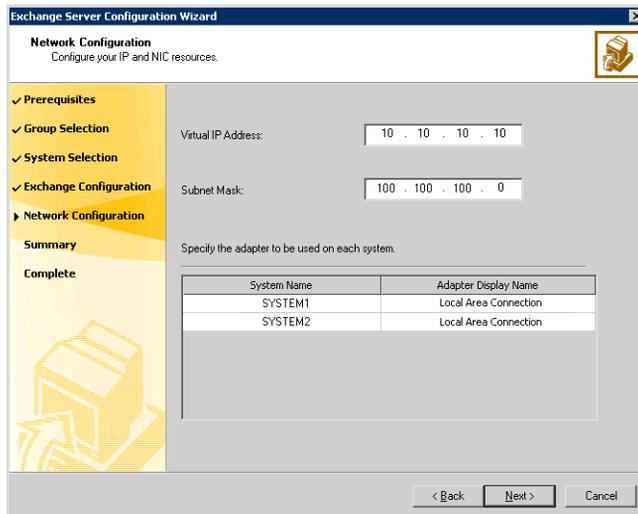
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:
 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- 6 On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.

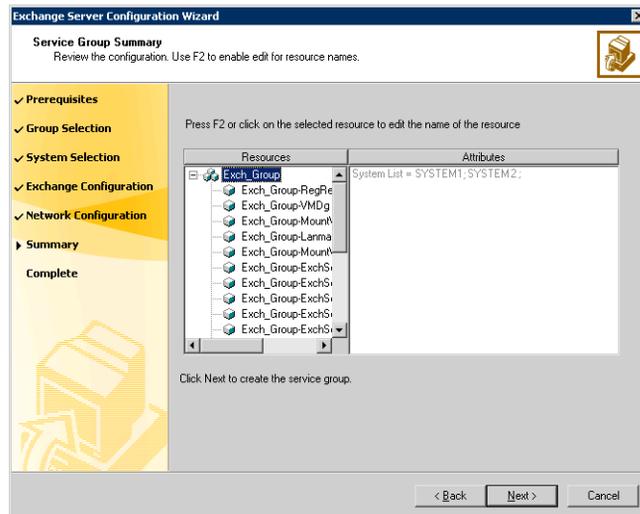


- 7 On the Network Configuration panel, specify information related to the network and then click **Next**:



- 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 Display Name** field to view the adapters associated with a node.
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.

8 Review the service group configuration, change the resource names, if desired, and then click **Next**:



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 the **Esc** key.
- 9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status

of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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 contains the following topics:

- [“Reviewing the requirements”](#) on page 292
- [“Configuring new nodes: prior to creating additional Exchange virtual server”](#) on page 300
- [“Specifying a common node for failover”](#) on page 338
- [“Verifying the cluster configuration”](#) on page 346

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.

Table 7-1 on page 290 outlines the high-level objectives and the tasks to complete each objective:

Table 7-1 Task List

Objective	Tasks
“Reviewing the requirements” on page 292	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 298	Understanding a basic any-to-any Exchange configuration, starting with an existing active-passive cluster and adding nodes.
“Configuring the storage hardware and network” on page 300	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 302	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 308	<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 to be added, completing the steps in “Adding a node to a cluster” on page 308.
“Configuring disk groups and volumes” on page 314	<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 (Continued)

Objective	Tasks
“Installing Exchange on the new nodes” on page 323	Installing Exchange on the new active Exchange node, a “first node”
“Moving Exchange databases to shared storage (EVS2)” on page 328	Moving the databases from the new active Exchange node to shared storage
“Installing Exchange on additional nodes” on page 332	Adding additional new failover nodes, if any
“Specifying a common node for failover” on page 338	<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 338.
“Verifying the cluster configuration” on page 346	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 298 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 45 to create a new active/passive cluster.
 - For an existing standalone Exchange server, refer to the chapter “[Deploying SFW HA for high availability: Standalone Exchange servers](#)” on page 121 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 virtual 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 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.

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

Table 7-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

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

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/302144>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agent for Microsoft Exchange.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2003 versions

The following table lists the Microsoft Exchange Server 2003 versions supported with SFW HA 5.1 Service Pack 1.

Table 7-4 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

Table 7-4 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Memory must be a minimum 256 MB of RAM per server for Exchange 2003; refer to the Microsoft documentation for more information.
- 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. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
See "[Best practices for SFW HA](#)" on page 297.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each Exchange Virtual Server (EVS).

Reviewing the requirements

- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone 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 Dynamic DNS (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 for SFW HA

The following permissions are required:

- 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 on 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 Exchange virtual server computer object 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, 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 for SFW HA

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 Server 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 2003.
- 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.

Best practices for SFW HA

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.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

Reviewing the configuration

Table 7-6 on page 298 details the systems in an active-passive configuration (SYSTEM1 and SYSTEM2) plus a new node (SYSTEM3) transformed into an any-to-any configuration:

Table 7-6 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

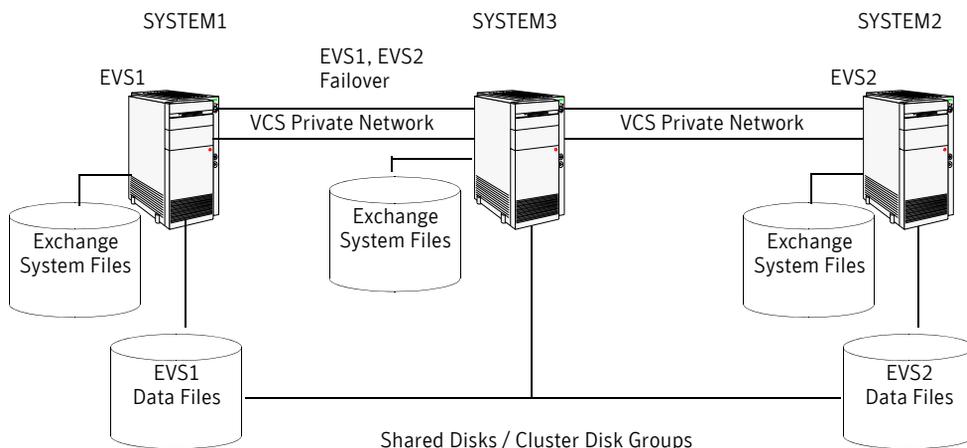
In an active-passive configuration, a separate failover system is required for each active Exchange node. In an any-to-any configuration, the active Exchange nodes can share failover nodes. Additional failover nodes can also exist in an any-to-any configuration.

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 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

Table 7-7 on page 299 describes the objects created and used during the installation and configuration tasks:

Table 7-7 Sample configuration

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3	Physical node names
EVS1, EVS2	Microsoft Exchange Virtual Servers
EVS1_GRP, EVS2_GRP	Microsoft Exchange service groups

Table 7-7 Sample configuration

Name	Object
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 storage hardware and network

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 and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network 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 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, 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

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 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, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

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

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 7-8 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 304.

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 Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The Select Product screen appears.
- 3 Review the links on the Select Product screen.
Links on this screen access Late Breaking News, the Configuration Checker, as well as begin the process to install Storage Foundation HA for Windows. Click on **Read Late Breaking News** for the latest information about updates, patches, and software issues regarding this release.
- 4 Click **Storage Foundation HA 5.1 SP1 for Windows**.
- 5 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 6 Review the Welcome message and the listed prerequisites. Ensure that any of the listed prerequisites applicable to your installation are met prior to proceeding.

Click **Next**.

- 7 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 AGREE TO the terms of the license agreement**, and then click **Next**.
- 8 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.
- 9 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.
- 10 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. Be sure to select the following as appropriate for your installation.

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, select the option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

- 11 Select the following 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.

When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

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`

- 12 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.
- 13 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.

14 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages.

If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning:

The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If applicable to your installation, perform the above procedure.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

If applicable to your installation, perform the above procedure.

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation 5.1 for Windows (Server Components) with the MSCS option, you may want MSCS Quorum Arbitration Time (Min. and Max) to be adjusted to ensure optimal functionality with Veritas dynamic volumes with MSCS.

For additional information, see the *Storage Foundation for Windows Administrator Guide* for details.

If applicable to your installation, perform the above procedure.

15 When finished reviewing the message or messages, click **OK**.

16 The Summary screen appears displaying an Install report. Review the information in the Install report. Click **Back** to make changes, if necessary. Click **Install** if information is validated.

17 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.

18 When the installation completes, review the summary screen and click **Next**.

- 19 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.
- 20 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 21 Review the log files and click **Finish**.
- 22 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 308.

Adding a node to a cluster

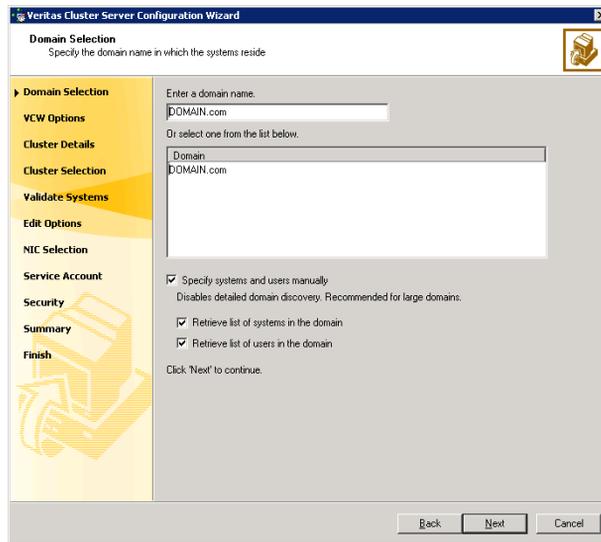
This procedure only applies to any-to-any clusters transformed from active-passive clusters that do not already have all the nodes for the any-to-any configuration in the cluster.

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.

Note: Run the VCS Cluster Configuration Wizard (VCW) from the standalone node or a node in the cluster.

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

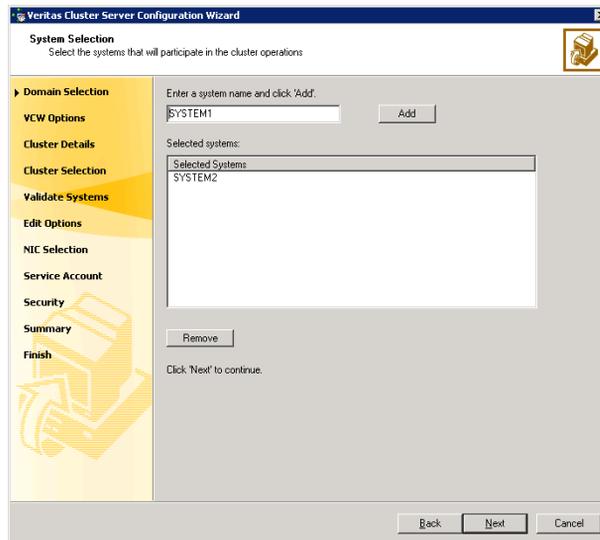


Do one of the following:

- To discover information about all the systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

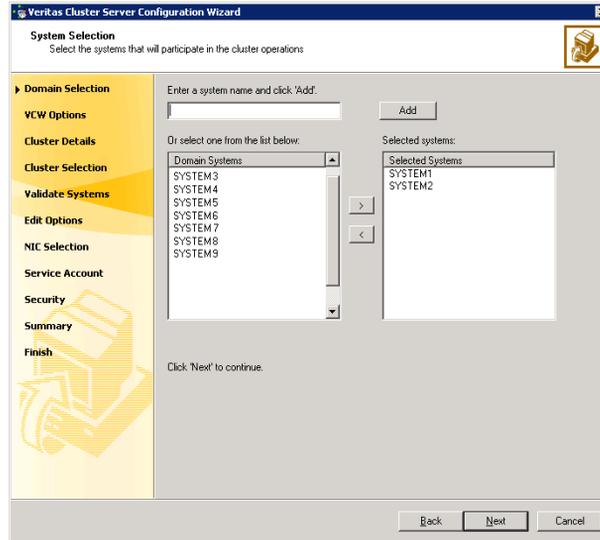
Proceed to [step 8](#) on page 312.

- 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 311. 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 8](#) on page 312.

- 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.

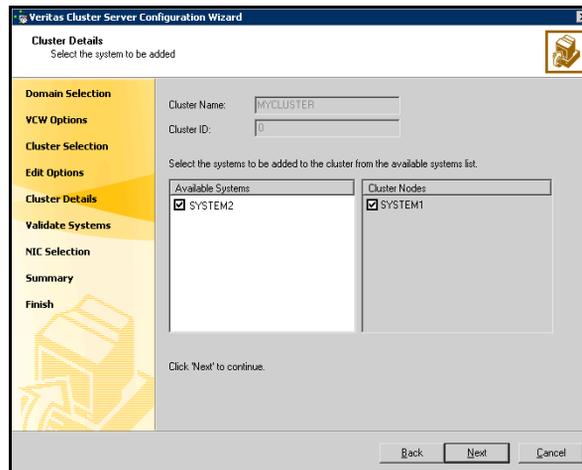
- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 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.
- 10 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, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.
- 11 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.

- 12 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.
- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**.
How you configure the VCS private network communication depends on

how it is configured in the cluster. If LLT is configured over ethernet, you have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - 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 both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.

The wizard will configure the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Check the **Configure LLT over UDP** check box.
 - Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
 - Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard will configure the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 14 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.

- 15 Specify the password for the user in whose context the VCS Helper service runs.
- 16 Review the summary information and click **Add**.
- 17 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

If you do not need to configure the VCS Web Console and notification, return to the task list in “[Configuring the storage hardware and network](#)” on page 300.

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.

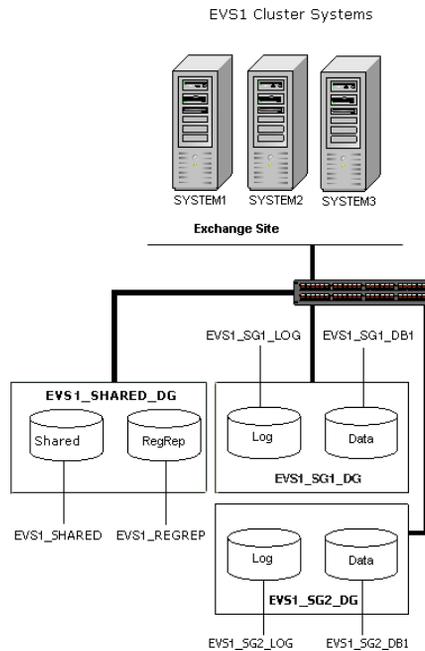
Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator’s Guide* for more information.

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

- The type of volume configurations that are required.
- The number of volumes or 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 the following 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 the following 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: If you are planning a disaster recovery configuration using Veritas Volume Replicator (VVR), you will need to allow additional disk space for a Storage Replicator Log volume (EVS1_REPLOG). This volume is created automatically when you run the wizard to set up replication. For more about VVR planning, see the *Veritas Volume Replicator, Administrator's Guide*.

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

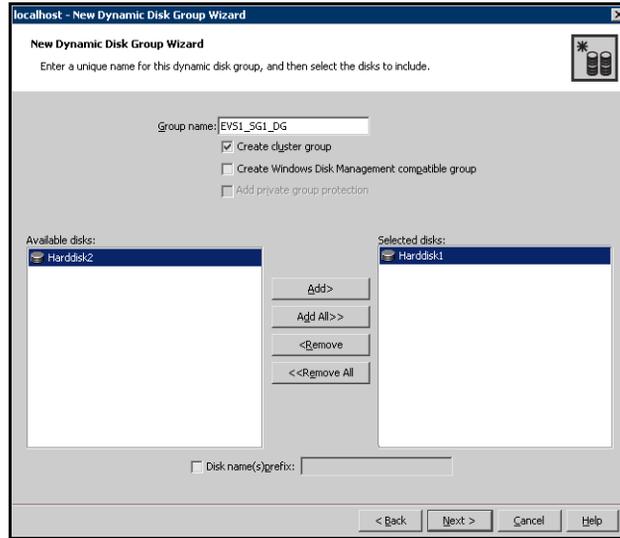
Steps to create a disk group are as follows.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

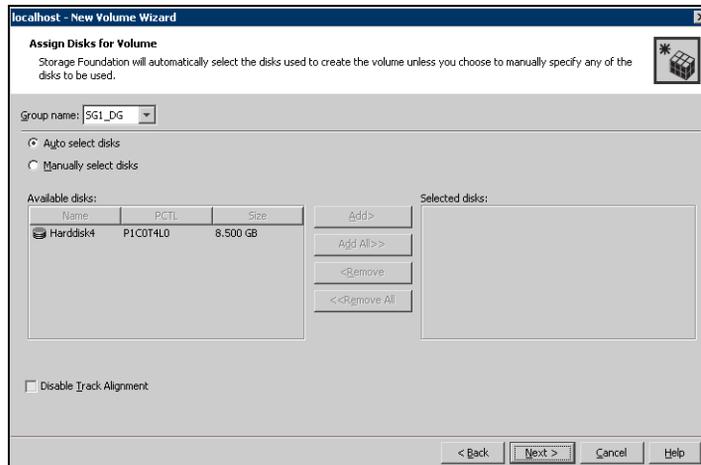
- Click **Next**.
- 7** Click **Next** to accept the confirmation screen with the selected disks.
 - 8** Click **Finish** to create the new disk group.

Creating volumes

This procedure assumes you are starting with the EVS1_SG1_DB1 volume.

To create dynamic volumes

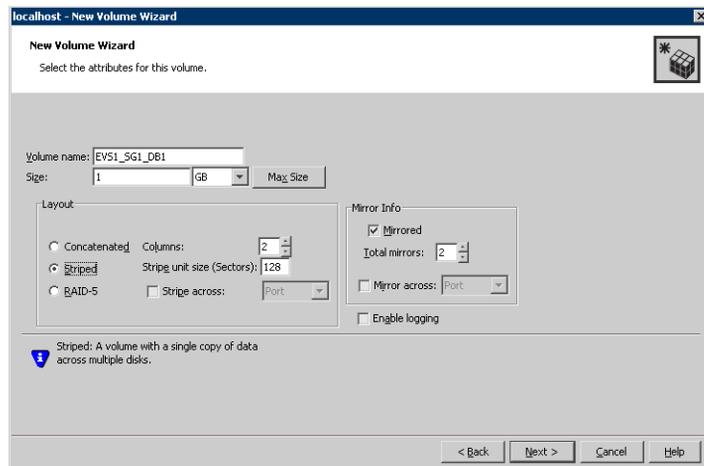
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu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

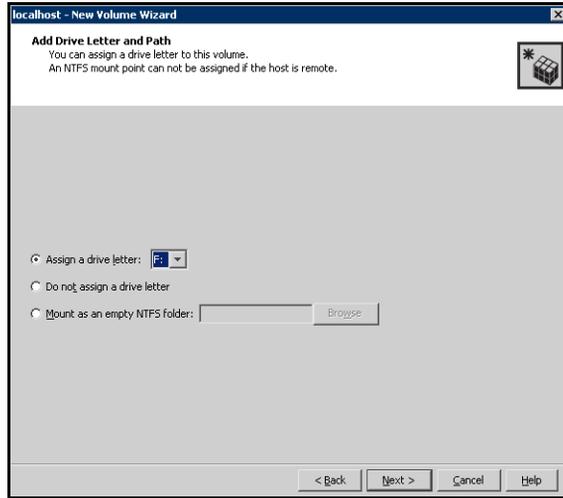
- 8 Click **Next**.
- 9 Specify the volume attributes.



- 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

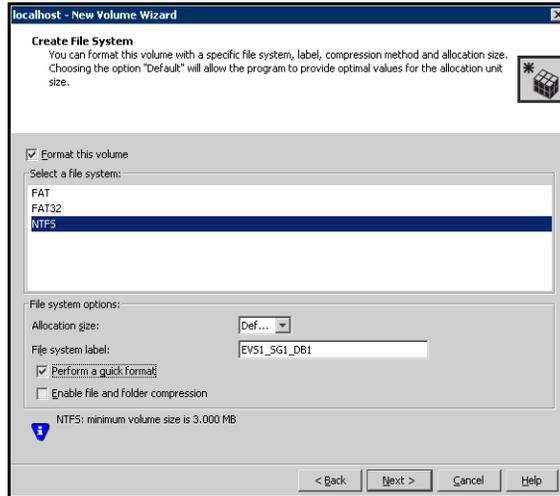
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

- 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 any additional volumes required. 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.
See [“Importing a disk group and mounting a shared volume”](#) on page 322 for instructions.
- Mount the volume containing the information for registry replication (EVS1_SG1_REGREP).
See [“Importing a disk group and mounting a shared volume”](#) on page 322 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - 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.
 - Either the logged-on user or the VCS Helper service user account must have write permissions on the DNS server to perform DNS updates.
 - Make sure the VCS Helper service 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

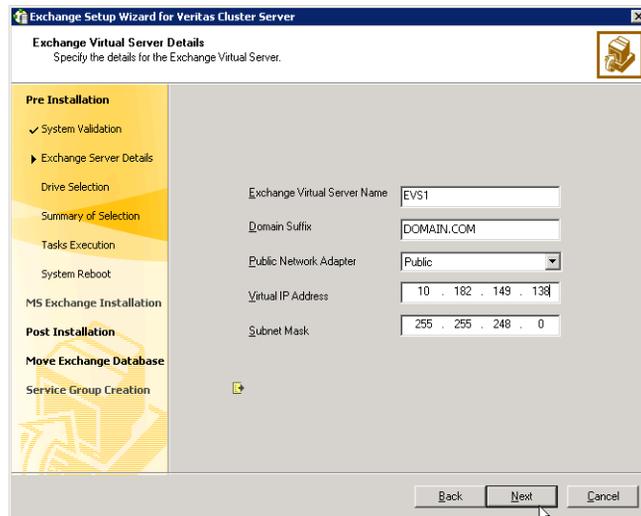
Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. You must install Exchange on a virtual node to facilitate high availability.

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 Tools > 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.

Warning: 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 unique on the network.
- 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.

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

If you wish to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: First node

Install Exchange on the node on which you performed the pre-installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```


The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```



```
C:\>net start had
```
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.

- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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.
Once the node is rebooted, move the databases created during the Exchange installation, from the local drive to the shared storage.

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.

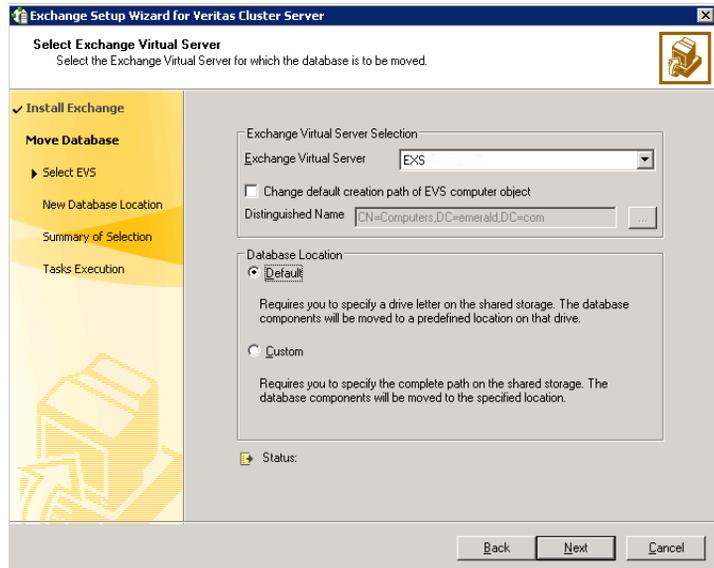
Complete the following tasks before moving the databases:

- 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 323.
- 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 314.
- 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 322.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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, click **Configure/Remove highly available Exchange Server** and then click **Next**.

- 4 In the Select Option dialog box, click **Move Exchange Databases** and then click **Next**.
- 5 In the Select Exchange Virtual Server dialog box, choose the Exchange virtual server and the database location option and then click **Next**.



Exchange Virtual Server

From the drop-down list, select the Exchange virtual server for which you want to move the databases.

Change default creation path of EVS computer object

Perform the following steps if you wish to change the default path for the Exchange virtual server object in Windows Active Directory:

- Check the **Change default creation path of EVS computer object** check box.
- Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**.
To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
The Lanman agent performs Windows AD updates. These settings are applicable to the Lanman resource in the service group.
By default, the Lanman resource adds the virtual server to the default container "Computers."

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

Default

Click **Default** if you wish to move the databases to predefined location on the shared storage. In the next step the wizard prompts you to specify the drive letter on the shared storage. The first mailbox store, public store, and MTA data are then moved to the generated default paths on the volumes that you specify.

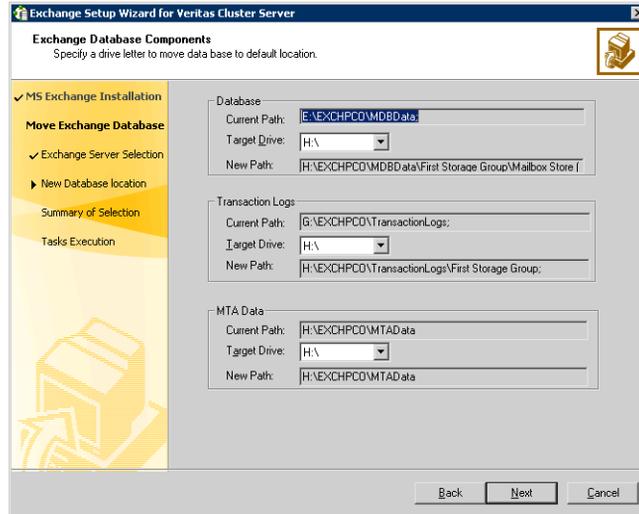
Caution: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

Custom

Click **Custom** if you wish to move the databases to a specific location on the shared storage. Choosing a custom location allows you to specify the Exchange database and streaming path. In the next step the wizard prompts you to specify the entire path of the location on the shared storage. The wizard then moves the databases to the specified directory.

If you chose the Default option, proceed to the next step. If you chose the Custom option, proceed to [step 7](#) on page 332.

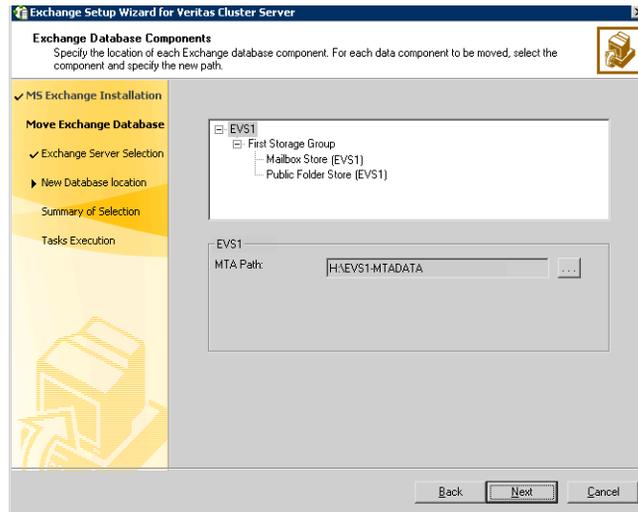
- 6 For the option of a default database location, specify the drives for moving the Exchange database components. The database components are then moved to a predefined location on that drive.



On the Exchange Database Components panel, complete the following steps:

- Specify a drive for moving the Exchange database.
- Specify a drive for moving the Exchange Transaction Logs.
- Specify a drive for moving the Exchange MTA Data.
- Click **Next** and proceed to [step 8](#) on page 332.

- 7 For the option of a custom database location, specify the location for specific Microsoft Exchange data components and then click **Next**.



For each data component that you wish to move, select the component and then click the ellipsis (...) to browse for the folder where you want to move it.

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

- 8 Review the summary of your selections and then click **Next**.
The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task.
- 9 After all the tasks are completed successfully, click **Next**.
- 10 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 338, 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 323.

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.

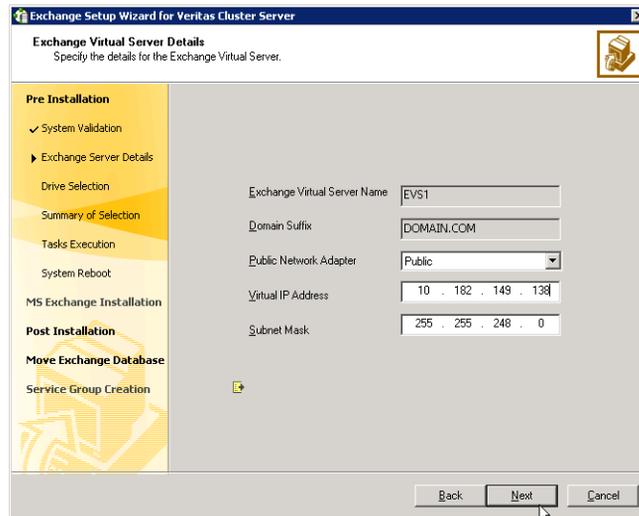
See “[Unmounting a volume and deporting a disk group](#)” on page 322.

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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: Additional nodes

Install Exchange on the additional node on which you performed the 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.

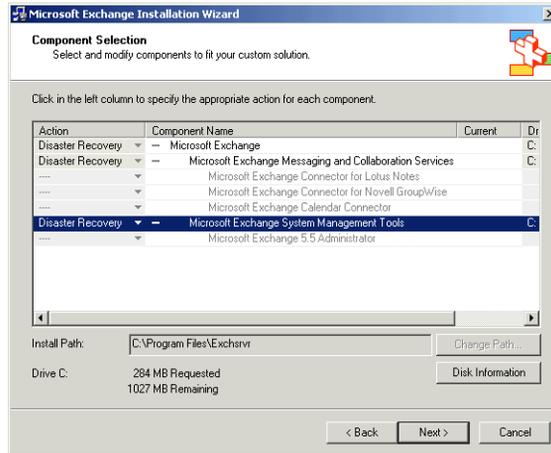
To install Exchange

- 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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:
`C:\>hasys -state`
The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`

- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

Specifying a common node for failover

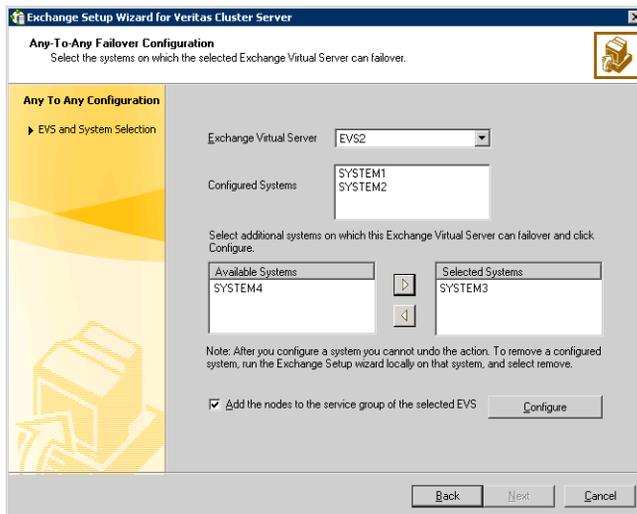
Specifying a common node for failover involves preparing the cluster with the Exchange Setup Wizard for VCS, and configuring the Exchange service group.

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 Tools > 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 332, 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. This privilege is required to configure service groups.
- 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.

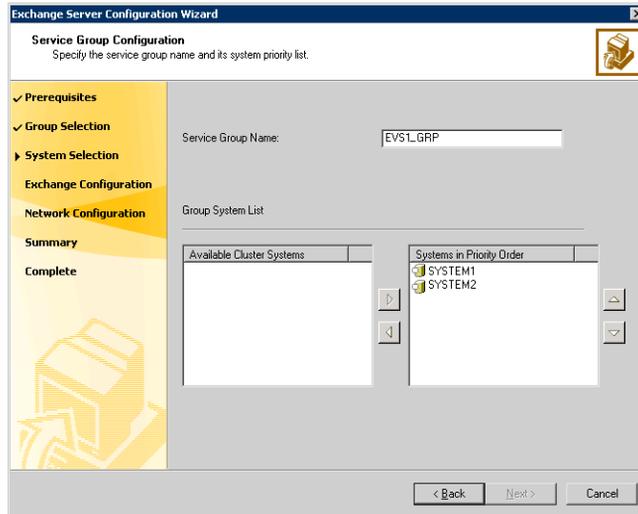
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk group and mount the shared volumes created to store 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 322. For instructions on unmounting, see “[Unmounting a volume and deporting a disk group](#)” on page 322.
- 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.

See *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on the VCS Exchange agent 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

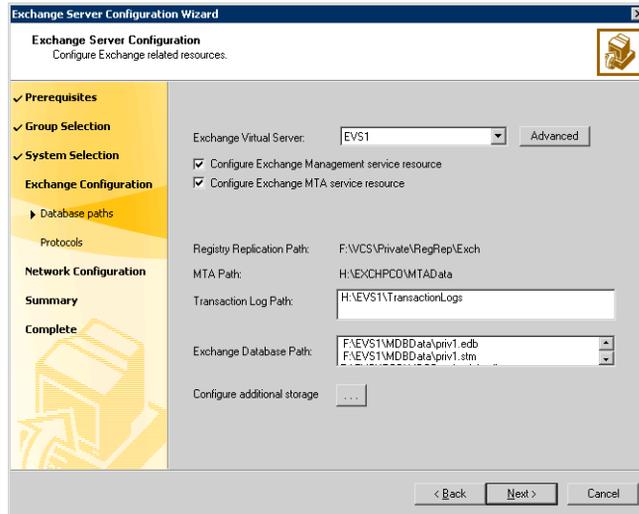
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**:

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



- Enter a name for the Exchange service group.
 If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.

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



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
- Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.

On the Lanman Advanced Configuration dialog box, complete the following:

- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
- Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.

If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

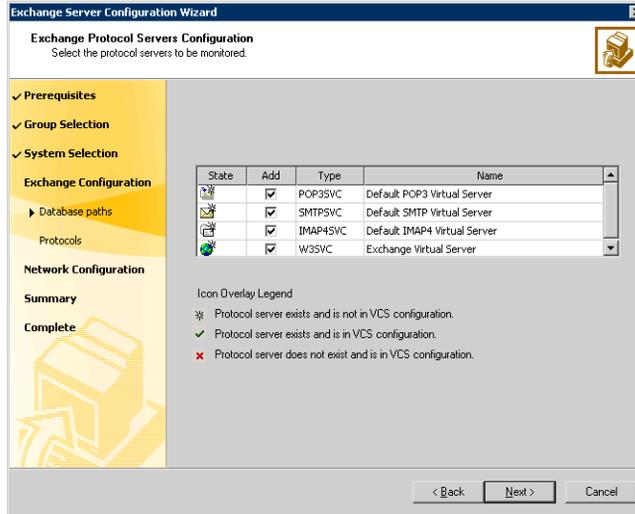
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.

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.

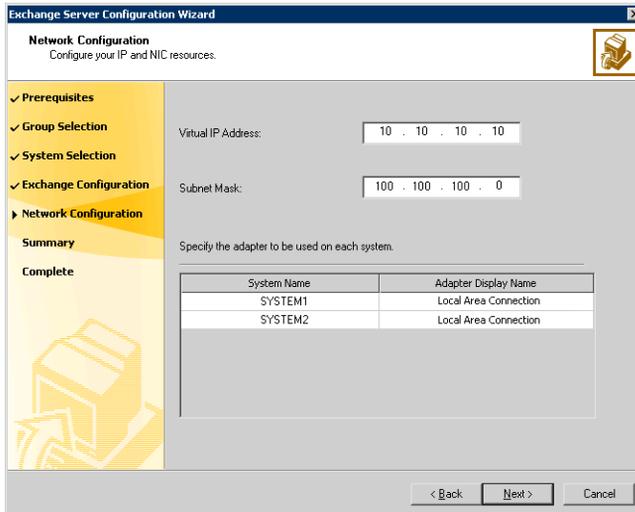
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:
 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- 6 On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.

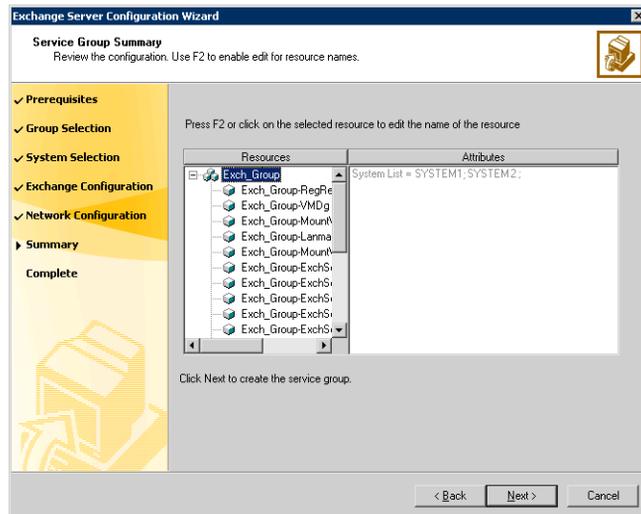


- 7 On the Network Configuration panel, specify information related to the network and then click **Next**:



- 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 Display Name** field to view the adapters associated with a node.
 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.

8 Review the service group configuration, change the resource names, if desired, and then click **Next**:



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 the **Esc** key.

9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status

of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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

This section contains the following chapters:

- [Chapter 8, “Campus cluster for Exchange: Overview”](#) on page 351
- [Chapter 9, “Deploying SFW HA for Campus Cluster: New Installation”](#) on page 355

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 chapters in this section to install and configure Exchange in an SFW HA campus cluster.

Campus cluster for Exchange: Overview

This chapter contains the following topics:

- [“What is a campus cluster?”](#) on page 352
- [“Why implement a campus cluster?”](#) on page 352
- [“What is high availability?”](#) on page 352
- [“Why implement a high availability solution?”](#) on page 353
- [“How the VCS application agent makes Microsoft Exchange highly available”](#) on page 353

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 355](#), 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 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.

How the VCS application agent makes Microsoft Exchange highly available

Deploying SFW HA for Campus Cluster: New Installation

This chapter contains the following topics:

- [“Reviewing the requirements”](#) on page 357
- [“Reviewing the configuration”](#) on page 363
- [“Configuring the network and storage”](#) on page 367
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 369
- [“Configuring the cluster”](#) on page 375
- [“Managing disk groups and volumes”](#) on page 401
- [“Preparing the forest and domain”](#) on page 402
- [“Installing Exchange on the first node”](#) on page 402
- [“Moving Exchange databases to shared storage”](#) on page 407
- [“Installing Exchange on additional nodes”](#) on page 412
- [“Configuring the Exchange service group for VCS”](#) on page 417
- [“Modifying the IP resource in the Exchange service group”](#) on page 424
- [“Verifying the campus cluster: switching the service group”](#) on page 426
- [“Possible tasks after creating the campus cluster”](#) on page 427

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 SFW HA DR solution with replication.

Table 9-1 on page 356 outlines the high-level objectives and the tasks to complete each objective.

Table 9-1 Task list: Exchange campus cluster configuration

Objective	Tasks
“ Reviewing the requirements ” on page 357	Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 363	Understanding a typical active-passive Exchange configuration in a two-node campus cluster
“ Configuring the network and storage ” on page 367	<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 369	<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 375	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard (VCW)
“ Configuring disk groups and volumes ” on page 392	<ul style="list-style-type: none"> ■ Creating disk groups ■ Creating the data, log, RegRep, and MTA volumes
“ Preparing the forest and domain ” on page 402	<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 402	<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 407	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: Exchange campus cluster configuration (Continued)

Objective	Tasks
“Installing Exchange on additional nodes” on page 412	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 417	Creating the Exchange service group using the VCS Exchange Configuration Wizard.
“Modifying the IP resource in the Exchange service group” on page 424	Modifying the Address and SubNetMask attributes if the sites are in different subnets.
“Possible tasks after creating the campus cluster” on page 427	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 Symantec recommended requirements may vary.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table 9-2](#) estimates disk space requirements for SFW HA.

Table 9-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

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

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/302144>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agent for Microsoft Exchange.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2003 versions

The following table lists the Microsoft Exchange Server 2003 versions supported with SFW HA 5.1 Service Pack 1.

Table 9-4 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

Table 9-4 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> <li data-bbox="796 317 1325 404">■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Memory must be a minimum 256 MB of RAM per server for Exchange 2003; refer to the Microsoft documentation for more information.
- 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. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
See "[Best practices for SFW HA](#)" on page 362.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each Exchange Virtual Server (EVS).

- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone 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 Dynamic DNS (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 for SFW HA

The following permissions are required:

- 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 on 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 Exchange virtual server computer object 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, 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 for SFW HA

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 Server 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 2003.
- 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.

Best practices for SFW HA

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.

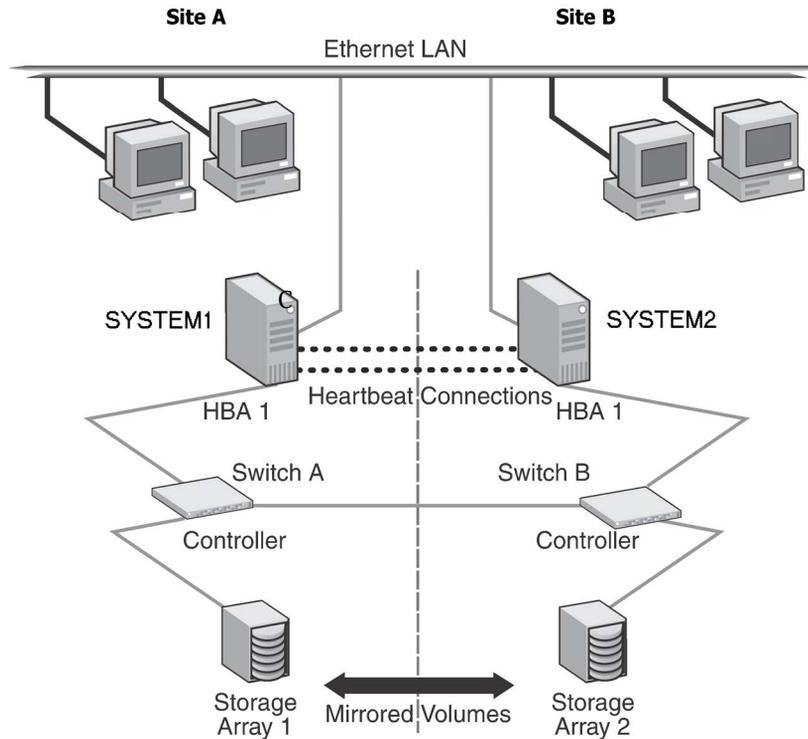
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

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.

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 9-6 Failure Situations

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
<p>1) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.</p>	<p>No interruption of service. Remaining disks in mirror are still accessible from the other node.</p>	<p>The Service Group does not failover. 50% of the mirrored disk is still available at remaining site.</p>
<p>2) Zone failure Complete Site failure, all accessibility to the servers and storage is lost.</p>	<p>Manual intervention required to online the Service Group at remaining site. Can not automatically import 50% of mirrored disk.</p>	<p>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.</p>
<p>3) Split-brain (loss of both heartbeats) 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>4) Storage interconnect lost 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>5) Split-brain and storage interconnect lost 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. 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>

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 and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network 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 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 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, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

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

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 9-7 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 370.

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 Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The Select Product screen appears.
- 3 Review the links on the Select Product screen.
Links on this screen access Late Breaking News, the Configuration Checker, as well as begin the process to install Storage Foundation HA for Windows. Click on **Read Late Breaking News** for the latest information about updates, patches, and software issues regarding this release.
- 4 Click **Storage Foundation HA 5.1 SP1 for Windows**.
- 5 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 6 Review the Welcome message and the listed prerequisites. Ensure that any of the listed prerequisites applicable to your installation are met prior to proceeding.
Click **Next**.
- 7 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 AGREE TO the terms of the license agreement**, and then click **Next**.
- 8 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.
- 9 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.
- 10 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. Be sure to select the following as appropriate for your installation.

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, select the option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

11 Select the following 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. When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

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

- 12 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.
- 13 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.
- 14 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages.

If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning:

The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If applicable to your installation, perform the above procedure.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

If applicable to your installation, perform the above procedure.

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation 5.1 for Windows (Server Components) with the MSCS option, you may want MSCS Quorum Arbitration Time (Min. and Max) to be adjusted to ensure optimal functionality with Veritas dynamic volumes with MSCS.

For additional information, see the *Storage Foundation for Windows Administrator Guide* for details.

If applicable to your installation, perform the above procedure.

- 15 When finished reviewing the message or messages, click **OK**.
- 16 The Summary screen appears displaying an Install report. Review the information in the Install report. Click **Back** to make changes, if necessary. Click **Install** if information is validated.
- 17 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.
- 18 When the installation completes, review the summary screen and click **Next**.
- 19 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.
- 20 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 21 Review the log files and click **Finish**.
- 22 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 (VCW) 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.

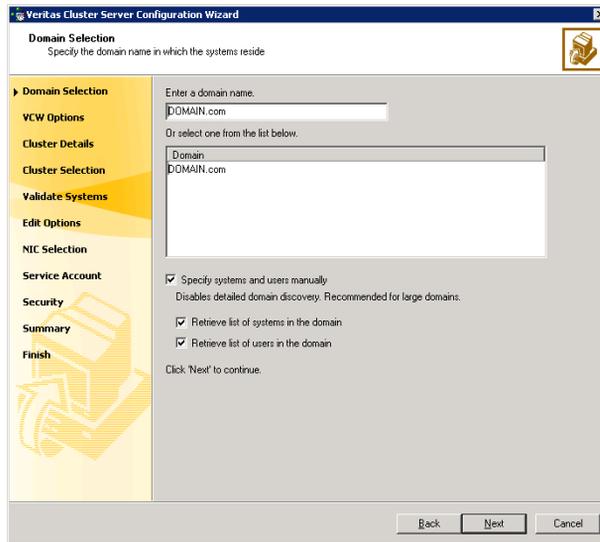
Complete the following tasks before creating a cluster:

- 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 HAD 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 HAD Helper service, you must know the password for the user account.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. 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. This account does not require domain admin privileges.
 - Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

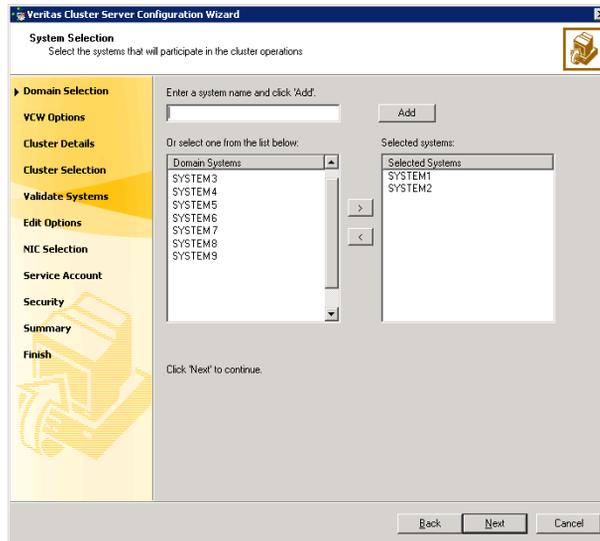


Do one of the following:

- 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 8](#) on page 378.
 - 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 377. 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 8](#) on page 378.

- 6 On the System Selection panel, specify the systems for the 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.

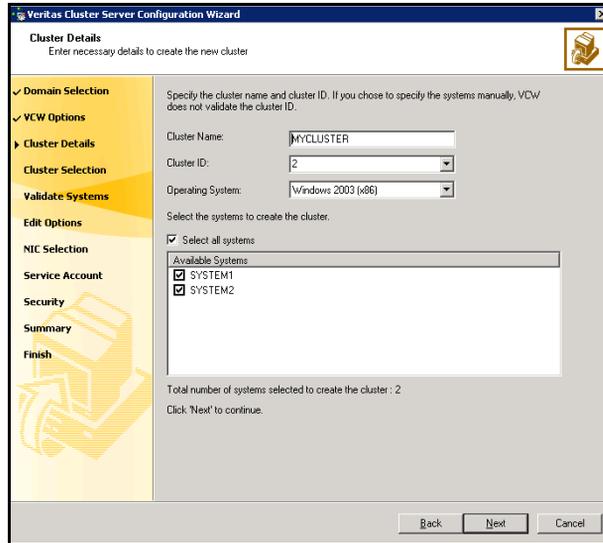
- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 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. The cluster ID can be any number from 0 to 255.

Caution: 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. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (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.

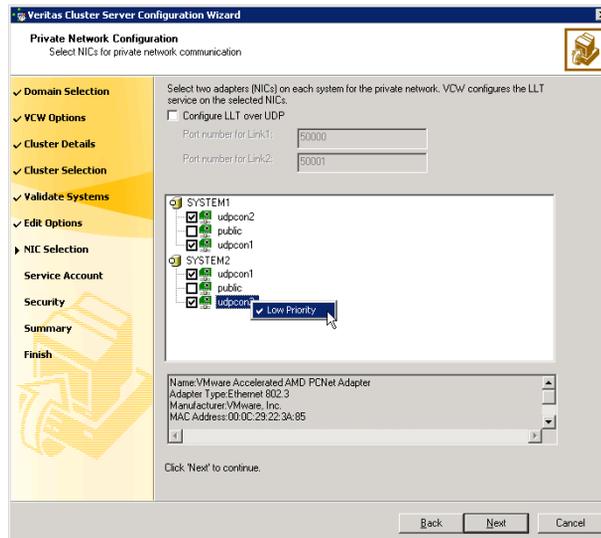
10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

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 the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 381.

11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



- 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 both public and private communication.

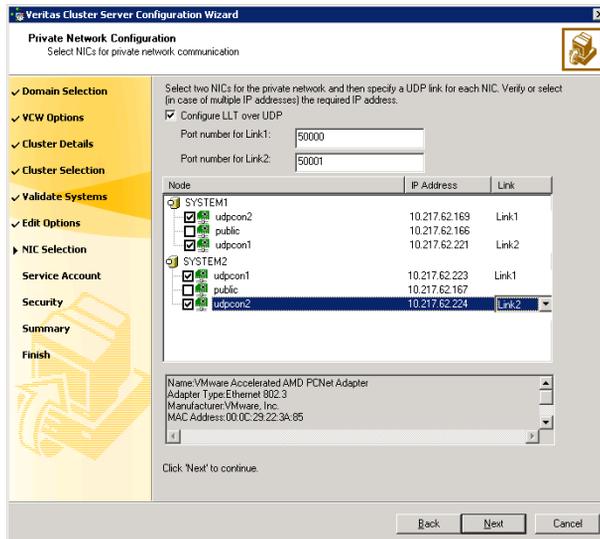
- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.

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.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to

65535. The default ports numbers are 50000 and 50001 respectively.

- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

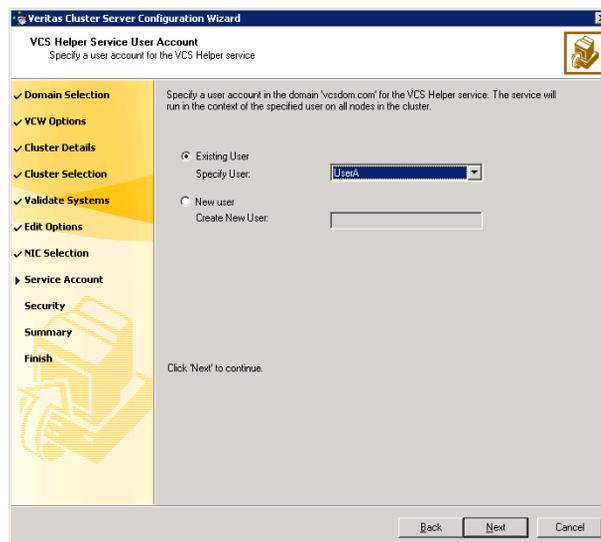
The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

This account does not require Domain Administrator privileges.



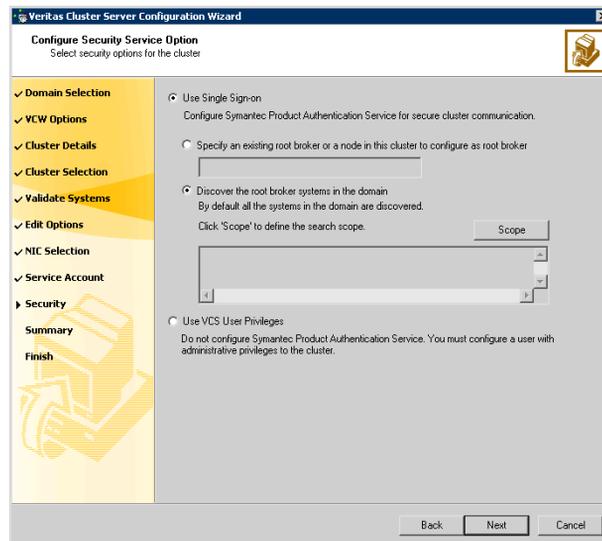
Specify a domain user as follows:

- 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 376, 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**.

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

Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability 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 an existing root broker or a node in this cluster to configure as root broker**, 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 then configures all nodes in the cluster as authentication brokers.

- If you want to search 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 a user *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**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 9-8 contains some more examples of search criteria.

Table 9-8 Search criteria examples

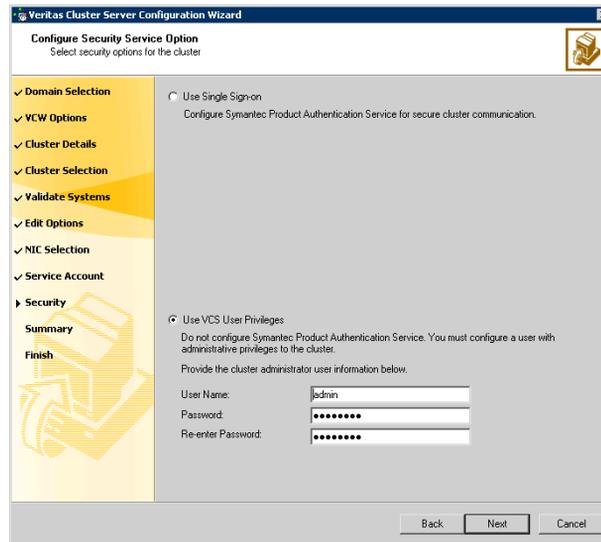
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- 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 a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS`Encrypt` utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password. After the cluster is configured, you can 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**.
- 14 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.

- 15 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**.

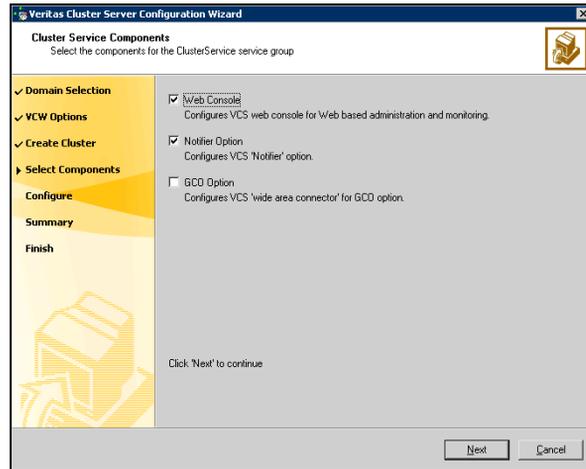
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

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.

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



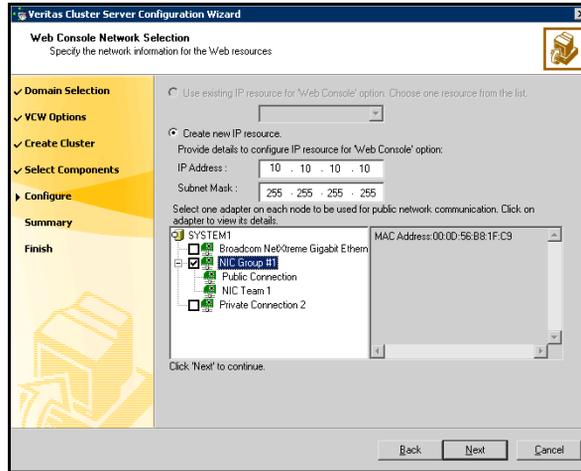
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 387.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 388.

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**.



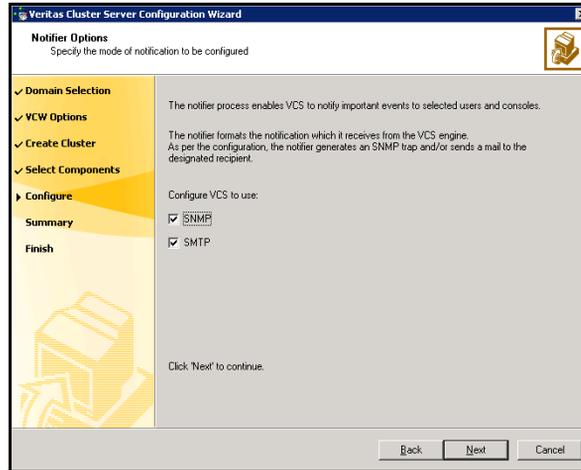
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 388. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

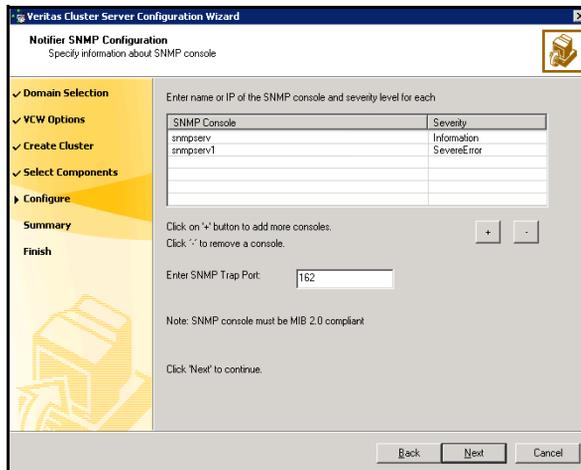
To configure notification

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

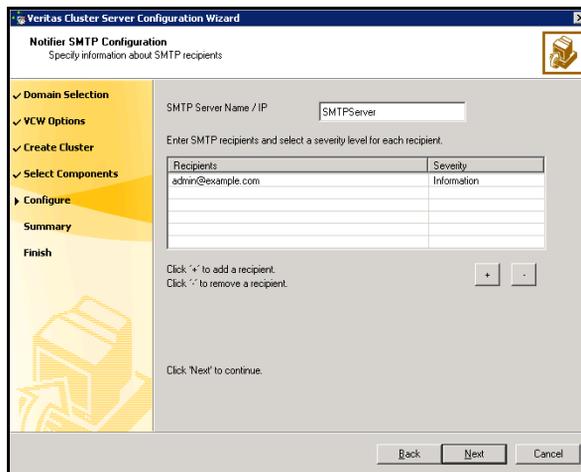


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

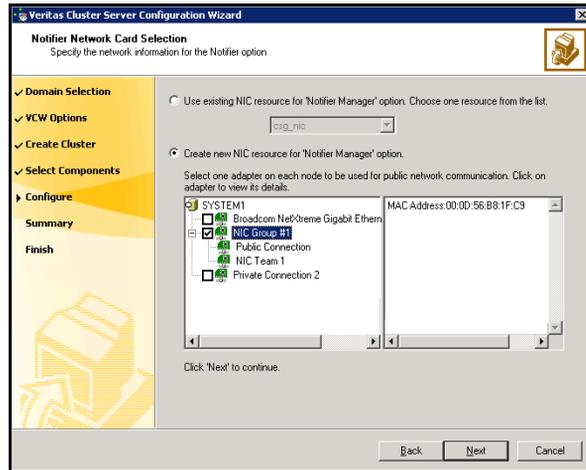


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - *Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

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

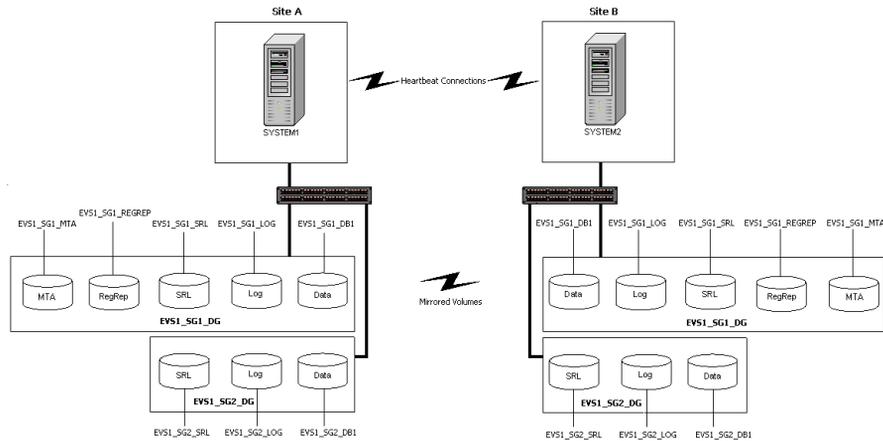
- The type of volume configurations that are required
- The number of volumes required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs that depend on the traffic load
- The disk groups and number of disks on each site
- Types 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. The MTA volumes in the illustration are applicable in case of Exchange Server 2003 only.

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 395
- [“Creating a volume”](#) on page 397

Considerations when creating new volumes

Use the Veritas Enterprise Administrator (VEA) console to view the available disk storage on a system.

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 select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

Creating a dynamic (cluster) disk group

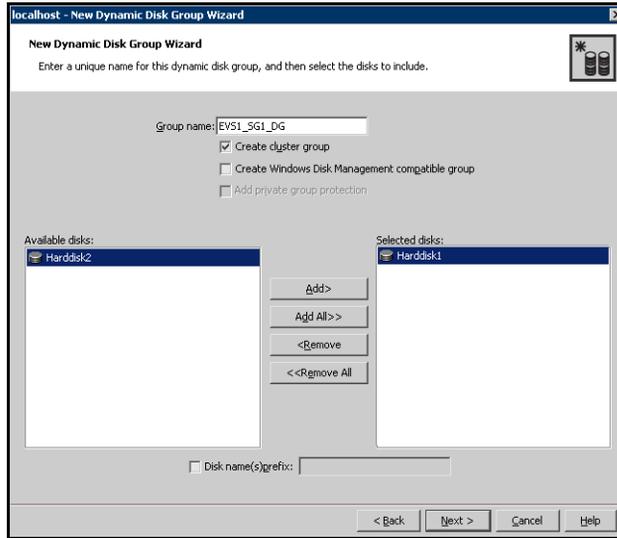
Use the following procedure to create a dynamic cluster disk group.

To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

- Click **Next**.

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the new disk group.

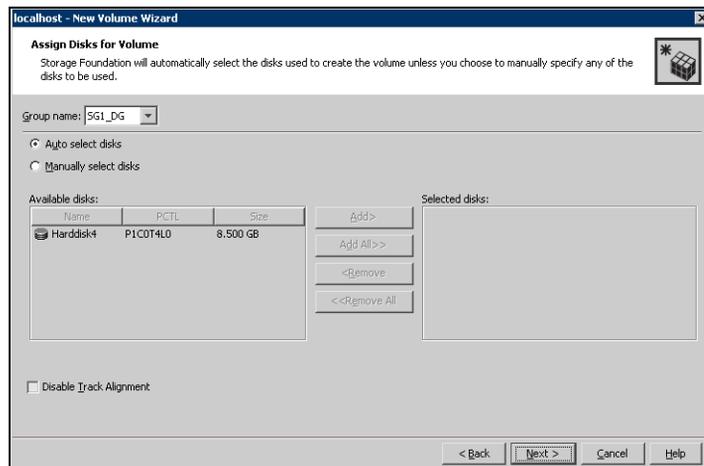
Proceed to create the appropriate volumes on each disk.

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.

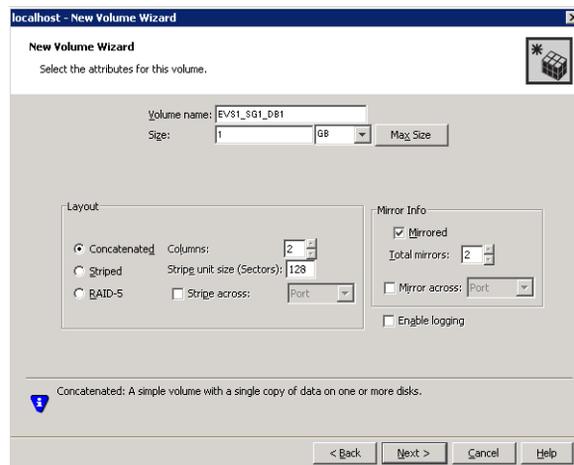


- 7 Select auto or manual disk selection and enable or disable track alignment.

- 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.

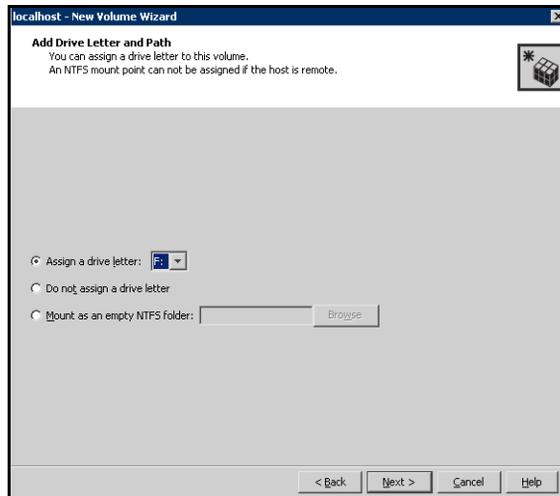
8 Click **Next**.

9 Specify the volume attributes.



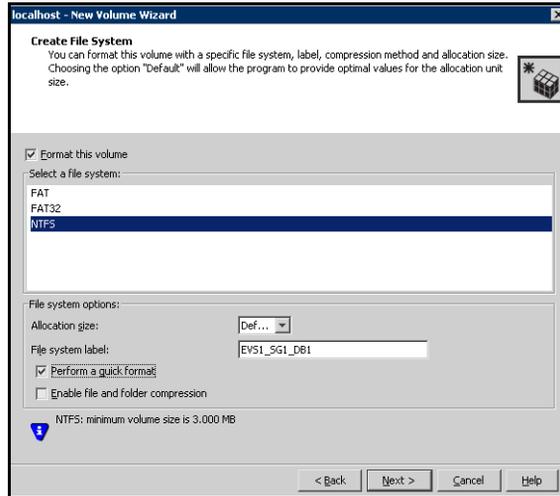
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume.
- If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that 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.
- 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 log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) For Exchange 2003, you could also create an MTA volume (EVS1_SG1_MTA).

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.

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.
See “[Preparing the forest and domain](#)” on page 402 for instructions.
- 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 401 for instructions.
- Mount the volume containing the information for registry replication (EVS1_SG1_REGREP).
See “[Unmounting a volume and deporting a disk group](#)” on page 401 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - Either the logged-on user or the VCS Helper service user account must have write permissions on the DNS server to perform DNS updates.
 - 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.
 - 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

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 Tools > 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.

The screenshot shows the 'Exchange Setup Wizard for Veritas Cluster Server' dialog box. 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, a navigation pane shows steps: 'Pre Installation' (with sub-steps: System Validation, Exchange Server Details, Drive Selection, Summary of Selection, Tasks Execution, System Reboot, MS Exchange Installation), 'Post Installation' (with sub-steps: Move Exchange Database, Service Group Creation). The 'Exchange Server Details' section is active, showing input fields for: 'Exchange Virtual Server Name' (EVS1), 'Domain Suffix' (DOMAIN.COM), 'Public Network Adapter' (Public), 'Virtual IP Address' (10 . 182 . 149 . 134), and 'Subnet Mask' (255 . 255 . 248 . 0). At the bottom, there are 'Back', 'Next', and 'Cancel' buttons.

- Enter a unique virtual name for the Exchange server.

Warning: 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 unique on the network.
- 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.

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

If you wish to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: First node

Install Exchange on the same node selected in “[Exchange pre-installation: First node](#)” on page 404.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```

The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```

```
C:\>net start had
```
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.

- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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.
Once the node is rebooted, move the databases created during the Exchange installation, from the local drive to the shared storage.

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.

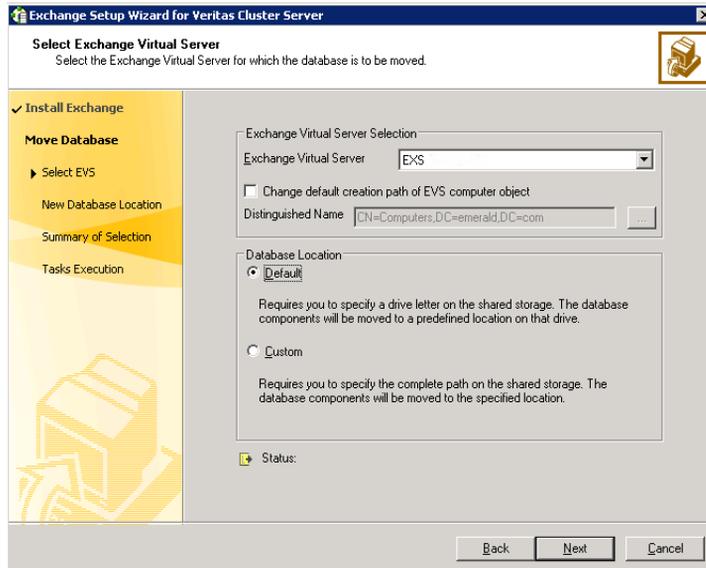
Complete the following tasks before moving the databases:

- 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 401 for instructions.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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, click **Configure/Remove highly available Exchange Server** and then click **Next**.
- 4 In the Select Option dialog box, click **Move Exchange Databases** and then click **Next**.

- 5 In the Select Exchange Virtual Server dialog box, choose the Exchange virtual server and the database location option and then click **Next**.



Exchange Virtual
Server

From the drop-down list, select the Exchange virtual server for which you want to move the databases.

Change default creation path of EVS computer object

Perform the following steps if you wish to change the default path for the Exchange virtual server object in Windows Active Directory:

- Check the **Change default creation path of EVS computer object** check box.
- Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**.
 To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
 The Lanman agent performs Windows AD updates. These settings are applicable to the Lanman resource in the service group.
 By default, the Lanman resource adds the virtual server to the default container "Computers."

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

Default

Click **Default** if you wish to move the databases to predefined location on the shared storage. In the next step the wizard prompts you to specify the drive letter on the shared storage. The first mailbox store, public store, and MTA data are then moved to the generated default paths on the volumes that you specify.

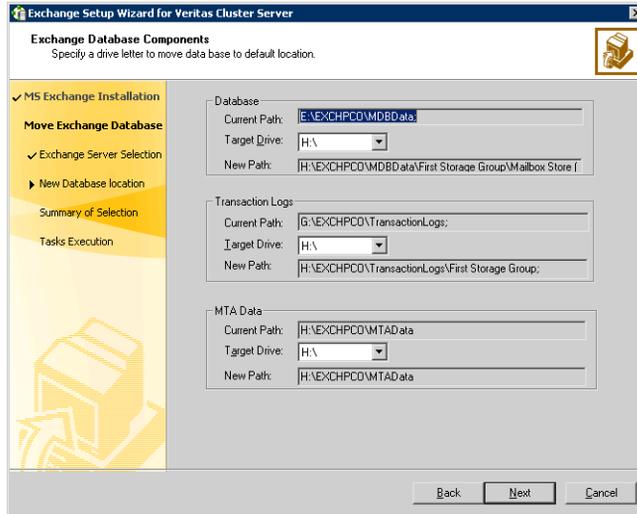
Caution: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

Custom

Click **Custom** if you wish to move the databases to a specific location on the shared storage. Choosing a custom location allows you to specify the Exchange database and streaming path. In the next step the wizard prompts you to specify the entire path of the location on the shared storage. The wizard then moves the databases to the specified directory.

If you chose the Default option, proceed to the next step. If you chose the Custom option, proceed to [step 7](#) on page 411.

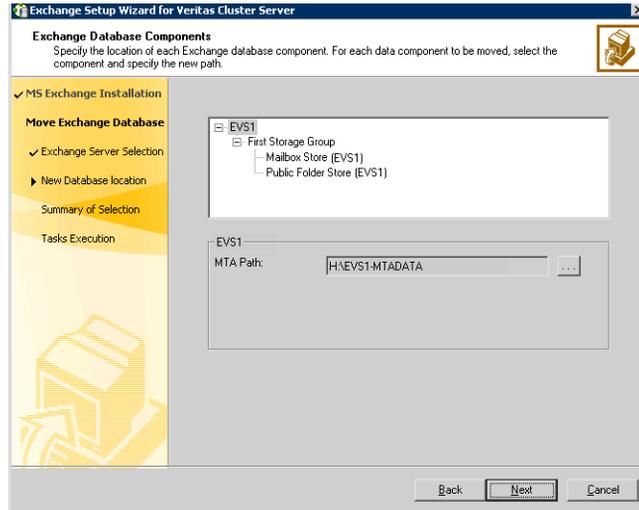
- For the option of a default database location, specify the drives for moving the Exchange database components. The database components are then moved to a predefined location on that drive.



On the Exchange Database Components panel, complete the following steps:

- Specify a drive for moving the Exchange database.
- Specify a drive for moving the Exchange Transaction Logs.
- Specify a drive for moving the Exchange MTA Data.
- Click **Next** and proceed to [step 8](#) on page 411.

- 7 For the option of a custom database location, specify the location for specific Microsoft Exchange data components and then click **Next**.



For each data component that you wish to move, select the component and then click the ellipsis (...) to browse for the folder where you want to move it.

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

- 8 Review the summary of your selections and then click **Next**.
The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task.
- 9 After all the tasks are completed successfully, click **Next**.
- 10 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 402.

Exchange pre-installation: Additional nodes

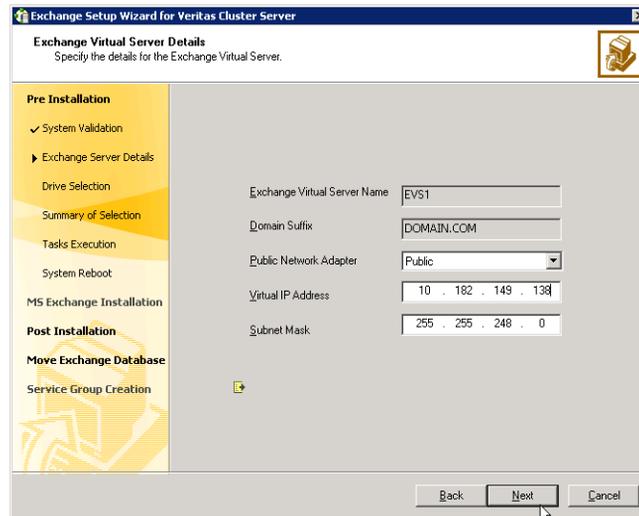
Note: Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes.

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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: Additional nodes

Install Exchange on the same node on which you performed the 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.

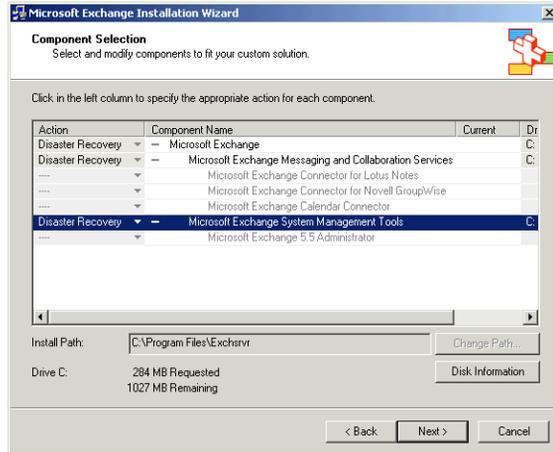
To install Exchange

- 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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
 Type the following on the command line:
`C:\>hasys -state`
 The state should display as RUNNING.
 If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`

- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

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.

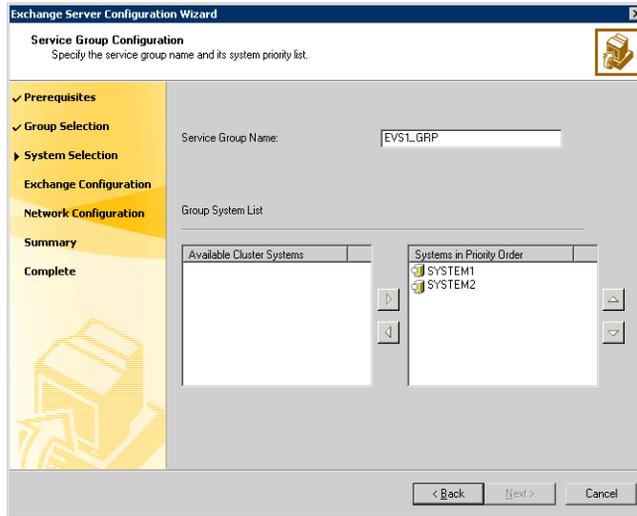
Prerequisites

- You must be a Cluster Administrator. This privilege is required to configure service groups.
- 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
- 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

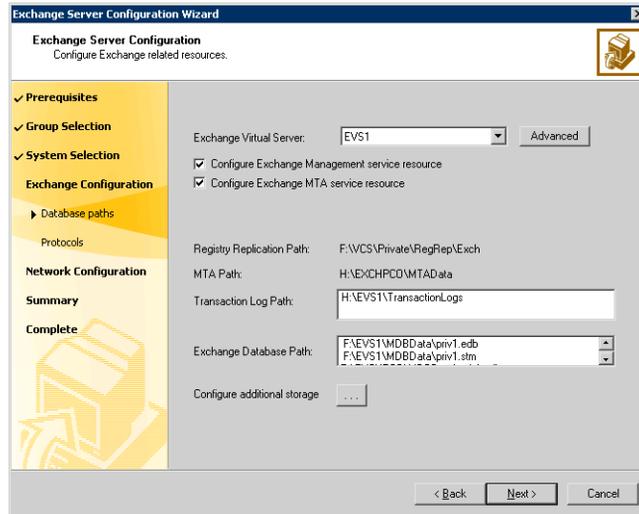
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**:

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



- Enter a name for the Exchange service group.
If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.

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



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
 - Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.
- On the Lanman Advanced Configuration dialog box, complete the following:
- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
 - Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.

If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

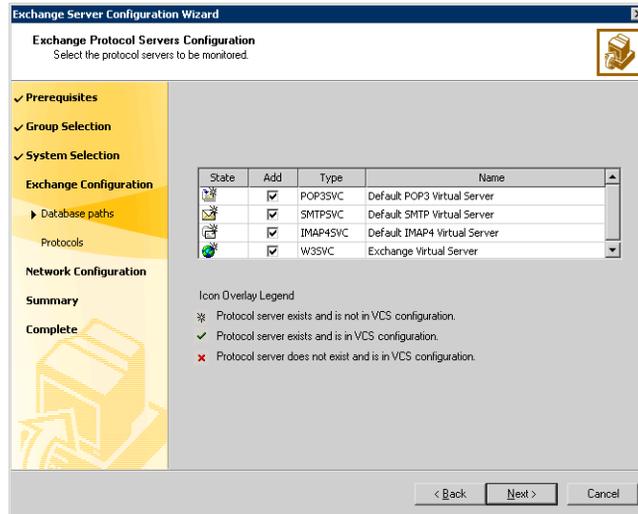
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.

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.

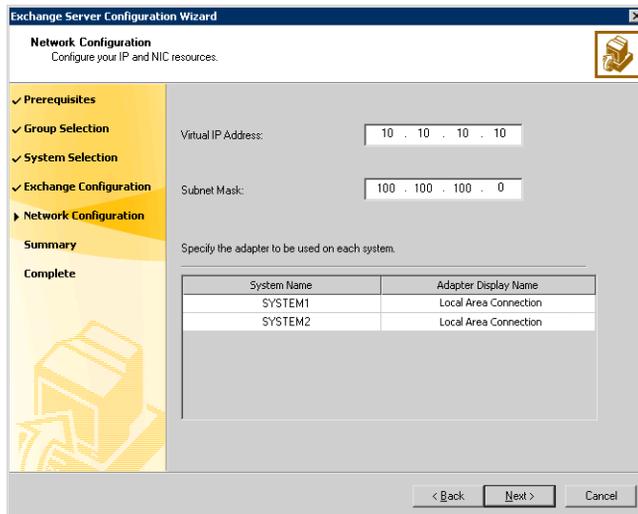
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:
 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.

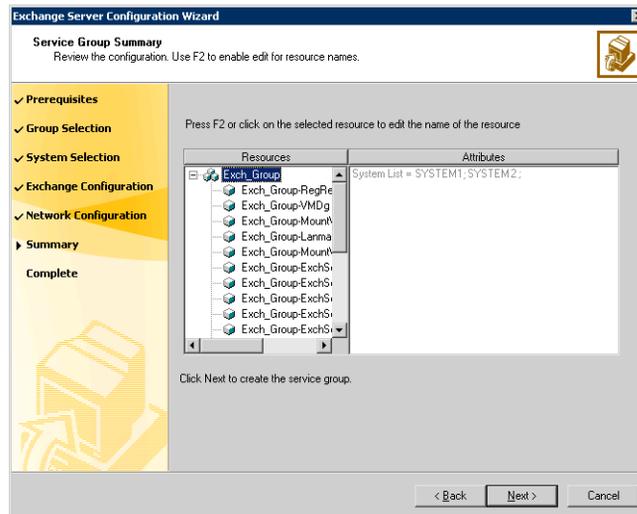


- On the Network Configuration panel, specify information related to the network and then click **Next**:



- 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 Display Name** field to view the adapters associated with a node.
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.

8 Review the service group configuration, change the resource names, if desired, and then click **Next**:



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 the **Esc** key.
- 9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status

of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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.

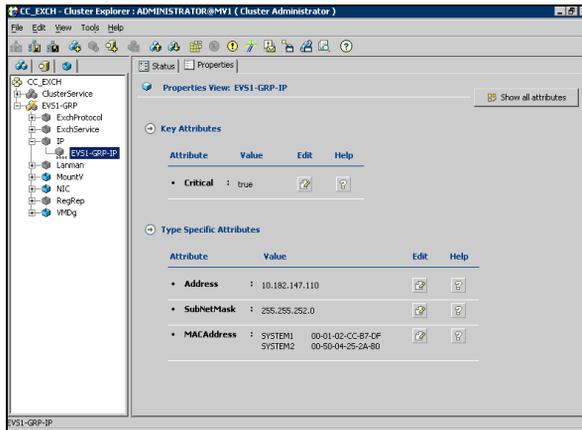
Modifying the IP resource in the Exchange service group

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.

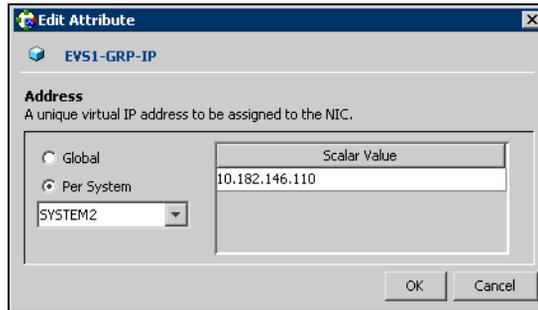
To modify the IP resource

- 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.

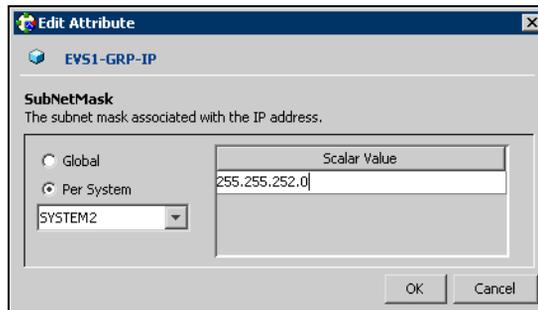
3 In the Edit Attribute dialog box:



- Select the **Per System** option.
- Select the system at Site B.
- Enter the virtual IP address at Site B.
- Click **OK**.

4 In the Properties View, click the **Edit** icon for the **SubNetMask** attribute.

5 In the Edit Attribute dialog box:



- Select the **Per System** option.
- Select the system at Site B.
- Enter the subnet mask at Site B.
- Click **OK**.

6 From the **File** menu of Cluster Explorer, click **Close Configuration**.

Verifying the campus cluster: switching the service group

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

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

To switch service groups

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

To shut down an active cluster node

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

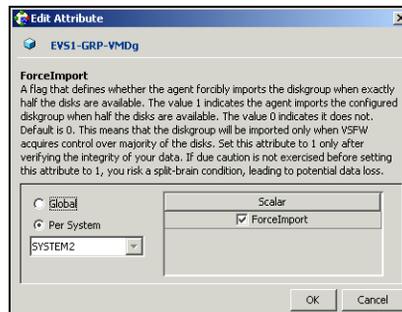
Possible tasks after creating the campus cluster

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 351](#), 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:



- 4 Select the **Per System** option.
- 5 Select the system in Site B.
- 6 Select the **ForceImport** check box.
- 7 Click **OK**.
- 8 From the **File** menu of Cluster Explorer, click **Close Configuration**.
- 9 After the failover takes place, revert the ForceImport attribute to its original value.

Replicated Data Clusters

This section contains the following chapters:

- [Chapter 10, “About Replicated Data Clusters”](#) on page 431
- [Chapter 11, “Configuring Replicated Data Clusters for Exchange”](#) on page 439

About Replicated Data Clusters

This chapter includes the following topics:

- [“About Replicated Data Clusters”](#) on page 431
- [“How VCS Replicated Data Clusters work”](#) on page 433
- [“Setting up a Replicated Data Cluster configuration”](#) on page 434
- [“Migrating the service group”](#) on page 436

About Replicated Data Clusters

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

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

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

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

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

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

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

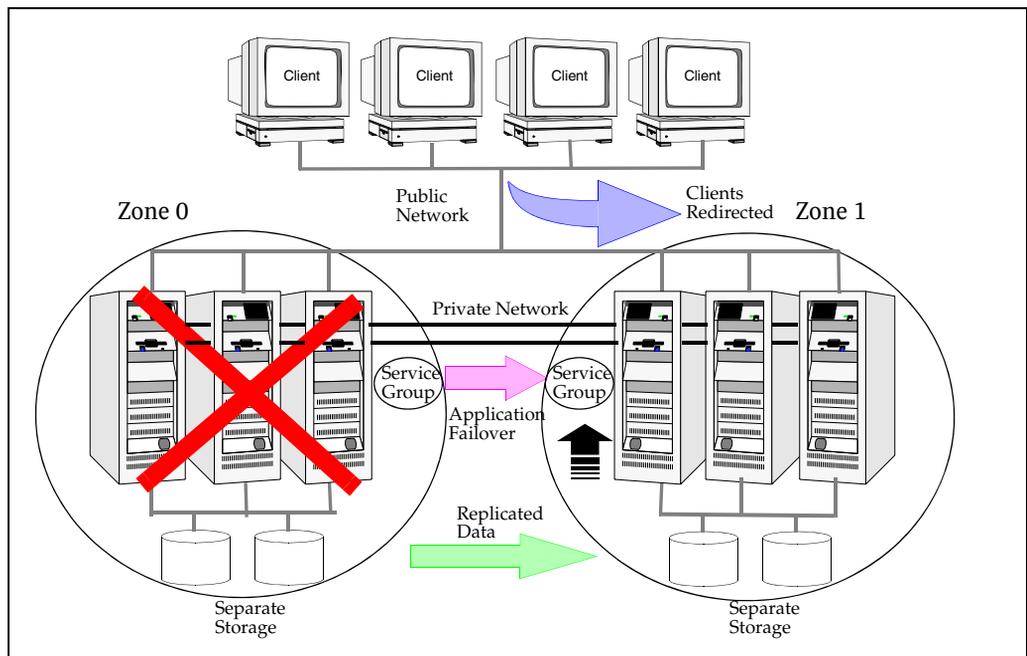
You must use dual dedicated LLT links between the replicated nodes.

How VCS Replicated Data Clusters work

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

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

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



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

Setting up a Replicated Data Cluster configuration

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

The process involves the following steps:

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

Setting up replication

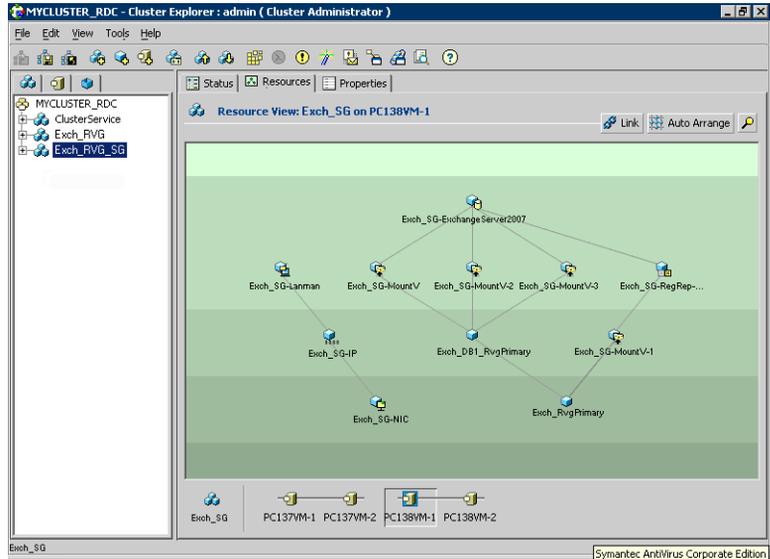
Set up replication between the shared disk groups. Use VVR to group the shared data volumes into a Replicated Volume Group and create the VVR Secondary on hosts in your secondary zone.

Create a Replicated Data Set (RDS) with the Primary RVG consisting of the shared volumes between the nodes in the first zone (zone 0) and Secondary RVG consisting of shared volumes between nodes in the second zone (zone 1). Use the same disk group and RVG names in both zones so that the MountV resources will mount the same block devices.

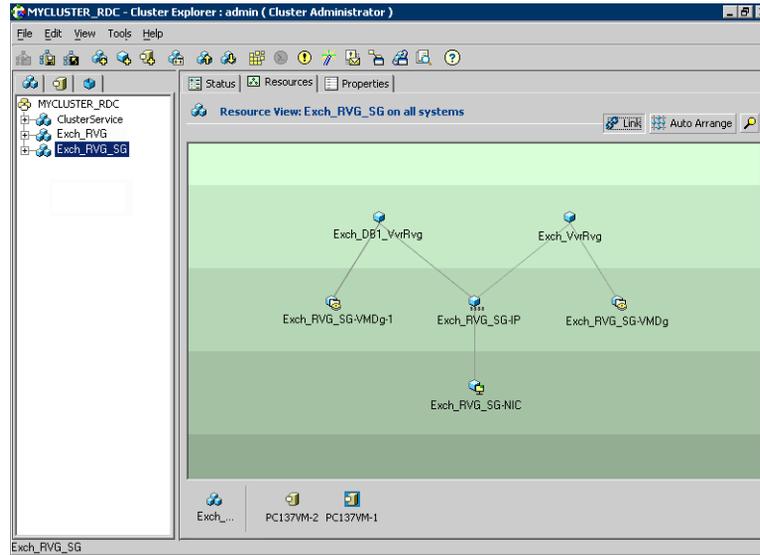
Configuring the service groups

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

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



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



Migrating the service group

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

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

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

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

To switch the service group

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

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

Configuring Replicated Data Clusters for Exchange

This chapter contains the following topics:

- [“Tasks for configuring Replicated Data Clusters for Exchange Server”](#) on page 440
- [“Reviewing the prerequisites”](#) on page 442
- [“Reviewing the configuration”](#) on page 449
- [“Configuring the storage hardware and network”](#) on page 450
- [“Preparing the forest and domain”](#) on page 451
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 452
- [“Configuring the cluster”](#) on page 459
- [“Configuring cluster disk groups and volumes”](#) on page 476
- [“Managing disk groups and volumes”](#) on page 484
- [“Installing Exchange on the first node”](#) on page 486
- [“Moving Exchange databases to shared storage”](#) on page 490
- [“Installing Exchange on additional nodes”](#) on page 495
- [“Configuring the Exchange service group for VCS”](#) on page 500
- [“Creating the primary system zone”](#) on page 507
- [“Verifying the installation in the primary zone”](#) on page 508
- [“Creating a parallel environment in the secondary zone”](#) on page 509

- [“Adding the systems in the secondary zone to the cluster”](#) on page 510
- [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 516
- [“Configuring a hybrid RVG service group for replication”](#) on page 528
- [“Setting a dependency between the service groups”](#) on page 538
- [“Adding the nodes from the secondary zone to the RDC”](#) on page 538
- [“Verifying the RDC configuration”](#) on page 544
- [“Additional instructions for GCO disaster recovery”](#) on page 545

Tasks for configuring Replicated Data Clusters for Exchange Server

Configure the high availability and Exchange Server components on the primary and secondary zones, then complete the Replicated Data Set solution by configuring the components for both zones.

For more information on VVR, see the *Veritas Volume Replicator Administrator’s Guide*.

[Table 11-1](#) outlines the high-level objectives and the tasks required to complete them.

Table 11-1 Tasks for configuring Replicated Data Clusters for Exchange

Objective	Tasks
“Reviewing the prerequisites” on page 442	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 449	<ul style="list-style-type: none">■ Understanding active/passive configuration and zone failover in a RDC environment■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 450	<ul style="list-style-type: none">■ Setting up the storage hardware for a cluster environment■ Verifying the DNS entries for the systems on which Exchange will be installed

Table 11-1 Tasks for configuring Replicated Data Clusters for Exchange

Objective	Tasks
“Installing Veritas Storage Foundation HA for Windows” on page 452	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation for Windows HA (automatic installation) ■ Selecting the option to install VVR; this also automatically installs the Veritas Cluster Server Agent for VVR ■ Selecting the option to install Veritas Cluster Server Agent for Microsoft Exchange ■ Configuring VxSAS
“Configuring the cluster” on page 459	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard (VCW) ■ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 476	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the Exchange data, log, RegRep, and MTA using the VEA
“Preparing the forest and domain” on page 451	<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 486	<ul style="list-style-type: none"> ■ Review the prerequisite checklist ■ Run the Exchange Setup Wizard for Veritas Cluster Server ■ Install Microsoft Exchange
“Moving Exchange databases to shared storage” on page 490	<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
“Installing Exchange on additional nodes” on page 495	Run the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Installation on additional nodes
“Configuring the Exchange service group for VCS” on page 500	Creating an Exchange service group using the VCS Exchange Configuration Wizard
“Creating the primary system zone” on page 507	<ul style="list-style-type: none"> ■ Creating the primary system zone ■ Adding the nodes to the primary zone

Table 11-1 Tasks for configuring Replicated Data Clusters for Exchange

Objective	Tasks
“ Verifying the installation in the primary zone ” on page 508	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes
“ Creating a parallel environment in the secondary zone ” on page 509	<ul style="list-style-type: none"> ■ Reviewing the prerequisites ■ Reviewing the configuration ■ Configuring the network and storage ■ Installing SFW HA ■ Adding the systems in the secondary zone to the cluster ■ Configuring disk groups and volumes for Exchange ■ Installing the application
“ Setting up the Replicated Data Sets (RDS) ” on page 516	<ul style="list-style-type: none"> ■ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones
“ Configuring a hybrid RVG service group for replication ” on page 528	<ul style="list-style-type: none"> ■ Creating a hybrid Replicated Volume Group (RVG) service group ■ Configuring the hybrid RVG service group
“ Setting a dependency between the service groups ” on page 538	Setting up a dependency between the RVG Service Group and the Exchange service group
“ Adding the nodes from the secondary zone to the RDC ” on page 538	<ul style="list-style-type: none"> ■ Adding the nodes from the secondary zone to the RVG service group ■ Configuring the IP resources for fail over ■ Adding the nodes from the secondary zone to the Exchange service group
“ Verifying the RDC configuration ” on page 544	Verifying that fail over occurs first within zones and then from the primary to the secondary zone

Reviewing the prerequisites

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

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

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/302144>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agent for Microsoft Exchange.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2003 versions

The following table lists the Microsoft Exchange Server 2003 versions supported with SFW HA 5.1 Service Pack 1.

Table 11-3 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

Table 11-3 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Memory must be a minimum 256 MB of RAM per server for Exchange 2003; refer to the Microsoft documentation for more information.
- 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. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
See "[Best practices for SFW HA](#)" on page 447.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each Exchange Virtual Server (EVS).

- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone 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 Dynamic DNS (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 for SFW HA

The following permissions are required:

- 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 on 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 Exchange virtual server computer object 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, 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 for SFW HA

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 Server 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 2003.
- 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.

Best practices for SFW HA

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.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

Note: Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space.

[Table 11-5](#) estimates disk space requirements for SFW HA.

Table 11-5 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

Reviewing the configuration

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

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

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

Sample configuration

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

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

Primary zone

SYSTEM1 & SYSTEM2	First and second nodes of the primary zone
EVS1	Microsoft Exchange Virtual Server name
EVS1_SG1	Microsoft Exchange service group
EVS1_SG1_DG	Cluster disk group names
EVS1_SHARED_DG	
EVS1_SG1_DATA	Volume for Microsoft Exchange Server database
EVS1_SG1_LOG	Volume for Microsoft Exchange Server database log file
EVS1_SG1_REGREP	Volume that contains the list of registry keys that must be replicated among cluster systems for the Exchange server
EVS1_SHARED_VOL	Volume for storing Microsoft Exchange Server MTA database, SMTP, and message tracking
EVS1_REPLOG	Replicator log volume required by VVR

Secondary zone

SYSTEM3 & SYSTEM4	First and second nodes of the secondary zone
-------------------	--

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

RDS and VVR Components

EVS1_RDS	RDS name for Exchange Server database
EVS1_RVG	RVG name for Exchange Server database
EVS1_RVG_SG	Replication service group for Exchange database and files

Configuring the storage hardware and network

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.

- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network 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 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, 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 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the options to install VVR and the Veritas Cluster Server Application Agent for Exchange. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 11-6](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 11-6 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.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 453.

To change the driver signing options on each local system

- 1 Open the Control Panel and click **System**.
- 2 Click the Hardware tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or one of the other options from the table, to allow installation to proceed.
- 4 Click **OK**.
- 5 Repeat for each computer.
If you do not change these options, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing options to their previous states.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

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 5.1 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 and click **Next**. Be sure to select the following.

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

- 10 Select the following 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 a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, 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

You can run the VVR Security Service Configuration (VxSAS) wizard after you install SFW HA on both the primary and secondary nodes. When you run the wizard, you can then specify the primary and secondary sites in one step.

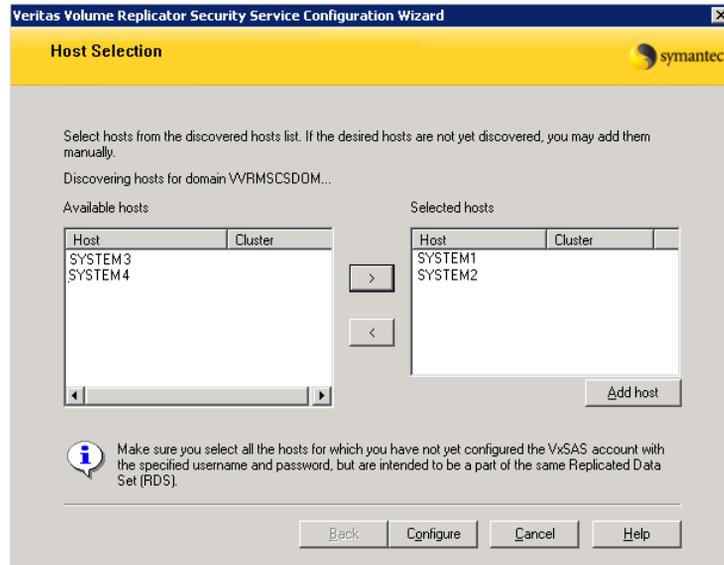
Complete the following procedure to configure the VxSAS service for VVR.

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 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*.

- 4 On the Host Selection panel, select the required hosts:



Selecting hosts

The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host 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, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with 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 Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

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

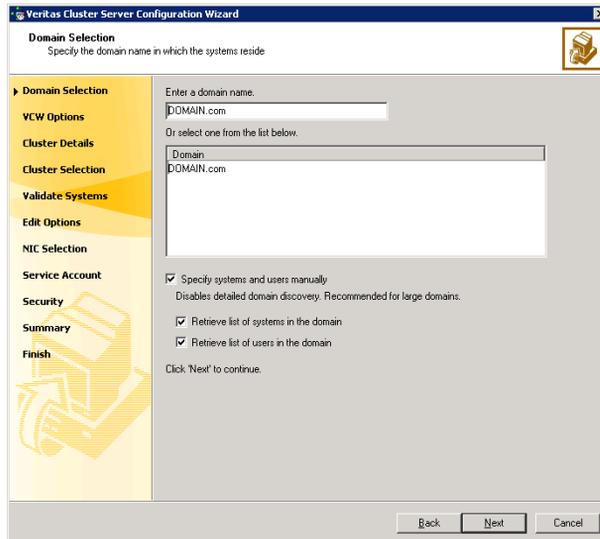
- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

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

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

To configure a VCS cluster

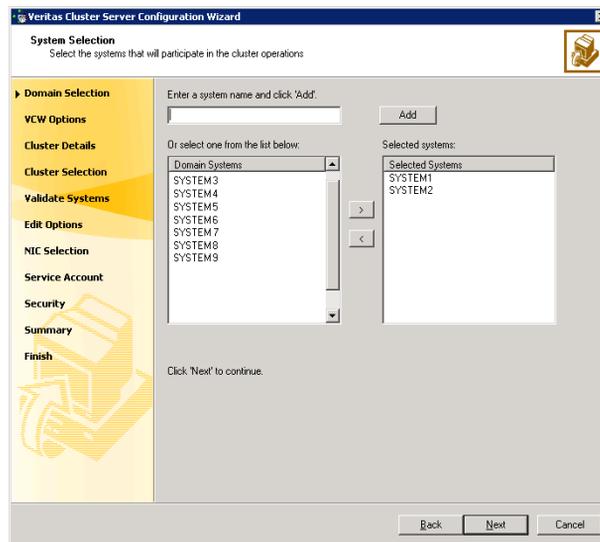
- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations**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



Do one of the following:

- 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 8](#) on page 462.
- 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 461. 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 8](#) on page 462.
- 6 On the System Selection panel, specify the systems for the 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.

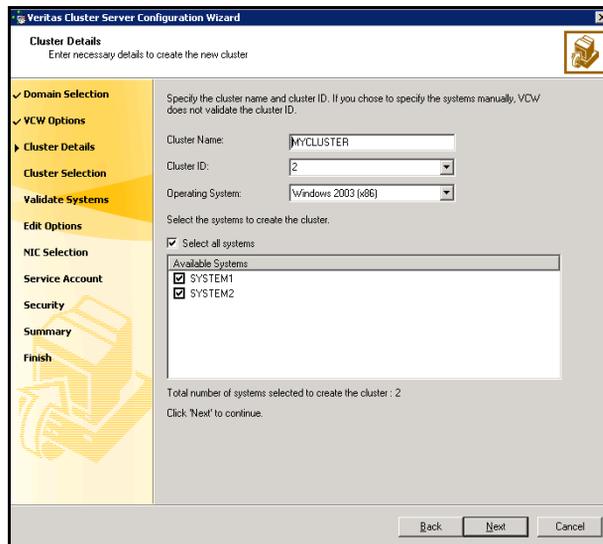
- The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

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

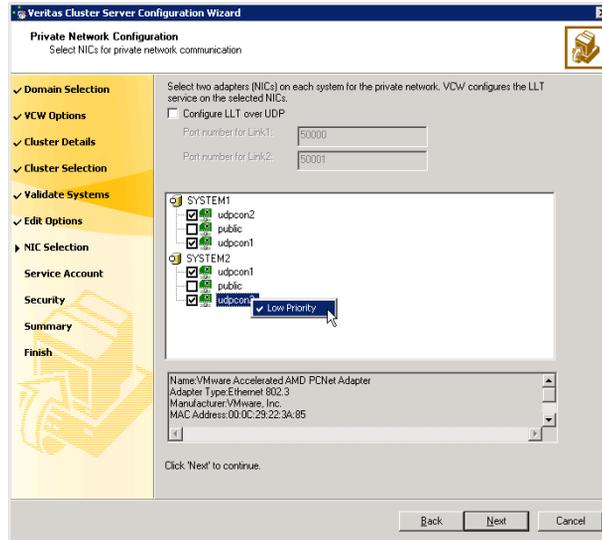


Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

Cluster ID	<p>Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.</p> <p>Caution: 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.</p>
Operating System	<p>From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.</p>
Available Systems	<p>Select the systems that you wish to configure in the cluster. Check the Select all systems check box to select all the systems simultaneously.</p> <p>The wizard discovers the network interface cards (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.</p>

- 10 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 the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 466.
- 11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

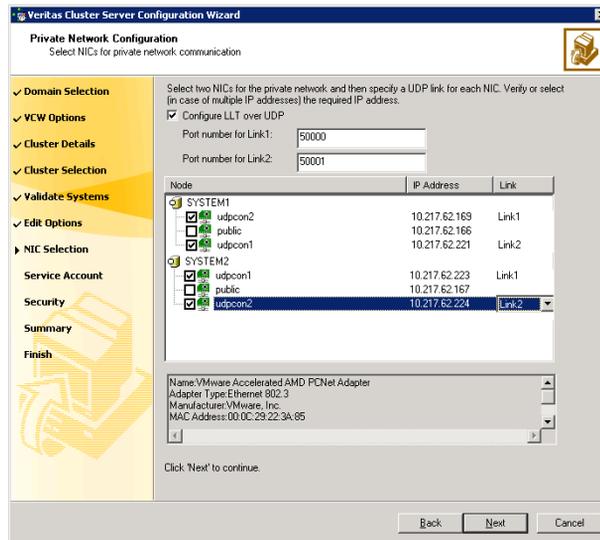
- To configure the VCS private network over the ethernet, complete the following steps:



- 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 both public and private communication.
- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



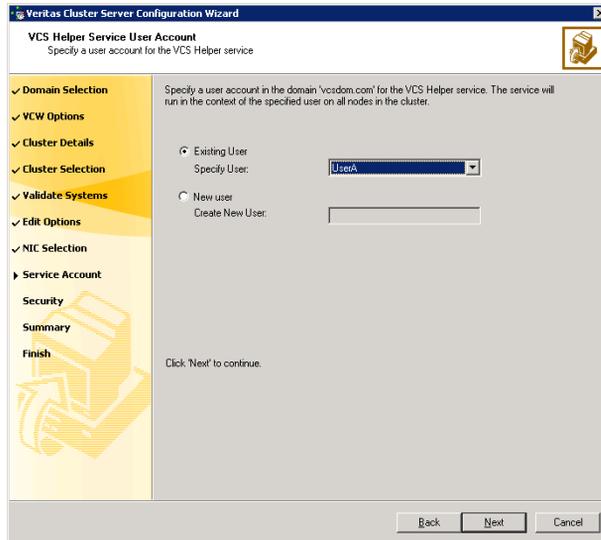
- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

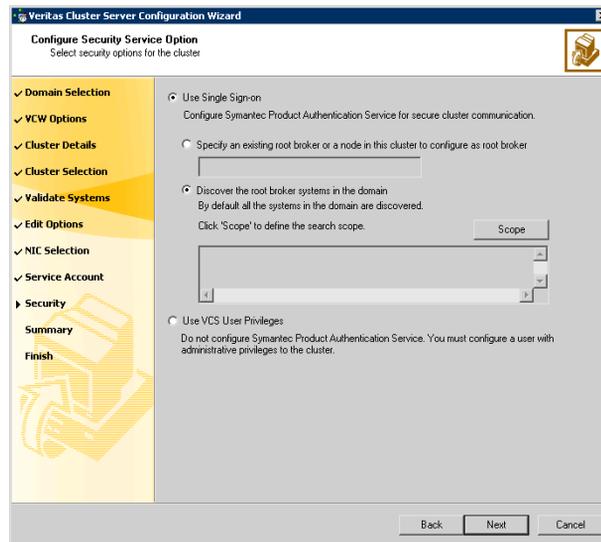
- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.



Specify a domain user as follows:

- 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 460, 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**.
- 13 On the Configure Security Service Option panel, specify the security options for the cluster and then click **Next**.
Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability 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 an existing root broker or a node in this cluster to configure as root broker**, 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 then configures all nodes in the cluster as authentication brokers.

- If you want to search 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 a user *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**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 11-7 contains some more examples of search criteria.

Table 11-7 Search criteria examples

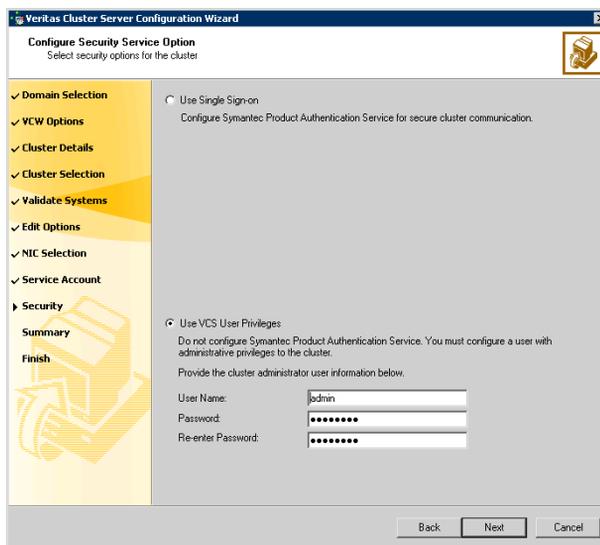
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- 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 a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS Encrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password. After the cluster is configured, you can 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**.

- 14 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.

- 15 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**.

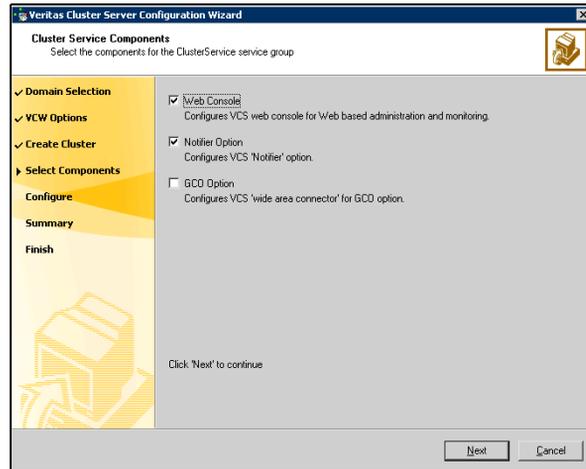
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

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.

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



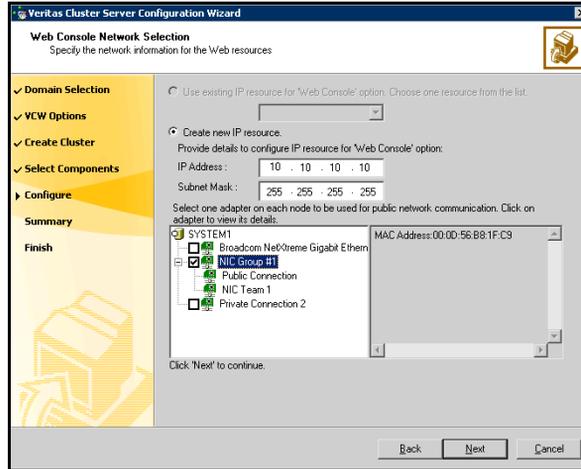
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See “[Configuring Web console](#)” on page 471.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 472.

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**.



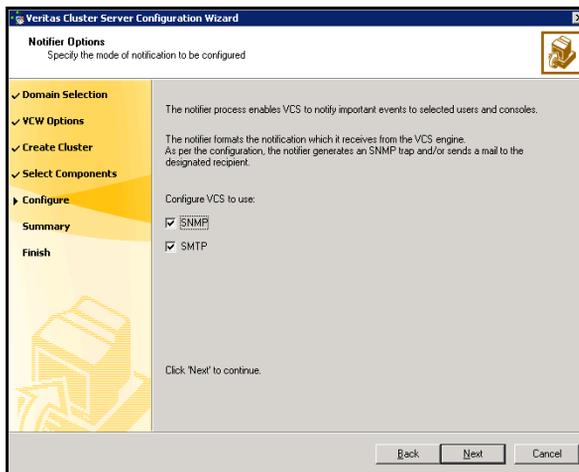
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 472. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

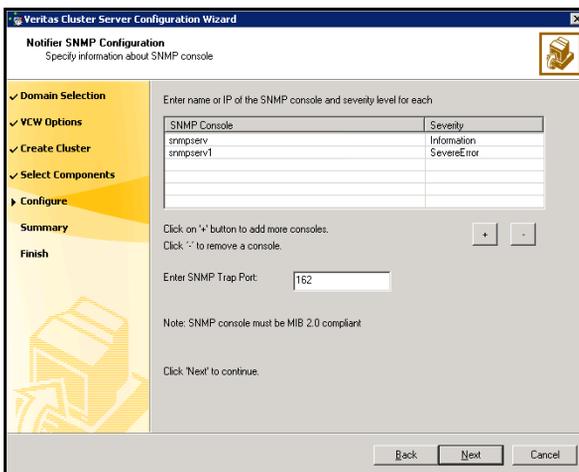
To configure notification

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



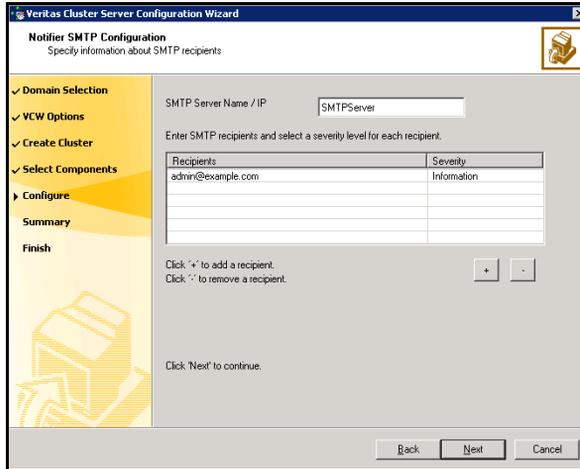
You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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



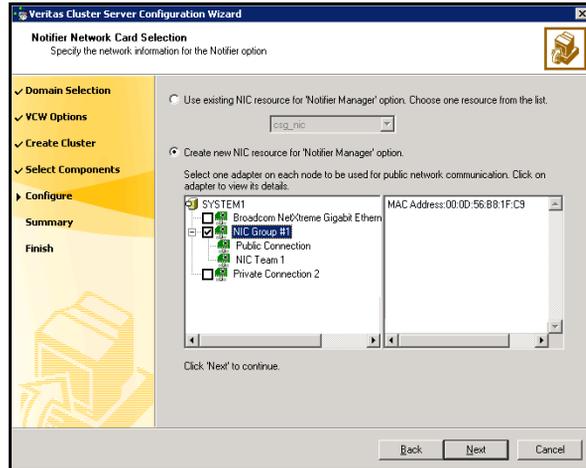
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.

- Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.
 - Enter an SNMP trap port. The default value is "162".
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



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

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



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

Configuring cluster disk groups and volumes

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

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - *Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

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

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

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

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

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

On the first node of the cluster you will first need to create the following cluster disk groups and volumes on shared disks:

Exchange disk group EVS1_SG1_DG should contain the following volumes:

- **EVS1_SG1_DATA:** Contains the Exchange Server 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 storage group **EVS1_SHARED_DG** should contain the following volumes:

- **EVS1_SG1_REGREP:** Contains the list of registry keys that must be replicated among cluster systems for the Exchange Server.
- **EVS1_SHARED_VOL:** Contains the MTA database, SMTP, and message tracking.
- **EVS1_REPLOG:** Contains the VVR Replicator Log.

You can create this volume later while setting up the Replicated Data Sets. See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 516.

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

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

Creating a cluster disk group

Create the required cluster disk groups on the first node of the cluster.

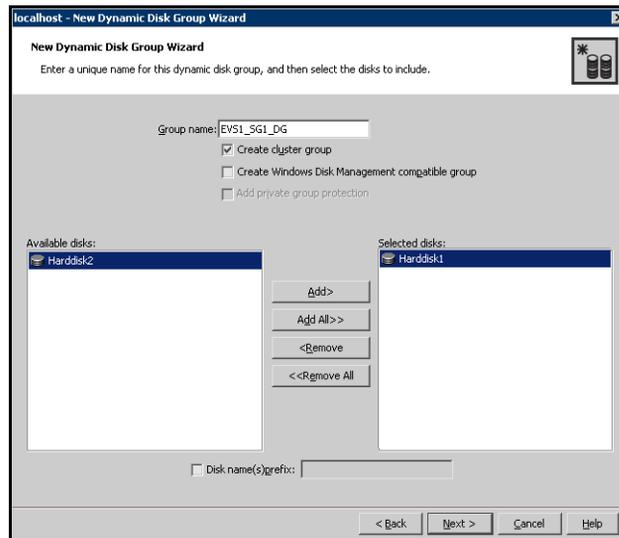
To create a dynamic (cluster) disk group

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

- 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

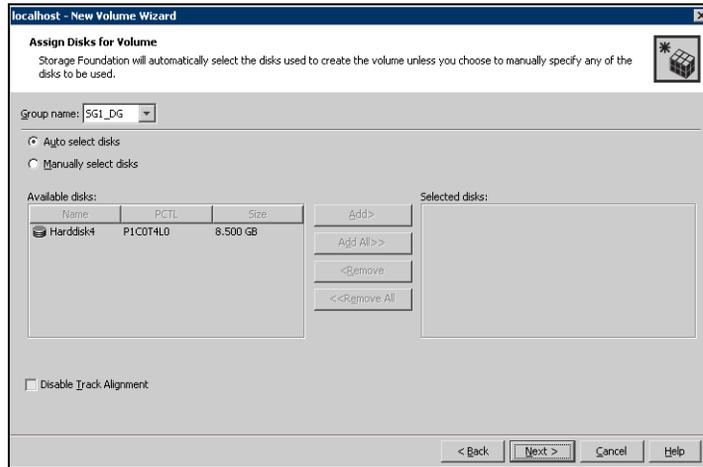
The following procedure describes how to create a volume on a cluster disk group. Repeat these steps to create the required volumes in the cluster disk groups (EVS1_SG1_DG, EVS1_SHARED_DG) on the first node of the cluster.

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

To create dynamic volumes

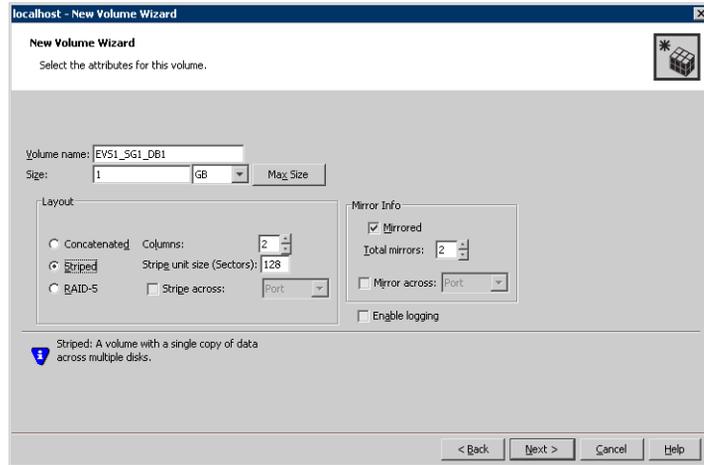
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example SG1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

6 Select the disks for the volume.



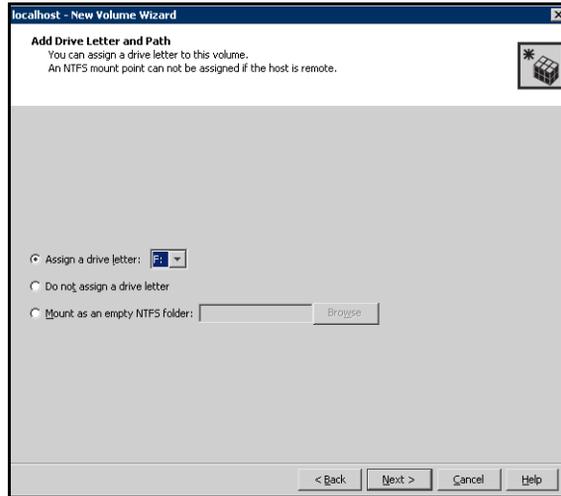
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected 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

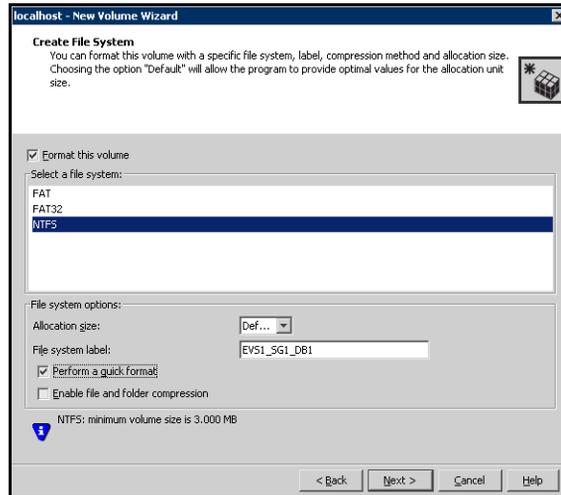
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- 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**.

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).

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 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 451.
- 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 484.
- Mount the volume containing the information for registry replication (EVS1_SG1_REGREP).
See “[Importing a disk group and mounting a shared volume](#)” on page 484.
- 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 Exchange installation. Exchange must be installed 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - 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.
 - Either the logged-on user or the VCS Helper service user account must have write permissions on the DNS server to perform DNS updates.

- Make sure the VCS Helper service 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.

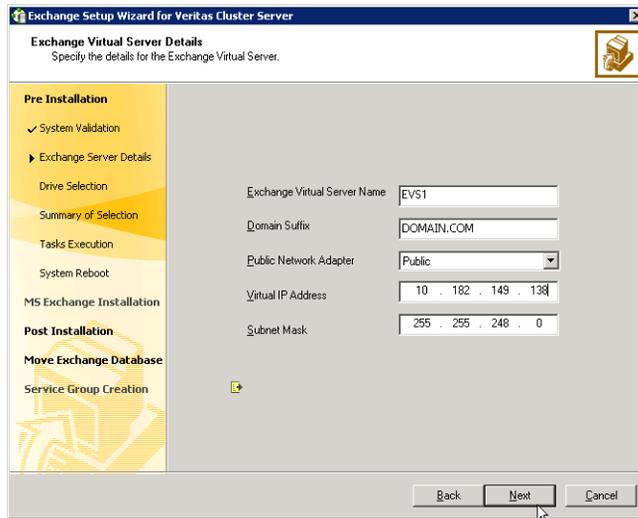
Exchange pre-installation: First node

Use the Exchange Setup Wizard for VCS to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. Exchange must be installed on a virtual node in order 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 Tools > 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.

Warning: 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 unique on the network.
- 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.

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

If you wish to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: First node

Install Exchange on the node on which you performed the pre-installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```

The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```

```
C:\>net start had
```
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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.
Once the node is rebooted, move the databases created during the Exchange installation, from the local drive to the shared storage.

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 storage managed by the cluster disk group. This is necessary to ensure proper failover operations in the cluster.

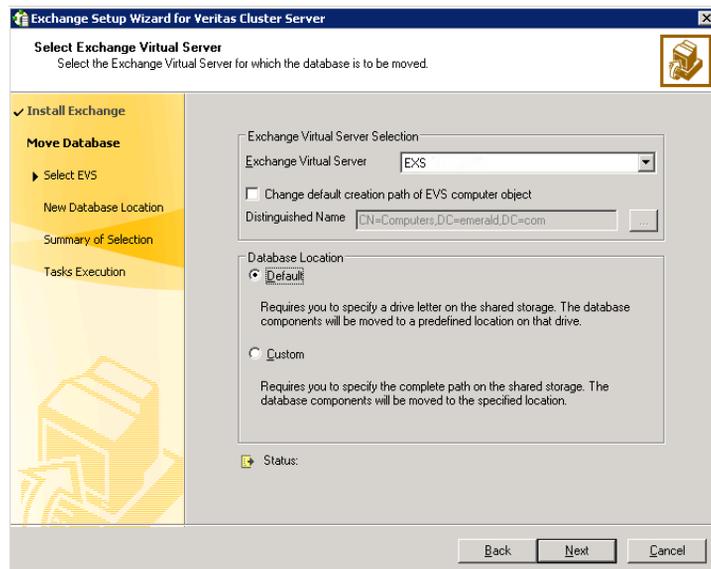
Complete the following tasks before moving the databases:

- 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 484.

To move Exchange databases

- 1 Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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, click **Configure/Remove highly available Exchange Server** and then click **Next**.
- 4 In the Select Option dialog box, click **Move Exchange Databases** and then click **Next**.
- 5 In the Select Exchange Virtual Server dialog box, choose the Exchange virtual server and the database location option and then click **Next**.



Exchange Virtual Server

From the drop-down list, select the Exchange virtual server for which you want to move the databases.

Change default creation path of EVS computer object

Perform the following steps if you wish to change the default path for the Exchange virtual server object in Windows Active Directory:

- Check the **Change default creation path of EVS computer object** check box.
- Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**.
To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
The Lanman agent performs Windows AD updates. These settings are applicable to the Lanman resource in the service group.
By default, the Lanman resource adds the virtual server to the default container "Computers."

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

Default

Click **Default** if you wish to move the databases to predefined location on the shared storage. In the next step the wizard prompts you to specify the drive letter on the shared storage. The first mailbox store, public store, and MTA data are then moved to the generated default paths on the volumes that you specify.

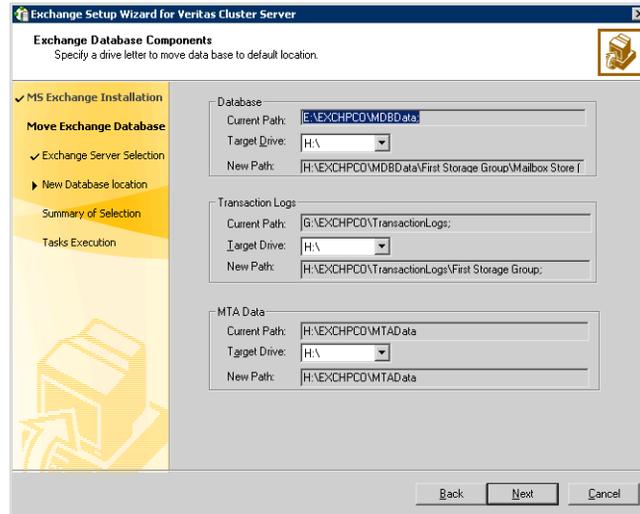
Caution: The Exchange data files and the MTA must be in different paths, and the MTA cannot be at the root level (for example K:\).

Custom

Click **Custom** if you wish to move the databases to a specific location on the shared storage. Choosing a custom location allows you to specify the Exchange database and streaming path. In the next step the wizard prompts you to specify the entire path of the location on the shared storage. The wizard then moves the databases to the specified directory.

If you chose the Default option, proceed to the next step. If you chose the Custom option, proceed to [step 7](#) on page 494.

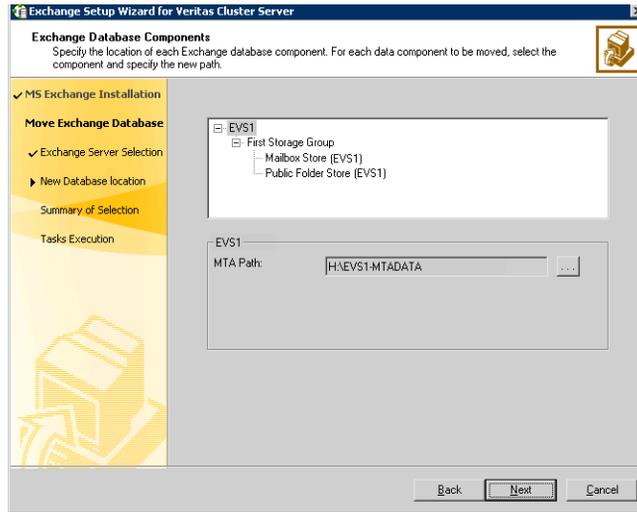
- 6 For the option of a default database location, specify the drives for moving the Exchange database components. The database components are then moved to a predefined location on that drive.



On the Exchange Database Components panel, complete the following steps:

- Specify a drive for moving the Exchange database.
- Specify a drive for moving the Exchange Transaction Logs.
- Specify a drive for moving the Exchange MTA Data.
- Click **Next** and proceed to [step 8](#) on page 494.

- 7 For the option of a custom database location, specify the location for specific Microsoft Exchange data components and then click **Next**.



For each data component that you wish to move, select the component and then click the ellipsis (...) to browse for the folder where you want to move it.

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

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

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.

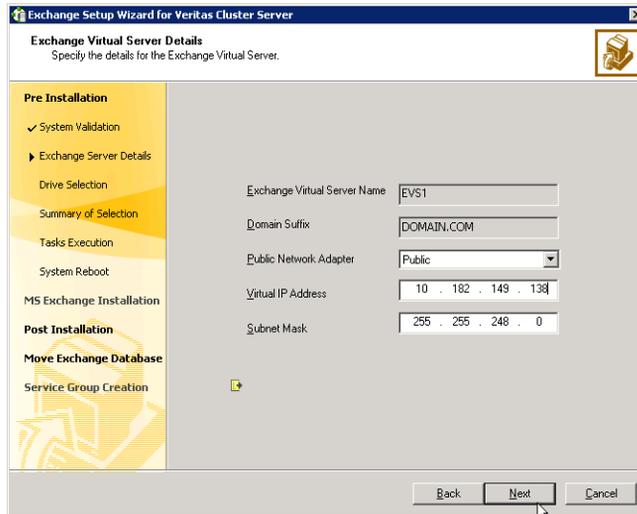
See “[Unmounting a volume and deporting a disk group](#)” on page 485.

Use the Exchange Setup Wizard for VCS 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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: Additional nodes

- Install Exchange on the additional node on which you performed the 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.

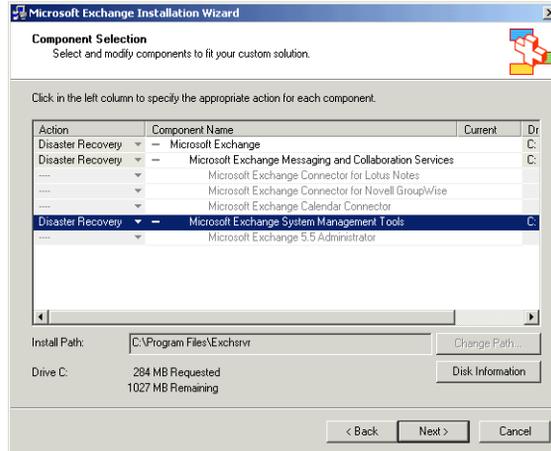
To install Exchange

- 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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:
`C:\>hasys -state`
The state should display as **RUNNING**.
If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`

- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

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. Refer to the Exchange documentation for instructions.

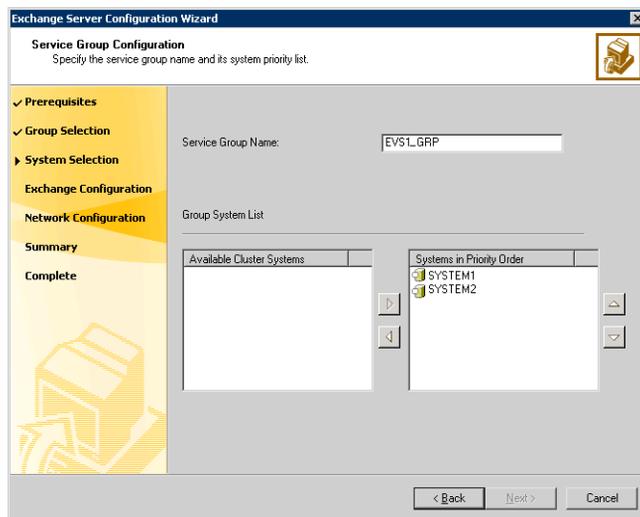
Prerequisites

- You must be a Cluster Administrator. This privilege is required to configure service groups.
- 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 volume](#)” on page 484 for instructions on mounting and “[Unmounting a volume and deporting a disk group](#)” on page 485 for instructions on unmounting.
- 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

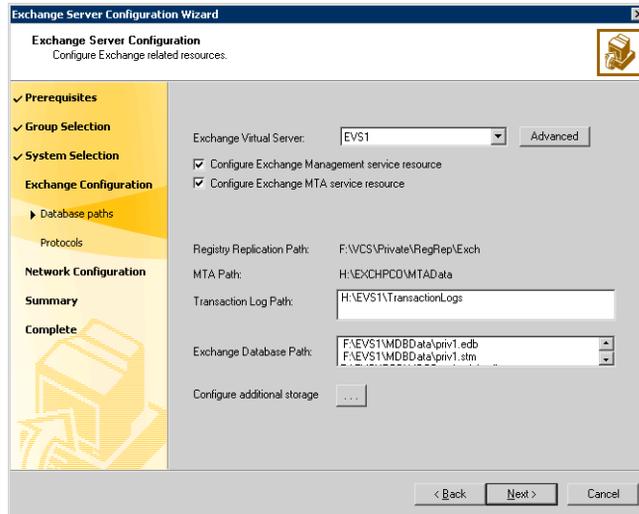
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**: The wizard starts validating your configuration. Various messages indicate the validation status.



- Enter a name for the Exchange service group.
 If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.

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



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
- Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.
On the Lanman Advanced Configuration dialog box, complete the following:
 - In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
 - Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.

If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

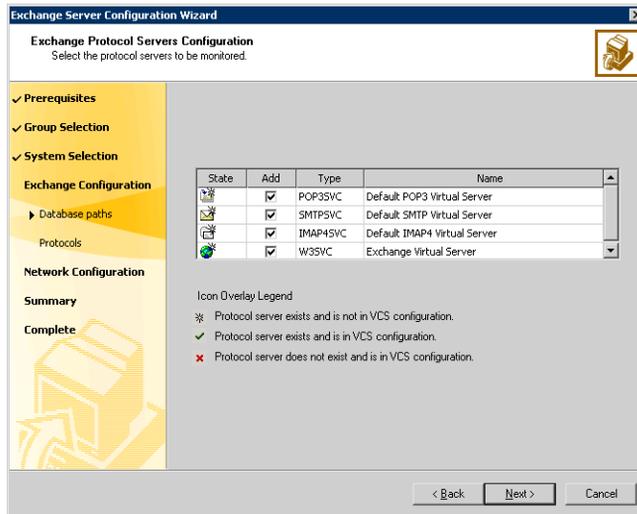
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.

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.

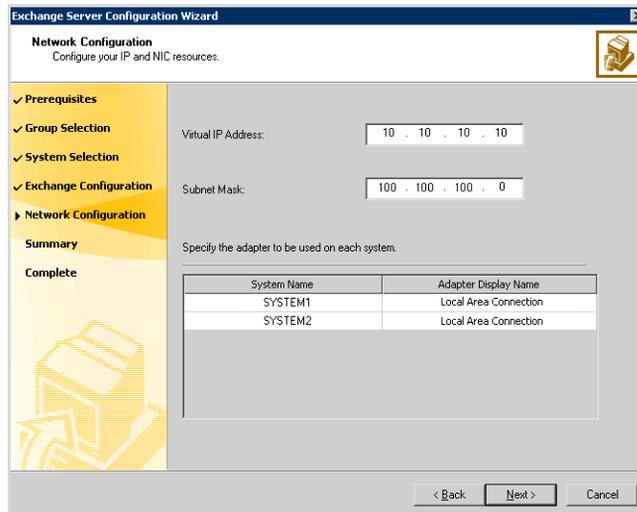
If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.

- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:
 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- 6 On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.



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



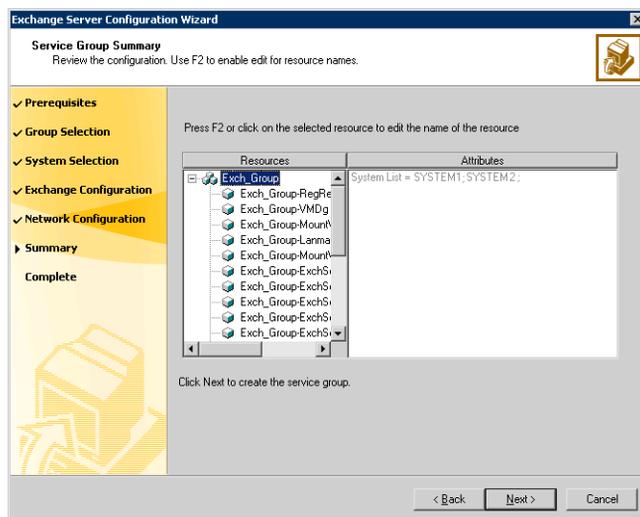
- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.

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 Display Name** field to view the adapters associated with a node.

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.

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



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 the **Esc** key.

- 9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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.

Creating the primary system zone

In the service group, set up systems in the primary zone (zone 0) to specify that initial fail over occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 Select the Exchange Server service group (EVS1_SG1) in the left pane and the Properties tab in the right pane.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.
- 7 Click **OK**.

Verifying the installation in the primary zone

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

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

The service group you selected is taken offline and brought online on the node that you selected.

Creating a parallel environment in the secondary zone

After setting up a SFW HA environment in the primary zone, use the guidelines in the following list to complete the same tasks in the secondary zone (Zone 1).

Before you begin to configure the secondary zone, offline the following resources in the Exchange service group in the primary zone:

- Exchange Server resource
- Exchange virtual server name resource
- Exchange virtual IP resource

The remaining resources should be online, including the VMDg resources and the MountV resources.

- [“Reviewing the prerequisites”](#) on page 442
- [“Reviewing the configuration”](#) on page 449
- [“Configuring the storage hardware and network”](#) on page 450
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 452
- [“Configuring VxSAS”](#) on page 456
- [“Adding the systems in the secondary zone to the cluster”](#) on page 510
- [“Configuring cluster disk groups and volumes”](#) on page 476

During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as the cluster at the primary zone:

- Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
 - [“Installing Exchange on additional nodes”](#) on page 495
- Follow the steps in this section to install Exchange on all nodes in the secondary zone. The instance name must be the same in the primary zone and secondary zone.

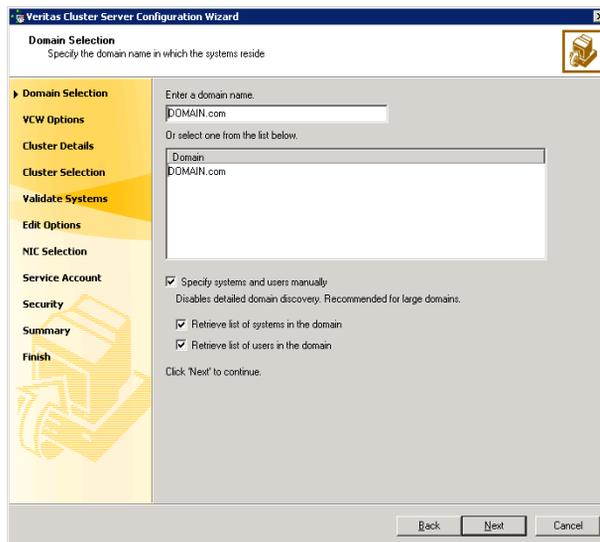
Note: After you install Exchange on the nodes in the secondary zone, make sure to use VEA to remove all the drive letters from the configured volumes to avoid conflicts during the configuration of the zones.

Adding the systems in the secondary zone to the cluster

Add the nodes in the secondary zone to the existing cluster with the following procedure.

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.



Do one of the following:

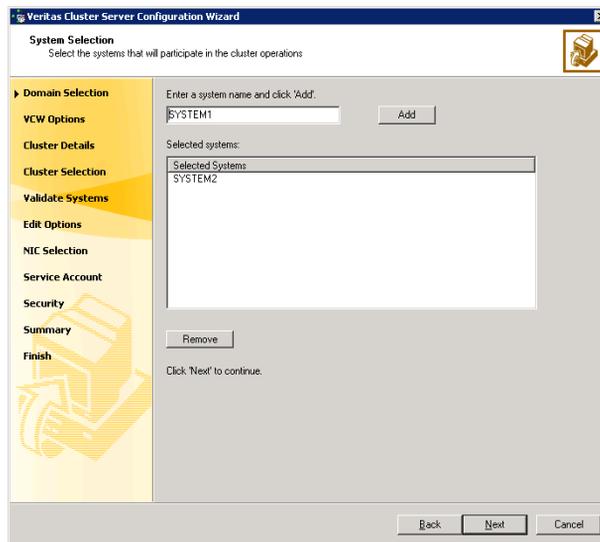
- To discover information about all the systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 513.

- 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 512. Otherwise proceed to the next step.

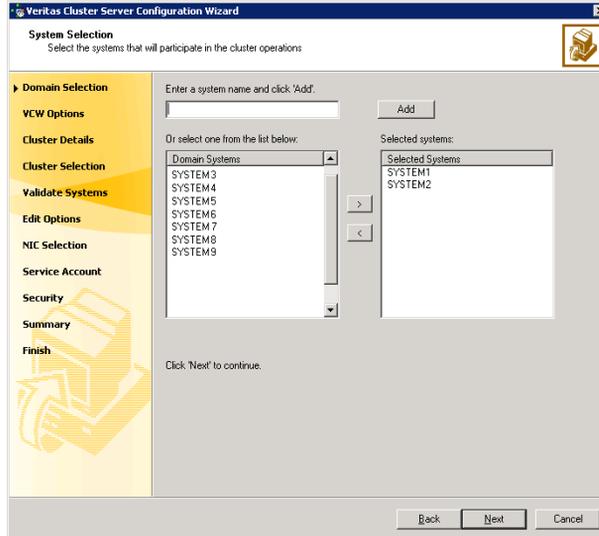
- 5 On the System Selection panel, complete the following and click **Next**.



- 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 8](#) on page 513.

- 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.

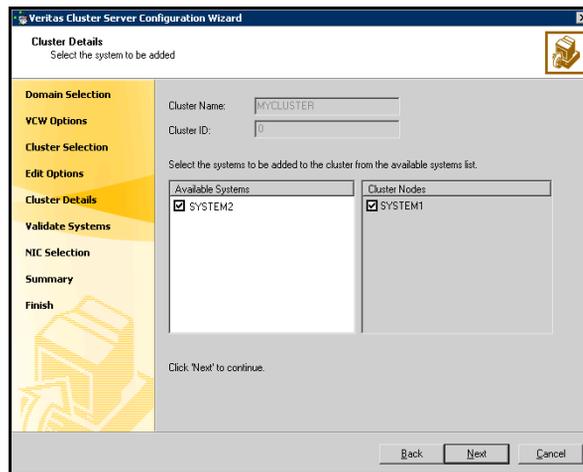
- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 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.
- 10 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, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.
- 11 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.

- 12 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.
- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**.
 How you configure the VCS private network communication depends on how it is configured in the cluster. If LLT is configured over ethernet, you

have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - 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 both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.

The wizard will configure the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Check the **Configure LLT over UDP** check box.
 - Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
 - Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard will configure the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 14 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.

- 15 Specify the password for the user in whose context the VCS Helper service runs.
- 16 Review the summary information and click **Add**.
- 17 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

Return to the task list “[Creating a parallel environment in the secondary zone](#)” on page 509.

Setting up the Replicated Data Sets (RDS)

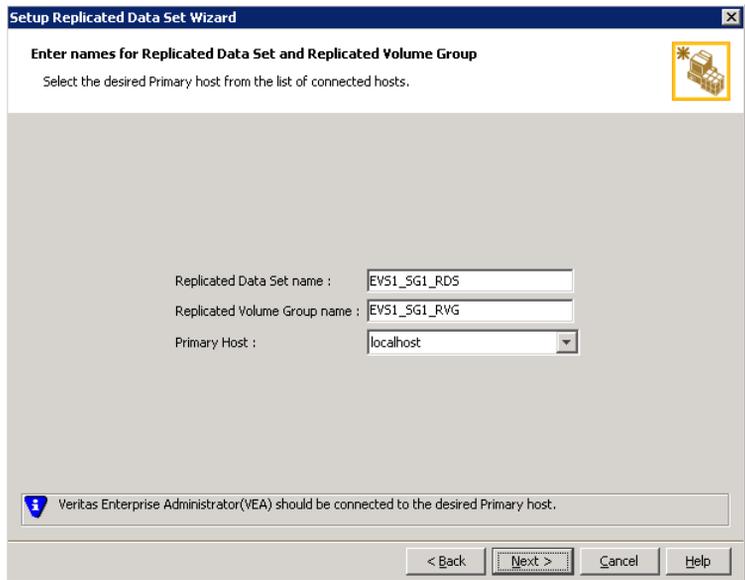
Set up the Replicated Data Sets (RDS) in the primary zone and secondary zone. You can configure an RDS using the Create RDS wizard for both zones.

- Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- Verify that the cluster disk group is imported and the volumes are mounted in the primary and secondary zone

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.

3 Read the Welcome page and click **Next**.

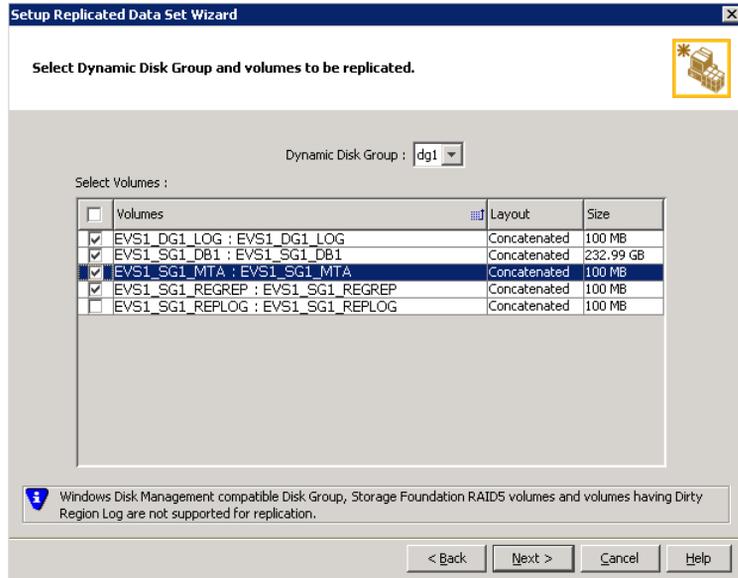


By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

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.

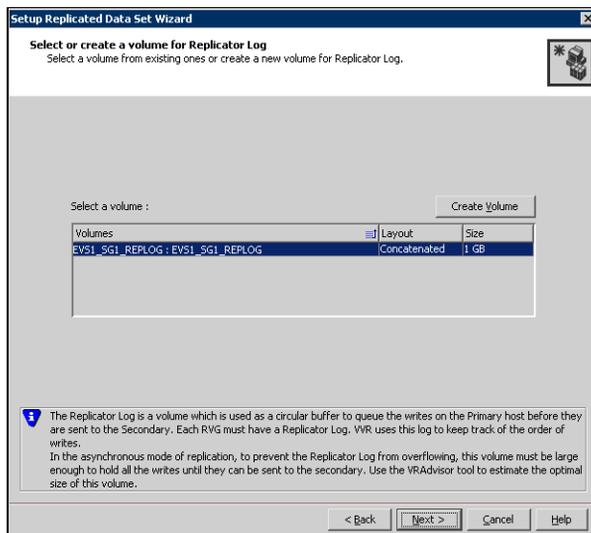


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (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

Enables you to specify the disk selection method.

- Enable the **Thin Provisioned Disks Only** checkbox to ensure that the Replicator Log volume is created only on Thin Provisioned (TP) disks.

Note: The checkbox will remain disabled if the diskgroup does not have any TP disk.

If this option is selected along with the **Select disks automatically** option, then the Replicator Log volume will be created only on TP disks. However, if you enable this checkbox along with **Select disks manually** option, then the user can select only TP disks from **Available Disks**.

For more information on Thin Provisioning refer to the *Veritas Storage Foundation Administrator's Guide*.

- Choose the **Select disks automatically** option if you want VVR to select the disks.
- Choose the **Select disks manually** option to use specific disks from the Available disks pane for creating the volume. Either double-click on 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 Primary RVG has been created successfully, VVR displays the following message:

```
RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?
```

Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.

Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.

10 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field. Click **Next**. If the Secondary host is not connected to VEA, the wizard tries to connect it when you click **Next**. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.

- 11 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.
 The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:
- the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the Primary
- Otherwise, the RDS setup wizard enables you to create the required volumes manually.
- Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
- 12 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.
 This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.
- If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
 - If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.
 Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.
 When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.
- If all the data volumes to be replicated meet the requirements, this screen does not occur.

- 13 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

The screenshot shows the 'Setup Replicated Data Set Wizard' dialog box, specifically the 'Edit replication settings' page. The title bar reads 'Setup Replicated Data Set Wizard'. Below the title bar, the text 'Edit replication settings' is displayed, followed by the instruction 'Edit replication settings or click next.' There is a small icon of a server rack with a star in the top right corner. The main area contains several configuration fields, each with a label and a drop-down menu or text box. The fields are: 'Primary side IP' with the value '10.217.53.214', 'Secondary side IP' with '10.217.53.215', 'Replication Mode' with 'Synchronous Override', 'Replicator Log Protection' with 'AutoDCM', 'Primary RLINK Name' with 'Pri_RLINK', and 'Secondary RLINK Name' with 'Sec_RLINK'. Below these fields is an 'Advanced' button. At the bottom of the dialog, there is a warning message: 'DHCP addresses are not supported by VWR.' and navigation buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Secondary side IP Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode	<p>Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous. The default is synchronous override.</p> <p>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.</p> <p>Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates.</p> <p>Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously.</p> <p>If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as <code>MISSING</code>.</p>
Replicator Log Protection	<p>The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.</p> <p>The DCM option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.</p> <p>The Off option disables Replicator Log Overflow protection.</p> <p>In the case of the Bunker node. Replicator Log protection is set to Off, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.</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.

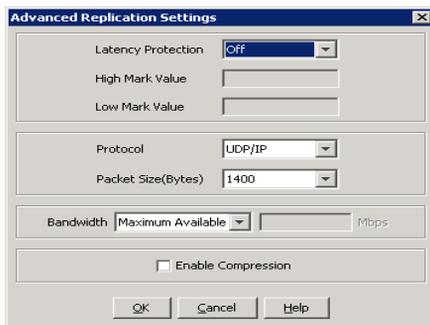
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

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	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- 14 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.

Enable Compression Enable this checkbox if you want to enable Compression for the secondary host.

Click **OK** to close the dialog box.

15 Click **Next**.

16 On the **Start Replication** page, select **Start Replication**.

Synchronize
Automatically

If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from
Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.

17 Review the information.

Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

If the additional Exchange storage groups are in a separate disk group, repeat the procedure “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 516 for the disk group that contains the Exchange storage groups. Provide unique names for the Replicated Data Set name, and the Replicated Volume Group name.

See “[Sample configuration](#)” on page 449 for a list of example names.

Configuring a hybrid RVG service group for replication

Create and configure a hybrid Replicated Volume Group (RVG) service group for replication. The RVG service group is hybrid because it behaves as a failover service group within a zone and as a parallel service group between zones.

For additional information about service group types, see the *Veritas Cluster Server 5.1 Administrator's Guide*.

Configure the RVG service group's resources manually by copying and modifying components of the Exchange Server service group. Then create new RVG resources and bring them online.

The RVG service group for RDC contains the following resources:

Table 11-8 Replication service group resources

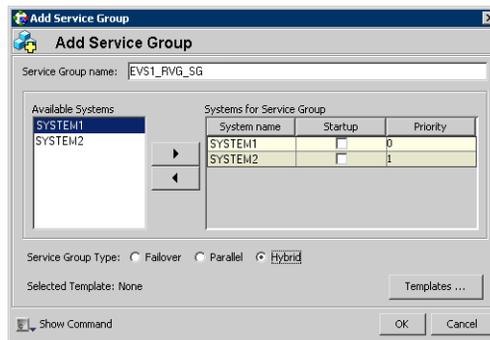
Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg for the disk group	Volume Manager disk group with Exchange database files
VvrRvg for the disk group	Replicated volume group with Exchange database files

Creating the RVG service group

Create a hybrid replicated volume (RVG) service group, to contain the resources for replication.

To create a hybrid RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.
- 3 In the **Add Service Group** window:



- Enter a name for the service group, for example EVS1_RVG_SG.
- Select the systems in the primary zone (Zone 0) and click the right arrow to add them to the service group.
- Select **Hybrid**.
- Click **OK**.

Configuring the RVG service group for RDC replication

Configure the RVG service group’s resources manually for RVG by completing the following tasks:

- **“Configuring the IP and NIC resources”**
Copy IP and NIC resources of the Exchange Server service group (EVS1_SG1), paste and modify them for the RVG service group (EVS1_RVG_SG).
- **“Configuring the VMDg resources for the disk groups”**
Copy the VMDg resources for all disk groups in the Exchange Server service group (EVS1_SG1), paste and modify them for the RVG service group (EVS1_RVG_SG).
- **“Adding the VVR RVG resources for the disk groups”**
Create the VVR RVG resources for all the disk groups and enter the attributes for each of the disk groups and the replication IP address.
- **“Linking the VVR RVG resources to establish dependencies”**
Link the VVR RVG resources to establish the dependencies between the VMDg resources, the IP resource for replication, and the VVR RVG resources for the disk groups. Configure the RVG service group’s VMDg resources to point to the disk groups that contain the RVGs.
- **“Deleting the VMDg resource from the Exchange Server service group”**
Delete the VMDg resources from the Exchange Server service group, because they depend on the replication and were configured in the RVG service group.

Configuring the IP and NIC resources

Configure the following resources and attributes for the IP and NIC:

Table 11-9 IP and NIC resources

Resource	Attributes to Modify
IP	Address
NIC	(none)

To create the IP resource and NIC resource

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EVS1_SG1) in the left pane.
- 2 On the **Resources** tab, right-click the IP resource (EVS1_SG1-IP), and click **Copy > Self and Child Nodes**.

- 3 In the left pane, select the RVG service group (EVS1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the names of the IP and NIC resources for the RVG service group.
- 6 Click **OK**.

To modify the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (EVS1_RVG_SG-IP) and select **View > Properties View**.
- 2 In the **Properties View** window, for the **Address** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, enter the VVR IP address for the Primary Zone as the scalar value.
- 4 Close the **Properties View** window.

To enable the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (EVS1_RVG_SG-IP) and select **Enabled**.
- 2 In the **Resources** tab display area, right-click the NIC resource (EVS1_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg resources for the disk groups

Create the VMDg resources in the Exchange Server service group, and clear the DGGuid attribute for the new VMDg.

Configure the following attributes in the Exchange Server service group for the MountV resource:

Table 11-10 MountV resources

Resource	Attributes to Modify
Resources for disk groups for the Exchange files:	
MountV (for the Exchange Server database volume)	VMDg Resource Name Volume Name
MountV (for the Exchange log volume)	VMDg Resource Name Volume Name
MountV (for the registry volume)	VMDg Resource Name Volume Name

To create the VMDg resource for a disk group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EVS1_SG1) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the disk group, with the Exchange files (EVS1_SG1-VMDg), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (EVS1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to EVS1_RVG_SG-VMDg.
- 6 Click **OK**.

To clear the DGGuid attribute for the new VMDg

- 1 In the **Resources** tab display area, right-click the new VMDg resource.
- 2 In the same **Properties View** window, for the **DGGuid** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, clear the scalar value for the **DGGuid** attribute.
- 4 Close the **Properties View** window.

To modify the MountV resources in the Exchange Server service group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EVS1_SG1) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the Exchange Server files (EVS1_SG1-MountV) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the Exchange Server database files (EVS1_SG1_DATA).
- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example EVS1_RVG_SG-VMDg.
- 6 Close the **Properties View** window.
- 7 Repeat these steps to modify the VMDGResName value for the additional MountV resources for the Exchange log volume and the Exchange registry replication volume.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (EVS1_RVG_SG).
- 2 In the **Resources** tab display area, right-click the VMDg resource (EVS1_RVG_SG-VMDg) and select **Enabled**.

Adding the VVR RVG resources for the disk groups

Add VVR RVG resources for replication of the disk groups. If the application has multiple disk groups, create a separate VvrRvg resource for each disk group. Configure the following attributes in the RVG service group for the VvrRvg resource:

Table 11-11 VvrRvg resources

Resource	Attributes to Modify
Resources for the disk group for the Exchange files:	
VvrRvg	VMDgResName IPResName

To create the VVR RVG resource for a disk group

- 1 In the left pane, select the RVG service group (EVS1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:
 - Enter the **Resource Name** for the VVR RVG resource, for example, EVS1_VvrRvg.
 - Select the **Resource Type** of VvrRvg.
- 3 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the RVG group that is being managed, for example EVS1_RVG.
- 5 Click **OK**.
- 6 In the **Add Resource** window, for the **VMDGResName** attribute, click **Edit**.
- 7 In the **Edit Attribute** window, enter the name of the disk group containing the RVG, for example EVS1_RVG_SG-VMDg.
- 8 Click **OK**.
- 9 In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, enter the name of the IP resource managing the IP address for replication, for example EVS1_RVG_SG-IP.

- 11 Click **OK**.
- 12 In the **Add Resource** window, verify that the attributes have been modified:
- 13 Click **OK**.

Linking the VVR RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the RVG service group to establish the dependencies between the resources. Start from the top parent and link the following resources:

Table 11-12 Dependencies for VVR RVG resources for RDC

Parent	Child
Resources for the disk group for the Exchange files:	
EVS1_VvrRvg	The IP for replication, for example EVS1_RVG_SG-IP
EVS1_VvrRvg	The VMDg for the Exchange files, for example EVS1_RVG_SG-VMDg

To link the VVR RVG resources

- 1 In the left pane, select the RVG service group (EVS1_RVG_SG).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example EVS1_DB1_VvrRvg.
- 4 Click the child resource, for example EVS1_RVG_SG-IP.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG resources:
 Notice that when you enable a resource and the state of the entity which it is monitoring is online, the corresponding VCS agent reports status for that resource as online. You do not have to bring the resource online manually.

Deleting the VMDg resource from the Exchange Server service group

The VMDg resources must now be manually deleted from the Exchange Server service group because they depend on replication and were configured in the RVG service group.

To delete the VMDg Resources from the Exchange Server service group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EVS1_SG1) from the left pane.

- 2 In the **Resources** tab display area, right-click the VMDg resource for the first disk group (EVS1_SG1-VMDg) and select **Delete**.
- 3 Click **Yes** to confirm that you want to delete it (even if it is online).
- 4 In the **Resources** tab display area, right-click the VMDg resource for any additional disk group, if configured, and select **Delete**.
- 5 Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

Add resources of type RVGPrimary to the Exchange Server service group for each of the Exchange Server disk groups and configure the attributes.

Set the value of the **RvgResourceName** attribute to the name of the RVG resource for the RVGPrimary agent.

Configure the following attributes in the Exchange service group for the RVG Primary resources:

Table 11-13 RVG Primary resources

Resource	Attributes to Modify
Resources for the disk group for the Exchange files:	
RVGPrimary	RvgResourceName

Creating the RVG Primary resources

For each disk group created for this Exchange Server, create a separate RVG Primary Resource for replication.

To create the RVG Primary resource for an Exchange Server disk group

- 1 In the VCS Cluster Explorer window, right-click the Exchange Server service group (EVS1_SG1) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:
 - Enter the **Resource Name** for the RVG Primary resource for the Exchange Server disk group, for example EVS1_RvgPrimary.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource that corresponds to the disk group, for example EVS1_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Veritas Cluster Server Administrator's Guide* for more information about the RVG Primary agent.
- 6 Verify that **Critical** and **Enabled** are both checked.
- 7 Click **OK**.

Linking the RVG Primary resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the Exchange Server service group (EVS1_SG1) to establish the dependencies between the resources for replication.

Link each MountV resource to the appropriate RVGPrimary resource.

Table 11-14 Dependencies for the RVG Primary resources for RDC

Parent	Child
EVS1_SG1-MountV	EVS1_RvgPrimary
EVS1_SG1-MountV-1	EVS1_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the Exchange Server service group (EVS1_SG1).
- 2 Click the **Link** button in the right pane.

- 3 Click the parent resource, for example EVS1_SG1-MountV.
- 4 Click the child resource, for example EVS1_RvgPrimary.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the Exchange Server service group (EVS1_SG1) online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the Exchange Server service group (EVS1_SG1).
- 2 In the right pane on the **Resources** tab, right-click the first RVG Primary resource (EVS1_RvgPrimary) and select **Online > SYSTEM1**.
- 3 In the right pane on the **Resources** tab, right-click any additional RVG Primary resource and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG

In the RVG service group, set up systems in the primary zone (Zone 0) to specify that initial fail over occurs to systems within the primary zone for the RVG service group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (EVS1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (Zone 0) for the primary zone.
- 7 Click **OK**.

Setting a dependency between the service groups

The RVG service group must be online on both the primary and secondary zones. However, if a failover occurs from one node to another within the same zone, the RVG service group must fail over along with the application service group.

To ensure that the Exchange Server service group and the RVG service group fail over and switch together, set up an online local hard dependency from the RVG service group to the Exchange Server service group.

The Exchange service group (for example, EVS1_SG1) is dependent on the replication service group (for example, EVS1_RVG_GRP).

To set up an online local hard dependency

- 1 From VCS Cluster Explorer, in the left pane, select the cluster (MYCLUSTER).
- 2 In the right pane, select the **Service Groups** tab.
- 3 Click the **Link** button to create a dependency between service groups.
- 4 Click the Exchange Server service group (the parent service group), for example EVS1_SG1.
- 5 Click the RVG service group (the child resource), for example EVS1_RVG_SG.
- 6 In the **Link Service Groups** window:
 - Select the **Relationship** of **online local**.
 - Select the **Dependency Type** of **hard**.
 - Click **OK**.

Adding the nodes from the secondary zone to the RDC

Configuration of the systems in the Primary Zone (Zone 0) is complete. The nodes in the Secondary Zone (Zone 1) can now be added to the RDC configuration.

Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

To add the nodes from the secondary zone to the RVG

- 1 From the active node of the cluster in the primary zone, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the **Welcome page**, and click **Next**.
- 3 In the **Wizard Options** dialog box:
 - Click **Modify an existing replication service group**. The existing replication service group is selected, by default.
 - Click **Next**.
- 4 If a VCS notice message appears, asking if you want to continue, click **Yes**.
- 5 Specify the system priority list:
 - In the **Available Cluster Systems** box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group's system list.
 To remove a node from the service group's system list, click the node in the **Systems in Priority Order** box, and click the left arrow icon.
 - 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.
 - Click **Next**.
- 6 If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
- 7 Review the Disk Group and Replicated Volume Group Configuration and click **Next**.
- 8 In the IP Resource Options dialog box, select **Modify IP resource** and click **Next**.
- 9 If a VCS error appears, click **OK**.
- 10 In the Network Configuration dialog box, verify that the selected adapters are correct and click **Next**.
- 11 Review the summary of the service group configuration.
 The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 12 Click **Next** to modify the replication service group. When prompted, click **Yes** to modify the service group.
- 13 Click **Finish**.

Configuring secondary zone nodes in the RVG service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (EVS1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Configuring the IP resources for fail over

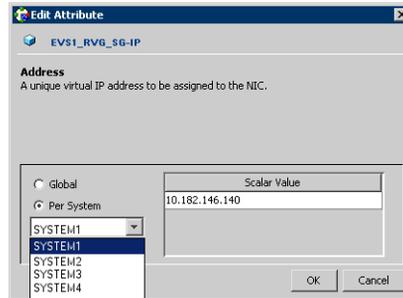
Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.

In the event of a system or Exchange failure, VCS attempts to fail over the Exchange Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone.

To modify the IP resources in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (EVS1_RVG_SG).
- 2 In the right pane, select the **Resources** tab.
- 3 Right-click the RVG IP resource (EVS1_RVG_SG-IP) and select **View > Properties View**.

- 4 In the Edit Attributes window, edit the Address attribute.



- **Select Per System.**
 - Select the first node in the primary zone and enter the virtual IP address for the primary zone.
 - Select the second node in the primary zone and enter the virtual IP address for the primary zone (the same IP address as the first node).
 - Repeat for all nodes in the primary zone.
 - Select the first node in the secondary zone (SYSTEM3) and enter the virtual IP address for the secondary zone.
 - Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - Repeat for all nodes in the secondary zone.
 - **Click OK.**
- 5 In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone have the same IP address, that is different from the IP address for the primary zone.
 - 6 Close the Properties View window.
 - 7 Since this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.

Adding the nodes from the secondary zone to the Exchange Server service group

Use the Exchange Server Configuration Wizard to add the nodes from the secondary zone to the Exchange Server service group.

To add the nodes from the secondary zone to the Exchange Server service group

- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Modify service group** option, select the Exchange service group (EVS1_SG1), and click **Next**.
- 4 On the Service Group Configuration panel, select the nodes in the secondary zone (Zone 1), from the Available Cluster Systems box select the systems and click the right-arrow icon to move the systems to the service group's system list.
- 5 Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.
- 6 On the Exchange Server Configuration panel, click **Next**.
- 7 On the Exchange Protocol Server panel, click **Next**.
- 8 On the Network Configuration panel, click **Next**.
- 9 On the Service Group Summary panel, review the service group configuration and click **Next**.
- 10 In the Completing the Exchange Configuration panel, clear the **Bring the service group online** check box and click **Finish**.

Configuring the zones in the Exchange Server service group

Specify Zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the Exchange Server service group

- 1 From VCS Cluster Explorer, in the left pane, select the Exchange Server service group (EVS1_SG1).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK** and close the Attributes View window.

Verifying the RDC configuration

After completing all the configuration tasks for the primary and secondary zones, you can bring the service group online, then verify the configuration.

Perform the following tasks:

- [Bringing the service group online](#)
- [Switching online nodes](#)

Bringing the service group online

After completing all configuration, ensure that the RVG service group is online in both the primary and secondary zone. Then you can bring the Exchange Server service group online in the primary zone.

To bring the Exchange service group online

- 1 From VCS Cluster Explorer, in the left pane, right-click the Exchange Server service group (EVS1_SG1).
- 2 Click **Online**.

Switching online nodes

Failover simulation is an important part of configuration testing. Test the failover by switching online nodes.

Note: This should never be tested on systems with live data. A reliable and tested backup should be available. A tested backup means that it has been tested successfully by a restore.

The RVG service group is online in both the primary and secondary zone. However, within a zone, if more than node is configured, the RVG service group should fail over with the application service group.

Switch the application service group between nodes using Veritas Cluster Manager (Java Console). When you complete the procedure, you will see the online system role shift from one system to another.

If you enter the system name manually from the Java Console, specify the name in upper case.

To switch online nodes

- 1 Open the Veritas Cluster Manager (Java Console) (**Start > All Programs > Veritas > Veritas Cluster Manager (Java Console)**).

- 2 Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, in the File menu, click **New Cluster**. In the **New Cluster - Connectivity Configuration** window, enter the computer name in the **Host name** field and click **OK**.
- 3 In the **Machinename - Login window**, enter your user name and password in the respective fields and click **OK**.
- 4 Right-click the service group in the left pane, and select an alternate system name from the **Switch To** entry.
- 5 In the **Question** dialog box, click **Yes** to confirm you do want to switch the service group to the other node.

Additional instructions for GCO disaster recovery

After completing the tasks for setting up a replicated data cluster for Exchange 2007, you can optionally create a secondary site for wide area disaster recovery using the SFW HA Global Cluster option (GCO).

With this option, 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.

To configure disaster recovery using a secondary site, you must install the SFW HA Global Cluster Option on all nodes on the primary (replicated data cluster) site cluster, as well as the secondary (DR) site cluster. GCO configuration also requires a static IP address available for each site.

You can use the Disaster Recovery (DR) wizard when setting up the secondary site. The secondary site is not configured as a replicated data cluster. There can be only one replicated data cluster in the DR environment. The DR wizard does the following tasks:

- Clones the storage
- Clones the application service group
- Sets up VVR replication for the secondary site
- Configures the primary and secondary site clusters as global clusters

See [“Deploying Disaster Recovery: New Exchange Server installation”](#) on page 551.

Disaster Recovery

This section contains the following chapters:

- [Chapter 12, “Disaster recovery for Exchange: Overview”](#) on page 549
- [Chapter 13, “Deploying Disaster Recovery: New Exchange Server installation”](#) on page 551
- [Chapter 14, “Deploying SFW HA for Disaster Recovery: Configuring any-to-any failover”](#) on page 655
- [Chapter 15, “Testing fault readiness by running a fire drill”](#) on page 677

Disaster recovery for Exchange: Overview

This chapter contains the following topics:

- [“What is a disaster recovery solution?”](#) on page 549
- [“Why implement a DR solution?”](#) on page 549
- [“Typical DR configurations for Exchange”](#) on page 550

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

You could implement any of the following DR configurations for Exchange:

- Using an active-passive configuration, create a new SFW HA environment with DR capabilities for Exchange on primary and secondary sites.
- 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.

Deploying Disaster Recovery: New Exchange Server installation

This chapter contains the following topics:

- [“Tasks for deploying a disaster recovery active-passive configuration of Microsoft Exchange”](#) on page 553
- [“Reviewing the requirements”](#) on page 556
- [“Reviewing the configuration”](#) on page 562
- [“Configuring the storage hardware and network”](#) on page 565
- [“Managing disk groups and volumes”](#) on page 567
- [“Setting up the secondary site: Installing SFW HA and configuring a cluster”](#) on page 569
- [“Verifying your primary site configuration”](#) on page 592
- [“Setting up your replication environment”](#) on page 593
- [“Assigning user privileges \(secure clusters only\)”](#) on page 600
- [“Configuring disaster recovery with the DR wizard”](#) on page 601
- [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 605
- [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 610
- [“Installing Exchange on the first node with DR option \(secondary site\)”](#) on page 613
- [“Installing Exchange on additional nodes \(secondary site\)”](#) on page 619

- [“Cloning the service group configuration on to the secondary site using the DR wizard”](#) on page 624
- [“Configuring replication and global clustering”](#) on page 627
- [“Verifying the disaster recovery configuration”](#) on page 643
- [“Establishing secure communication within the global cluster \(optional\)”](#) on page 645
- [“Adding multiple DR sites \(optional\)”](#) on page 647
- [“Recovery procedures for service group dependencies”](#) on page 647

Tasks for deploying a disaster recovery active-passive configuration of Microsoft Exchange

Before setting up disaster recovery at the secondary site, you must complete the high availability active-passive configuration on the primary site.

See [Chapter 4, “Deploying SFW HA for high availability: New installation” on page 45](#).

See [Chapter 5, “Deploying SFW HA for high availability: Standalone Exchange servers” on page 121](#).

You can also configure disaster recovery for a primary site that is configured as a replicated data cluster.

See [Chapter 11, “Configuring Replicated Data Clusters for Exchange” on page 439](#).

After setting up the SFW HA environment for Exchange on a primary site, you can create a secondary or “failover” site for disaster recovery.

The Disaster Recovery (DR) wizard helps you to clone the storage and service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

After service group configuration, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [Chapter 2, “Using the Solutions Configuration Center” on page 27](#).

Note: If you want to configure the secondary site manually, without using the DR wizard, see [Appendix A, “Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation” on page 705](#).

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

Table 13-1 Task list for a DR active-passive Exchange configuration

Objective	Tasks
“ Reviewing the requirements ” on page 556	Verifying hardware and software prerequisites and configuration requirements
“ Reviewing the configuration ” on page 562	<ul style="list-style-type: none"> ■ Understanding active-passive configuration and site failover in a DR environment ■ Understanding supported disaster recovery configurations for service group dependencies
“ Configuring the storage hardware and network ” on page 565	<ul style="list-style-type: none"> ■ Setting up the storage hardware 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 569	<ul style="list-style-type: none"> ■ Installing SFW HA ■ Configuring the cluster using the VCS Cluster Configuration Wizard (VCW)
“ Verifying your primary site configuration ” on page 592	Verifying that Exchange has been configured for high availability at the primary site
“ Setting up your replication environment ” on page 593	Ensuring replication prerequisites for your selected method of replication are met before running the DR wizard
“ Assigning user privileges (secure clusters only) ”	For a secure cluster only, assigning user privileges
“ Configuring disaster recovery with the DR wizard ” on page 601	<ul style="list-style-type: none"> ■ Reviewing prerequisites for the DR wizard ■ Starting the DR wizard and selecting a primary site system, the service group, the secondary site system, and the replication method
“ Cloning the storage on the secondary site using the DR wizard (VVR replication option) ” on page 605	(VVR replication option) Cloning the storage configuration on the secondary site
“ Creating temporary storage on the secondary site using the DR wizard (array-based replication) ” on page 610	(EMC SRDF, Hitachi TrueCopy, or GCO only replication option) Using the DR wizard to create temporary storage for application installation on the secondary site

Table 13-1 Tasklist for a DR active-passive Exchange configuration (Continued)

Objective	Tasks
“Installing Exchange on the first node with DR option (secondary site)” on page 613	<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
“Installing Exchange on additional nodes (secondary site)” on page 619	Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes in the same virtual server
“Cloning the service group configuration on to the secondary site using the DR wizard” on page 624	Cloning the service group configuration from the primary to the secondary site using the DR wizard
“Configuring replication and global clustering” on page 627	<ul style="list-style-type: none"> ■ (VVR replication) Using the wizard to configure replication and global clustering ■ (EMC SRDF replication) Setting up replication and then using the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Setting up replication and then using the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Using the wizard to configure global clustering, and then setting up replication
“Verifying the disaster recovery configuration” on page 643	Verifying the disaster recovery configuration
“Establishing secure communication within the global cluster (optional)” on page 645	Adding secure communication between local clusters within the global cluster (optional task)
“Adding multiple DR sites (optional)” on page 647	Optionally, adding additional DR sites to a VVR environment
“Recovery procedures for service group dependencies” on page 647	Reviewing actions required for disaster recovery if there are service group dependencies

Table 13-1 Task list for a DR active-passive Exchange configuration (Continued)

Objective	Tasks
“Possible task after creating the DR environment: Adding a new failover node to a VVR environment” on page 651	Completing required tasks when adding a new failover system to either the primary or secondary site in a VVR environment

Reviewing the requirements

This DR solution requires a primary site and secondary site.

Review the following installation and configuration requirements for your systems before SFW HA installation. Minimum requirements and Symantec recommended requirements may vary.

- “Disk space requirements” on page 556
- “Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:” on page 556

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space.

Table 13-2 estimates disk space requirements for SFW HA.

Table 13-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

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

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/302144>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agent for Microsoft Exchange.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2003 versions

The following table lists the Microsoft Exchange Server 2003 versions supported with SFW HA 5.1 Service Pack 1.

Table 13-4 Supported Microsoft Exchange Server 2003 versions

Exchange Server 2003	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Memory must be a minimum 256 MB of RAM per server for Exchange 2003; refer to the Microsoft documentation for more information.
- 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. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs. See "[Best practices for SFW HA](#)" on page 561.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each Exchange Virtual Server (EVS).

- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone 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 Dynamic DNS (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 for SFW HA

The following permissions are required:

- 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 on 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 Exchange virtual server computer object 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, 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 for SFW HA

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 Server 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 2003.
- 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.

Best practices for SFW HA

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.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

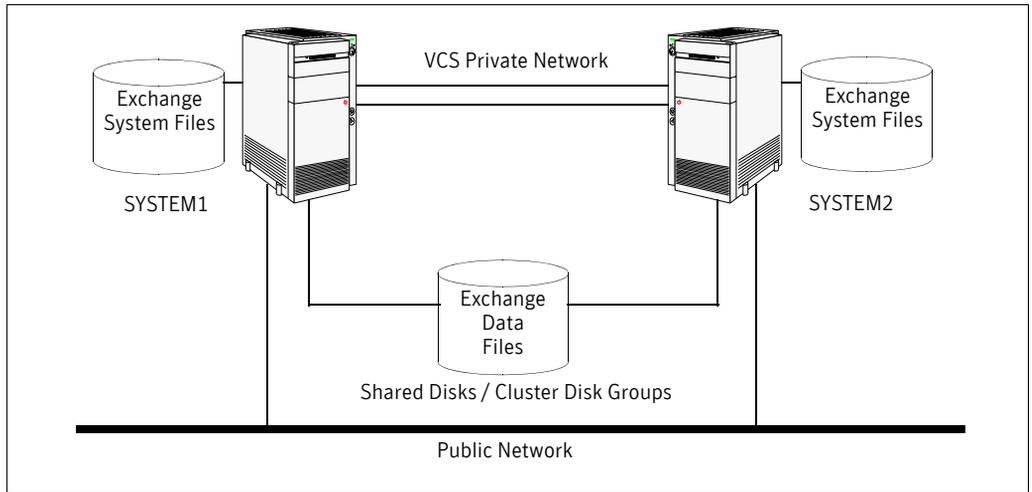
Reviewing the configuration

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.

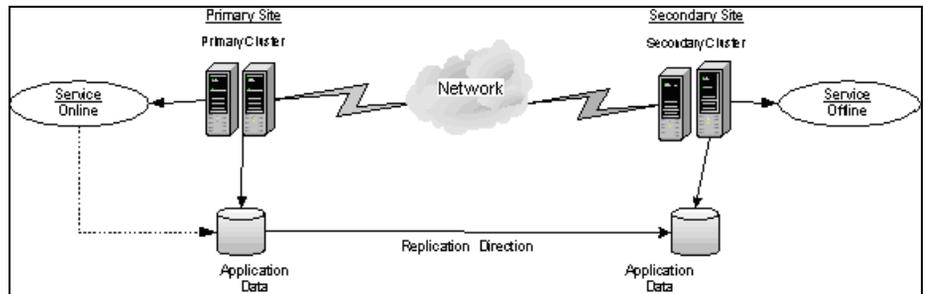
[Figure 13-1](#) provides a view of a cluster configuration on the primary site:

Figure 13-1 Cluster configuration on the primary site



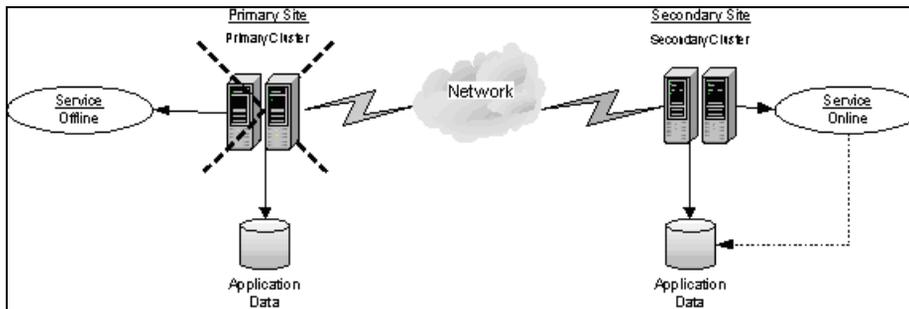
In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. [Figure 13-2](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 13-2 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 13-3](#) illustrates this type of failure:

Figure 13-3 Application services restored after primary site failure



You can choose to configure replication using VVR or an agent-supported array-based hardware replication. You can use the DR wizard to configure VVR replication or required options for the VCS agents for EMC SRDF or Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

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 and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network 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 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, 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

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - *Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

When installing the application on multiple nodes, you need to manage the disk groups and volumes.

Managing disks groups and volumes includes the following procedures:

- Importing a disk group and mounting a 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.

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.

Before you begin, review the requirements.

- See “[Reviewing the requirements](#)” on page 556.

Installing SFW HA

The product installer enables you to install the software for Veritas Storage Foundation HA 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, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

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

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 13-6 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 571.

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 Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The Select Product screen appears.
- 3 Review the links on the Select Product screen.
Links on this screen access Late Breaking News, the Configuration Checker, as well as begin the process to install Storage Foundation HA for Windows. Click on **Read Late Breaking News** for the latest information about updates, patches, and software issues regarding this release.
- 4 Click **Storage Foundation HA 5.1 SP1 for Windows**.
- 5 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 6 Review the Welcome message and the listed prerequisites. Ensure that any of the listed prerequisites applicable to your installation are met prior to proceeding.

Click **Next**.

- 7 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 AGREE TO the terms of the license agreement**, and then click **Next**.
- 8 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.
- 9 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.
- 10 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. Be sure to select the following as appropriate for your installation.

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, select the option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

- 11 Select the following 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.

When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

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

- 12 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.
- 13 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.

14 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages.

If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning:

The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If applicable to your installation, perform the above procedure.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

If applicable to your installation, perform the above procedure.

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation 5.1 for Windows (Server Components) with the MSCS option, you may want MSCS Quorum Arbitration Time (Min. and Max) to be adjusted to ensure optimal functionality with Veritas dynamic volumes with MSCS.

For additional information, see the *Storage Foundation for Windows Administrator Guide* for details.

If applicable to your installation, perform the above procedure.

15 When finished reviewing the message or messages, click **OK**.

16 The Summary screen appears displaying an Install report. Review the information in the Install report. Click **Back** to make changes, if necessary. Click **Install** if information is validated.

17 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.

18 When the installation completes, review the summary screen and click **Next**.

- 19 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.
- 20 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 21 Review the log files and click **Finish**.
- 22 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 Cluster Configuration Wizard (VCW) 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.
- When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. 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. This account does not require domain admin privileges.
- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

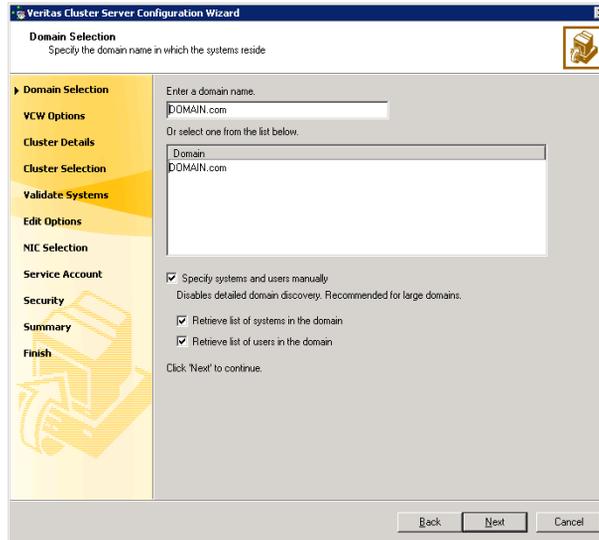
Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

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



Do one of the following:

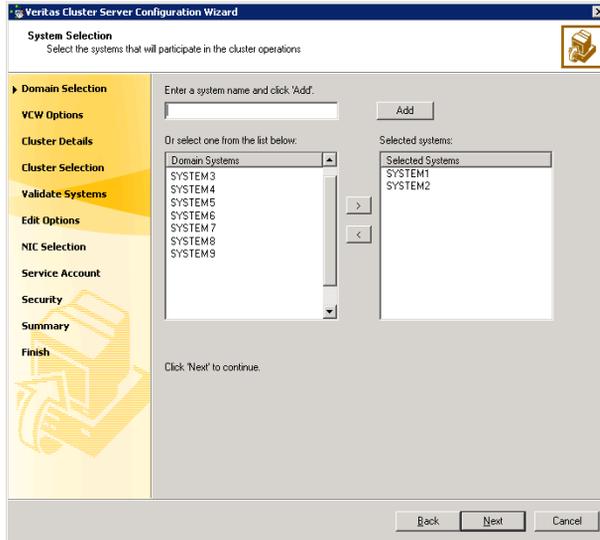
- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 578.
- 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 578. 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 8](#) on page 578.

- 6 On the System Selection panel, specify the systems for the 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 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

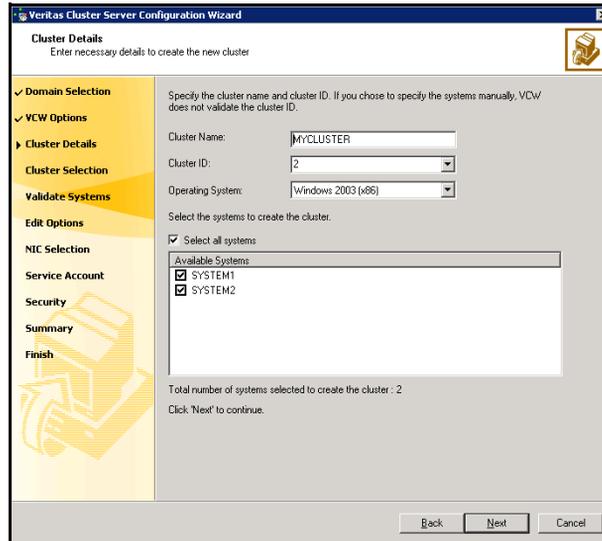
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.

- 9 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. The cluster ID can be any number from 0 to 255.

Caution: 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. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (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.

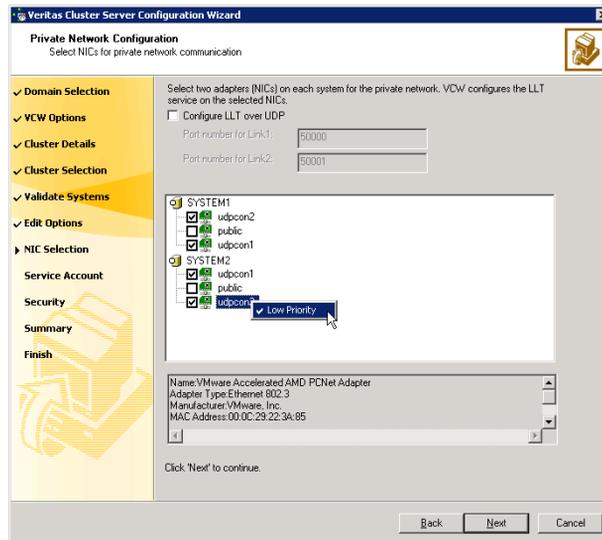
10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

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 the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 582.

11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



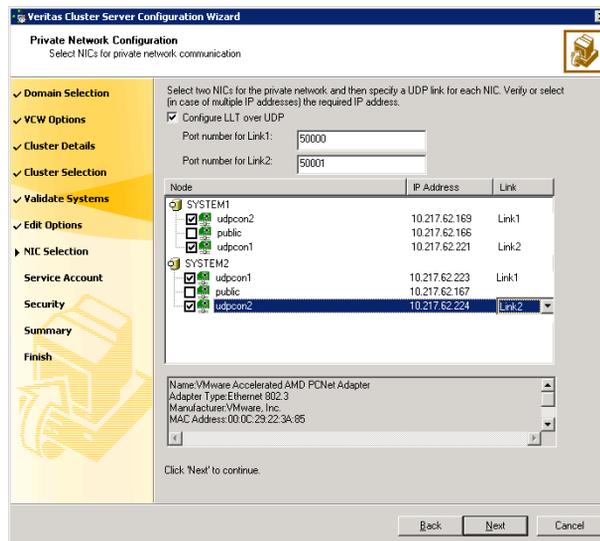
- Select the check boxes next to the two NICs to be assigned to the private network.

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 both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as 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.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



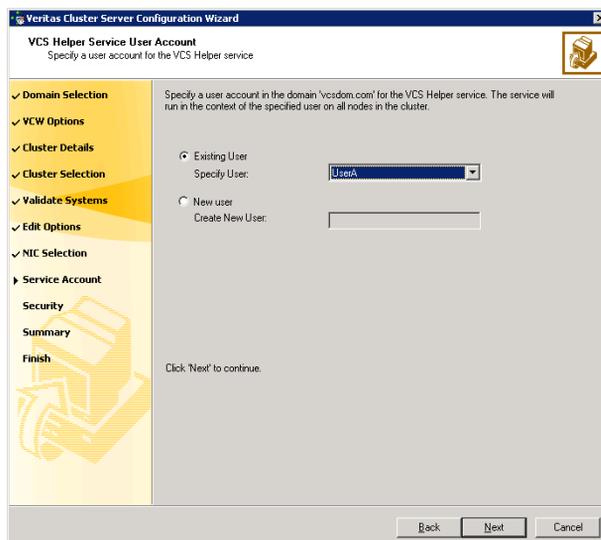
- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to

65535. The default ports numbers are 50000 and 50001 respectively.

- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.



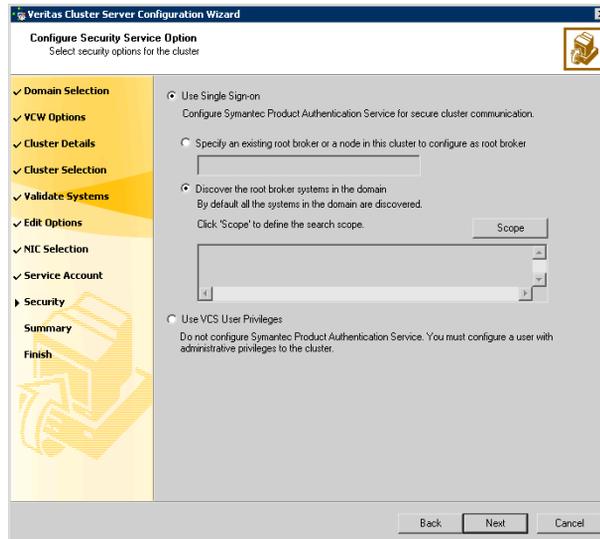
Specify a domain user as follows:

- 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 577, 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**.

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

Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability 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 an existing root broker or a node in this cluster to configure as root broker**, 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 then configures all nodes in the cluster as authentication brokers.

- If you want to search 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 a user *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**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 13-7 contains some more examples of search criteria.

Table 13-7 Search criteria examples

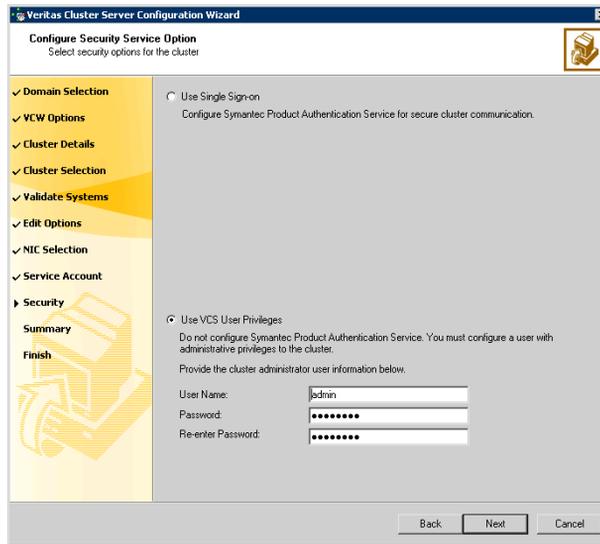
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- 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 a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS Encrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password. After the cluster is configured, you can 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**.
- 14 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.

- 15 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**.

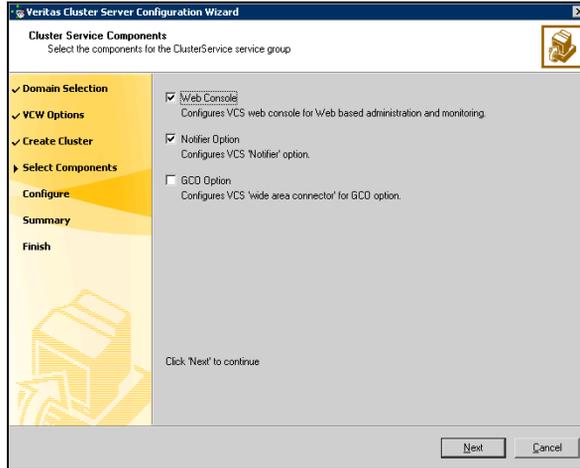
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

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.

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



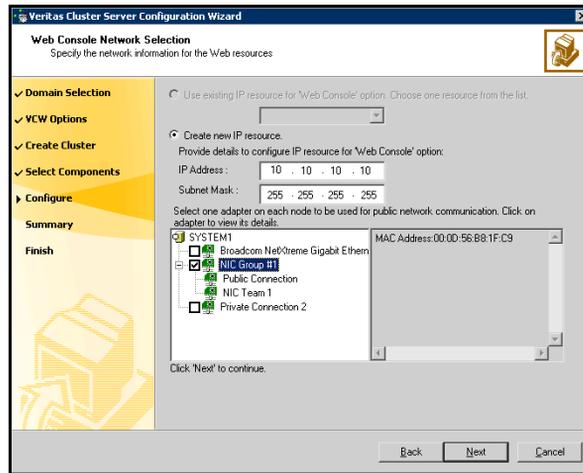
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 588.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 589.

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**.



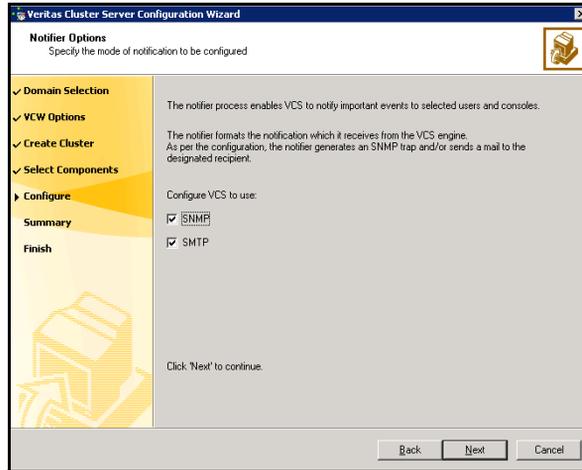
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 589. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

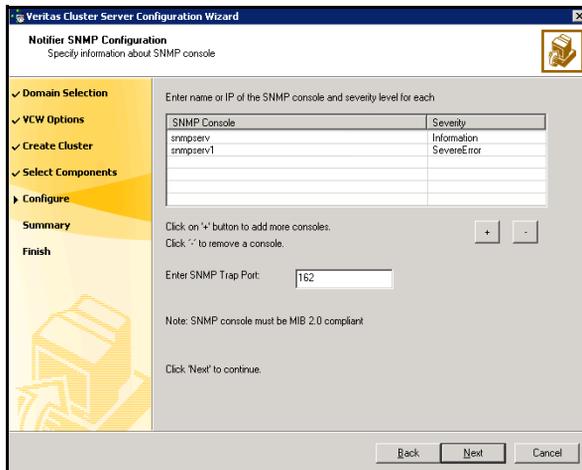
To configure notification

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

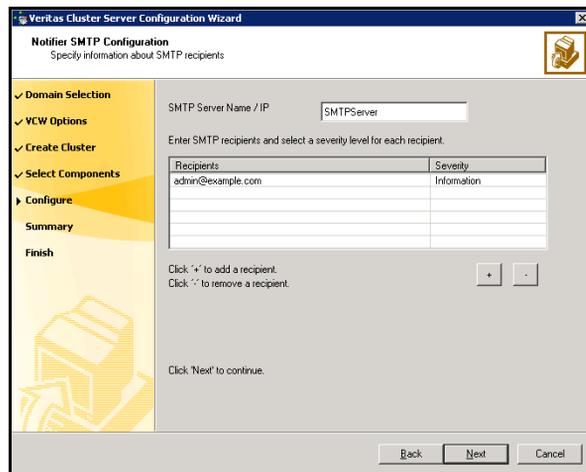


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

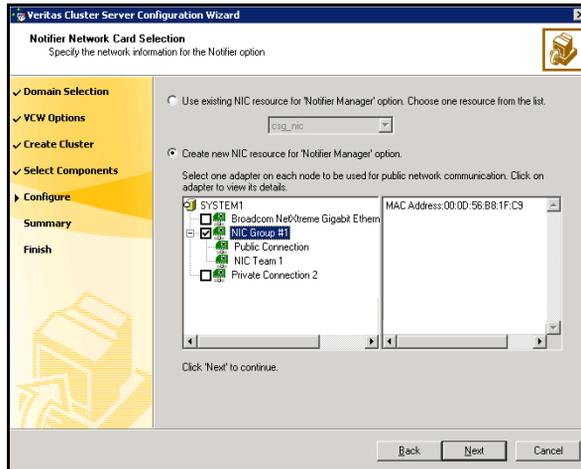


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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

See [Section 2, “High Availability” on page 41](#).

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.

Note: If you are setting up a replicated data cluster at the primary site, use the replicated data cluster instructions rather than the high availability configuration steps in the Solutions Configuration Center. See [Chapter 11, “Configuring Replicated Data Clusters for Exchange” on page 439](#).

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Veritas Volume Replicator (VVR)
- EMC SRDF
- Hitachi TrueCopy

For array-based hardware replication, you can use any replication agent supported by Veritas Cluster Server. The DR wizard can help with configuring the methods listed above. If you choose a different replication method, you must run the wizard first to complete configuring global clustering; then afterwards, you configure replication separately.

See [“Configuring global clustering only” on page 641](#).

Before configuring replication with the wizard, ensure that you set up the replication environment prerequisites. Choose from the following topics, depending on which replication method you are using:

- [“Setting up security for VVR” on page 593](#)
- [“Requirements for EMC SRDF array-based hardware replication” on page 596](#)
- [“Requirements for Hitachi TrueCopy array-based hardware replication” on page 598](#)

Setting up security for VVR

If you are using Veritas Volume Replicator (VVR) replication, you must configure the VVR Security Service (VxSAS) on all cluster nodes on both the primary and secondary sites.

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 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 `vxascfg.exe` from the command prompt of the required machine.

Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information panel as follows:

Account name Enter the administrative account name.
(domain\account)

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts. Click **Next**.

- 3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

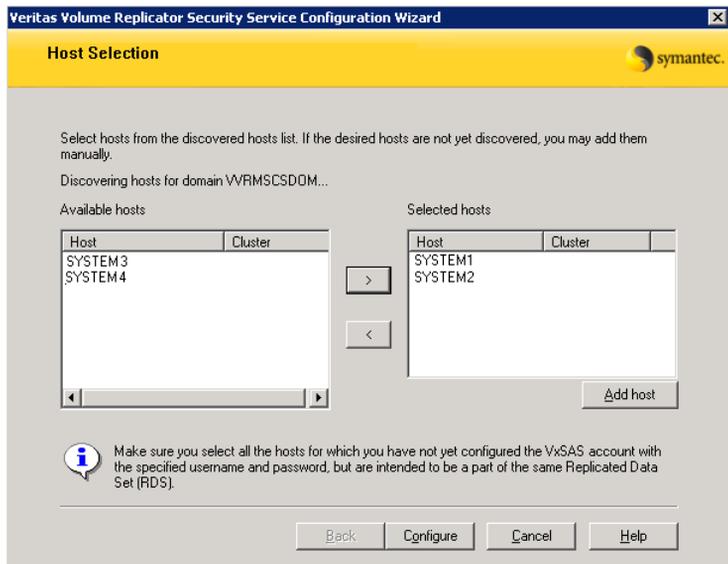
Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

- 4 On the Host Selection panel, select the required hosts:



Selecting hosts The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host 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, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with 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.

Requirements for EMC SRDF array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for EMC SRDF. The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also configures the Symm heartbeat. Optional resource settings are left in the default state.

For more information about the EMC SRDF agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for configuring EMC SRDF”](#) on page 596
- [“Replication requirements for EMC SRDF”](#) on page 596

Software requirements for configuring EMC SRDF

The EMC SRDF agent supports SYMCLI versions that EMC recommends for the firmware on the array. The agent supports SRDF on all microcode levels on all Symmetrix arrays, provided that the host/HBA/array combination is in EMC's hardware compatibility list.

To use the DR wizard to configure the required agent settings for EMC SRDF, ensure that the following software requirements are met:

- The EMC Solutions Enabler is installed on all cluster nodes.
- The SYMCLI version that is installed supports the generation of XML output.
- The SYMCLI version and the microcode level support dynamic swapping.
- The VCS EMC SRDF agent is installed on all cluster nodes.

Replication requirements for EMC SRDF

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that no devices are RDF2.
- On the secondary site, the wizard verifies that no devices are RDF1.

Otherwise, the wizard displays an invalid configuration message and is unable to proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All disks in SFW disk groups must belong to the same device group.
- The device group must not span more than one array (no composite device groups).
- A device group can contain one or more disk groups.
- Dynamic swap must be enabled on both sites.
- On the primary site:
 - All devices must be RDF1 and part of an RDF1 device group.
 - Devices must have write access.
- On the secondary site:
 - All devices must be RDF2 and part of an RDF2 device group.
 - Write access must be disabled.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the SRDF resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the SRDF resource, not to the array configuration. However, the SRDF resource will be unable to come online in the service group until replication has been configured correctly.

In addition, note the following agent requirement:

- Device group configuration must be the same on all nodes of the cluster.

Requirements for Hitachi TrueCopy array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for Hitachi TrueCopy. The wizard configures the required settings for the HTC resource in the VCS application service group. Optional settings are left in the default state.

For more information about the Hitachi TrueCopy agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for Hitachi TrueCopy”](#) on page 598
- [“Replication requirements for Hitachi TrueCopy”](#) on page 598

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

For details, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

To use the DR wizard to configure the required agent settings for Hitachi TrueCopy, ensure that the following requirements are met:

- RAID Manager is installed in the same location on all nodes on a site.
- Enter the primary and secondary site file paths for the horcm files on the Hitachi TrueCopy Path Information panel in the wizard. The default location is:
`System Driver\Windows`
- The horcm files are named `horcmnn.conf` (where *nn* is a positive number without a leading zero, for example, `horcm1.conf` but not `horcm01.conf`).

Replication requirements for Hitachi TrueCopy

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that all devices are the same type, but not S-SWS or SSUS.
- On the secondary site, the wizard verifies that all devices are the same type, but not P-VOL or PSUS.

Otherwise, the wizard displays an invalid configuration message and does not proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All configured instances are running.
- No disks in the SFW disk group span across the Device Group.
- A device group can contain one or more disk groups.
- The device group does not span more than one array.
- At the primary site, all devices are of the type P-VOL.
- At the secondary site, all devices are of the type S-VOL.
- All device groups at the primary site are paired to an IP address which must be online on the secondary node.
- Device group and device names include only alphanumeric characters or the underscore character.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the HTC resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the HTC resource, not to the array configuration. However, the HTC resource will be unable to come online in the service group until replication has been configured correctly.

Assigning user privileges (secure clusters only)

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`

Configuring disaster recovery with the DR wizard

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (VVR replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure VVR replication, or configure the VCS hardware replication agent settings for EMC SRDF or Hitachi TrueCopy
- Configure global clustering

Warning: To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment that is not configured by the wizard, you must first run the wizard to configure global clustering before configuring replication.

You will need to exit the wizard after the storage cloning task to install the Exchange application. The wizard allows you to exit after the logical completion of each task.

Each time you re-start the wizard, you specify the primary site system, service group, secondary site system, and replication method, as described in the following procedure. Clicking **Next** then takes you to the start page of the process following the one that you had last completed.

The DR Wizard list of service groups shows only those that contain a MountV resource. For a dependent service group to be listed, the parent service group must also contain a MountV resource.

Warning: Once you have completed configuring replication and global clustering with the DR wizard, you cannot use the wizard to change the method of replication.

Before running the DR wizard to configure disaster recovery, ensure that you meet the following prerequisites:

- SFW HA is installed and a cluster is configured at the secondary site. Ensure that the name assigned to the secondary site cluster is different than the name assigned to the primary site cluster.
- Your application or server role is configured for HA at the primary site and all required services are running at the primary site.

- The clusters taking part in the DR configuration should have distinct names.
- Enough free disk space is available at the secondary site to duplicate the storage configuration at the primary site.
- One static IP address is available per application service group to be cloned.
- If using VVR for replication, a minimum of one static IP address per site is available for each application instance running in the cluster.
- Global Cluster Option (GCO) is installed at the primary and secondary site, and one static IP address is available at each site for configuring GCO.
- A VCS user is configured with the same name and privileges in each cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall is set to allow both ingoing and outgoing TCP requests on port 7419.

Note: The DR wizard does not support VVR configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- [“Setting up security for VVR”](#) on page 593
- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 596
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 598

To start configuring disaster recovery with the DR wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Server tab and 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 In the Welcome panel, review the prerequisites to ensure that they are met 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 virtual server is online at the primary site, you can also specify `localhost` to connect to the system.

Click **Next**.

4 In the Service Group Selection panel, select the service group 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. Only online and local dependencies are supported, in soft, firm, or hard configurations. The wizard can configure only one level of dependency. In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported.

The panel lists only service groups that contain a MountV resource.

Click **Next**.

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 In the Replication Options panel, select the replication method. Although you must select the replication method now, configuring replication and the global cluster option is done later, after service group cloning.

Configure Veritas Volume Replicator (VVR) and the Global Cluster Option (GCO)	<p>Select this option if you want to configure VVR replication.</p> <p>Select this option even if you plan to configure VVR replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a VVR environment.</p> <p>The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p> <p>You cannot mix replication methods. That is, if your primary site is using array-based replication, and you select the VVR option, the wizard will warn you that you cannot use VVR replication for the disaster recovery site.</p>
Configure EMC SRDF and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS EMC SRDF agent. All disks used for the service group on the primary site must belong to an EMC SRDF array.</p> <p>Select this option even if you plan to configure EMC SRDF replication or the GCO option manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>
Configure Hitachi TrueCopy and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS Hitachi TrueCopy agent. All disks used for the service group on the primary site must belong to a Hitachi TrueCopy array.</p> <p>Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>

Configure the Global Cluster Option (GCO) only If you select this option, the DR wizard does not configure any replication settings. It configures the global cluster option.

Select this option if you want to use the wizard in an array-based replication environment that is not supported by this wizard. You must configure replication manually after you finish the wizard.

If you select the GCO only option, the DR wizard sets up the storage and service group configuration on the secondary site for an array-based hardware replication environment. Therefore, you cannot use this option to clone the storage and service group for a VVR replication environment.

Click **Next**.

- 7 Continue with the next DR configuration task.
For VVR replication, see [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 605.
For array-based replication, see [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 610.

Cloning the storage on the secondary site using the DR wizard (VVR replication option)

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at the secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, the wizard can correct this and then complete the configuration.

If you have not yet started the wizard, see the following topic before continuing with the storage cloning procedure:

- [“Configuring disaster recovery with the DR wizard”](#) on page 601.

To clone the storage configuration from the primary site to the secondary site (VVR replication method)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system. In the Replication Options panel, select the VVR replication method and click **Next**.

- 2 Review the information in the Storage Validation Results panel. This panel compares the configuration at the secondary site with that on the primary. If the storage is already configured identically on both sites, the panel shows that results are identical. Otherwise, the panel shows the differences and recommended actions. You can toggle between a summary and detailed view of information about the differences.

The detailed view shows the following:

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.
Mount	Displays the mount to be assigned the volume on the secondary site.
Recommended Action	Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary. <ul style="list-style-type: none">■ If the volume does not exist, a new volume will be created.■ If the volume exists but is of a smaller size than that on the primary, the volume will be expanded to the required size.■ If the volume is of a greater size than that on the primary, the volume will be recreated using the appropriate size.■ If the volume is the same as that on the primary, the message indicates that the volumes are identical and no action is required.

The summary view shows the following:

Disk groups that do not exist	Displays the names of any disk groups that exist on the primary but do not exist on the secondary.
Existing disk groups that need modification	Displays the names of any disk groups on the secondary that need to be modified to match the primary.
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 primary site configuration on the secondary, you

can free some disks on the secondary or add more storage. Then click **Refresh/Validate** to have the wizard update its information about the secondary storage configuration.

You continue with the wizard to provide information for the recommended actions. Before proceeding to the service group configuration, the wizard ensures that the configuration of the disk groups and volumes for the service group is the same at the primary and secondary site.

Click **Next**.

- 3 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.
	Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures.

Click **Next**.

- 4 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 disk name or the >> option to move the disks 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**.

- 5 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.
- 6 In the Implementation panel, wait until 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**.
- 7 In the Storage Cloning Configuration Result screen, view the results and click **Next**.
- 8 In the Exchange Installation panel, review the information. Do one of the following:
 - Click **Finish** to exit the wizard and proceed with installing the application on the required nodes on the secondary site. Before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.
 - If the DR wizard is run from a remote node, you can keep the wizard running on that node. You can then install the application locally on each of the required nodes.
 - If you are running the DR wizard from a local system and need to install the Exchange application on that system, the system gets restarted when the application installation is complete. You can then restart the wizard.
 - Click **Next** to continue with service group cloning if the application is already installed on the required nodes.

When restarting the Disaster Recovery wizard, continue through the wizard from the Welcome panel, specifying the primary site system, the service group, the secondary site system, and the replication method. The wizard

proceeds to the storage cloning panel. If it detects that the storage is identical on the secondary site, it proceeds to the next task.

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

To enable you to install applications, the DR wizard can create a temporary disk group, DR_APP_INSTALL_DG, which contains the volumes and mount points for use in application installation. The temporary configuration uses 500 MB volumes or the volume size at the primary site, depending on which is smaller. The wizard deletes the temporary configuration after application installation.

If you have already installed the application on all nodes, you can skip this storage cloning step by unchecking the Perform storage cloning check box on the Storage Cloning panel.

If you are starting the wizard for the first time, see the following topic before continuing with the storage cloning procedure:

- [“Configuring disaster recovery with the DR wizard”](#) on page 601.

To create temporary storage for application installation (array-based replication)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system.
- 2 In the Replication Options panel, select the array-based replication method you plan to use and click **Next**:
 - EMC SRDF
 - Hitachi TrueCopy
 - Global Cluster Option only (select if you are using another agent-supported array-based replication method)
- 3 If you selected Hitachi TrueCopy replication, the Hitachi TrueCopy File Paths panel is displayed. The wizard populates the fields if it locates the files in the default location. Otherwise, fill in the file path information for both the primary and secondary sites as follows:

RAID Manager bin path	Path to the RAID Manager Command Line interface Default: C:\HORCM\etc where C is the system drive.
-----------------------	--

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

HORCM files location	Path to the horcm configuration files (horcm nn .conf) Default: C:\Windows where C is the system drive An horcm configuration file is required by the RAID Manager on all nodes; however the wizard does not validate this.
----------------------	--

- 4 In the Storage Cloning panel, choose one of the following:
 - If you have not yet installed the application on all nodes, leave **Perform storage cloning** checked and click **Next**. Continue with the next step in this procedure.
 - If you have already installed the application on all nodes, uncheck **Perform storage cloning** and click **Next**. Continue with the procedure for service group cloning.

- 5 The Storage Validation Results panel shows the temporary storage configuration that the wizard will configure at the secondary site. You can click **Show Summary** to toggle to a summary view and toggle back to a detailed view by clicking **Show Details**.

The detailed view shows the following:

Disk Group	Displays the name of the single disk group required on the secondary site for temporary storage: DR_APP_INSTALL_DG
Volume	Displays the list of volumes required at the secondary site.
Size	Displays the size of the volumes required on the secondary site.
Mount	Displays the mounts required at the secondary site.
Recommended Action	Indicates the action that the wizard will take at the secondary site.

The summary view shows the following:

Existing configuration	Displays the existing secondary configuration.
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 on the secondary are inadequate, you can free some disks on the secondary or add more storage. Then click **Refresh/Validate** so that the wizard can update its information about the secondary storage configuration.

Click **Next**.

- 6 In the Disk Selection for Storage Cloning panel, a default disk selection is shown for the temporary storage at the secondary site. You can change the selection by moving disks to and from the Available Disks and Selected Disks pane. Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures. Click **Next**.
- 7 The Volume Layout for Secondary Site Storage panel shows a default volume layout for the temporary storage based on the primary site volume layout. Optionally, you can change the default disk assignment and layout for any volume:

Disk Group	Displays the DR_APP_INSTALL__DG disk group.
Volume (Volume Size)	Displays the name and the size of the volume to be created on the secondary.
Available Disks	Displays the disks that are available for the volumes. To select a disk, either double-click on the host name or click the >> button to move the hosts into the Selected Disks pane.
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 selected for the volume. To remove a disk from the list, select it and click the << button.
View Primary Layout	Displays the volume layout at the primary site.

Click **Next**.

- 8 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 temporary storage configuration at the secondary site.
- 9 In the Implementation panel, wait until 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**.

- 10 In the Storage Configuration Cloning Result screen, view the results and click **Next**.
- 11 In the Exchange Installation panel, review the information and do one of the following:
 - Before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.
 - If you are running the DR Wizard from a local system and need to install the Exchange application on that system, click **Finish** to exit the wizard and proceed with installing the application on the required nodes.

After completing the application installation, you can launch the DR Wizard again to proceed with service group cloning. At this point the temporary cloned storage is no longer needed. Before beginning service group cloning, the wizard displays the Temporary Storage Deletion panel to confirm the deletion of the temporary storage.
 - If the DR Wizard is run from a remote node, you can keep the wizard running on that node. You can then install the Exchange application locally on each of the required nodes.

After completing the application installation, click **Next** to proceed with service group cloning. At this point the temporary cloned storage is no longer needed. Before beginning service group cloning, the wizard displays the Temporary Storage Deletion panel to confirm the deletion of the temporary storage.

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, Exchange virtual server EVS1 will fail over from SYSTEM4 to SYSTEM5.

See the following topics:

- [“Prerequisites for installing Exchange Server”](#) on page 614

- [“Exchange pre-installation on first node \(secondary site\)”](#) on page 615
- [“Exchange installation on first node \(secondary site\)”](#) on page 617
- [“Exchange post-installation on first node \(secondary site\)”](#) on page 618

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 567.
- Mount the volume containing the information for registry replication.
- 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 binaries to be installed at the same location on all the nodes. For example, if you install Exchange at drive C of one node, installations on all other nodes must occur on drive C.
- Make sure to use the same drive letters employed on the primary site.
- Set the required permissions:
 - You must be a domain user.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - 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.

- Either the logged-on user or the VCS Helper service user account 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 take the Exchange service group offline on the primary site; otherwise, the wizard will prompt you to take the service group offline.
- Verify that you have completed the forest and domain preparation. See “[Preparing the forest and domain](#)” on page 58

Exchange pre-installation on first node (secondary site)

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.

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 Tools > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Available Option panel, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 Specify the name or IP address of the cluster node at the primary site and click **Next**.
- 6 Select the name of the Exchange Virtual Server for which you want to set up a secondary site and click **Next**.
- 7 In the Select Option panel, choose the **Create a failover node for Exchange disaster recovery setup** option and click **Next**.
- 8 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.

- 9 Specify information related to the network.
 - Verify the virtual server name and domain suffix.
 - 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. 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**.
- 10 Select a drive where the registry replication data will be stored and click **Next**.
- 11 Review the summary of your selections and click **Next**.
- 12 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 13 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**.
- 14 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.

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

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 615.

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

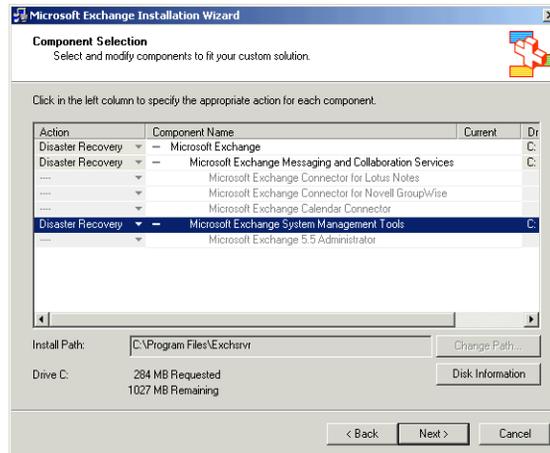
The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:
`C:\>hasys -state`
The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.

- 8 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 [“Before installing Veritas Storage Foundation High Availability for Windows \(SFW HA\), ensure that you review the following:”](#) on page 556.
- 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 567.

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.

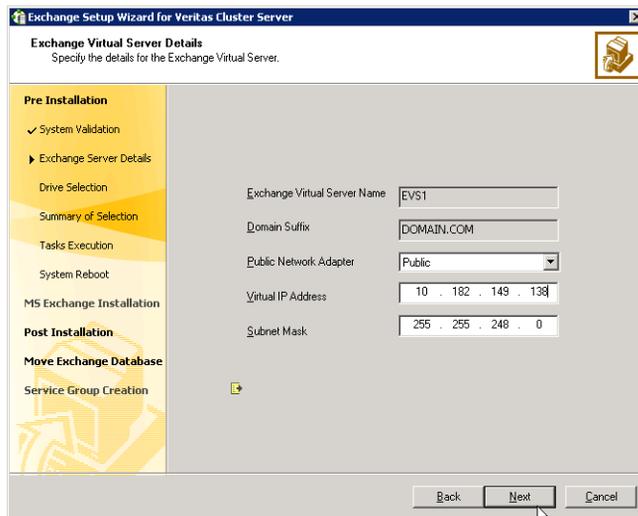
Exchange pre-installation: Additional nodes

Use the following procedure to perform Exchange pre-installation on additional nodes for the same EVS.

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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: Additional nodes

Install Exchange on the same node selected in “[Installing Exchange on additional nodes \(secondary site\)](#)” on page 619.

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

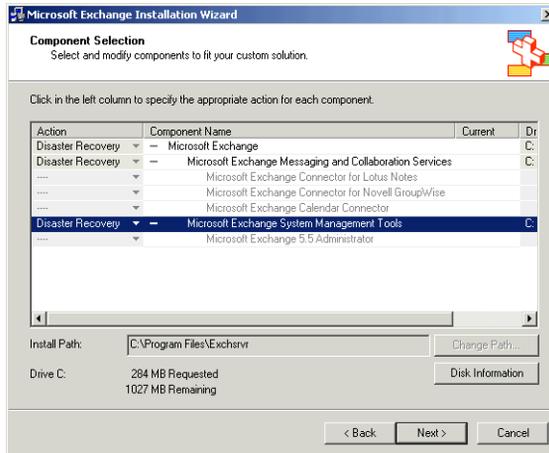
The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```

The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```

```
C:\>net start had
```
- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

If you wish to add more failover nodes later, use the Exchange service group configuration wizard. For information on adding a failover node to a VVR environment, see [“Possible task after creating the DR environment: Adding a new failover node to a VVR environment”](#) on page 651.

Cloning the service group configuration on to the secondary site using the DR wizard

Before cloning a service group on the secondary site, verify the following:

- On the secondary site, the application is installed.
- If the secondary site is in a different subnet, ensure that a static IP address is available on the secondary site to assign to the virtual server.

If you are launching the wizard for the first time, see the following topic for additional information:

- [“Configuring disaster recovery with the DR wizard”](#) on page 601.

Note: Although you can view the cloning progress in the VCS Java Console, do not save and close the configuration while cloning is in progress. Otherwise, the cloning fails and you have to delete the service group on the secondary site and run the wizard again.

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**. Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 In the Welcome panel, click **Next** and continue through the wizard, providing the requested information for the primary site system, the service group, the secondary site system, and the replication method.
If you selected the VVR replication method, the wizard proceeds to the storage cloning task and notifies you if it detects that the storage is identical. Click **Next** until you reach the Service Group Analysis panel.
If you selected an array-based replication method (EMC SRDF, HTC, or GCO only), the temporary storage is no longer needed once the application is installed and the wizard confirms whether or not to delete it.
- 4 (Array-based replication method only) In the Temporary Storage Deletion panel, confirm whether or not to delete the cloned storage:

- If the application is already installed on the required nodes, leave **Delete cloned storage** checked and click **Next**. When the wizard prompts you to confirm deleting the shared storage, click **Yes**.
 - If you want to delete the cloned storage manually later, uncheck **Delete cloned storage** and click **Next**.
- 5 (Array-based replication method only) If you selected to delete the cloned storage, the wizard shows the progress of the tasks in the Implementation panel. If the storage deletion fails, the wizard will show a failure summary page. Otherwise, when it shows the tasks are complete, click **Next**.
 - 6 Review the following information displayed in the Service Group Analysis 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.

Service Group Details on the Primary Displays the resource attributes for the service group at the primary site. These include:

- Cluster
- IP Resource: consists of the IP address and the subnet mask
 - NIC Resource: is the MAC address

Service Group Details on the Secondary Cluster Displays a message to indicate whether the service group or the corresponding attributes have been configured at the secondary site.

- 7 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	<p>Displays a list of available systems on the secondary cluster that are not yet selected for service group cloning.</p> <p>Select any additional secondary systems on which you want the wizard to clone the application service group configuration.</p> <p>Either double-click on the system name or use the > option to move the hosts into the Selected Systems pane.</p> <p>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.</p>
Selected Systems	<p>Displays the list of selected systems. The secondary system that you selected earlier in the wizard is listed by default.</p>

Click **Next**.

- 8 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	<p>Displays the list of resources that exist on the primary cluster.</p>
Attribute Name	<p>Displays the attribute name associated with each of the resources displayed in the Resource Name column.</p> <p>If you need to edit additional attributes that are not shown, you must edit them manually on the secondary site service group once service group cloning is complete.</p>
Primary Cluster	<p>Displays the primary attribute values for each of the displayed attributes.</p>
Secondary Cluster	<p>The default is the same as the primary cluster. The same virtual IP address can be used if both sites exist on the same network segment. You can specify different attributes depending on your environment. For the MACAddress attribute select the appropriate public NIC from the drop-down list.</p>

Click **Next**.

- 9 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 site.

- 10 In the Implementation panel, wait until all the tasks are completed. The progress bar indicates the status of the tasks. Successful tasks are marked with a check symbol. If some task could not be completed successfully, the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**
- 11 If the cloning failed, review the troubleshooting information. Otherwise, click **Next** to continue with the replication and GCO configuration, or with GCO only, depending on which option you selected.
Optionally, you can exit the wizard at this point and launch the wizard again later. When you launch the wizard again, continue through the wizard, specifying the primary site system, the service group, the secondary site system, and the replication method. Click **Next** to continue to the replication and/or GCO configuration task.

Configuring replication and global clustering

After creating the identical service group configuration on both sites, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication.

If you are using an array-based replication that is not supported by the wizard, you configure global clustering only. In this case, you must complete configuring global clustering before configuring replication.

The following topics cover the steps required for each replication method:

- [“Configuring VVR replication and global clustering”](#) on page 627
- [“Configuring EMC SRDF replication and global clustering”](#) on page 635
- [“Configuring Hitachi TrueCopy replication and global clustering”](#) on page 638
- [“Configuring global clustering only”](#) on page 641

Configuring VVR replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure VVR replication and global clustering.

Note: By default, in an Exchange or SQL Server environment, the DR wizard configures all the volumes in a disk group under one Replicated Volume Group (RVG). If you require a different organization, you should configure it using the Veritas Enterprise Administrator (VEA) rather than the DR wizard. For information on setting up VVR replication with the VEA, see [Appendix A, “Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation”](#) on page 705.

You can then return to the wizard to configure global clustering.

Before you begin, ensure that you have met the following prerequisites:

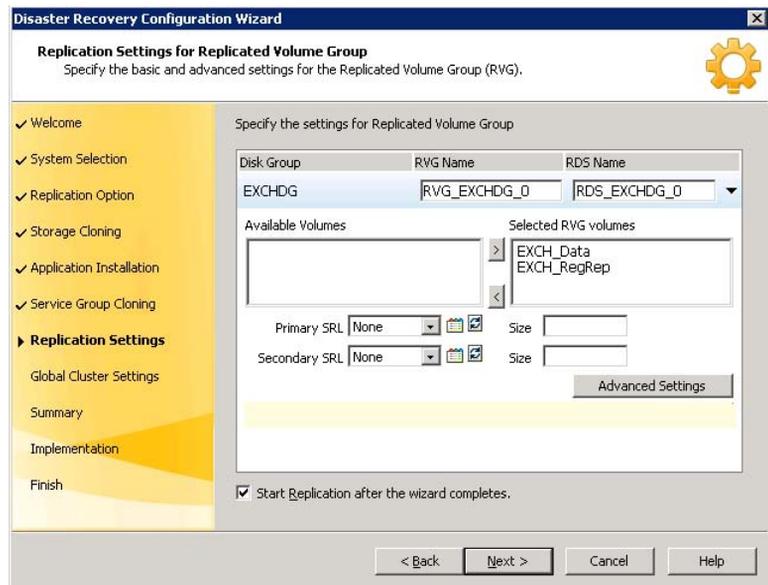
- Ensure that Veritas Volume Replicator is installed at the primary and secondary site.
- Ensure that Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- Ensure that VVR Security Service (VxSAS) is configured at the primary and secondary site. See the following topic:
 - “[Setting up security for VVR](#)” on page 593
- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that, if using secure clusters, you configure a VCS user with the same name and privileges in each cluster.

Use the following procedure to configure VVR replication and global clustering with the DR wizard.

To configure VVR replication and GCO

- 1 Verify that the application server service group is online at the primary site and the appropriate disk groups are imported at the secondary site.
- 2 If the wizard is still open after the previous wizard task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.

- On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - On the Replication Methods panel, click **Configure VVR and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**. If not, click **Cancel** and restart the wizard after meeting the requirements.
 - 4 In the Replication Settings for Replicated Volume Group panel, specify the requested information. If you are adding a DR site to an existing DR configuration, fields that must match the existing settings, such as the RVG or RDS name, are dimmed so that you cannot change them.

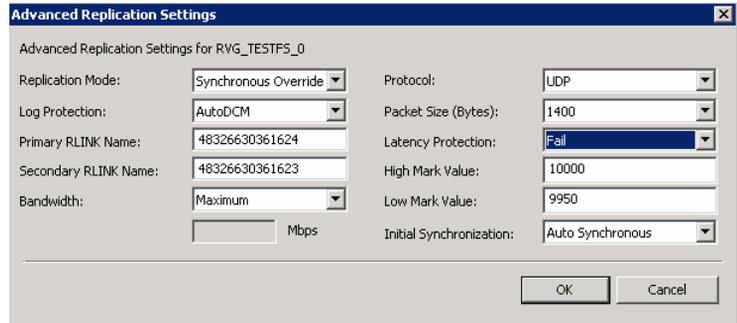


Disk Group	The left column lists the disk groups. By design, an RVG is created for each disk group.
RVG Name	Displays the default RVG name. If required, change this to a name of your choice.

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	If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the name, size, and disk. Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.
Secondary SRL	If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the same name and size as you specified for the primary SRL. Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.
Start Replication after the wizard completes	Select this check box to start replication automatically after the wizard completes the necessary configurations. Once replication is configured and running, deselecting the checkbox does not stop replication.

- 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 **AutoDCM** is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **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	<p>By default, VVR replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit.</p> <p>The default unit is Mega bits per second (Mbps) and the minimum allowed value is 1 Mbps.</p>
Protocol	Choose TCP or UDP. UDP/IP is the default replication protocol.
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	<p>By default, latency protection is set to Off.</p> <p>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.</p> <p>This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail option when the Secondary is connected.</p>
High Mark Value	<p>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.</p> <p>To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.</p>
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 is 9950.

Initial Synchronization

If you are doing an initial setup, then use the **Auto Synchronous** 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.

To apply changes to advanced settings, click **OK**. On the Replication Settings for Replicated Volume Group panel click **Next**.

- 5 In the Replication Attribute Settings panel, specify required replication attribute information for the cluster at the primary and secondary site. Click the arrow icon to expand an RVG row and display the replication attribute fields. If you are configuring an additional secondary site (multiple DR sites), some fields are disabled.

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 above network settings to any additional RVGs that are listed on this screen. If there is only one RVG, this option does not apply.

After specifying the replication attributes for each of the RVGs, click **Next**.

- 6 In 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. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

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. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters. If you have a printer installed, you can click the printer icon at the bottom of the scrollable list to print the settings.
Click **Next** to implement the settings.
- 8 In the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for global clustering. If a task could not be completed successfully, it is marked with an (x) symbol. For any critical errors, the wizard displays an error message. For less critical errors, the Information column displays a brief description

about the task failure and the next screen displays additional information on what action you can take to remedy it. Click **Next**.

- 9 In the Finish panel, review the displayed information. If a 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.

Configuring EMC SRDF replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the SRDF resource in the application service group.

Ensure that you have met the prerequisites for replication. See the following topic:

- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 596

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings. See the following topic:

- [“Optional settings for EMC SRDF”](#) on page 637

To configure EMC SRDF replication and GCO

- 1 Verify that you have brought the application service group online at the primary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Servertab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.

- In the Replication Methods panel, click **Configure EMC SRDF and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.

- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the SRDF resource cannot come online in the service group.

- 4 In the SRDF Resource Configuration panel, the wizard populates the required resource fields if replication has been configured. Otherwise, you must enter the required resource settings manually.

Symmetrix Array ID (SID) Specify the array ID for the primary site and for the secondary site.

Device Group name Specify the name of the Symmetrix device group that contains the disks of the disk group for the selected instance.

Available VMDG Resources Select the disk groups associated with the selected application instance.

- 5 If you want to configure an additional SRDF resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In 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. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

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. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

Click **Next**.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a 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.
- 10 Proceed with configuring additional optional settings for the SRDF resource if desired, and then verifying the disaster recovery configuration.

Optional settings for EMC SRDF

The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also detects and configures the SymHome attribute.

Other settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for EMC*

SRDF, Configuration Guide. If you change any settings, ensure that you edit the resource on both the primary and secondary sites.

The wizard also detects and configures the SymHome attribute.

The optional settings use the following defaults:

Option	Default setting
DevFOTime	2 seconds per device required for a device to fail over
AutoTakeover	The default is 1; the agent performs a read-write enable on partitioned devices in the write-disabled state during a failover, if devices are consistent.
SplitTakeover	The default is 1; the agent brings service groups online on the R2 side even if the devices are in the split state because they are read-write enabled.

Configuring Hitachi TrueCopy replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the HTC resource in the application service group.

Ensure that you have met the prerequisites. See the following topic:

- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 598

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings. It uses defaults for optional settings. See the following topic:

- [“Optional settings for HTC”](#)

To configure Hitachi TrueCopy replication and GCO

- 1 Verify that you have brought the application server service group online at the primary site
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server**

> **Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the HTC resource cannot come online in the service group.

- 4 In the HTC Resource Configuration panel, the wizard populates the required resource fields if the horcm file is configured properly. If not, you can configure the horcm file and click **Refresh** to populate the fields. Alternatively, enter the required resource settings manually:

Instance ID	Specify the instance number of the device group. Multiple device groups may have the same instance number.
Device Group name	Specify the name of the Hitachi device group that contains the disk group for the selected instance. The device group name must be the same on both the primary and secondary sites.
Available VMDG Resources	Select the disk groups associated with the selected application instance.
Add, Remove, Reset buttons	Click Add or Remove to display empty fields so that you can manually add or remove additional resources. Click Refresh to repopulate all fields from the current horcm file.

- 5 If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In 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. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

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. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a 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.

- 10 Proceed with configuring additional optional settings for the HTC resource if desired, and then verifying the disaster recovery configuration.

Optional settings for HTC

The wizard configures the required settings for the HTC resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

The optional settings use the following defaults:

Option	Default setting
LinkMonitor	The default is 0; the agent does not periodically attempt to resynchronize the S-VOL side if the replication link is disconnected. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the pairresync command.
SplitTakeover	The default is 0; the agent does not permit a failover to S-VOL devices if the replication link is disconnected; that is, if P-VOL devices are in the PSUE state.

Configuring global clustering only

If you are using a replication method that the DR wizard does not configure, you must select the replication option to configure global clustering only.

For the GCO only option, you use the wizard to complete all DR tasks except the replication configuration task. You must complete the final wizard task of configuring global clustering before configuring replication.

Before configuring GCO:

- Ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- If you created secure clusters at the primary site and secondary site, ensure that you have configured a VCS user with the same name and privileges in each cluster, and the user must be added in the Administrator role.

The following procedure assumes that you have completed the earlier wizard tasks through the service group cloning task and are continuing with the final step of configuring global clustering.

To configure GCO only

- 1 If the wizard is still open after the service group cloning task, continue with the GCO Setup panel. Otherwise, launch the wizard and proceed to the GCO Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for MicrosoftExchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Global Cluster Option (GCO) only**. Click **Next** and continue to the GCO Setup panel.
- 2 In the GCO Setup panel, review the requirements. If you have met the requirements, click **Next**.
- 3 In 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. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

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. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes.

Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 4 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified. Click **Next**.
- 5 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (x) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 6 In the Finish panel, review the displayed information. If a 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

The steps you need to take to verify your DR configuration depend on the type of replication you are using.

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- For VVR replication, 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.
- For VVR replication:
 - 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 VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.
- Confirm that the RVG service groups are online at the primary and secondary sites.
- Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- For array-based replication, verify that the required array resource is created in the primary and secondary cluster's application service group and that a dependency is set between the VMDg resource and the array resource.
- For EMC SRDF replication, verify that the SRDF resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, verify that the HTC resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, you must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the disk groups online. This must be performed only once, after which the failover works uninterrupted. For more information, see *Veritas™ Cluster Server Hardware Replication Agent for Hitachi TrueCopy Installation and Configuration Guide*.
- 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 configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- 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 for VVR-based replication.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. 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.

Adding multiple DR sites (optional)

In a Veritas Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. Veritas Cluster Server supports up to four DR sites. In other replication environments, additional DR sites require manual configuration.

Run the DR wizard and on the Secondary System selection panel, select the new site.

Before you start the wizard on the task of configuring replication and global clustering, ensure that the cluster service group is online at the existing primary and secondary sites. This enables the wizard to configure GCO not only between the selected primary site and the new secondary site but also between the new site and the earlier configured secondary site. Otherwise, the wizard displays a warning message after the global clustering task.

When configuring the VVR replication settings with the wizard for the additional site, fields that must match existing settings are dimmed so that you cannot change them. For example, you cannot change the RVG name or RVG layout on the Replication Settings panel. Similarly, on the Global Cluster Settings panel, GCO has already been configured at the primary site, so the primary site fields are dimmed.

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 565.

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, used for VVR replication, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, soft dependency link.

Table 13-8 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 13-9 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. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Leave the RVG group online at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ 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. 	<p>1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.</p> <p>2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p>

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, hard dependency link.

Table 13-10 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. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Do not take the RVG group offline at the primary site.</p>

Table 13-10 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 to a VVR environment

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 651
- [“Preparing the existing DR environment”](#) on page 651
- [“Installing Exchange on the new node”](#) on page 652
- [“Modifying the replication and Exchange service groups”](#) on page 652
- [“Reversing replication direction”](#) on page 653

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 569 for installation instructions.
- 2 Use the **Cluster Operations** option of the VCS Configuration wizard (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.
 - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.

- Click **Switch To**, and click **Remote switch**.
 - In the **Switch global group** dialog box:
 - Click the cluster at the secondary site you want to switch the group to.
 - Click the specific system where you want to bring the global 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 567.
- 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 619; 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 > Symantec >**

- Veritas Cluster Server > Configuration Tools > 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 > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server 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 651, 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:
 - Click the cluster to switch the group to.
 - Click the specific system where you want to bring the global Exchange service group online.
 - Click **OK**.

Possible task after creating the DR environment: Adding a new failover node to a VVR environment

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

This chapter contains the following topics:

- [“Tasks for deploying a disaster recovery any-to-any configuration of Microsoft Exchange”](#) on page 656
- [“Reviewing the configuration”](#) on page 658
- [“Configuring disaster recovery for the first Exchange virtual server”](#) on page 662
- [“Verifying your primary site configuration for an additional Exchange virtual server”](#) on page 663
- [“Adding the user to the service group \(secure clusters only\)”](#) on page 663
- [“Configuring disaster recovery for the second Exchange virtual server”](#) on page 664
- [“Cloning the storage on the secondary site using the DR wizard”](#) on page 664
- [“Installing Exchange on the first node of an additional EVS \(secondary site\)”](#) on page 665
- [“Specifying a common node for failover”](#) on page 670
- [“Cloning the service group configuration on to the secondary site using the DR wizard”](#) on page 672

- [“Configuring replication and global clustering”](#) on page 672
- [“Verifying the disaster recovery configuration”](#) on page 672
- [“Establishing secure communication within the global cluster \(optional\)”](#) on page 674
- [“Possible tasks after creating the DR environment”](#) on page 676

Tasks for deploying a disaster recovery any-to-any configuration of Microsoft Exchange

You can set up disaster recovery for a high availability environment that involves an any-to-any configuration with multiple failover nodes.

Before setting up disaster recovery at the secondary site, you must complete the high availability any-to-any configuration on the primary site.

See [Chapter 6, “Deploying SFW HA for high availability: Configuring a new any-to-any failover”](#) on page 199 .

See [Chapter 7, “Deploying SFW HA for high availability: Converting an existing installation to any-to-any failover”](#) on page 289.

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.

The Disaster Recovery (DR) wizard helps you to clone the storage configuration and the service group configuration from the primary site to the secondary site. After service group configuration, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 27.

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

Table 14-1 Tasks for deploying a new any-to-any DR configuration

Objective	Tasks
“ Reviewing the configuration ” on page 658	Understanding any-to-any configuration and site failover in a DR environment
“ Configuring disaster recovery for the first Exchange virtual server ” on page 662	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 663	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 663	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 664	Cloning the storage configuration on the secondary site
“ Installing Exchange on the first node of an additional EVS (secondary site) ” on page 665	<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 672	Cloning the service group configuration from the primary to the secondary site using the DR wizard

Table 14-1 Tasks for deploying a new any-to-any DR configuration

Objective	Tasks
“Configuring replication and global clustering” on page 672	<ul style="list-style-type: none"> ■ (VVR replication) Using the wizard to configure replication and global clustering ■ (EMC SRDF replication) Setting up replication and then using the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Setting up replication and then using the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Using the wizard to configure global clustering, and then setting up replication
“Verifying the disaster recovery configuration” on page 672	Verifying the disaster recovery configuration
“Establishing secure communication within the global cluster (optional)” on page 674	Adding secure communication between local clusters within the global cluster (optional task)
“Possible tasks after creating the DR environment” on page 676	<ul style="list-style-type: none"> ■ 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:

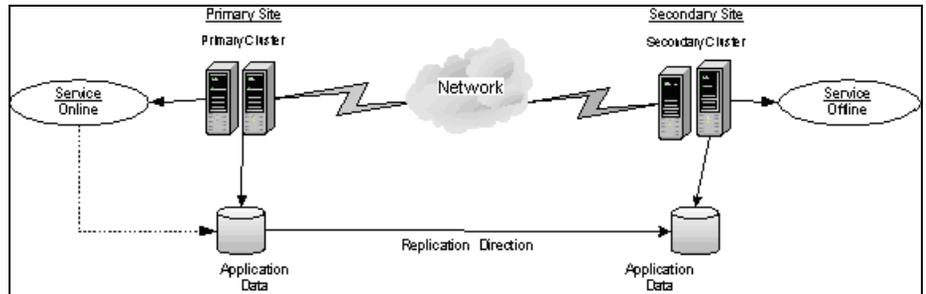
- [“Disaster recovery configuration”](#) on page 658
- [“Any-to-any configuration”](#) on page 660
- [“Sample any-to-any configuration for disaster recovery”](#) on page 661

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 14-1](#) displays an

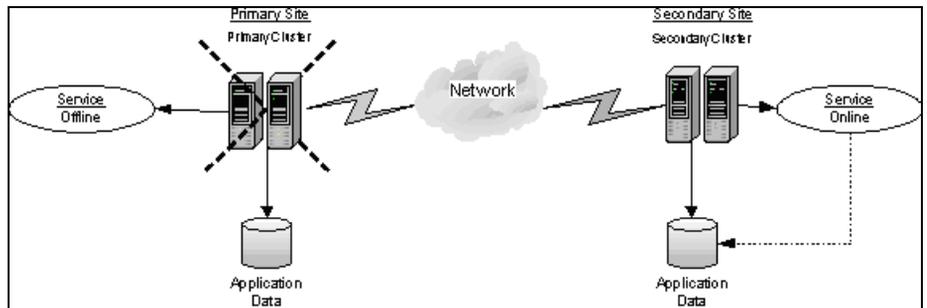
environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 14-1 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 14-2](#) illustrates this type of failure:

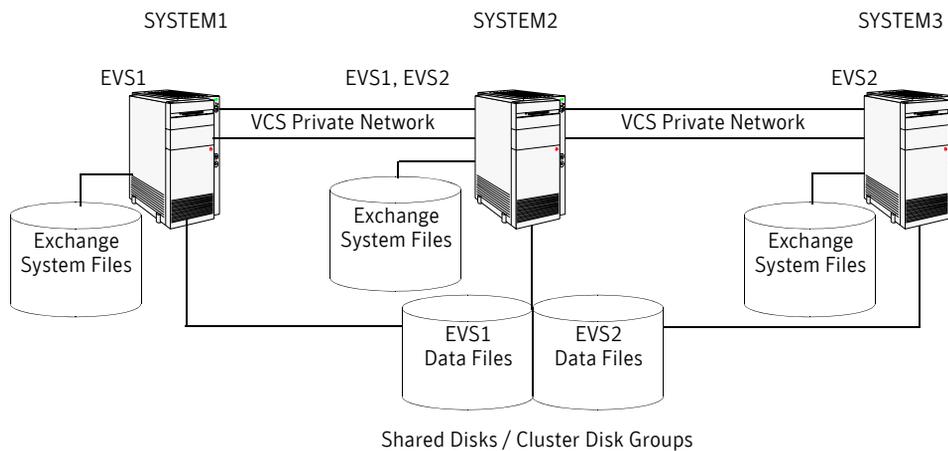
Figure 14-2 Application services restored after primary site failure



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 14-3](#) show an example of an any-to-any configuration.

Figure 14-3 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 14-2 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 14-3 Sample Configuration

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3 (Primary Site) SYSTEM4, SYSTEM5, SYSTEM6 (Secondary Site)	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_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 14-3 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 13, “Deploying Disaster Recovery: New Exchange Server installation”](#) on page 551.

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 569.
- 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 575.
- 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 593.

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 266 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 663.

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. You clone the storage separately for each EVS.

To clone the storage for an additional EVS, you follow the same storage cloning procedure as when you configured EVS1. The storage cloning procedure you follow depends on whether you are using Veritas Volume Replicator or an array-based hardware replication.

See the following topics in [Chapter 13, “Deploying Disaster Recovery: New Exchange Server installation”](#):

- [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 605
- [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 610

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 567.
- Mount the volume containing the information for registry replication.
- 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 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - 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.
 - Either the logged-on user or the VCS Helper service user account 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 prompt you to take the service group offline.
- Verify that you have completed the forest and domain preparation. See “[Preparing the forest and domain](#)” on page 58

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 Tools > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Available Option panel, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 Specify the name or IP address of the cluster node at the primary site and click **Next**.
- 6 Select the name of the Exchange Virtual Server for which you want to set up a secondary site and click **Next**.
- 7 In the Select Option panel, choose the **Create a failover node for Exchange disaster recovery setup** option and click **Next**.
- 8 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.
- 9 Specify information related to the network.
 - Verify the virtual server name and domain suffix.

- 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. 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**.

- 10 Select a drive where the registry replication data will be stored and click **Next**.
- 11 Review the summary of your selections and click **Next**.
- 12 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 13 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**.
- 14 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.

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

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 666.

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

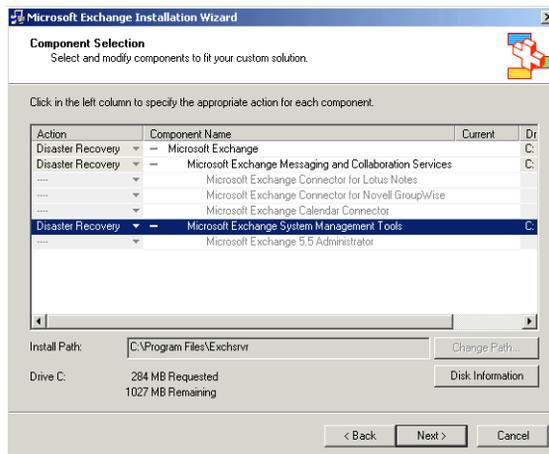
The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
 Type the following on the command line:
`C:\>hasys -state`
 The state should display as **RUNNING**.
 If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
 If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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 662. 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.

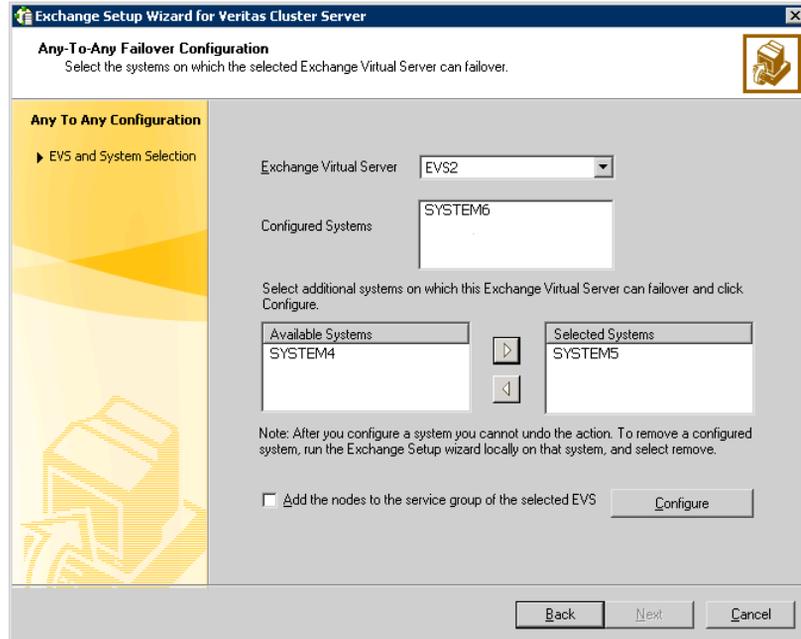
In the example, EVS1 is already configured with SYSTEM5 as a failover node, so you launch this wizard from EVS2 only.

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 Tools > 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

The DR wizard enables you to clone the service group configuration present at the primary site on to the secondary site. You clone the service group separately for each EVS.

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.

To clone the service group for an additional EVS, you follow the same procedure as when you cloned the service group for EVS1.

See the following topic in [Chapter 13, “Deploying Disaster Recovery: New Exchange Server installation”](#):

- [“Cloning the service group configuration on to the secondary site using the DR wizard”](#) on page 624

Configuring replication and global clustering

To configure replication and global clustering for an additional EVS, you follow the same procedure as when you configured replication and global clustering for EVS1.

See the following topic in [Chapter 13, “Deploying Disaster Recovery: New Exchange Server installation”](#):

- [“Configuring replication and global clustering”](#) on page 627

Verifying the disaster recovery configuration

The steps you need to take to verify your DR configuration depend on the type of replication you are using.

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- For VVR replication, 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.
- For VVR replication:
 - 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 VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.
 - Confirm that the RVG service groups are online at the primary and secondary sites.
 - Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- For array-based replication, verify that the required array resource is created in the primary and secondary cluster's application service group and that a dependency is set between the VMDg resource and the array resource.
- For EMC SRDF replication, verify that the SRDF resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, verify that the HTC resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, you must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the disk groups online. This must be performed only once, after which the failover works uninterrupted. For more information, see *Veritas™ Cluster Server Hardware Replication Agent for Hitachi TrueCopy Installation and Configuration Guide*.

- 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 configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- 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 for VVR-based replication.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. 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 647.
- Adding a new failover node for a VVR environment
See [“Possible task after creating the DR environment: Adding a new failover node to a VVR environment”](#) on page 651.

Testing fault readiness by running a fire drill

Topics in this chapter include:

- [About disaster recovery fire drills](#)
- [About the Fire Drill Wizard](#)
- [About post-fire drill scripts](#)
- [Tasks for configuring and running fire drills](#)
- [Prerequisites for a fire drill](#)
- [Preparing the fire drill configuration](#)
- [Running a fire drill](#)
- [Recreating a fire drill configuration that has changed](#)
- [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.

A fire drill is performed at the secondary site using a special service group for fire drills. The fire drill service group uses a copy of the data that is used by the application service group.

About the Fire Drill Wizard

Veritas Storage Foundation HA for Windows (SFW HA) provides a Fire Drill Wizard to help you set up and run a fire drill on a disaster recovery environment. You launch the Fire Drill Wizard from the Solutions Configuration Center.

The Fire Drill Wizard tests the fault readiness of a disaster recovery 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 Fire Drill Wizard supports conducting a fire drill for a disaster recovery site that uses Veritas Volume Replicator (VVR) or that uses Hitachi TrueCopy or EMC SRDF hardware replication.

About Fire Drill Wizard general operations

The Fire Drill Wizard performs the following operations:

- Prepares for the fire drill by creating a fire drill service group on the secondary site
The fire drill service group is a copy of the application service group. When creating the fire drill service group, the wizard uses the application service group name, with the suffix `_fd`. The Exchange fire drill service group contains only the VMDg and mountV resources, and in the case of hardware replication, the HTCSnap or SRDFSnap resource. The wizard renames the fire drill service group resources with a prefix `FDnn` and changes attribute values as necessary to refer to the FD resources.
The wizard also supports fire drill service groups created under a different naming convention by an earlier version of the wizard.
- Runs the fire drill by bringing the fire drill service group online on the secondary site
Optionally the wizard runs Eseutil to check for data consistency as part of the fire drill test. The fire drill tests the replication and consistency of the data to verify that the data will be available if the Exchange service group fails over and comes 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.

Note: The fire drill service group for Exchange contains only the data resources, not the application, so that the Exchange application does not itself come online during a fire drill.

- Restores the fire drill configuration, taking the fire drill service group offline

After you complete the fire drill, you run the wizard to restore the fire drill configuration to a prepared state. Otherwise, the fire drill service group remains online. If you run a fire drill on one service group, restore that service group before you continue with a fire drill on another service group. You must also restore the fire drill configuration before you can delete it.

Warning: 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. Therefore, always use the wizard to restore the fire drill configuration to a prepared state as soon as possible after completing the fire drill testing for a service group.

See “[Restoring the fire drill system to a prepared state](#)” on page 699.

- Deletes the fire drill configuration

The details of some Fire Drill Wizard operations are different depending on the replication environment.

See “[About Fire Drill Wizard operations in a VVR environment](#)” on page 679.

See “[About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment](#)” on page 681.

About Fire Drill Wizard operations in a VVR environment

The general operations of the Fire Drill Wizard are the same in all replication environments.

- Prepares for the fire drill, creating a fire drill service group on the secondary site
- Runs the fire drill, bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration, taking the fire drill service group offline
- Deletes the fire drill configuration

See “[About the Fire Drill Wizard](#)” on page 678.

However, the following additional Fire Drill Wizard operations are specific to a Veritas Volume Replicator (VVR) environment.

Preparing the fire drill configuration

In a VVR environment, when preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, replaces the RVGPrimary resources with VMDg resources
- Uses the SFW HA VxSnap feature to prepare snapshot mirrors for use during the fire drill
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.

Running the fire drill

In a VVR environment, when running the fire drill, the wizard does the following:

- Detaches the mirrors from the original volumes to create point-in-time snapshots of the production data
- Creates a fire drill disk group on the secondary site with a snapshot of the application data to use for testing purposes

Restoring the fire drill configuration

In a VVR environment, when restoring the fire drill system to a prepared state, the wizard does the following:

- 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 to the application service group disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

Deleting the fire drill configuration

In a VVR environment, when deleting the fire drill configuration, the wizard does the following:

- Deletes the fire drill service group and any associated registry entry
- Performs the snap abort operation on the snapshot mirrors to free up the disk space

About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment

The general operations of the Fire Drill Wizard are the same in all replication environments.

- Prepares for the fire drill, creating a fire drill service group on the secondary site
- Runs the fire drill, bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration, taking the fire drill service group offline
- Deletes the fire drill configuration

See [“About the Fire Drill Wizard”](#) on page 678.

However, additional Fire Drill Wizard operations are specific to a Hitachi TrueCopy or EMC SRDF replication environment.

In a Hitachi TrueCopy or EMC SRDF replication environment, the wizard performs the following additional actions during preparation, running of the fire drill, restoring the configuration, and deleting the configuration. You must configure the ShadowImage (for Hitachi) or BCV (for SRDF) pairs before running the wizard.

Preparing the fire drill configuration

When preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, the wizard creates HTCSnap or SRDFSnap resources for each HTC and SRDF resource in the application service group. The SRDFSnap and HTCSnap resources from the fire drill service group are linked to the respective resources configured in the main application service group.
- The wizard configures the Snap resource. The following Snap resource attributes are set to a value of 1:
 - UseSnapshot (take a local snapshot of the target array)
 - RequireSnapshot (require a successful snapshot for the Snap resource to come online)
 - MountSnapshot (use the snapshot to bring the fire drill service group online)

Running the fire drill

When running the fire drill, the wizard brings the HTCSnap or SRDFSnap agent online. The HTCSnap or SRDFSnap agent manage the replication and mirroring functionality according to the attribute settings. The Snap agents take a consistent snapshot of the replicating data using the mirroring technology provided by the array vendor. The Snap agents also import the disk group present on the snapshot devices with a different name.

In more detail, the Snap agent does the following:

- Suspends replication to get a consistent snapshot
- For HTCSnap, takes a snapshot of the replicating application data on a ShadowImage device
- For SRDFSnap, takes a snapshot of the replicating application data on a BCV device
- Resumes replication
- Modifies the disk group name in the snapshot

Restoring the fire drill configuration

When restoring the fire drill configuration to a prepared state, the wizard does the following:

- Takes the fire drill service group offline, thus also taking offline the SRDF and HTC Snap agents
This action reattaches the hardware mirrors to the replicating secondary devices and resynchronizes them.

Deleting the fire drill configuration

When deleting the fire drill configuration, the wizard does the following:

- Deletes the fire drill service group
- Deletes any associated registry entry

If you want to remove the hardware mirrors, you must do so manually.

For more information about the Hitachi TrueCopy Snap agent functions, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

For more information about the EMC SRDF Snap agent functions, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

About post-fire drill scripts

You can specify a script for the Fire Drill Wizard to run on the secondary site at the end of the fire drill.

For the wizard to run the script, the script must exist on the secondary system that you are specifying for the fire drill.

Note: The wizard does not support using script commands to launch a user interface window. In such a case, the process is created but the UI window does not display.

Optionally, you can specify to run a Windows PowerShell cmdlet. To run a cmdlet, create a .bat file with the following entry:

```
%windir%\system32\WindowsPowerShell\v1.0\PowerShell.exe -command "$ScriptName"
```

Where

ScriptName = .ps1 script (fully qualified) / cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\PowerShell.exe -command C:\myTest.ps1
```

Specify the name of the .bat file as the script to run.

For Exchange 2007, Symantec recommends using scripts rather than cmdlets as Exchange is not brought online at the secondary site, and the Exchange Management shell should be run under the virtual environment context, which is not available on a node where Exchange is offline.

To specify a cmdlet for Exchange 2007, use the following entry in the .bat file:

```
%windir%\system32\WindowsPowerShell\v1.0\PowerShell.exe -PSConsoleFile "$ExchDir"\bin\exshell.psc1 -command "$ScriptName"
```

Where

ExchDir equals HKLM -> SOFTWARE -> Microsoft -> Exchange -> Setup:: Services key.

ScriptName = .ps1 script (fully qualified) / cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\PowerShell.exe -PSConsoleFile "D:\Program Files\Microsoft\Exchange Server\bin\exshell.psc1" -command C:\myTest.ps1
```

Tasks for configuring and running fire drills

While running the Fire Drill Wizard, the following sequence of actions are available:

- Prepare the fire drill configuration
- Run the fire drill or delete the configuration
- Restore the fire drill configuration after running a fire drill
- Run another fire drill or delete the configuration

In addition, you have the option to recreate a fire drill configuration that has changed.

After an action is complete, the next action becomes available in the wizard. You can select the next action or exit the wizard and perform the next action later.

[Table 15-1](#) gives more details of the process of configuring and running fire drills with the wizard.

Table 15-1 Process for configuring and running fire drills

Action	Description
Verify the hardware and software prerequisites	Before running the wizard, review the prerequisites and make sure that they are met. See “Prerequisites for a fire drill” on page 686.
Prepare the fire drill configuration	Use the wizard to configure the fire drill. See “Preparing the fire drill configuration” on page 688.
Recreate a fire drill configuration that has changed	If a fire drill configuration exists for the selected service group, the wizard checks for differences between the fire drill service group and the application service group. If differences are found, the wizard can recreate the fire drill configuration before running the fire drill. See “Recreating a fire drill configuration that has changed” on page 696.

Table 15-1 Process for configuring and running fire drills

Action	Description
Run the fire drill	<p>Use the wizard to run the fire drill. Running the fire drill brings the fire drill service group online. Optionally you can specify a script to be run once the fire drill is complete.</p> <p>See “Running a fire drill” on page 694.</p> <p>Confirm the resources are online and replicated data is available</p> <p>Note: After completing the fire drill testing, run the wizard again as soon as possible to restore the configuration. Otherwise the fire drill service group remain online. Be sure to restore one fire drill service group to a prepared state before running a fire drill on another service group.</p>
Restore the fire drill configuration to a prepared state	<p>Use 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</p> <p>This is a required action after running the fire drill.</p> <p>See “Restoring the fire drill system to a prepared state” on page 699.</p> <p>This operation reattaches snapshot mirrors and takes the fire drill service group offline.</p>
Delete the fire drill configuration	<p>If a fire drill service group is no longer needed, or if you want to free up resources, use the wizard to remove the fire drill configuration</p> <p>See “Deleting the fire drill configuration” on page 700.</p> <p>The wizard deletes the service group on the secondary site. In a VVR environment, the wizard performs a snap abort to delete the snapshot mirrors created on the secondary site for use in the fire drill. In hardware replication environments, you can delete these manually.</p> <p>If a fire drill has been run, the wizard ensures that you first restore the fire drill configuration to a prepared state before this option becomes available. This ensures that mirrors are reattached and the fire drill service group is offline before the configuration is deleted.</p>

Prerequisites for a fire drill

Before running the Fire Drill Wizard make sure that you meet the following general requirements:

- 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.
- 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.
- If a firewall exists between the wizard and any systems it needs access to, the firewall must be set to allow both ingoing and outgoing TCP requests on port 7419.
- 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 you specify for the fire drill wizard to run Eseutil, the output files are placed by default in the system's TEMP environment variable folder (for example, C:\Windows\Temp). If you want the output files to go to another folder, use the WINSOL_ESEUTIL_OUT_DIR environment variable to define the output file location.

Additional requirements apply to specific replication environments.

See [“Prerequisites for a fire drill in a VVR environment”](#) on page 686.

See [“Prerequisites for a fire drill in a Hitachi TrueCopy environment”](#) on page 687.

See [“Prerequisites for a fire drill in an EMC SRDF environment”](#) on page 688.

Prerequisites for a fire drill in a VVR environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 686.

Make sure that the following additional prerequisites are met before configuring and running a fire drill in a Veritas Volume Replicator (VVR) environment:

- 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 secondary site cluster.

- 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. Snapshot mirrors take up the same amount of space as the original volumes. In addition, two disk change object (DCO) volumes are created for each snapshot mirror, one for the source volume and one for the snapshot volume. The two DCO volumes must be on different disks. Allow 2 MB additional space for each DCO volume.
The empty 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.
- All disk groups in the service group must be configured for replication. The Fire Drill wizard does not support a VVR configuration in which disk groups are excluded from replication. However, you can exclude individual volumes within a disk group from replication.

Prerequisites for a fire drill in a Hitachi TrueCopy environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 686.

Make sure that the following prerequisites are met before configuring and running a fire drill in a Hitachi TrueCopy environment:

- The primary and secondary sites must be fully configured with Hitachi TrueCopy replication and the global cluster option. The configuration must follow the applicable instructions in the Veritas Storage Foundation HA for Windows documentation for configuring disaster recovery with Hitachi TrueCopy.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- Make sure that Hitachi RAID Manager/Command Control Interface (CCI) is installed.
- ShadowImage for TrueCopy must be installed and configured for each LUN on the secondary site target array. ShadowImage pairs must be created to allow for mirroring at the secondary site.

- The name of the ShadowImage device group must be the same as the replicated device group for both replicated and non-replicated LUNs that are to be snapshot. The instance number should be different.
- Make sure the HORCM instance managing the S-VOLs runs continuously; the agent does not start this instance.

Prerequisites for a fire drill in an EMC SRDF environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment. General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 686.

Make sure that the following prerequisites are met before configuring and running a fire drill in an EMC SRDF environment:

- The primary and secondary sites must be fully configured with EMC SRDF replication and the global cluster option. The configuration must follow the applicable instructions in the Veritas Storage Foundation HA for Windows documentation for configuring disaster recovery with EMC SRDF.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- The infrastructure to take snapshots at the secondary site must be properly configured between the secondary site source and target arrays. This process involves associating Symmetric Business Continuance Volumes (BCVs) and synchronizing them with the secondary site source.
- If you plan to run a fire drill on SRDF/A devices, you must have a TimeFinder/CG license. Make sure TimeFinder for SRDF is installed and configured at the target array.
- To take snapshots of R2 devices, BCVs must be associated with the RDF2 device group and fully established with the devices.
- To take snapshots of non-replicated devices, create a EMC Symmetrix device group with the same name as the SFW disk group. The device group must contain the same devices as in the disk group and have the corresponding BCVs associated.

Preparing the fire drill configuration

Preparing the fire drill configuration creates a fire drill service group. 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.

For a Veritas Volume Replicator (VVR) environment, the preparation step also prepares snapshot mirrors of production data at the specified node on the secondary site.

Note: Preparing the snapshot mirrors takes some time to complete.

Before you prepare the fire drill configuration with the Fire Drill Wizard, make sure that you meet the prerequisites.

See [“Prerequisites for a fire drill”](#) on page 686.

To prepare the fire drill configuration

- 1 Open the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**).
- 2 Start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, 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 a system in the primary site cluster and click **Next**.
See [“System Selection panel details”](#) on page 691.
- 5 In the Service Group Selection panel, select the service group that you want to use for the fire drill and click **Next**.
See [“Service Group Selection panel details”](#) on page 691.
- 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**.
See [“Secondary System Selection panel details”](#) on page 691.
- 7 If the Fire Drill Prerequisites panel is displayed, review the information and ensure that all prerequisites are met. Click **Next**.
See [“Prerequisites for a fire drill”](#) on page 686.
Otherwise, if a fire drill service group already exists on this system for the specified service group, one of the following panels is displayed:

If the Run Fire Drill option or Delete Fire Drill options are shown, 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 694.
--	---

If the Fire Drill Restoration panel is displayed, the fire drill service group remains online from a previous fire drill.

Follow the procedure for restoring the fire drill configuration to a prepared state. This must be done before running a new fire drill.

See [“Restoring the fire drill system to a prepared state”](#) on page 699.

If the Recreate Fire Drill Service Group panel is displayed, a fire drill service group has already been prepared but is not up to date.

You can choose to recreate the fire drill configuration to bring it up to date.

See [“Recreating a fire drill configuration that has changed”](#) on page 696.

Or you can clear the check box to recreate the configuration and run the fire drill on the existing configuration.

- 8 The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

VVR replication Disk Selection panel

See [“Disk Selection panel details”](#) on page 691.

Hitachi TrueCopy replication Horcm Files Path Selection panel

See [“Hitachi TrueCopy Path Information panel details”](#) on page 692.

HTCSnap Resource Configuration panel

See [“HTCSnap Resource Configuration panel details”](#) on page 693.

EMC SRDF replication SRDFSnap Resource Configuration panel

See [“SRDFSnap Resource Configuration panel details”](#) on page 693.

Click **Next**.

- 9 In the Fire Drill Preparation panel, the wizard shows the status of the preparation tasks.
 - 10 The Summary panel displays the message that preparation is complete.
- See [“Fire Drill Preparation panel details”](#) on page 694.
- When preparation is complete, click **Next**.
- To run the fire drill now, click **Next**. Continue with the procedure to run the fire drill.
- See [“Running a fire drill”](#) on page 694.

To run the fire drill later, click **Finish**. The fire drill preparation remains in place.

System Selection panel details

Use the System Selection panel of the wizard to specify a system in the primary site cluster.

All systems containing online global service groups are available to select. The default system is the node where you launched the wizard (localhost) if a global service group is online on that system. When selecting a system you can specify either a fully qualified host name or IP address.

Service Group Selection panel details

Use the Service Group Selection panel of the wizard to select the service group that you want to use for the fire drill. You can select only one service group at a time for a fire drill.

Secondary System Selection panel details

Use the Secondary System Selection panel of the wizard to select the cluster and the system to be used for the fire drill at the secondary site.

The selected system must have access to the replicated data.

The system must have access to disks for the snapshots that will be created for the fire drill.

Disk Selection panel details

During fire drill preparation in a VVR replication environment, you must ensure that information is available to the wizard for creating the fire drill snapshots.

Use the Disk Selection panel of the wizard to review the information on disks and volumes and make the selections for the fire drill snapshots, as follows:

Volume	<p>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.</p> <p>Note: The Disk Selection panel also appears if the wizard is recreating a fire drill service group to which volumes have been added. In that case, only the new volumes are shown for selection.</p>
Disk Group	<p>Shows the name of the disk group that contains the original volumes. This field is display only.</p>
Fire Drill DG	<p>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>.</p>
Disk	<p>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.</p> <p>You can store multiple snapshot volumes on the same disk, if the production volumes reside on disks in the same disk group.</p> <p>If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard.</p>
Mount Details	<p>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.</p>

Hitachi TrueCopy Path Information panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the Hitachi TrueCopy Path Information panel is displayed.

The wizard populates the path field with the customary default location:

C:\Windows

where C is the system drive.

If the horcm configuration files are in a different location, edit the field to specify that location.

HTCSnap Resource Configuration panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the wizard discovers the HTC resources and non-replicating SFW disk groups in the application service group

This information is used to configure the HTCSnap resources.

The wizard lists each HTCSnap resource that will be configured. You can clear the HTCSnap resource name check box if you do not want to include its dependent disk groups in the fire drill.

You must specify the ShadowImage instance.

The HTCSnap Resource Configuration panel shows the following:

Target Resource Name	The panel shows the HTC resource name in the case of a Replication Device Group or the disk group resource name in the case of a non-replicating disk group.
ShadowImage Instance ID	For every HTC resource, specify the ID of the ShadowImage instance associated with the replicating secondary devices.
Refresh	If you click the Refresh button, the wizard rediscovers and validates the HTC configuration.

More information about HTCSnap resource configuration and operation is available.

See “[About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment](#)” on page 681.

SRDFSnap Resource Configuration panel details

During fire drill preparation in an EMC SRDF replication environment, the wizard validates whether at least one BCV device is attached to every device (RDF2) of the SRDF device group. If not, the wizard displays an informational message on this panel. The panel shows as the Target Resource Name the name of the resource that is managing the LUNs that you want to snapshot. For data being replicated from the primary site, the Target Resource Name is the name of the SRDF resource. For data that is not replicated, the Target Resource Name is the name of the disk group resource.

The wizard lists each SRDFSnap resource that will be configured. You can clear the SRDFSnap resource name check box if you do not want to include its dependent disk groups in the fire drill.

More information about SRDFSnap resource configuration and operation is available.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 681.

Fire Drill Preparation panel details

After you enter the information required to prepare a fire drill configuration, the Fire Drill Preparation panel is displayed. You wait while the wizard completes the preparation tasks.

The fire drill service group is created on the secondary site (but remains offline).

In addition, for a VVR replication environment, 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. If the wizard is completing the preparation steps as part of recreating a fire drill configuration, snapshot mirrors are prepared only for new volumes.

See [“Recreating a fire drill configuration that has changed”](#) on page 696.

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
- Enables the firedrill resources
- Brings the fire drill service group online
- Optionally runs Eseutil with the /g option
- Optionally, executes a specified command to run a script
See [“About post-fire drill scripts”](#) on page 683.

For details on the operations that occur when running a fire drill, see the following topics:

- [“About Fire Drill Wizard operations in a VVR environment”](#) on page 679
- [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 681

Warning: After running the fire drill, the fire drill service group remains online. After you verify the fire drill results, run the wizard again to restore the system to the prepared state. Otherwise, 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. See [“Restoring the fire drill system to a prepared state”](#) on page 699.

To run a fire drill

- 1 If you completed the initial preparation and have not exited the wizard, or if you are returning to this procedure after recreating a fire drill service group, go to [step 8](#). Otherwise, if you need to restart the wizard, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, 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 a system in the primary site cluster 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 and click **Next**.
If the fire drill configuration is in a prepared state, the wizard compares the resources of the fire drill service group with the resources of the application service group.
- 7 If the application service group changed since the fire drill configuration was prepared, the wizard displays the Recreate Fire Drill Service Group panel, showing the differences. Choose one of the following:
 - Leave the option checked to recreate the configuration before running the fire drill and click **Next**. You complete additional steps in the wizard before running the fire drill.
For more information, see [“Recreating a fire drill configuration that has changed”](#) on page 696.
 - To run the fire drill on the existing configuration, clear the option to recreate the fire drill service group and click **Next**.
- 8 In the Fire Drill Mode Selection panel, click **Run Fire Drill** and click **Next**.

- 9 In the Post Fire Drill Script panel, you have the option to specify the full path to a script for the wizard to run after the running the fire drill. In addition, you can specify to run the Eseutil consistency check. See “[Post fire drill operations panel details](#)” on page 696.
- 10 In the Fire Drill Implementation screen, wait until all fire drill tasks are performed and click **Next**. The Summary panel displays the message that the fire drill is complete. You can leave the wizard running while you verify the results or exit the wizard. To exit the wizard, click **Finish**.
- 11 Run your own tests to verify the fire drill results.

Warning: You should always restore the fire drill system to a prepared state immediately after completing fire drill testing on a service group.

- 12 Restore the fire drill configuration to the prepared state. See “[Restoring the fire drill system to a prepared state](#)” on page 699.

Post fire drill operations panel details

In the Post Fire Drill Script panel, the wizard displays options for the following actions that it can perform after bringing the fire drill service group online:

- 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.
For more information, see “[About post-fire drill scripts](#)” on page 683.
- Check the **Run Eseutil** check box if you want the Eseutil consistency check run on the fire drill snapshots once they are created.
The Eseutil output files are placed by default in the system’s TEMP environment variable folder unless you defined another location by using the WINSOL_ESEUTIL_OUT_DIR environment variable.

Recreating a fire drill configuration that has changed

When you run the Fire Drill wizard, a fire drill service group may already exist for the selected application service group. However, the application service group may have changed since the fire drill service group was created. Therefore, the wizard compares the resource names of the two service groups. If differences are found, the wizard lists them on the Recreate Fire Drill Service Group panel.

You have the following choices from the Recreate Fire Drill Service Group panel:

- Leave the option checked to recreate the fire drill service group.
 Proceed with using the wizard to recreate the configuration to match the application service group.
 The wizard deletes the existing fire drill configuration first, before creating the new one.
 For a VVR replication environment, the wizard handles existing volumes as follows: It does not delete the mirrors for volumes that still exist. When it recreates the fire drill configuration, it prepares new mirrors only for new volumes. If volumes have been removed, the wizard displays an additional option to snap abort the obsolete snapshot volumes to free up disk space.
- Clear the option to recreate the fire drill service group. You can then proceed with using the wizard to do either of the following:
 - Run the fire drill, ignoring the differences.
 - Delete the entire fire drill configuration. Then start over with preparing the fire drill configuration.

Note: The wizard does not check for changes in volume attributes, such as the MountPath attribute. For example, if you have a MountV resource with an attribute that points to drive Y and you change that attribute to point to drive X, the wizard does not identify this change and does not give the option to recreate the fire drill service group.

You can choose whether to manually edit the fire drill service group for such changes and then run the fire drill, ignore the differences, or delete the configuration and start over.

The following procedure describes the choice of recreating the fire drill configuration.

To recreate the fire drill configuration if the service group has changed

- 1 In the Recreate Fire Drill Service Group panel, leave the option checked to recreate the configuration before running the fire drill.
 For a VVR replication environment, if volumes have been removed, optionally select to snap abort the volumes.
 Click **Next**.
- 2 In the Fire Drill Mode Selection panel, Delete Fire Drill Configuration is selected. Click **Next**, and click **Yes** to confirm the deletion.
- 3 The Fire Drill Deletion panel shows the progress of the deletion.
 For a VVR replication environment, the wizard leaves the existing fire drill snapshot volumes so that those snapshot mirrors do not have to be

prepared again. If volumes were removed and you selected the option to snap abort, the wizard snap aborts the snapshots of those volumes.

Warning: If you close the wizard after deleting the fire drill configuration without continuing on to the fire drill preparation step, the information of the existing snapshot volumes is lost.

When all tasks are complete, click **Next**.

- 4 In the Fire Drill Prerequisites panel, review the information and ensure that all prerequisites are met. Click **Next**.
See “[Prerequisites for a fire drill](#)” on page 686.
- 5 The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

VVR replication If volumes have been added, the Disk Selection panel is displayed. Specify the information for the added volumes.

If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard.

See “[Disk Selection panel details](#)” on page 691.

Hitachi TrueCopy replication Horcm Files Path Selection panel

See “[Hitachi TrueCopy Path Information panel details](#)” on page 692.

HTCSnap Resource Configuration panel

See “[HTCSnap Resource Configuration panel details](#)” on page 693.

EMC SRDF replication SRDFSnap Resource Configuration panel

See “[SRDFSnap Resource Configuration panel details](#)” on page 693.

Click **Next**.

- 6 The Fire Drill Preparation panel is displayed. Wait while the wizard recreates the fire drill service group.
For VVR replication environments, wait while the wizard starts mirror preparation.
Mirror creation can take some time. You may want to minimize the wizard while the task runs in the background. You can also close the wizard and track the mirror preparation progress in the VEA.

- 7 Once preparation is complete, click **Next**. The Summary page is displayed. To continue with running the fire drill, click **Next**. See “[Running a fire drill](#)” on page 694.

Restoring the fire drill system to a prepared state

After running a fire drill and verifying the results, use the Fire Drill Wizard as soon as possible to restore the fire drill system at the secondary site to a prepared state. A prepared state is the initial fire drill configuration created by the wizard, in which the fire drill service group has been prepared but is offline.

Restoring the fire drill system to a prepared state is required for any of the following:

- Making the secondary system available for failover of the application service group at the primary site.
- Running another fire drill.
- Deleting the fire drill configuration after a fire drill has been run.

For details on the operations that occur when restoring a fire drill configuration, see the following topics:

- “[About Fire Drill Wizard operations in a VVR environment](#)” on page 679
- “[About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment](#)” on page 681

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 8](#). Otherwise, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, 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 a system in the primary site cluster 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 Restoration Information panel, review the requirements for restoration and click **Next**.
- 8 In the Fire Drill Restoration screen, wait until the screen shows the restoration tasks are completed and click **Next**.
- 9 In the Summary screen, click **Next** if you want to delete the fire drill configuration. Otherwise 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.

In a VVR replication environment, deleting a fire drill configuration also 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.

In a Hitachi TrueCopy or EMC SRDF environment, you could manually remove mirrors after the deletion is complete.

To delete a fire drill configuration

- 1 If you have just used the wizard to prepare or restore a fire drill configuration and have not exited the wizard, go to [step 10](#). Otherwise continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, 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 a system in the primary site cluster 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 If the wizard detects that the fire drill service group is different from the application service group, it displays the Recreate Fire Drill Service Group panel. Clear the option to recreate the fire drill service group and click **Next**.

- 8 If the wizard detects that the fire drill service group is still online, the Fire Drill Restoration panel is displayed. Review the requirements for restoration and click **Next**.
- 9 In the Restore Fire Drill screen, wait until the screen shows the restoration tasks are completed. Then click **Next**.
- 10 In the Fire Drill Mode Selection panel, click **Delete Fire Drill Configuration** and click **Next**, and click **Yes** to confirm the deletion.
- 11 The Fire Drill Deletion panel shows the progress of the deletion. Wait until all tasks are complete and then click **Next**.
If errors occur while deleting the fire drill configuration, the wizard will list any incomplete steps so that you can complete them manually.
- 12 The Summary panel is displayed. Click **Finish**.

702 | Testing fault readiness by running a fire drill
| **Deleting the fire drill configuration**

Appendices

This appendice contains the following appendix:

- [Appendix A, “Deploying Disaster Recovery: Manual implementation of a new Exchange Server installation”](#) on page 705

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

This chapter contains the following topics:

- [“Tasks for configuring disaster recovery”](#) on page 706
- [“Reviewing the configuration”](#) on page 709
- [“Verifying the primary site configuration”](#) on page 712
- [“Setting up the SFW HA environment \(secondary site\)”](#) on page 712
- [“Setting up security for VVR”](#) on page 713
- [“Installing Exchange on the first node and additional nodes \(secondary site\)”](#) on page 716
- [“Installing Exchange on the first node with DR Option \(secondary site\)”](#) on page 716
- [“Installing Exchange on additional nodes \(secondary site\)”](#) on page 721
- [“Copying the .CRK file to the primary site”](#) on page 727
- [“Backing up and restoring the Exchange disk group”](#) on page 727
- [“Configuring the Exchange service group for VCS \(secondary site\)”](#) on page 728

- [“Verifying the cluster configuration”](#) on page 734
- [“About configuring the DR components \(VVR and GCO\)”](#) on page 735
- [“Reviewing the prerequisites for configuring DR”](#) on page 736
- [“Setting up the replicated data sets \(RDS\) for VVR”](#) on page 736
- [“Creating the VVR RVG service group”](#) on page 748
- [“Configuring the global cluster option for wide-area failover”](#) on page 751
- [“Administering global service groups”](#) on page 757
- [“Establishing secure communication within the global cluster \(optional\)”](#) on page 762
- [“Possible task after creating the DR Environment: Adding a new failover node”](#) on page 764

Tasks for configuring disaster recovery

You configure disaster recovery in the following sequence:

- Configure the primary site for high availability and disaster recovery
These tasks are covered in the High Availability section.
While configuring the primary site, you select certain options to prepare the site for disaster recovery.
See [“Verifying the primary site configuration”](#) on page 712.
- Create a secondary “failover” site for disaster recovery
This chapter provides information on how to install and configure the secondary site for disaster recovery using SFW HA and Veritas Volume Replicator (VVR). It uses a scenario of an active/passive configuration with one to one failover capabilities.

Note: This chapter covers the “manual” method of deploying disaster recovery. A newer method uses the Solutions Configuration Center and the Disaster Recovery (DR) wizard to clone storage configuration and service groups and set up replication. See [Chapter 13, “Deploying Disaster Recovery: New Exchange Server installation”](#) on page 551.

Table A-1 on page 707 outlines the high-level objectives and the tasks to complete each objective:

Table A-1 Task list for configuring disaster recovery

Objective	Tasks
“Reviewing the configuration” on page 709	Understanding active-passive configuration and site failover in a DR environment
“Verifying the primary site configuration” on page 712	Verify that the primary site has been fully configured for high availability
“Setting up the SFW HA environment (secondary site)” on page 712	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites ■ Setting up the network and storage for a cluster environment ■ Installing SFW HA ■ Configuring the cluster using the VCS Cluster Configuration Wizard (VCW) ■ 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
“Setting up security for VVR” on page 713	Configure the VxSAS service for VVR, specifying the cluster nodes at both primary and secondary sites
“Installing Exchange on the first node and additional nodes (secondary site)” on page 716	<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 A-1 Task list for configuring disaster recovery (Continued)

Objective	Tasks
“Copying the .CRK file to the primary site” on page 727	Copying the public cryptographic key of the Exchange virtual server from the secondary site to the primary site
“Backing up and restoring the Exchange disk group” on page 727	Backing up the Exchange disk group on the primary site and restoring it on the secondary site
“Configuring the Exchange service group for VCS (secondary site)” on page 728	Configuring the Exchange service group for VCS
“Verifying the cluster configuration” on page 734	Verifying the cluster configuration by switching service groups and shutting down an active cluster node
“Reviewing the prerequisites for configuring DR” on page 736	Verifying HA prerequisites for DR components
“Setting up the replicated data sets (RDS) for VVR” on page 736	<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 748	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 751	<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

Table A-1 Task list for configuring disaster recovery (Continued)

Objective	Tasks
“Administering global service groups” on page 757	Performing administrative tasks on global service groups
“Possible task after creating the DR Environment: Adding a new failover node” on page 764	Reviewing required tasks when adding a new failover system to either the primary or secondary site

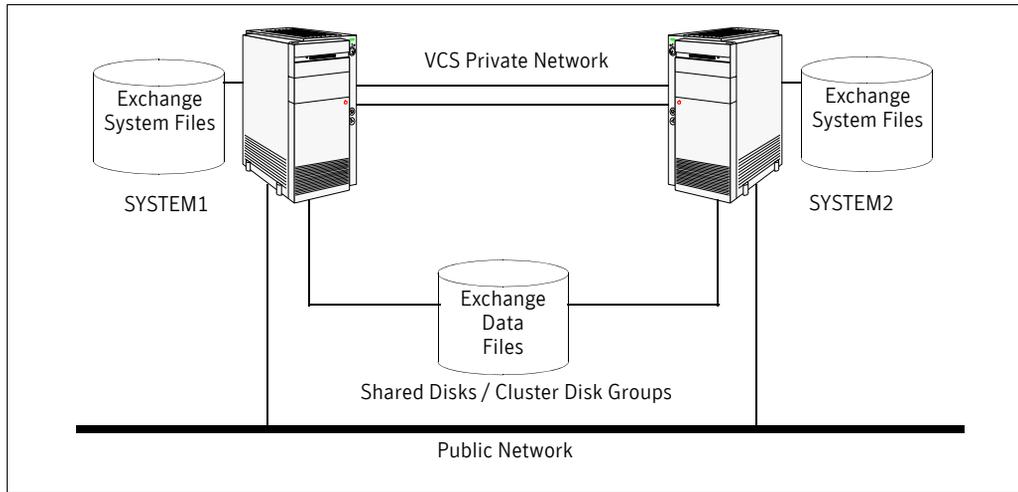
Reviewing the configuration

This configuration 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.

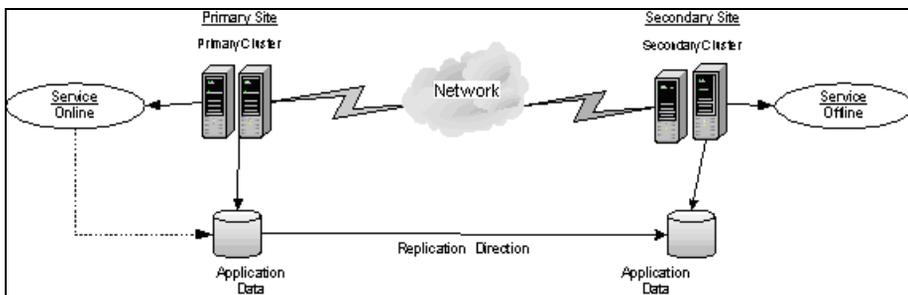
Figure A-1 Cluster configuration on the primary site



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.

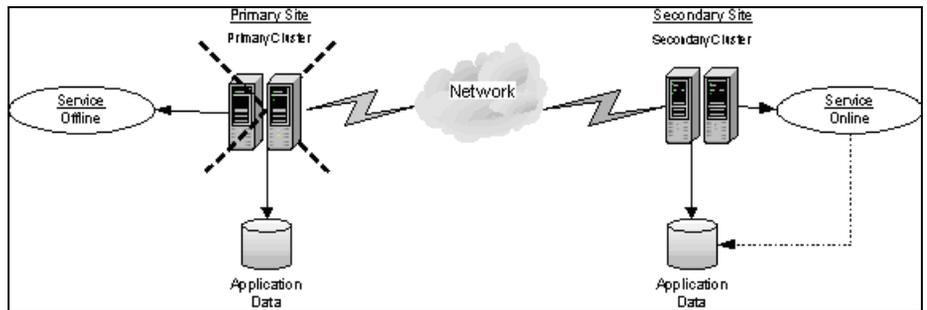
[Figure A-2](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure A-2 Data replication in a disaster recovery environment



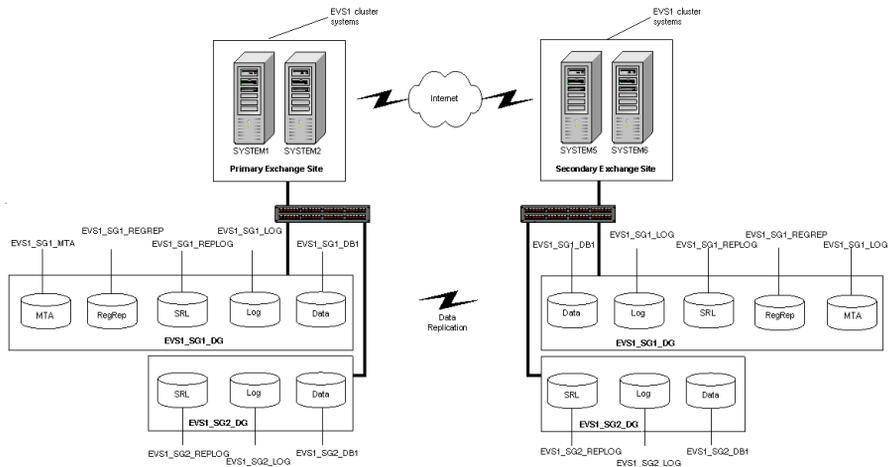
When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients. [Figure A-3](#) illustrates this type of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

Figure A-3 Application services restored after primary site failure



Below is an example of disk groups and volumes on both the primary and secondary sites of a disaster recovery environment using VVR replication.

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



Exchange storage group EVS1_SG1_DG contains the following 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 only required in a DR solution using VVR

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.

Verifying the primary site configuration

If you have not yet set up the primary site, follow the instructions in the chapter on deploying SFW HA for a new high availability installation.

See [Chapter 4, “Deploying SFW HA for high availability: New installation”](#).

The instructions include specifying the disaster recovery configuration options that are required on a primary site. Make sure that you complete the following tasks:

- When installing SFW HA, follow the instructions to install the GCO option. If you are using VVR for replication, ensure that you install VVR.
- When configuring disk groups and volumes, if you are using VVR for replication, allow sufficient disk space for a Storage Replicator Log (SRL) volume for each storage group.
- When running the Veritas Cluster Server Configuration Wizard to configure the cluster, follow the instructions to select the GCO option to configure the Global Cluster Option resource for the cluster.

Setting up the SFW HA environment (secondary site)

On the secondary site, begin by repeating the same tasks used to configure SFW HA on the primary site prior to the Exchange installation. Use the following instructions from [Chapter 4, “Deploying SFW HA for high availability: New installation”](#) to set up SFW HA on the secondary site before continuing with the tasks in this chapter:

- [“Reviewing the requirements”](#) on page 48
- [“Configuring the storage hardware and network”](#) on page 56
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 59
 - Ensure that you install the GCO option and (if using VVR for replication) the VVR option.
- [“Configuring disk groups and volumes”](#) on page 65

- 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.
- Ensure that you allow sufficient disk space to create a volume for the VVR Storage Replicator Log for each storage group. You can create the volume now, or later, when running the wizard to create replicated data sets. See “[Setting up the replicated data sets \(RDS\) for VVR](#)” on page 736.
- “[Configuring the cluster](#)” on page 75
 - When running the Veritas Cluster Server Configuration Wizard, ensure that you select the GCO option to configure the Global Cluster Option resource for the cluster.

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 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.

Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information panel as follows:

Account name Enter the administrative account name.
(domain\account)

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

Click **Next**.

- 3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

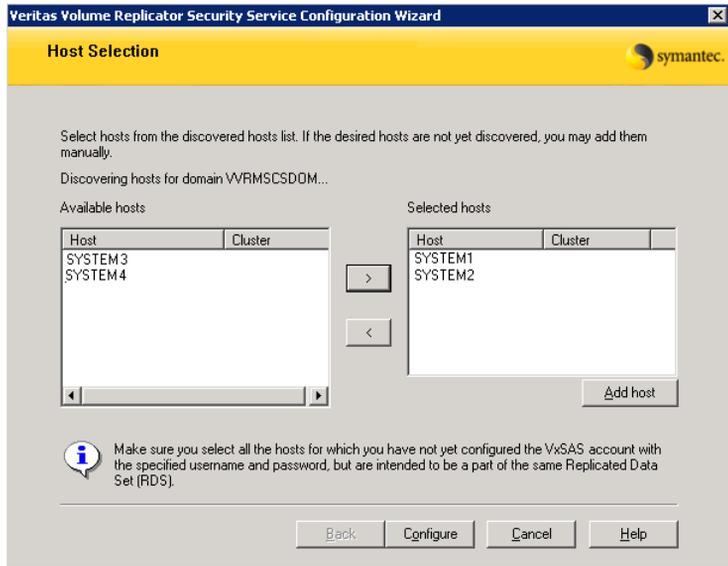
Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

- 4 On the Host Selection panel, select the required hosts:



Selecting hosts

The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host 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, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with 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.

Installing Exchange on the first node and additional nodes (secondary site)

When installing Exchange, you complete the following tasks:

- Review the prerequisite checklist
- Install Exchange on the first node
You run the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation on the first node.
Make sure to perform the first node pre-installation, installation, and post-installation procedures.
- Install Exchange on additional nodes
You run 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.

If you are familiar with installing Exchange for HA on the primary site, you will find that the procedures for installing Exchange are the same on the secondary site except for the pre-installation procedure for the first node.

In the pre-installation procedure for the first node, you must select the wizard option to create a failover node for Exchange disaster recovery setup, instead of the option to create a new Exchange virtual 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 73.
- 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.
 - For installing Exchange, you must be an Exchange Full Administrator and a member of the Exchange Domain Servers group.
 - If the logged-on user does not have Domain Administrator privileges, then the Exchange Domain Servers group must be managed by the VCS Helper service user account.
 - The logged-on user must either have permissions on the Exchange Domain Servers group in Active Directory or the Exchange Domain Servers group in the Active Directory must be managed by the VCS Helper service user account (In Active Directory Users and Computers, expand the domain entry, click Microsoft Exchange Security Groups and then specify the user account on the Managed By tab on the Exchange Servers Properties dialog box).
 - 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.
 - Either the logged-on user or the VCS Helper service user account 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 the VCS Helper service 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.
- 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 to take the service group offline.

Exchange pre-installation on first node (secondary site)

Use the following procedure to perform Exchange pre-installation.

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.

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 Tools > Exchange Server Setup Wizard** to start the Exchange Setup Wizard for VCS.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Available Option panel, choose the **Install Exchange Server for High Availability** option and click **Next**.
- 5 Specify the name or IP address of the cluster node at the primary site and click **Next**.
- 6 Select the name of the Exchange Virtual Server for which you want to set up a secondary site and click **Next**.
- 7 In the Select Option panel, choose the **Create a failover node for Exchange disaster recovery setup** option and click **Next**.
- 8 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.
- 9 Specify information related to the network.
 - Verify the virtual server name and domain suffix.
 - 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. 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**.

- 10 Select a drive where the registry replication data will be stored and click **Next**.
- 11 Review the summary of your selections and click **Next**.
- 12 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.
- 13 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**.
- 14 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.

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

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 on which you performed the pre-installation.

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

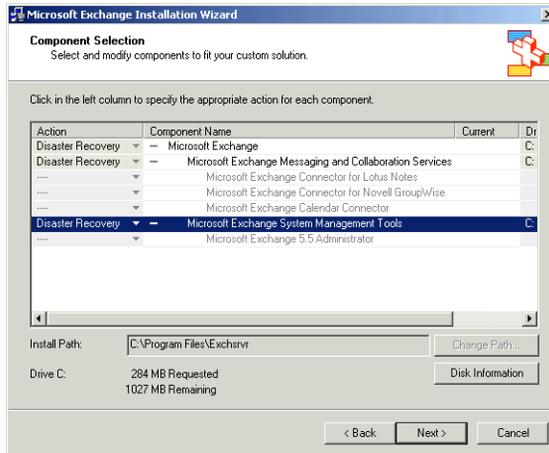
The procedure below is based on Exchange 2003. This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

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 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 that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
Type the following on the command line:

```
C:\>hasys -state
```


The state should display as `RUNNING`.
If HAD is not running, start it. Type the following on the command line:

```
C:\>net stop had
```



```
C:\>net start had
```
- 2 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#).
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Next**.
- 7 Click **Finish**.
- 8 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 in 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 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 716.
- 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 73.

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.

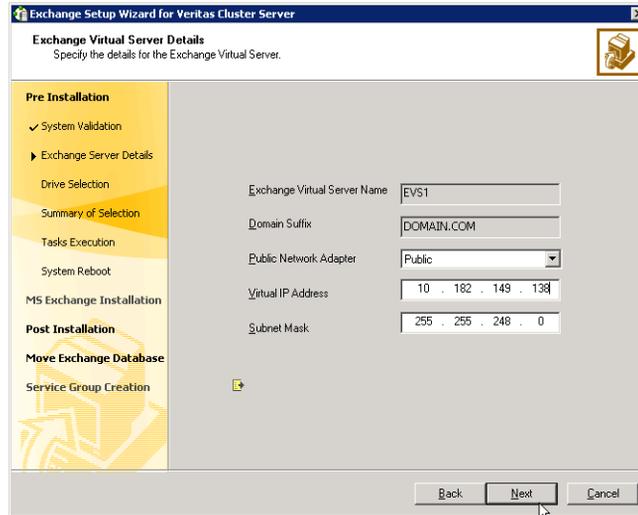
Exchange pre-installation: Additional nodes

Use the following procedure to perform Exchange pre-installation on additional nodes for the same EVS.

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 > Programs > Symantec > Veritas Cluster Server > Configuration Tools > 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.

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

If you want to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

Exchange installation: additional nodes

Install Exchange on the additional node on which you performed the 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.

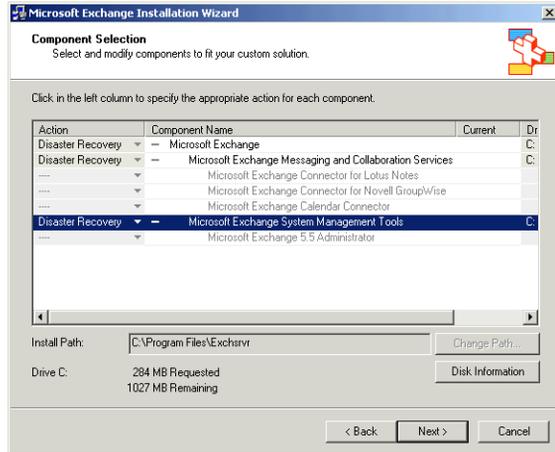
To install Exchange

- 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 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 Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.
 Type the following on the command line:
`C:\>hasys -state`
 The state should display as RUNNING.
 If HAD is not running, start it. Type the following on the command line:
`C:\>net stop had`
`C:\>net start had`

- 2 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.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to [step 5](#). If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
- 4 Review the information in the Welcome dialog box and click **Next**.
- 5 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.
- 6 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
- 7 If service groups are already configured for the EVS, specify whether you want to add the node to the system list of the service group for the EVS selected in the Exchange pre-installation step. You can also add nodes to a system list when configuring or modifying a service group with the Exchange service group configuration wizard.
- 8 Click **Finish**.
- 9 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.
- 10 If you are using the Disaster Recovery wizard, return to the Disaster Recovery wizard to continue with disaster recovery configuration.

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

A DR installation of Microsoft Exchange does not create Exchange data files. Therefore, after installing Exchange on the secondary site, you must back up the Exchange disk group on the primary site and then restore it on the secondary site.

Complete the following tasks before configuring the VCS Exchange service group on the secondary site:

- On the primary site, back up all volumes in the the Exchange disk group (EVS1_SG1_DG).
- Restore the group in the corresponding location on the secondary site.

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. This privilege is required to configure service groups.
- 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.
- Verify that you have backed up the Exchange disk group on the primary site and restored it on the secondary site.
See [“Backing up and restoring the Exchange disk group”](#) on page 727.
- 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 73.
- 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.

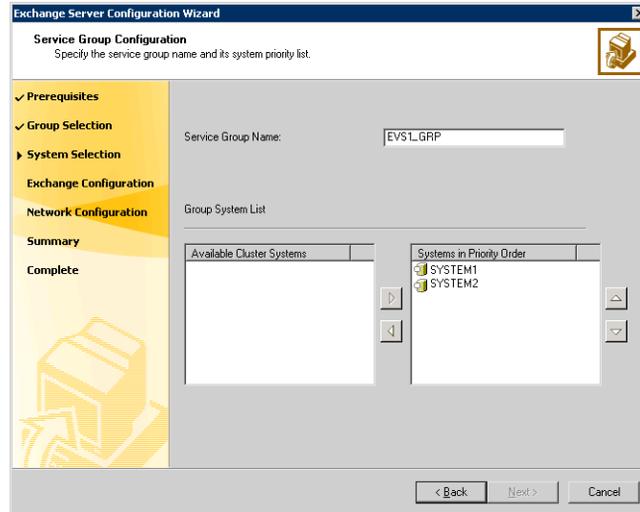
Refer to the *Veritas Cluster Server Application Agent for Microsoft Exchange, Configuration Guide* for information on the Exchange agent resource types,

attribute definitions, resource dependencies, and sample service group configurations.

Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on how to add additional resources to the EVS1_SG1_DG1 disk group.

To configure the Exchange service group

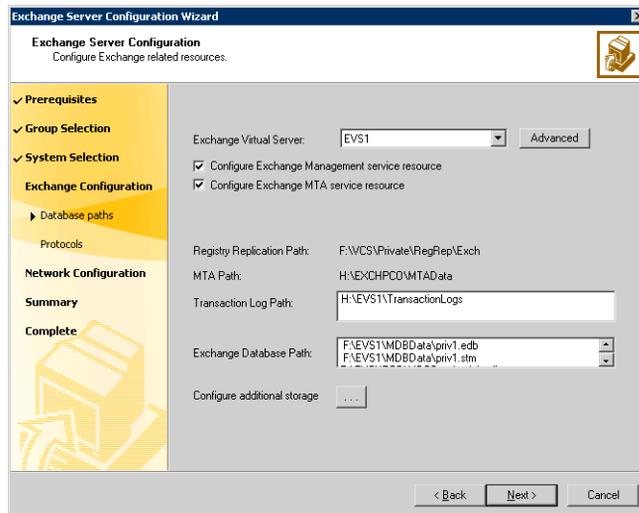
- 1 Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**
- 2 Review the information in the Welcome dialog box and click **Next**.
- 3 In the Wizard Options panel, choose the **Create service group** option and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will be part of the service group and then click **Next**: The wizard starts validating your configuration. Various messages indicate the validation status.



- Enter a name for the Exchange service group.
If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
- 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.
- 5 Verify the Exchange virtual server name and the drives or folder mounts created to store Exchange data:



Complete the following steps:

- Select the Exchange Virtual Server name from the drop-down list.
- Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.

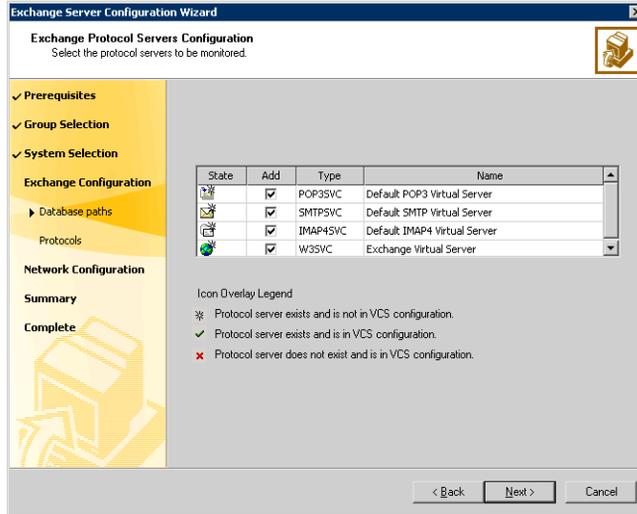
On the Lanman Advanced Configuration dialog box, complete the following:

- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ... (ellipsis) button and search for the OU using the Windows Find Organization Units dialog box. By default, the

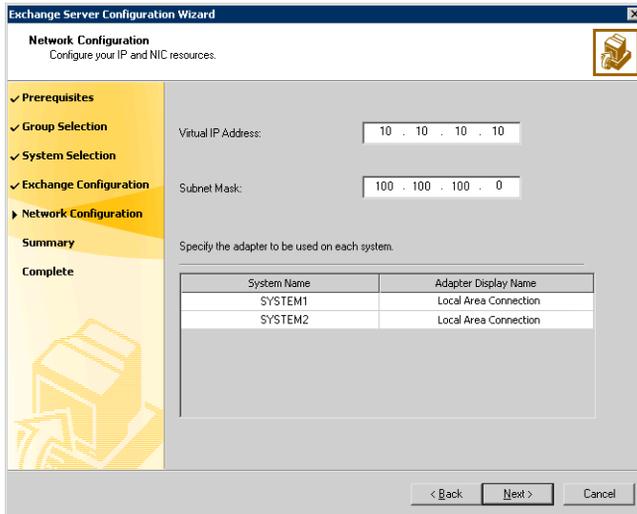
Lanman resource adds the virtual server to the default container "Computers."

- Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- Check the **Configure Exchange Management service resource** check box if you want to configure a resource for the Exchange Management service, in the Exchange service group.
 If you are running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Check the **Configure Exchange MTA service** resource check box to configure a resource for the Exchange Message Transfer Agent service, in the Exchange service group.
 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 running the wizard to modify an Exchange service group, unchecking this check box will remove the ExchangeMTA service resource from the service group configuration.
- Verify the registry replication path for the selected Exchange virtual server.
- Verify the MTA path for the selected Exchange virtual server.
- Verify the Transaction Log Path for the selected Exchange virtual server.
- To configure additional storage, click the ... (ellipsis) button and complete the following on the Additional Storage Configuration dialog box:
 - In the Available Volumes box, select a volume that you wish to add and click the right-arrow button to move the volume to the Selected Volumes box.
 - To remove a volume, select the volume in the Selected Volumes box, and click the left-arrow button.
 - Click **OK**. The wizard will configure resources required for the additional storage as child resources of the Microsoft Exchange System Attendant (MSEExchangeSA) service resource.
- Click **Next**.

- 6 On the Exchange Protocol Servers Configuration panel, check the protocol check boxes next to the protocol servers to be monitored and then click **Next**.

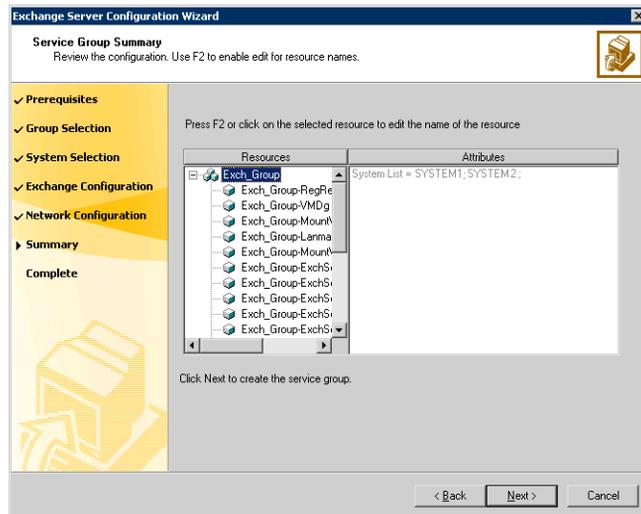


- 7 On the Network Configuration panel, specify information related to the network and then click **Next**:



- 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 Display Name** field to view the adapters associated with a node.
 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.

8 Review the service group configuration, change the resource names, if desired, and then click **Next**:



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 the **Esc** key.

9 Click Yes on the message that prompts you that the wizard will run commands to create the service group. Various messages indicate the status

of these commands. After the commands are executed, the completion dialog box appears.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then 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 then create the new storage groups and mailbox stores in Exchange System Manager. Run the Exchange Configuration Wizard again 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.

About configuring the DR components (VVR and GCO)

After configuring high availability and Exchange components on the primary and secondary sites, you configure the DR components for both sites. You

configure 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.

Note: You also have the choice of using array-based hardware replication for your disaster recovery solution. For information on configuring array-based hardware replication with VCS, see the VCS hardware agent documentation for the particular array you want to configure.

This section covers the following topics:

- [Reviewing the prerequisites for configuring DR](#)
- [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](#)
- [Possible task after creating the DR Environment: Adding a new failover node](#)

Reviewing the prerequisites for configuring DR

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

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.

Ensure that the following prerequisites are met:

- Verify that the data and replicator log volumes are *not* of the following types as VVR does not support these types of volumes:

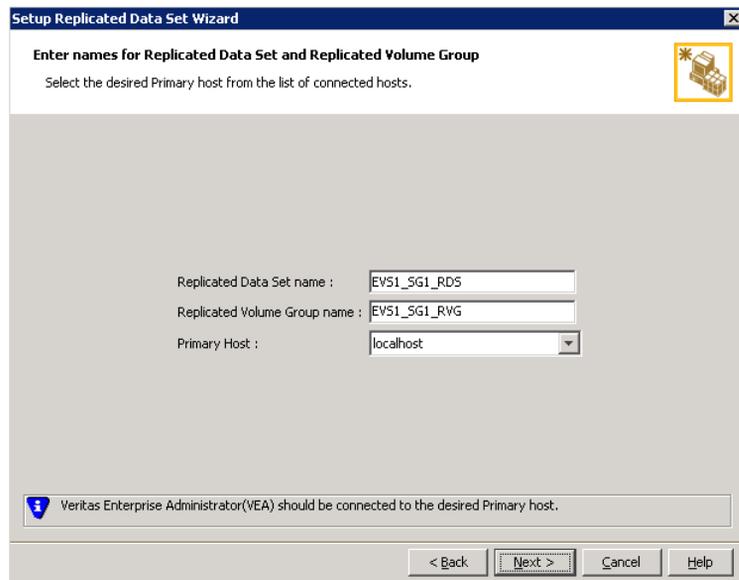
- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with a Dirty Region Log (DRL)
- Volumes that are already part of another RVG
- Volumes names containing a comma
- Verify that the Replicator Log volume does not have a DCM.
- Verify that the Replicator log volume does not have a drive letter assigned.
- Verify that the cluster disk group is imported on the primary and secondary site

Note: If you have not yet created the Storage Replicator Log volume, you can create it while setting up the Replicated Data Sets.

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.

- 3 Read the Welcome page and click **Next**.



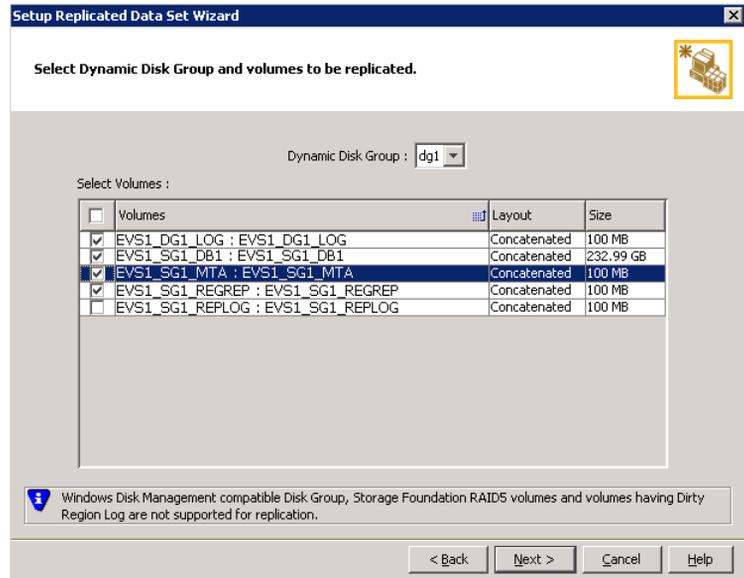
The screenshot shows a Windows-style dialog box titled "Setup Replicated Data Set Wizard". The main heading is "Enter names for Replicated Data Set and Replicated Volume Group". Below this, a sub-heading reads "Select the desired Primary host from the list of connected hosts." There is a yellow icon with a star and a server rack in the top right corner. The form contains three input fields: "Replicated Data Set name" with the text "EV51_SG1_RDS", "Replicated Volume Group name" with "EV51_SG1_RVG", and "Primary Host" with a dropdown menu showing "localhost". At the bottom left, there is a blue information icon and a message: "Veritas Enterprise Administrator(VEA) should be connected to the desired Primary host." At the bottom right, there are four buttons: "< Back", "Next >", "Cancel", and "Help".

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.

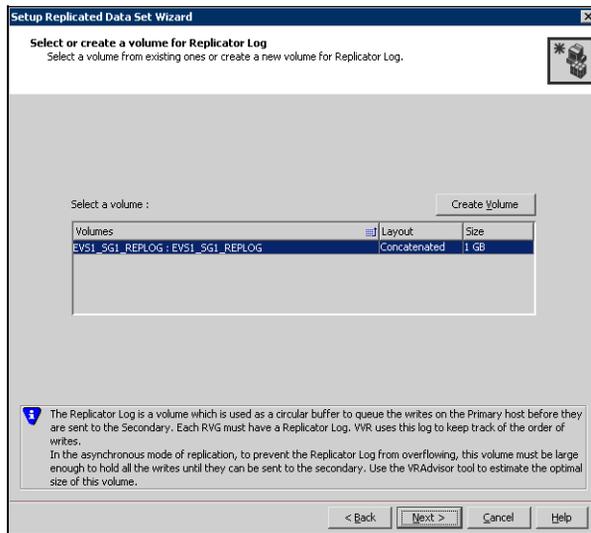


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (EVSI_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

Enables you to specify the disk selection method.

- Enable the **Thin Provisioned Disks Only** checkbox to ensure that the Replicator Log volume is created only on Thin Provisioned (TP) disks.

Note: The checkbox will remain disabled if the diskgroup does not have any TP disk.

If this option is selected along with the **Select disks automatically** option, then the Replicator Log volume will be created only on TP disks. However, if you enable this checkbox along with **Select disks manually** option, then the user can select only TP disks from **Available Disks**.

For more information on Thin Provisioning refer to the *Veritas Storage Foundation Administrator's Guide*.

- Choose the **Select disks automatically** option if you want VVR to select the disks.
- Choose the **Select disks manually** option to use specific disks from the Available disks pane for creating the volume. Either double-click on 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 Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?

Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.

Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.

10 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field. Click **Next**. If the Secondary host is not connected to VEA, the wizard tries to connect it when you click **Next**. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.

- 11 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.
The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:
 - the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the PrimaryOtherwise, the RDS setup wizard enables you to create the required volumes manually.
 - Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
- 12 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.
This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.
 - If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
 - If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.
Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.
When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.
 - If all the data volumes to be replicated meet the requirements, this screen does not occur.

- 13 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

The screenshot shows the 'Setup Replicated Data Set Wizard' window with the 'Edit replication settings' tab selected. The window title is 'Setup Replicated Data Set Wizard'. Below the title bar, there is a sub-header 'Edit replication settings' and a note 'Edit replication settings or click next.' To the right of the note is a small icon of a server rack. The main area contains several configuration fields, each with a label and a value in a drop-down or text box:

- Primary side IP: 10.217.53.214
- Secondary side IP: 10.217.53.215
- Replication Mode: Synchronous Override
- Replicator Log Protection: AutoDCM
- Primary RLINK Name: Pri_RLINK
- Secondary RLINK Name: Sec_RLINK

Below these fields is an 'Advanced' button. At the bottom of the window, there is a warning message: 'DHCP addresses are not supported by VVR.' and navigation buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Secondary side IP Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode	<p>Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous. The default is synchronous override.</p> <p>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.</p> <p>Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates.</p> <p>Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously.</p> <p>If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as <code>MISSING</code>.</p>
Replicator Log Protection	<p>The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.</p> <p>The DCM option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.</p> <p>The Off option disables Replicator Log Overflow protection.</p> <p>In the case of the Bunker node. Replicator Log protection is set to Off, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.</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.

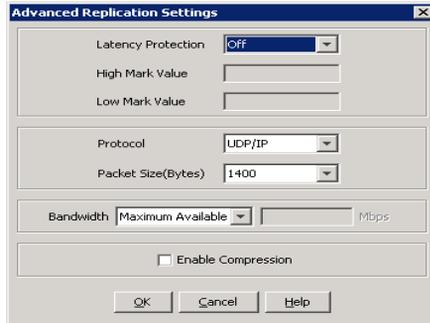
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

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	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- 14 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.

Enable Compression Enable this checkbox if you want to enable Compression for the secondary host.

Click **OK** to close the dialog box.

15 Click **Next**.

16 On the **Start Replication** page, select **Start Replication**.

Synchronize
Automatically

If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from
Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.
- 17 Review the information.
Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

Creating the 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. You create a replication service group, also known as an RVG service group. Before creating the service group verify the following:

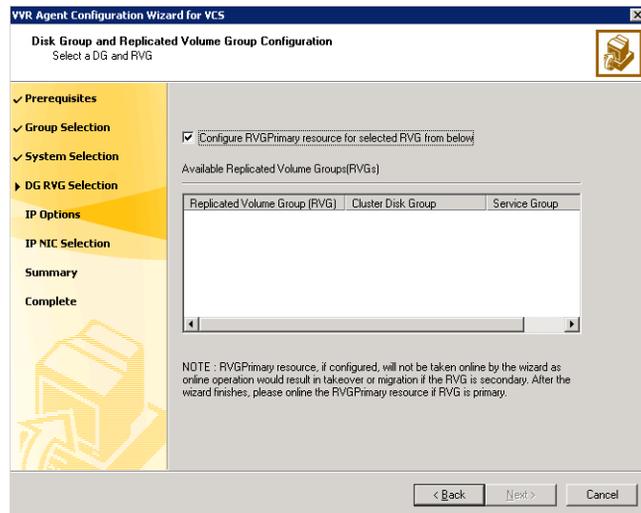
- Verify that the disk group is imported on the node on which you want to create the Replication Service Group.
- Verify VCS is running, by running the following command on the host on which the you intend to run the Volume Replicator Agent Configuration Wizard.

```
> 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 > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Review the requirements on the Welcome page and click **Next**.

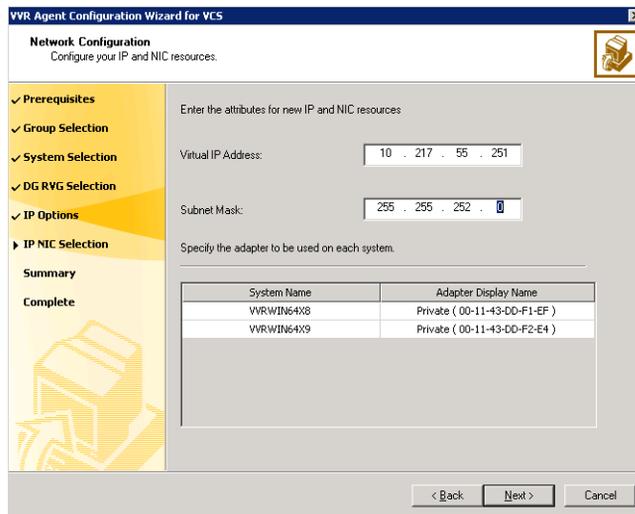
- 3 In the Wizard Options panel, click **Create a new replication service group** and click **Next**.
- 4 Specify the service group name and system priority list as follows:
 - Enter the service group name (EVS1_RVG_GRP).
 - In the Available Cluster Systems box, click the nodes on which to configure 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.
 - 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.
 - 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 panel, make the following selections:



- Select **Configure RVGPrimary resource for selected RVG**.

This resource is required when you want to configure your setup to automatically enable takeover in case of a failure of the Primary cluster. The `RVGPrimary` resource is created in the application service group and replaces the `VMDg` resource.

- Select the replicated volume group for which you want to configure the RVG primary resource.
For example, select `EVS1_SG1_RVG`.
 - Click **Next**.
- 7 In the IP Resource Options panel, select **Create a new IP resource** and click **Next**.
 - 8 In the Network Configuration panel, enter the network information as follows:



- Verify or enter the virtual IP address; use the IP address specified as the primary IP address when you configured the RDS.
 - Specify the subnet mask.
 - Specify the adapters for each system in the configuration.
 - Click **Next**.
- 9 Review the summary of the service group configuration as follows:
 - The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the Attributes box.
 - 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**.

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.
The name for the application service group must be the same on both sites. 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.

You can use the VCS Java Console or Web Console to perform global cluster operations; this guide provides procedures for the Java Console.

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 Cluster Manager (Java Console) provides a wizard to create global clusters by linking standalone clusters or bringing a standalone cluster into an existing global environment.

Note the following uses of the wizard:

- 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 VCS Cluster Manager:

- 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.

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 then click **Next**.

- 3 In the Wizard Options panel, click **Add Cluster**, then click **Next**.
- 4 In the New Cluster Details panel, enter the details of the new cluster.
 If the cluster is not running in secure mode, specify the following:
 - 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.
 - If necessary, change the default port number.
 - Enter the user name and the password.
 - Click **Next**.
 If the cluster is running in secure mode, specify the following:
 - 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.
 - Verify the port number.
 - 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.
 - 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 3.

- 2 Review the information required for the Global Group Configuration wizard and click **Next**.
- 3 Enter the details of the service group to modify, as follows:
 - Click the name of the service group that will be converted from a local group to a global group, or vice versa.
 - 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.
 - Select the policy for cluster failover as follows:

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.

- 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, as follows:

5 Click **Next**, then click **Finish**.

- | | |
|----------------------------|--|
| Cluster not in secure mode | <ul style="list-style-type: none"> ■ Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system. ■ Verify the port number. ■ Enter the user name. ■ Enter the password. ■ Click OK. ■ Repeat these steps for each cluster in the global environment. |
| Cluster in secure mode | <ul style="list-style-type: none"> ■ Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system. ■ Verify the port number. ■ 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. ■ Click OK. ■ Repeat these steps for each cluster in the global environment. |

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, specify the following:
 - Click the remote cluster to bring the group online.

- Click the specific system, or click **Any System**, to bring the group online.
- 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:
 - Click the remote cluster to take the group offline.
 - Click the specific system, or click **All Systems**, to take the group offline.
 - 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:
 - Click the cluster to switch the group.
 - Click the specific system, or click **Any System**, to take the group offline.
 - 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.

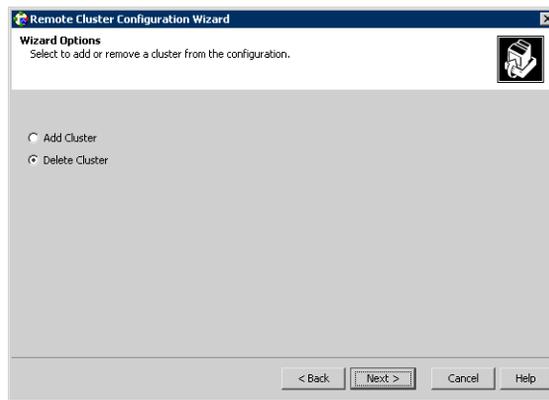
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **wac** resource under the **Application** type in the **ClusterService** group.
or
 Click a service group in the configuration tree, click the **Resources** tab, and right-click the **wac** resource in the 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

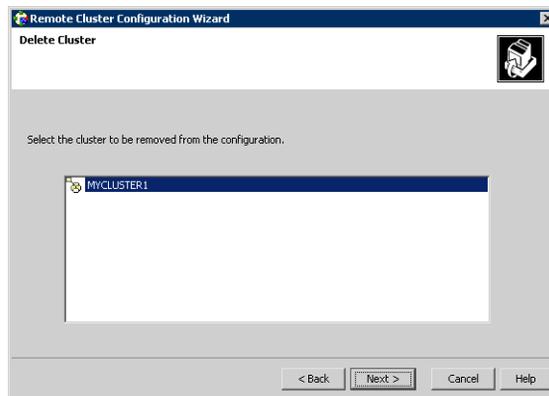
- 1 From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
- 2 Click **Next**.
- 3 Enter the details of the service group to modify:
 - Click the name of the service group.
 - For global to local cluster conversion, click the left arrow to move the cluster name from the cluster list back to the **Available Clusters** box.
 - 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:
 - Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
 - Verify the port number.
 - Enter the user name.
 - Enter the password.
 - Click **OK**.
 If the cluster is running in secure mode:
 - Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
 - Verify the port number.
 - Choose to connect to the remote cluster using the connected cluster's credentials, or enter new credentials, including the user name, password, and domain.
 - Click **OK**.
- 5 Click **Next**.
- 6 Click **Finish**.

To delete a remote cluster from the local cluster

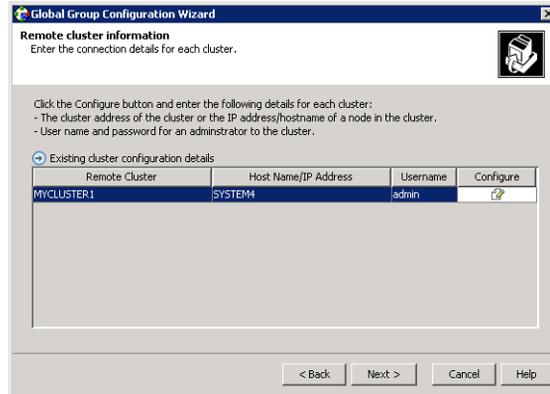
- 1 From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.
or
From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Clusters**.
- 2 Review the required information for the **Remote Cluster Configuration** wizard and click **Next**.
- 3 On the **Wizard Options** panel, click **Delete Cluster**, then click **Next**.



- 4 In the Delete Cluster panel, click the name of the remote cluster to delete, then 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:

- Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
- Verify the port number.
- Enter the user name.
- Enter the password.
- Click **OK**.

If the cluster is running in secure mode:

- Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
- Verify the port number.
- 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.

- Click **OK**.

- 6 Click **Finish**.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.
- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

To add the -secure option to the StartProgram resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add `-secure` switch to the path of the executable Scalar Value.
 For example:

```
"C:\Program Files\Veritas\Cluster Server\bin\wac.exe" -secure
```
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.

Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:

```
vssat setuptrust --broker <host:port> --securitylevel <low|medium|high> [--hashfile <filename> | --hash <root hash in hex>]
```

For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:

from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low
```

from RB2, type:

```
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the ClusterService-Proc (wac) resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

Possible task after creating the DR Environment: Adding a new failover node

The following procedure describes how to add an additional node to the cluster at either the primary or secondary site after your disaster recovery environment is in operation. The clusters at each site are not required to have the same number of nodes or the same failover configuration.

Preparing the new node

Install SFW HA on the new system and then add the system to the cluster.

To install SFW HA and add the system to the cluster

- 1 Refer to “[Installing Veritas Storage Foundation HA for Windows](#)” on page 59 for installation instructions.
- 2 Use the **Cluster Operations** option of the VCS Configuration wizard (**Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > 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

You prepare the DR environment by taking the global Exchange service group and VVR replication service group offline.

However, to add a failover node to the secondary site, you must first 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 are adding a failover node to the secondary site, switch the roles of the primary and secondary sites as follows:
 - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
 - Click **Switch To**, and click **Remote switch**.
 - In the **Switch global group** dialog box:
 - Click the cluster at the secondary site you want to switch the group to.
 - Click the specific system where you want to bring the global 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 73.
- 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”](#) on page 103 or [“Installing Exchange on additional nodes \(secondary site\)”](#) on page 721; reboot when prompted in these procedures.
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 > Configuration Tools > Volume Replicator Agent Configuration Wizard**) to add the new node to the system list for the replication service group.
See 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 >**

Configuration Tools > Exchange Server 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.

See the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for more 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 764, 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:
 - Click the cluster to switch the group to.
 - Click the specific system where you want to bring the global Exchange service group online.
 - Click **OK**.

Index

A

- any-to-any HA
 - configuration 202, 210, 298
 - configuring new nodes 300
 - creating a new cluster 210
 - disk space requirements 205, 292
 - new installation 199, 289
 - process overview 199, 289
 - sample configuration 203, 299
 - specifying a common node for failover 258, 338

C

- campus cluster
 - configuration 363
 - defined 352
 - disk space requirements 357
 - Exchange service group, modifying the IP resource 424
 - failover using the forceimport attribute 365
 - forceimport 427
 - handling site failure 427
 - new installation 355, 357
 - overview 351, 352
 - process overview 355
 - site failure 427
- cconfiguration
 - Exchange service group for VCS HA 108, 500
- cloning for DR
 - Exchange service group 624
 - secondary storage (array-based replication) 610
 - secondary storage (VVR replication) 605
- cluster
 - adding a node
 - any-to-any HA 308
 - configure LLT over ethernet 79, 160, 222, 379, 464, 580
 - configure LLT over UDP 80, 161, 223, 380, 465, 581

- configuring
 - any-to-any HA 308
 - HA 75
- preparing
 - with the any-to-any option 338, 670
- verifying configuration
 - campus cluster 426
 - DR 734
 - HA 115, 197, 265, 287, 346
- clusters
 - assigning user privileges 600, 663
 - configuring (RDC) 459
 - creating a disk group (RDC) 477
 - preconditions for a cluster disk group (RDC) 476
 - switching online nodes 544
- configuration
 - any-to-any DR example 661
 - any-to-any HA 202, 210, 298
 - any-to-any HA example 203, 299
 - campus cluster 363
 - DR 709
 - Exchange service group for VCS
 - any-to-any HA 258, 339
 - DR secondary site 728
 - standalone to HA 190
 - HA 53
 - standalone Exchange server
 - to existing cluster 131
 - to HA 129, 132
 - to new cluster 129
- configure
 - LLT over ethernet 79, 160, 222, 379, 464, 580
 - LLT over UDP using VCW 80, 161, 223, 380, 465, 581
- configure cluster
 - ethernet 79, 160, 222, 379, 464, 580
 - UDP 80, 161, 223, 380, 465, 581
- configuring
 - cluster (RDC) 459
 - Exchange service group for VCS
 - campus cluster 417

- configuring SFW HA
 - prior to installing Exchange
 - DR primary site 58
- configuring VSFW HA
 - prior to installing Exchange
 - DR secondary site 716

D

- disaster recovery
 - configuration 658
 - defined 549
 - new installation 705
 - overview 549
 - see also DR
- disaster recovery (DR)
 - cloning secondary storage (VVR replication) 605
 - cloning SQL Server 2005 service group 624
 - configuring GCO with DR wizard 641
 - configuring replication with DR wizard 627
 - creating temporary storage (array-based replication) 610
 - DR wizard overview 601
 - DR wizard requirements 601
 - multiple sites 647
- disk groups
 - cloning for secondary site (array-based replication) 610
 - cloning for secondary site (VVR replication) 605
 - creating
 - any-to-any HA 316
 - campus cluster 395
 - HA 67, 237
 - standalone to HA 145
 - creating a cluster disk group (RDC) 477
 - deporting
 - any-to-any HA 322
 - campus cluster 401
 - HA 73, 211, 302, 484
 - importing
 - any-to-any HA 322
 - campus cluster 401
 - HA 73, 211, 302, 484
 - overview
 - any-to-any HA 314
 - campus cluster 392
 - HA 65, 212
 - standalone to HA 142

- preconditions for a cluster disk group (RDC) 476
- disk groups and volumes
 - configuring
 - any-to-any HA 314
 - campus cluster 392
 - campus cluster for site A and site B 394
 - HA 65, 212
 - standalone to HA 142
 - managing
 - any-to-any HA 322
 - campus cluster 401
 - HA 73, 211, 302, 484
- disk space requirements
 - DR 556, 709
 - HA 48
 - standalone to HA 124
- DR
 - adding a new failover node
 - DR 651, 764
 - any-to-any
 - sample configuration 661
 - components
 - configuring on primary and secondary sites 735
 - copying the .crk file to the primary site 727
 - defined 549
 - disk space requirements 556, 709
 - new configuration 709
 - new installation 705
 - process overview 705
- DR wizard
 - cloning secondary storage (array-based replication) 610
 - cloning secondary storage (VVR replication) 605
 - cloning SQL Server 2005 service group 624
 - configuring replication and GCO 627
 - overview 601
 - requirements 601
- driver signing options
 - any-to-any HA 308
 - resetting 64, 142, 217, 308, 374, 459, 575
 - HA 64

E

- EMC SRDF
 - configure SRDF replication with the DR wizard 635

- requirements for DR wizard 596
 - Exchange
 - converting standalone servers to HA 121
 - disaster recovery overview 549
 - HA configurations 44
 - high availability overview 43
 - installing on the first node 486
 - tasks for setting up a secondary site 509
 - Exchange databases
 - moving from standalone to shared storage
 - standalone to HA 179
 - moving to shared storage
 - any-to-any HA 328
 - campus cluster 407
 - HA 98, 248, 274, 490
 - Exchange disk group, backing up and restoring (DR) 727
 - Exchange high availability, VCS application
 - agent 44, 353
 - Exchange installation
 - additional nodes
 - any-to-any HA 332, 335
 - campus cluster 412, 414
 - DR 621, 724
 - DR secondary site 619, 721
 - HA 103, 105, 253, 255, 495, 497
 - standalone to HA 184, 187
 - first node
 - any-to-any HA 327
 - campus cluster 402, 406
 - DR primary site 716
 - DR secondary site 617, 668, 716, 719
 - HA 93, 97, 244, 247, 273, 489
 - first node and additional nodes
 - DR secondary site 716
 - new any-to-any nodes
 - any-to-any HA 323
 - Exchange post-installation
 - additional nodes
 - any-to-any HA 336
 - campus cluster 415
 - DR 622, 725
 - HA 106, 256, 498
 - standalone to HA 188
 - first node
 - any-to-any HA 327
 - campus cluster 406
 - DR secondary site 618, 669, 720
 - HA 97, 247, 273, 489
 - Exchange pre-installation
 - additional nodes
 - any-to-any HA 333
 - campus cluster 412
 - DR 619, 722
 - HA 103, 253, 495
 - standalone to HA 185
 - first node
 - any-to-any HA 324
 - campus cluster 404
 - DR secondary site 615, 666, 718
 - HA 94, 245, 271, 487
 - Exchange service group
 - configuring
 - any-to-any HA 258, 339
 - campus cluster 417
 - DR 728
 - HA 108, 500
 - standalone to HA 190
 - configuring for an additional Exchange virtual server
 - HA 280
 - IP resource, modifying 424
 - prerequisites
 - any-to-any HA 258, 280, 339
 - campus cluster 417
 - DR 728
 - HA 108, 500
 - standalone to HA 190
- ## F
- Fire Drill Wizard
 - actions 688
 - changing a fire drill configuration 696
 - deleting the configuration 700
 - overview 678
 - preparing the configuration 688
 - prerequisites for a fire drill 686, 687, 688
 - restoring the prepared configuration 699
 - running a fire drill 694
 - forceimport
 - attribute for campus cluster 365
 - defined 365
 - setting after a site failure 427
 - forest and domain, preparing for Exchange
 - campus cluster 402
 - HA 58, 211, 302, 451

G**GCO**

- adding a remote cluster to a local cluster 752
- bringing a global service group online 755
- configuring for wide-area failover 751
- converting a local service group to a global service group 753
- defined 751
- prerequisites 736, 751

Global Cluster Option

- secure configuration 645, 674, 762

global cluster option

- overview 751
- see also GCO

Global Cluster Option (GCO)

- configuring with the DR wizard 641
- prerequisites 641

global service group

- defined 751

H**HA**

- defined 43
- disk space requirements 48
- installing
 - any-to-any 302
 - campus cluster 369
 - HA configuration 58, 244
 - standalone Exchange server
 - conversion 136
- IPaddresses required 55
- new configuration 53
- new installation 45
- process overview 45
- Sample configuration 56

HA configurations, Exchange 44**hardware, configuring 450****high availability**

- defined 43
- new installation 45
- overview 43, 44, 353
- see also HA

Hitachi TrueCopy

- configure replication with the DR wizard 638
- requirements for DR wizard 598

I**installing**

- Exchange on the first node 486
- Exchange, secondary site tasks 509
- SFW HA (RDC) 452
- verifying the installation 508

L**LLT over ethernet**

- configuring using VCW 79, 160, 222, 379, 464, 580

LLT over UDP

- configuring using VCW 80, 161, 223, 380, 465, 581

M**multiple DR sites 647****N****network and storage, configuring**

- any-to-any HA 300
- campus cluster 367
- DR (primary site) 55, 569
- HA 210
- standalone to HA 134

network, configuring 450**P****permissions requirements 51, 127, 208, 296, 361, 446, 560****prerequisites**

- Exchange service group
 - any-to-any HA 258, 280, 339
 - campus cluster 417
 - DR 728
 - HA 108, 500
 - standalone to HA 190

R**replicated data clusters**

- setting up 434

replicated data clusters (RDC)

- setting up a secondary Exchange site 509

replication

- configuring for EMC SRDF with the DR wizard 635
- configuring for HTC with the DR wizard 638
- configuring for VVR with DR wizard 627

- setting up a Replicated Data Set (RDS) 516
- requirements
 - any-to-any HA installation 205, 292
 - DR new installation 709
 - HA new installation 48
 - permissions 51, 127, 208, 296, 361, 446, 560
 - see also prerequisites
 - SQL 2008 software 49, 125, 206, 293, 358, 443, 557
 - standalone Exchange server to HA 123
 - VSW HA standalone 124
- requirements, additional for SFW HA 52, 128, 209, 297, 362, 447, 561
- requirements, network 50, 126, 207, 295, 360, 445, 559
- requirements, system 50, 126, 207, 295, 360, 445, 559
- resetting
 - driver signing options 64, 142, 217, 308, 374, 459, 575

S

- secondary site
 - setting up for disaster recovery 601
- secure clusters
 - assigning user privileges 600, 663
- secure GCO, establishing 645, 674, 762
- Security Services
 - configuring 82, 163, 225, 382, 466, 583
- setting bandwidth
 - using RDS wizard 632
- SFW for HA
 - installing (RDC) 452
- SFW HA
 - additional requirements 52, 128, 209, 297, 362, 447, 561
 - best practices 52, 128, 209, 297, 362, 447, 561
 - installing
 - HA 58
 - network requirements 50, 126, 207, 295, 360, 445, 559
 - system requirements 50, 126, 207, 295, 360, 445, 559
- site failure, forceimport attribute 427
- software requirements
 - SQL 2008 49, 125, 206, 293, 358, 443, 557
- Solutions Configuration Center
 - context sensitivity 29
 - overview 27

- running wizards remotely 35
- starting 28
- wizard descriptions 35
- SRDF
 - configuring replication with the DR wizard 635
- standalone Exchange conversion
 - HA disk space requirements 124
- standalone Exchange server
 - adding a new node
 - HA 155
 - adding nodes
 - HA 156
 - adding nodes to an existing cluster
 - HA 173
 - adding to a cluster
 - HA 155
 - configuration
 - HA 129
 - converting
 - to HA 121
 - converting to a “clustered” Exchange server
 - HA 152
 - creating a new cluster
 - HA 156
 - prerequisites for a new cluster
 - HA 155
 - prerequisites for adding nodes to an existing cluster
 - HA 172
 - process overview
 - HA 121
- storage cloning with the DR wizard
 - for array-based replication 610
 - for VVR replication 605
- switching online nodes 544

U

- user privileges
 - assigning 600, 663

V

- VCS
 - configuring the cluster for RDC 459
 - switching online nodes 544
- VCS Application Agent 44, 353
- verifying
 - cluster configuration for campus cluster 426

- cluster configuration for DR 734
- cluster configuration for HA 115, 197, 265, 287, 346
- volumes
 - creating
 - any-to-any HA 318
 - campus cluster 397
 - HA 69, 239
 - standalone to HA 147
 - creating on a cluster disk group 479
 - mounting
 - any-to-any HA 322
 - campus cluster 401
 - HA 73, 211, 302, 484
 - overview
 - any-to-any HA 314
 - campus cluster 392
 - HA 65, 212
 - standalone to HA 142
 - preconditions on a cluster disk group 476
 - unmounting
 - any-to-any HA 322
 - campus cluster 401
 - HA 73, 211, 302, 484
- VSW HA
 - installing
 - any-to-any HA 302
 - campus cluster 369
 - HA 244
 - standalone to HA 136
- VVR
 - configuring replication with DR wizard 627
 - creating replicator log volumes 736
 - creating the VVR RVG service group 748
 - prerequisites 736
 - setting up RDS (VCS) 516
 - setting up the replicated data sets 736
 - VxSAS service 593
- VxSAS service 593