

Cluster Server Implementation Guide for Microsoft Exchange 2007

Windows

7.0

Cluster Server Implementation Guide for Microsoft Exchange 2007

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

Product version: 7.0

Document version: 7.0 Rev 0

Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, the Veritas Logo, InfoScale, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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

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

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

Symantec Corporation

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

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apj@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

Contents

Technical Support	4	
Chapter 1	Introducing the VCS agents for Exchange and NetApp	11
	About the VCS agents for Exchange and NetApp	11
	About VCS support for Exchange Server 2007	12
	About the VCS application agent for Microsoft Exchange	13
	About the Exchange Service agent	13
	About the VCS hardware replication agent for NetApp	15
	About the NetApp Filer agent	15
	About the NetApp SnapDrive agent	16
	About the NetApp SnapMirror agent	18
	How the agents make Microsoft Exchange highly available	22
	Local cluster configuration	22
	Disaster recovery configuration	23
	Typical Exchange configurations in a VCS cluster	23
	Active-Passive failover configuration	23
	Any-to-Any failover configuration	24
	Disaster recovery configuration	25
Chapter 2	Installing the product and configuring a VCS cluster	26
	About installing the VCS agents	26
	Configuring the cluster using the Cluster Configuration Wizard	26
	Configuring notification	35
	Configuring Wide-Area Connector process for global clusters	38
Chapter 3	Installing Microsoft Exchange	41
	About installing Exchange 2007 in a VCS environment	41
	Prerequisites for installing Exchange in a VCS environment	42
	Privileges for installing Exchange	43
	Verifying DNS settings for Exchange hosts	44
	Configuring Microsoft iSCSI initiator	44

Managing storage using NetApp filer	45
Connecting virtual disks to the cluster node	46
Disconnecting virtual disks from the cluster nodes	47
Managing storage using Windows Logical Disk Manager	47
Reserving disks (if you use Windows LDM)	48
Creating volumes (if you use Windows LDM)	49
Mounting volumes (if you use Windows LDM)	49
Unassigning a drive letter	50
Releasing disks (if you use Windows LDM)	51
Installing Exchange 2007 on the first node	51
Performing Exchange pre-installation: first node	52
Performing Exchange Server installation: first node	53
Performing Exchange post-installation: first node	54
Moving Exchange databases to shared storage	55
Installing Exchange 2007 on additional nodes	59
Performing Exchange pre-installation: additional nodes	60
Performing Exchange Server installation: additional nodes	62
Performing Exchange post-installation: additional nodes	63
Chapter 4	
Configuring the Exchange service group	65
About configuring the Exchange service group	65
Prerequisites for configuring the Exchange service group	66
Configuring the Exchange service group using the wizard	67
Running SnapManager for Exchange	74
About the modifications required for tagged VLAN or teamed network	74
About verifying the service group configuration	75
Bringing the service group online	75
Taking the service group offline	76
Switching the service group	76
About modifying the Exchange service group configuration	77
Prerequisites for modifying an Exchange service group	77
Modifying the Exchange service group	77
Deleting the Exchange service group	78
Chapter 5	
Making a standalone Exchange server highly available	79
About configuring a standalone Exchange Server in a VCS environment	79
Renaming the existing Exchange Server and assigning a name to the physical node	81

Chapter 6	Configuring any-to-any failover	84
	About any-to-any configuration	84
	Specifying a common node for failover	86
Chapter 7	Deploying agents for disaster recovery	88
	About disaster recovery configuration	88
	Setting up disaster recovery configuration	88
	Configure replication using NetApp SnapMirror	91
	Configure NetAppSnapMirror resources at the primary site	92
	About managing failover in a disaster recovery environment	93
	Managing a successful remote failover	93
	Managing failover in response to a network outage	93
	About switching the service group back to the local cluster	94
Chapter 8	Removing the software	96
	About removing the software components	96
	Remove Microsoft Exchange	97
	Removing a node without removing Microsoft Exchange	97
	Removing a node and removing Microsoft Exchange	98
Chapter 9	Troubleshooting the agents	100
	About troubleshooting VCS agents	100
	VCS logging	100
	VCS Cluster Configuration Wizard (VCW) logs	102
	VCWsilent logs	102
	NetApp agents error messages	103
	Exchange Service agent error messages	104
	Troubleshooting Microsoft Exchange uninstallation	106
	About troubleshooting Exchange Setup Wizard issues	107
	Exchange Setup Wizard may fail to rename the node	107
	Exchange Setup Wizard may fail while moving the Exchange databases	108
	Exchange database status appears as unknown	108
Appendix A	Resource type definitions	109
	About resource type definitions	109
	Exchange Service agent	109
	Exchange Service agent resource type definition	109
	Exchange Service agent attribute definitions	110

Dependency graph for an Exchange local cluster
 configuration 111

Dependency graph for an Exchange disaster recovery
 configuration 112

Detail monitoring and agent behavior 113

Appendix B **Sample configuration** 115

About Exchange sample configurations 115

Sample configuration for an Exchange cluster (local cluster
 configuration) 115

Introducing the VCS agents for Exchange and NetApp

This chapter includes the following topics:

- [About the VCS agents for Exchange and NetApp](#)
- [About VCS support for Exchange Server 2007](#)
- [About the VCS application agent for Microsoft Exchange](#)
- [About the VCS hardware replication agent for NetApp](#)
- [How the agents make Microsoft Exchange highly available](#)
- [Typical Exchange configurations in a VCS cluster](#)

About the VCS agents for Exchange and NetApp

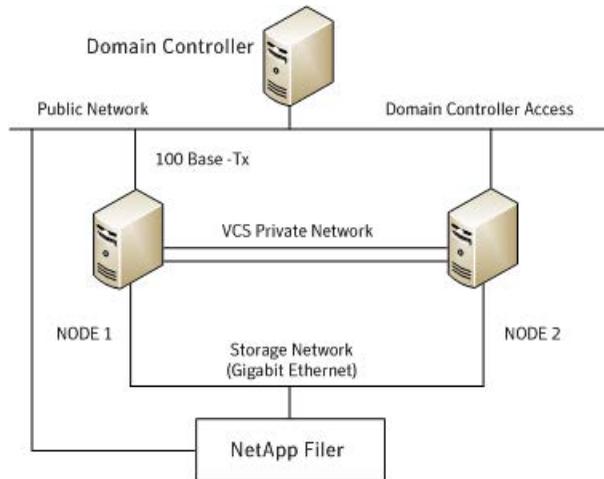
The Cluster Server (VCS) application agent for Exchange provides high availability to Microsoft Exchange Server. The VCS application agent for NetApp SnapMirror enables configuring NetApp filers over an iSCSI or Fibre Channel (FC) connection in a VCS cluster environment. Both agents work together to provide high availability and disaster recovery to Exchange Server in environments using NetApp filers for shared storage. The agents also support disaster recovery configurations set up using the VCS Global Cluster Option and NetApp SnapMirror for data replication.

In a typical configuration, the agents are installed on each node in the cluster. The nodes are connected to the NetApp filers through a dedicated (private) storage

network. VCS nodes are physically attached to the NetApp filer via an ethernet cable supporting iSCSI or Fibre Channel (FC) as the transport protocol.

The following figure illustrates a typical VCS cluster configuration in a NetApp storage environment.

Figure 1-1 Typical VCS configuration in a NetApp storage environment



For more information about the agents refer to their resource type definitions and attribute definitions.

See [“About resource type definitions”](#) on page 109.

About VCS support for Exchange Server 2007

VCS support for Exchange Server 2007 includes the following features:

- High availability 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 VCS 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. VCS provides a utility, the VCS Application Manager (VAM), that you can use to

launch the Exchange Management Shell under the context of the virtual server name.

On the SCC, under Tools, click VCS Application Manager to start VAM. Then double-click the Exchange resource to launch the Exchange Management Shell in the virtual server context.

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.

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 VCS 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.

About the VCS application agent for Microsoft Exchange

The VCS application agent for Microsoft Exchange contains the Exchange Service agent that monitors Exchange services in a VCS cluster, brings them online, and takes them offline. The agent provides high availability for Exchange Server 2007 in a VCS cluster.

Note: High availability support for Microsoft Exchange Server 2007 is available for the Mailbox Server role only.

See [“About the Exchange Service agent”](#) on page 13.

About the Exchange Service agent

The Exchange Service agent brings the Exchange services online, monitors their status, and takes them offline.

Each Exchange Server service is configured as a VCS resource of type ExchService2007.

The Exchange services are as follows:

- 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 Mail Submission (MSExchangeMailSubmission)

This service submits messages from the Mailbox Server to the Hub Transport Server.

- Microsoft Exchange System Attendant (MSEExchangeSA)
The Exchange component responsible for monitoring, maintenance and Active Directory lookup services, and ensuring that operations run smoothly.

In addition, you can also configure the agent to monitor the following optional services:

- Microsoft Exchange Mailbox Assistants (MSEExchangeMailboxAssistants)
This service performs background processing of mailboxes in the Exchange store.
- Microsoft Exchange Monitoring (MSEExchangeMonitoring)
This service allows applications to call the Exchange diagnostic cmdlets (pronounced "command-lets").
- Microsoft Exchange Replication Service (MSEExchangeRepl)
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 (MSEExchangeSearch)
This service performs indexing of mailbox content, which improves the performance of content search.
- Microsoft Exchange Service Host (MSEExchangeServiceHost)
This service provides a host for several Microsoft Exchange services.
- Microsoft Exchange Transport Log Search (MSEExchangeTransportLogSearch)
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.

Exchange Service agent functions

The Exchange Service agent functions are as follows:

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 113.

Exchange Service agent state definitions

The Exchange Service agent state definitions are as follows:

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.

About the VCS hardware replication agent for NetApp

The VCS hardware replication agent for NetApp provides failover support and recovery in environments employing NetApp filers for storage and NetApp SnapMirror for replication.

The agent monitors and manages the state of replicated filer devices and ensures that at a time only one system has safe and exclusive access to the configured devices.

The agent can be used in local clusters, single VCS replicated data clusters, and multi-cluster environments set up using the VCS Global Cluster Option (GCO).

The VCS agents for NetApp are as follows:

- NetAppFiler agent
- NetAppSnapDrive agent
- NetAppSnapMirror agent

About the NetApp Filer agent

The NetApp Filer agent monitors the state of the filer device. The agent is represented by the NetAppFiler resource type in VCS. NetAppFiler resources are persistent, meaning that they are not brought online or taken offline.

NetApp Filer agent function

The NetApp Filer agent function is as follows:

Monitor	<p>Performs the following tasks:</p> <ul style="list-style-type: none"> ■ Verifies the state of the filer attached to the host by sending an ICMP ping command to the filer. If the filer does not respond, the agent reports the state of the filer as faulted. ■ Opens a filer connection and checks if ONTAPI version is supported by the filer. If the connection fails or the ONTAPI version is not supported, the agent reports the state as offline.
---------	---

The NetApp Filer agent resource type definition and attribute definitions are as follows. This information will assist you during the agent configuration.

NetAppFiler agent resource type definition

The NetApp Filer agent is configured as a resource of type NetAppFiler.

```
type NetAppFiler (
    static int MonitorInterval = 30
    static i18nstr ArgList[] = { FilerName, StorageIP }
    static str Operations = None
    str FilerName
    str StorageIP
)
```

NetAppFiler agent attribute definitions

The following table describes the NetApp Filer agent attributes.

Attribute	Description
FilerName	<p>DNS-resolvable name or IP address of the locally attached filer.</p> <p>Type and dimension: string-scalar</p>
StorageIP	<p>The private storage IP address of the filer.</p> <p>Type and dimension: string-scalar</p>

About the NetApp SnapDrive agent

The NetApp SnapDrive agent monitors, connects, and disconnects filer volumes. You can configure the agent to use the iSCSI or the FC protocol.

NetApp SnapDrive agent functions

The NetApp SnapDrive agent functions are as follows:

Online	Connects a virtual disk (LUN) using an iSCSI or an FC initiator. The agent presents the LUN as a locally-attached drive to the host. The agent also removes LUN-host mappings made before the online operation.
Offline	Disconnects the virtual disk (LUN) from the host.
Monitor	Verifies that the specified virtual disk (LUN) is connected to the host.
Open	Verifies that there is connectivity to the filer. It also checks that the VCS Helper service is running with the same privileges as the SnapDrive service.
Clean	Attempts to forcibly disconnect a virtual disk (LUN).

NetApp SnapDrive agent resource type definition and attribute definitions are as follows. This information will assist you during the agent configuration.

NetAppSnapDrive agent resource type definition

NetApp SnapDrive agent is configured as a resource of type NetAppSnapDrive.

```
type NetAppSnapDrive (
    static int MonitorInterval = 30
    static int NumThreads = 1
    static i18nstr ArgList[] = { FilerResName,
        "FilerResName:FilerName", "FilerResName:StorageIP",
        VolumeName, ShareName, LUN, MountPath,
        Initiator, InitiatorMonitorInterval }
    str FilerResName
    str VolumeName
    str ShareName
    str LUN
    str MountPath
    str Initiator[]
    int InitiatorMonitorInterval = 30
)
```

NetAppSnapDrive agent attribute definitions

The following table describes the NetApp SnapDrive agent attributes.

Attribute	Description
FilerResName	Name of the VCS NetAppFiler-type resource in the service group. Type and dimension: string-scalar
VolumeName	Name of the volume containing the virtual disk. Define the volume name in the same case as on the filer. Type and dimension: string-scalar
ShareName	Name of the CIFS share containing the virtual disk. Type and dimension: string-scalar
LUN	Name of the LUN (virtual disk) on the filer that is presented to the host for mounting. Define the LUN name in the same case as on the filer. Type and dimension: string-scalar
MountPath	Drive letter to be assigned to the virtual disk. Type and dimension: string-scalar
Initiator	Name of iSCSI or FC initiator the host uses to connect virtual disks. You can retrieve this value from the Disk Management console. Type and dimension: string-vector

About the NetApp SnapMirror agent

The NetApp SnapMirror agent monitors the replication state of filer devices. When a failover occurs, the agent reverses the direction of replication. The agent supports the replication modes supported by NetApp. The agent supports asynchronous, semi-synchronous, and synchronous modes of replication. You can set the mode of replication using the SyncMode agent attribute.

NetApp SnapMirror agent functions

The NetApp SnapMirror agent functions are as follows:

Online	<p>If the state of the local filer device is SOURCE, the agent creates a lock file to indicate that the resource can come online. This effectively makes the devices writable for the application.</p> <p>If the state of the local filer is SNAPMIRRORED, the agent attempts to reverse the direction of replication by changing the state of the local filer to SOURCE and that of the original source to SNAPMIRRORED.</p> <p>If the original source filer is down, the agent performs a mirror breakoff to enable local write access, if the filer is not already broken off.</p> <p>If the original source returns to life, you must resynchronize the data manually. The online function touches a lock file if read-write access is enabled successfully.</p>
Offline	<p>Removes the lock file. The agent does not perform any filer operations because an offline entry point does not necessarily indicate an intention to give up the devices.</p>
Monitor	<p>Verifies that the lock file exists. If the lock file exists, the monitor function reports the status of the resource as online. If the lock file does not exist, the monitor function reports the status of the resource as offline.</p>
Open	<p>Removes the lock file thereby preventing potential concurrency violation if the group fails over to another node.</p> <p>Note: The agent does not remove the lock file if the agent is started after an <code>hastop -force</code> command.</p>
Clean	<p>Removes the lock file. No filer operations are performed as taking the resource offline does not indicate a pending role swap.</p>

Action function

Use the Action function to perform predefined actions on a resource. To perform an action on a resource, type the following command:

```
hares -action <SnapMirror_resname> <token> [-actionargs <arg1> ...]
[-sys <system>] [-clus <cluster> ]
```

The following table lists the action supported by the NetAppSnapMirror agent.

Table 1-1 Actions supported by NetAppSnapMirror agent

Token for Action	Description
fbsync	<p>Resynchronises an original source volume with a broken-off volume. After synchronization, the original source volume becomes the target volume.</p> <p>The broken-off volume was initially the target volume, but was broken off as a result of a take over.</p>

To synchronize volumes, type the following at the command prompt:

```
hares -action SnapMirror_resname fbsync -sys node_name
```

Where, *SnapMirror_resname* represents the name of the SnapMirror resource and *node_name* represents the node on which the service group is online.

Run the action for each SnapMirror resource.

You can also add custom actions for the agents. Refer to the *Cluster Server Agent Developer's Guide* for more information.

NetApp SnapMirror agent resource type definition and attribute definitions are as follows. This information will assist you during the agent configuration.

NetAppSnapMirror agent resource type definition

NetApp SnapMirror agent is configured as a resource of type NetAppSnapMirror.

```
type NetAppSnapMirror (
    static keylist SupportedActions = { fbsync }
    static int MonitorInterval = 300
    static int NumThreads = 1
    static i18nstr ArgList[] = { FilerResName,
        "FilerResName:FilerName", "FilerResName:StorageIP",
        VolumeName, SnapMirrorArguments, SnapMirrorSchedule,
        AppResName, VisibilityFrequency, SyncMode }
    str FilerResName
    str VolumeName
    str SnapMirrorArguments
    str SnapMirrorSchedule
    str AppResName
    int VisibilityFrequency = 180
    str SyncMode = async
)
```

NetAppSnapMirror agent attribute definitions

The following table describes the NetApp SnapMirror agent attributes.

Attribute	Description
FilerResName	<p>Name of the VCS NetAppFiler-type resource in the group.</p> <p>Type and dimension: string-scalar</p>
VolumeName	<p>Name of the filer volume containing the virtual disk. This is the volume that is to be mounted. Define the volume name in the same case as on the filer.</p> <p>Type and dimension: string-scalar</p>
SnapMirrorArguments	<p>Specifies the SnapMirror arguments such as maximum transfer speed and restart mode.</p> <p>The format for this attribute is:</p> <p><RestartMode> MaxSpeed</p> <p>MaxSpeed is an optional parameter. The RestartMode parameter can have the following values:</p> <ul style="list-style-type: none"> ■ Always ■ Never ■ Default <p>Setting the RestartMode parameter value to default indicates schedule priority.</p> <p>Default value of RestartMode parameter is 'Default'</p> <p>Example: Always MaxSpeed</p> <p>Type and dimension: string-scalar</p>
SnapMirrorSchedule	<p>Specifies the schedule the destination uses for updating data. Do not assign a value for this attribute if you use SnapManager.</p> <p>The schedule is in the following format:</p> <p>minute hour dayofmonth dayofweek</p> <p>Each field is separated by a space.</p> <p>Refer to the NetApp documentation for more details on the rules for each of these schedule fields.</p> <p>By default, this attribute does not have any value.</p> <p>Type and dimension: string-scalar</p>

Attribute	Description
AppResName	<p>Name of the resource configured to monitor the application being made highly available.</p> <p>Type and dimension: string-scalar</p>
SyncMode	<p>Specifies the mode of replication for the mirror.</p> <p>This attribute can have the following values:</p> <ul style="list-style-type: none"> ■ <code>async</code>: Indicates that the mirror should be configured in the asynchronous mode. ■ <code>semi-sync</code>: Indicates that the mirror should be configured in the semi-synchronous mode. ■ <code>sync</code>: Indicates that the mirror should be configured in the synchronous mode. <p>The default is <code>async</code> (asynchronous) mode.</p> <p>Type and dimension: string-scalar</p>
VisibilityFrequency	<p>Specifies how often the source snapshot will be visible on the destination mirror. It controls the value of <code>visibility_interval</code> in the <code>snapmirror.conf</code> file.</p> <p>The default value is 180 seconds.</p> <p>This attribute is applicable only if the mirror is configured in synchronous or semi-synchronous mode.</p> <p>Type and dimension: string-scalar</p>

How the agents make Microsoft Exchange highly available

The VCS application agent for Microsoft Exchange detects an application failure if a configured Exchange service is not running or if a configured virtual server is not available. The NetApp agents ensure consistent data access to the node on which Exchange Server is running.

This section describes how the agents migrate Exchange Server to another node in local clusters and in global disaster recovery environments.

Local cluster configuration

When the Exchange agent detects an application or host failure, VCS attempts to fail over the Exchange service group to the next available system in the service group's SystemList.

The NetApp agents connects the virtual disks (LUNs) containing Exchange data to the new node.

The configured Exchange services and virtual servers are started on the new node, thus ensuring continuous availability for Exchange data, including configured mailboxes.

Disaster recovery configuration

In a disaster recovery configuration, VCS first attempts to fail over the Exchange service group to a node in the local cluster. If all nodes in the local cluster are unavailable, or if a disaster strikes the site, VCS attempts to fail over the Exchange service group to the remote site.

This involves the following steps:

- Connecting the virtual disks (LUNs) to the target hosts (using the NetAppSnapDrive agent)
- Performing a mirror break, which enables write access to the target (using the NetAppSnapMirror agent)
- Reversing the direction of replication by demoting the original source to a target, and begin replicating from the new source (using the NetAppSnapMirror agent)
- Starting the Exchange services on the remote node (using the VCS agents for Exchange Server)
See [“About managing failover in a disaster recovery environment”](#) on page 93.

Typical Exchange configurations in a VCS cluster

The VCS application agent for Microsoft Exchange supports the Active-Passive and the Any-to-Any configurations. It also supports the Disaster Recovery configuration.

Active-Passive failover configuration

An Active-Passive setup involves one to one failover capabilities. For example, if you have two nodes (SYSTEM1 and SYSTEM2), SYSTEM1 can fail over to SYSTEM2.

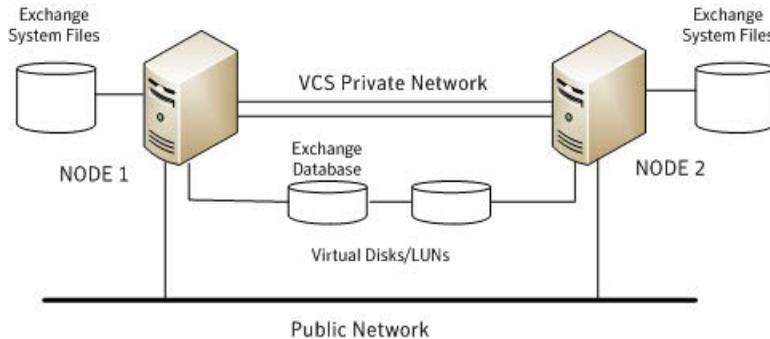
In an Active-Passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

In a typical two-node configuration, Microsoft Exchange and VCS application agent for Microsoft Exchange are installed on both nodes. The Exchange database is on

shared storage. The shared storage can be managed using Windows Logical Disk Management or the NetApp suite of products.

The following figure illustrates an Active-Passive configuration.

Figure 1-2 Active-Passive fail over configuration

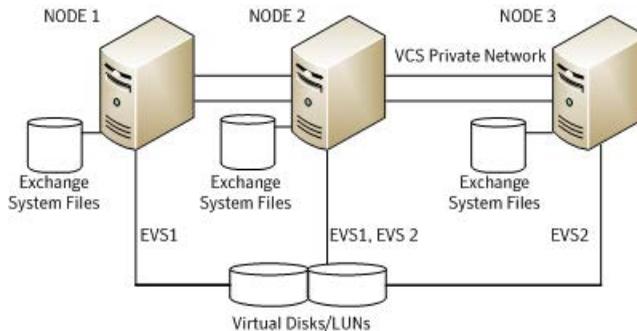


Any-to-Any failover configuration

In an Any-to-Any configuration, each Exchange virtual server in the cluster can be 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. In other words, you must ensure that an Exchange service group does not fail over to a node on which another Exchange service group is online.

The following figure illustrates an Any-to-Any configuration.

Figure 1-3 Any-to-Any fail over 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 nodes 1 and 2 host the EVS1 service group and nodes 2 and 3 host the

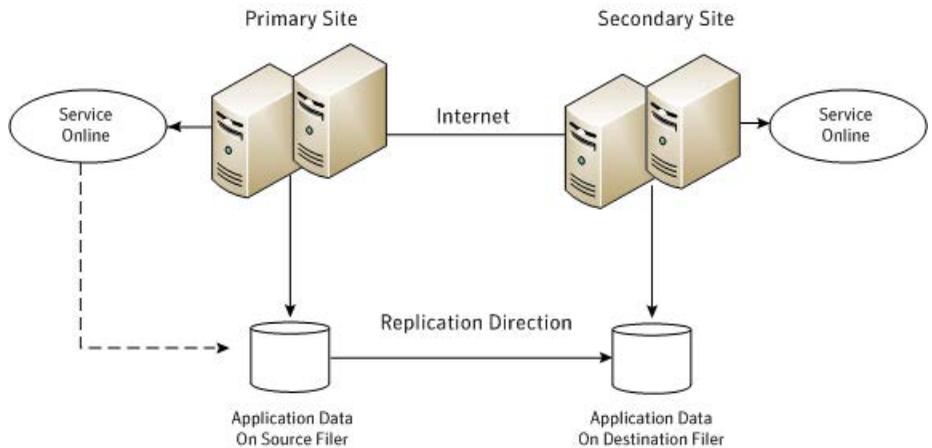
EVS2 service group. If node 1 (or node 3) fails, the service group containing the EVS resources is failed over to node 2.

Disaster recovery configuration

A Disaster Recovery (DR) configuration enables you to restore application data and services in the event of a catastrophic failure. A typical DR solution requires primary and secondary sites, and clusters within those sites. The cluster at the primary site provides data and services during normal operation, and the cluster at the secondary site provides data and services if the primary site fails.

The following figure illustrates a Disaster Recovery configuration.

Figure 1-4 Disaster Recovery configuration



The illustration displays a disaster recovery configuration in a NetApp storage environment. In this case, the primary site is replicating its application data to the secondary site.

When a failure occurs, such as 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.

Installing the product and configuring a VCS cluster

This chapter includes the following topics:

- [About installing the VCS agents](#)
- [Configuring the cluster using the Cluster Configuration Wizard](#)

About installing the VCS agents

Install Veritas InfoScale Availability on all the systems where you want to configure the application. During installation, the product installer installs the VCS agents required for making the applications highly available.

Refer to the *Veritas InfoScale Installation and Upgrade Guide* for instructions.

For the latest information about supported operating systems and software, see the Software Compatibility List at:

<http://www.veritas.com/docs/000025350>

Configuring the cluster using the Cluster Configuration Wizard

After installing the software, set up the components required to run Cluster Server. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, the user account for the VCS Helper service, and provides an option for configuring the VCS Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for notification and global clusters (GCO). You can also use VCW to modify or delete cluster configurations.

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

Note the following prerequisites before you proceed:

- The required network adapters, and SCSI controllers are installed and connected to 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 auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Symantec recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.
- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Symantec recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Symantec recommends that you disable TCP/IP from private NICs to lower system overhead.

Note: If you wish to use Windows NIC teaming, you must select the Static Teaming mode. Only the Static Teaming mode is currently supported.

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.
- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.

- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.
- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.
Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.
- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

- 1 Start the VCS Cluster Configuration Wizard from **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 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, do the following:

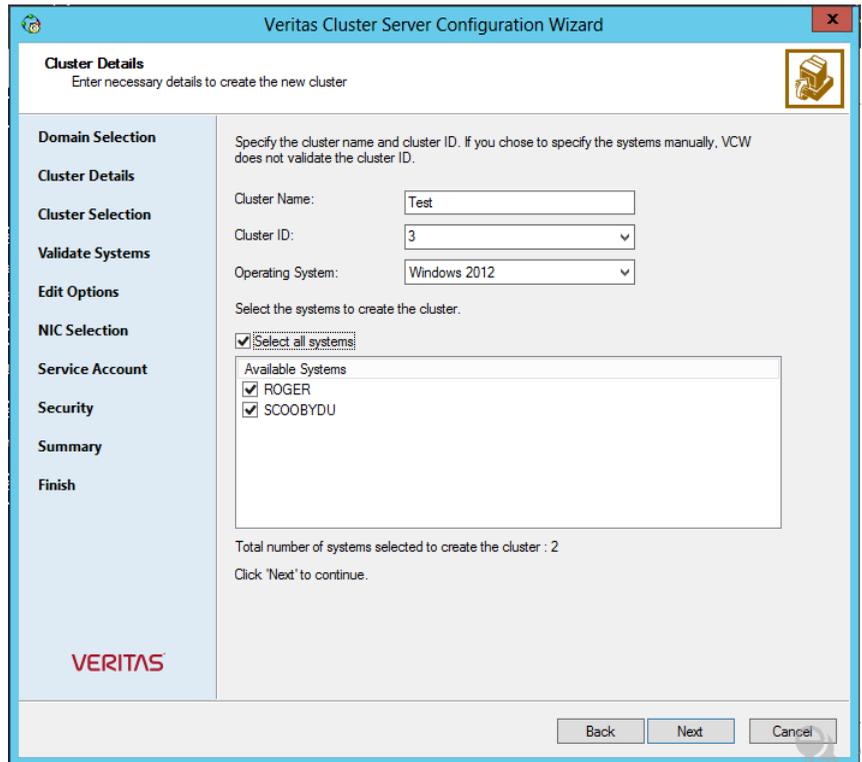
- Clear **Specify systems and users manually**.
- Click **Next**.
Proceed to step 8.

To specify systems and user names manually (recommended for large domains), do the following:

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

- Click **Next**.
If you chose to retrieve the list of systems, proceed to step 6. Otherwise, proceed to the next step.
- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**.
Do not specify systems that are part of another cluster.
Proceed to step 8.
- 6 On the System Selection panel, specify the systems for the cluster and then click **Next**.
Do not select systems that are part of another cluster.
Enter the name of the system and click **Add** to add the system to the Selected Systems list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.
- 7 The System Report panel displays the validation status, whether Accepted or Rejected, of all the systems you specified earlier. Review the status and then click **Next**.
Select the system to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.
A system can be rejected for any of the following reasons:
 - System is not pingable.
 - WMI access is disabled on the system.
 - Wizard is unable to retrieve the system architecture or operating system.
 - Product is either not installed or there is a version mismatch.
- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and then click **Next**.

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



Specify the cluster details as follows:

- | | |
|--------------|--|
| Cluster Name | Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 65535. |
- Note:** 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.

All the systems in the cluster must have the same operating system and architecture. For example, you cannot configure a Windows Server 2008 R2 system and a Windows Server 2012 system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster.

Check the **Select all systems** check box to select all the systems simultaneously.

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

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

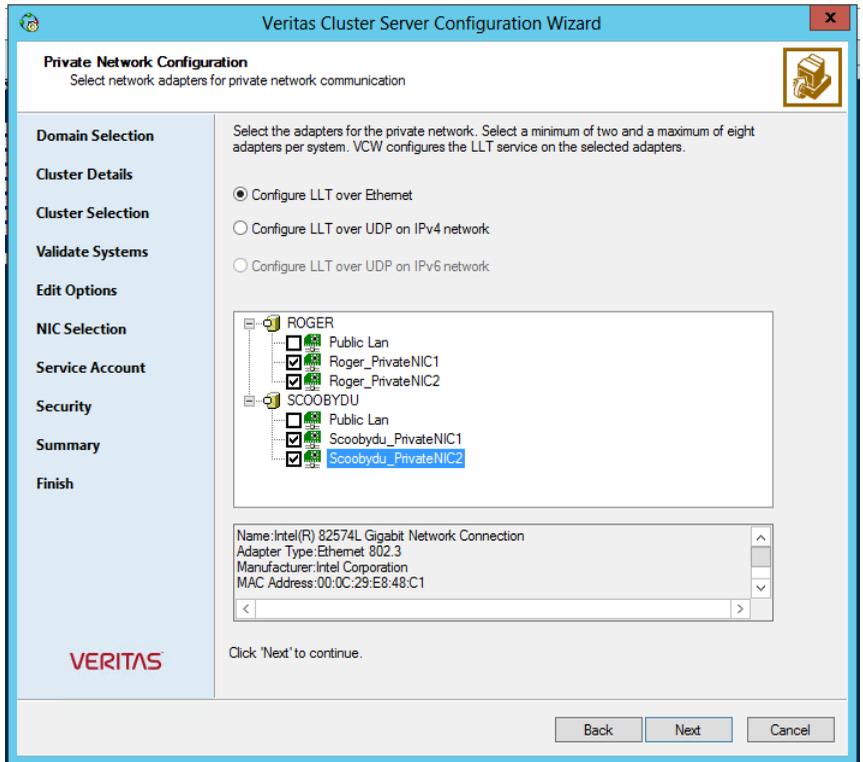
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in step 9, proceed to the next step. Otherwise, proceed to step 12.

- 11** On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer using IPv4 or IPv6 network.

Do one of the following:

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



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

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
- Select **Configure LLT over UDP on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use. The IPv6 option is disabled if the network does not support IPv6.
- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. In case of IPv4, each IP address can be in a different subnet.

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

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.

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

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

- 12** On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service.

The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

Specify the domain user details as follows:

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

- To specify a new user, click **New user** and type a valid user name in the Create New User field and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
 - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13** On the Configure Security Service Option panel, specify security options for the cluster communications and then click **Next**.

Do one of the following:

- To use VCS cluster user privileges, click **Use VCS User Privileges** and then type a user name and password.
The wizard configures this user as a VCS Cluster Administrator. In this mode, communication between cluster nodes and clients, including Cluster Manager (Java Console), occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password.
The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.
Symantec recommends that you specify a new user name and password.
- To use the single sign-on feature, click **Use Single Sign-on**.
In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The Veritas High Availability Engine (HAD) and Veritas Command Server run in secure mode.
The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

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

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

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

- 15** On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService group; this group is required to set up components for 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.

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

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.

See [“Configuring notification”](#) on page 35.

- 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.

Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.

You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.

See [“Configuring Wide-Area Connector process for global clusters”](#) on page 38.

Configuring notification

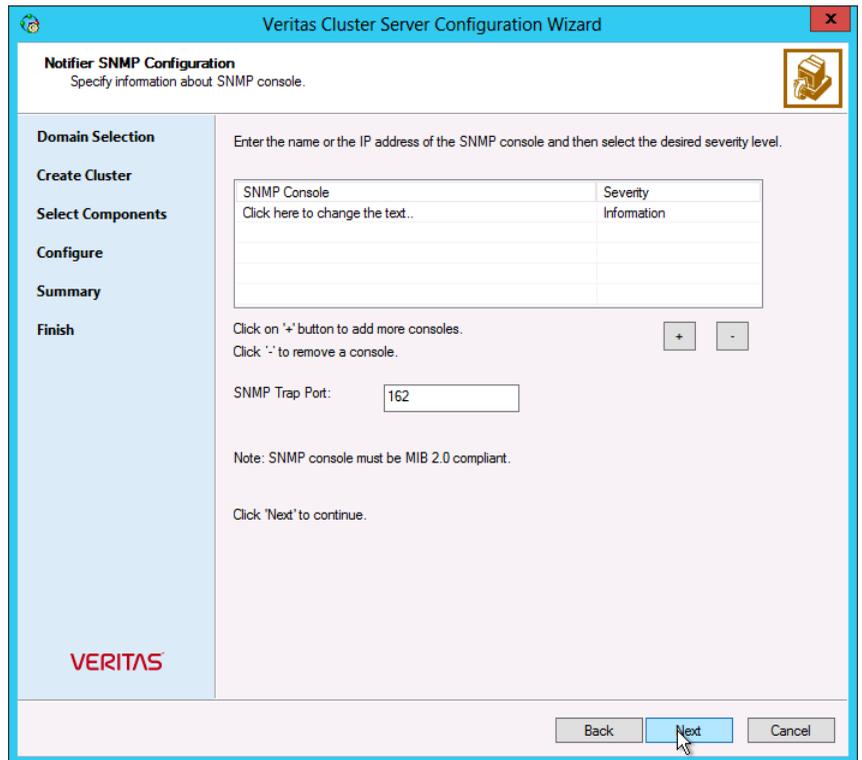
This section describes steps to configure notification.

To configure notification

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

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

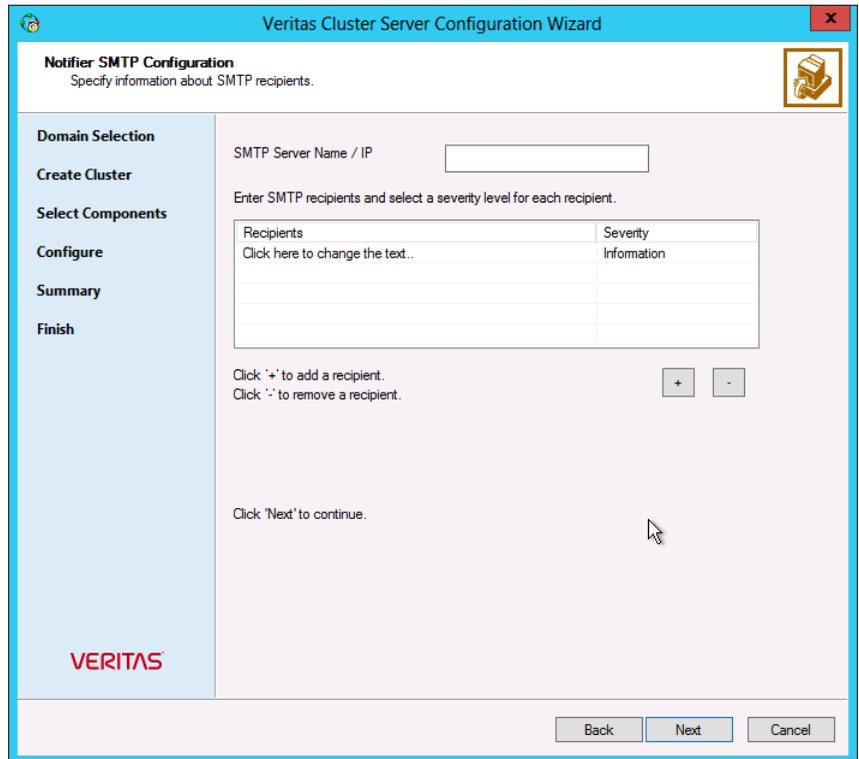
- 2 If you chose to configure SNMP, specify information about the SNMP console and then click **Next**.



Do the following:

- 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 the + icon to add a field; click the - icon 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 then click **Next**.



Do the following:

- 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 the + icon to add fields; click the - icon to remove a field.
- 4 On the Notifier Network Card Selection panel, specify the network information and then click **Next**.

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that 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 starts and click **Configure**.
 - 6 Click **Finish** to exit the wizard.

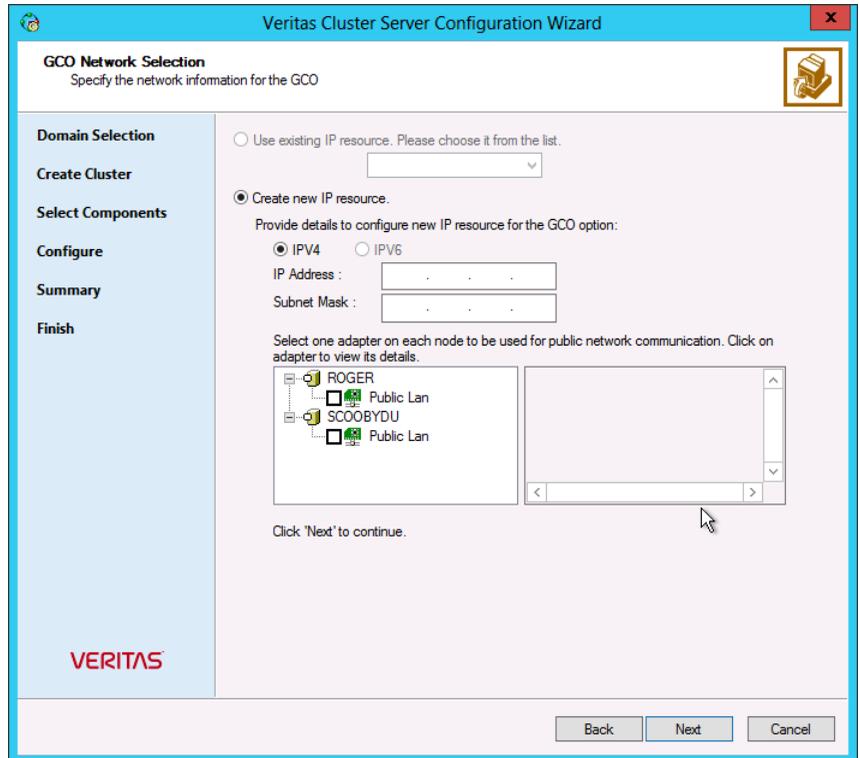
Configuring Wide-Area Connector process for global clusters

Configure the Wide-Area Connector process only if you are configuring a disaster recovery environment. The GCO option configures the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication. Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.

You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.

To configure the wide-area connector process for global clusters

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



If the cluster has a ClusterService group configured, you can use the IP address configured in the service group or configure a new IP address.

Do the following:

- To specify an existing IP address, select **Use existing IP resource** and then select the IP address from the drop-down list.
- To use a new IP address, do the following:
 - In case of IPv4, select **IPV4** and then enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
 - In case of IPv6, select **IPV6** and select the IPv6 network from the drop-down list.
 The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.
 The IPv6 option is disabled if the network does not support IPv6.
- 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 WAC resources online when VCS starts and then click **Configure**.
- 3 Click **Finish** to exit the wizard.

Installing Microsoft Exchange

This chapter includes the following topics:

- [About installing Exchange 2007 in a VCS environment](#)
- [Prerequisites for installing Exchange in a VCS environment](#)
- [Managing storage using NetApp filer](#)
- [Managing storage using Windows Logical Disk Manager](#)
- [Installing Exchange 2007 on the first node](#)
- [Moving Exchange databases to shared storage](#)
- [Installing Exchange 2007 on additional nodes](#)

About installing Exchange 2007 in a VCS environment

This chapter describes how to install Exchange Server and configure a VCS cluster.

If you already have a standalone Exchange Server setup and you want to configure it for high availability.

See [“About configuring a standalone Exchange Server in a VCS environment ”](#) on page 79.

Prerequisites for installing Exchange in a VCS environment

Prerequisites for installing Exchange in a VCS environment are as follows:

- Verify InfoScale Availability is installed on the node.
- Verify you have configured a VCS cluster using VCS Cluster Configuration Wizard (VCW).
See ["Configuring the cluster using the Cluster Configuration Wizard"](#) on page 26.
- Verify the DNS and Active Directory Services are available. Make sure that a reverse lookup zone is created in the DNS.
Refer to Microsoft Exchange documentation for instructions on creating a reverse lookup zone.
- Symantec recommends that the Dynamic Update option for the DNS server be set to "Secure Only."
- Verify the DNS settings for all systems on which Microsoft Exchange will be installed.
See ["Verifying DNS settings for Exchange hosts"](#) on page 44.
- VCS requires Microsoft Exchange to be installed on the same local drive on all nodes. For example if you install Exchange on drive C of one node, installations on all other nodes must be on their respective C drives. Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- In a NetApp storage environment, while configuring Exchange on additional nodes, if there are any other LUNs (not necessarily for Exchange) mounted on the node, ensure that you unmount them before you perform the Exchange post-installation tasks on that node.
- Exchange 2007 installer requires that the Exchange database drive should be mounted on the node when installing Exchange 2007 in RecoverServer install mode. After the Exchange installation is complete, you must unmount the Exchange database LUN before you perform the post-installation tasks on the node.
- If using iSCSI, verify that the Microsoft iSCSI Initiator is configured to establish a non-persistent connection between the NetApp filer and the cluster nodes. Set the *"Automatically restore this connection when the system boots"* option to False.
Symantec recommends that you use non-persistent iSCSI connections to avoid service group concurrency violation issue in a NetApp storage environment. See the Microsoft documentation for instructions.

- If using FC, verify that you install the NetApp FCP Attach Kit or Windows Host Utilities on all the cluster nodes.
Refer to the NetApp documentation for more information.
- Symantec recommends that you create volumes or LUNs (virtual disks), one each for the following:
 - Exchange database
 - Registry replication information
 - Transaction logs for the first storage group
- Verify that the volume or LUN created to store the registry replication information is mounted on the node where you install Microsoft Exchange and unmounted from other nodes in the cluster.
- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.

Privileges for installing Exchange

You must have the following privileges:

- The logged-on user must be a domain user.
- The logged-on user must be logged on with either the Exchange Organization Administrator role or have been delegated the permission to install the server through Setup's server provisioning process.
- The logged-on user must be a part of the Account Operators group in the domain. If the logged-on user account is not a Domain Administrator then the Exchange Servers group must be managed by the logged-on user account or the VCS Helper service user account.
- The logged-on user must be a member of the local Administrators group on all nodes where you are installing Microsoft Exchange and must have write permissions for objects corresponding to these nodes in the Active Directory.
- Either the logged-on user or the VCS Helper service domain user account must have write permissions on the DNS server to perform DNS updates.
- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- If a computer object corresponding to the Exchange virtual server exists in the Active Directory, you must have delete permissions on the object.

- The same user, or a user with the same privileges must perform the pre-installation, installation, and post-installation tasks for Microsoft Exchange.

Verifying DNS settings for Exchange hosts

Complete the following steps to verify the DNS settings.

To verify the DNS settings for Exchange hosts

- 1 Open the Network Connections applet in Control Panel.
- 2 Double-click the adapter.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 3 From the Local Area Connection Status window, click **Properties**.
- 4 On the General tab, check the **Internet Protocol (TCP/IP)** check box and then click **Properties**.
- 5 Select the **Use the following DNS server addresses** option.
- 6 Verify that the correct values for the DNS server IP address and domain name are entered and then click **Advanced**.
- 7 On the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 8 Make sure the correct domain suffix is entered in the DNS suffix for this connection field.

Configuring Microsoft iSCSI initiator

The Microsoft iSCSI initiator enables communication between Windows systems and NetApp Filers. The initiator uses the iSCSI protocol to present the filer volume as a local block device to the system.

Perform the following steps after you have mounted the required LUNs using the NetApp SnapDrive agent.

To configure Microsoft iSCSI initiator on a Windows Server 2008 system:

- 1 Start the Microsoft iSCSI initiator.
- 2 On the Target Portals dialog box, specify the DNS name or IP address for the NetApp Filer.
- 3 On the Favorite Targets tab, remove the corresponding entry.
- 4 Click OK.

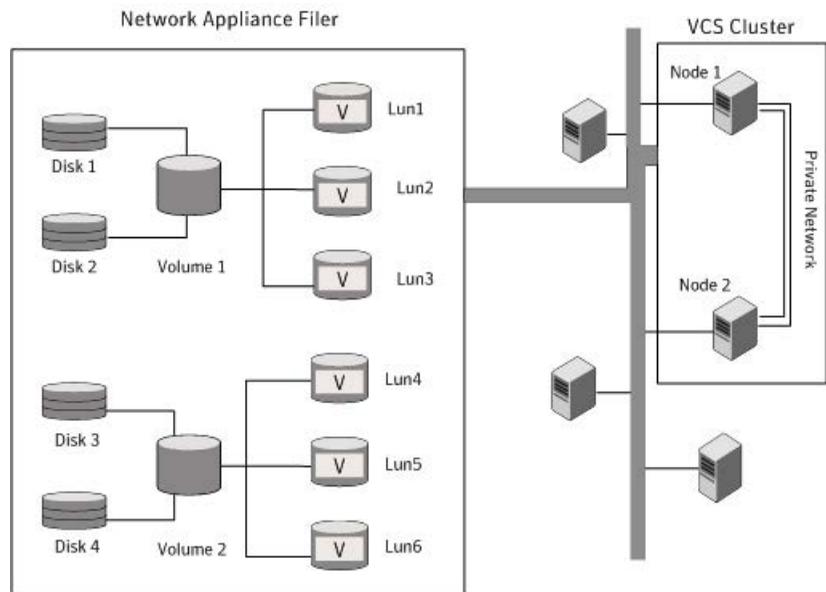
Managing storage using NetApp filer

NetApp manages data by creating volumes on physical disks. These volumes can further be divided into LUNs (Logical Unit Numbers). The LUNs are accessible from the cluster nodes, provided the nodes have Microsoft iSCSI Initiator and NetApp SnapDrive installed. However, if you plan to use Fibre Channel (FC) for connecting the LUNs, ensure that filer is connected to the nodes and the LUNs are shared between all the cluster nodes.

Refer to the NetApp documentation for more information.

The following figure illustrates a typical VCS cluster in a NetApp storage environment.

Figure 3-1 VCS cluster in a NetApp storage environment



Symantec recommends that you create separate LUNs (virtual disks) for the following:

- Exchange database
- Transaction logs for the first storage group
- Registry replication information

If the Exchange database and registry replication files are configured on the same volume, there are potential chances of data corruption after you upgrade Exchange with the latest service pack.

These LUNs must be accessible from all cluster nodes.

Perform the following tasks to create LUNs on the NetApp filer and to make them accessible from cluster nodes:

- Add the filer storage system to the SnapDrive Storage System Management snap-in on the cluster nodes.
- Create volumes on the NetApp filer.
- Share the volumes.
- Create LUNs or virtual disks on the shared volumes.
Refer to NetApp documentation for instructions on performing these tasks.

Connecting virtual disks to the cluster node

Once the virtual disks are created on the NetApp filer, they must be connected (if not connected already) to the cluster nodes using NetApp SnapDrive.

To connect virtual disks to the cluster node

- 1 On the cluster node where you want to connect the LUN, launch the Computer Management MMC from **Start > All Programs > Administrative Tools > Computer Management** or, on Windows Server 2012 operating systems, click **Administrative Tools** from the **Start** screen.
- 2 From the left pane, expand **Storage** and double-click **SnapDrive**.
- 3 Right-click **Disks** and then click **Connect Disk** to launch the Connect Disk wizard.
- 4 Click **Next** on the Welcome page.
- 5 Specify the path of the virtual disk that you wish to connect to the cluster node and then click **Next**.
- 6 Select **Dedicated** as the Virtual Disk Type and then click **Next**.
- 7 Click **Assign a Drive Letter** and then choose a drive letter from the drop-down list.
- 8 On the Select Initiator panel, specify the initiator(s) for the virtual disk and then click **Next**.
- 9 On the igroup Management Type panel, choose the option that allows SnapDrive to perform igroup management automatically and then click **Next**.
- 10 Click **Finish** to begin connecting the specified virtual disk to the cluster node.

Disconnecting virtual disks from the cluster nodes

Perform the following steps to disconnect the virtual disks from a cluster node.

To disconnect virtual disks

- 1 On the cluster node where you want to disconnect the LUNs, launch the Computer Management MMC from **Start > All Programs > Administrative Tools > Computer Management** or, on Windows Server 2012 operating systems, click **Administrative Tools** from the **Start** screen .
- 2 From the left pane, expand **Storage** and double-click **SnapDrive**.
- 3 Double-click **Disks** to see the LUNs that are connected to the node.
- 4 Right-click the LUN you want to disconnect and then click **Disconnect Disk**.
- 5 In the Disconnect Disk alert box, click **OK**.

Managing storage using Windows Logical Disk Manager

If your configuration uses shared disks and volumes that are managed using Windows Logical Disk Manager (LDM), use the VCS Mount and DiskReservation (DiskRes) agents. If you use LDM to manage non-shared local storage, use the VCS Mount and NativeDisks agents.

Before configuring the storage, review the resource types and attribute definitions of these VCS storage agents (Mount, DiskRes, NativeDisks) described in the *Cluster Server Bundled Agents Reference Guide*.

The following restrictions apply for storage managed using LDM:

- Mount, DiskRes, and NativeDisks agents are supported on VCS for Windows only. These agents are not supported if the storage is managed using Storage Foundation (SFW).
- If you are using shared storage, your storage devices must be configured to use SCSI-2 disk reservations. SCSI-3 is not supported. SCSI support is not required if you are using non-shared storage.
- LDM support is not applicable for Disaster Recovery configurations. Currently only HA configurations are supported.

Symantec recommends that you create separate volumes for the following:

- Exchange database
- Transaction logs for the first storage group
- Registry replication information

If the Exchange database and registry replication files are configured on the same volume, there are potential chances of data corruption after you upgrade Exchange with the latest service pack.

If you are using a shared storage configuration, ensure that these volumes are created on shared storage and are accessible from all cluster nodes.

If you are using a non-shared storage configuration, create these volumes separately on the local storage attached to each cluster node.

Perform the following tasks to configure your storage:

- Reserve disks
See [“Reserving disks \(if you use Windows LDM\)”](#) on page 48.
- Create volumes
See [“Creating volumes \(if you use Windows LDM\)”](#) on page 49.
- Mount volumes
See [“Mounting volumes \(if you use Windows LDM\)”](#) on page 49.
- Unassign the drive letter
See [“Unassigning a drive letter”](#) on page 50.
- Release the disks
See [“Releasing disks \(if you use Windows LDM\)”](#) on page 51.

Reserving disks (if you use Windows LDM)

Complete the following steps to reserve the disks on the node on which you are going to perform the application installation.

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

To reserve the disks

- 1 To display all the disks, type the following on the command line:

```
C:\>haval -scsitest /l
```

Make a note of the disk numbers (Disk# column in the table). You will need it in the next step.

- 2 To reserve a disk, type the following on the command line:

```
C:\>haval -scsitest /RES:<disk #>
```

For example, to reserve disk #4, type:

```
C:\>haval -scsitest /RES:4
```

Make a note of the disk number and the corresponding signature. You will require these details to identify and reserve the disks during installation and while configuring the service group, on additional nodes in the cluster.

Creating volumes (if you use Windows LDM)

Perform the following steps to create volumes.

To create volumes

- 1 Use the Windows Disk Management tool to verify that the disks are visible on the cluster nodes, and then create volumes on the disks.
- 2 In case of shared storage, after creating the required volumes on a node, release the reserved disks from that node.

See [“Releasing disks \(if you use Windows LDM\)”](#) on page 51.
- 3 3. In case of shared storage, rescan the disks on all the remaining nodes in the cluster.

Refer to Microsoft Windows documentation for more information about the Disk Management tool.

Mounting volumes (if you use Windows LDM)

Perform the following steps to mount volumes on a cluster node.

To mount a volume

- 1 Use the Windows Disk Management tool to mount the volumes that you created earlier.
- 2 After mounting the volumes on a cluster node, run the CHKDSK command and verify that there are no errors on the mounted volumes.
- 3 Make a note of the drive letters that you assign to the mounted volumes.

Use the same drive letters while mounting these volumes on the remaining cluster nodes.

Refer to Microsoft Windows documentation for more information about the CHKDSK command and the Disk Management tool.

Unassigning a drive letter

In case of a shared storage configuration, while installing an application on multiple nodes, you must first unassign drive letters and release the disks from one node, and then reserve the disks, mount the volumes using the same drive letters and then install the application on the failover node.

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

Note: You must run Disk Management on all systems each time you add a shared disk. This ensures each disk has a valid signature written to it, and that the device paths and symbolic links are updated.

Complete these steps to unassign the drive letters from a node.

To unassign drive letter

- 1 Log in as Administrator.
- 2 Open Disk Management. Type the following at the command prompt:

```
C:\> diskmgmt.msc
```

- 3 Right-click the partition or logical drive and click **Change Drive Letter and Path**.
- 4 In the **Change Drive Letter and Paths** dialog box, click the drive letter and click **Remove**.

Releasing disks (if you use Windows LDM)

Perform the following steps to release reserved disks from a cluster node.

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

To release disks

- 1 To display all the disks, type the following on the command line:

```
C:\>haval -scsitest /l
```

Make a note of the disk numbers (Disk# column in the table) of the disk that you wish to release. You will need it in the next step.

- 2 To release a reserved disk, type the following on the command line:

```
C:\>haval -scsitest /REL:<disk #>
```

For example, to release disk 4, type:

```
C:\>haval -scsitest /REL:4
```

Make a note of the disk number and the corresponding signature. You may require these details to identify and reserve the disks later.

Installing Exchange 2007 on the first node

Installing Exchange 2007 on the first node involves the following three stages:

- Pre-installation
As a part of the pre-installation, you must run Exchange Setup Wizard to change the physical name of the node to a virtual name.
See [“Performing Exchange pre-installation: first node”](#) on page 52.
- Installation
As a part of the installation, you must install Exchange on the node where the Exchange Setup Wizard for VCS was run for the pre-installation tasks.
See [“Performing Exchange Server installation: first node”](#) on page 53.
- Post-installation
As a part of the post-installation, you must revert the node name to original name and set the startup type of the Exchange services to manual.
See [“Performing Exchange post-installation: first node”](#) on page 54.

Performing Exchange pre-installation: first node

Use the Exchange Setup Wizard for Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. After you have run the wizard, you will be prompted 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 or LUN created to store the registry replication information is mounted on or connected to this node and dismounted or disconnected from other nodes in the cluster.
- 2 Navigate to `C:\Program Files\Veritas\cluster server\bin\Exch2007Setup` directory and run the `Exch2007Setup.exe` to launch the **Exchange Server 2007 Setup Wizard**.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 On the Available Option panel, click **Install Exchange 2007 Mailbox Server role for High Availability** and click **Next**.
- 5 On the Select Option panel, click **Create New Exchange Virtual Server** and click **Next**.
- 6 Specify information related to the network.
Specify Exchange Virtual Server details as follows:
 - 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 Server from the VCS environment and reinstall it using the Exchange Server Setup Wizard for VCS.
 - Enter the 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.
 - Specify a unique virtual IP address Exchange server as follows:
 - For IPv4 networks, select **IPv4** and then enter a unique IP address and its associated subnet mask. Ensure that the IP address is currently not being used on your network, but is in the same subnet as the current node.
 - For IPv6 networks, select **IPv6** and then select the IPv6 network from the drop-down list. The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.

- Click **Next**.

The installer verifies that the selected node meets the Exchange requirements and checks whether the Exchange virtual server name is unique on the network.

7 Select a drive for storing the registry replication data and then click **Next**.

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

9 Click **Yes** on the message dialog that informs you that the system will be renamed and restarted after you quit the wizard.

The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task.

10 After all the commands are executed, click **Next**.

11 Click **Reboot**.

When prompted to reboot the node, click **Yes**.

After you reboot the node, the value specified for the Exchange virtual server is temporarily assigned to the node. So, all network connections to the node must be made using the temporary name.

After you reboot the node, the Exchange Setup Wizard for VCS is launched automatically with a message that Pre-Installation is complete. Review the information in the wizard dialog box and proceed to installing Microsoft Exchange Server. Do not click **Continue** at this time. Wait until after the Exchange installation is complete. If you wish to undo all actions performed by the wizard during the pre-installation procedure, click **Revert**.

See [“Performing Exchange Server installation: first node”](#) on page 53.

After installing Microsoft Exchange you must run the Exchange Setup Wizard for VCS again. The wizard will assign the original name to the node.

Performing Exchange Server installation: first node

Install Exchange on the node where the Exchange Setup Wizard for VCS was run for the pre-installation tasks.

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. You can also install Exchange server using the Setup’s server provisioning process. Refer to Microsoft documentation for details.

To install Exchange

- 1 Install Exchange Server using the Microsoft Exchange installation program.
See the Microsoft Exchange documentation for instructions.
- 2 Reboot the node if prompted to do so.
- 3 Install Exchange service packs listed in the requirements.

Performing Exchange post-installation: first node

After completing the installation, use the Exchange Setup Wizard for VCS to complete the post-installation tasks. This process reverts the node name to original name and sets the startup type of the Exchange services to manual, so that they can be controlled by VCS.

Note: If you do not unmount the LUNs before the post-installation tasks, then after post-installation, you may not be able to mount those LUNs on any node. In such a case, you have to then manually remove the corresponding LUN mappings on the NetApp filer and then mount the LUN.

To perform Exchange post-installation

- 1 Make sure that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.

Type the following on the command line:

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

The state should display as RUNNING.

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

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

- 2 Make sure that the volume or LUN containing the registry replication information is mounted on or connected to the node on which you will perform the post-installation.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 5.
If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard for VCS is launched automatically.
- 4 Review the information in the Welcome panel and click **Next**.

- 5 Click **Yes** on the message dialog that informs you that the system will be renamed and restarted after you quit the wizard.

This sets the node name back to its physical host name. The wizard starts performing the post-installation tasks. Various messages indicate the status.

- 6 After all the commands are executed, click **Next** and then click **Finish**.
- 7 When prompted to reboot the node, click **Yes**.

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

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

You cannot use the Exchange Management Console to move the Exchange database created during installation. You must use the Exchange Setup Wizard for VCS to move the Exchange database.

Moving Exchange databases to shared storage

After completing Microsoft Exchange installation on the first node, move the Exchange databases from the local drive of the first node to a location on shared storage.

You must use the Exchange Setup Wizard for VCS to move the Exchange database; you cannot use the Exchange Management Console to move the Exchange database. However, if an Exchange service group is already configured in the cluster, you can also move the Exchange database by running the appropriate cmdlets in the Exchange Management Shell under the virtual server context. If using cmdlets, ensure that the Exchange service group is online.

See [“About VCS support for Exchange Server 2007”](#) on page 12.

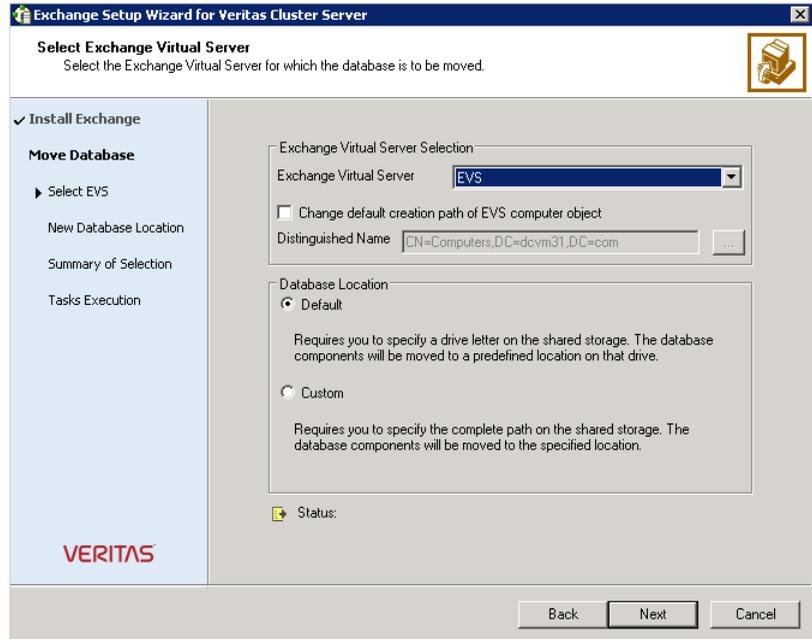
Complete the following tasks before moving the Exchange databases:

- Make sure that the volumes or LUNs created to store the Exchange database, transaction logs, and registry replication information are connected. Disconnect the volumes or LUNs from other systems in the cluster.
- 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 database to a shared storage

- 1 Navigate to `C:\Program Files\Veritas\cluster server\bin\Exch2007Setup` directory and run the **Exch2007Setup.exe** to launch the Exchange Server 2007 Setup Wizard.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Available Option panel, click **Configure/Remove highly available Exchange Server** and then click **Next**.
- 4 In the Select Option panel, click **Move Exchange Databases** and then click **Next**.

- 5 In the Select Exchange Virtual Server panel, select the Exchange virtual server and specify whether you want to move the Exchange databases to a default or a custom location and then click **Next**.

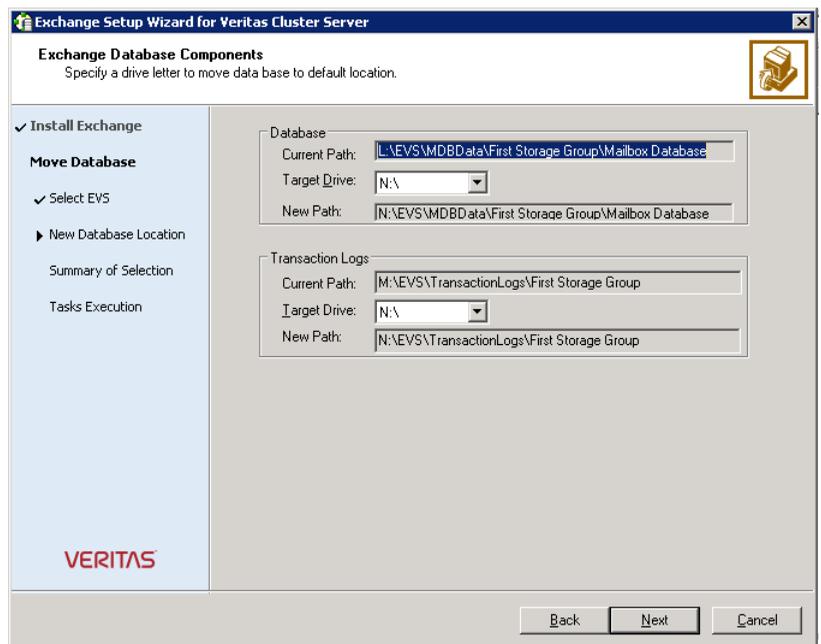


Do the following:

- | | |
|---|---|
| Exchange Virtual Server | From the drop-down list, select the Exchange virtual server for which you want to move the database components. |
| Change default creation path of EVS computer object | Check the Change default creation path of EVS computer object check box if you want to specify a different path for creating the Exchange virtual server computer object. |
| | Then, in the Distinguished Name field type the distinguished name of the Organizational Unit for the virtual server in the format CN=containername,DC=domainname,DC=com . To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. |
| | By default, the Exchange virtual server object is added to the default container "Computers." |

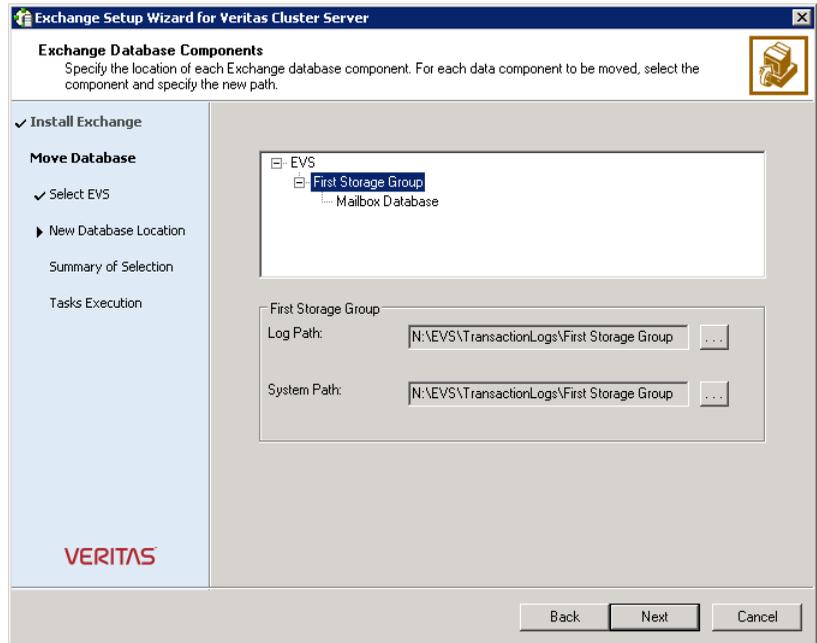
- Default** Select this option if you want to move the database to a default location on the shared storage.
- After you click **Next**, the wizard prompts you for the drive letter on the shared storage. The database components will be moved to a pre-defined location on the drive letter that you select.
- Custom** Select this option if you want to move the database to a custom location on the shared disk.
- After you click **Next**, the wizard prompts you for the drive letter and the complete path on the shared storage. The database components will be moved to the location that you specify.

- 6** On the Exchange Database Components panel, complete the following and then click **Next**:
- If you chose to move the Exchange database to a default location, do the following:



- For the Exchange database, select a drive letter from the respective drop-down list.
- For the Transaction logs, select a drive letter from the respective drop-down list.

- If you chose to move the Exchange database to a custom location, do the following:



- For each data component to be moved, select the component and specify the path to designate the new location of the component. Click the ellipsis button to browse for folders.
 Make sure that 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.
- 7 Review the summary of your selections and click **Next**.
 The wizard starts performing tasks to move the Exchange databases. Various messages indicate the status of each task.
 - 8 After all the tasks are complete, click **Next**.
 - 9 Click **Finish** to exit the wizard.

Installing Exchange 2007 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. Before you proceed with the installation, ensure that you disconnect or dismount the LUNs or

volumes created to store the registry replication information for the Exchange virtual server from the first cluster node where you installed Exchange.

Installing Exchange on additional nodes involves the following 3 stages:

- **Pre-installation**
As a part of pre-installation, use the Exchange Setup Wizard for VCS to change the physical name of the node to a virtual name.
See [“Performing Exchange pre-installation: additional nodes”](#) on page 60.
- **Installation**
As a part of installation, install Exchange on the node where the Exchange Setup Wizard was run for the pre-installation tasks.
See [“Performing Exchange Server installation: additional nodes”](#) on page 62.
- **Post-installation**
As a part of post-installation, use the Exchange Setup Wizard to revert the node name to the original name.
See [“Performing Exchange post-installation: additional nodes”](#) on page 63.

Performing Exchange pre-installation: additional nodes

Use the Exchange Setup Wizard for Cluster Server to complete the pre-installation phase. This process changes the physical name of the node to a virtual name. Before adding a node to the Exchange cluster, make sure you meet the prerequisites for installing Exchange.

Ensure that you disconnect or dismount the LUNs or volumes created to store the registry replication information for the Exchange virtual server from the first cluster node where you installed Exchange.

See [“Prerequisites for installing Exchange in a VCS environment”](#) on page 42.

To perform Exchange pre-installation

- 1** Verify the volume or LUN created to store the registry replication information for this Exchange virtual server is mounted or connected to this node and dismounted or disconnected from other nodes in the cluster.
- 2** From the node to be added to an Exchange cluster, navigate to `C:\Program Files\Veritas\cluster server\bin\Exch2007Setup` and run the `Exch2007Setup.exe` to launch the **Exchange Server 2007 Setup Wizard**.
On Windows 2012 operating systems, use the **Apps** menu in the **Start** screen.
- 3** Review the information in the Welcome panel and click **Next**.
- 4** In the Available Option panel, click **Install Exchange 2007 Mailbox Server role for High Availability** and click **Next**.

- 5 In the Select Option panel, click **Create a failover node for existing Exchange Virtual Server** and click **Next**.
- 6 Select the Exchange virtual server for which you are adding the failover node and click **Next**.
- 7 Specify the network details 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.

Specify network information as follows:

- 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.
- You should not have to change the virtual IP address that is automatically generated when setting up an additional failover node for the virtual server in the same cluster. By default, the text box displays the IP address assigned when the Exchange Virtual Server was created on the first node.

If you want to change the address, do the following:

- For IPv4 networks, select **IPv4** and then enter a unique IP address and its associated subnet mask. Ensure that the IP address is currently not being used on your network, but is in the same subnet as the current node.
 - For IPv6 networks, select IPv6 and then select the **IPv6** network from the drop-down list. The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.
- Click **Next**.
 - 8 Review the summary of your selections and click **Next**.
 - 9 Click **Yes** on the message dialog that informs you that the system will be renamed and restarted after you quit the wizard.

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**.

After you reboot the node, the value specified for the Exchange virtual server is temporarily assigned to the node. So, all network connections to the node must be made using the temporary name. After installing Microsoft Exchange, you must rerun this wizard to assign the original name to the node.

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

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

See [“Performing Exchange Server installation: additional nodes”](#) on page 62.

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

Performing Exchange Server installation: additional nodes

Install Exchange on the node where the Exchange Setup Wizard was run for the pre-installation tasks. This is a standard Microsoft Exchange Server installation.

Refer to the Microsoft documentation for details on this installation.

While installing Exchange, note the following:

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

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
```

Here <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.

Performing Exchange post-installation: additional nodes

After completing the Microsoft Exchange installation, use the Exchange Setup Wizard to complete the post-installation tasks. This process reverts the node name to original name.

Note the following before you proceed:

- Exchange 2007 installer requires that the Exchange database drive should be mounted on the node when installing Exchange 2007 in RecoverServer install mode. After the Exchange installation is complete, you must unmount the Exchange database LUN before you perform the post-installation tasks on the node.
- While configuring Exchange on additional nodes, if there are any other LUNs (not necessarily for Exchange) mounted on the node, ensure that you unmount them before you perform the post-installation tasks on that node.
- If you do not unmount the LUNs before the post-installation tasks, then after post-installation, you may not be able to mount those LUNs on any node. In such a case, you have to then manually remove the corresponding LUN mappings on the NetApp filer and then mount the LUN.

To perform Exchange post-installation

- 1 Make sure that the Veritas High Availability Engine (HAD) is running on the node on which you plan to perform the post-installation tasks.

Type the following on the command line:

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

The state should display as RUNNING.

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

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

- 2 Make sure that the volumes or LUNs containing the registry replication information and the Exchange database are mounted or connected to the node on which you will perform the post-installation.
- 3 If the Exchange installation did not prompt you to reboot the node, click **Continue** from the Exchange Setup Wizard and proceed to step 5.

If you rebooted the node after Microsoft Exchange installation, the Exchange Setup Wizard for VCS is launched automatically.
- 4 Review the information in the Welcome panel and click **Next**.
- 5 Click **Yes** on the message dialog that informs you that the system will be renamed and restarted after you quit the wizard.

The wizard starts performing the post-installation tasks.
- 6 After all commands are executed, click **Next**.
- 7 Specify whether you want to add the node to the SystemList of the service group for the EVS selected in the Exchange pre-installation step.

Select this option only if service groups are already configured for the EVS.

If you wish to add the nodes later, you can do so by using the Exchange service group Configuration Wizard.
- 8 Click **Finish**.
- 9 The wizard prompts you to reboot the node. Click **Yes**.

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

Configuring the Exchange service group

This chapter includes the following topics:

- [About configuring the Exchange service group](#)
- [Running SnapManager for Exchange](#)
- [About the modifications required for tagged VLAN or teamed network](#)
- [About verifying the service group configuration](#)
- [About modifying the Exchange service group configuration](#)

About configuring the Exchange service group

Configuring the Exchange service group involves creating the required VCS resources for Exchange. VCS provides several ways of configuring a service group, including the service group configuration wizard, Cluster Manager (Java Console), and the command line. This section provides instructions on configuring an Exchange service group using the Exchange Server Configuration Wizard.

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.

See [“About modifying the Exchange service group configuration”](#) on page 77.

Review the resource types, the attribute definitions of the agents, the sample configuration and resource dependency graph of the Exchange service group before configuring the agents.

See [“About resource type definitions”](#) on page 109.

Prerequisites for configuring the Exchange service group

Ensure that the following prerequisites are met before configuring the service group:

- Verify that InfoScale Availability is installed on all cluster nodes.
Refer to the *Veritas InfoScale Installation and Upgrade Guide*.
- Verify the cluster is configured.
See [“Configuring the cluster using the Cluster Configuration Wizard”](#) on page 26.
- Verify that Exchange is installed and configured identically on all the cluster nodes.
See [“About installing Exchange 2007 in a VCS environment ”](#) on page 41.
- Verify your DNS server settings. Make sure a static DNS entry maps the virtual IP address with the virtual computer name.
- You must be a Cluster Administrator to create and configure service groups.
- You must be a local Administrator on the node where you run the wizard.
- You must be an Administrator for the NetApp Filer containing the LUNs created to store Exchange data components.
- Verify that Command Server service (CmdServer) 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.
- On the node on which you run the wizard, select Services on the Administrative Tools menu and verify that the Veritas High Availability Engine service is running.
- Verify the volumes or the virtual disks (LUNs) created to store the following data components are mounted or connected to the node where you run the wizard and dismounted or disconnected from other nodes in the cluster:
 - Exchange database
 - Transaction logs for the first storage group
 - Registry replication information
 See [“Managing storage using NetApp filer”](#) on page 45.
 See [“Managing storage using Windows Logical Disk Manager”](#) on page 47.
- 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.
- If you have configured a firewall, add the required ports and services to the Firewall Exception list.
For a detailed list of services and ports used, refer to the *Veritas InfoScale Installation and Upgrade Guide*.

- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.

Configuring the Exchange service group using the wizard

The following steps describe how to configure an Exchange service group using the Exchange Configuration Wizard.

Note: Microsoft Exchange Server and Microsoft SQL Server can exist in the same cluster but cannot run on or fail over to the same system. If a SQL Server service group is configured in the cluster, make sure to select a distinct set of systems in the SystemList attribute for each application's service group.

To configure an Exchange service group

- 1 Navigate to `C:\Program Files\Veritas\cluster server\bin\` directory and run the `Exch2007Wizard.exe` to launch the **Exchange Server 2007 Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, click **Create service group** and click **Next**.

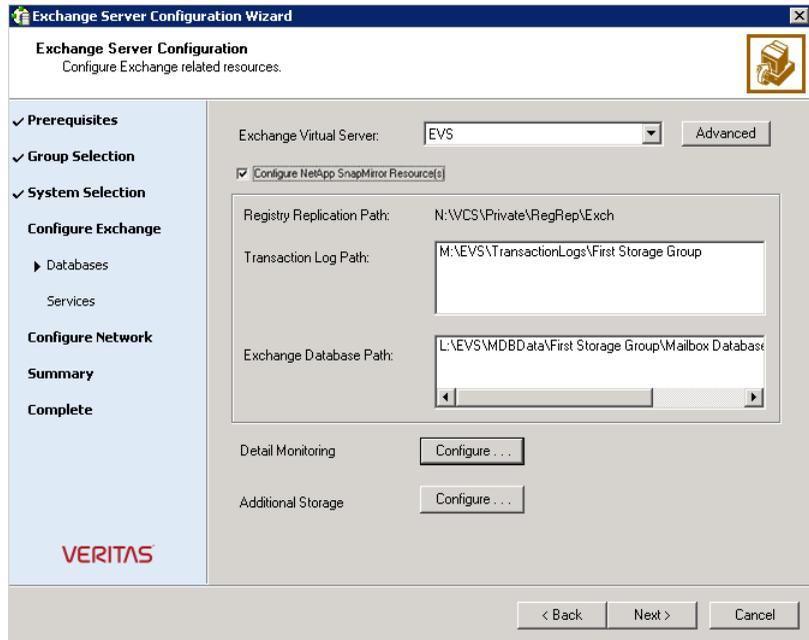
- 4 On the Service Group Configuration panel, specify the service group name and the systems that will part of the service group and then click **Next**.

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

Service Group Name	Type a name for the Exchange service group. If you are configuring the service group on the secondary site, ensure that the name matches the service group name on the primary site.
Available Cluster Systems	Select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order box.
Systems in Priority Order	This list represents the service group's system list. To remove a system from the service group's system list, select the a system and click the left arrow. To change a system's priority in the service group's system list, select the system and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox. For information about the AutoStartList attribute, see the <i>Cluster Server Administrator's Guide</i> .

- 5 On the Exchange Server Configuration panel, verify the Exchange virtual server name and paths to the LUNs created to store Exchange data.

An informational message appears if you chose to configure a SnapMirror resource without configuring replication between NetApp filers at primary and secondary sites. Review the message and click **Yes** to continue. You must always click Yes if you encounter this message while configuring a service group at the secondary site.



Do the following:

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

If you are running the wizard to modify a service group, unchecking this check box will remove the NetApp SnapMirror resource from the service group configuration

- 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:

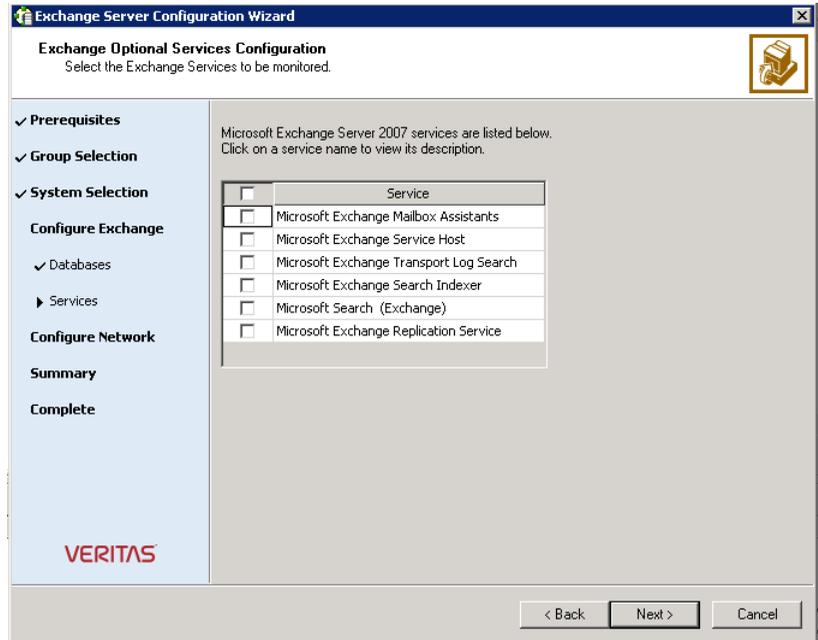
- In the Available Databases box, select the databases for detail monitoring and double-click, or click the right-arrow button to move them to the Selected Databases box. To remove a database, select the database in the Selected Databases box, and double-click or click the left-arrow button.
- Check **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 dismantled, check the **Fault service group if any of the selected database is dismantled** 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 On the Exchange Optional Services Configuration panel, select the optional Exchange services to be monitored and then click **Next**.

Each optional service that is selected will be configured as a VCS resource of type ExchService2007.



- 7 On the Initiator Selection panel, select the initiator(s) for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.

If you are configuring MPIO over FC, you must select at least two FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

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

Do the following:

- For IPv4 networks, the IP Address details displayed are the values entered while installing Exchange. You can keep the displayed values or enter new values.

Select **IPv4** and then enter a unique IP address and its associated subnet mask. Ensure that the IP address is currently not being used on your network, but is in the same subnet as the current node.

If you change the virtual IP address, create a static entry in the DNS server mapping the new virtual IP address to the virtual server name. Ensure that the IP address is currently not being used on your network, but is in the same subnet as the current node.

- For IPv6 networks, select **IPv6** and then select the IPv6 network from the drop-down list. The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.

The IPv6 option is disabled if the network does not support IPv6.

If you select another network, create a static entry in the DNS server mapping the new IPv6 address created by the wizard to the virtual server name.

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

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network.

Note: If you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that have the same MAC address, then you must edit the “MACAddress” attribute of the NIC agent and the IP agent, after you configure the application service group.

See [“About the modifications required for tagged VLAN or teamed network”](#) on page 74.

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

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

The wizard assigns unique names to resources. To edit a resource name, select the resource, click the resource or press the **F2** key. Edit the resource, and press the **Enter** key to confirm the changes. To cancel editing a resource name, press the **Esc** key.

- 10 Click **Yes** on the message that prompts you that the wizard will run commands to modify the service group configuration. Various messages indicate the status of these commands.

- 11** 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 node and click **Finish**.

After bringing the service group online, run the Exchange Management Console to modify the database settings such that all the stores are automatically mounted on start-up.

If you need to configure additional storage groups or mailbox stores, you should do that now.

Perform the following tasks in the order given:

- Create the new storage groups and mailbox stores.
- Mount or connect the volumes or LUNs that have been created for the additional storage groups or mailbox stores.
- You cannot move the Exchange database to the shared storage, using the Exchange Management Console. Move the new storage groups and mailbox stores on the shared storage using either of the following ways:
 - Using the Move Exchange Databases option in the Exchange Setup Wizard for VCS.
 See [“Moving Exchange databases to shared storage”](#) on page 55.
 or
 - Running appropriate cmdlets in the Exchange Management Shell under the virtual server context.
 See [“About VCS support for Exchange Server 2007”](#) on page 12.
- Run the Exchange Configuration Wizard to bring the new storage groups and mailbox stores under VCS control.
 You do not need to run the wizard if you already designated the additional volumes or LUNs when you ran the configuration wizard the first time while installing and configuring Exchange.

- 12** If you want to bring the Offline Address Book under VCS control, perform the following additional steps:

- Using the Java Console, bring the Exchange service group offline.
- Open the registry key
 "HKLM\System\CurrentControlSet\Services\MSEExchangeSA\Parameters"
 and create a registry string value "OAB Folder Location".
- Set the value of "OAB Folder Location" to some path on the shared volume.
 The Offline Address Book gets created at this location.
- Configure the storage resources for the Offline Address Book.

- Create the resource dependency such that the System Attendant resource depends on the storage resources created in the above step.
- Navigate to “%ProgramFiles%\Microsoft\Exchange Server” and remove the “ExchangeOAB” share, if it already exists. Perform this task on each of the cluster node.
- Bring the Exchange service group online.

Running SnapManager for Exchange

Run the SnapManager configuration wizard on the node on which the service group is online, to complete the configuration process and to schedule backups for the Exchange database.

If you are scheduling backups in a VCS cluster, schedule them on the node on which the service group is online.

If the Exchange virtual server fails over to another node, you must set up the backup schedule again on the new node.

Review details about running SnapManager for Exchange. See the NetApp documentation.

You must adhere to the following requirements while running SnapManager for Exchange:

- Make sure the Exchange service group is online.
- Do not move the Exchange database components already moved using the Exchange Server Setup Wizard.

About the modifications required for tagged VLAN or teamed network

Perform this task only if you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that share the same MAC address.

After you configure the application service group, you must edit the “MACAddress” attribute of the VCS NIC agent and the IP agent.

During the application service group configuration, you are required to select a network adapter for each cluster system and specify the virtual IP address for the virtual server. The application configuration wizard internally retrieves the MAC address of the specified interface and the MAC address of the interface to which

the specified IP address is assigned. It then sets these MAC Addresses as the value of the "MACAddress" attribute of the VCS NIC and IP agent respectively.

If the selected interface or the interface to which the specified IP is assigned shares the MAC address with other logical interfaces, then the following issues may occur:

- NIC agent may begin to monitor an interface other than the one selected.
- The IP agent may assign the specified virtual IP address or the virtual server name to an interface other than the one selected. As a result, the IP agent may monitor an IP address other than the one specified.

As a workaround, use the VCS Java Console to edit the "MACAddress" attribute and specify its value as the interface name instead of the MAC address. You must enter the interface name in double quotes. For example, MACAddress = "InterfaceName"

Notes:

- After you specify the interface name as the "MACAddress" attribute value, if you want to use the VCS wizards to modify any settings, then you must first reset the value of the "MACAddress" attribute to the MAC address of the interface. Failing this, the VCS wizard may fail to identify and populate the selected interface. Use the VCS Java Console to edit the attribute values.
- If you change the interface name, you must update the "MACAddress" attribute value to specify the new name. Failing this, the NIC resource will go in an UNKNOWN state.
- While editing the "MACAddress" attribute to specify the interface name, you must specify the name of only one interface.

About verifying the service group configuration

This section provides steps to verify a service group configuration by bringing the service group online, taking it offline, and switching the service group to another cluster node.

Bringing the service group online

Perform the following steps to bring the service group online from the VCS Cluster Manager (Java Console).

To bring a service group online from the Java Console

- 1 In the Cluster Explorer configuration tree, select the Exchange service group to be taken online.
- 2 Right-click the service group and select to online the service group on the system. (Right-click > Online > *system_name*)

Taking the service group offline

Perform the following steps to take the service group offline from the VCS Cluster Manager (Java Console).

To take a service group offline from the Java Console

- 1 On the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.

or

Select the cluster in the Cluster Explorer configuration tree, select the Service Groups tab, and right-click the service group icon in the view panel.
- 2 Choose **Offline**, and choose the appropriate system from the pop-up menu. (Right-click > Offline > *system_name*)

Switching the service group

The process of switching a service group involves taking it offline on its current system and bringing it online on another system. Perform the following steps to switch the service group from the VCS Cluster Manager (Java Console).

To switch a service group from the Java Console

- 1 On the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.

or

Select the cluster in the Cluster Explorer configuration tree, select the Service Groups tab, and right-click the service group icon in the view panel.
- 2 Choose **Switch To**, and choose the appropriate system from the pop-up menu. (Right-click > Switch To > *system_name*)

About modifying the Exchange service group configuration

You can dynamically modify the Exchange service group configuration in several ways, including the Exchange Server Configuration Wizard, and the command line. The following steps describe how to modify the service group using the configuration wizard.

Prerequisites for modifying an Exchange service group

Prerequisites for modifying an Exchange service group are as follows:

- If the Exchange service group is online, you must run the wizard from the node on which the service group is online. You can then use the wizard to add and remove resources. You cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the NetAppFiler and NetAppSnapDrive resources for the service group must be online on the node where you run the wizard and offline on all other nodes.
So, the LUNs created to store the Exchange database, the registry replication information and the transaction logs should be connected to node where you run the wizard.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard from the node being removed.
- After the application service group configuration, if you have manually edited any of the resource attributes, then you must reset them to their default values. Failing this, the wizard may fail to identify and populate the resources involved in the service group configuration.
After you modify the service group configuration you can again edit the resource attributes to set the desired value.

Modifying the Exchange service group

The following steps describe how to modify an Exchange service group using the Exchange Server configuration wizard. If you run the wizard to add a system to an online service group, resources having local attributes may go in an UNKNOWN state for a short duration. These resources will come out of the UNKNOWN state in the next monitor cycle.

To modify an Exchange service group

- 1 Navigate to `C:\Program Files\Veritas\cluster server\bin\` directory and run the `Exch2007Wizard.exe` to launch the **Exchange Server 2007 Configuration Wizard**.

On Windows 2012 operating systems, use the **Apps** menu in the **Start** screen.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Modify service group**, click the service group to be modified, and click **Next**.
- 4 Follow the wizard instructions and make desired modifications to the service group configuration.

See [“About configuring the Exchange service group”](#) on page 65.

Deleting the Exchange service group

The following steps describe how to delete an Exchange service group using the configuration wizard.

To delete an Exchange service group

- 1 Navigate to `C:\Program Files\Veritas\cluster server\bin\` directory and run the `Exch2007Wizard.exe` to launch the **Exchange Server 2007 Configuration Wizard**.

On Windows 2012 operating systems, use the **Apps** menu in the **Start** screen.
- 2 Review the information on the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, click **Delete service group**, click the service group to be deleted and click **Next**.
- 4 In the Service Group Summary panel, click **Next**.

A message appears informing you that the wizard will run commands to delete the service group.
- 5 Click **Yes** to delete the service group.
- 6 Click **Finish**.

Making a standalone Exchange server highly available

This chapter includes the following topics:

- [About configuring a standalone Exchange Server in a VCS environment](#)
- [Renaming the existing Exchange Server and assigning a name to the physical node](#)

About configuring a standalone Exchange Server in a VCS environment

Configuring an existing Exchange Server 2007 in a VCS environment involves the following tasks. This scenario considers an active/passive configuration with one to one failover capabilities.

Note: High availability 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.

The following table lists the tasks for configuring a standalone Exchange Server for high availability.

Table 5-1 Tasks for configuring a standalone Exchange Server for high availability

Task	Description
Install InfoScale Availability on the standalone Exchange Server	Refer to the <i>Veritas InfoScale Installation and Upgrade Guide</i> for instructions. During installation, the product installer installs the VCS agents required for making the application highly available.
Change the node name of the Exchange Server and assign a name for the Exchange Virtual Server	See “Renaming the existing Exchange Server and assigning a name to the physical node” on page 81.
Install InfoScale Availability on additional nodes and configure a cluster	<p>Note: Perform this task only if you plan to configure a new cluster. Skip this task if you have already configured the cluster.</p> <p>Refer to the <i>Veritas InfoScale Installation and Upgrade Guide</i>. To configure a cluster, refer to, See “Configuring the cluster using the Cluster Configuration Wizard” on page 26.</p>
Add the standalone Exchange Server node to the cluster	<p>After converting the standalone Exchange server into a virtual server, use the VCS Cluster Configuration Wizard (VCW) to add the node to the cluster.</p> <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 26.</p>
Move the Exchange databases	<p>Use the Exchange Server Setup Wizard for VCS with the Move Exchange Database option to move the Exchange database from the clustered Exchange server to a shared disk.</p> <p>See “Moving Exchange databases to shared storage” on page 55.</p>
Install Exchange on the additional nodes	<p>After moving the Exchange databases to shared storage, install Exchange on additional nodes in the cluster for the same Exchange virtual server, if required. You must run pre-installation, installation, and post-installation procedures for each additional node.</p> <p>See “Installing Exchange 2007 on additional nodes” on page 59.</p>

Table 5-1 Tasks for configuring a standalone Exchange Server for high availability (*continued*)

Task	Description
Configure the Exchange service group	<p>Configure an Exchange service group in the cluster using the Exchange Server Configuration Wizard. Configuring the Exchange service group involves creating an Exchange service group and defining the attribute values for its resources.</p> <p>After the Exchange service group is created, you must configure the databases to mount automatically at startup.</p> <p>See “About configuring the Exchange service group” on page 65.</p>

Renaming the existing Exchange Server and assigning a name to the physical node

Use the Exchange Setup Wizard to convert a standalone Exchange server into a "clustered" Exchange server. In this process, 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.

Before proceeding, make sure you meet the following requirements:

- You have installed InfoScale Availability on the system that hosts Exchange. See [“About installing the VCS agents”](#) on page 26.
- HA support for Exchange Server 2007 is available for the Mailbox Server role only. 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 you proceed.
- The system hosting the Exchange server to be made highly available is not configured as a root broker.
- The system hosting the Exchange server does not have VCS configured.

To convert a standalone Exchange Server into a "clustered" Exchange server

- 1 From the node having the standalone Exchange server installed, navigate to `C:\Program Files\Veritas\cluster_server\bin\Exch2007Setup` directory and run the `Exch2007Setup.exe` to launch the **Exchange Server 2007 Setup Wizard**.

On Windows 2012 operating systems, use the **Apps** menu in the **Start** screen.

- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Available Option panel, click **Make a standalone Exchange Server highly available** and click **Next**.
- 4 Specify information related to your network, as follows:
 - Enter a name for the node.
This name will be permanently assigned to the node on which the wizard is being run. 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 name of the domain suffix.
 - Select the appropriate public NIC from the drop-down list.
The wizard lists the public adapters and low-priority TCP/IP enabled private adapters on the system.
 - Enter a unique virtual IP address for the Exchange virtual server.
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 Select a drive where the registry replication data will be stored and click **Next**.
Make sure to select the volumes or the virtual disks (LUNs) to store registry replication information.
- 6 Review the summary of your selections and click **Next**.
- 7 Click **Yes** on the message dialog that informs you that the system will be renamed and restarted after you quit the wizard.

The wizard starts running commands to set up the VCS environment. Various messages indicate the status of each task.

Renaming the existing Exchange Server and assigning a name to the physical node

8 After all the commands are executed, click **Next**.

9 Click **Finish**.

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.

Configuring any-to-any failover

This chapter includes the following topics:

- [About any-to-any configuration](#)
- [Specifying a common node for failover](#)

About any-to-any configuration

An any-to-any Exchange Server configuration could have many Exchange virtual servers in a cluster, each configured in a separate service group. Each service group can fail over to any configured node in the cluster, provided no other Exchange virtual server is online on that node.

Configuration of failover nodes for Exchange instances 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.

The following table lists the tasks for configuring any-to-any configuration

Table 6-1 Tasks for configuring any-to-any configuration

Task	Description
Install InfoScale Availability and configure the cluster	<p>Refer to the <i>Veritas InfoScale Installation and Upgrade Guide</i> for instructions. During installation, the product installer installs the VCS agents required for making the application highly available.</p> <p>Run the VCS Cluster Configuration Wizard (VCW) to configure the components required to run the VCS cluster.</p> <p>See “About any-to-any configuration” on page 84.</p>
Install Exchange	<p>See “Installing Exchange 2007 on the first node” on page 51.</p>
Configure the Exchange service group for the first Exchange virtual server	<p>Configure an Exchange service group in the cluster using the Exchange Server Configuration Wizard. Configuring the Exchange service group involves creating an Exchange service group and defining the attribute values for its resources.</p> <p>After the Exchange service group is created, you must configure the databases to mount automatically at startup.</p> <p>See “About configuring the Exchange service group” on page 65.</p>
Move the Exchange databases	<p>Use the Exchange Server Setup Wizard for VCS with the Move Exchange Database option to move the Exchange database from the clustered Exchange server to a shared disk.</p> <p>See “Moving Exchange databases to shared storage” on page 55.</p>
Specify a common node for failover	<p>See “Specifying a common node for failover” on page 86.</p>
Install Exchange on the additional nodes	<p>After moving the Exchange databases to shared storage, install Exchange on additional nodes in the cluster for the same Exchange virtual server, if required. You must run pre-installation, installation, and post-installation procedures for each additional node.</p> <p>See “Installing Exchange 2007 on additional nodes” on page 59.</p>

Table 6-1 Tasks for configuring any-to-any configuration (*continued*)

Task	Description
Configure the Exchange service group for the second Exchange virtual server	<p>A new Exchange service group must be configured for the second Exchange virtual server. Configure an Exchange service group in the cluster using the Exchange Server Configuration Wizard.</p> <p>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.</p> <p>See “About configuring the Exchange service group” on page 65.</p>

Specifying a common node for failover

Exchange was installed on the node that will be the common failover node during the installation process of the first Exchange virtual server. You do not need to install Exchange a second time on the common failover node. You must run the Exchange Setup Wizard for VCS to set up the any-to-any failover configuration in the Exchange cluster.

The first Exchange virtual server is already configured with the common failover node. You must run this wizard for the second Exchange virtual server only.

To specify a common node for failover

- 1 From any node configured to host an Exchange service group, navigate to `C:\Program Files\Veritas\cluster_server\bin\Exch2007Setup` directory and run the `Exch2007Setup.exe` to launch the **Exchange Server 2007 Setup Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Available Option panel, click **Configure/Remove highly available Exchange Server** and click **Next**.
- 4 In the Select Options panel, click **Configure any-to-any failover** and click **Next**.
- 5 Select systems to be configured for any-to-any failover and then click **Next**.
Do the following:
 - Select the Exchange virtual server to which you want to add the additional failover nodes.

- From the Available Systems box, select the systems to be configured for any-to-any failover.
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.
The Existing Systems box lists the nodes on which the Exchange Server service group can fail over.
- Click the right arrow to move the selected systems to the Selected Systems box. To remove a system from the box, select the system and click the left arrow.
- Specify whether you want to add the systems to the SystemList of the service group for the selected EVS.
- Click **Configure**.
Repeat these steps for all the Exchange virtual servers for which you want to configure any-to-any failover.

6 Click **Finish**.

Deploying agents for disaster recovery

This chapter includes the following topics:

- [About disaster recovery configuration](#)
- [Setting up disaster recovery configuration](#)
- [Configure replication using NetApp SnapMirror](#)
- [Configure NetAppSnapMirror resources at the primary site](#)
- [About managing failover in a disaster recovery environment](#)

About disaster recovery configuration

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.

Setting up disaster recovery configuration

The following table lists the tasks for disaster recovery configuration for Exchange Server.

Table 7-1 Tasks for configuring a disaster recovery setup

Tasks	Description
Set up the primary site	<p>Make sure you have completed the following tasks at the primary site:</p> <ul style="list-style-type: none"> ■ Install InfoScale Availability with the GCO option and then configure the cluster. While configuring the cluster, ensure that you select the GCO option to configure the wide area connector (WAC) resource in the cluster. For details refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>. For details on configuring a cluster refer to <i>Cluster Server Administrator's Guide</i>. ■ Install Microsoft Exchange Server on the primary sites cluster nodes. See "About installing Exchange 2007 in a VCS environment" on page 41. ■ Configure the Exchange service group See "About configuring the Exchange service group" on page 65.
Set up the secondary site	<p>Perform the following tasks to set up the cluster on the secondary site:</p> <ul style="list-style-type: none"> ■ Install InfoScale Availability and configure the cluster. ■ Offline the service group in the primary site cluster. ■ Connect to the LUNs created to store the registry replication information using the same drive letters and LUN names used at the primary site. ■ Install Microsoft Exchange Server on the cluster nodes. See "About installing Exchange 2007 in a VCS environment" on page 41.

Table 7-1 Tasks for configuring a disaster recovery setup (*continued*)

Tasks	Description
<p>Back up and restore the Exchange data files</p>	<p>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 volumes or LUNs on the primary site and then restore it on the secondary site. Make sure you restore the data at the same path at the secondary site as on the primary site.</p> <p>Note: Perform this step only once at the secondary site.</p> <p>Complete the following tasks:</p> <ul style="list-style-type: none"> ■ On the primary site, back up all volumes or LUNs specified for Exchange data files. ■ Restore the group in the corresponding location on the secondary site. See the NetApp documentation for instructions on restoring data.
<p>Configure the Exchange service group at the secondary site</p>	<p>Configuring the Exchange service group involves creating an Exchange service group and defining the attribute values for its resources.</p> <p>See “About configuring the Exchange service group” on page 65.</p> <p>Note the following before configuring the Exchange service group:</p> <ul style="list-style-type: none"> ■ Make sure the service group has the same name as in the primary site cluster. ■ In case of a NetApp storage environment, make sure you configure NetApp SnapMirror resources in the service group. ■ Do not bring the service group online. <p>Note that the service group may be partially online because the LUNs are connected to the node.</p>
<p>Configure replication using NetApp SnapMirror</p>	<p>You can replicate Exchange data by establishing a SnapMirror relationship between the filers at the primary and secondary sites.</p> <p>Before configuring replication, make sure the service group is offline at the secondary site.</p> <p>SnapMirror replicates snapshots taken on a filer and applies them to a remote filer over a wide area network; these snapshots can be used by the target host to provide rapid failover in case of a disaster.</p> <p>You can transfer the initial base snapshot image from the primary to secondary via tape, and then set up incremental SnapMirror updates to the destination filer.</p> <p>Refer to NetApp documentation for more information.</p>

Table 7-1 Tasks for configuring a disaster recovery setup (*continued*)

Tasks	Description
Configure NetApp SnapMirror resources at the primary site	Configure NetAppSnapMirror resources at the primary site to monitor data replication from the primary site to the secondary site. See “Configure NetAppSnapMirror resources at the primary site” on page 92.
Link clusters at primary and secondary sites	Once all the setup tasks are completed at the primary and secondary sites, you must link the clusters at both the sites. The VCS Java Console provides a wizard to create global cluster by linking standalone clusters. For instructions, review the chapter on Administering Global Clusters from Cluster Manager (Java Console). See the <i>Cluster Server Administrator’s Guide</i> .
Make the Exchange service group global	After linking the clusters at the primary and secondary sites, use the Global Group Configuration Wizard from the Java Console to convert the Exchange service group from a local service group to a global service group. This will enable the Exchange service group to fail over across clusters. For instructions, review the chapter on Administering Global Clusters from Cluster Manager (Java Console). See the <i>Cluster Server Administrator’s Guide</i> .

Configure replication using NetApp SnapMirror

You can replicate Exchange data by establishing a SnapMirror relationship between the filers at the primary and secondary sites.

Before configuring replication, make sure the service group is offline at the secondary site.

SnapMirror replicates snapshots taken on a filer and applies them to a remote filer over a wide area network; these snapshots can be used by the target host to provide rapid failover in case of a disaster.

You can transfer the initial base snapshot image from the primary to secondary via tape, and then set up incremental SnapMirror updates to the destination filer.

Refer to NetApp documentation for more information.

Configure NetAppSnapMirror resources at the primary site

Configure NetAppSnapMirror resources at the primary site to monitor data replication from the primary site to the secondary site. The following steps describe how to add the resources using the Exchange Server Configuration Wizard.

You may want to repeat this procedure and create a NetAppSnapMirror resource at the secondary site.

This is required in the following cases:

- The service group is online at the secondary site (either it is failed over or switched to the secondary site) and the filer should replicate from secondary to primary site.
- If you want to fail over or switch the service group from the secondary to the primary site.

To configure SnapMirror resource using Exchange Server Configuration Wizard

- 1 Verify the volumes or LUNs created to store the registry replication information and the Exchange database are mounted on or connected to this node and dismounted or disconnected from other nodes in the cluster.
- 2 Navigate to `C:\Program Files\Veritas\cluster server\bin\` directory and the `Exch2007Wizard.exe` to launch the **Exchange Server 2007 Configuration Wizard**.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 In the Wizard Options panel, click **Modify service group**, click the service group to be modified, and click **Next**.
- 5 In the Service Group Configuration panel, verify the list of systems in the service group and click **Next**.
- 6 In the Exchange Server Configuration panel, check **Configure the NetApp SnapMirror resource(s)** and click **Next**.
- 7 Accept default values in the subsequent dialog boxes and click **Next** till you reach the wizard completion panel.
- 8 In the Completing the Exchange Configuration panel, uncheck the **Bring the service group online** check box and click **Finish**.

About managing failover in a disaster recovery environment

In a disaster recovery configuration, VCS first attempts to fail over the application to a node in the local cluster. If all nodes in the local cluster are unavailable, or if a disaster strikes the site, VCS attempts to fail over the application to the remote site.

Remote failover involves starting the Exchange services on a node in the remote cluster. In case of an administrative failover, this also involves reversing the direction of replication by demoting the original source to a target, and replicating from the new source.

Managing a successful remote failover

For a successful failover, you must perform the following tasks after the service group comes online at the remote site.

To manage a successful failover

- 1 Freeze the Exchange service group at the remote site.
- 2 Restore Exchange data from the latest valid database snapshot using the NetApp SnapManager Restore utility.
- 3 Unfreeze the Exchange service group.

Managing failover in response to a network outage

In the event that the public network or the private storage network at the local cluster fails, the application fails over to the remote site. Perform the following tasks to ensure a proper failover.

To ensure a proper failover

- 1 Freeze the service group at the local site.
- 2 Restore the network connections. You may see concurrency violation errors in the engine log. Ignore these errors.
- 3 Unfreeze the service group.
- 4 Take the service group offline at the local site.
- 5 Freeze the service group at the remote site.
- 6 Restore Exchange data from the latest valid database snapshot using the NetApp SnapManager Restore utility.
- 7 Unfreeze the Exchange service group.

About switching the service group back to the local cluster

When the application fails over to a remote site, switching the application back to the local site involves the following additional tasks, depending on whether the failover was administrative or in response to a disaster.

Switching back the service group in case of Administrative failover

In case of an administrative failover, VCS brings the service group online at the remote site and reverses the direction of replication.

To switch the service group back to the local cluster

- 1 Back up the Exchange data using NetApp SnapManager.
See the NetApp documentation for instructions.
- 2 Switch the service group, as follows:
 - In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
 - Click **Switch To**, and click **Remote switch**.
 - Select a system at the local site and click **OK**.

Switching back the service group after a disaster

In the event that a disaster strikes the local cluster and the application fails over to the remote site, data is written to the LUNs at the remote site. When the local site comes up again, the Exchange data at both sites is out-of-sync.

To switch the service group back to the local cluster

- 1 Synchronize the Exchange data at both sites by running the fbsync action at the site at which the service group is online.

```
# hares -action SnapMirror_resname fbsync -sys node_name
```

The variable `SnapMirror_resname` represents the name of the SnapMirror resource; `node_name` represents the node on which the service group is online.

Run the action for each SnapMirror resource.

- 2 Back up the Exchange data using NetApp SnapManager.
See the NetApp documentation for instructions.
- 3 Switch the service group, as follows:

- In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- Click **Switch To**, and click **Remote switch**.
- Select a system at the local site and click **OK**.

Removing the software

This chapter includes the following topics:

- [About removing the software components](#)
- [Remove Microsoft Exchange](#)

About removing the software components

You can remove a node from the cluster without removing Microsoft Exchange or remove Microsoft Exchange and then remove the node from the cluster.

To remove InfoScale Availability without removing Microsoft Exchange, remove the node from the configuration using VCW and then use the Veritas Product Installer or the command line to remove the software.

To remove the node from the cluster configuration refer to *Cluster Server Administrator's Guide*.

To uninstall InfoScale Availability refer to *Veritas InfoScale Installation and Upgrade Guide*.

To remove Microsoft Exchange refer to, Microsoft product documentation.

Note the following tasks to remove the software

- Verify that you have local administrative privileges on the node where you are removing the software components.
- Verify that all Exchange service groups are offline on the node you want to remove.
- Remove the node from the Exchange service group SystemList.
- Verify that the user mailboxes and routing group connectors are deleted from the system where you are removing Microsoft Exchange.

Remove Microsoft Exchange

The Exchange Server Setup Wizard for VCS performs the following tasks for removing Microsoft Exchange from a node:

- If the node being removed is configured to host other Exchange virtual servers, the wizard removes the node from the SystemList of the service group for the specified Exchange virtual server. The wizard does not remove Microsoft Exchange from the node.
See [“Removing a node without removing Microsoft Exchange”](#) on page 97.
- If the node being removed is not configured to host other Exchange virtual servers, the wizard removes the node from the SystemList of the service group for the specified Exchange virtual server. The wizard also removes Microsoft Exchange from the node by launching the Microsoft Exchange Installation wizard. See [“Removing a node and removing Microsoft Exchange”](#) on page 98.
If you are uninstalling Microsoft Exchange from all nodes in the cluster, delete the service group after taking it offline.

Removing a node without removing Microsoft Exchange

Complete these steps if the node being removed is configured to host other Exchange virtual servers. The wizard removes the node from the SystemList of the service group for the specified Exchange virtual server. The wizard does not remove Microsoft Exchange from the node.

These steps describe how to remove a node without removing Microsoft Exchange.

To remove a node without removing Microsoft Exchange

- 1 Navigate to `C:\Program Files\Veritas\cluster server\bin\Exch2007Setup` directory and run the `Exch2007Setup.exe` to launch the **Exchange Server 2007 Setup Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Available Option panel, click **Configure/Remove highly available Exchange Server** and click **Next**.
- 4 In the Select Option panel, click **Remove Exchange Server** and click **Next**.

If an Exchange service group is configured on the node, the wizard prompts you to remove the system from the service group's SystemList attribute. Resolve the error and re-run the Exchange Server Setup Wizard.

- 5 Select the Exchange virtual server for which you are removing the failover node and click **Next**.

The wizard starts running commands to set up the VCS environment for removing the node from the Exchange service group. Various messages indicate the status of each command.

- 6 Once all the commands are executed, click **Next**.
- 7 Click **Finish**.

Proceed to uninstall InfoScale Availability.

For details refer to the *Veritas InfoScale Installation and Upgrade Guide*.

Removing a node and removing Microsoft Exchange

These steps describe how to remove a node and remove Microsoft Exchange.

To remove a node and remove Microsoft Exchange

- 1 Navigate to `C:\Program Files\Veritas\cluster server\bin\Exch2007Setup` directory and run the `Exch2007Setup.exe` to launch the **Exchange Server 2007 Setup Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Available Option panel, click **Configure/Remove highly available Exchange Server** and click **Next**.
- 4 In the Select Option panel, click **Remove Exchange Server** and click **Next**. If an Exchange service group is configured on the node, the wizard prompts you to remove the system from the service group's SystemList attribute. Resolve the error and rerun the Exchange Server Setup Wizard.
- 5 Select the Exchange virtual server for which you are removing the failover node and click **Next**.
- 6 Click **Yes** on the message dialog that informs you that the system will be renamed and restarted after you quit the wizard.

The wizard starts running commands to set up the VCS environment for removing the node from the Exchange service group.

- 7 Once all the commands are executed, click **Next**.

If the node is the last node in the Exchange service group, the wizard prompts you to choose whether you want to retain the entry for the EVS in the Active Directory. Click **Yes** to remove the entry or **No** to retain the entry.

8 Click Reboot.

The wizard prompts you to restart the node. Click **Yes** to restart the node.

If you have other applications running, click **No**, close all applications, and restart the node manually.

9 Restarting the node automatically launches the Exchange Server Setup Wizard for VCS. Review the information in the Welcome dialog box and click Next.

10 Click Yes on the message dialog that informs you that the system will be renamed and restarted after you quit the wizard.

11 In the Microsoft Exchange Installer Welcome panel, read the welcome information and click Next.

12 In the Server Role Selection panel, clear the Mailbox Role check box and the Management Tools check box if applicable, and then click Next. The Exchange installer will proceed with the uninstallation.

13 In the completion panel, click Finish.

Do not reboot the node at this stage. The Exchange Server Setup Wizard for VCS must complete its operations before the node is rebooted.

The Exchange Server Setup Wizard for VCS will be launched automatically. The wizard performs the post-uninstallation tasks.

14 Once all the tasks are complete, click Next.

15 Click Finish.

The wizard prompts you to restart the node. Click **Yes** to restart the node.

If you have other applications running, click **No**, close all applications, and restart the node manually.

Troubleshooting the agents

This chapter includes the following topics:

- [About troubleshooting VCS agents](#)
- [VCS logging](#)
- [NetApp agents error messages](#)
- [Exchange Service agent error messages](#)
- [Troubleshooting Microsoft Exchange uninstallation](#)
- [About troubleshooting Exchange Setup Wizard issues](#)

About troubleshooting VCS agents

This chapter describes how to troubleshoot common problems in the VCS agents for NetApp and 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. The letter A indicates the first log file, B indicates the second, C indicates 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 (YYYY/MM/DD) | Mnemonic | Severity | UMI | Agent Type | Resource Name | Entry Point | Message Text.*

The following table describes the agent log message components and their descriptions.

Table 9-1 Log message components and their description

Log message component	Description
Timestamp	Denotes the date and time when the message was logged.
Mnemonic	Denotes which Symantec product logs the message. For Cluster Server, the mnemonic is 'VCS'.
Severity	<p>Denotes the severity of the message. Severity is classified into the following types:</p> <ul style="list-style-type: none"> ■ 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 the user that VCS has initiated an action. ■ INFO informs the user of various state messages or comments. <p>Among these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.</p>
UMI or Unique Message ID	<p>UMI 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 ExchProtocol agent is 20023 while that for ExchService agent is 20024. Message ID is a unique number assigned to the message text.</p>
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 written to the Windows event log.

A typical agent log resembles:

```
ExchService2007:d1-ExchService2007-MSEExchangeIS:online:Service
(MSEXCHANGEIS) is taking longer to start. Timeout = 10 seconds
```

VCS Cluster Configuration Wizard (VCW) logs

The VCS Cluster Configuration Wizard (VCW) log is located at

```
%allusersprofile%\Veritas\Cluster Server\vcw.log.
```

Here, %allusersprofile% is the file system directory containing application data for all users. A typical path is C:\ProgramData\.

The format of the wizard log is of the format *ThreadID | Message Text*.

ThreadID is the ID of the thread initiated by the wizard and Message Text is the actual message generated by the wizard.

A typical wizard log resembles the following:

```
00000576-00000264: ExecMethod return 00000000.
00000576-00000110: CRegistry::Query for VCS License failed.
Error=0x00000000
00000576-00000264: ExecMethod return 00000000.
00000576-00000264: ExecMethod return 00000001.
00000576-00000127: QueryDWORDValue returned 0x00000001
00000576-00000132: CRegistry::Query for VxSS Root information
failed. Error=0x00000001
```

VCWsilent logs

The VCWsilent log is located at <currentdirectory>\vcwsilent.log.

Here, <currentdirectory> is the directory from where the VCWsilent.exe is run.

A typical VCWsilent log resembles the following:

```
00005540-00000064: 5540: STARTING - Discovering NICs on the
selected machines...
00009956-00000064: 9956: STARTING - Generating private network
related files...
00009956-00000048: 9956: COMPLETED - Generating LLT host
files...
00009956-00000048: 9956: COMPLETED - Generating GAB tab files...
00009956-00000048: 9956: COMPLETED - Generating main.cf file...
00009956-00000064: 9956: STARTING - Configuring LLT on all the
nodes.
00009956-00000048: 9956: COMPLETED - Configuring LLT on all the
nodes.
```

NetApp agents error messages

The following table contains a list of error messages for the VCS agents for NetApp.

Table 9-2 NetApp agents error messages

Message	Description
Failed to open connection to filer %s.	<p>Make sure that the VCS Helper service account has a domain user and is part of the administrator's group on the local host and the filer.</p> <p>Make sure the private network is functioning properly. Verify you can ping the IP used for the private storage network. This is the IP defined the StorageIP attribute of the NetAppFiler resource.</p>
Failed to initialize ONTAPI on system	The agent could not find the file NTAPADMIN.DLL on the system. Verify the file exists in the %VCS_HOME%\bin directory
Invalid attributes exist in the configuration	Some agent attributes have not been defined or have been defined incorrectly. Verify the configuration definition for the agent.
ONTAP API called failed for object_name on filer_name.	The specified API failed on the specified object. See the NetApp ONTAP API documentation for information about the associated error message
Volume %s on filer %s is not a SnapMirror replicated volume	Verify replication is set up on the specified volume.
Multiple snapmirror destinations for a volume is not supported by this agent. 'snapmirror status' for volume %s on filer %s returned multiple status entries. Administrative intervention required	There should be only one destination per source volume.
Initialize VLibNetAppHost::Initialize() failed. (error_type: %s, error_code: 0x%s)	<p>The agent could not detect the iSCSI or the FC Initiator on the host.</p> <p>Make sure that you have installed and configured Microsoft iSCSI Initiator or an FC Initiator on each node.</p>

Table 9-2 NetApp agents error messages (*continued*)

Message	Description
Failed to connect/disconnect virtual disk. (error_type: %s, error_code: 0x%s, error_message: %s)	<p>This could occur because one or more of the following parameters are defined incorrectly in the VCS configuration:</p> <ul style="list-style-type: none"> ■ Filer name ■ Volume name/LUN name ■ Share name ■ Storage IP <p>Verify the configuration definition of the resource. Make sure each attribute is defined correctly.</p>
Unable to create/delete online lock file %s. Error code %s,	Make sure you have write permissions on the specified directory.

Exchange Service agent error messages

The following table describes the error messages for the Exchange Service agent.

Table 9-3 Exchange Service agent error messages

Message	Description
Failed to find the service object. Please check the 'Service' attribute.	<p>The value specified for the "Service" attribute is incorrect.</p> <p>Solution: Provide a valid value for the Lanman resource. If the value is correct, see error type and error code for further information.</p>
Failed to open the service object.(Service = service name). Error Type, Error Code.	<p>The agent failed to open the service object.</p> <p>Solution: See the associated Windows error type and error code for more information.</p>
Failed to get the state of the service (service name). Error Type, Error Code.	<p>The agent failed to retrieve the state of the service.</p> <p>Solution: See the associated Windows error type and error code for more information.</p>
Failed to start the service (service name) Error Type, Error Code.	<p>The agent failed to start the specified service.</p> <p>Solution: See the associated Windows error type and error code for more information.</p>

Table 9-3 Exchange Service agent error messages (*continued*)

Message	Description
Failed to stop the service (service name). Error Type, Error Code.	The agent failed to stop the service. 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. 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. Solution: Specify a valid value for the attribute.
Configuration error. 'LanmanResName' attribute is not configured.	No value specified for the "LanManResName" attribute. Solution: Specify a valid value for the attribute.
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. 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. The agent needs to set the environment block for starting the service in the context of the virtual server name. 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. While taking the resource offline, the agent stops the service and removes the environment block. 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. Solution: Specify a valid value for the attribute.

Table 9-3 Exchange Service agent error messages (*continued*)

Message	Description
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. Solution: Select the databases for the detail monitoring.
FaultOnMountFailure flag is true.\"Auto Mount\" on database: (database names) is enabled but database is dismantled. 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 dismantled. So the agent will fault the service group.
\"Auto Mount\" on database: (database names) is enabled but database is dismantled. Agent will return status as Unknown.	Databases that are set to mount automatically on startup are dismantled. 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. 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, complete the following steps to resolve the error.

To troubleshoot Microsoft Exchange uninstallation issues

- 1 Start the following Exchange services manually using the Service Control Manager:
 - MExchangeSA
 - MExchangeIS
- 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 Server Setup Wizard for VCS and select the **Remove Exchange** option.
Note that you must uninstall Exchange only by using the Exchange Server Setup Wizard for VCS.

About troubleshooting Exchange Setup Wizard issues

This section describes some of the issues related to the Exchange Setup Wizard for VCS.

Exchange Setup Wizard may fail to rename the node

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.

Exchange Setup Wizard may fail while moving the Exchange databases

When moving the Exchange databases to shared storage, the Exchange Setup Wizard may fail. This may happen if the DNS entry for the Exchange virtual server does not exist. Check if the DNS entry exists.

To resolve this issue, you must add the DNS entry manually.

Exchange database status appears as unknown

The status of the Exchange databases appears as unknown. This may happen even if the Microsoft Exchange Information Store (MSExchangeIS) service is running.

To resolve this issue, check if the DNS entry for the Exchange virtual server exists. If it does not, add the DNS entry manually.

Resource type definitions

This appendix includes the following topics:

- [About resource type definitions](#)
- [Exchange Service agent](#)
- [Detail monitoring and agent behavior](#)

About resource type definitions

This appendix lists the resource type definitions and attribute definitions of the agents. The resource type represents the VCS configuration definition of the agent and specifies how the agent is defined in the 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.

Exchange Service agent

The Exchange Service agent resource type definition and attribute definitions are as follows. This information will assist you during the agent configuration.

Exchange Service agent resource type definition

The Exchange Service agent is configured as a resource of type ExchService2007.

```
type ExchService2007 (
    static int IMF{} = { Mode=2, MonitorFreq=5, RegisterRetryLimit=3 }
    static i18nstr IMFRegList[] = { Service }
    static i18nstr ArgList[] = { Service, "LanmanResName:VirtualName",
    DetailMonitor, FaultOnMountFailure, DBList }
    str Service
```

```

    str LanmanResName
    int DetailMonitor
    boolean FaultOnMountFailure = 0
    i18nstr DBList[]
)

```

Exchange Service agent attribute definitions

The following table describes the Exchange Service agent required attributes.

Table A-1 Exchange Service agent required attributes

Required Attributes	Type - Dimension	Definition
Service	string-scalar	<p>The name of the Exchange service to be monitored. This attribute could take any of the following values:</p> <ul style="list-style-type: none"> ■ MExchangeADTopology ■ MExchangeIS ■ MExchangeMailSubmission ■ MExchangeSA ■ MExchangeMailboxAssistants ■ MExchangeServiceHost ■ MExchangeTransportLogSearch ■ MExchangeSearch ■ msftesql-Exchange ■ MExchangeMonitoring ■ MExchangeRepl
LanmanResName	string-scalar	The name of the LANMAN resource on which the ExchService resource depends.

The following table describes the Exchange Service agent optional attributes.

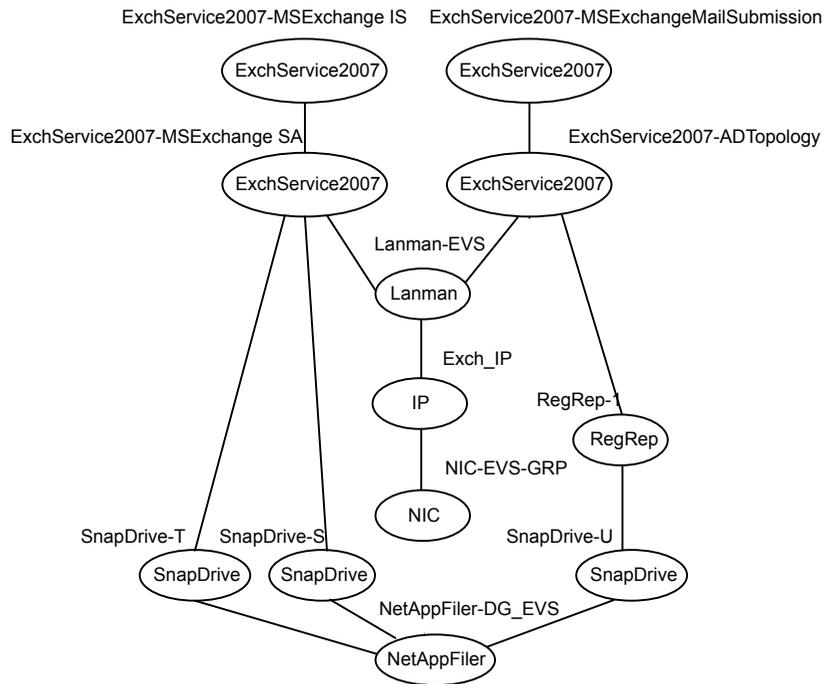
Table A-2 Exchange Service agent optional attributes

Optional Attribute	Type - Dimension	Definition
DetailMonitor	integer-scalar	<p>The interval at which the agent performs detail monitoring on the databases specified in the DBList attribute.</p> <p>The default value 5 indicates that the agent performs detail monitoring on every 5th monitor cycle.</p> <p>Setting this value to 0 disables detail monitoring.</p>
FaultOnMountFailure	boolean-scalar	<p>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.</p> <p>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.</p> <p>The default value is 0 (false).</p>
DBList	string-vector	List of databases for which the agent will perform detail monitoring.

Dependency graph for an Exchange local cluster configuration

The following figure illustrates the dependency graph of a VCS service group that has NetApp related resources.

Figure A-1 Local cluster configuration dependency graph



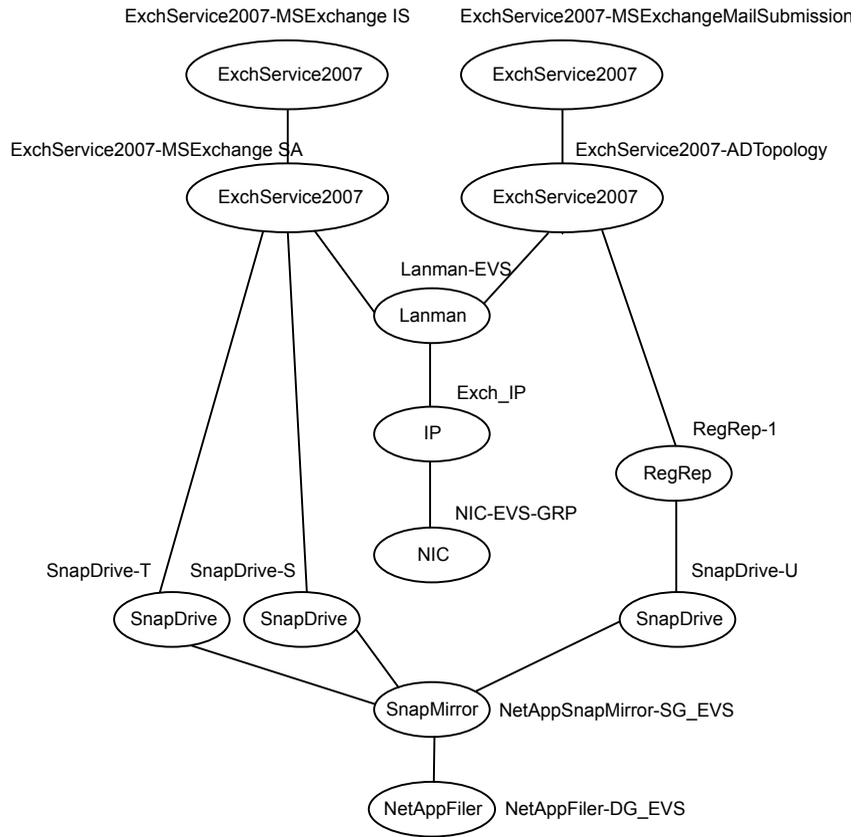
Note the following:

- Three LUNs are mounted as drives T, S, and U respectively, using the SnapDrive agent.
 The Registry Replication resource, configured to replicate the Exchange registry keys, is set up on LUN U of the cluster disk group. The Exchange database is installed on LUN S.
- The Exchange services (MSEExchangeSA, MSEExchangeIS, MSEExchangeADTopology, and MSEExchangeMailSubmission) 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.

Dependency graph for an Exchange disaster recovery configuration

The following figure illustrates the dependency graph for a VCS service group in a cluster that is a part of a global cluster.

Figure A-2 Disaster Recovery configuration dependency graph



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.

The following table 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 (Service group will fault)
		0 (False)	Unknown (Administrative action required)
No	Mounted	Does not matter	Online
No	Dismounted	Does not matter	Unknown (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.

Sample configuration

This appendix includes the following topics:

- [About Exchange sample configurations](#)
- [Sample configuration for an Exchange cluster \(local cluster configuration\)](#)

About Exchange 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 NetApp filers to manage shared storage.

The sample configuration graphically depicts the resource types, resources, and resource dependencies within the service group. The sample configuration files (main.cf) are also included for your reference.

See the *Cluster Server Bundled Agents Reference Guide*.

Sample configuration for an Exchange cluster (local cluster configuration)

A sample cluster configuration file is included as a reference.

```
include "types.cf"

cluster clc12 (
    SecureClus = 1
)

system VCSW2K287 (
```

```

Limits = { ExchLoad = 10 }
)

system VCSW2K288 (
    Limits = { ExchLoad = 10 }
)

group A2 (
    SystemList = { VCSW2K288 = 0, VCSW2K287 = 1 }
    Prerequisites = { ExchLoad = 10 }
)

ExchService2007 A2-ExchService2007-MSEExchangeSA (
    Service = MSEExchangeSA
    LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSEExchangeIS (
    Service = MSEExchangeIS
    LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSEExchangeMailSubmission (
    Service = MSEExchangeMailSubmission
    LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSEExchangeADTopology (
    Service = MSEExchangeADTopology
    LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSEExchangeMailboxAssistants (
    Service = MSEExchangeMailboxAssistants
    LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSEExchangeServiceHost (
    Service = MSEExchangeServiceHost
    LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSEExchangeTransportLogSearch (

```

Sample configuration for an Exchange cluster (local cluster configuration)

```
Service = MSExchangeTransportLogSearch
LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSExchangeSearch (
Service = MSExchangeSearch
LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-msftesql-Exchange (
Service = msftesql-Exchange
LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSExchangeMonitoring (
Service = MSExchangeMonitoring
LanmanResName = A2-Lanman
)

ExchService2007 A2-ExchService2007-MSExchangeRepl (
Service = MSExchangeRepl
LanmanResName = A2-Lanman
)

IP A2-IP (
Address = "10.121.59.48"
SubNetMask = "255.255.253.0"
MACAddress @VCSW2K288 = "00-11-43-DD-F4-36"
MACAddress @VCSW2K287 = "00-11-43-DD-F4-80"
)

Lanman A2-Lanman (
VirtualName = ANY2
IPResName = A2-IP
DNSUpdateRequired = 1
ADUpdateRequired = 1
ADCriticalForOnline = 1
)

NIC A2-NIC (
MACAddress @VCSW2K288 = "00-11-43-DD-F4-36"
MACAddress @VCSW2K287 = "00-11-43-DD-F4-80"
)
```

Sample configuration for an Exchange cluster (local cluster configuration)

```

NetAppFiler A2-NetAppFiler (
  FilerName = winvm-netapp02
  StorageIP = "10.121.57.8"
)

NetAppSnapDrive A2-NetAppSnapDrive (
  FilerResName = A2-NetAppFiler
  VolumeName = vol8
  ShareName = data08
  LUN = "d85.lun"
  MountPath = T
  Initiator @VCSW2K288 = { "21:00:00:e0:8b:1c:cc:45",
    "21:00:00:e0:8b:90:74:1a" }
  Initiator @VCSW2K287 = { "21:00:00:e0:8b:0e:3a:6a" }
)

NetAppSnapDrive A2-NetAppSnapDrive-1 (
  FilerResName = A2-NetAppFiler
  VolumeName = vol8
  ShareName = data08
  LUN = "d84.lun"
  MountPath = S
  Initiator @VCSW2K288 = { "21:00:00:e0:8b:1c:cc:45",
    "21:00:00:e0:8b:90:74:1a" }
  Initiator @VCSW2K287 = { "21:00:00:e0:8b:0e:3a:6a" }
)

RegRep A2-RegRep (
  MountResName = A2-NetAppSnapDrive
  ReplicationDirectory = "\\VCS\\Private\\RegRep\\Exch"
  Keys =
  {"HKLM\\SYSTEM\\CurrentControlSet\\Services\\
MSExchangeADTopology" = "",
  "HKLM\\SYSTEM\\CurrentControlSet\\Services\\
MSExchangeIS" = "",
  "HKLM\\SYSTEM\\CurrentControlSet\\Services\\
MSExchangeMailSubmission" = "",
  "HKLM\\SYSTEM\\CurrentControlSet\\Services\\
MSExchangeMailboxAssistants" = "",
  "HKLM\\SYSTEM\\CurrentControlSet\\Services\\
MSExchangeMonitoring" = "",
  "HKLM\\SYSTEM\\CurrentControlSet\\Services\\

```

Sample configuration for an Exchange cluster (local cluster configuration)

```

MSExchangeRepl" = "",
"HKLM\SYSTEM\CurrentControlSet\Services\
MSExchangeSA" = "",
"HKLM\SYSTEM\CurrentControlSet\Services\
MSExchangeSearch" = "",
"HKLM\SYSTEM\CurrentControlSet\Services\
MSExchangeServiceHost" = "",
"HKLM\SYSTEM\CurrentControlSet\Services\
MSExchangeTransportLogSearch" = "",
"HKLM\SYSTEM\CurrentControlSet\Services\
msftesql-Exchange" = "" }
RestoreLocally = 1
)

```

```

A2-RegRep requires A2-NetAppSnapDrive
A2-NetAppSnapDrive requires A2-NetAppFiler
A2-Lanman requires A2-IP
A2-NetAppSnapDrive-1 requires A2-NetAppFiler
A2-ExchService2007-MSExchangeSA requires A2-RegRep
A2-ExchService2007-MSExchangeSA requires A2-Lanman
A2-ExchService2007-MSExchangeSA requires A2-NetAppSnapDrive-1
A2-ExchService2007-MSExchangeIS requires
A2-ExchService2007-MSExchangeSA
A2-ExchService2007-MSExchangeMailSubmission requires
A2-ExchService2007-MSExchangeADTopology
A2-ExchService2007-MSExchangeADTopology requires A2-RegRep
A2-ExchService2007-MSExchangeADTopology requires A2-Lanman
A2-ExchService2007-MSExchangeMailboxAssistants requires
A2-ExchService2007-MSExchangeADTopology
A2-ExchService2007-MSExchangeServiceHost requires
A2-ExchService2007-MSExchangeADTopology
A2-ExchService2007-MSExchangeTransportLogSearch requires
A2-ExchService2007-MSExchangeADTopology
A2-ExchService2007-MSExchangeSearch requires
A2-ExchService2007-msftesql-Exchange
A2-ExchService2007-msftesql-Exchange requires
A2-ExchService2007-MSExchangeADTopology
A2-ExchService2007-MSExchangeMonitoring requires
A2-ExchService2007-MSExchangeADTopology
A2-ExchService2007-MSExchangeRepl requires
A2-ExchService2007-MSExchangeADTopology
A2-IP requires A2-NIC

```