# Veritas Storage Foundation™ and High Availability Solutions HA and DR Solutions Guide for Microsoft Exchange Server 2007

Windows Server 2003 (x64)

5.0

# Veritas Storage Foundation and HA Solutions HA and DR Solutions Guide for Microsoft Exchange Server 2007

Copyright © 2007 Symantec Corporation. All rights reserved.

Storage Foundation 5.0 for Windows HA

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

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

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

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

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

## Third-party legal notices

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

Windows is a registered trademark of Microsoft Corporation.

## Licensing and registration

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

## Technical support

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

# Contents

## Chapter 4 Deploying SFW HA for high availability: Configuring a new any-to-any failover

# Section 3 Disaster Recovery

## Chapter 6 Disaster recovery for Exchange: Overview



## Chapter 7 Deploying Disaster Recovery: New Exchange Server installation

Section 1

# Introduction

This section contains the following chapter:

Chapter 1

# Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange

This chapter includes the following topics:



## About the solutions guides

This guide contains solutions for Exchange 2007:

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

Separate guides are available 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.

# About disaster recovery

Wide area disaster recovery (DR) provides the ultimate protection for data and applications in the event of a disaster. If a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away. Veritas Storage Foundation HA for Windows (SFW HA) provides the capability for implementing disaster recovery.

# About SFW HA support for Exchange Server 2007

Support for Microsoft Exchange Server 2007 has been released in the form of patches. These patches are supported with SFW HA 5.0 only.

### SFW patch for Exchange Server 2007

This patch enables SFW support of VSS-based backup and restore operations with Exchange Server 2007.

Refer to the readme file accompanying this patch for more information on installation, supported features, and known issues.

### SFW HA patch for Exchange Server 2007

This patch contains the Veritas Cluster Server agent for Exchange Server 2007. The agent enables you to make Exchange Server 2007 highly available in a VCS cluster environment. The following features are provided with this patch:

- HA support for Mailbox Server role only
  High availability support for Exchange Server 2007 is available for the Mailbox Server role only. While installing Exchange, ensure that you do not install any other server role on the system on which you install the Mailbox Server role. If you have already installed the Mailbox Server role along with the other server roles on the same server, you will have to remove the other server roles before configuring Exchange in a SFW HA environment.

- Exchange Management Shell in the virtual server context

  The Exchange Management Shell provides a command-line interface that enables automation of administrative tasks for Exchange Server 2007. SFW HA provides a shortcut to launch the Exchange Management Shell under the context of the virtual server name.

  To launch the Exchange Management Shell in the virtual server context, click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server 2007 > Exchange Server 2007 Shell Launcher**.

  You must run the Exchange Management Shell under the virtual server context if you wish to administer a clustered Exchange Server 2007 using cmdlets. Ensure that the Exchange service group is online before using the Exchange Management Shell in the virtual server context.

  ---

  **Note:** The Exchange Management Shell in the virtual server context is provided to run cmdlets for administering Exchange in a VCS cluster environment only. Do not run SFW HA executable files or commands in this shell.

  ---

  Refer to the Exchange Server 2007 documentation for more information on server roles, the Exchange Management Shell and cmdlets.

Refer to the readme files accompanying the patches for the latest information on supported features and known issues.

Contact Symantec technical support for further information.

See http://entsupport.symantec.com.

# How this guide is organized

This guide contains sections on High Availability and Disaster Recovery. Each section contains chapters that describe setup, installation, and configuration information for specific Exchange Server 2007 HA and DR configurations in an SFW HA environment.

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 Appendix section includes the VCS agent for Exchange Server 2007 resource type definitions, attribute descriptions, service group dependency graph, and sample configuration.

Section 2

# High Availability

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

Refer to the following chapters to install and configure a clustered Exchange Server 2007 environment using Veritas Storage Foundation HA for Windows:

Chapter 2

# High availability for Exchange: Overview

This chapter covers the following topics:



## 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.

**Note:** HA for Exchange 2007 is available for the Mailbox Server role only.

# Typical HA configurations for Exchange

Typical HA configurations for Exchange are as follows:

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

Chapter 3

# Deploying SFW HA for high availability: Configuring a new active/passive failover

This chapter covers the following topics:

# Tasks for a new HA installation of Microsoft Exchange Server 2007

You can install and configure a new Veritas Storage Foundation high availability environment for Exchange Server 2007. This environment involves an active/passive configuration with one-to-one failover capabilities.

The procedures are slightly different if you are setting up an any-to-any configuration.

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

For information on setting up a secondary site for disaster recovery after configuring high availability, see the disaster recovery section.

**Warning:** For the SFW HA patch release for Exchange 2007, you cannot use the Solutions Configuration Center to configure Exchange 2007 for high availability or disaster recovery. Instead use the instructions in this Solutions Guide.

Table 3-1 outlines the high-level objectives and the tasks to complete each objective.

**Table 3-1** Task list for a new HA installation of Exchange 2007

| Objective | Tasks |
|---|---|
| "Reviewing the requirements" on page 26 | ■ Verify hardware and software prerequisites |
| "Reviewing the configuration" on page 31 | ■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster |
| "Configuring the storage hardware and network" on page 34 | ■ Set up the network and storage for a cluster environment<br>■ Verify the DNS entries for the systems on which Exchange will be installed |

**Table 3-1** Task list for a new HA installation of Exchange 2007 (continued)

| Objective | Tasks |
|---|---|
| "Configuring SFW HA: Prior to installing Exchange" on page 36 | ■ Verify the driver signing options for Windows 2003 systems<br>■ Install SFW HA<br>■ Restore driver signing options for Windows 2003 systems |
| "Installing the SFW HA patch for Exchange Server 2007" on page 54 | ■ Install the SFW HA patch; involves installing the VCS agent for Exchange Server 2007<br>■ If required, install the SFW patch to perform VSS-based backup and restore operations with Exchange Server 2007 |
| "Configuring disk groups and volumes" on page 44 | ■ Use the VEA console to create disk groups<br>■ Use the VEA console to create the data, log, RegRep, and Shared volumes<br>■ Manage disk groups and volumes, with instructions for mounting and unmounting volumes |
| "Configuring the cluster" on page 55 | ■ Verify static IP addresses and name resolution configured for each node<br>■ Configure cluster components using the Veritas Cluster Server Configuration Wizard |
| "Installing Exchange on the first node" on page 71 | ■ Review the prerequisites checklist<br>■ Run the Exchange Setup Wizard for Veritas Cluster Server and Microsoft Exchange Server installation |

**Table 3-1** Task list for a new HA installation of Exchange 2007 (continued)

| Objective | Tasks |
|---|---|
| "Moving Exchange databases to shared storage" on page 75 | ■ 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 79 | ■ Run the Exchange Setup Wizard for Veritas Cluster Server and Microsoft Exchange Server installation for additional nodes |
| "Configuring the Exchange service group for VCS" on page 83 | ■ Create the Exchange service group using the Exchange Configuration Wizard for Veritas Cluster Server |
| "Verifying the cluster configuration" on page 90 | ■ Verify the cluster configuration by switching service groups and shutting down an active cluster node |

# Reviewing the requirements

Before installation, review these product installation requirements for your systems. Minimum requirements and Symantec recommended requirements may vary.

## Disk space requirements

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

Table 3-2 estimates disk space requirements for SFW HA.

**Table 3-2** Disk space requirements

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

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

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

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

## Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA) with the Veritas Cluster Server Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:
  - Microsoft Exchange Server 2007 Standard Edition or Enterprise Edition
    with
    Windows Server 2003 x64 Standard Edition, Enterprise Edition, Datacenter Edition (SP1 required for all editions, SP2 supported)
    or
    Windows Server 2003 R2 x64 Standard Edition, Enterprise Edition, Datacenter Edition

## System requirements

Systems must meet the following requirements:

- Processor: x64 architecture-based computer with Intel processor that supports Intel Extended Memory 64 Technology (Intel EM64T) or AMD processor that supports the AMD64 platform; Intel Itanium family IA64 processors are not supported.
- Memory: minimum 2 GB of RAM per server.
- File format: Disk partitions must be formatted for the NTFS file system.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs.
  See “Best practices” on page 30.

- All servers must run the same operating system, service pack level, and system architecture.

## Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 domain.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.
- Static IP addresses for the following purposes:
  - One static IP address available per site for each Exchange Virtual Server (EVS)
  - One IP address for each physical node in the cluster
  - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
  - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
  Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
  See the *Veritas Cluster Server Bundled Agents Reference Guide*.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

## Permission requirements

This section lists the following permission requirements:

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

- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

## Additional requirements

Please review the following additional requirements:

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

## Best practices

Symantec recommends that you perform the following tasks:

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

# Reviewing the configuration

The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server if the active node fails.

In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. In this case, EVS1 can fail over from SYSTEM1 or SYSTEM2.

Figure 3-1 illustrates an active/passive failover configuration with an Exchange virtual server.

**Figure 3-1** Active/Passive failover configuration

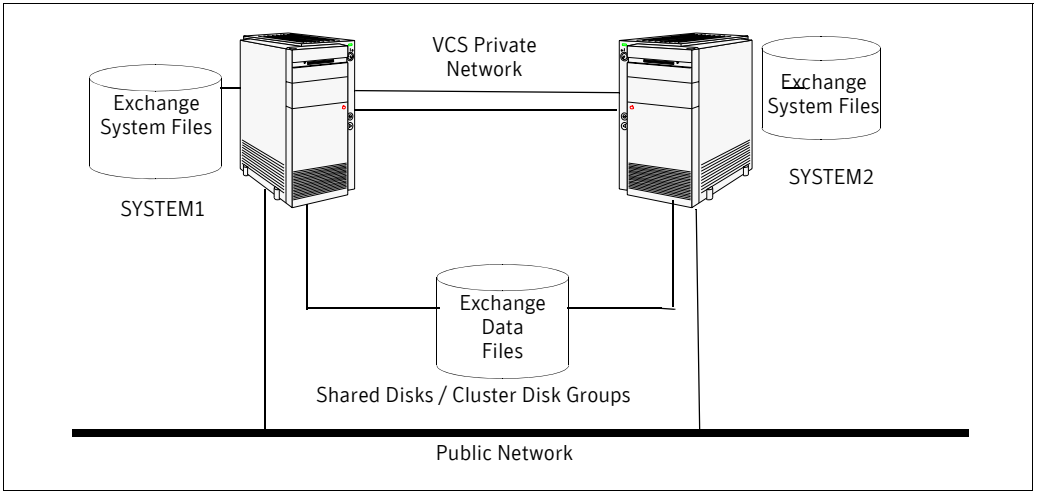


In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails.

Figure 3-2 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 3-2** Disaster Recovery environment

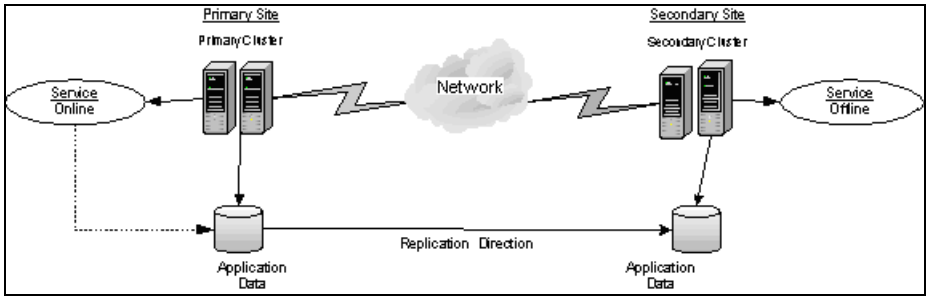


When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients. Figure 3-3 illustrates this type of failure:

**Figure 3-3** Failure in a disaster recovery environment

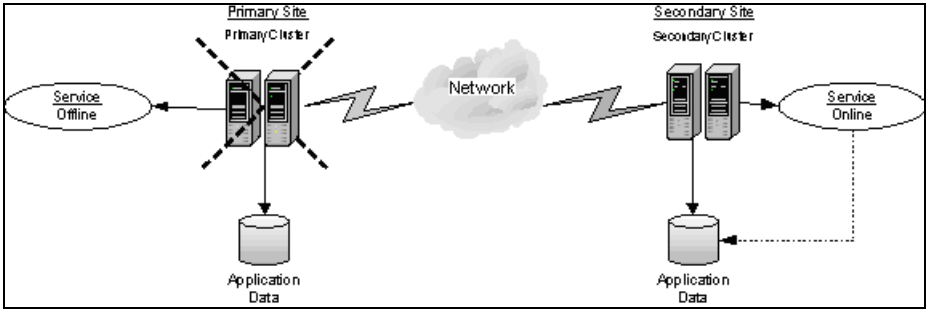

## IP addresses required during configuration

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

| | |
|---|---|
| Exchange virtual server | The virtual IP address for the Exchange server.<br>For a disaster recovery configuration, the virtual IP address for the Exchange server at the primary and disaster recovery site can only be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site. |
| Cluster IP address | Used by Veritas Cluster Management Console (Single Cluster Mode), also referred to as Web Console.<br>Used by VCS notifier.<br>For a disaster recovery configuration, used by the Global Cluster Option.<br>For a disaster recovery configuration, a separate IP address is required for the secondary site. |
| Replication IP address (disaster recovery configuration only) | For a disaster recovery configuration, an IP address is required for each Replicated Data Set (RDS), one for the for the primary site and one for the secondary site.<br>Two IP addresses are required per Replicated Volume Group (RVG). |

## Sample configuration

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

**Table 3-3** Sample configuration

| Name | Object |
|---|---|
| SYSTEM1, SYSTEM2 | Physical node names |
| EVS1 | Microsoft Exchange Virtual Server |
| EVS1_SG1 | Microsoft Exchange service group |

**Table 3-3** Sample configuration (continued)

| Name | Object |
|---|---|
| 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_REGREP | Volume that contain the list of registry keys that must be replicated among cluster systems for the Exchange server |

# Configuring the storage hardware and network

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

### To configure the hardware

1. Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
2. Connect the network adapters on each system.
   - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
   - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
3. Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
4. Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk.
   Use Windows Disk Management (**Start > Control Panel > Administrative Tools > Computer Management)** or Veritas Enterprise Administrator (VEA) on each system to verify that the attached shared disks are visible.

**To verify the DNS settings and binding order for all systems**

1. Open the Control Panel (**Start>Control Panel**).
2. Right-click on Network Connections and click **Open**.
3. Ensure the public network adapter is the first bound adapter:
   - From the Advanced menu, click **Advanced Settings**.
   - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
   - Click **OK**.
4. In the Network and Dial-up Connections window, double-click the adapter for the public network.
   When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
5. From the status window, click **Properties**.
6. In the General tab:
   - Select the **Internet Protocol (TCP/IP)** check box.
   - Click **Properties**.
7. Select the **Use the following DNS server addresses** option.
8. Verify the correct value for the IP address of the DNS server.
9. Click **Advanced**.
10. In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
11. Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
12. Click **OK**.

# 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 36.
- Set up a VCS environment.
  See "Configuring the cluster" on page 55
- Create the required disk groups and volumes.
  See "Configuring disk groups and volumes" on page 44
  See "Managing disk groups and volumes" on page 52.

# Installing Veritas Storage Foundation HA for Windows

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

## Setting Windows driver signing options

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

Table 3-4 describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

**Table 3-4** Installation behavior with unsigned drivers

| Driver Signing Setting | Installation behavior on the local system | Installation behavior on remote systems |
|---|---|---|
| Ignore | Always allowed | Always allowed |
| Warn | Warning message, user interaction required | Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation. |

**Table 3-4** Installation behavior with unsigned drivers (continued)

| Driver Signing Setting | Installation behavior on the local system | Installation behavior on remote systems |
|---|---|---|
| Block | Never allowed | Never allowed |

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

**To change the driver signing options on each system**

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

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

**To install the product**

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

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


4 Do one of the following:

- Click **Complete/Custom** to begin installation.
- Click the **Administrative Console** link to install only the client components.

5 Review the Welcome message and click **Next**.

6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.

7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.

8 Repeat for additional license keys. Click **Next**

- To remove a license key, click the key to select it and click **Remove**.
- To see the license key's details, click the key.

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

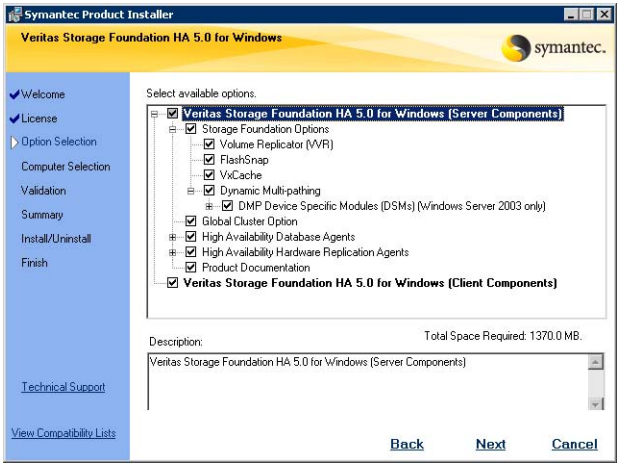


The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

| | |
|---|---|
| Veritas Volume Replicator | If you plan to use VVR for replication, you must also select the option to install VVR. |
| Global Cluster Option | Required for a disaster recovery configuration only. |
| Client Components | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console. |

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

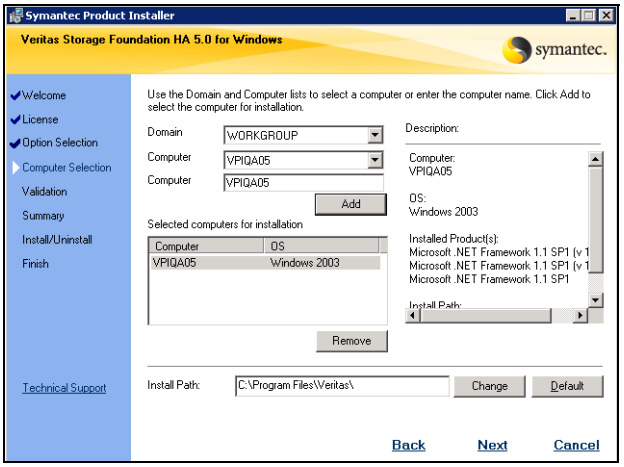


| | |
|---|---|
| Domain | Select a domain from the list.<br>Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate. |
| Computer | To add a computer for installation, select it from the Computer list or type the computer's name in the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

11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.

12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.

13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
- For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
- For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.

14 Click **OK**.

15 Review the information and click **Install**. Click **Back** to make changes, if necessary.

16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.

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

18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.

19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.

20 Review the log files and click **Finish**.

21 Click **Yes** to reboot the local node.

# Configuring VxSAS

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server 2003 environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

---

**Note:** The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

---

### To configure the VxSAS service

1. To launch the wizard, select **Start** > **All Programs** > **Symantec** > **Veritas Storage Foundation** > **Configuration Wizards** > **VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.
   Read the information provided on the Welcome page and click **Next**.
2. Complete the Account Information panel as follows:

   | | |
   |---|---|
   | Account name (domain\account) | Enter the administrative account name. |
   | 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.<br>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:

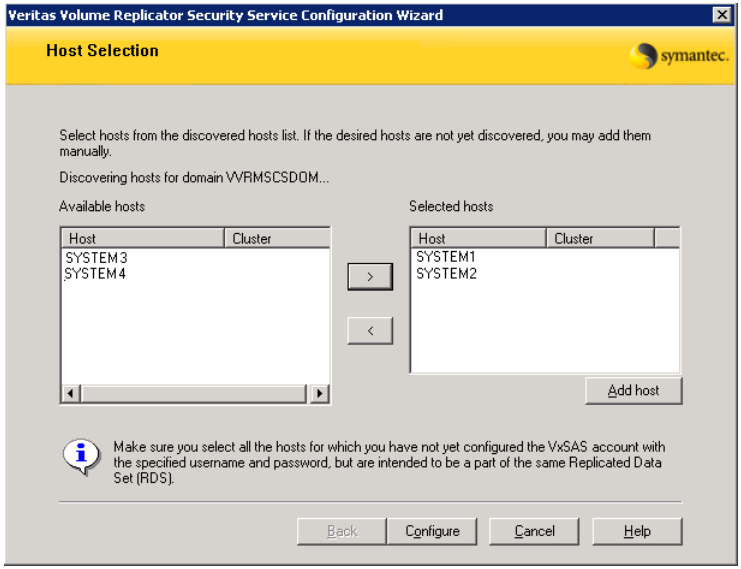


| | |
|---|---|
| Selecting hosts | The Available hosts pane lists the hosts that are present in the specified domain.<br>Move the appropriate name from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts. |

| | |
|---|---|
| Adding a host | If the host name you require is not displayed, 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 the **Configure** button 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 disk groups and volumes

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

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

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

- The size of databases and logs.

Typically, a SFW disk group corresponds to an Exchange storage group. Figure 3-4 displays a detailed view of the disk groups and volumes in an HA environment.

**Figure 3-4** Disk groups and volumes for Exchange virtual server EVS1 in HA setup

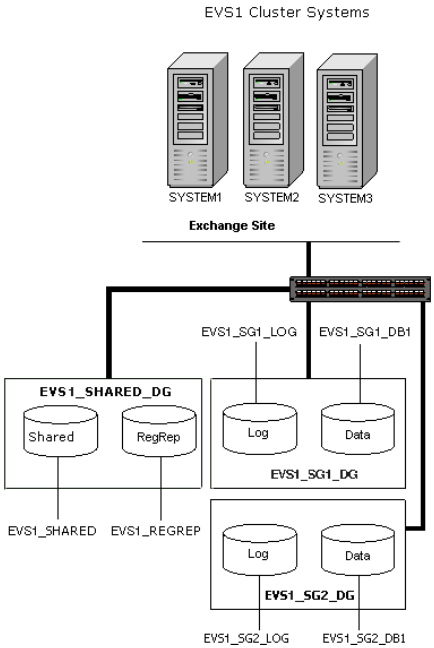


Use the following procedures to create the appropriate disk groups and volumes. The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

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 create contains the following:

- EVS1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.

For a disaster recovery configuration only, you will need a volume for the VVR Storage Replicator Log (SRL):

- EVS1_REPLOG

The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups.

Additional storage groups (for example, EVS1_SG2_DG) only contain the data, log, and VVR Storage Replicator Log volumes; the other volumes are included in the first storage group.

---

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

---

## Creating a disk group

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

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

6 Provide information about the cluster disk group:

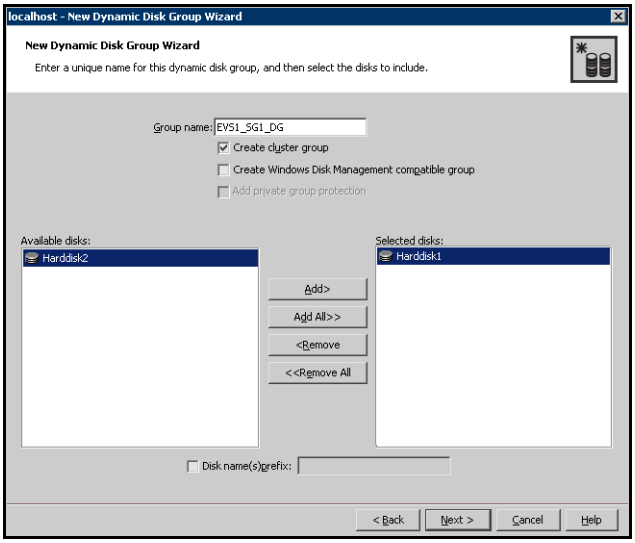


- Enter the name of the disk group (for example, EVS1_SG1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
  Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
- Click **Next**.

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the new disk group.

# Creating volumes

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

---

**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.

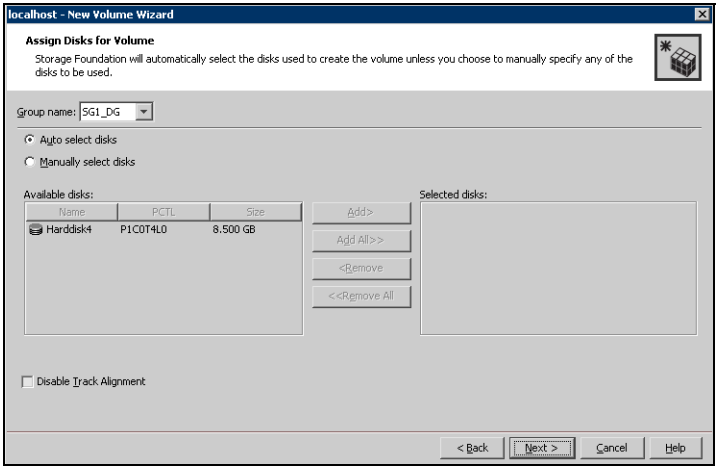


7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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.

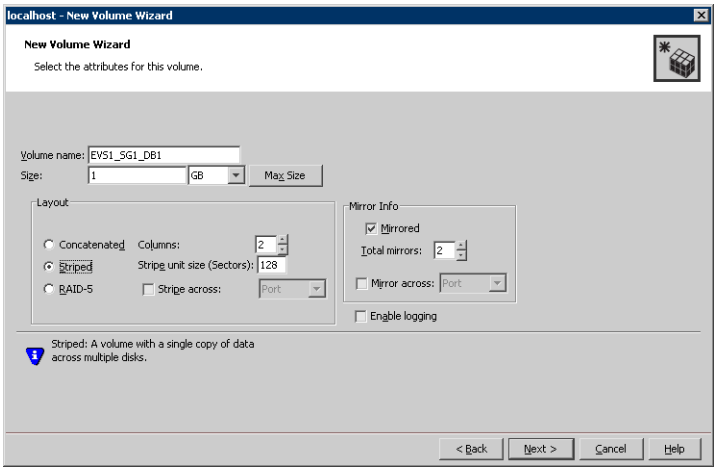


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
- If you are creating a striped volume, the **Columns**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

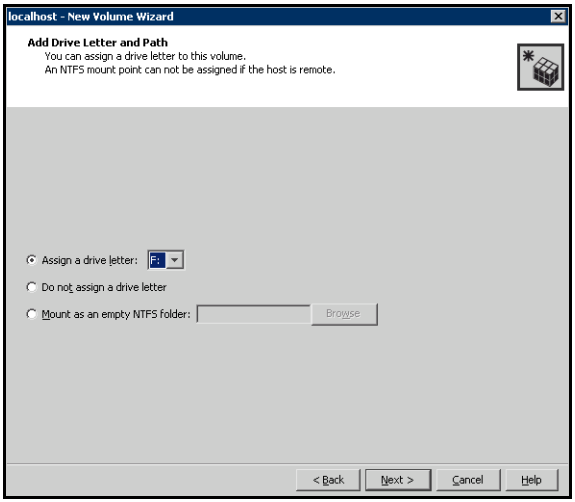


11 Click **Next**.

12 Create an NTFS file system.

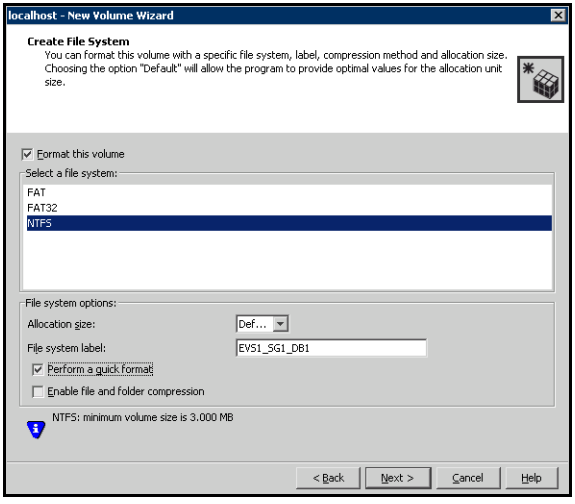


- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.

- 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 both the RegRep volume (EVS1_REGREP) and the EVS1_SHARED volumes in the EVS1_SHARED_DG disk group.

15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3).

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

If you are configuring an any-to-any environment, you can also create similar disk groups and volumes for the other Exchange servers. For example, create disk group (EVS2_SG1_DG) and volumes (EVS2_SG1_DB1, EVS2_REGREP, EVS2_SG1_LOG, and EVS2_SHARED).

# Managing disk groups and volumes

This section includes the following procedures:

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

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

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

## Importing a disk group and mounting a volume

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

**To import a disk group**

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

**To mount a volume**

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

## Unmounting a volume and deporting a disk group

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

**To unmount a volume and deport the dynamic disk group**

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

# Installing the SFW HA patch for Exchange Server 2007

The SFW HA patch for Exchange Server 2007 contains the new VCS agent for Exchange Server 2007. The patch zip file contains the following files:

- **vrtsvcsexch.msi** - the VCS agent msi file
- **InstallVCSExch2007.bat** - batch file to install the VCS agent
- **UninstallVCSExch2007.bat** - batch file to remove the VCS agent

Extract these files to a temporary location on the system. Ensure that the agent .msi file and the .bat file are at the same level in a directory.

### To install the SFW HA patch

1. If the system is part of a cluster, complete this step. If not, proceed to step 2.
   - Make sure that all the service groups are offline in the cluster.
   - Save and close the cluster configuration. Type the following on the command prompt:
     ```
     C:\> haconf -dump -makero
     ```
   - Stop the Veritas High Availability engine (HAD) on all the cluster nodes. Type the following on the command prompt:
     ```
     C:\> hastop -all
     ```
2. On the system, double-click **InstallVCSExch2007.bat**. The .msi will install the VCS agent for Exchange Server 2007 on the system.
3. Repeat step 2 on all the systems where you want to install the patch.
4. If the system is part of a cluster, complete this step.
   Start HAD on the system on which you installed the patch. Type the following on the command prompt:
   ```
   C:\> hastart
   ```

# Installing the SFW patch for Exchange Server 2007

The SFW patch for Exchange Server 2007 enables SFW support for performing VSS-based backup and restore operations with Exchange Server 2007.

This step is optional. If required, you can install the SFW patch now. Refer to the readme file accompanying the SFW patch for the installation steps.

# Configuring the cluster

After installing the VCS agent for Exchange Server 2007, set up the components required to run a cluster. The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures the Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

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

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

### To configure a VCS cluster

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

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

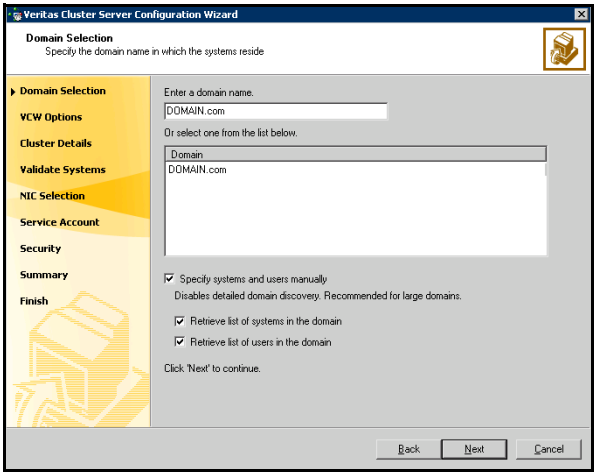


To discover information about all systems and users in the domain:

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

  Proceed to step 7 on page 58.

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 57. Otherwise proceed to the next step.

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

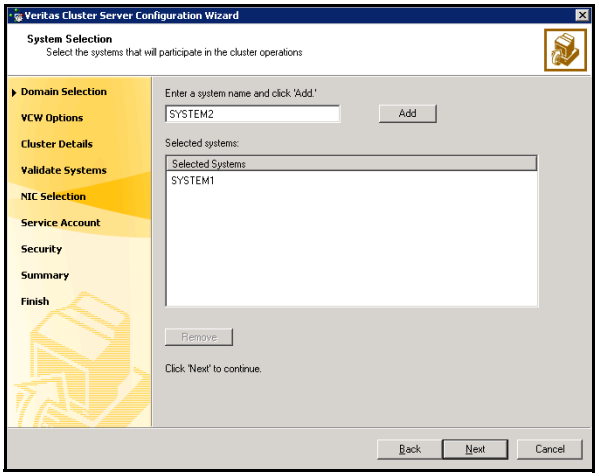


Proceed to step 7 on page 58.

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

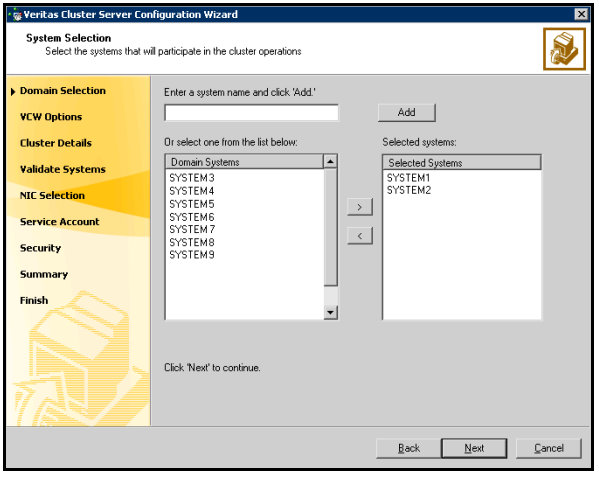


Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

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

8 On the Cluster Details panel, specify the details for the cluster and then click **Next**.

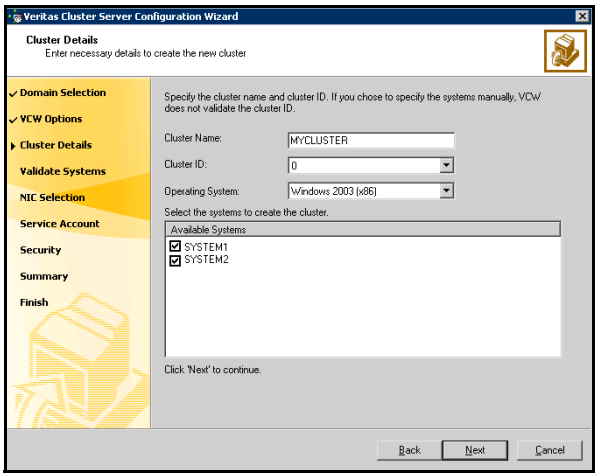


| | |
|---|---|
| Cluster Name | Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. |

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

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

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

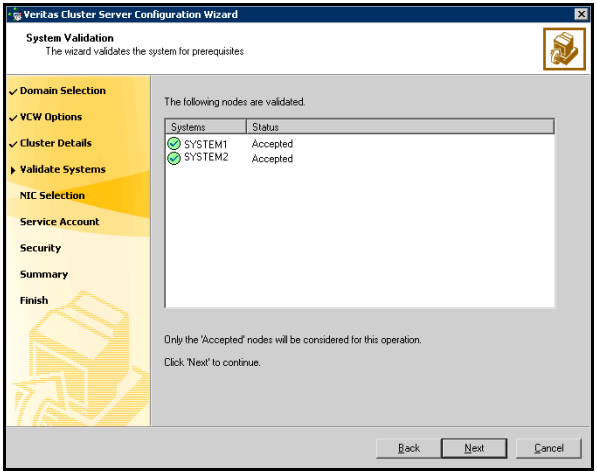


If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in step 8 on page 58, proceed to the next step. Otherwise, proceed to step 11 on page 60.

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

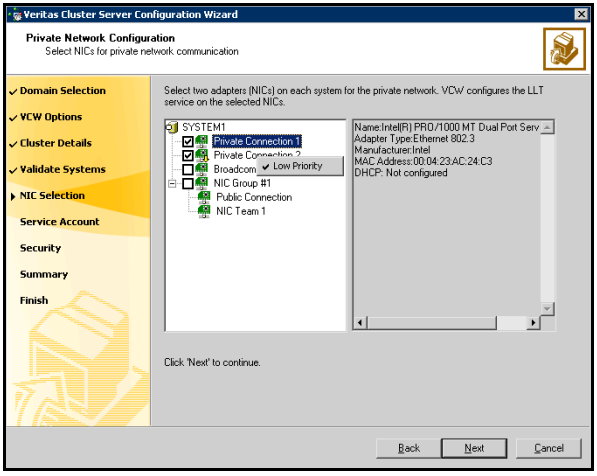


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of the NIC that is used for public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as NIC Group #N where N is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

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

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

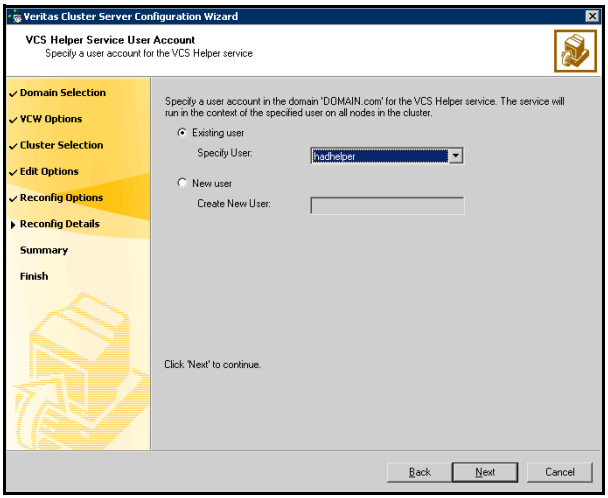


- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list
  - If you chose not to retrieve the list of users in step 4 on page 56, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

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

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

Do one of the following:

- To use the single sign-on feature

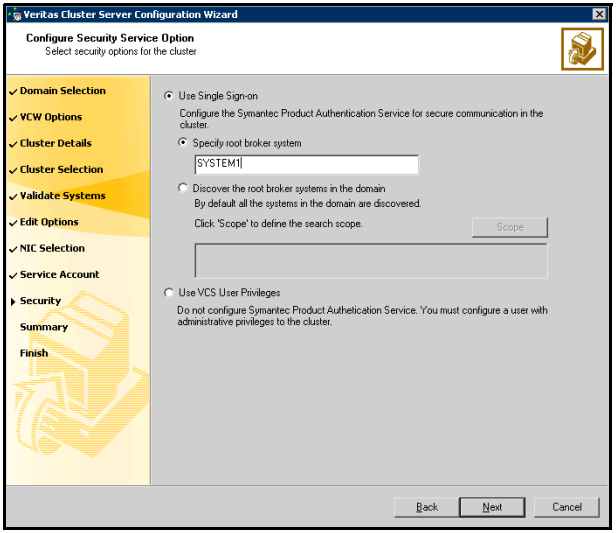


  - Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

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

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

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

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

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

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

- To use VCS user privilege

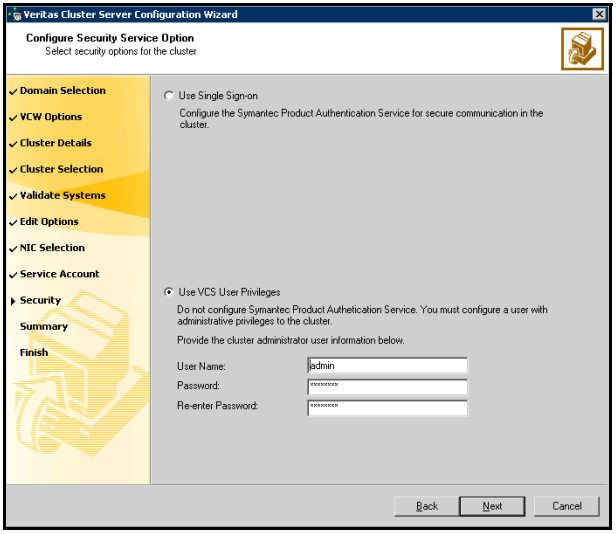


- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

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

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

- Click **Next**.

13 Review the summary information on the Summary panel, and click **Configure**.

The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.

The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

14 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Web Console, notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

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

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

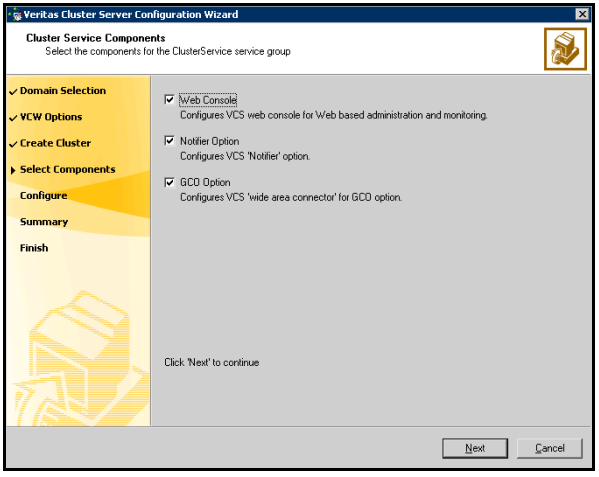

- Check the **Web Console** check box to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
  The GCO Option applies only if you are configuring a Disaster Recovery environment.

## Configuring the Web Console

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

### To configure the Web console

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

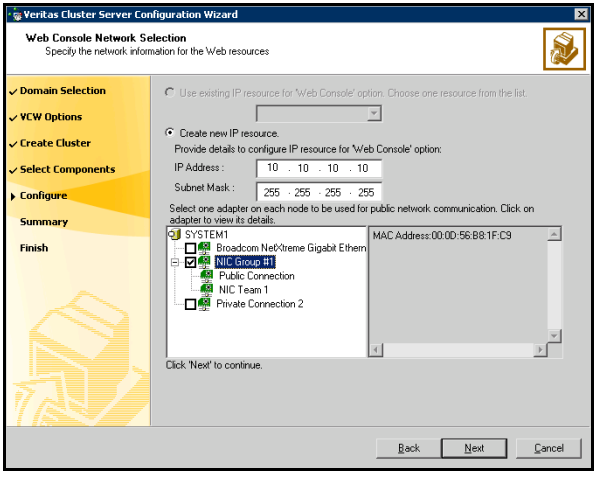


- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
- If you choose to configure a new IP address, type the IP address and associated subnet mask.

- Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.

2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.

3 If you chose to configure a Notifier resource, proceed to "Configuring notification" on page 66.
If you chose to configure global cluster components, proceed to "Configuring the wide-area connector process for global clusters" on page 70.
Otherwise, click **Finish** to exit the wizard.

## Configuring notification

This section describes steps to configure notification.

### To configure notification

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

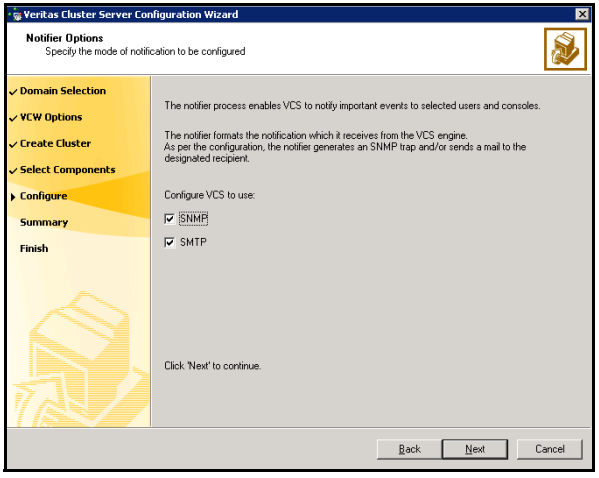


You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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

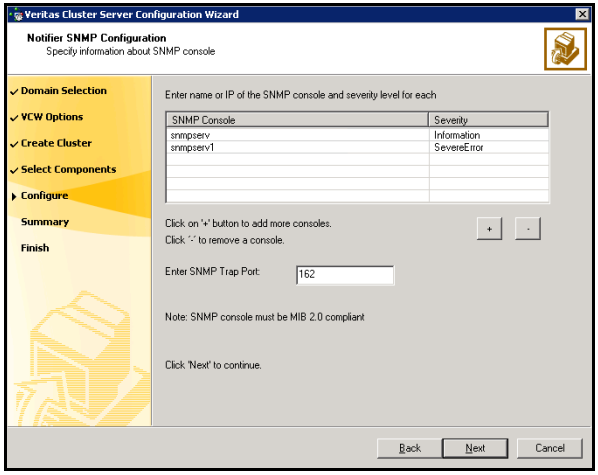


- Click a field in the SNMP Console column and enter the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click **+** to add a field; click **-** to remove a field.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

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

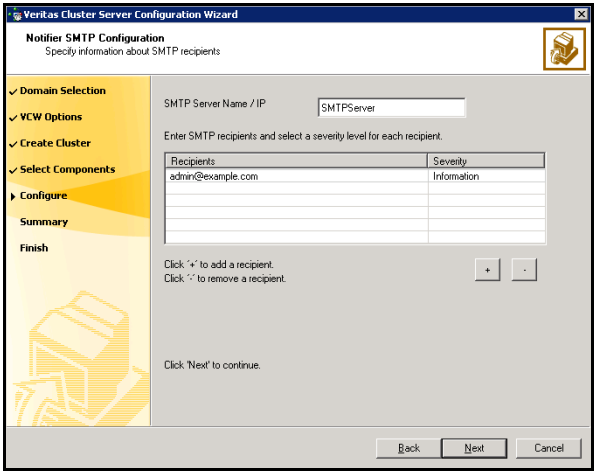


- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity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

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

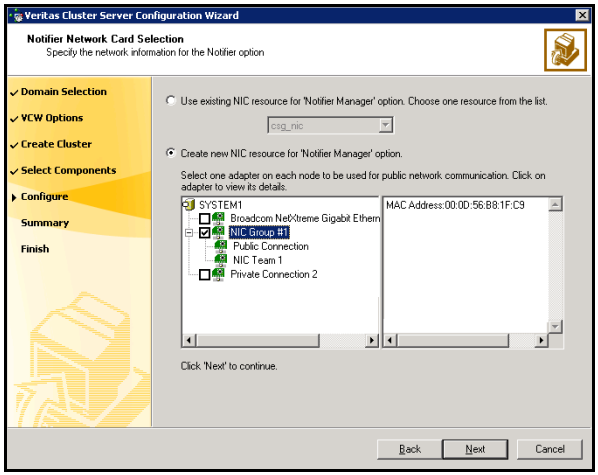


- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
- If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.

5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started and click **Configure**.

6 If you chose to configure global cluster components, proceed to "Configuring the wide-area connector process for global clusters" on page 70. Otherwise, click **Finish** to exit the wizard.

## Configuring the wide-area connector process for global clusters

This section describes steps to configure the wide-area connector resource required for global clusters.

### To configure the wide-area connector process for global clusters

1 On the GCO Network Selection panel, specify the network information and click **Next**.

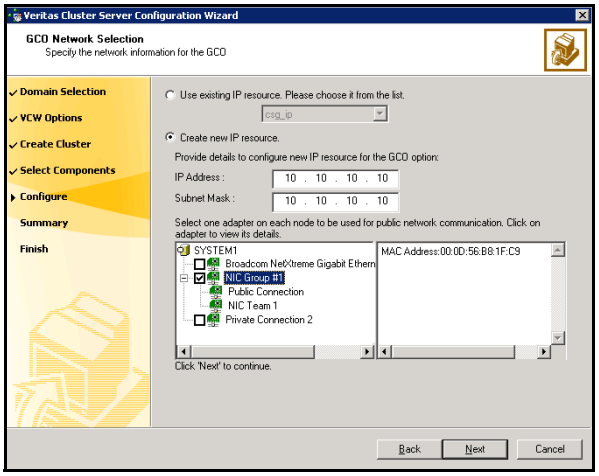


- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
- If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.

2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.

3 Click **Finish** to exit the wizard.

# 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:

- 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 52.
- Mount the volume containing the information for registry replication (EVS1_REGREP).
  See "Importing a disk group and mounting a shared volume" on page 52.
- Verify that all systems on which Exchange Server will be installed have IIS installed. You must install WWW services on all systems.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
  - You must be a domain user.
  - You must be an Exchange Full Administrator.
  - You must be a member of the Exchange Servers group.
  - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
  - You must have write permissions on the DNS server to perform DNS updates.
  - Make sure the HAD Helper domain user account has the "Add workstations to domain" privilege enabled in the Active Directory. To verify this, click Start > Administrative Tools > Local Security Policy on the domain controller to launch the security policy display. Click Local Policies > User Rights Management and make sure the user account has this privilege.
  - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
  - The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

# Exchange pre-installation: First node

Use the Exchange Setup Wizard for Veritas Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. You must install Exchange on a virtual node to facilitate high availability. After you have run the wizard, you will be requested to restart the node. So, close all open applications and save your data before running the wizard.

### To perform Exchange pre-installation

1. Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
2. Click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 2007 Mailbox Server role 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. Specify information related to the network.

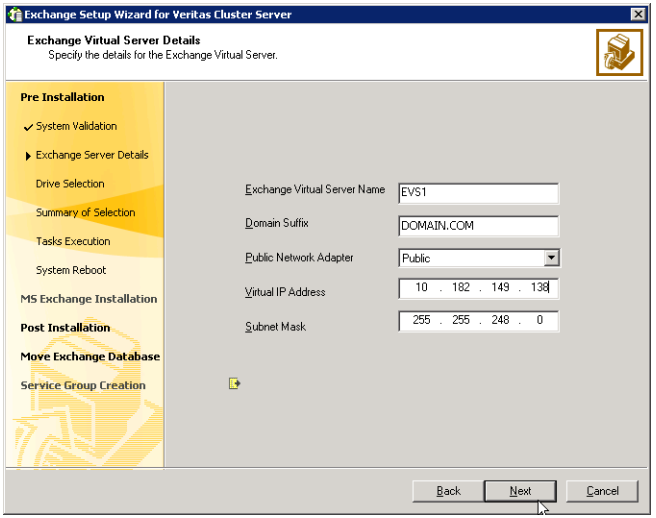


- Enter a unique virtual name for the Exchange server.

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

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

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

8 Review the summary of your selections and click **Next**.

9 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.

10 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**.

11 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.

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

## Exchange installation: First node

Install Exchange on the node on which you performed the pre-installation. HA support for Exchange Server 2007 is available for the Mailbox Server role. While installing Exchange, ensure that you install the Mailbox Server role only.

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.

## Exchange post-installation: First node

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

### To perform Exchange post-installation

1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.

2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 4.
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

3 Review the information in the Welcome dialog box and click **Next**.

4 A message appears indicating that the system will be renamed and restarted after you quit the wizard.This sets the node back to its physical host name. Click **Yes** to continue.

5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.

6 Click **Finish**.

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

8 Once the node is rebooted, move the databases created during the Exchange installation from the local drive to the shared drive.

# Moving Exchange databases to shared storage

After completing the Exchange installation and the post-installation tasks 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 to import the disk group and mount the volumes for the Exchange database, and transaction logs.
  See "Managing disk groups and volumes" on page 52.
- The Exchange Setup Wizard for VCS cannot move the Exchange storage groups until local continuous replication (LCR) is suspended for those storage groups. Please suspend LCR using the Exchange Management Console or the Exchange Management Shell, before moving the Exchange databases.
  Refer to the Microsoft Exchange documentation for information on how to suspend LCR.

### To move Exchange databases

1 Click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 Setup Wizard** to start the Exchange Setup Wizard for VCS.

2 Review the information in the Welcome dialog box and click **Next**.

3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.

4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.

5 In the Select Exchange Virtual Server dialog box:

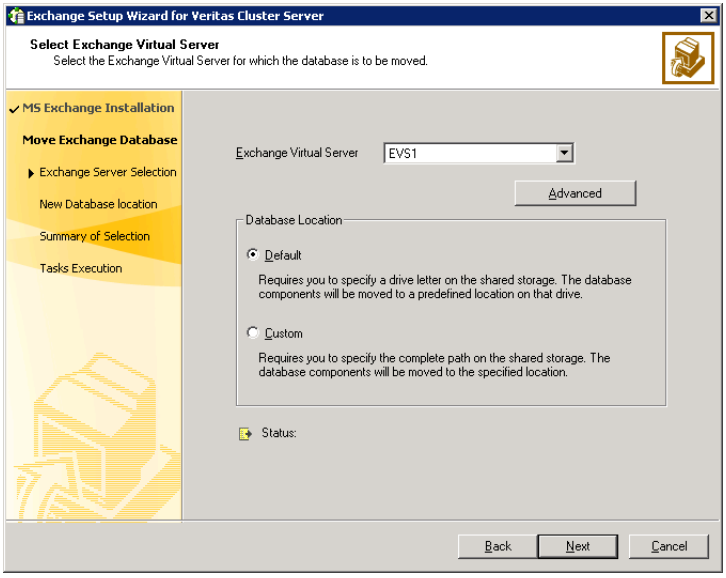


- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

---

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

---

- Click **OK**.
- Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
- Click **Next**.

6 Do one of the following:

- If you would like to move the default mailbox store, and the public store to the generated default paths on the volumes for these components, proceed to the next step.
- Otherwise, proceed to step 8 on page 78 to specify the path location on the volumes that you will designate for these components.

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

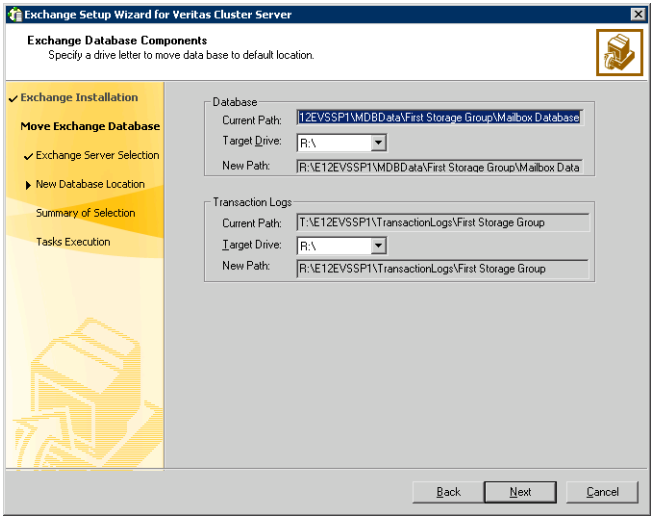


- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Click **Next** and proceed to step 9 on page 78.

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

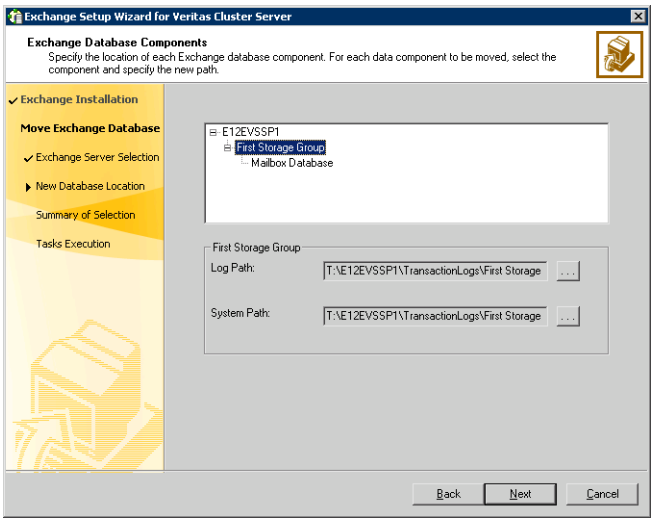


For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the paths for the Exchange database components are not the root of a drive. You must select a directory on the specified drive.

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

9 Review the summary of your selections and click **Next**.

10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.

11 Click **Finish** to exit the wizard.

# Installing Exchange on additional nodes

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

See "Installing Exchange on the first node" on page 71.

## 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 52.

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 Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 2007 Mailbox Server role 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 Specify network information for the Exchange virtual server.

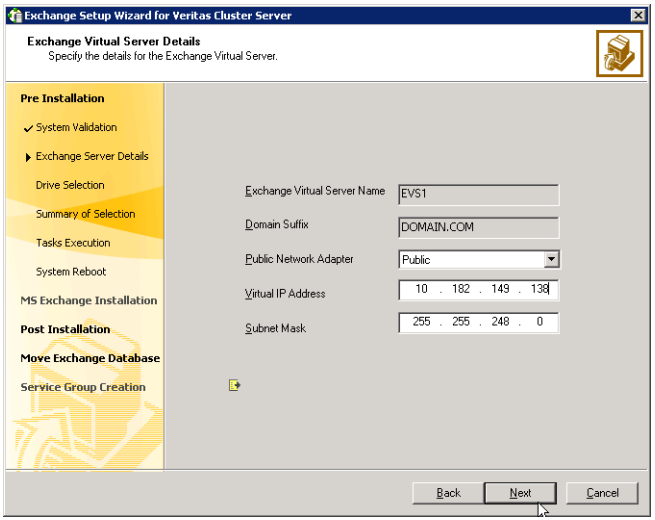


The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.

- 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**.

8 Review the summary of your selections and click **Next**.

9 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.

10 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**.

11 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.

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: additional nodes

- Install Exchange on the additional node on which you performed the pre-installation. HA support for Exchange Server 2007 is available for the Mailbox Server role. While installing Exchange, ensure that you install the Mailbox Server role only.
- Install the same Exchange version and components on all nodes.

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 RecoverServer as the install mode:
   ***<drive letter>***`:\setup.com /mode:recoverserver`
   where ***<drive letter>*** is the location where the Exchange software is located.
2. Setup copies the setup files locally to the computer on which you are installing Exchange 2007 and then checks the prerequisites, including all prerequisites specific to the server roles that you are installing. If you have not met all of the prerequisites, Setup fails and returns an error message that explains the reason for the failure. If you have met all of the prerequisites, Setup installs Exchange 2007.
3. Verify that the installation completed successfully. Refer to the Microsoft documentation for more information.

# Exchange post-installation: additional nodes

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

### To perform Exchange post-installation

1. Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
2. If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 4.
   If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
3. Review the information in the Welcome dialog box and click **Next**.
4. A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
5. The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
6. 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.

7 Click **Finish**.

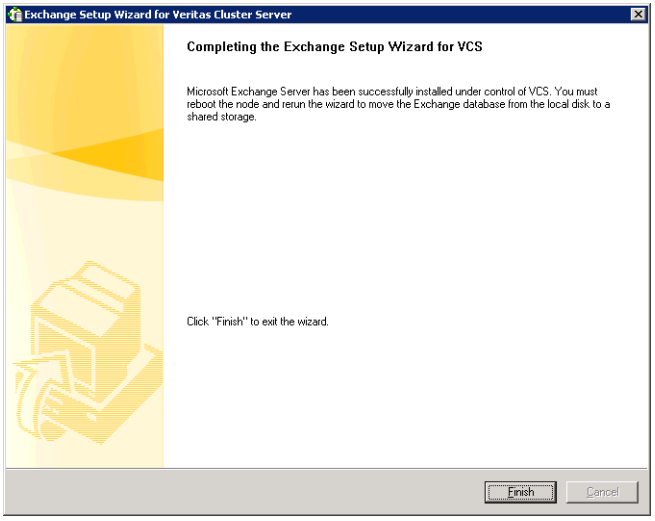


8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

9 Changes made during the post-installation steps do not take effect till you reboot the node.

# 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 start-up. Refer to the Exchange documentation for instructions.

## Prerequisites

- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.

- Verify that the Veritas High Availability Daemon (HAD) is running on the node on which you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk group and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
  - Exchange database
  - Registry changes related to Exchange
  - Transaction logs for the first storage group

  See "Importing a disk group and mounting a shared volume" on page 52 for instructions on mounting and "Unmounting a volume and deporting a disk group" on page 52 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 Appendix A, "VCS agent for Exchange Server 2007" for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* for more 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 Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 Specify the service group name and system list:

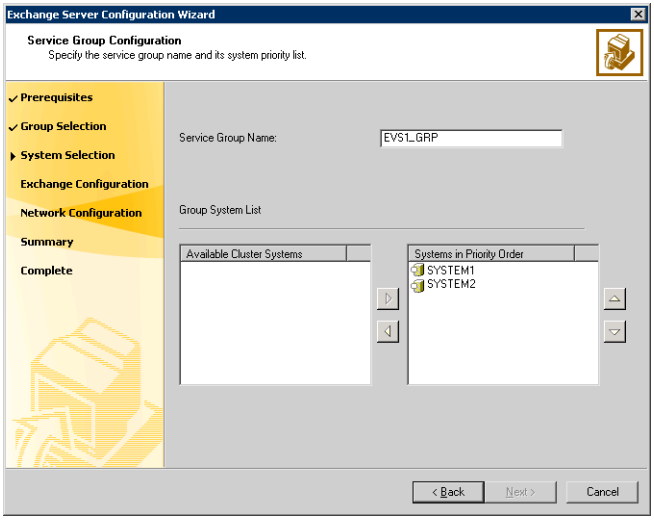


- Enter a name for the Exchange service group.
  If you are configuring the service group on the secondary site, ensure that the name matches the service group on the primary site.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow 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.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

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

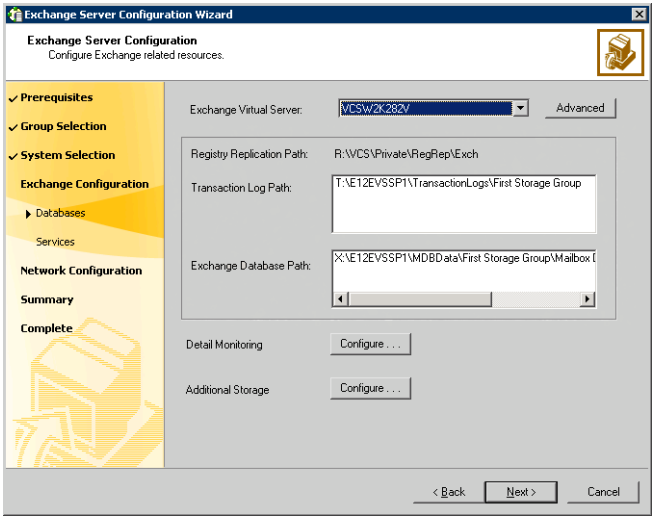


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
  - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.

- To configure Detail Monitoring for Exchange databases, click **Configure...**.

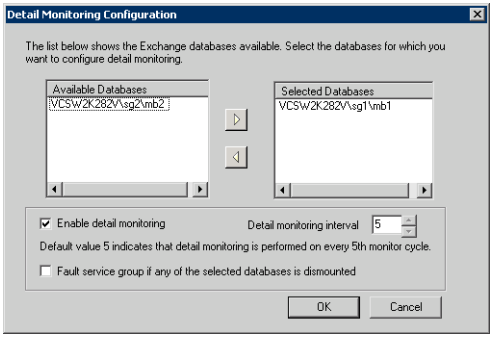


  On the Detail Monitoring Configuration dialog box, complete the following:

  - In the Available Databases box, select the databases for detail monitoring and double-click, or click the right-arrow button to move them to the Selected Databases box. To remove a database, select the database in the Selected Databases box, and double-click or click the left-arrow button.
  - Check **Enable detail monitoring** check box, and specify the monitoring interval in the **Detail monitoring interval** field.
  - If you want the VCS agent to fault the service group if a database selected for detail monitoring is dismounted, check the **Fault service group if any of the selected database is dismounted** check box.

  See the VCS agent attribute descriptions in the Appendix, for more information on detail monitoring and VCS agent behavior.

  - Click **OK**.

- To configure additional storage, click **Configure...**. On the Additional Storage Configuration dialog box, complete the following:

  - 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 (MSExchangeSA) service resource.

- Click **Next**.

6 Select the optional Exchange services to be monitored and click **Next**. Each optional service that is selected will be configured as a VCS resource of type ExchService2007.

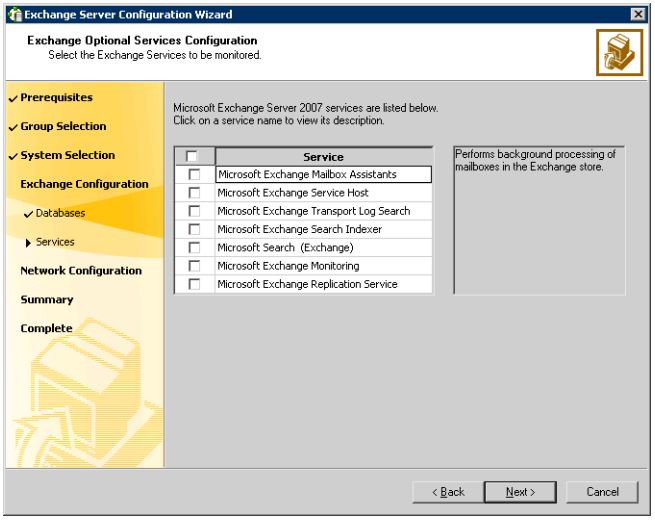


7 Specify information related to the network:

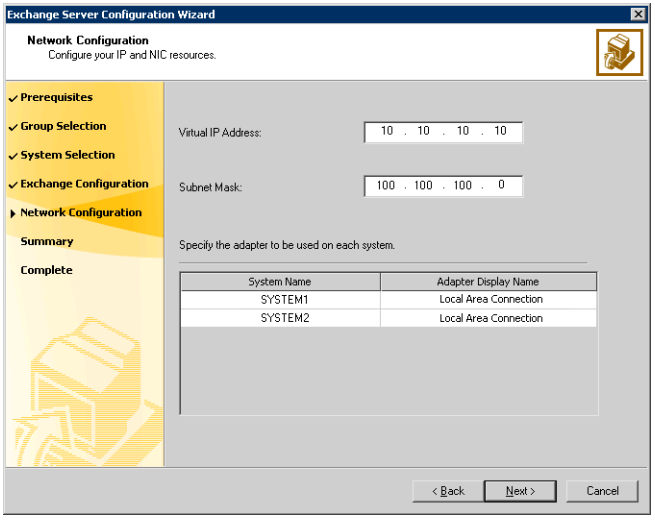

- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.

  If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

---

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

---

- Click **Next**.

8 Review the service group configuration and change the resource names, if desired:

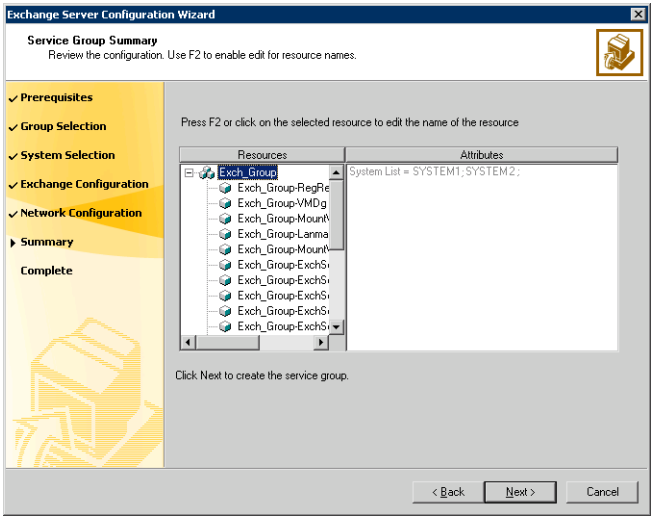


The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

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

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

- Click **Next**.
- A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

9 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system.

10 Click **Finish**.

After bringing the service group online, you must run the Exchange Management Console so that all the stores are automatically mounted on start-up.

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

# Verifying the cluster configuration

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

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

### To switch service groups

1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.

- Click **Switch To**, and click the appropriate node from the menu.

- In the dialog box, click **Yes.** The service group you selected is taken offline on the original node and brought online on the node you selected.
  If there is more than one service group, you must repeat this step until all the service groups are switched.

2 Verify that the service group is online on the node you selected to switch to in step 1.

3 To move all the resources back to the original node, repeat step 1 for each of the service groups.

**To shut down an active cluster node**

1 Gracefully shut down or restart the cluster node where the service group is online.

2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.

3 Verify that the service group has failed over successfully, and is online on the next node in the system list.

4 If you need to move all the service groups back to the original node:

- Restart the node you shut down in step 1.
- Click **Switch To**, and click the appropriate node from the menu.
- In the dialog box, click **Yes.**
  The service group you selected is taken offline and brought online on the node that you selected.

# Configuring the Cluster Management Console connection

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

CMC comprises of the following components:

- *Management Server*
  The management server accepts and processes the operational commands and the configuration inputs that users enter through CMC. The management server communicates with the VCS High Availability engine (HAD). Install the CMC Management Server only if you plan to centrally manage multiple clusters. You must install the management server on a

standalone system that is outside any cluster but available on the local network.

- *Cluster Connector*
  The cluster connector is an agent that enables the management server to communicate with clusters through intervening firewalls. You must install the cluster connector on each cluster that is separated from the management server by a firewall. If there are no firewalls between the management server and the clusters, you can configure the clusters to use direct connection instead.
  In each cluster, the cluster connector runs on one node at a time, but is installed on all nodes and is configured for failover.

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

## Prerequisites for installing the cluster connector

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

- The root hash of the management server. Use the `vssat showbrokerhash` command and copy the root hash of the management server. Note that you must run this command from the `C:\Program Files\Veritas\Security\Authentication\bin` directory on the management server.

## Installing the cluster connector on Windows clusters

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

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

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

(HTTPS), and 2994 (DBMS; this port can be shared with other Symantec products)

10 In the Services Account Password dialog box:

- Enter a password for the user account that the cluster connector uses for management server communications
- Enter the root hash of the authentication broker used by the authentication broker installed on the management server

The password is the password that was entered for the cluster connector service account during management server installation.

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

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

The output of this command looks similar to the following:

```
Root Hash:      9dfde3d9aaebee084f8e35819c1fed7e6b01d2ae
```

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

11 In the Summary dialog box, review the information you have specified and, if satisfactory, click **Next** to accept it and start the installation.

The Installing Veritas Cluster Management Console dialog box displays a progress bar and a status message window for the installation.

12 After the installation is complete, click **Next**.

13 In the Completed the Symantec Veritas Cluster Management Console Installation Manager dialog box, click **Finish**.

The installer creates log files at `C:\Documents and Settings\All Users\Application Data\Veritas\Cluster Management Console`. The file names are Install_GUI_0.log and Install_MSI_0.log. The installer creates Install_GUI_0.log on the system from which you run the cluster connector installation. The installer creates Install_MSI_0.log on the target systems.

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

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

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

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

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

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

1. On a cluster node, obtain a command prompt and change to the following directory:
   `Veritas\Security\Authentication\bin`
   This directory may be in one of the following paths:
   `C:\Program Files`
   or
   `C:\Program Files\Common Files`
2. Set up a trust relationship between the authentication broker on the management server and the authentication broker on the local cluster node. Type the following command:
   ```
   vssat setuptrust --broker MS_IPAddress:[2821 (optional)]--
   securitylevel high --hash Hash_From_MS
   ```
3. Authenticate the CMC_CC@CMC_SERVICES account on the local node. Type the following command:
   ```
   "vssat authenticate --domain vx:CMC_SERVICES --prplname CMC_CC
   --password  password_for_CMC_CC_user_created_during_MS_install
   --broker MS_IPAddress:2821
   ```
   Usage for this command is
   ```
   vssat authenticate  --domain <type:name> [--prplname <prplname>
   [--password <password>]] [--broker <host:port>]
   ```

Repeat these steps on each node in the cluster.

## Uninstalling the cluster connector

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

**To uninstall the cluster connector from Windows clusters**

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

6 Follow the prompts in the uninstallation wizard. When available, click **Finish** to close the wizard.

Chapter 4

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

This chapter covers the following topics:



You can either install and configure a new “any-to-any” SFW HA environment for Exchange Server 2007 to provide a production node with multiple failover nodes or, you can transform an existing active/passive SFW HA environment for Exchange Server 2007 into an any-to-any environment.

Table 4-1 outlines the high-level objectives to create a new any-to-any environment and the tasks to complete each objective:

**Table 4-1** Task list

| Objective | Tasks |
|---|---|
| "Reviewing the configuration" on page 100 | ■ Understanding a basic any-to-any Exchange configuration |
| "Reviewing the requirements" on page 103 | ■ Verifying hardware and software prerequisites |
| "Configuring the storage hardware and network" on page 106 | ■ Setting up the network and storage for a cluster environment<br>■ Verifying the DNS entries for the systems on which Exchange will be installed |
| "Installing Veritas Storage Foundation HA for Windows" on page 108 | ■ Verifying the driver signing options for Windows 2003 systems<br>■ Installing SFW HA<br>■ Restoring driver signing options for Windows 2003 systems |
| "Installing the SFW HA patch for Exchange Server 2007" on page 113 | ■ Install the SFW HA patch; involves installing the VCS agent for Exchange Server 2007<br>■ If required, install the SFW patch to perform VSS-based backup and restore operations with Exchange Server 2007 |
| "Configuring the cluster" on page 114 | ■ Verifying static IP addresses and name resolution configured for each node<br>■ Configuring cluster components using the Veritas Cluster Server Configuration Wizard |

**Table 4-1** Task list (continued)

| Objective | Tasks |
|---|---|
| "Managing disk groups and volumes" on page 136 | ■ Using the VEA console to create disk groups<br>■ Using the VEA console to create the Data, Log, RegRep, and SHARED volumes<br>■ Managing disk groups and volumes, with instructions for mounting and unmounting volumes |
| "Installing Exchange on the first node" on page 138 | ■ Reviewing the prerequisite checklist<br>■ Running the Exchange Setup Wizard for Veritas Cluster Server and Microsoft Exchange Server installation<br>■ Performing this "First Node" installation on each of the active Exchange nodes in the final configuration.<br>■ 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 142 | ■ Moving databases on the first node from the local drive to the shared drive using the Exchange Setup Wizard for Veritas Cluster Server<br>■ Repeating this task for each of the active Exchange nodes in the final configuration, making sure that each of the active Exchange servers has a separate area for its databases. Do not share databases between separate Exchange servers. |

**Table 4-1** Task list (continued)

| Objective | Tasks |
|---|---|
| "Installing Exchange on additional nodes" on page 145 | ■ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes<br>■ Perform this task for all of the failover systems. |
| "Configuring the Exchange service group for VCS" on page 150 | ■ Preparing the cluster for any-to-any failover using the Exchange Setup Wizard for Veritas Cluster Server. This step must be completed on each of the Exchange Virtual Servers.<br>■ 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 157 | 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 4-2:

**Table 4-2** 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 4-1 shows an example of a three-node cluster in an any-to-any configuration.

**Figure 4-1** Three-node cluster in an any-to-any configuration

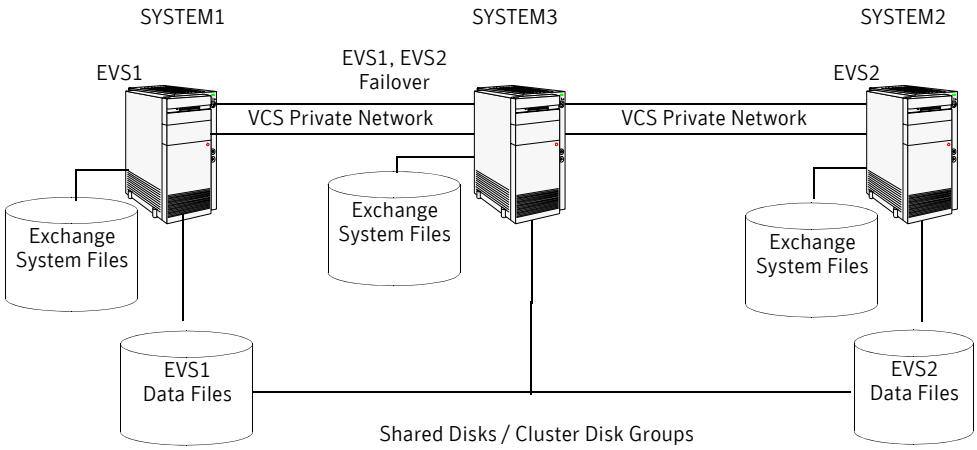


For example, consider a three-node cluster hosting two Exchange Virtual Servers, EVS1 and EVS2. The virtual servers are configured in VCS in two service groups such that SYSTEM1 has first priority for the EVS1 service group and SYSTEM2 has first priority for the EVS2 service group, while SYSTEM3 is shared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

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

## 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 170.

## Sample configuration

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

**Table 4-3** Sample configuration

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

# Reviewing the requirements

Two or more Exchange virtual servers can exist in an any-to-any configuration. Refer to “Reviewing the configuration” on page 100.

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

## Disk space requirements

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

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

**Table 4-4** Disk space requirements

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

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

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

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

### Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA) with the Veritas Cluster Server Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:

- Microsoft Exchange Server 2007 Standard Edition or Enterprise Edition
  *with*
  Windows Server 2003 x64 Standard Edition, Enterprise Edition, Datacenter Edition (SP1 required for all editions, SP2 supported)
  *or*
  Windows Server 2003 R2 x64 Standard Edition, Enterprise Edition, Datacenter Edition

## System requirements

Systems must meet the following requirements:

- Processor: x64 architecture-based computer with Intel processor that supports Intel Extended Memory 64 Technology (Intel EM64T) or AMD processor that supports the AMD64 platform; Intel Itanium family IA64 processors are not supported.
- Memory: minimum 2 GB of RAM per server.
- File format: Disk partitions must be formatted for the NTFS file system.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs.
  See "Best practices" on page 106.
- All servers must run the same operating system, service pack level, and system architecture.

## Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 domain.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.
- Static IP addresses for the following purposes:
  - One static IP address available per site for each Exchange Virtual Server (EVS)
  - One IP address for each physical node in the cluster

  - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option (VVR only). The same IP address may be used for all options.
  - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
  Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
  See the *Veritas Cluster Server Bundled Agents Reference Guide*.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

## Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Servers group.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

### Additional requirements

Please review the following additional requirements:

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

### Best practices

Symantec recommends that you perform the following tasks:

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

## Configuring the storage hardware and network

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

#### To configure the hardware

1. Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
2. Connect the network adapters on each system.
   - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.

- Symantec recommends removing TCP/IP from private NICs to lower system overhead.

3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.

4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk.
Use Windows Disk Management (**Start > Control Panel > Administrative Tools > Computer Management)** or Veritas Enterprise Administrator (VEA) on each system to verify that the attached shared disks are visible.

### To verify the DNS settings and binding order for all systems

1 Open the Control Panel (**Start>Control Panel**).

2 Right-click on Network Connections and click **Open**.

3 Ensure the public network adapter is the first bound adapter:

- From the Advanced menu, click **Advanced Settings**.
- In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
- Click **OK**.

4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

5 From the status window, click **Properties**.

6 In the General tab:

- Select the **Internet Protocol (TCP/IP)** check box.
- Click **Properties**.

7 Select the **Use the following DNS server addresses** option.

8 Verify the correct value for the IP address of the DNS server.

9 Click **Advanced**.

10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.

11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

12 Click **OK**.

# Installing Veritas Storage Foundation HA for Windows

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

## Setting Windows driver signing options

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

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

**Table 4-5** Installation behavior with unsigned drivers

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

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

**To change the driver signing options on each system**

1 Log on locally to the system.

2 Open the Control Panel and click **System**.

3 Click the **Hardware** tab and click **Driver Signing**.

4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.

5 Click **OK**.

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

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

### To install the product

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

2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.

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


4 Do one of the following:

- Click **Complete/Custom** to begin installation.
- Click the **Administrative Console** link to install only the client components.

5 Review the Welcome message and click **Next**.

6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.

7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.

8 Repeat for additional license keys. Click **Next**
- To remove a license key, click the key to select it and click **Remove**.
- To see the license key's details, click the key.

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

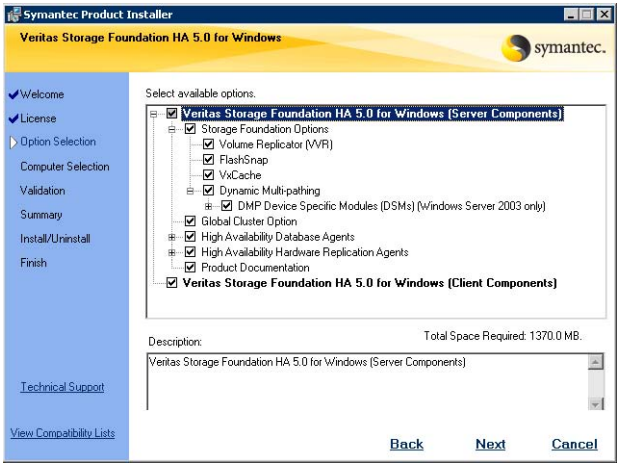


The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

| | |
|---|---|
| Veritas Volume Replicator | If you plan to use VVR for replication, you must also select the option to install VVR. |
| Global Cluster Option | Required for a disaster recovery configuration only. |
| Client Components | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console. |

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

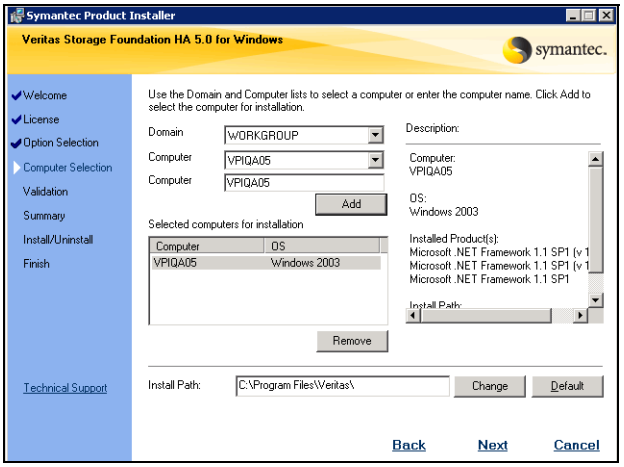


| | |
|---|---|
| Domain | Select a domain from the list.<br>Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate. |
| Computer | To add a computer for installation, select it from the Computer list or type the computer's name in the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

11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.

12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.

13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
- For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
- For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.

14 Click **OK**.

15 Review the information and click **Install**. Click **Back** to make changes, if necessary.

16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.

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

18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.

19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.

20 Review the log files and click **Finish**.

21 Click **Yes** to reboot the local node.

## Resetting the driver signing options

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

**To reset the driver signing options**

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

# Installing the SFW HA patch for Exchange Server 2007

The SFW HA patch for Exchange Server 2007 contains the new VCS agent for Exchange Server 2007. The patch zip file contains the following files:

- **vrtsvcsexch.msi** - the VCS agent msi file
- **InstallVCSExch2007.bat** - batch file to install the VCS agent
- **UninstallVCSExch2007.bat** - batch file to remove the VCS agent

Extract these files to a temporary location on the system. Ensure that the agent .msi file and the .bat file are at the same level in a directory.

**To install the SFW HA patch**

1. If the system is part of a cluster, complete this step. If not, proceed to step 2.
   - Make sure that all the service groups are offline in the cluster.
   - Save and close the cluster configuration. Type the following on the command prompt:
     ```
     C:\> haconf -dump -makero
     ```
   - Stop the Veritas High Availability engine (HAD) on all the cluster nodes. Type the following on the command prompt:
     ```
     C:\> hastop -all
     ```
2. On the system, double-click **InstallVCSExch2007.bat**. The .msi will install the VCS agent for Exchange Server 2007 on the system.
3. Repeat step 2 on all the systems where you want to install the patch.

4 If the system is part of a cluster, complete this step.
Start HAD on the system on which you installed the patch. Type the following on the command prompt:

```
C:\> hastart
```

# Installing the SFW patch for Exchange Server 2007

The SFW patch for Exchange Server 2007 enables SFW support for performing VSS-based backup and restore operations with Exchange Server 2007.

This step is optional. If required, you can install the SFW patch now. Refer to the readme file accompanying the SFW patch for the installation steps.

# Configuring the cluster

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

Complete the following tasks before creating a cluster:

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

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

### To configure a VCS cluster

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

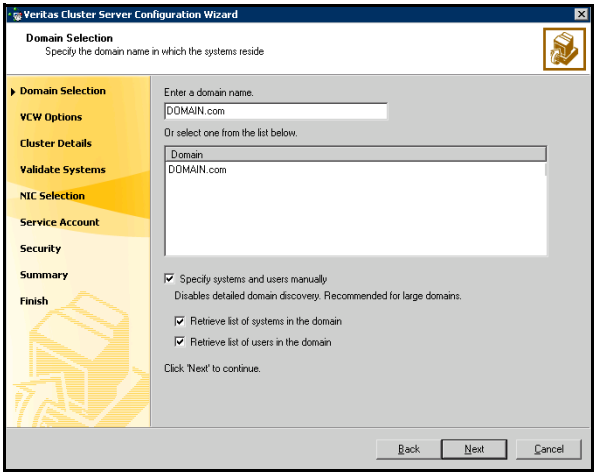


To discover information about all systems and users in the domain:

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

  Proceed to step 7 on page 117.

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 116. Otherwise, proceed to the next step.

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

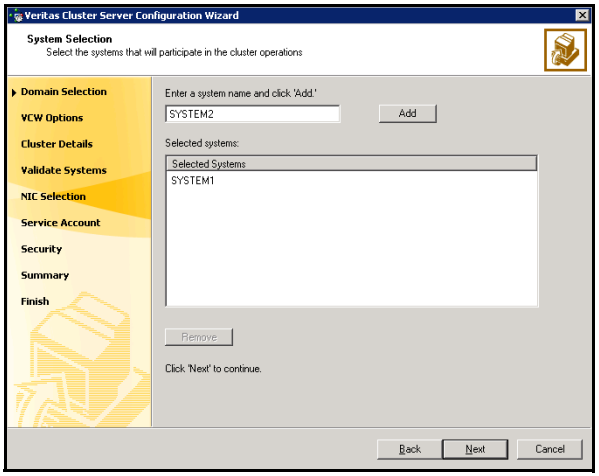


Proceed to step 7 on page 117.

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

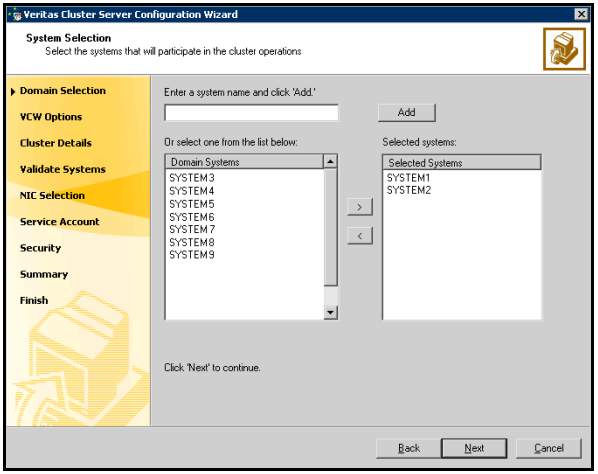


Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

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

8 On the Cluster Details panel, specify the details for the cluster and then click **Next**.

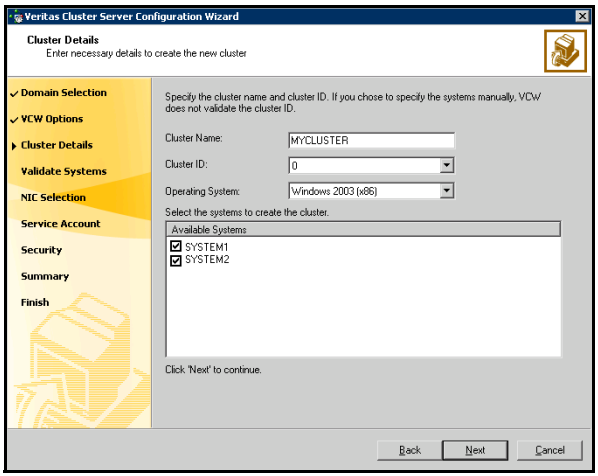


| | |
|---|---|
| Cluster Name | Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. |

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

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

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

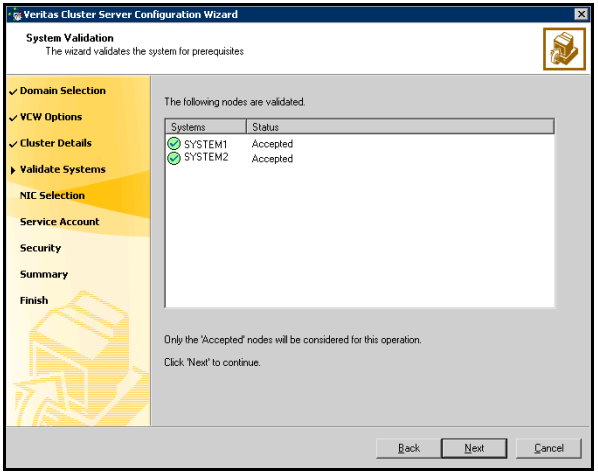


If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in step 8 on page 117, proceed to the next step. Otherwise, proceed to step 11 on page 119.

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

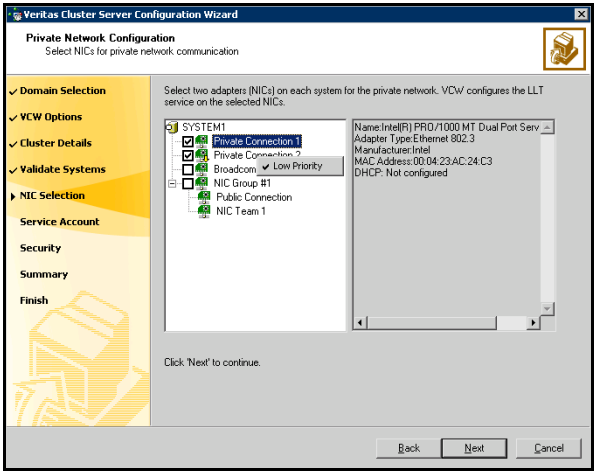


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs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

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

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

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

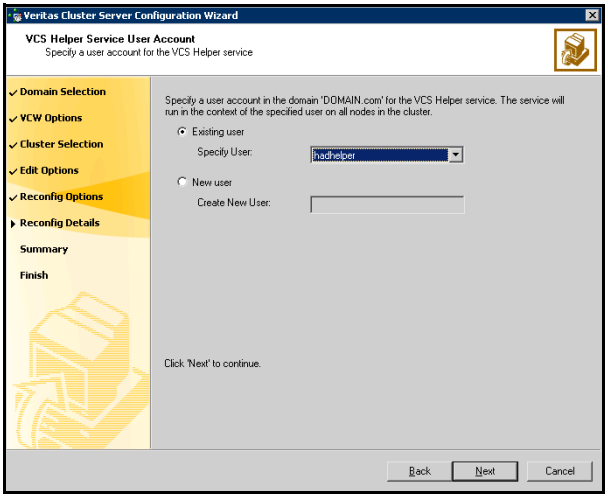


- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list,
  - If you chose not to retrieve the list of users in step 4 on page 115, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

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

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

Do one of the following:

- To use the single sign-on feature

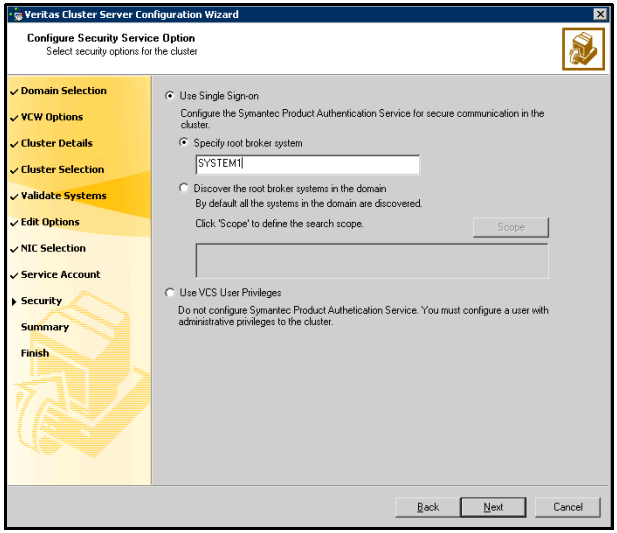


  - Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

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

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

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

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

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

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

- To use VCS user privilege:

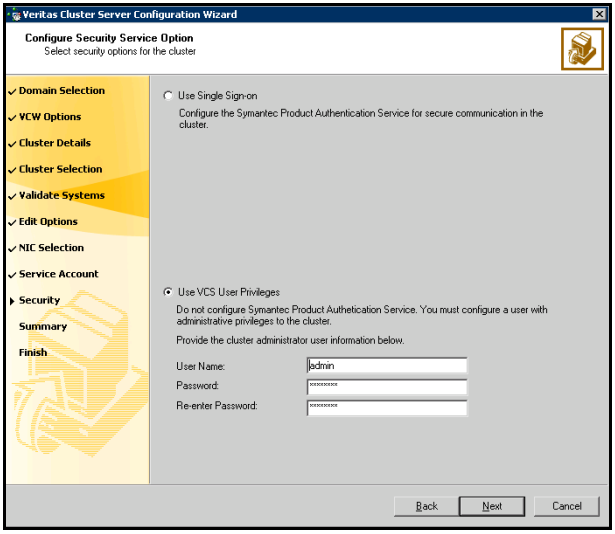


- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

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

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

- Click **Next**.

13 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network.

If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.

The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

14 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

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

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

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

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

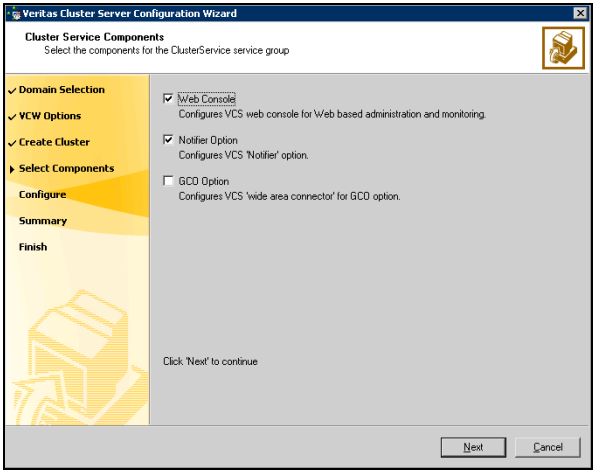


- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
  See "Configuring Web console" on page 125.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
  See "Configuring notification" on page 126.

# 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**.

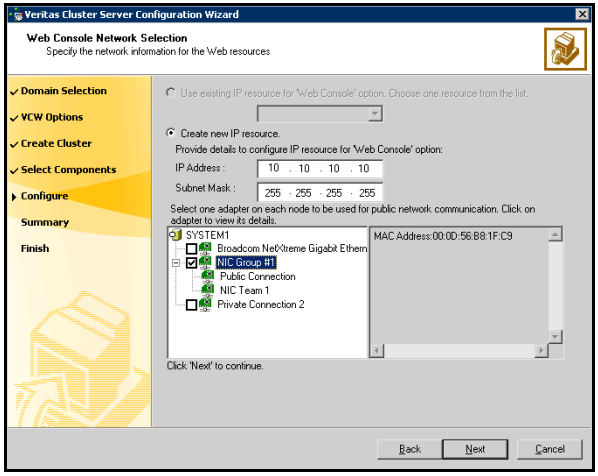


- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
- If you choose to configure a new IP address, type the IP address and associated subnet mask.
- Select a network adapter for each 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 126.
Otherwise, click **Finish** to exit the wizard.

# Configuring notification

This section describes steps to configure notification.

### To configure notification

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

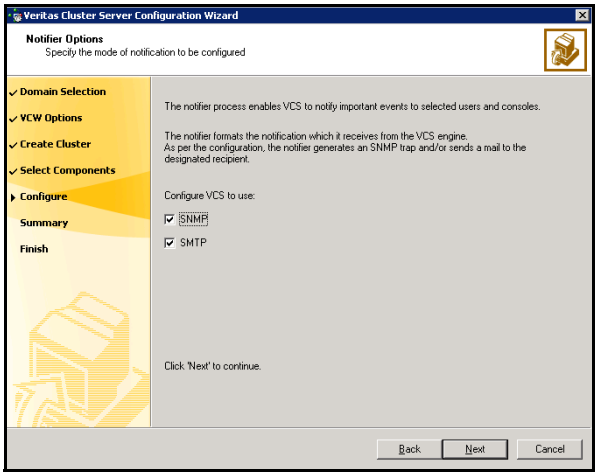


You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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

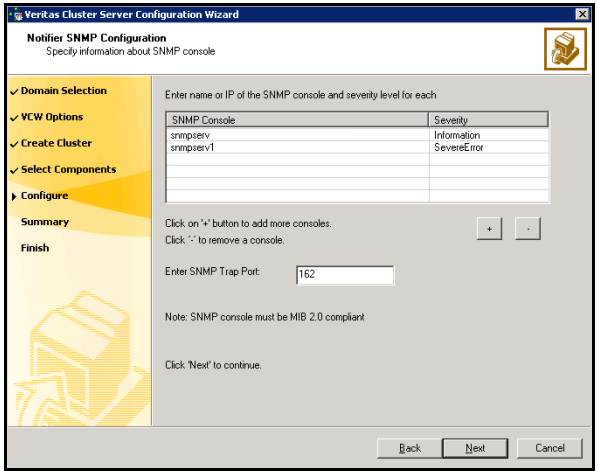


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '**+**' to add a field; click '**-**' to remove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

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

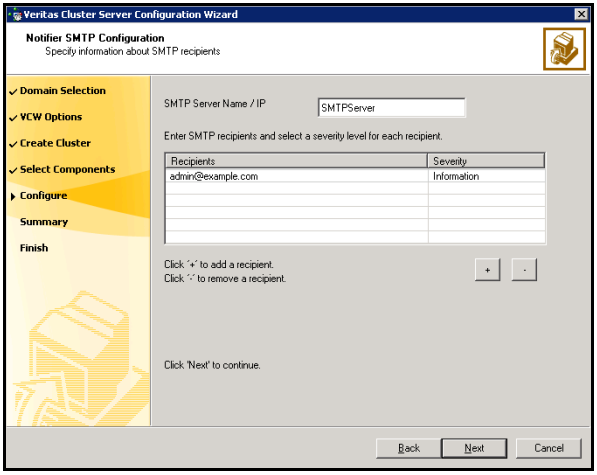


- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity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

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

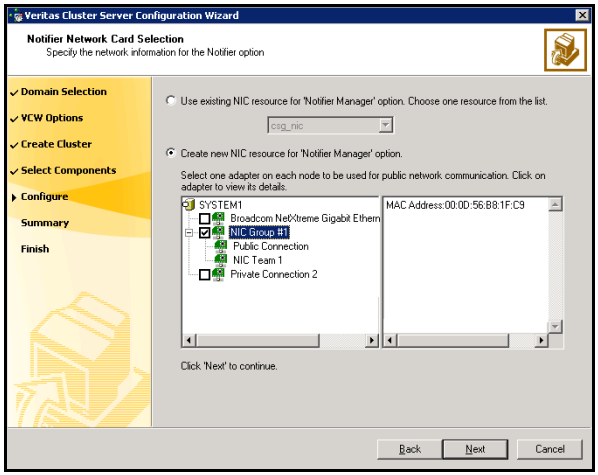


- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
- If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.

5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.

6 Click **Configure**.

7 Click **Finish** to exit the wizard.

# Configuring the first Exchange Virtual Server

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

See “Reviewing the configuration” on page 100.

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

# Configuring disk groups and volumes

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

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

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

Typically, a SFW disk group corresponds to an Exchange storage group. Figure 4-2 is a detailed view of the disk groups and volumes in an HA environment.

**Figure 4-2** Disk groups and volumes for Exchange virtual server EVS1 in HA setup

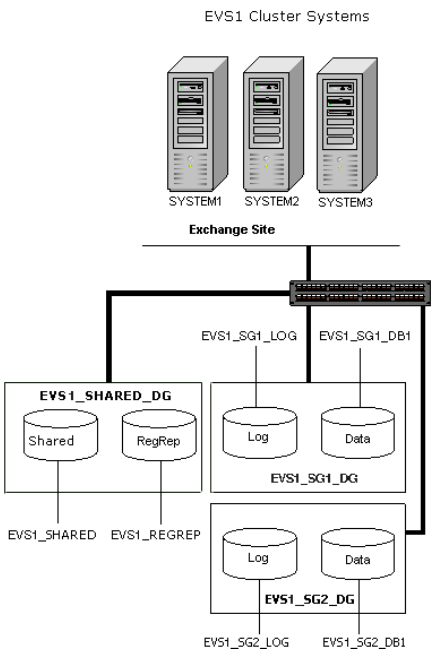

Use the following procedures to create the appropriate disk groups and volumes. The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange disk group EVS1_SG1_DG contains two volumes:

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

Exchange disk group EVS1_SHARED_DG contains the following volume:

- EVS1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.

---

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

---

## Creating a disk group

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

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

6 Provide information about the cluster disk group:

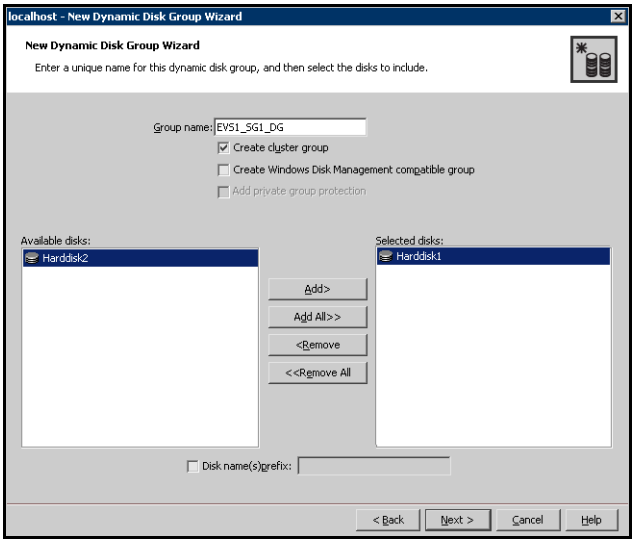


- Enter the name of the disk group (for example, EVS1_SG1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
  Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
- Click **Next**.

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the new disk group.

## Creating volumes

This procedure assumes you are starting with the EVS1_SG1_DB1 volume.

### To create dynamic volumes

1. If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
2. Click **Connect to a Host or Domain**.
3. In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
   To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
4. To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
   You can right-click the disk group you have just created, for example EVS1_SG1_DG.
5. At the New Volume wizard opening screen, click **Next**.
6. Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

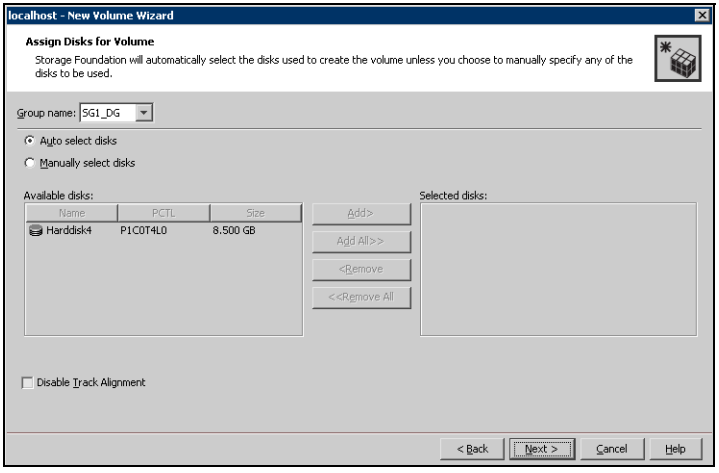


7. Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

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

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

8 Click **Next**.

9 Specify the volume attributes.

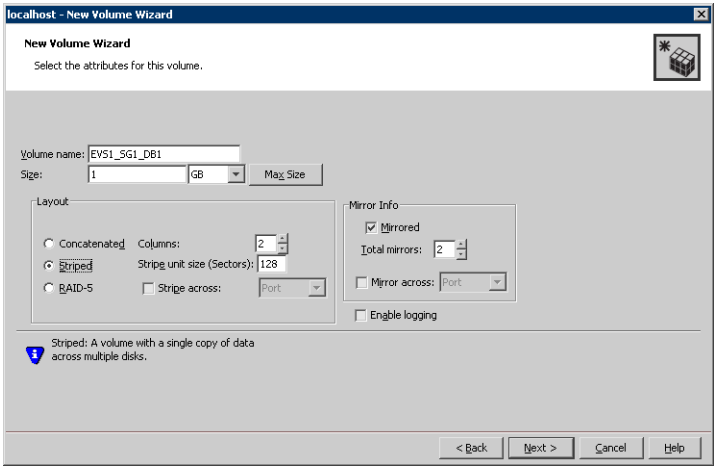


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
- If you are creating a striped volume, the **Columns**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

10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

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

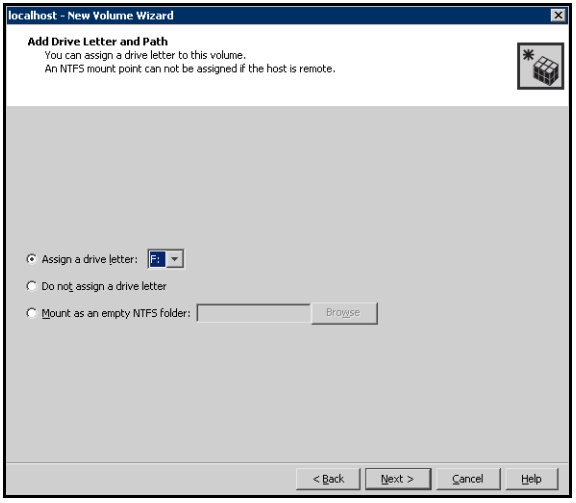


11 Click **Next**.

12 Create an NTFS file system.

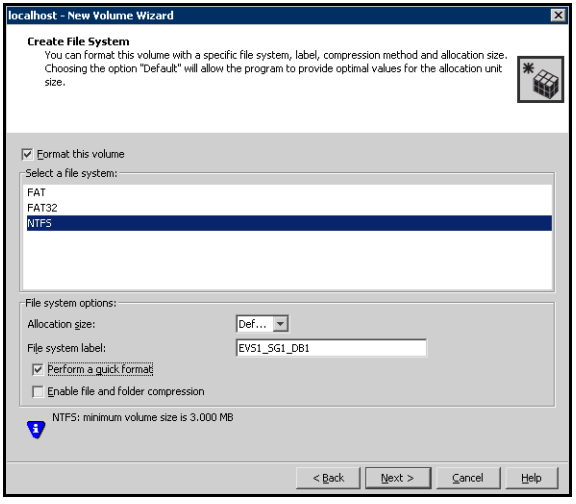


- Make sure the **Format this volume** checkbox is checked and click **NTFS**.

- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

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

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

15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3).

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

## Managing disk groups and volumes

This section includes the following procedures:

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

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

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

## Importing a disk group and mounting a volume

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

**To import a disk group**

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

**To mount a volume**

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

## Unmounting a volume and deporting a disk group

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

**To unmount a volume and deport the dynamic disk group**

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

## Installing Exchange on the first node

Installing Exchange on the first node of EVS1 is described in three stages that involve 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 136.
- Mount the volume containing the information for registry replication (EVS1_REGREP).
  See "Managing disk groups and volumes" on page 136.
- Verify that all systems on which Exchange Server will be installed have IIS installed; you must install WWW services on all systems.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
  - You must be a domain user.
  - You must be an Exchange Full Administrator.
  - You must be a member of the Exchange Servers group.
  - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
  - You must have write permissions on the DNS server to perform DNS updates.
  - Make sure the HAD Helper domain user account has the "Add workstations to domain" privilege enabled in the Active Directory. To verify this, click **Start > Administrative Tools > Local Security Policy** on the domain controller to launch the security policy display. Click **Local Policies > User Rights Management** and make sure the user account has this privilege.
  - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.
  - The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.

## Exchange pre-installation: First node

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

### To perform Exchange pre-installation

1. Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
2. Click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 2007 Mailbox Server role 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. Specify information related to the network.

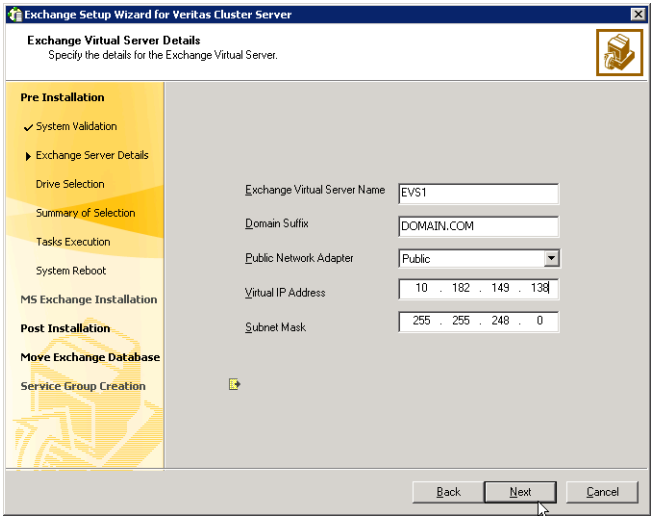


- Enter a unique virtual name for the Exchange server.
  Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you

must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.

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

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

8 Review the summary of your selections and click **Next**.

9 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.

10 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**.

11 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.

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

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

## Exchange installation: First node

Install Exchange on the node on which you performed the pre-installation. HA support for Exchange Server 2007 is available for the Mailbox Server role. While installing Exchange, ensure that you install the Mailbox Server role only.

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.

## Exchange post-installation: First node

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

### To perform Exchange post-installation

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

8 Once the node is rebooted, move the databases created during the Exchange installation from the local drive to the shared drive.

## 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 to import the disk groups and mount the volumes for the Exchange database, and transaction logs. Refer to "Managing disk groups and volumes" on page 136 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.
- The Exchange Setup Wizard for VCS cannot move the Exchange storage groups until local continuous replication (LCR) is suspended for those storage groups. Please suspend LCR using the Exchange Management Console or the Exchange Management Shell, before moving the Exchange databases.
  Refer to the Microsoft Exchange documentation for information on how to suspend LCR.

### To move Exchange databases

1 Click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 Setup Wizard** to start the Exchange Setup Wizard for VCS.

2 Review the information in the Welcome dialog box and click **Next**.

3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.

4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.

5 In the Select Exchange Virtual Server dialog box:

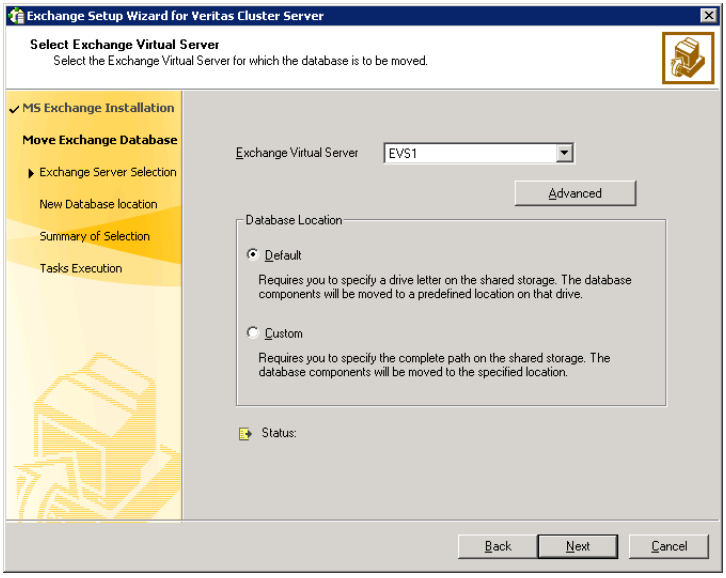


- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

---

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

---

- Click **OK**.
- Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
- Click **Next**.

6 Do one of the following:

- If you would like to move the default mailbox store, and the public store to the generated default paths on the volumes for these components, proceed to the next step.
- Otherwise, proceed to step 8 on page 145 to specify the path location on the volumes that you will designate for these components.

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

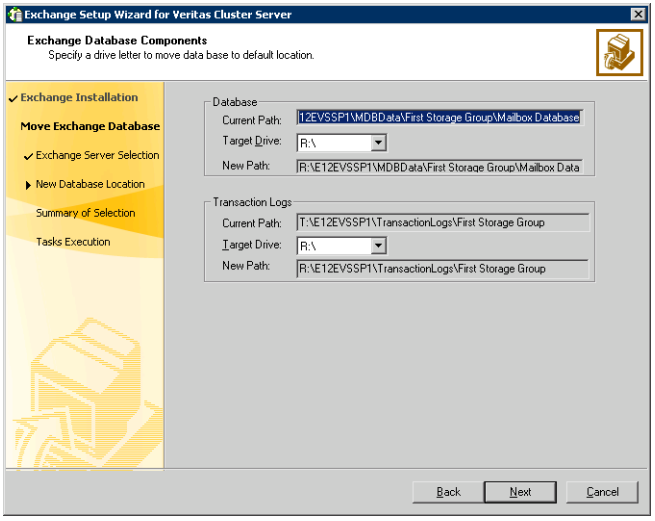


- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Click **Next** and proceed to step 9 on page 145.

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

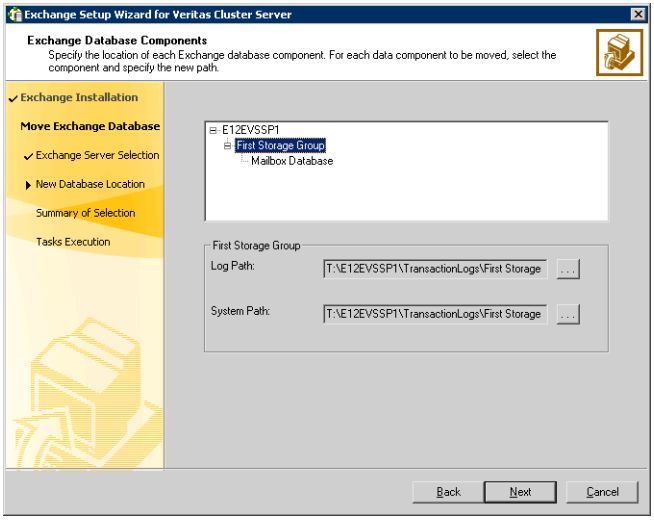


For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the paths for the Exchange database components are not the root of a drive. You must select a directory on the specified drive.

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

9 Review the summary of your selections and click **Next**.

10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.

11 Click **Finish** to exit the wizard.

## Installing Exchange on additional nodes

After moving the Exchange databases to shared storage, install Exchange on all failover nodes in the cluster for the same Exchange virtual server (EVS1). You must run 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 138.

## Exchange pre-installation: additional nodes

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

**Note:** Use the VEA console to unmount the Exchange volumes and deport the cluster disk group from the first node before beginning the Exchange Pre-Installation on additional nodes. See "Managing disk groups and volumes" on page 136 for instructions.

### To perform Exchange pre-installation

1. Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
2. From the node to be added to an Exchange cluster, click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 2007 Mailbox Server role 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 Specify network information for the Exchange virtual server.

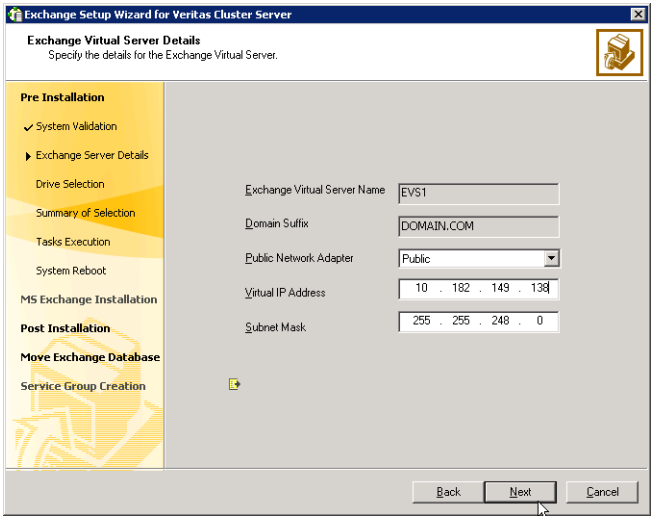


The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.

- 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**.

8 Review the summary of your selections and click **Next**.

9 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.

10 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**.

11 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.

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: additional nodes

Install Exchange on the additional node on which you performed the pre-installation. HA support for Exchange Server 2007 is available for the Mailbox Server role. While installing Exchange, ensure that you install the Mailbox Server role only.

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

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 RecoverServer as the install mode:

   ***<drive letter>***`:\setup.com /mode:recoverserver`

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

2. Setup copies the setup files locally to the computer on which you are installing Exchange 2007 and then checks the prerequisites, including all prerequisites specific to the server roles that you are installing. If you have not met all of the prerequisites, Setup fails and returns an error message that explains the reason for the failure. If you have met all of the prerequisites, Setup installs Exchange 2007.

3. Verify that the installation completed successfully. Refer to the Microsoft documentation for more information.

## Exchange post-installation: additional nodes

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

### To perform Exchange post-installation

1. Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
2. If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 4.
   If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
3. Review the information in the Welcome dialog box and click **Next**.
4. A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
5. The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
6. 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.

7 Click **Finish**.

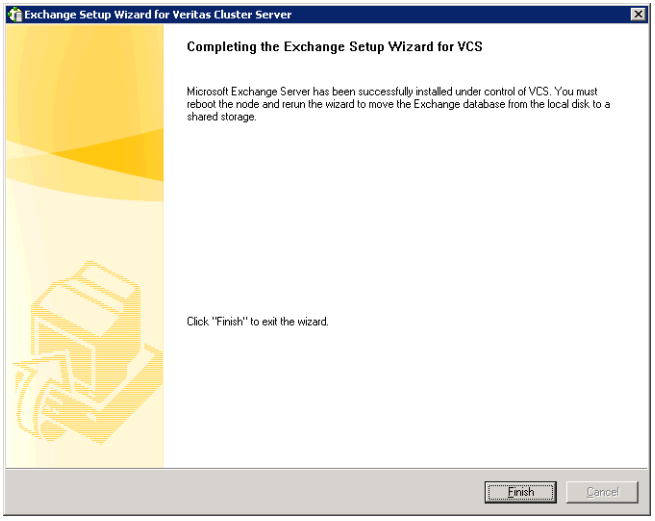


8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

9 Changes made during the post-installation steps do not take effect till you reboot the node.

## 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. Refer to the Exchange documentation for instructions.

### 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

  See "Managing disk groups and volumes" on page 136 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 Appendix A, "VCS agent for Exchange Server 2007" for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

Refer to the *Veritas Cluster Server Administrator's Guide* for more 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 Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 Specify the service group name and system list:

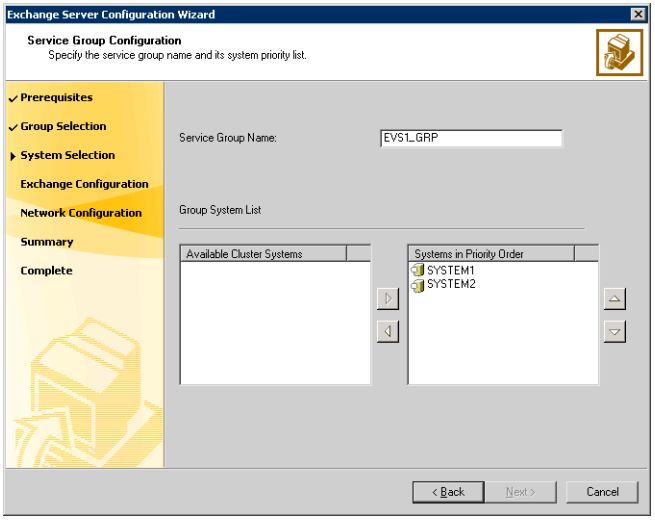


- Enter a name for the Exchange service group.
  If you are configuring the service group on the secondary site, ensure that the name matches the service group on the primary site.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow 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.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

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

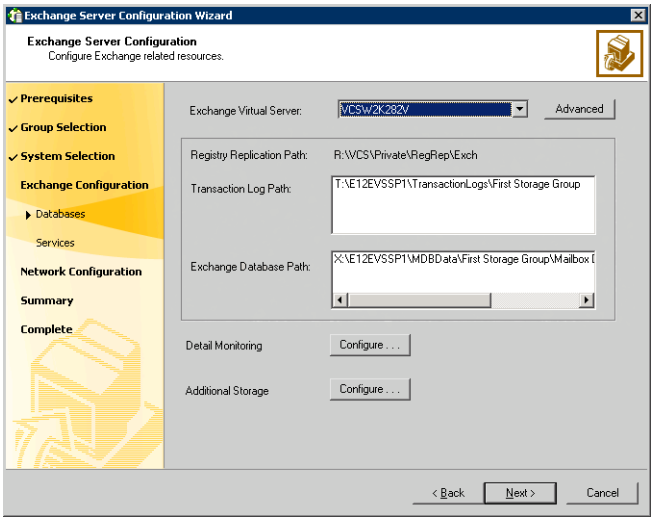


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
  - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.

- To configure Detail Monitoring for Exchange databases, click **Configure...**.

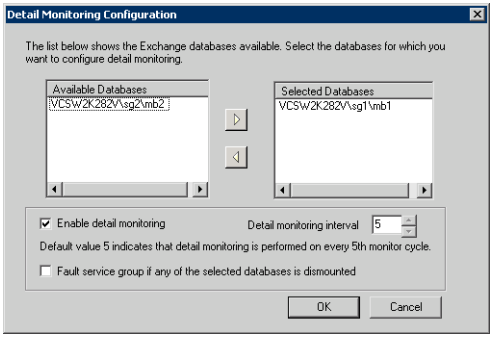


  On the Detail Monitoring Configuration dialog box, complete the following:

  - In the Available Databases box, select the databases for detail monitoring and double-click, or click the right-arrow button to move them to the Selected Databases box. To remove a database, select the database in the Selected Databases box, and double-click or click the left-arrow button.
  - Check **Enable detail monitoring** check box, and specify the monitoring interval in the **Detail monitoring interval** field.
  - If you want the VCS agent to fault the service group if a database selected for detail monitoring is dismounted, check the **Fault service group if any of the selected database is dismounted** check box.

  See the VCS agent attribute descriptions in the Appendix, for more information on detail monitoring and VCS agent behavior.

  - Click **OK**.

- To configure additional storage, click **Configure...**. On the Additional Storage Configuration dialog box, complete the following:

  - 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 (MSExchangeSA) service resource.

- Click **Next**.

6 Select the optional Exchange services to be monitored and click **Next**. Each optional service that is selected will be configured as a VCS resource of type ExchService2007.

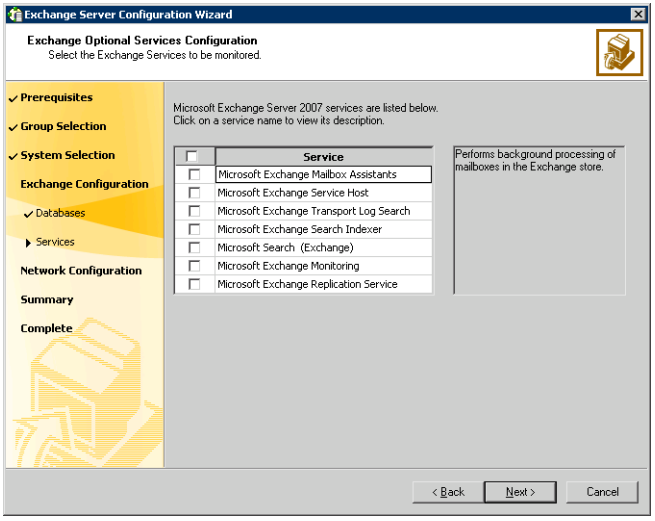


7 Specify information related to the network:

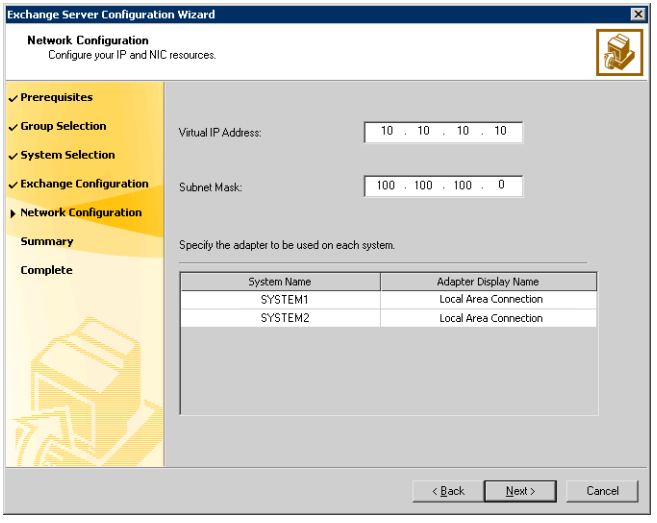

- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
  If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual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

---

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

---

- Click **Next**.

8 Review the service group configuration and change the resource names, if desired:

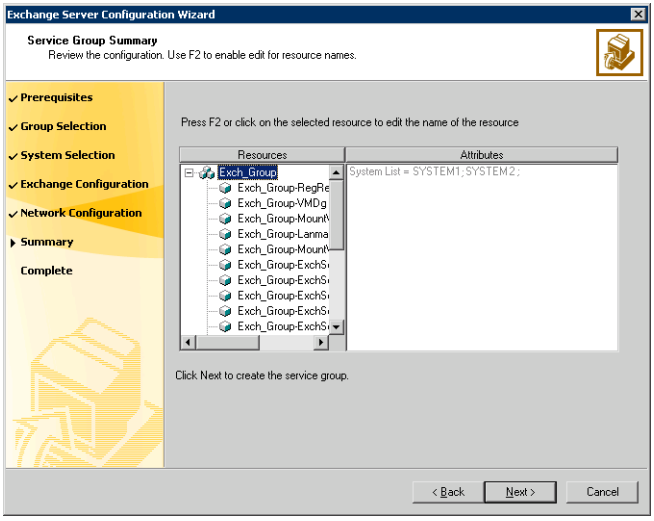


The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

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

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

- Click **Next**.
- A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

9 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system.

10 Click **Finish**.

After bringing the service group online, you must run the Exchange Management Console so that all the stores are automatically mounted on start-up.

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

## Verifying the cluster configuration

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

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

### To switch service groups

1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.

- Click **Switch To**, and click the appropriate node from the menu.

- In the dialog box, click **Yes.** The service group you selected is taken offline on the original node and brought online on the node you selected.
  If there is more than one service group, you must repeat this step until all the service groups are switched.

2 Verify that the service group is online on the node you selected to switch to in step 1.

3 To move all the resources back to the original node, repeat step 1 for each of the service groups.

**To shut down an active cluster node**

1 Gracefully shut down or restart the cluster node where the service group is online.

2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.

3 Verify that the service group has failed over successfully, and is online on the next node in the system list.

4 If you need to move all the service groups back to the original node:

- Restart the node you shut down in step 1.
- Click **Switch To**, and click the appropriate node from the menu.
- In the dialog box, click **Yes.**
  The service group you selected is taken offline and brought online on the node that you selected.

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

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

See “Reviewing the configuration” on page 100.

## 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 4-3 shows a detailed view of the disk groups and volumes in an HA environment.

**Figure 4-3** Disk groups and volumes for Exchange virtual server EVS2 in HA setup

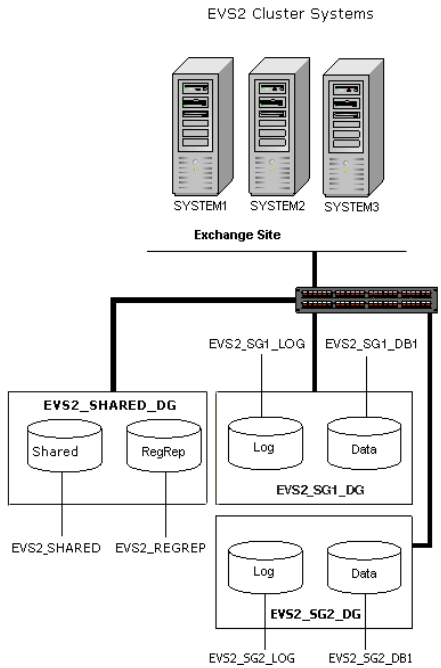


Use the following procedures to create the appropriate disk groups and volumes.The general guidelines for disk group and volume setup for EVS2_SG2_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange disk group EVS2_SG2_DG create contains two volumes:

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

Exchange storage group EVS2_SHARED_DG create contains the following volume:

- EVS2_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.

For instructions on creating disk groups:

see "Creating a disk group" on page 131

For instructions on creating volumes:

see "Creating volumes" on page 133.

## Managing disk groups and volumes

This section includes the following procedures:

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

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

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

## Importing a disk group and mounting a volume

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

### To import a disk group

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

### To mount a volume

1. If the disk group is not imported, import it.
2. To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
3. Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.

4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.

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

5 Click **OK**.

## Unmounting a volume and deporting a disk group

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

**To unmount a volume and deport the dynamic disk group**

1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.

2 In the Drive Letter and Paths dialog box, click **Remove.** Click **OK** to continue.

3 Click **Yes** to confirm.

4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.

5 Click **Yes**.

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

Installing Exchange on the first node of EVS2 is described in three stages that involve 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 136.
- Mount the volume containing the information for registry replication (EVS2_REGREP).
  See "Managing disk groups and volumes" on page 136.
- Verify that all systems on which Exchange Server will be installed have IIS installed; you must install WWW services on all systems. .
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation

to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.

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

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

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

### To perform Exchange pre-installation

1. Verify the volume created to store the registry replication information is mounted on this node and unmounted from other nodes in the cluster.
2. Click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 2007 Mailbox Server role 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 Specify information related to the network.

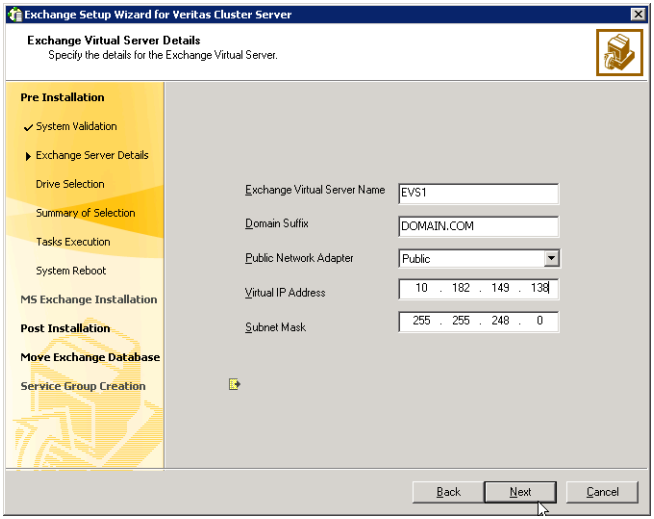


- Enter a unique virtual name for the Exchange server.
  Once you have assigned a virtual name to the Exchange server, you cannot change the virtual name later. To change the virtual name, you must uninstall Exchange from the VCS environment and again install it using the Exchange Setup Wizard for VCS.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

and the wizard will proceed with offlining the Exchange virtual server at the primary site. When all requirements are validated and met. Click **Next.**

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

8 Review the summary of your selections and click **Next**.

9 A warning message appears indicating, that the system will be renamed and rebooted when you exit the wizard. Click **Yes** to continue.

10 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**.

11 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.

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

## Exchange installation: first node of an additional Exchange Virtual Server

Install Exchange on the node on which you performed the pre-installation. HA support for Exchange Server 2007 is available for the Mailbox Server role. While installing Exchange, ensure that you install the Mailbox Server role only.

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.

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

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

#### To perform Exchange post-installation

1 Make sure the registry replication volume is online and mounted on the node on which you plan to perform the post-installation tasks.

2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 4.
If you reboot the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

3 Review the information in the Welcome dialog box and click **Next**.

4 A message appears indicating that the system will be renamed and restarted after you quit the wizard.This sets the node back to its physical host name. Click **Yes** to continue.

5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After all the commands are executed, click **Continue**.

6 Click **Finish**.

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

8 Once the node is rebooted, move the databases created during the Exchange installation from the local drive to the shared drive.

## 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 to import the disk group and mount the volumes for the Exchange database, and transaction logs.
  See “Managing disk groups and volumes” on page 136.

- Start VEA and go to SYSTEM1. Select the storageagent and import the disk groups. Make sure the volumes have been assigned a drive letter.
- The Exchange Setup Wizard for VCS cannot move the Exchange storage groups until local continuous replication (LCR) is suspended for those storage groups. Please suspend LCR using the Exchange Management Console or the Exchange Management Shell, before moving the Exchange databases.
  Refer to the Microsoft Exchange documentation for information on how to suspend LCR.

### To move Exchange databases

1. Click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 Setup Wizard** to start the Exchange Setup Wizard for VCS.
2. Review the information in the Welcome dialog box and click **Next**.
3. In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
4. In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.
5. In the Select Exchange Virtual Server dialog box:

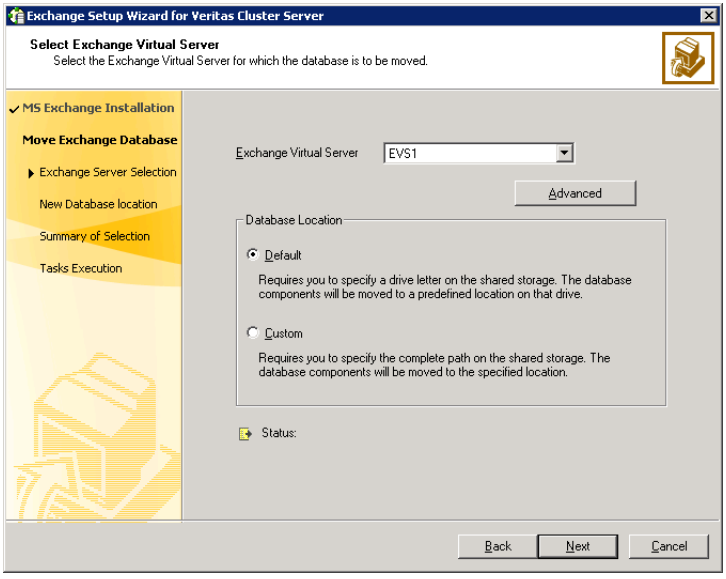


   - Select the Exchange virtual server.

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

---

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

---

- Click **OK**.
- Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
- Click **Next**.

6 Do one of the following:

- If you would like to move the default mailbox store, and the public store to the generated default paths on the volumes for these components, proceed to the next step.
- Otherwise, proceed to step 8 on page 145 to specify the path location on the volumes that you will designate for these components.

7 For the option of a default database location, specify the drives where the Exchange database components will be moved. The database components

will be moved to a default location on that drive. In the Exchange Database Components dialog box:

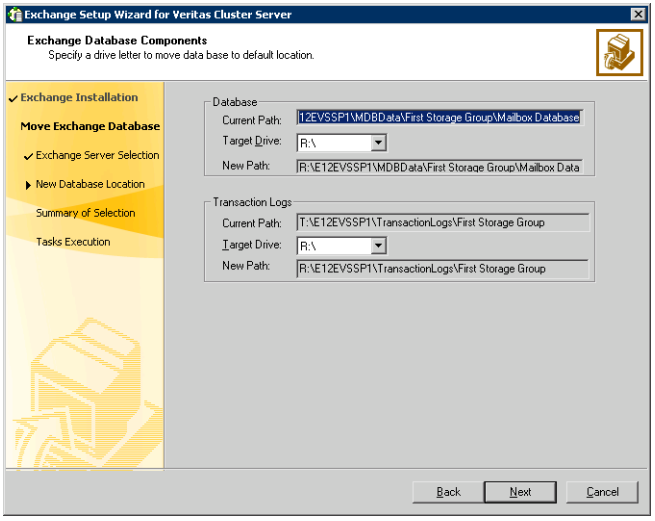


- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Click **Next** and proceed to step 9 on page 145.

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

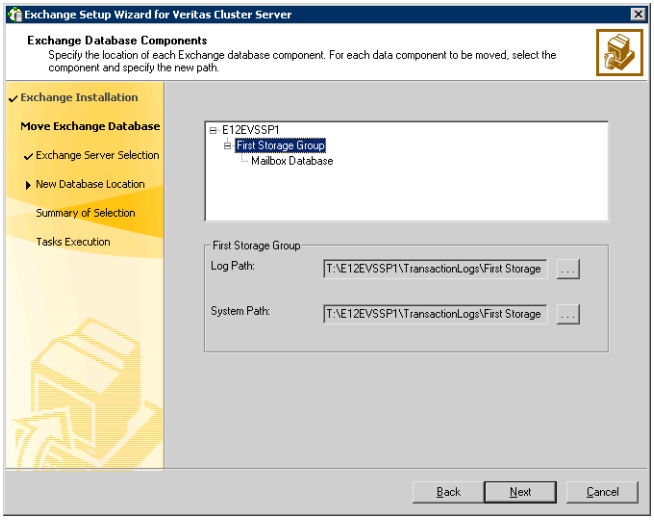


For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the paths for the Exchange database components are not the root of a drive. You must select a directory on the specified drive.

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

9 Review the summary of your selections and click **Next**.

10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.

11 Click **Finish** to exit the wizard.

# Specifying a common node for failover

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

The failover node for the first Exchange virtual server, EVS1, was specified when the EVS1 service group was created. After the designated Exchange virtual servers have been installed in the cluster, run the Exchange Setup Wizard for VCS from any system in the cluster. Ensure that you select the any-to-any option in the wizard.

Repeat the following task for each additional Exchange virtual server.

---

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

---

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

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

1. Start the Exchange Setup Wizard for VCS from any node configured to host an Exchange service group. Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server 2007 > Exchange Server 2007 Setup Wizard**.
2. Review the information in the Welcome dialog box and click **Next**.
3. In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.
4. In the Select Options dialog box, choose the **Configure any-to-any failover** option and click **Next**.

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

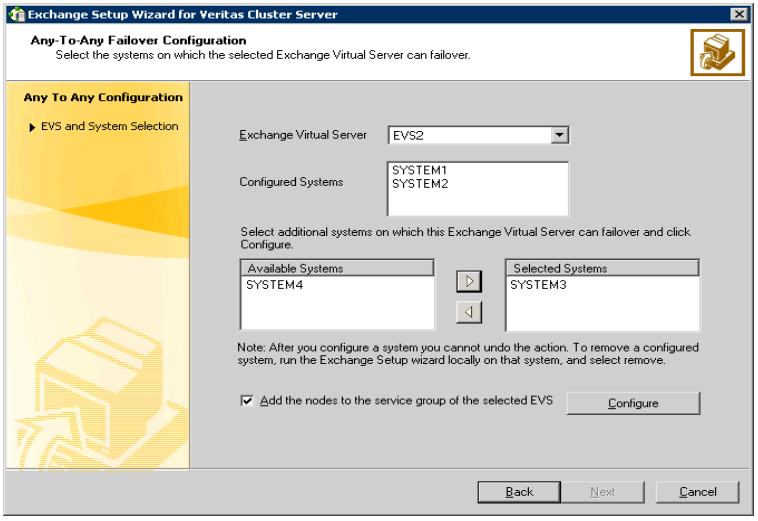


- Select the Exchange virtual server to which you want to add the additional failover nodes.
- The Configured Systems box displays the nodes on which the Exchange Server has been installed.
- From the **Available Systems** box, select the systems to be configured for any-to-any failover.
- The **Available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

6 Click **Finish**.

# Configuring the Exchange service group for an additional Exchange Virtual Server

A new Exchange service group must be configured for the new Exchange virtual server, EVS2. Configuring the Exchange service group involves creating an Exchange service group and defining the attribute values for its resources. After the Exchange service group is created, you must configure the databases to mount automatically at startup.

## Prerequisites

- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk groups and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
  - Exchange database
  - registry changes related to Exchange
  - transaction logs for the first storage group

  See "Managing disk groups and volumes" on page 136.
- 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 Appendix A, "VCS agent for Exchange Server 2007" 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 Wizards > Application**

**Agent for Exchange Server 2007> Exchange Server 2007 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 Specify the service group name and system list:

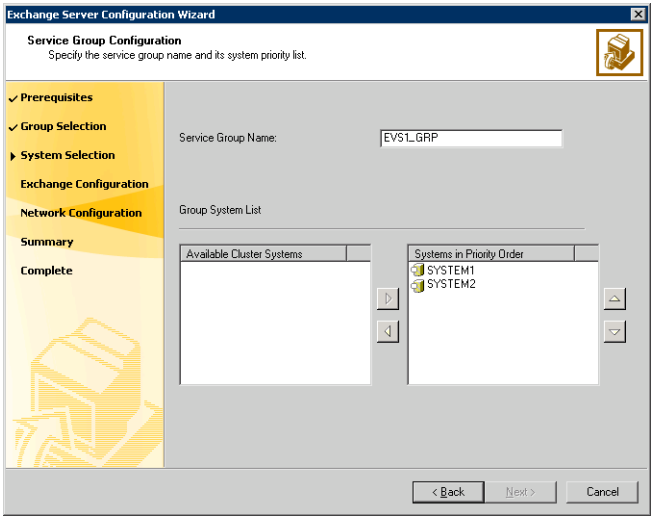


- Enter a name for the Exchange service group.
  If you are configuring the service group on the secondary site, ensure that the name matches the service group on the primary site.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow 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.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

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

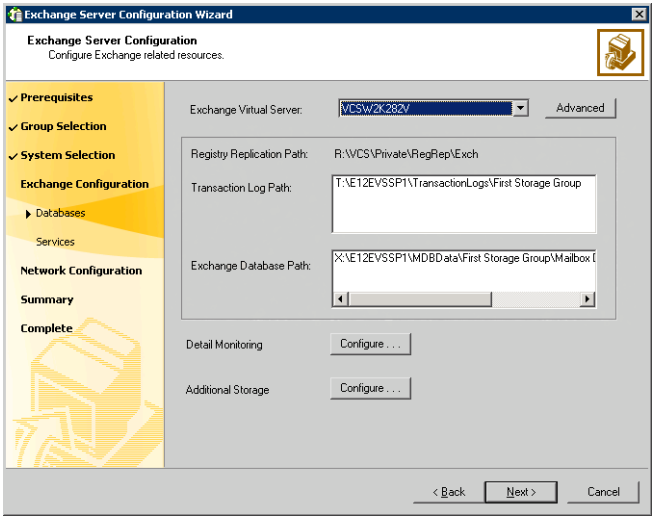


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
  - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.

- To configure Detail Monitoring for Exchange databases, click **Configure...**.

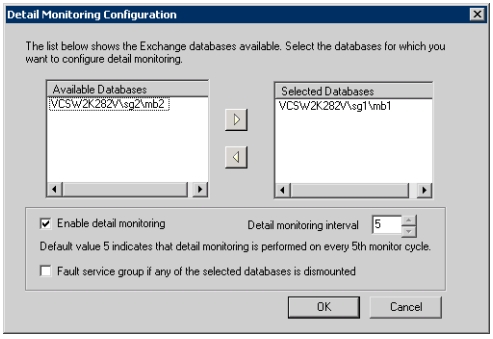


  On the Detail Monitoring Configuration dialog box, complete the following:

  - In the Available Databases box, select the databases for detail monitoring and double-click, or click the right-arrow button to move them to the Selected Databases box. To remove a database, select the database in the Selected Databases box, and double-click or click the left-arrow button.
  - Check **Enable detail monitoring** check box, and specify the monitoring interval in the **Detail monitoring interval** field.
  - If you want the VCS agent to fault the service group if a database selected for detail monitoring is dismounted, check the **Fault service group if any of the selected database is dismounted** check box.

  See the VCS agent attribute descriptions in the Appendix, for more information on detail monitoring and VCS agent behavior.

  - Click **OK**.

- To configure additional storage, click **Configure...**. On the Additional Storage Configuration dialog box, complete the following:
  - 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 (MSExchangeSA) service resource.

- Click **Next**.

6 Select the optional Exchange services to be monitored and click **Next**. Each optional service that is selected will be configured as a VCS resource of type ExchService2007.

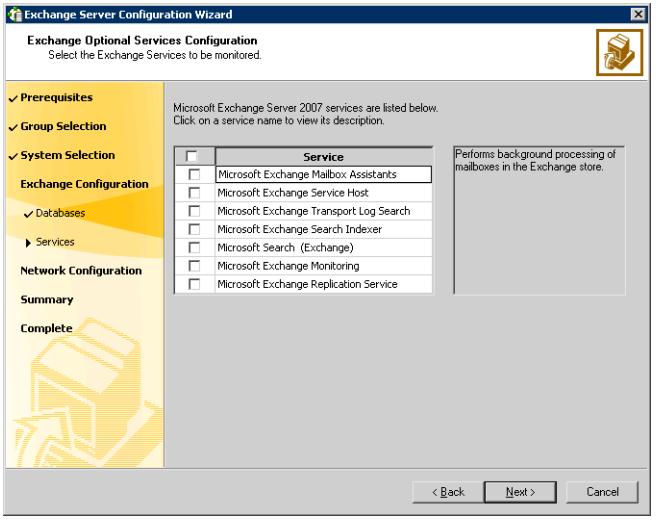


7 Specify information related to the network:

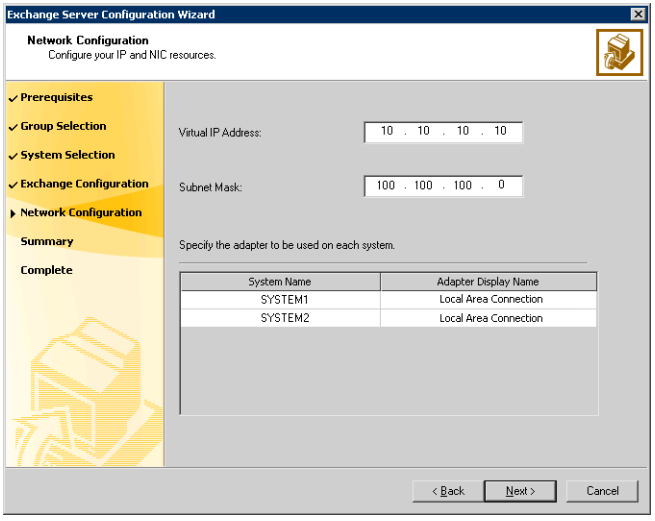

- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.

  If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

---

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

---

- Click **Next**.

8 Review the service group configuration and change the resource names, if desired:

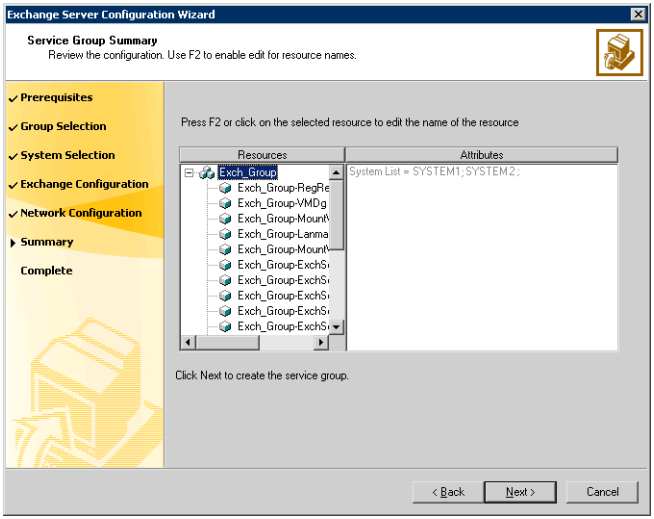


The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

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

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

- Click **Next**.
- A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

9 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system.

10 Click **Finish**.

After bringing the service group online, you must run the Exchange Management Console so that all the stores are automatically mounted on start-up.

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

## 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.

Chapter 5

# Deploying SFW HA for high availability: Configuring a standalone Exchange server

This chapter covers the following topics:

# Tasks for converting a standalone Exchange server into a clustered server

You can 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.

HA support for Exchange Server 2007 is available for the Mailbox Server role only. If you have installed other server roles on the server where you have installed the Mailbox Server role, remove those server roles before you proceed.

**Warning:** You cannot use the Solutions Configuration Center to convert a standalone Exchange server into a clustered server. Instead use the instructions in this Solutions Guide.

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

**Table 5-1** Task list for converting a standalone Exchange server into a clustered server

| Objective | Tasks |
|---|---|
| "Reviewing the requirements" on page 184 | ■ Verifying hardware and software prerequisites |
| "Reviewing the configuration" on page 188 | ■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster |
| "Configuring the network and storage" on page 192 | ■ Setting up the network and storage for a cluster environment<br>■ Verifying the DNS entries for the systems on which Exchange will be installed |

**Table 5-1** Task list for converting a standalone Exchange server into a clustered server (continued)

| Objective | Tasks |
|---|---|
| "Installing Veritas Storage Foundation HA for Windows" on page 193 | ■ Checking the prerequisites<br>■ Verifying the driver signing options for Windows 2003 systems<br>■ Installing SFW, and VCS<br>■ Restoring driver signing options for Windows 2003 systems |
| "Installing the SFW HA patch for Exchange Server 2007" on page 201 | ■ Installing the SFW HA patch; involves installing the VCS agent for Exchange Server 2007<br>■ If required, installing the SFW patch to perform VSS-based backup and restore operations with Exchange Server 2007 |
| "Configuring disk groups and volumes" on page 202 | ■ Using the VEA console to create disk groups<br>■ Using the VEA console to create the data, log, and RegRep volumes |
| "Managing disk groups and volumes" on page 208 | ■ 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 211 | ■ Converting the standalone Exchange server into a cluster node using the Exchange Setup Wizard for Veritas Cluster Server |

**Table 5-1** Task list for converting a standalone Exchange server into a clustered server (continued)

| Objective | Tasks |
|---|---|
| "Adding the standalone Exchange server to a cluster" on page 212 | ■ Configuring the cluster<br>■ For a new cluster, creating the cluster, "Creating a new cluster and adding nodes" on page 214<br>■ For an existing cluster, adding the new nodes to the cluster, "Adding nodes to an existing cluster" on page 231 |
| "Moving Exchange databases to shared storage" on page 241 | ■ 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 244 | ■ 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 250 | ■ Creating the Exchange service group using the VCS Exchange Configuration Wizard |
| "Verifying the cluster configuration" on page 257 | ■ Verifying the cluster configuration by switching service groups and shutting down an active cluster node |

# Reviewing the requirements

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

## Disk space requirements

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

Table 5-2 estimates disk space requirements for SFW HA.

**Table 5-2** Disk space requirements

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

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

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

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

### Supported software

- Veritas Storage Foundation HA 5.0 for Windows (SFW HA) with the Veritas Cluster Server Application Agent for Microsoft Exchange
- Microsoft Exchange servers and their operating systems:
  - Microsoft Exchange Server 2007 Standard Edition or Enterprise Edition
    *with*
    Windows Server 2003 x64 Standard Edition, Enterprise Edition, Datacenter Edition (SP1 required for all editions, SP2 supported)
    *or*
    Windows Server 2003 R2 x64 Standard Edition, Enterprise Edition, Datacenter Edition

## System requirements

Systems must meet the following requirements:

- Processor: x64 architecture-based computer with Intel processor that supports Intel Extended Memory 64 Technology (Intel EM64T) or AMD processor that supports the AMD64 platform; Intel Itanium family IA64 processors are not supported.
- Memory: minimum 2 GB of RAM per server.
- File format: Disk partitions must be formatted for the NTFS file system.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs.
  See "Best practices" on page 188.
- All servers must run the same operating system, service pack level, and system architecture.

## Network requirements

This section lists the following network requirements:

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

- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
  Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
  See the *Veritas Cluster Server Bundled Agents Reference Guide*.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

## Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be an Exchange Full Administrator.
- You must be a member of the Exchange Servers group.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- You must have delete permissions on the object if a computer object corresponding to the Exchange virtual server exists in the Active Directory.
- The user for the pre-installation, installation, and post-installation phases for Exchange must be the same user.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

## Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications

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

### Best practices

Symantec recommends that you perform the following tasks:

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

## Reviewing the configuration

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

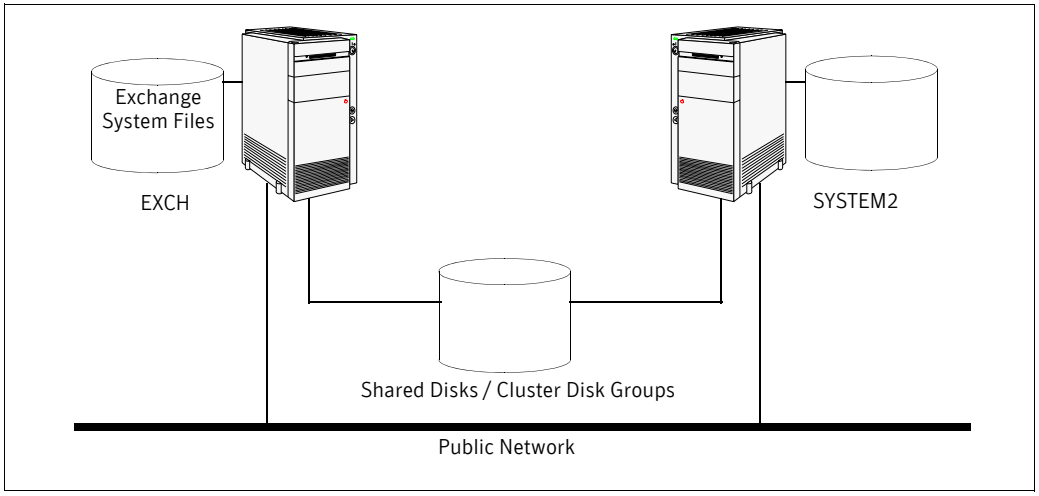


During the following procedures, the initial standalone Exchange server will become part of a new cluster which includes SYSTEM2, be renamed, and become an Exchange virtual server, allowing failover capabilities.

**Figure 5-2** Standalone to active / passive completed configuration

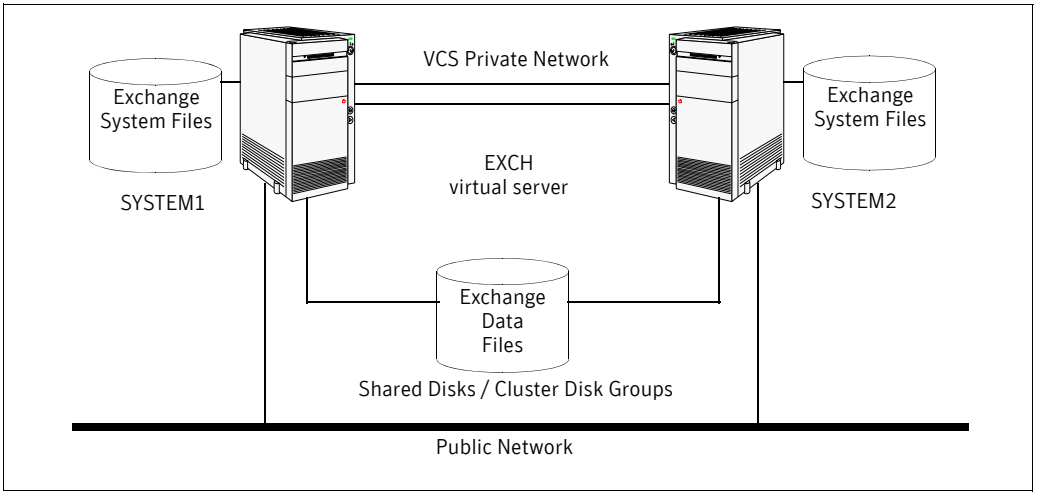


In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group

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

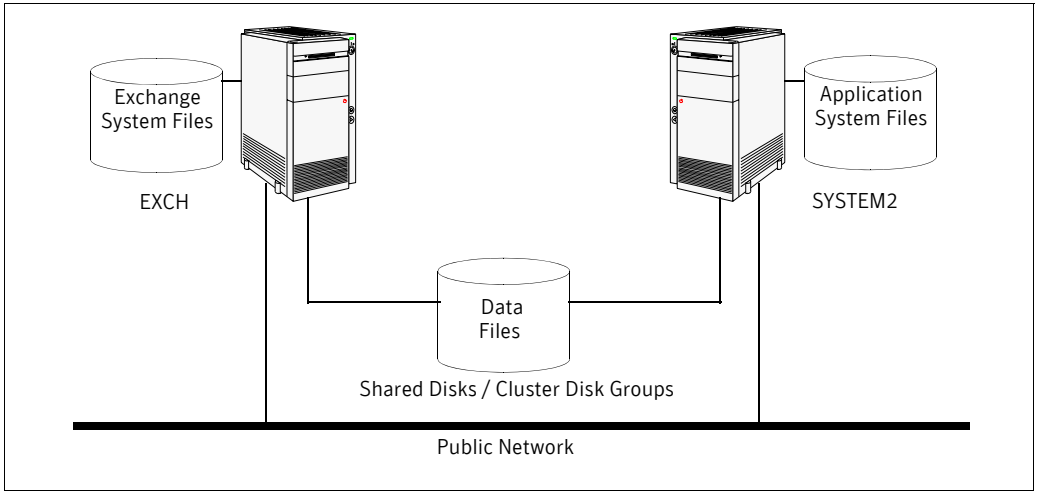


During the following procedures, the initial standalone Exchange server will receives a new physical node name and the original physical node name becomes the name of the Exchange virtual server, allowing failover capabilities within the existing cluster.

**Figure 5-4** Standalone to active / passive completed configuration

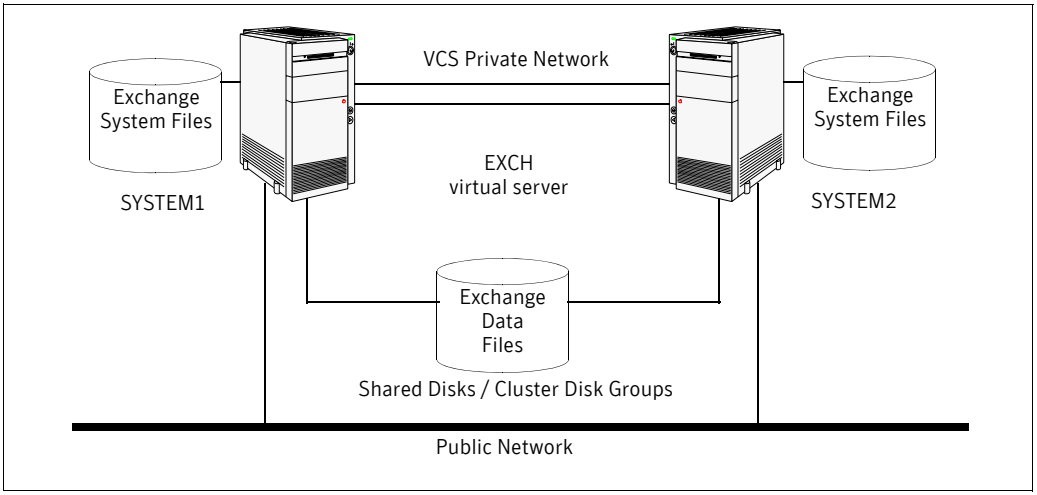


In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. In this case, the Exchange virtual server can fail over from SYSTEM1 to SYSTEM2 and vice versa.

## Sample configuration

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

**Table 5-3** Sample configuration

| Name | Object |
|---|---|
| (EXCH) SYSTEM1, SYSTEM2 | Physical node names; SYSTEM1 was EXCH standalone. |
| EVS1 (EXCH) | Microsoft Exchange Virtual Server |
| EVS1_GRP | Microsoft Exchange service group |
| EVS1_SG1_DG | Cluster disk group name |
| EVS1_SG1_DB1 | volume for storing the Microsoft Exchange Server database |

**Table 5-3** Sample configuration

| Name | Object |
|---|---|
| EVS1_SG1_LOG | Volume for storing a Microsoft Exchange Server database log file |
| EVS1_REGREP | Volume that contain the list of registry keys that must be replicated among cluster systems for the Exchange server |

# Configuring the network and storage

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

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

### To configure the hardware

1. Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
2. Connect the network adapters on each system.
3. Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
4. Verify that each system can access the storage devices.
5. Reboot each system. Verify that each system recognizes the attached shared disk.
6. Use Windows Disk Management on each system to verify that the attached shared disks are visible.

### To verify the DNS settings for all systems that will run Exchange

1. Open the Control Panel (**Start > Control Panel**).
2. Open Network and Dial-up Connections.
3. Ensure the public network adapter is the first bound adapter:

- From the **Advanced** menu, click **Advanced Settings**.
- In the **Adapters and Bindings** tab, verify the public adapter is the first adapter in the **Connections** list. If necessary, use the arrow button to move the adapter to the top of the list.
- Click **OK**.

4 In the Network and Dial-up Connections window, double-click the adapter for the public network. When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

5 From the status window, click **Properties**.

6 In the General tab:

- Select the **Internet Protocol (TCP/IP)** check box.
- Click **Properties**.

7 Select the **Use the following DNS server addresses** option.

8 Verify the correct value for the IP address of the DNS server.

9 Click **Advanced**.

10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.

11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

12 Click **OK**.

# Installing Veritas Storage Foundation HA for Windows

Make sure to review the prerequisites for permissions in "Reviewing the requirements" on page 184.

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 188, SFW HA must be installed on both the standalone Exchange server and the node that will serve as the failover node.

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

## Setting Windows driver signing options

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

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

**Table 5-4** Installation behavior with unsigned drivers

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

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

**To change the driver signing options on each system**

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

4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.

5 Click **OK**.

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

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

### To install the product

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

2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.

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


4 Do one of the following:

- Click **Complete/Custom** to begin installation.
- Click the **Administrative Console** link to install only the client components.

5 Review the Welcome message and click **Next**.

6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.

7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.

8 Repeat for additional license keys. Click **Next**
- To remove a license key, click the key to select it and click **Remove**.
- To see the license key's details, click the key.

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

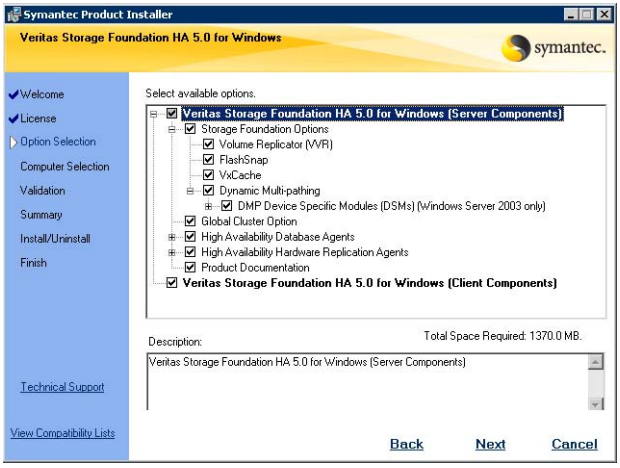


The bottom of the screen displays the total hard disk space required for the installation and a description of an option.

| | |
|---|---|
| Veritas Volume Replicator | If you plan to use VVR for replication, you must also select the option to install VVR. |
| Global Cluster Option | Required for a disaster recovery configuration only. |
| Client Components | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console. |

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

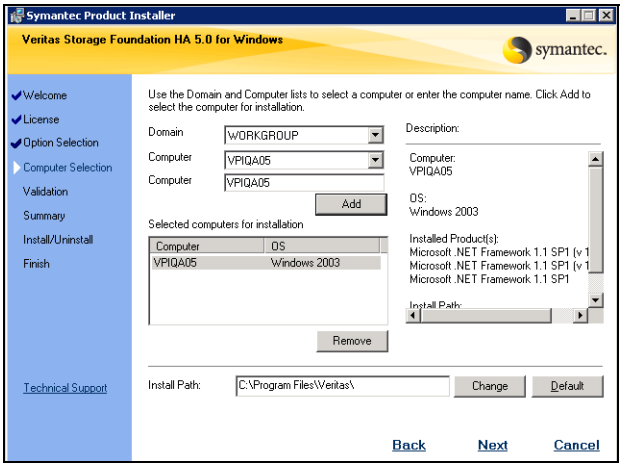


Domain

Select a domain from the list.

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

Computer

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

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

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.
  The default path is:
  `C:\Program Files\Veritas`
  For 64-bit installations, the default path is:
  `C:\Program Files (x86)\Veritas`

11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.

12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.

13 If you are using multiple paths and selected DMP ASLs or a specific DSM you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning:
- For DMP ASLs installations—make sure that you have disconnected all but one path of the multipath storage to avoid data corruption.
- For DMP DSMs installations—the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.

14 Click **OK**.

15 Review the information and click **Install**. Click **Back** to make changes, if necessary.

16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.

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

18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.

19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.

20 Review the log files and click **Finish**.

21 Click **Yes** to reboot the local node.

## Resetting the driver signing options

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

**To reset the driver signing options**

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

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 the SFW HA patch for Exchange Server 2007

The SFW HA patch for Exchange Server 2007 contains the new VCS agent for Exchange Server 2007. The patch zip file contains the following files:

- **vrtsvcsexch.msi** - the VCS agent msi file
- **InstallVCSExch2007.bat** - batch file to install the VCS agent
- **UninstallVCSExch2007.bat** - batch file to remove the VCS agent

Extract these files to a temporary location on the system. Ensure that the agent .msi file and the .bat file are at the same level in a directory.

### To install the SFW HA patch

1. If the system is part of a cluster, complete this step. If not, proceed to step 2.
   - Make sure that all the service groups are offline in the cluster.
   - Save and close the cluster configuration. Type the following on the command prompt:
     ```
     C:\> haconf -dump -makero
     ```
   - Stop the Veritas High Availability engine (HAD) on all the cluster nodes. Type the following on the command prompt:
     ```
     C:\> hastop -all
     ```
2. On the system, double-click **InstallVCSExch2007.bat**. The .msi will install the VCS agent for Exchange Server 2007 on the system.
3. Repeat step 2 on all the systems where you want to install the patch.
4. If the system is part of a cluster, complete this step.
   Start HAD on the system on which you installed the patch. Type the following on the command prompt:
   ```
   C:\> hastart
   ```

# Installing the SFW patch for Exchange Server 2007

The SFW patch for Exchange Server 2007 enables SFW support for performing VSS-based backup and restore operations with Exchange Server 2007.

This step is optional. If required, you can install the SFW patch now. Refer to the readme file accompanying the SFW patch for the installation steps.

# Configuring disk groups and volumes

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

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

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

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

Figure 5-5 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

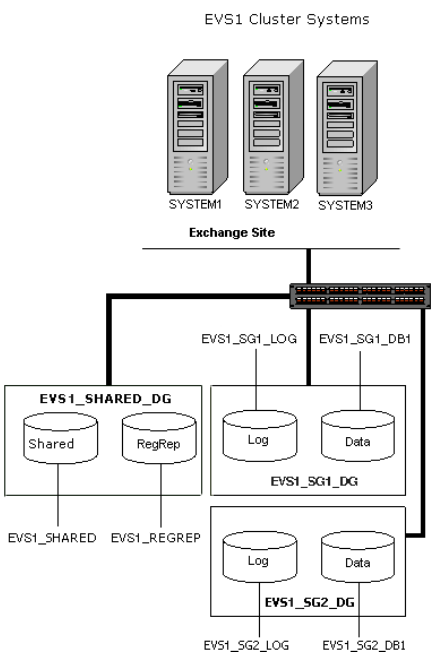

Use the following procedures to create the appropriate disk groups and volumes.The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups. The procedures in this section assume you are using one Exchange database.

Exchange storage group EVS1_SG1_DG create contains two volumes:

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

- EVS1_REGREP: Contains the list of registry keys that must be replicated among cluster systems for the Exchange server.

Additional storage groups (for example, EVS1_SG2_DG) only contain the data, and log volumes.

## Creating a disk group

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

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

6 Provide information about the cluster disk group:

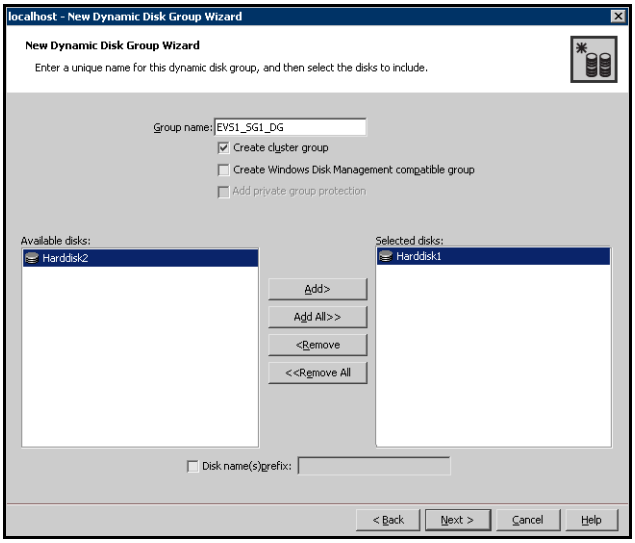


- Enter the name of the disk group (for example, EVS1_SG1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
  Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
- Click **Next**.

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the new disk group.

## Creating volumes

This procedure assumes you are starting with the EXCH_SG1_DB1 volume.

### To create dynamic volumes

1. If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
2. Click **Connect to a Host or Domain**.
3. In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
   To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
4. To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
   You can right-click the disk group you have just created, for example EVS1_SG1_DG.
5. At the New Volume wizard opening screen, click **Next**.
6. Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

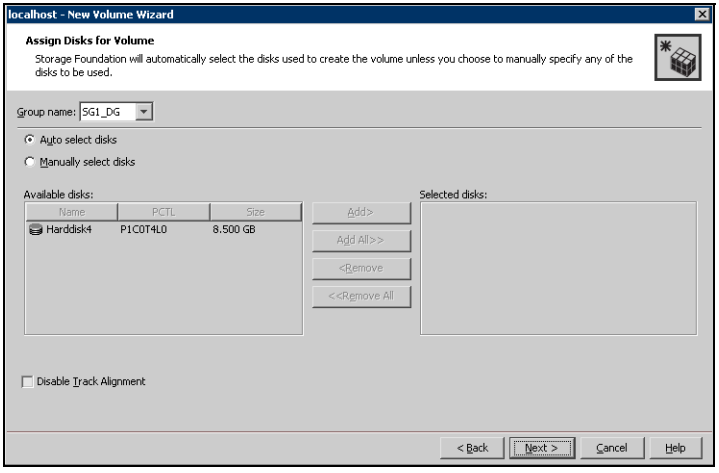


7. Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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

8 Click **Next**.

9 Specify the volume attributes.

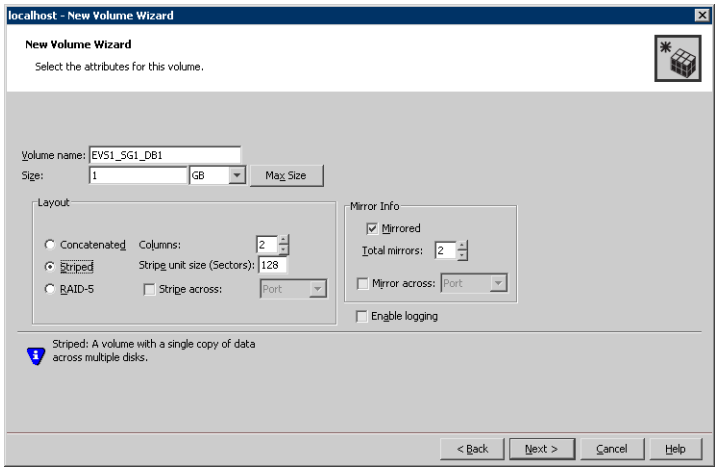


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
- If you are creating a striped volume, the **Columns**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

10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

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

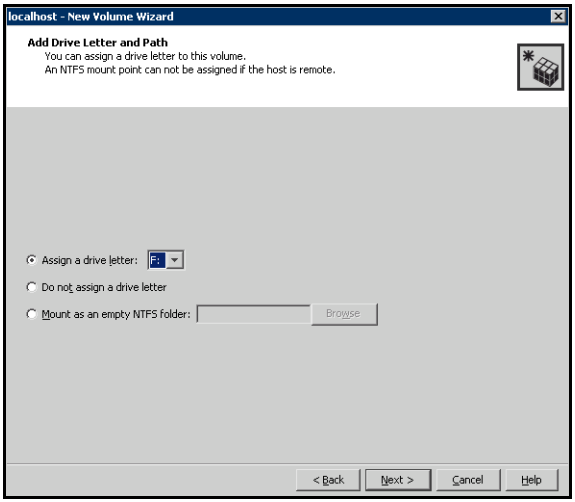


11 Click **Next**.

12 Create an NTFS file system.

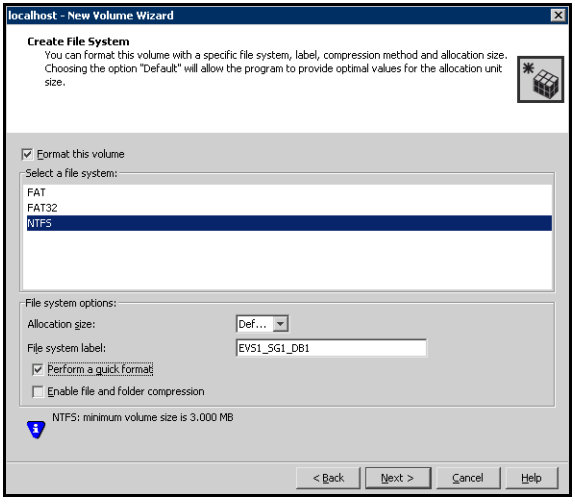


- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.

- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

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

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

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

# Managing disk groups and volumes

This section includes the following procedures:

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

---

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

---

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

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

## Importing a disk group and mounting a volume

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

**To import a disk group**

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

**To mount a volume**

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

## Unmounting a volume and deporting a disk group

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

**To unmount a volume and deport the dynamic disk group**

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

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

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

## Importing a disk group and mounting a volume

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

### To import a disk group

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

### To mount a volume

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

## Unmounting a volume and deporting a disk group

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

**To unmount a volume and deport the dynamic disk group**

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

# Converting the standalone Exchange server into a “clustered” Exchange server

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

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

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

---

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

---

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

1. Start the Exchange Setup Wizard for VCS from the node having the standalone Exchange server installed.
   Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server 2007 > Exchange Server 2007 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 213 and then use "Creating a new cluster and adding nodes" on page 214 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 230 and then continue with the procedure "Adding nodes to an existing cluster" on page 231 to add any new nodes, including the standalone Exchange server, to the cluster.

## Prerequisites for a new cluster

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

Complete the following tasks before creating a cluster:

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

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

If no cluster exists, continue with "Creating a new cluster and adding nodes" on page 214 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 230 and then continue with the

procedure “Adding nodes to an existing cluster” on page 231 to add any new nodes, including the standalone Exchange server, to the cluster.

## Creating a new cluster and adding nodes

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

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

Complete the following tasks before creating a cluster:

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

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

#### To configure a VCS cluster

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

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

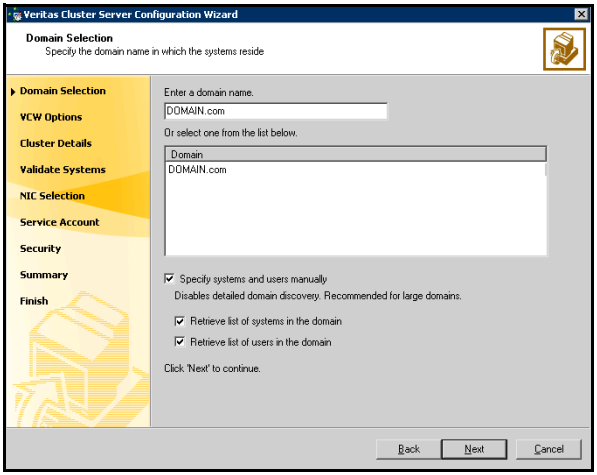


To discover information about all systems and users in the domain:

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

  Proceed to step 7 on page 217.

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 216. Otherwise proceed to the next step.

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

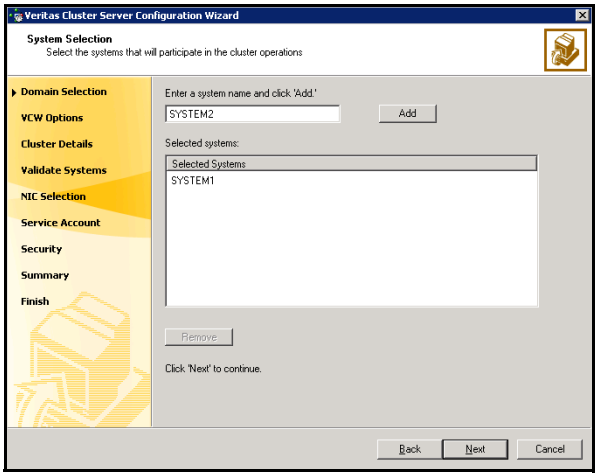


Proceed to step 7 on page 217.

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

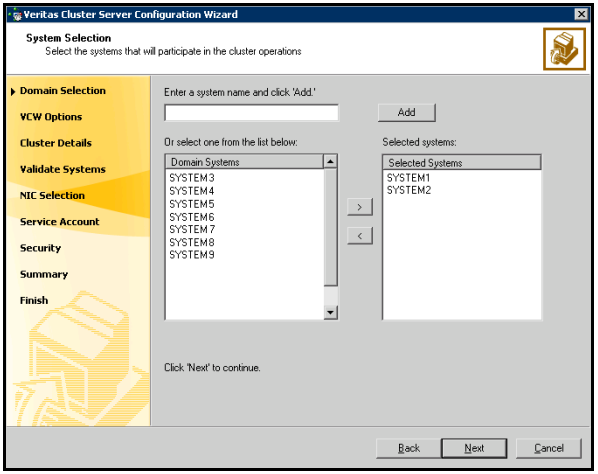


Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

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

8 On the Cluster Details panel, specify the details for the cluster and then click **Next**.

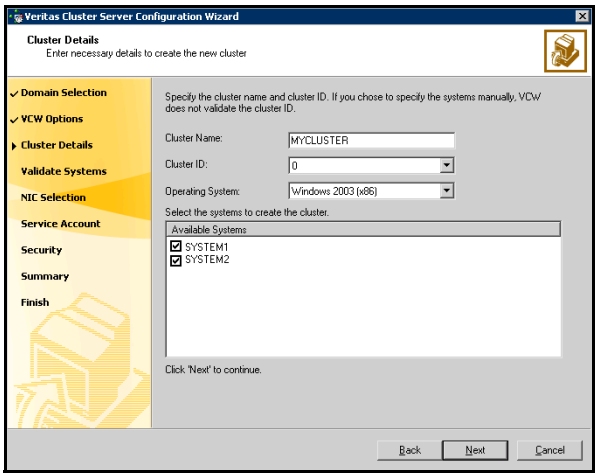


| | |
|---|---|
| Cluster Name | Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. |

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

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

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

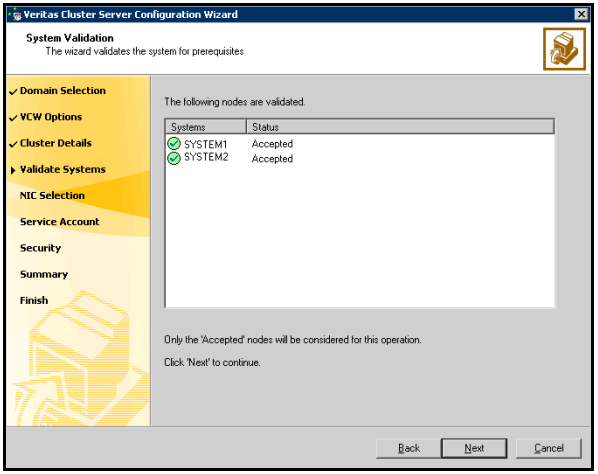


If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in step 8 on page 217, proceed to the next step. Otherwise, proceed to step 11 on page 219.

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

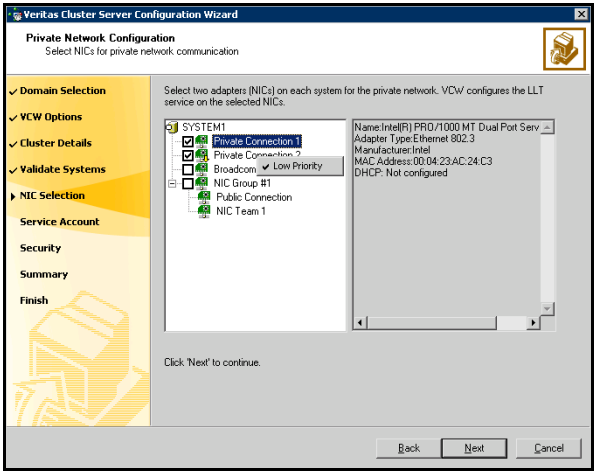


- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of the NIC that is used for public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as NIC Group #N where N is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

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

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

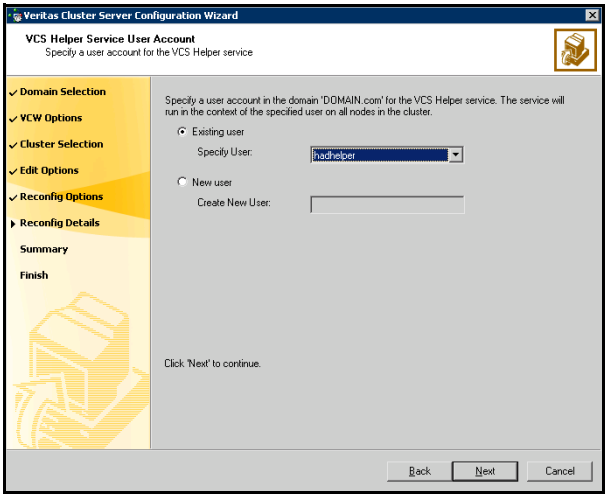


- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list
  - If you chose not to retrieve the list of users in step 4 on page 215, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

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

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

Do one of the following:

- To use the single sign-on feature

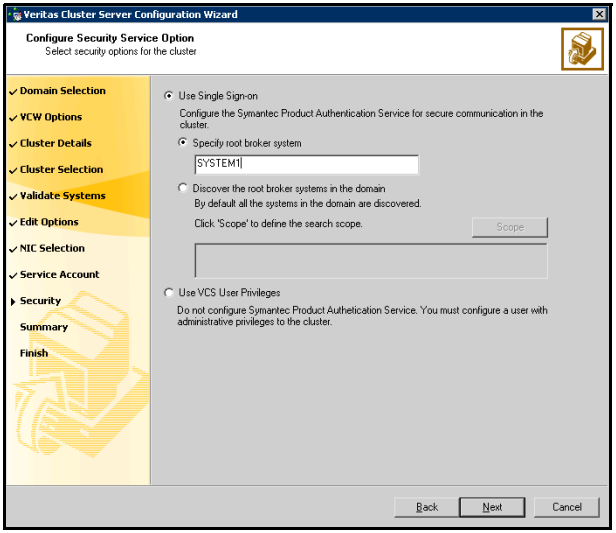


  - Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

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

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

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

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

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

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

- To use VCS user privilege

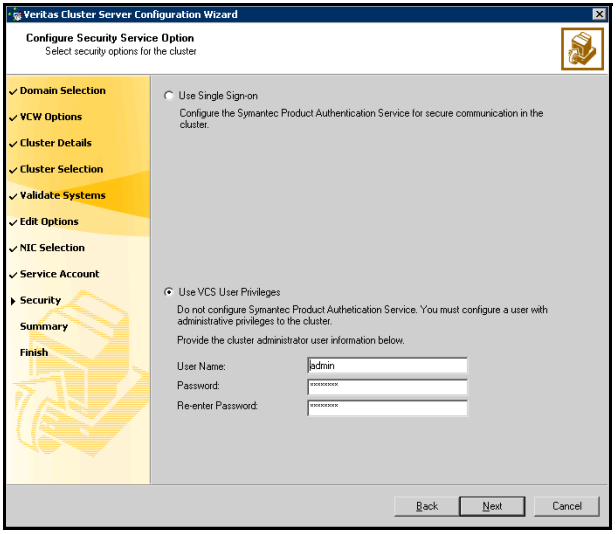


- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

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

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

- Click **Next**.

13 Review the summary information on the Summary panel, and click **Configure**.

The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.

The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

14 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Web Console, notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

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

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

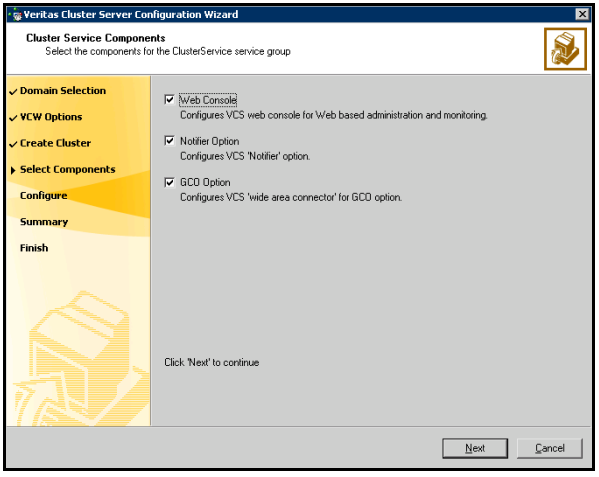

- Check the **Web Console** check box to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
  The GCO Option applies only if you are configuring a Disaster Recovery environment.

## Configuring the Web Console

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

### To configure the Web console

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

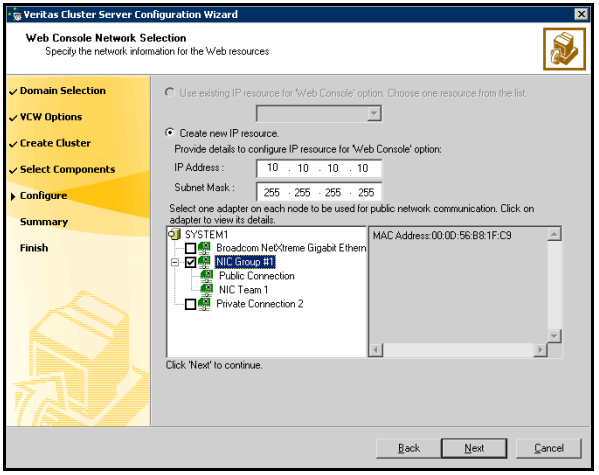


- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
- If you choose to configure a new IP address, type the IP address and associated subnet mask.

- Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.

2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.

3 If you chose to configure a Notifier resource, proceed to "Configuring notification" on page 225.
If you chose to configure global cluster components, proceed to "Configuring the wide-area connector process for global clusters" on page 229.
Otherwise, click **Finish** to exit the wizard.

## Configuring notification

This section describes steps to configure notification.

### To configure notification

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

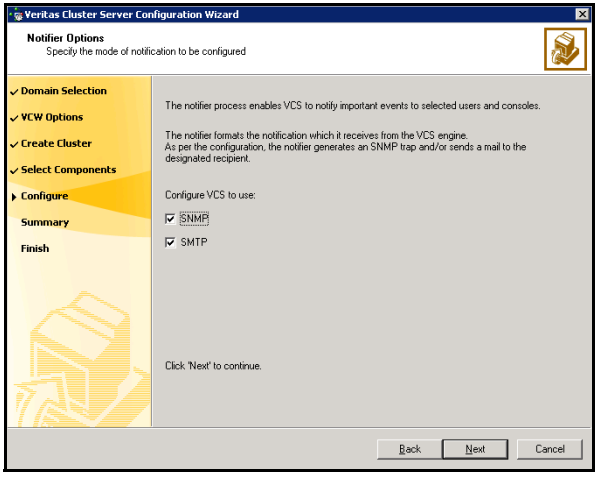


You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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

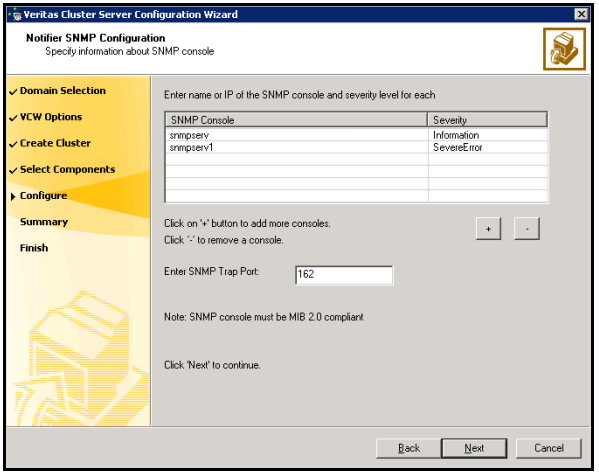


- Click a field in the SNMP Console column and enter the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click **+** to add a field; click **-** to remove a field.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

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

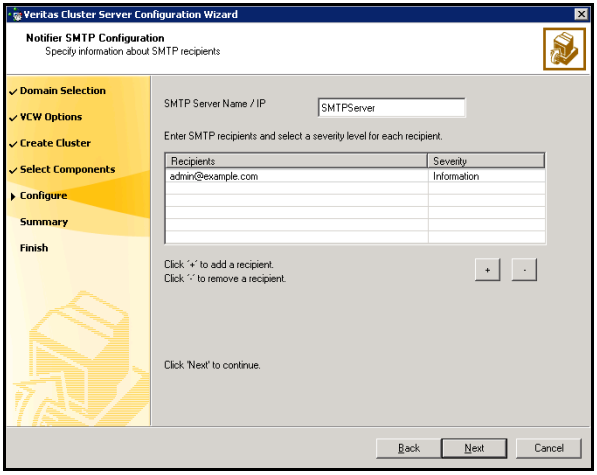


- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity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

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

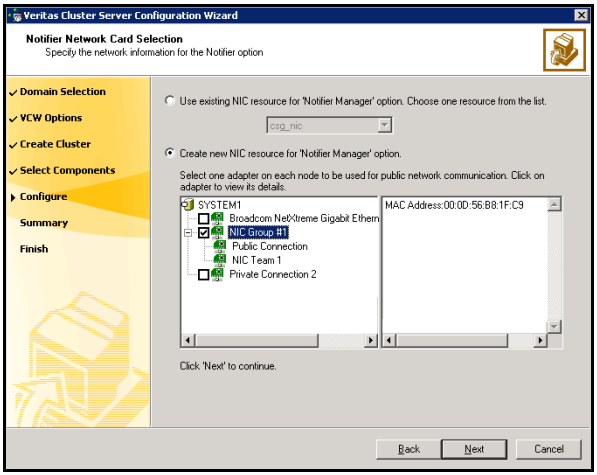


- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
- If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.

5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started and click **Configure**.

6 If you chose to configure global cluster components, proceed to "Configuring the wide-area connector process for global clusters" on page 229. Otherwise, click **Finish** to exit the wizard.

## Configuring the wide-area connector process for global clusters

This section describes steps to configure the wide-area connector resource required for global clusters.

### To configure the wide-area connector process for global clusters

1 On the GCO Network Selection panel, specify the network information and click **Next**.

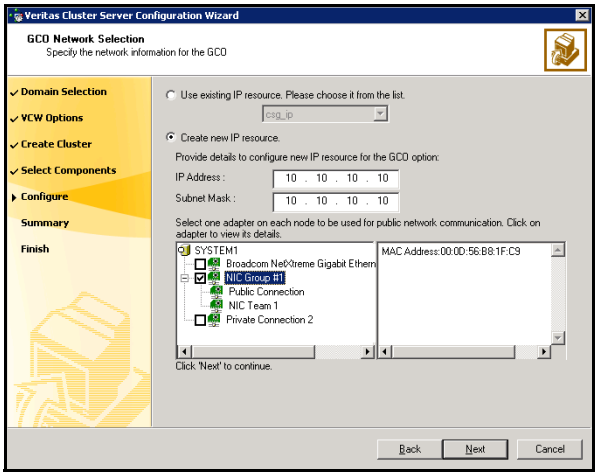


- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
- If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.

2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.

3 Click **Finish** to exit the wizard.

If you have completed the previous procedure, skip to "Moving Exchange databases to shared storage" on page 241.

## Prerequisites for adding nodes to an existing cluster

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

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

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

- You must be a Cluster Administrator.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that Command Server is running on all nodes. Select **Services** on the **Administrative Tools** menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node from where you run the wizard. Select **Services** on the **Administrative Tools** menu and verify that the Veritas High Availability Daemon is running.
- Import the disk group and mount the shared volumes created to store the following data; unmount the drives from other nodes in the cluster:
  - Exchange database
  - Registry changes related to Exchange
  - Transaction logs for the first storage group

  See "Importing a disk group and mounting a shared volume" on page 208 for instructions on mounting and "Unmounting a volume and deporting a disk group" on page 208 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.

See Appendix A, "VCS agent for Exchange Server 2007" on page 313 for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to the EVS1_SG1_DG disk group.

## Adding nodes to an existing cluster

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

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

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

### To add a node to a cluster

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

### To add a node to a VCS cluster

1. Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server> Configuration Wizards > Cluster Configuration Wizard**)
   Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.
2. Read the information on the Welcome panel and click **Next**.
3. On the Configuration Options panel, click **Cluster Operations** and click **Next**.

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

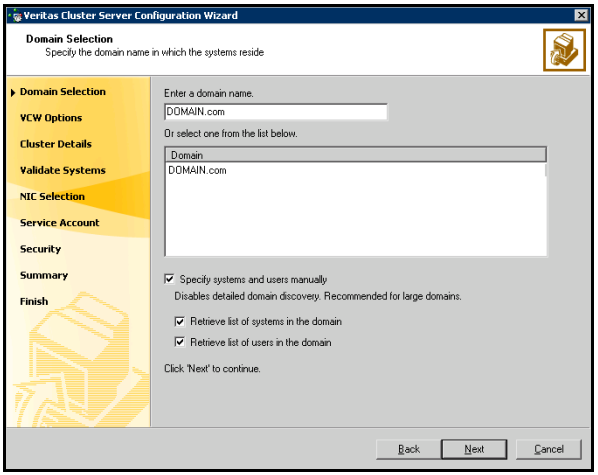


To discover information about all the systems and users in the domain:

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

  Proceed to step 7 on page 235.

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 234. Otherwise proceed to the next step.

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

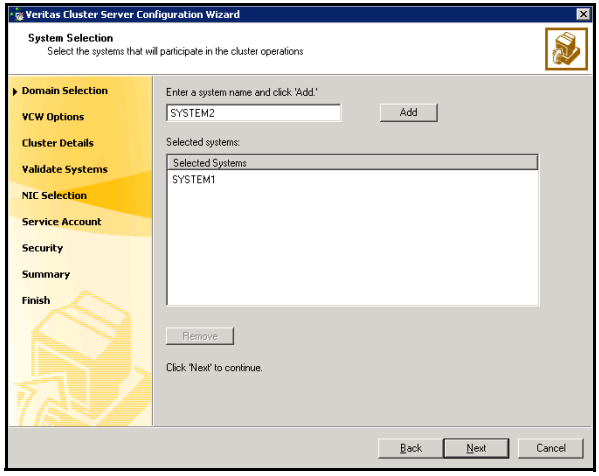


- Type the name of a node in the cluster and click **Add**.
- Type the name of the system to be added to the cluster and click **Add**.

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

Proceed to step 7 on page 235.

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

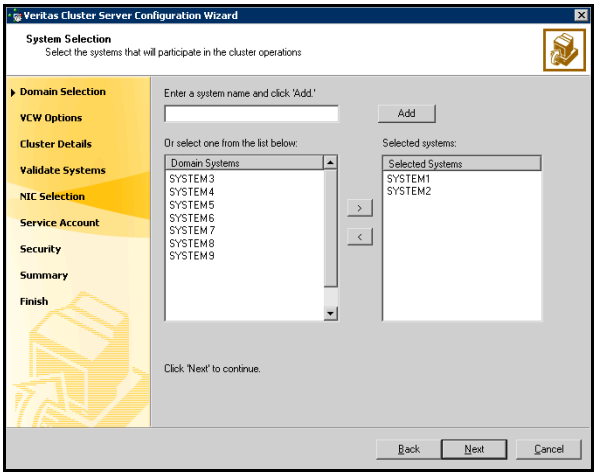


Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

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

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

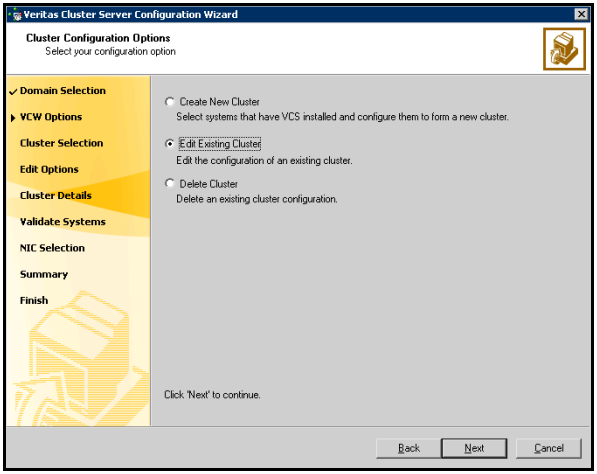


8 On the Cluster Selection panel, select the cluster to be edited and click **Next**.

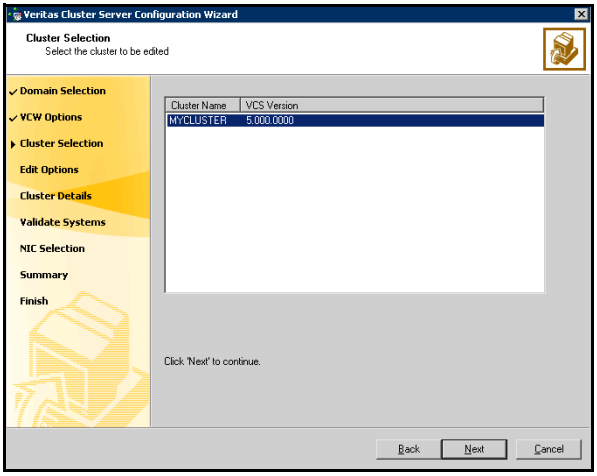


If you chose to specify the systems manually in step 4, only the clusters configured with the specified systems are displayed.

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

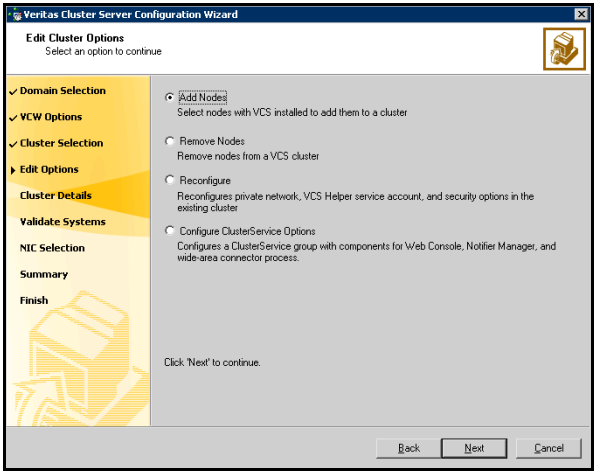


In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

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

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

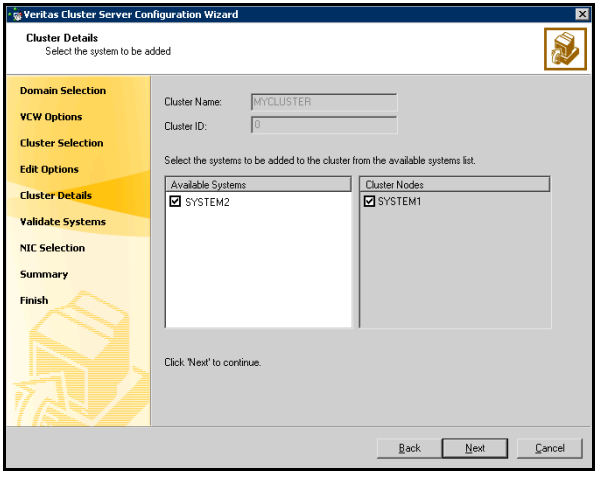

The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

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

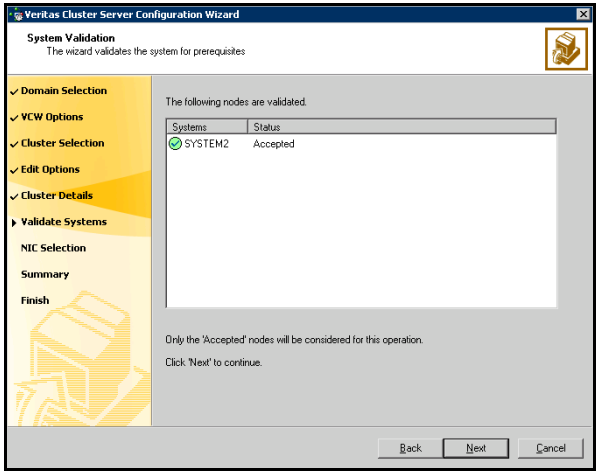


If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.

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

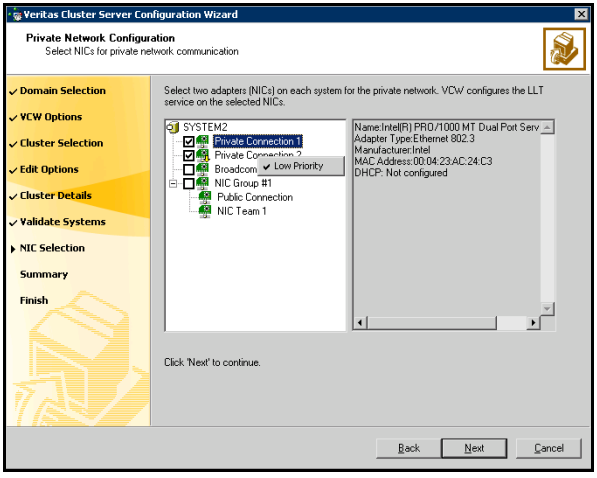

- Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of the NIC that is used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

13 On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.

14 Specify the credentials for the user in whose context the VCS Helper service runs.

15 Review the summary information and click **Add**.

16 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

## Modifying values for ClusterService group attributes

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

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

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

If you need the VCS Web Console or notification for the cluster, proceed to the next procedure, “Modifying the ClusterService group for VCS” on page 239.

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 241.

If a new ClusterService group needs created, be sure to complete the procedure, "Configuring the Exchange service group for VCS" on page 250 when this procedure appears in the sequence.

## Modifying the ClusterService group for VCS

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

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

---

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

---

### To configure the ClusterService group

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

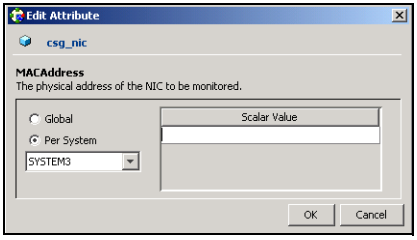


   - Select the **per system** option, and select the newly added system.
   - Enter the scalar value. To obtain the MAC address, run the `ipconfig /all` command from the command prompt on that system.
   - Click **OK**.

If you are only configuring notification, skip to the tasks "Moving Exchange databases to shared storage" on page 241.

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

5 From the Service Groups tab of the Cluster Explorer configuration tree, expand the IP resource type and click the **csg_ip** resource.

6 In the Properties tab of the view panel, click the **Edit** icon for the **MACAddress** attribute.

7 In the Edit Attribute dialog box:

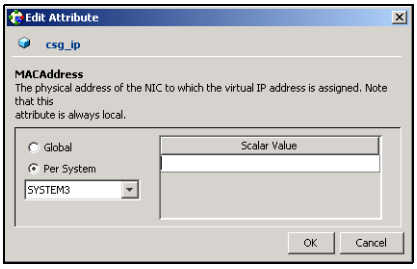


- Select the **per system** option, and select the newly added system.
- Enter the scalar value. To obtain the MAC address, run the `ipconfig /all` command from the command prompt on that system.
- Click **OK**.

8 From the **Service Groups** tab of the Cluster Explorer configuration tree, expand the VRTSWebApp resource type and select the **VCSweb** resource.

9 In the **Properties** tab of the view panel, click the **Edit** icon for the **InstallDir** attribute.

10 In the Edit Attribute dialog box:

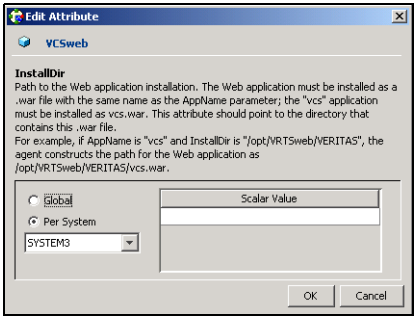


- Select the **per system** option, and select the system.

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

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

# Moving Exchange databases to shared storage

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

- Make sure to import the disk group and mount the volumes for the Exchange database,, and transaction logs.
  See “Managing disk groups and volumes” on page 208.
- The Exchange Setup Wizard for VCS cannot move the Exchange storage groups until local continuous replication (LCR) is suspended for those storage groups. Please suspend LCR using the Exchange Management Console or the Exchange Management Shell, before moving the Exchange databases.
  Refer to the Microsoft Exchange documentation for information on how to suspend LCR.

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

### To move Exchange databases

1 Click **Start > Programs > Symantec > Veritas Cluster Server >Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 Setup Wizard** to start the Exchange Setup Wizard for VCS.

2 Review the information in the Welcome dialog box and click **Next**.

3 In the Available Option dialog box, choose the **Configure/Remove highly available Exchange Server** option and click **Next**.

4 In the Select Option dialog box, choose the **Move Exchange Databases** option and click **Next**.

5 In the Select Exchange Virtual Server dialog box:

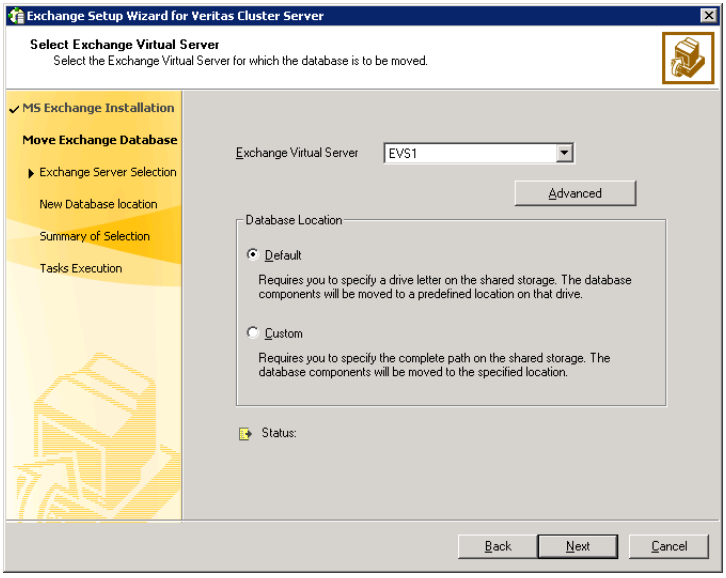


- Select the Exchange virtual server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

---

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

---

- Click **OK**.
- Specify whether you want to move the Exchange databases to a default or custom location. Choosing a custom location allows you to specify the Exchange database and streaming path.
- Click **Next**.

6 Do one of the following:

- If you would like to move the default mailbox store, and the public store to the generated default paths on the volumes for these components, proceed to the next step.
- Otherwise, proceed to step 8 on page 244 to specify the path location on the volumes that you will designate for these components.

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

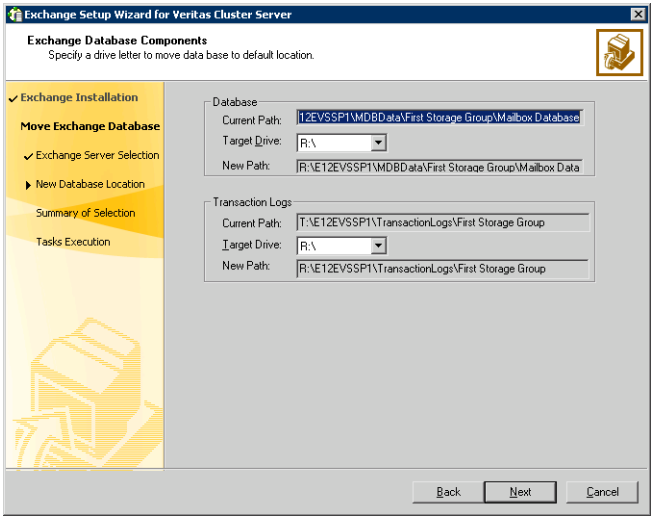


- Specify the drive where the Exchange database will be moved.
- Specify the drive where the Exchange Transaction Logs will be moved.
- Click **Next** and proceed to step 9 on page 244.

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

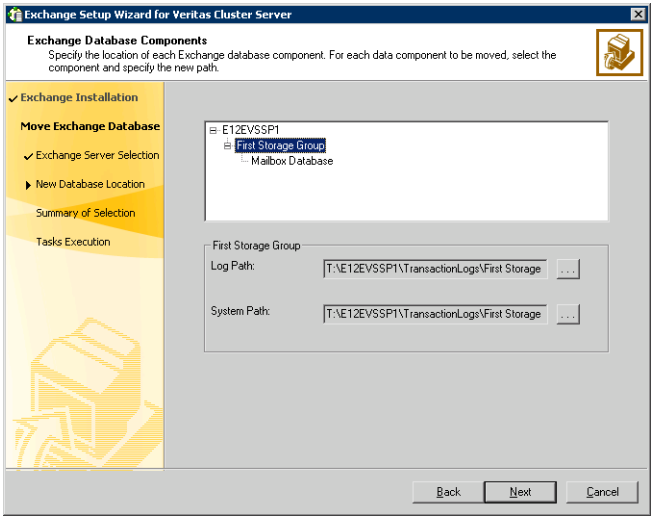


For each data component to be moved select the component and specify the path to which the component is to be moved. Click the ellipsis (...) to browse for folders. Click **Next**.

Make sure the paths for the Exchange database components are not the root of a drive. You must select a directory on the specified drive.

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

9 Review the summary of your selections and click **Next**.

10 The wizard performs the tasks to move the Exchange databases. Messages indicate the status of each task. After all the tasks are completed successfully, click **Next**.

11 Click **Finish** to exit the wizard.

# Installing Exchange on additional nodes

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

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

- Verify the disk group is imported on the existing Exchange node of the cluster. Refer to "Importing a disk group and mounting a shared volume" on page 208 for instructions.
- Mount the volume containing the information for registry replication (EVS1_SG1_REGREP). Refer to "Importing a disk group and mounting a shared volume" on page 208 for instructions.
- Verify that all systems on which Exchange Server will be installed have IIS installed. You must install WWW services on all systems.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
  - You must be a domain user.
  - You must be an Exchange Full Administrator.
  - You must be a member of the Exchange  Servers group.
  - You must be a member of the Local Administrators group for all nodes on which Exchange will be installed. You must have write permissions for objects corresponding to these nodes in the Active Directory.
  - You must have write permissions on the DNS server to perform DNS updates.
  - Make sure the HAD Helper domain user account has "Add workstations to domain" privilege enabled in the Active Directory. To verify this, click **Start > Administrative Tools > Local Security Policy** on the domain controller to launch the security policy display. Click **Local Policies > User Rights Management** and make sure the user account has this privilege.
  - If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.

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

# Exchange pre-installation: additional nodes

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

See "Unmounting a volume and deporting a disk group" on page 208.

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 Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 2007 Mailbox Server role 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 Specify network information for the Exchange virtual server.

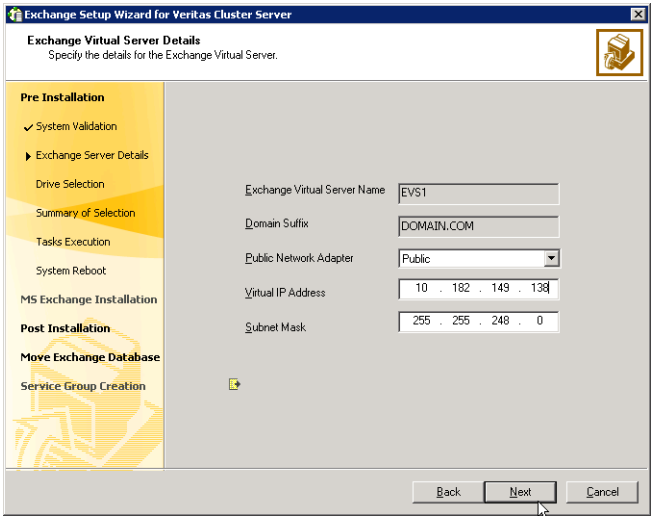


The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.

- 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**.

8 Review the summary of your selections and click **Next**.

9 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.

10 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**.

11 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.

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: additional nodes

Install Exchange on the same node selected in "Installing Exchange on additional nodes" on page 244.

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

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 RecoverServer as the install mode:

   ***<drive letter>***`:\setup.com /mode:recoverserver`

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

2. Setup copies the setup files locally to the computer on which you are installing Exchange 2007 and then checks the prerequisites, including all prerequisites specific to the server roles that you are installing. If you have not met all of the prerequisites, Setup fails and returns an error message that explains the reason for the failure. If you have met all of the prerequisites, Setup installs Exchange 2007.

3. Verify that the installation completed successfully. Refer to the Microsoft documentation for more information.

# Exchange post-installation: additional nodes

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

### To perform Exchange post-installation

1. Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.
2. If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 4.
   If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.
3. Review the information in the Welcome dialog box and click **Next**.
4. A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.
5. The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.
6. 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.

7 Click **Finish**.

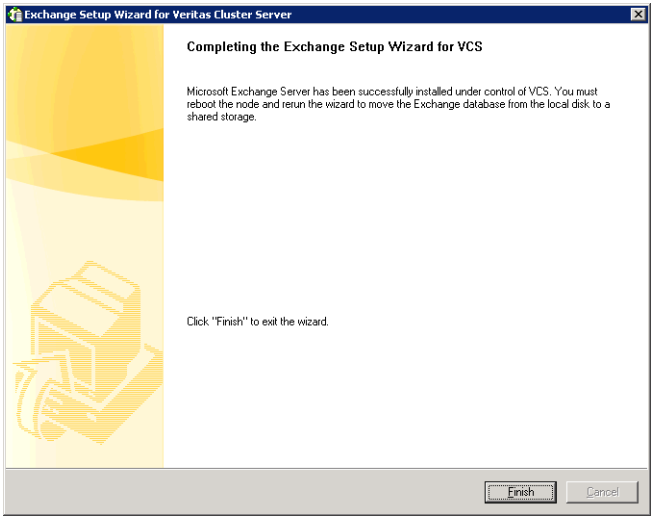


8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

9 Changes made during the post-installation steps do not take effect till you reboot the node.

# 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.
- 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

  See "Importing a disk group and mounting a shared volume" on page 208 for instructions on mounting and "Unmounting a volume and deporting a disk group" on page 208 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 Appendix A, "VCS agent for Exchange Server 2007" on page 313 for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to the EVS1_SG1_DG disk group.

### To configure the Exchange service group

1. Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 Specify the service group name and system list:

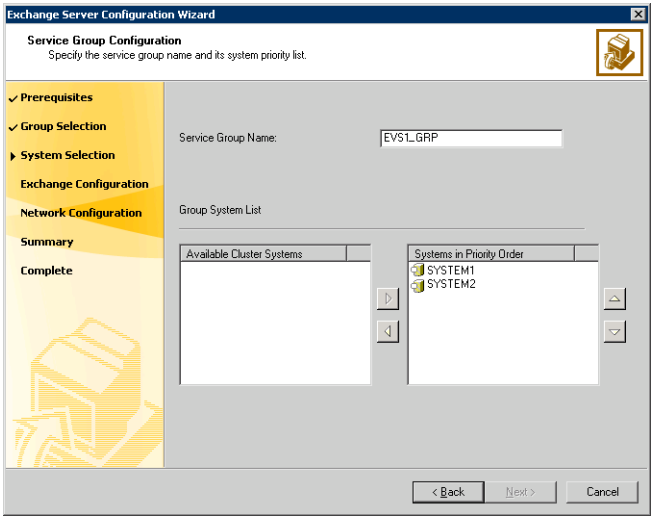


- Enter a name for the Exchange service group.
  If you are configuring the service group on the secondary site, ensure that the name matches the service group on the primary site.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow 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.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

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

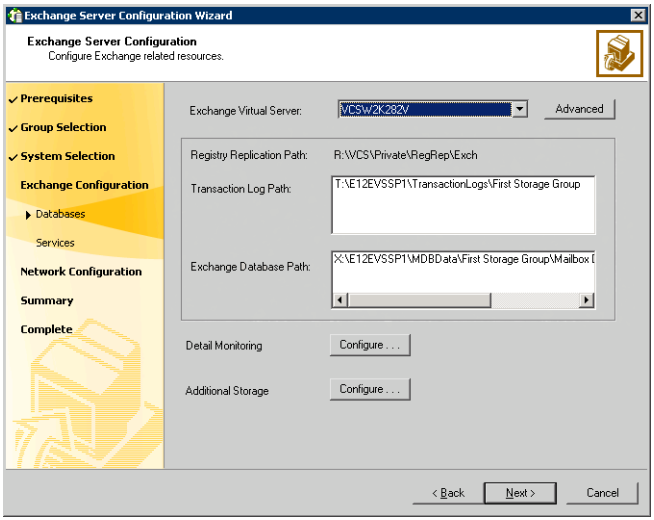


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
  - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.

- To configure Detail Monitoring for Exchange databases, click **Configure...**.

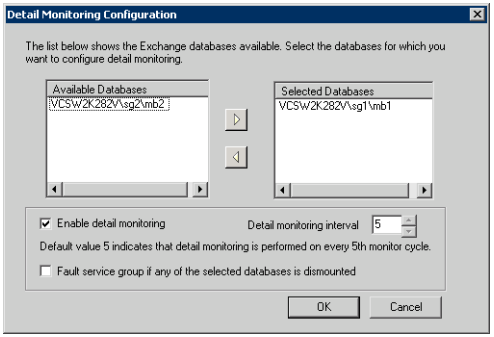


  On the Detail Monitoring Configuration dialog box, complete the following:

  - In the Available Databases box, select the databases for detail monitoring and double-click, or click the right-arrow button to move them to the Selected Databases box. To remove a database, select the database in the Selected Databases box, and double-click or click the left-arrow button.
  - Check **Enable detail monitoring** check box, and specify the monitoring interval in the **Detail monitoring interval** field.
  - If you want the VCS agent to fault the service group if a database selected for detail monitoring is dismounted, check the **Fault service group if any of the selected database is dismounted** check box.

  See the VCS agent attribute descriptions in the Appendix, for more information on detail monitoring and VCS agent behavior.

  - Click **OK**.

- To configure additional storage, click **Configure...**. On the Additional Storage Configuration dialog box, complete the following:

  - 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 (MSExchangeSA) service resource.

- Click **Next**.

6 Select the optional Exchange services to be monitored and click **Next**. Each optional service that is selected will be configured as a VCS resource of type ExchService2007.

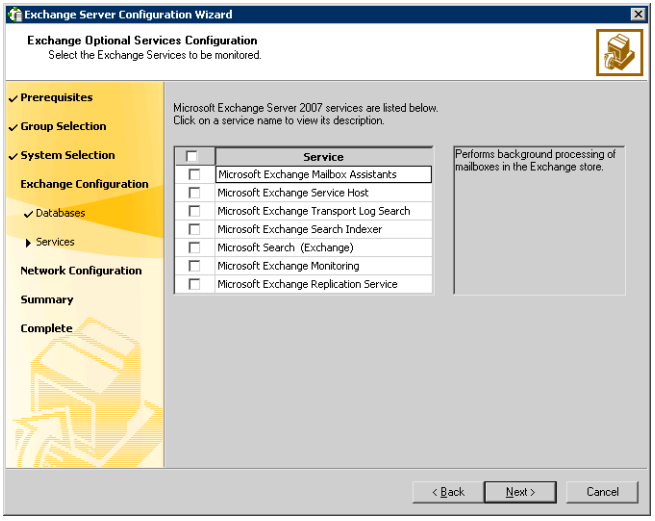


7 Specify information related to the network:

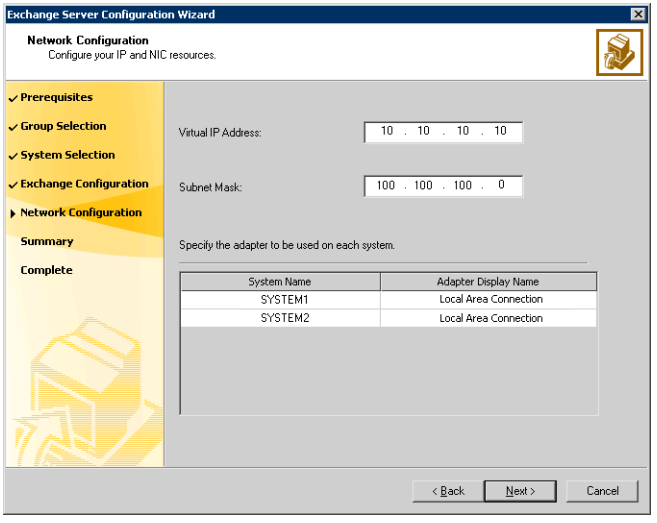

- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.
  If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual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

---

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

---

- Click **Next**.

8 Review the service group configuration and change the resource names, if desired:

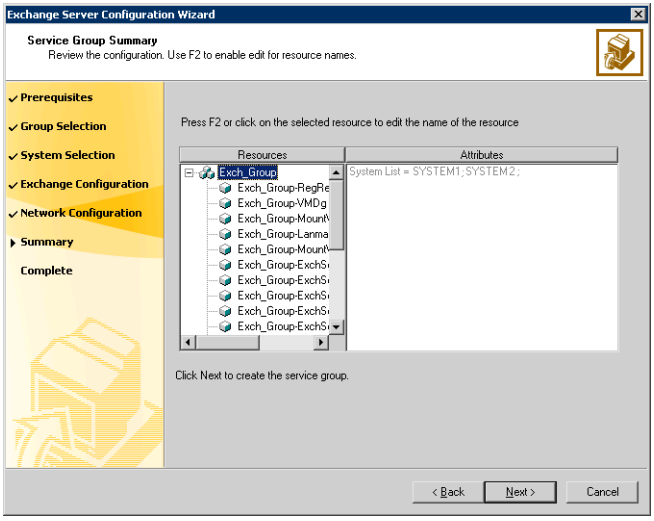


The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

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

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

- Click **Next**.
- A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

9 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system.

10 Click **Finish**.

After bringing the service group online, you must run the Exchange Management Console so that all the stores are automatically mounted on start-up.

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

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.

Section 3

# Disaster Recovery

This section includes the following chapters:

Chapter 6

# Disaster recovery for Exchange: Overview

This chapter includes the following topics:



## 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.

Chapter 7

# Deploying Disaster Recovery: New Exchange Server installation

This chapter covers the following topics:

- Adding a new failover node after DR environment is in operation

# 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.
  See "Tasks for a new HA installation of Microsoft Exchange Server 2007" on page 24.
  While configuring the primary site, you select certain options to prepare the site for disaster recovery.
  See "Verifying the primary site configuration" on page 266.
- 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.
  Table 7-1 outlines the high-level objectives and the tasks to complete each objective for configuring the secondary site.

**Table 7-1** Task List for configuring the secondary site for disaster recovery

| Objective | Tasks |
|---|---|
| "Reviewing the configuration" on page 266 | ■ Understanding Active/Passive configuration and site failover in a DR environment |
| "Verifying the primary site configuration" on page 266 | ■ Verifying the steps to set up HA and DR on the primary site |
| "Setting up the SFW HA environment (secondary site)" on page 267 | ■ Installing SFW HA<br>■ Configuring the disk groups and volumes<br>■ Configuring the cluster using the Veritas Cluster Server Configuration Wizard<br>■ Installing the SFW HA patch and the SFW patch (if required) for Exchange Server 2007<br>■ Configuring the cluster |

**Table 7-1** Task List for configuring the secondary site for disaster recovery

| Objective | Tasks |
|---|---|
| "Installing Exchange on the first node (secondary site)" on page 268 | ■ Reviewing the prerequisite checklist<br>■ 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 273 | ■ Running the Exchange Setup Wizard for Veritas Cluster Server and the Microsoft Exchange Server installation for additional nodes in the same virtual server |
| "Backing up and restoring the Exchange disk group" on page 277 | ■ 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 277 | ■ Creating the Exchange service group using the VCS Exchange Configuration Wizard. |
| "Verifying the cluster configuration (secondary site)" on page 284 | ■ Verifying the cluster configuration by switching service groups and shutting down an active cluster node |
| "Reviewing the prerequisites for configuring DR" on page 286 | ■ Verifying HA prerequisites for DR components |
| "Setting up the replicated data sets (RDS) for VVR" on page 286 | ■ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites<br>■ Using the Setup Replicated Data Set Wizard to create Replicator Log volumes for the primary and secondary sites |

**Table 7-1** Task List for configuring the secondary site for disaster recovery

| Objective | Tasks |
|---|---|
| "Creating the VVR RVG service group" on page 293 | ■ 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 296 | ■ Linking clusters (adding a remote cluster to a local cluster<br>■ Converting the application service group that is common to all the clusters to a global service group<br>■ Converting the local service group to a global group<br>■ Bringing the global service group online |

# Reviewing the configuration

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.

# 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 "Tasks for a new HA installation of Microsoft Exchange Server 2007" on page 24.

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 and VVR option, and ensure that you configure the VVR security service (VxSAS).
- When configuring disk groups and volumes, create 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 to set up SFW HA on the secondary site before continuing with the tasks in this chapter:

- "Reviewing the requirements" on page 26
- "Configuring the storage hardware and network" on page 34
- "Installing Veritas Storage Foundation HA for Windows" on page 36
  - Ensure that you install the GCO option and VVR option.
  - Ensure that you configure the VVR security service (VxSAS).
- "Configuring disk groups and volumes" on page 44
  - Ensure that you create a volume for the VVR Storage Replicator Log for each storage group.
- "Installing the SFW HA patch for Exchange Server 2007" on page 54
- "Configuring the cluster" on page 55
  - 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.

# Installing Exchange: Overview (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 (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.

All the procedures are repeated here for your convenience.

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 52.
- Mount the volume containing the information for registry replication.
- Verify that on all systems on which Exchange Server is to be installed, the services required by Exchange are installed.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation. VCS requires the Exchange installation to take place on the same local drive on all nodes. For example, if you install Exchange on drive C of one node, installations on all other nodes must occur on drive C.
- Set the required permissions:
  - You must be a domain user.
  - You must be an Exchange Full Administrator.
  - You must be a member of the Exchange Servers group.

- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- 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.
- You must have write permissions on the DNS server to perform DNS updates.
- Make sure the HAD Helper domain user account has the "Add workstations to domain" privilege enabled in the Active Directory. To verify this, click Start > Administrative Tools > Local Security Policy on the domain controller to launch the security policy display. Click Local Policies > User Rights Management and make sure the user account has this privilege.

- 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.

## 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 > All Programs > Veritas > Veritas Cluster Server > Application Agent for Exchange 2007 > Exchange Server 2007 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 panel, choose the **Install Exchange 2007 Mailbox Server role for High Availability** option and click **Next**.

5 In the Select Option panel, choose the **Create a failover node for Exchange disaster recovery setup** option and click **Next**.

6 In the Primary Node Selection panel, type the name of the node at the primary site in the **System from the primary site** field and then click **Next**.

7 The wizard validates the system for prerequisites. Various messages indicate the validation status. Once all the validations are done, click **Next**.

8 In the Select Exchange Virtual Server panel, select the Exchange virtual server for disaster recovery and click **Next.**

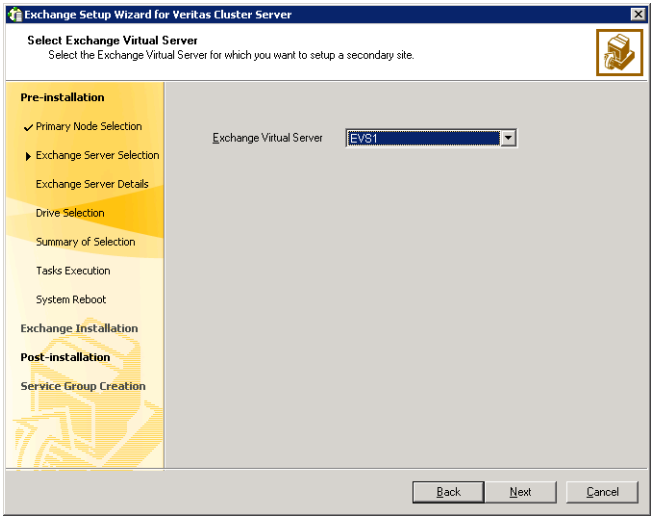


9 The installer verifies that the selected node meets the Exchange requirements. If the service group on the primary node has not been taken offline the installer prompts you to do so without exiting the installer, or you can cancel the installation wizard and take the service group offline manually. Click **Next**

10 Enter the name of a failover node, and click **Next**.

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

- Verify the virtual computer name for the server.
- Verify the domain suffix.
- Enter a unique virtual IP address for the virtual server, or use the same IP address as the virtual server on the primary site.
- Enter the subnet to which the virtual IP address belongs.

- Select the public NIC.

**Warning:** The installer displays all TCP/IP enabled adapters on a node, including the private network adapters. Make sure that you select the adapters for the public network, and not those assigned to the private network.

- Click **Next**.

12. In the Registry Replication Drive panel, select the same drive letter (or directory in the case of folder mounts) as the one used on the primary site for registry replication, and click **Next.**
13. Review the summary of selections and click **Next**.
14. After reviewing the warning of the renaming and rebooting of the system, click **Yes**.
15. After the installer performs configuration tasks, click **Next**.
16. If the wizard could not locate a DNS entry for the specified Exchange server and IP address, click **OK** to create one.
17. Click **Reboot**.
18. Click **Yes** to reboot the node.

**Warning:** After you reboot the node, the values specified for the Exchange virtual server are temporarily assigned to the node. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you rerun this wizard to assign the original name to the node.

19. 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.
    Do not click **Continue** at this time. Wait until after you install Exchange installation.
    If you need to undo all actions performed by the wizard during the pre-installation phase, click **Revert**.

## Exchange installation on first node (secondary site)

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

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

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 RecoverServer as the install mode:
   `<drive letter>:\setup.com /mode:recoverserver`
   where *<drive letter>* is the location where the Exchange software is located.
2. Setup copies the setup files locally to the computer on which you are installing Exchange 2007 and then checks the prerequisites, including all prerequisites specific to the server roles that you are installing. If you have not met all of the prerequisites, Setup fails and returns an error message that explains the reason for the failure. If you have met all of the prerequisites, Setup installs Exchange 2007.
3. Verify that the installation completed successfully. Refer to the Microsoft documentation for more information.

## Exchange post-installation on first node (secondary site)

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

**To perform Exchange post-installation**

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

8 Once the node is rebooted, move the databases created during the Exchange installation from the local drive to the shared drive.

# 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.

Make sure to complete the following tasks before the Exchange installation:

- Review the prerequisites for permissions.
  See "Installing Exchange on the first node (secondary site)" on page 268.
- 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 52.

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 Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 2007 Mailbox Server role 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 Specify network information for the Exchange virtual server.

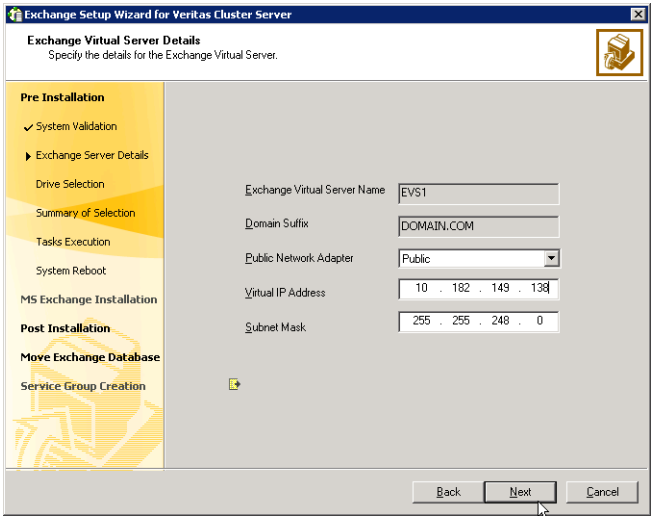


The wizard discovers the Exchange virtual server name and the domain suffix from the Exchange configuration. Verify this information and provide values for the remaining text boxes.

- 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**.

8 Review the summary of your selections and click **Next**.

9 A message appears informing you that the system will be renamed and restarted after you quit the wizard. Click **Yes** to continue.

10 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**.

11 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.
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: Additional nodes

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

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

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 RecoverServer as the install mode:
***<drive letter>***`:\setup.com /mode:recoverserver`
where ***<drive letter>*** is the location where the Exchange software is located.

2 Setup copies the setup files locally to the computer on which you are installing Exchange 2007 and then checks the prerequisites, including all prerequisites specific to the server roles that you are installing. If you have

not met all of the prerequisites, Setup fails and returns an error message that explains the reason for the failure. If you have met all of the prerequisites, Setup installs Exchange 2007.

3 Verify that the installation completed successfully. Refer to the Microsoft documentation for more information.

## Exchange post-installation: Additional nodes

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

### To perform Exchange post-installation

1 Make sure that the volume created to store the registry replication information is mounted on this node and unmounted on other nodes in the cluster.

2 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 4.
If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard is launched automatically.

3 Review the information in the Welcome dialog box and click **Next**.

4 A message appears informing you that the system will be renamed and restarted after you quit the wizard. The system will be renamed to the physical host name. Click **Yes** to continue.

5 The wizard starts performing the post-installation tasks. Various messages indicate the status. After the commands are executed, click **Next**.

6 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.

7 Click **Finish**.

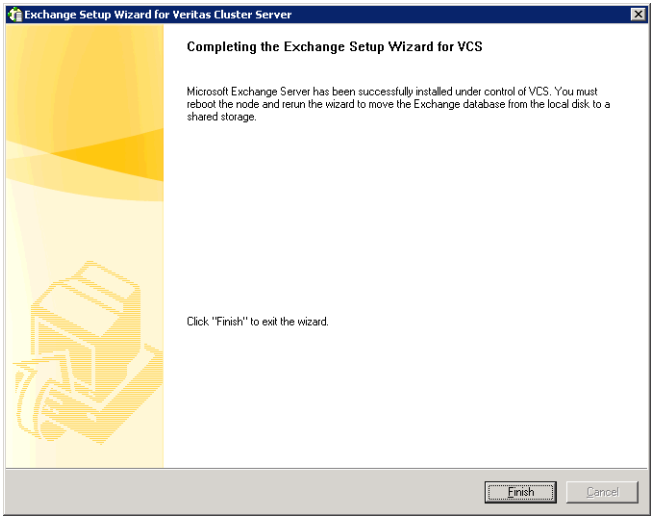


8 The wizard prompts you to reboot the node. Click **Yes** to reboot the node.

9 Changes made during the post-installation steps do not take effect till you reboot the node.

# 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:

- On the primary site, back up all volumes in 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.
- 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 277.
- 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

  See "Managing disk groups and volumes" on page 52.
- 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 Appendix A, "VCS agent for Exchange Server 2007" for information on resource types, attribute definitions, resource dependencies, and sample service group configurations. Refer to the *Veritas Cluster Server Administrator's Guide* to add additional resources to the EVS1_SG1_DG1 disk group.

### To configure the Exchange service group

1. Start the Exchange Server Configuration Wizard. **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Agent for Exchange Server 2007> Exchange Server 2007 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 Specify the service group name and system list:

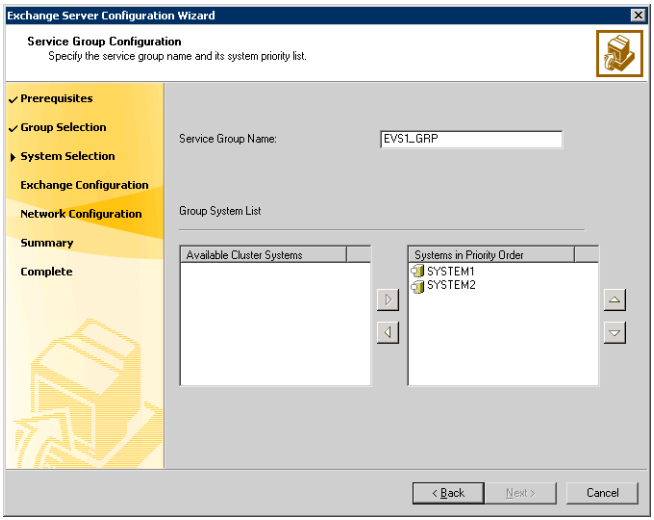


- Enter a name for the Exchange service group.
  If you are configuring the service group on the secondary site, ensure that the name matches the service group on the primary site.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow 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.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

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

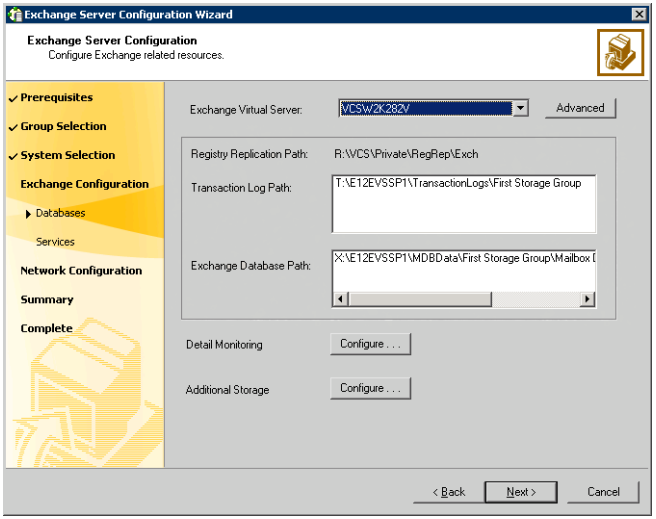


- Specify the Exchange Virtual Server.
- If desired, click **Advanced** to specify details for the Lanman resource.
  - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
  - Click **OK**.
- Verify the Exchange Database Path.
- Verify the Transaction Log Path.

- To configure Detail Monitoring for Exchange databases, click **Configure...**.

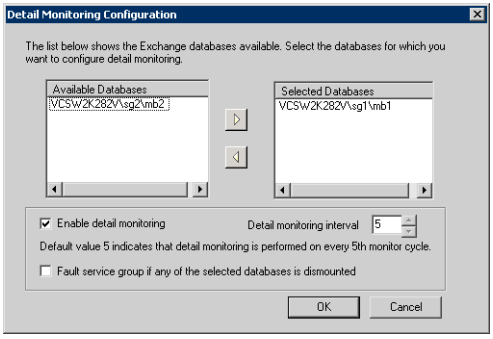


  On the Detail Monitoring Configuration dialog box, complete the following:

  - In the Available Databases box, select the databases for detail monitoring and double-click, or click the right-arrow button to move them to the Selected Databases box. To remove a database, select the database in the Selected Databases box, and double-click or click the left-arrow button.
  - Check **Enable detail monitoring** check box, and specify the monitoring interval in the **Detail monitoring interval** field.
  - If you want the VCS agent to fault the service group if a database selected for detail monitoring is dismounted, check the **Fault service group if any of the selected database is dismounted** check box.

  See the VCS agent attribute descriptions in the Appendix, for more information on detail monitoring and VCS agent behavior.

  - Click **OK**.

- To configure additional storage, click **Configure...**. On the Additional Storage Configuration dialog box, complete the following:

  - 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 (MSExchangeSA) service resource.

- Click **Next**.

6 Select the optional Exchange services to be monitored and click **Next**. Each optional service that is selected will be configured as a VCS resource of type ExchService2007.

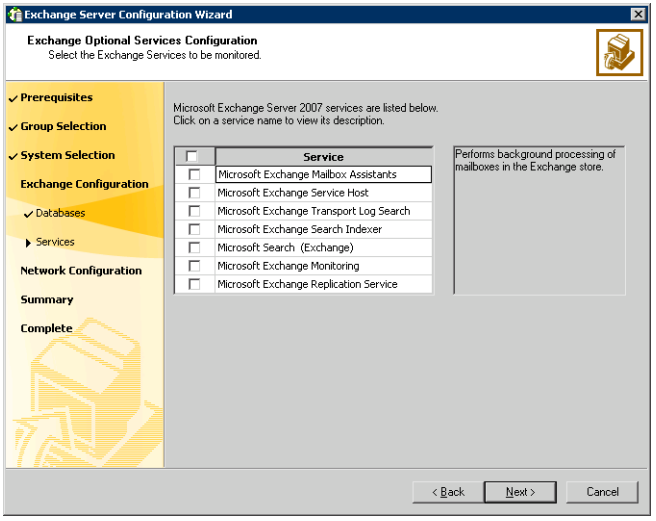


7 Specify information related to the network:

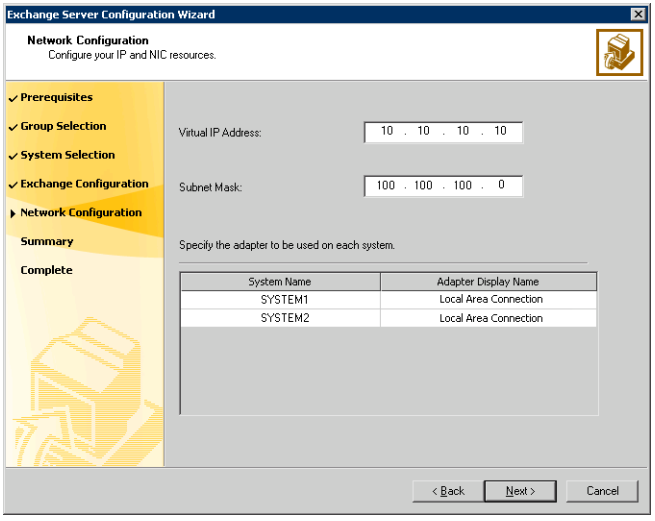

- The Virtual IP Address and the Subnet Mask text boxes display the values entered while installing Exchange. You can keep the displayed values or enter new values.

  If you change the virtual IP address, you must create a static entry in the DNS server mapping the new virtual IP address to the virtual server name.
- For each system in the cluster, select the public network adapter name. Select the Adapter Name field to view the adapters associated with a node.

---

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

---

- Click **Next**.

8 Review the service group configuration and change the resource names, if desired:

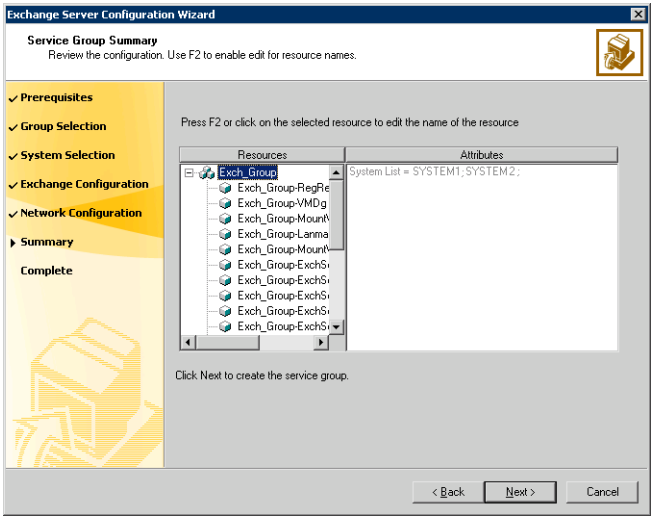


The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

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

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

- Click **Next**.
- A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

9 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system.

10 Click **Finish**.

After bringing the service group online, you must run the Exchange Management Console so that all the stores are automatically mounted on start-up.

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

# Verifying the cluster configuration (secondary site)

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.

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
- Adding a new failover node after DR environment is in operation

# 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 site where the cluster disk group is imported, use the VEA console to launch the Setup Replicated Data Set Wizard. Right-click **Replication Network** on the Management Host configuration tree, and click **Setup Replicated Data Set**.
2. Read the Welcome page and click **Next**.
3. Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

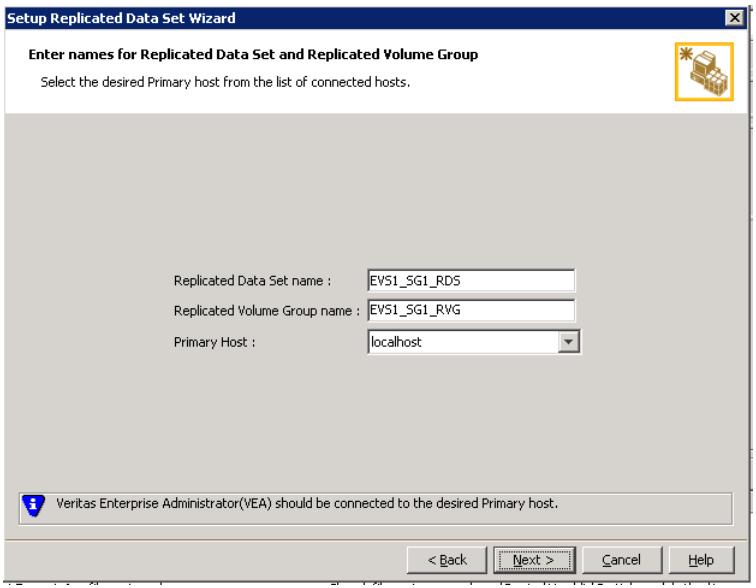


   By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.
   If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.
4. Click **Next**.
5. Select from the table the dynamic disk group and data volumes that will undergo replication.
   To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

6 Click **Next**.

7 Select or create a volume for the Replicator Log:

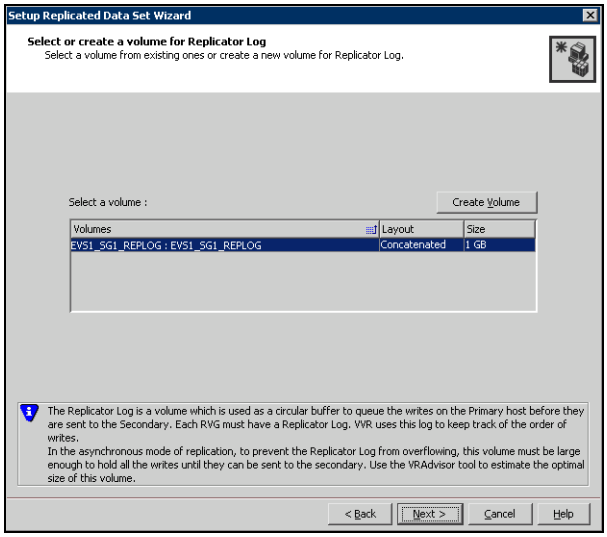


**To select an existing volume**

- Select the volume for the Replicator Log in the table (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** | ■ Choose **Select disks automatically** if you want VVR to select the disks for the Replicator Log.<br>■ Choose **Select disks manually** to use specific disks from the available disks pane for creating the Replicator Log volume.<br>Either double-click the disk to select it, or select **Add** to move the disks into the selected disks pane. |

- Click **OK** to create the Replicator Log volume.
- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.

8 Review the information on the summary page and click **Create Primary RVG**.

9 After the RVG for the primary site is successfully created, click **Yes** to add the secondary host to the RDS for replication.

10 Specify the name of the host where the disk group is imported on the secondary site. If necessary, specify the fully qualified domain name.

11 Click **Next**.

12 If the Veritas Enterprise Administrator console is not already connected to the secondary host, the connection process starts when you click **Next**. Enter valid user credentials, click **OK**, and click **Next** again.

13 The configuration for these volumes on the primary and secondary sites must be identical and meet VVR configuration requirements. If a Replicator Log volume does not exist on the secondary site, it can be created with this procedure:

- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.
  Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.
  When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.
- If all the data volumes to be replicated meet the requirements, this screen does not occur.

**14** If necessary, edit the replication settings for a secondary host.

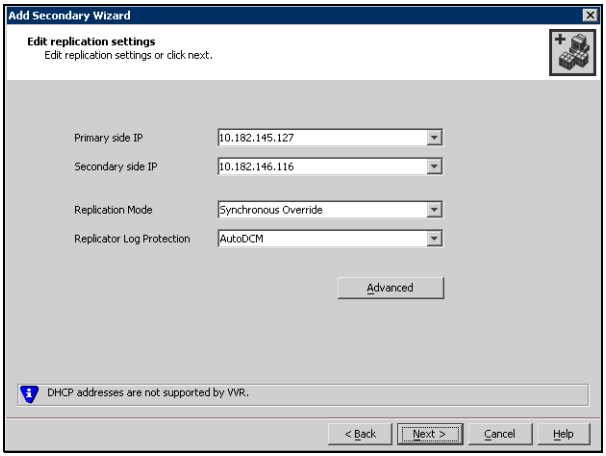


- Enter the virtual IP address for the Primary IP resource that will be used for replication.
- Select or specify an IP address for the Secondary IP resource.
- Specify the replication mode.

| | |
|---|---|
| **Synchronous Override** | Enables Synchronous updates under typical operating conditions. If the secondary site is disconnected from the primary site, and write operations occur on the primary site, the mode of replication temporarily switches to **Asynchronous**. |
| **Synchronous** | Determines updates from the application on the primary site are completed only after the secondary site successfully receives the updates. |
| **Asynchronous** | Determines updates from the application on the primary site are completed after VVR stores the updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously |

- Specify the replicator log overflow protection property.

| | |
|---|---|
| **AutoDCM** | Is the default option and enables the DCM when the Replicator Log overflows even though the secondary site is connected. All data volumes in the primary RVG must have a DCM log to use this option. The wizard attempts to ensure all volumes under the primary RVG have a DCM log. |
| **DCM** | Enables Replicator Log protection for the secondary site. DCM is enabled when the Replicator Log overflows and the secondary site is disconnected from the primary site. All data volumes in the primary RVG must have a DCM log to use this option. The wizard attempts to ensure all volumes under the primary RVG have a DCM log. |
| **Off** | Disables Replicator Log overflow protection. |
| **Override** | Enables log protection. If the secondary site is still connected and the Replicator Log is about to overflow, VVR stalls write operations until five percent or 20 megabytes (whichever is smaller) of the space on the Replicator Log becomes available.<br>If the secondary site becomes inactive because of a connection failure or administrative action, VVR disables Replicator Log protection and causes the Replicator Log to overflow. |
| **Fail** | Enables log protection. When the Replicator Log is about to overflow, VVR stalls write operations until five percent or 20 megabytes (whichever is smaller) of the space on the Replicator Log becomes available. If the connection between the primary RVG and secondary RVG is broken, subsequent write operations to the primary RVG fail. |

15 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.

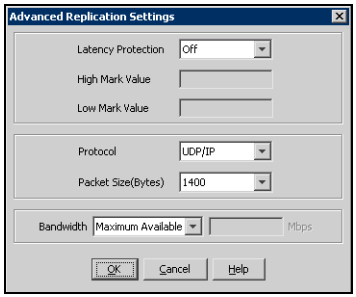


**Latency protection** Determines the extent of stalling write operations on the primary site to allow the secondary site to "catch up" with the updates before new write operations can occur.

- **Off** is the default option and disables latency protection.
- **Fail** enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.
- **Override** enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

**Caution:** Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

**High Mark Value** Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

**Low Mark Value** Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can "catch up" to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

**Caution:** When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

**Protocol** UDP/IP is the default protocol for replication.

**Packet Size** Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

**Bandwidth** By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

16 Click **OK** to close the dialog box.

17 Click **Next**.

18 On the **Start Replication** page, accept the **Synchronize Automatically** option, which is the default recommended for initial setup.

19 Select **Start Replication**, which is the default.
If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.
If the virtual IPs have been created, select **Start Replication** to start synchronization immediately.
If replication must be started later, use the **Start Replication** option of VEA to begin replication. Refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for additional details.

20 Click **Next**.

21 Review the specifications and click **Finish** to add the host on the secondary site to the RDS. Click **Back** to change any information.
Replication physically starts when the IP address is created.

# Creating the VVR RVG service group

Run the wizard from the system that has the Exchange service group online.

The procedure uses EVS1 as an example for all Exchange virtual servers.

Prerequisites:

- Verify that the disk group is imported on the node on which you want to create the Replication Service Group.
- Verify VCS is running, by running the following command on the host on which the you intend to run the Volume Replicator Agent Configuration Wizard.

```
> hasys -state
```

### To create a replication service group

1. From the active node of the cluster at the primary site, click **Start > All Programs > Symantec > Veritas Cluster Server > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
2. 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:

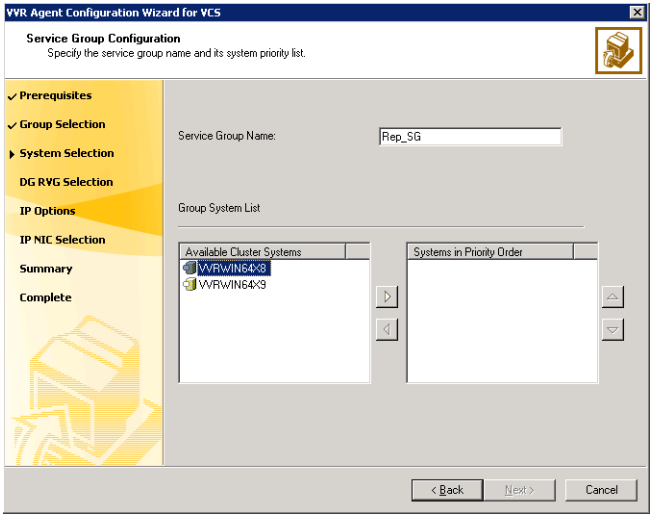


   - Enter the service group name (EVS1_RVG_GRP).
   - In the **Available Cluster Systems** box, click the nodes on which to configure the service group, and click the right-arrow icon to move the nodes to the service group's system list. Make sure that the set of nodes selected for the replication service 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.
- 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 Java Console provides a wizard to create global clusters by linking standalone clusters or bringing a standalone cluster into an existing global environment.

- If you are creating a global cluster environment for the first time with two standalone clusters, run the wizard from either the cluster on the primary site or the cluster on the secondary site.
- If you are adding a standalone cluster to an existing global cluster environment, run the wizard from a cluster already in the global cluster environment.

The following information is required for the Remote Cluster Configuration Wizard in Cluster Explorer:

- The active host name or IP address of each cluster in the global configuration and of the cluster being added to the configuration.
- The user name and password of the administrator for each cluster in the configuration.
- The user name and password of the administrator for the cluster being added to the configuration.

**Note:** Symantec does not support adding a cluster that is already part of a global cluster environment. To merge the clusters of one global cluster environment (for example, cluster A and cluster B) with the clusters of another global environment (for example, cluster C and cluster D), separate cluster C and cluster D into standalone clusters and add them one by one to the environment containing cluster A and cluster B.

### To add a remote cluster in Cluster Explorer

1. From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.
   *or*
   From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.
2. Review the required information for the **Remote Cluster Configuration Wizard** and click **Next**.
3. In the **Wizard Options** panel, click **Add Cluster**, then click **Next**.
4. In the New Cluster Details panel, enter the details of the new cluster as follows:
   If the cluster is not running in secure mode:

- 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.
- Enter the password.
- Click **Next**.

If the cluster is running in secure mode:

- 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 3b.

2 Review the information required for the Global Group Configuration wizard and click **Next**.

3 Enter the details of the service group to modify:

- 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:

| | |
|---|---|
| Cluster not 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.<br>■ Verify the port number.<br>■ Enter the user name.<br>■ Enter the password.<br>■ Click **OK**.<br>■ Repeat these steps for each cluster in the global environment. |
| Cluster 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.<br>■ Verify the port number.<br>■ 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.<br>■ If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.<br>■ Click **OK**.<br>■ Repeat these steps for each cluster in the global environment. |

5 Click **Next**, then click **Finish**.
At this point, you must bring the global service group online from Cluster Explorer.

## Bringing a global service group online

After converting the local service group that is common to the global clusters to a global group, use the Cluster Explorer to bring the global service group online.

**To bring a remote global service group online from Cluster Explorer**

1 In the **Service Groups** tab of the configuration tree, right-click the service group.
*or*
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.

2 Click **Online**, and click **Remote online**.

3 In the Online global group dialog box:
■ Click the remote cluster to bring the group online.

- Click the specific system, or click **Any System**, to bring the group online.
- Click **OK**.

# Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. If the local clusters are running in secure mode, you may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.
- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

### To take the ClusterService-Proc (wac) resource offline on all clusters

1. From Cluster Monitor, log on to a cluster in the global cluster.
2. In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.

3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.

4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

### To add the -secure option to the StartProgram resource

1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.

2 Click **View**, and then **Properties** view.

3 Click the Edit icon to edit the **StartProgram** attribute.

4 In the Edit Attribute dialog box, add `-secure` switch to the path of the executable Scalar Value. For example:

```
C:\Program Files\Veritas\Cluster Server\bin\wac.exe -
secure
```

5 Repeat step 4 for each system in the cluster.

6 Click **OK** to close the Edit Attribute dialog box.

7 Click the **Save and Close Configuration** icon in the tool bar.

8 Repeat step 1 to step 7 for each cluster in the global cluster.

### To establish trust between root brokers if there is more than one root broker

- Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.
  Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:
  **vssat setuptrust --broker *<host:port>* --securitylevel <low|medium|high> [--hashfile *<filename>* | --hash *<root hash in hex>*]**
  For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:
  from RB1, type:
  **vssat setuptrust --broker RB2:14141 --securitylevel low**
  from RB2, type:
  **vssat setuptrust --broker RB1:14141 --securitylevel low**

### To bring the ClusterService-Proc (wac) resource online on all clusters

1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.

2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.

3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

# 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**.

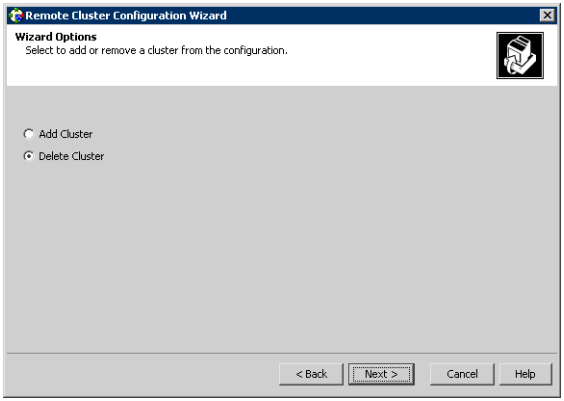


4 In the Delete Cluster panel, click the name of the remote cluster to delete, then click **Next**.

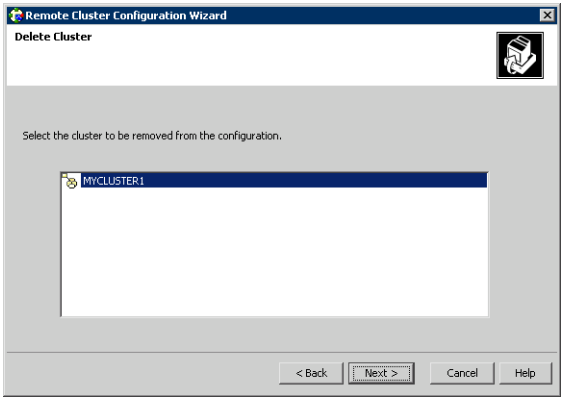

5 Review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:

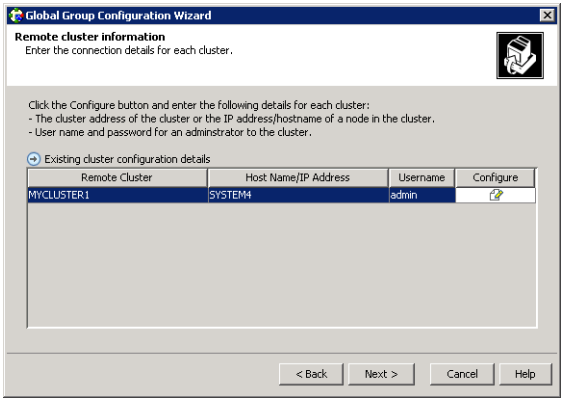


If the cluster is not running in secure mode:

- 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**.

# Adding a new failover node after DR environment is in operation

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 36 for installation instructions.
2. Use the **Cluster Operations** option of the VCS Configuration wizard (**Start > All Programs > Veritas > Veritas Cluster Server > VCS Configuration Wizard**) to add the new system to the cluster. If necessary, refer to the *Veritas Cluster Server Administrator's Guide* for information on this procedure.

## Preparing the existing DR environment

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 52.

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 273; 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** > **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 > Application Agent for Exchange > Configuration Wizard)** to add the new

node to the system list for the Exchange service group. Check the check box to bring the service group online after the wizard completes.
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 Management Console 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 308, 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**.

Section 4

# Appendices

Appendix A

# VCS agent for Exchange Server 2007

This appendix describes the VCS application agent for Exchange Server 2007 and lists the resource type definition and attribute definition of the agent. The resource type represents the VCS configuration definition of the agent and specifies how the agent is defined in the cluster configuration file main.cf. The Attribute Definitions lists the attributes associated with the agent. The Required attributes table lists the attributes that must be configured for the agent to function properly.

## VCS application agent for Microsoft Exchange

The VCS application agent for Microsoft Exchange monitors Exchange services in a VCS cluster, brings them online, and takes them offline.

The VCS application agent for Microsoft Exchange contains the following agent:

- Exchange Service agent—Monitors core Exchange services.

The agent provides high availability for Microsoft Exchange Server 2007 in a VCS cluster.

## Exchange Service agent

The Exchange Service (ExchService2007) agent brings the following Exchange services online, monitors their status, and takes them offline:

- Microsoft Exchange AD Topology service (MSExchangeADTopology):
  This service provides Active Directory topology information to the Exchange services. If this service is stopped, most Exchange services are unable to start.
- Microsoft Exchange System Attendant (MSExchangeSA):
  The Exchange component responsible for monitoring, maintenance, and Active Directory lookup services, and ensuring that operations run smoothly.
- Microsoft Exchange Information Store (MSExchangeIS):
  The Exchange storage used to hold messages in users' mailboxes and in public folders.
- Microsoft Exchange Mail Submission (MSExchangeMailSubmission):
  This service submits messages from the Mailbox Server to the Hub Transport Server.

In addition, you can also configure the agent to monitor the following optional services:

- Microsoft Exchange Mailbox Assistants (MSExchangeMailboxAssistants):
  This service performs background processing of mailboxes in the Exchange store.
- Microsoft Exchange Monitoring (MSExchangeMonitoring):
  This service allows applications to call the Exchange diagnostic cmdlets (pronounced "command-lets").
- Microsoft Exchange Replication Service (MSExchangeRepl):
  This service provides replication functionality for Mailbox Server role databases and is used by Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR).
- Microsoft Exchange Search Indexer (MSExchangeSearch):
  This service performs indexing of mailbox content, which improves the performance of content search.
- Microsoft Exchange Service Host (MSExchangeServiceHost):
  This service provides a host for several Microsoft Exchange services.
- Microsoft Exchange Transport Log Search (MSExchangeTransportLogSearch):

This service provides remote search capability for Microsoft Exchange Transport log files.

- Microsoft Search (msftesql-Exchange):
  This service creates full-text indexes on content and properties of structured and semi-structured data to allow fast linguistic searches on the data.

Each Microsoft Exchange service is configured as a VCS resource of type ExchService2007.

---

**Note:** The agent does not support the Active Directory Connector and the Site Replication Service. Do not run these services on systems that are part of the VCS Exchange cluster.

---

## Agent Operations

- Online—Starts the configured Exchange service.
- Offline—Stops the configured Exchange service.
- Monitor—Determines the state of the configured Exchange service by querying the Service Control Manager (SCM).
  The agent monitors and verifies the state of all the databases that are selected for detail monitoring. The agent behavior varies depending on how the attributes are configured.
  See "Detail monitoring and agent behavior" on page 317 for more information.

## State definition

- Online—Indicates that the configured Exchange service has started.
- Offline—Indicates that the configured Exchange service has stopped.
- Unknown—Indicates that the agent is unable to determine the state of the configured Exchange service.

## Resource type definition

The Exchange Service agent is represented by the ExchService2007 resource type.

```
type ExchService2007 (
    static i18nstr ArgList[] = {Service,
    "LanmanResName:VirtualName", DetailMonitor,
    FaultOnMountFailure, DBList}
    str Service
    str LanmanResName
    int DetailMonitor = 0
    boolean FaultOnMountFailure = 0
    i18nstr DBList[]
)
```

## Attribute definitions

Review the following information to familiarize yourself with the required agent attributes for an ExchService2007 resource type. This information will assist you during the agent configuration.

**Table A-1** Exchange Service agent required attributes

| Required Attributes | Type and Dimension | Definition |
|---|---|---|
| Service | string-scalar | The name of the Exchange service to be monitored. This attribute could take any of the following values:<br>■ MSExchangeADTopology<br>■ MSExchangeIS<br>■ MSExchangeMailSubmission<br>■ MSExchangeSA<br>■ MSExchangeMailboxAssistants<br>■ MSExchangeServiceHost<br>■ MSExchangeTransportLogSearch<br>■ MSExchangeSearch<br>■ msftesql-Exchange<br>■ MSExchangeMonitoring<br>■ MSExchangeRepl |
| LanmanResName | string-scalar | The name of the Lanman resource on which the ExchService2007 resource depends. |

**Table A-2** Exchange Service agent optional attributes

| Optional Attribute | Type and Dimension | Definition |
|---|---|---|
| DetailMonitor | integer-scalar | The interval at which the agent performs detail monitoring on the databases specified in the DBList attribute.<br>The default value 5 indicates that the agent performs detail monitoring on every 5th monitor cycle.<br>Setting this value to 0 disables detail monitoring. |
| FaultOnMountFailure | boolean-scalar | This flag is used to control the agent behavior in case of detail monitoring. It is applicable to Exchange databases that are selected for detail monitoring.<br>If this flag is set to true and a database that is set to mount automatically on startup is dismounted, the agent will fault the service group.<br>The default value is 0 (false). |
| DBList | string-vector | List of databases for which the agent will perform detail monitoring. |

## Detail monitoring and agent behavior

You can configure the VCS agent for Exchange Server 2007 to perform detail monitoring on Exchange databases by specifying the desired databases in the DBList attribute. The frequency at which the agent monitors the database is determined by the Detail Monitor attribute.

If you have selected certain databases but do not want the agent to perform detail monitoring on those databases, you do not have to delete the selected databases from the DBList attribute. You can disable detail monitoring by just setting the value of the Detail Monitor attribute to 0. That way, you do not have to select the databases again.

Table A-3 describes the agent behavior depending on the state of the databases selected for detail monitoring and the FaultonMountFailure attribute settings.

**Table A-3** Detail monitoring and agent behavior

| Exchange database set to mount on startup | Exchange database state | FaultonMountFailure attribute value | Agent state |
|---|---|---|---|
| Yes | Mounted | Does not matter | Online |
| Yes | Dismounted | 1 (True) | Offline<br>(Service group will fault) |
| | | 0 (False) | Unknown<br>(Administrative action required) |
| No | Mounted | Does not matter | Online |
| No | Dismounted | Does not matter | Unknown<br>(Administrative action required) |

You may want to dismount the Exchange databases for performing certain administrative operations. In such cases, to avoid the agent from faulting the service group, you can set the FaultonMountFailure attribute value to 0 (false), and then dismount the database and perform the operations.

Once done, you can again mount the databases, set the FaultonMountFailure attribute to 1 (true) and restore the agent behavior to fault the service group if a database is dismounted.

Appendix B

# Troubleshooting

This chapter describes how to troubleshoot common problems in the VCS application agent for Microsoft Exchange. The chapter lists the error messages, and describes the problem associated with the agent. Recommended solution is included, where applicable.

## VCS logging

VCS generates two error message logs: the engine logs and the agent logs. Log file names are appended by letters. Letter A indicates the first log file, B the second, C the third, and so on.

The agent log is located at %VCS_HOME%\log*agent_A*.txt. The format of agent log messages is:

Timestamp (Year/MM/DD) | Mnemonic | Severity | UMI | Agent Type | Resource Name | Entry Point | Message Text

A typical agent log resembles:

```
2003/12/19 15:09:22 VCS INFO V-16-20024-13
ExchService2007:d1-ExchService2007-MSExchangeIS:online:Service
(MSEXCHANGEIS) is taking longer to start. Timeout = 10 seconds
```

Here,

- Timestamp denotes the date and time when the message was logged.
- Mnemonic denotes which Symantec product logs the message. For VCS application agent for Microsoft Exchange, mnemonic is ‘VCS’.
- Severity denotes the seriousness of the message. Severity of the VCS error messages is classified into the following types:
  - CRITICAL indicates a critical error within a VCS process. Contact Technical Support immediately.
  - ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action.

- WARNING indicates a warning or error, but not an actual fault.
- NOTE informs that VCS has initiated an action.
- INFO informs about various state messages or comments.
  Of these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.

- UMI or Unique Message ID is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the ExchService agent would resemble: `V-16-20024-13`
  Originator ID for all VCS products is 'V-16.' Category ID for ExchService agent is 20024. Message ID is a unique number assigned to the message text.
- Message text denotes the actual message string.

You can view these message logs using Notepad or any text editor. All messages are logged to the engine and the agent logs. Messages of type CRITICAL and ERROR are also written to the Windows event log.

The following table lists the messages of type ERROR and WARNING.

# Exchange Service agent error messages

Table B-4 lists the Exchange Service agent error messages and their descriptions.

**Table B-4** Exchange Service agent error messages

| Message | Description |
|---|---|
| Failed to find the service object. Please check the 'Service' attribute. | The value specified for the “Service” attribute is incorrect.<br>Solution: Provide a valid value for the Lanman resource. If the value is correct, see error type and error code for further information. |
| Failed to open the service object.(Service = *service name*). *Error Type*, *Error Code*. | The agent failed to open the service object.<br>Solution: See the associated Windows error type and error code for more information. |
| Failed to get the state of the service (*service name*). *Error Type*, *Error Code*. | The agent failed to retrieve the state of the service.<br>Solution: See the associated Windows error type and error code for more information. |
| Failed to start the service (*service name*) *Error Type*, *Error Code*. | The agent failed to start the specified service.<br>Solution: See the associated Windows error type and error code for more information. |
| Failed to stop the service (*service name*). *Error Type*, *Error Code*. | The agent failed to stop the service.<br>Solution: See the associated Windows error type and error code for more information. |
| Failed to kill the service (*service name*) *Error Type*, *Error Code*. | The agent failed to terminate the service.<br>Solution: See the associated Windows error type and error code for more information. |
| Configuration error. 'Service' attribute is not configured. | No value specified for the “Service” attribute.<br>Solution: Specify a valid value for the attribute. |
| Configuration error. 'LanmanResName' attribute is not configured. | No value specified for the “LanManResName” attribute.<br>Solution: Specify a valid value for the attribute. |

**Table B-4** Exchange Service agent error messages (continued)

| Message | Description |
|---|---|
| The (*service name*) service is in STARTED state but is not running under the context of Virtual Server (*virtual server name*). | The Exchange service is already running, but not in the context of the virtual server name.<br>Solution: Stop the service and bring the corresponding ExchService2007 resource online. |
| Failed to set the virtual environment for service: (*service name*). *Error Type*, *Error Code*. | The agent failed to set the environment block for the service.<br>The agent needs to set the environment block for starting the service in the context of the virtual server name.<br>Solution: See the associated Windows error type and error code for more information. |
| Failed to remove virtual environment for Service = (*service name*). *Error Type*, *Error Code*. | The agent failed to remove the environment block for the service.<br>While taking the resource offline, the agent stops the service and removes the environment block.<br>Solution: See the associated Windows error type and error code for more information. |
| Configuration error. \"LanmanResName\" attribute is not configured. | No value specified for the “LanmanResName” attribute.<br>Solution: Specify a valid value for the attribute. |
| Configuration error. DetailMonitoringInterval attribute is greater than zero but DBList is empty. No database is specified for detail monitoring. | Detail monitoring for databases is enabled and the monitoring interval (DetailMonitor attribute) is also specified. But there are no databases selected. The DBList attribute is empty.<br>Solution: Select the databases for the detail monitoring. |
| FaultOnMountFailure flag is true.\"Auto Mount\" on database: (*database names*) is enabled but database is dismounted. Agent will return status as offline." | The attribute FaultOnMountFailure is set to True for databases that are set to mount automatically on startup. But these databases are dismounted. So the agent will fault the service group. |

**Table B-4** Exchange Service agent error messages (continued)

| Message | Description |
|---|---|
| \"Auto Mount\" on database: (*database names*) is enabled but database is dismounted. Agent will return status as Unknown. | Databases that are set to mount automatically on startup are dismounted. If these databases are selected for detail monitoring, the agent will return an Unknown status and appropriate administrative action is required. |
| Failed to add computer account to 'Exchange Servers' group *Error Type*, *Error Code*. | Unable to add the computer account to the Exchange Servers group.<br>Solution: Make sure that the user has permissions to add computer accounts to the Exchange Servers group. If the user has those permissions, see the error type and error code for further information. |

# Troubleshooting Microsoft Exchange uninstallation

You might encounter errors while removing Microsoft Exchange if any of the following requirements are not adhered to:

- User mailboxes exist.
- The Exchange Server to be uninstalled has routing group connectors configured.
- Public folder databases exist.

In any of the above scenarios, carry out the following steps to resolve the error.

1. Start the following Exchange services manually using the Service Control Manager:
   - MSExchangeSA
   - MSExchangeIS
2. Move or delete user mailboxes. See the Exchange documentation for instructions.
3. Move or delete public folder. See the Exchange documentation for instructions.
4. Stop all Exchange services started in Step 1.
5. Start the Exchange Setup Wizard for VCS and select the *Remove Exchange* option. Note that you must uninstall Exchange *only* by using the Exchange Setup Wizard for VCS.

# Troubleshooting Exchange Setup Wizard issues

When adding a failover node to an existing Exchange cluster, the Exchange Setup Wizard may fail to rename the node during the pre-installation phase, and report the following error message:

```
Failed to rename the node. Refer to the log file for further
details.
```

This can happen if the Exchange Setup Wizard is unable to delete the Exchange Virtual Server computer object in the Active Directory.

To resolve this issue, you must manually delete the Exchange Virtual Server computer object from the AD, and run the wizard again.

Appendix C

# Sample configurations

The sample configurations in this appendix describe typical service groups configured to monitor the state of the Exchange Server in a VCS cluster.

The appendix lists the sample configuration for clusters using SFW to manage shared storage. See “Active/Passive failover configuration” on page 328 for information.

The sample configuration graphically depicts the resource types, resources, and resource dependencies within the service group. For more information about these resource types, see the chapter VCS Resource Types and Agents in the *Veritas Cluster Server Administrator’s Guide*.

# Active/Passive failover configuration

This section lists the sample configuration for clusters employing Storage Foundation for Windows to manage shared storage. In the sample configuration shown in the dependency graph below, the shared disk group is configured using the Volume Manager Diskgroup (VMDg) agent.

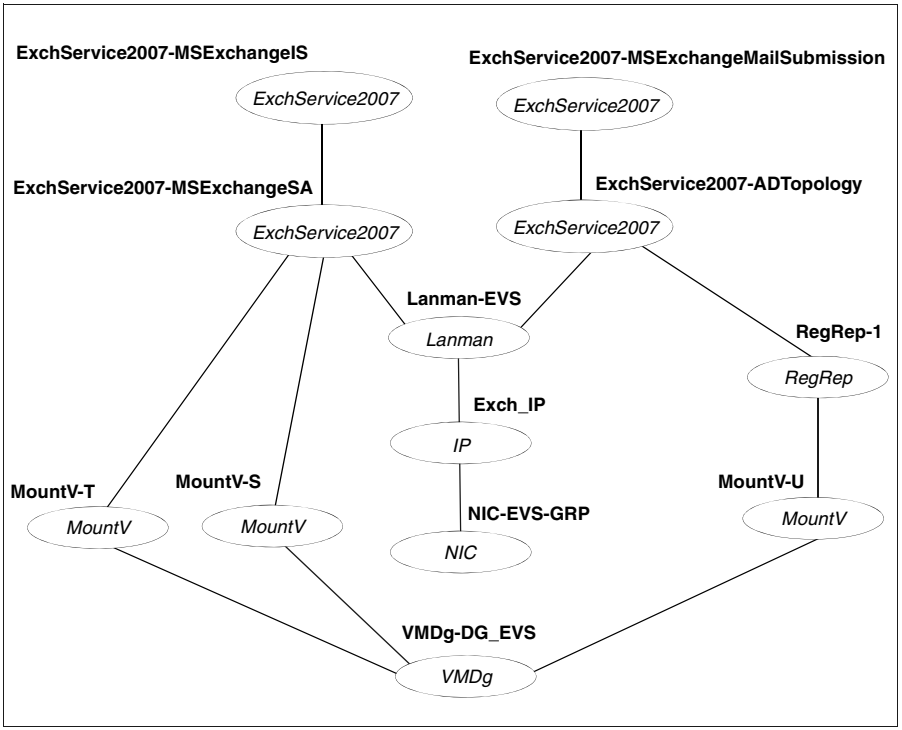


- Three volumes are mounted as drives T, S, and U respectively, using the MountV agent.
  The Registry Replication resource, configured to replicate the Exchange registry keys, is set up on volume U of the cluster disk group. The Exchange database is installed on volume S.
- The Exchange services (MSExchangeSA, MSExchangeIS, MSExchangeADTopology, and MSExchangeMailSubmission) are configured as resources of type ExchService2007.
- The virtual name for the server is created using the Lanman resource.
  The service group virtual IP address for the server is configured using the IP and NIC resource types.

## Sample configuration file

The sample configuration file (main.cf) is included for your reference.

```
include "types.cf"
cluster prim (
    UserNames = { a = cNNh }
    ClusterAddress = "10.217.119.62"
    Administrators = { a }
    )
heartbeat Icmp (
    AYATimeout = 30
    )

system CNODE1 (
    Limits = { ExchLoad = 10 }
    )

system CNODE2 (
    Limits = { ExchLoad = 10 }
    )

group ClusterService (
    SystemList = { CNODE1 = 0, CNODE2 = 1 }
    AutoStartList = { CNODE1, CNODE2 }
    )

    IP csg_ip (
        Address = "10.217.119.62"
        SubNetMask = "255.255.252.0"
        MACAddress @CNODE1 = "4C:00:10:71:B3:FE"
        MACAddress @CNODE2 = "00:0E:A6:C9:47:A6"
        )

    NIC csg_nic (
        MACAddress @CNODE1 = "4C:00:10:71:B3:FE"
        MACAddress @CNODE2 = "00:0E:A6:C9:47:A6"
        )

    Process wac (
        StartProgram @CNODE1 = "C:\\Program Files\\Veritas\\Cluster
Server\\bin\\wac.exe"
        StartProgram @CNODE2 = "C:\\Program Files\\Veritas\\Cluster
Server\\bin\\wac.exe"
        StopProgram @CNODE1 = "C:\\Program Files\\Veritas\\Cluster
Server\\bin\\wacstop.exe"
        StopProgram @CNODE2 = "C:\\Program Files\\Veritas\\Cluster
Server\\bin\\wacstop.exe"
        MonitorProgram @CNODE1 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacmonitor.exe"
```

```
        MonitorProgram @CNODE2 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacmonitor.exe"
        )

    wac requires csg_ip
    csg_ip requires csg_nic

    // resource dependency tree
    //
    //  group ClusterService
    //  {
    //  Process wac
    //      {
    //      IP csg_ip
    //          {
    //          NIC csg_nic
    //          }
    //      }
    //  }

group Group1 (
    SystemList = { CNODE1 = 0, CNODE2 = 1 }
    Prerequisites = { ExchLoad = 10 }
    )

    ExchService2007 Group1-ExchService2007-MSExchangeSA (
        Service = MSExchangeSA
        LanmanResName = Group1-Lanman
        )

    ExchService2007 Group1-ExchService2007-MSExchangeIS (
        Service = MSExchangeIS
        LanmanResName = Group1-Lanman
        )

    ExchService2007 Group1-ExchService2007-MSExchangeMailSubmission
(
        Service = MSExchangeMailSubmission
        LanmanResName = Group1-Lanman
        )

    ExchService2007 Group1-ExchService2007-MSExchangeADTopology (
        Service = MSExchangeADTopology
        LanmanResName = Group1-Lanman
        )

    ExchService2007
Group1-ExchService2007-MSExchangeMailboxAssistants (
        Service = MSExchangeMailboxAssistants
        LanmanResName = Group1-Lanman
        )
```

```
    ExchService2007 Group1-ExchService2007-MSExchangeServiceHost (
        Service = MSExchangeServiceHost
        LanmanResName = Group1-Lanman
        )

    ExchService2007
Group1-ExchService2007-MSExchangeTransportLogSearch (
        Service = MSExchangeTransportLogSearch
        LanmanResName = Group1-Lanman
        )

    ExchService2007 Group1-ExchService2007-MSExchangeSearch (
        Service = MSExchangeSearch
        LanmanResName = Group1-Lanman
        )

    ExchService2007 Group1-ExchService2007-msftesql-Exchange (
        Service = msftesql-Exchange
        LanmanResName = Group1-Lanman
        )

    ExchService2007 Group1-ExchService2007-MSExchangeMonitoring (
        Service = MSExchangeMonitoring
        LanmanResName = Group1-Lanman
        )

    ExchService2007 Group1-ExchService2007-MSExchangeRepl (
        Service = MSExchangeRepl
        LanmanResName = Group1-Lanman
        )

    IP Group1-IP (
        Address = "10.217.119.90"
        SubNetMask = "255.255.252.0"
        MACAddress @CNODE1 = "4C-00-10-71-B3-FE"
        MACAddress @CNODE2 = "00-0E-A6-C9-47-A6"
        )

    Lanman Group1-Lanman (
        VirtualName = EVS1
        IPResName = Group1-IP
        DNSUpdateRequired = 1
        ADUpdateRequired = 1
        ADCriticalForOnline = 1
        )

    MountV Group1-MountV (
        MountPath = "K:"
        VolumeName = REGREP
        VMDGResName = Group1-VMDg
        )
```

```
    MountV Group1-MountV-1 (
        MountPath = "I:"
        VolumeName = DATA
        VMDGResName = Group1-VMDg
        )

    MountV Group1-MountV-2 (
        MountPath = "J:"
        VolumeName = LOG
        VMDGResName = Group1-VMDg
        )

    NIC Group1-NIC (
        MACAddress @CNODE1 = "4C-00-10-71-B3-FE"
        MACAddress @CNODE2 = "00-0E-A6-C9-47-A6"
        )

    RegRep Group1-RegRep (
        MountResName = Group1-MountV
        ReplicationDirectory = "\\VCS\\Private\\RegRep\\Exch"
        Keys = {

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeADTopology" =
"",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeIS" = "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeMailSubmissio
n" = "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeMailboxAssist
ants" = "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeMonitoring" =
"",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeRepl" = "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeSA" = "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeSearch" = "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeServiceHost"
= "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\MSExchangeTransportLogS
earch" = "",

"HKLM\\SYSTEM\\CurrentControlSet\\Services\\msftesql-Exchange" = ""
}
        RestoreLocally = 1
        )
```

```
    VMDg Group1-VMDg (
        DiskGroupName = sacDG1
        DGGuid = f8b39e92-9de1-4baa-b8a3-edb9b93e2100
        )

    Group1-RegRep requires Group1-MountV
    Group1-MountV requires Group1-VMDg
    Group1-Lanman requires Group1-IP
    Group1-MountV-1 requires Group1-VMDg
    Group1-ExchService2007-MSExchangeSA requires Group1-RegRep
    Group1-ExchService2007-MSExchangeSA requires Group1-Lanman
    Group1-ExchService2007-MSExchangeSA requires Group1-MountV-1
    Group1-ExchService2007-MSExchangeSA requires Group1-MountV-2
    Group1-ExchService2007-MSExchangeIS requires
Group1-ExchService2007-MSExchangeSA
    Group1-ExchService2007-MSExchangeMailSubmission requires
Group1-ExchService2007-MSExchangeADTopology
    Group1-ExchService2007-MSExchangeADTopology requires
Group1-RegRep
    Group1-ExchService2007-MSExchangeADTopology requires
Group1-Lanman
    Group1-MountV-2 requires Group1-VMDg
    Group1-ExchService2007-MSExchangeMailboxAssistants requires
Group1-ExchService2007-MSExchangeADTopology
    Group1-ExchService2007-MSExchangeServiceHost requires
Group1-ExchService2007-MSExchangeADTopology
    Group1-ExchService2007-MSExchangeTransportLogSearch requires
Group1-ExchService2007-MSExchangeADTopology
    Group1-ExchService2007-MSExchangeSearch requires
Group1-ExchService2007-msftesql-Exchange
    Group1-ExchService2007-msftesql-Exchange requires
Group1-ExchService2007-MSExchangeADTopology
    Group1-ExchService2007-MSExchangeMonitoring requires
Group1-ExchService2007-MSExchangeADTopology
    Group1-ExchService2007-MSExchangeRepl requires
Group1-ExchService2007-MSExchangeADTopology
    Group1-IP requires Group1-NIC

    // resource dependency tree
    //
    //  group Group1
    //  {
    //  ExchService2007 Group1-ExchService2007-MSExchangeIS
    //      {
    //      ExchService2007 Group1-ExchService2007-MSExchangeSA
    //          {
    //          RegRep Group1-RegRep
    //              {
    //              MountV Group1-MountV
    //                  {
    //                  VMDg Group1-VMDg
```

```
    //                  }
    //              }
    //          Lanman Group1-Lanman
    //              {
    //              IP Group1-IP
    //                  {
    //                  NIC Group1-NIC
    //                  }
    //              }
    //          MountV Group1-MountV-1
    //              {
    //              VMDg Group1-VMDg
    //              }
    //          MountV Group1-MountV-2
    //              {
    //              VMDg Group1-VMDg
    //              }
    //          }
    //      }
    //  ExchService2007
Group1-ExchService2007-MSExchangeMailSubmission
    //      {
    //      ExchService2007
Group1-ExchService2007-MSExchangeADTopology
    //          {
    //          RegRep Group1-RegRep
    //              {
    //              MountV Group1-MountV
    //                  {
    //                  VMDg Group1-VMDg
    //                  }
    //              }
    //          Lanman Group1-Lanman
    //              {
    //              IP Group1-IP
    //                  {
    //                  NIC Group1-NIC
    //                  }
    //              }
    //          }
    //      }
    //  ExchService2007
Group1-ExchService2007-MSExchangeMailboxAssistants
    //      {
    //      ExchService2007
Group1-ExchService2007-MSExchangeADTopology
    //          {
    //          RegRep Group1-RegRep
    //              {
    //              MountV Group1-MountV
    //                  {
    //                  VMDg Group1-VMDg
```

```
    //                  }
    //              }
    //          Lanman Group1-Lanman
    //              {
    //              IP Group1-IP
    //                  {
    //                  NIC Group1-NIC
    //                  }
    //              }
    //          }
    //      }
    //  ExchService2007 Group1-ExchService2007-MSExchangeServiceHost
    //      {
    //      ExchService2007
Group1-ExchService2007-MSExchangeADTopology
    //          {
    //          RegRep Group1-RegRep
    //              {
    //              MountV Group1-MountV
    //                  {
    //                  VMDg Group1-VMDg
    //                  }
    //              }
    //          Lanman Group1-Lanman
    //              {
    //              IP Group1-IP
    //                  {
    //                  NIC Group1-NIC
    //                  }
    //              }
    //          }
    //      }
    //  ExchService2007
Group1-ExchService2007-MSExchangeTransportLogSearch
    //      {
    //      ExchService2007
Group1-ExchService2007-MSExchangeADTopology
    //          {
    //          RegRep Group1-RegRep
    //              {
    //              MountV Group1-MountV
    //                  {
    //                  VMDg Group1-VMDg
    //                  }
    //              }
    //          Lanman Group1-Lanman
    //              {
    //              IP Group1-IP
    //                  {
    //                  NIC Group1-NIC
    //                  }
    //              }
```

```
    //          }
    //      }
    //  ExchService2007 Group1-ExchService2007-MSExchangeSearch
    //      {
    //      ExchService2007 Group1-ExchService2007-msftesql-Exchange
    //          {
    //          ExchService2007
Group1-ExchService2007-MSExchangeADTopology
    //              {
    //              RegRep Group1-RegRep
    //                  {
    //                  MountV Group1-MountV
    //                      {
    //                      VMDg Group1-VMDg
    //                      }
    //                  }
    //              Lanman Group1-Lanman
    //                  {
    //                  IP Group1-IP
    //                      {
    //                      NIC Group1-NIC
    //                      }
    //                  }
    //              }
    //          }
    //      }
    //  ExchService2007 Group1-ExchService2007-MSExchangeMonitoring
    //      {
    //      ExchService2007
Group1-ExchService2007-MSExchangeADTopology
    //          {
    //          RegRep Group1-RegRep
    //              {
    //              MountV Group1-MountV
    //                  {
    //                  VMDg Group1-VMDg
    //                  }
    //              }
    //          Lanman Group1-Lanman
    //              {
    //              IP Group1-IP
    //                  {
    //                  NIC Group1-NIC
    //                  }
    //              }
    //          }
    //      }
    //  ExchService2007 Group1-ExchService2007-MSExchangeRepl
    //      {
    //      ExchService2007
Group1-ExchService2007-MSExchangeADTopology
    //          {
```

```
//          RegRep Group1-RegRep
//              {
//              MountV Group1-MountV
//                  {
//                  VMDg Group1-VMDg
//                  }
//              }
//          Lanman Group1-Lanman
//              {
//              IP Group1-IP
//                  {
//                  NIC Group1-NIC
//                  }
//              }
//          }
//      }
//  }
```

# Index

## E

## T



## V