

Cluster Server 7.4.1 Implementation Guide for Microsoft Exchange 2010 - Windows

Last updated: 2019-02-01

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas 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. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introducing the VCS agents for Exchange and NetApp	7
	About VCS support for Exchange 2010 and NetApp	7
	About the VCS database agent for Exchange 2010	8
	Agent functions	9
	Agent state definitions	10
	About the VCS hardware replication agent for NetApp	10
	About the NetApp Filer agent	10
	About the NetApp SnapDrive agent	11
	About the NetApp SnapMirror agent	13
	How the VCS agents make Exchange highly available	17
	Local cluster configuration	18
	Disaster recovery configuration	18
	Typical Exchange configurations in a VCS cluster	18
	Active-Active failover configuration	18
	Disaster recovery configuration for Exchange	19
Chapter 2	Installing the product and configuring a VCS cluster	21
	About installing the VCS agents	21
	Configuring the cluster using the Cluster Configuration Wizard	21
	Configuring notification	30
	Configuring Wide-Area Connector process for global clusters	33
Chapter 3	Installing Microsoft Exchange Server 2010	36
	About installing Exchange 2010 in a VCS environment	36
	Before you install Exchange 2010	37
	Configuring Microsoft iSCSI initiator	37
	Managing storage using NetApp filer	37
	Connecting virtual disks to the cluster node	39
	Disconnecting virtual disks from the cluster nodes	39
	Managing storage using Windows Logical Disk Manager	40
	Reserving disks (if you use Windows LDM)	41

	Creating volumes (if you use Windows LDM)	41
	Mounting volumes (if you use Windows LDM)	42
	Unassigning a drive letter	42
	Releasing disks (if you use Windows LDM)	43
	Installing Exchange Server 2010	43
	Creating mailbox databases on shared storage	44
Chapter 4	Configuring the Exchange database service group	45
	About configuring the Exchange service group	45
	Prerequisites for configuring the Exchange database service group	46
	Creating the Exchange database service group	47
	Running SnapManager for Exchange	50
	About the modifications required for tagged VLAN or teamed network	51
	About verifying the service group configuration	52
	Bringing the service group online	52
	Taking the service group offline	52
	Switching the service group	52
	About modifying the Exchange database service group configuration	53
	Prerequisites for modifying an Exchange database service group	53
	Modifying the Exchange database service group	54
	Deleting the Exchange service group	54
Chapter 5	Making a standalone Exchange server highly available	55
	High availability configuration for a standalone server	55
	Moving mailbox databases to shared storage	56
Chapter 6	Deploying Disaster Recovery for Exchange Server	57
	About disaster recovery configuration	57
	Setting up disaster recovery configuration	57
	Configure replication using NetApp SnapMirror	60
	Configure NetAppSnapMirror resources at the primary site	60
	Configure NetAppSnapMirror resources at the secondary site	61

Chapter 7	Removing the software components	63
	About removing the software components	63
	Remove Microsoft Exchange	64
	Removing a node without removing Microsoft Exchange	64
	Removing a node and removing Microsoft Exchange	65
Chapter 8	Troubleshooting	67
	About troubleshooting VCS agents for Exchange	67
	VCS logging	67
	VCS Cluster Configuration Wizard (VCW) logs	68
	VCWsilent logs	69
	NetApp agents error messages	69
Appendix A	Resource type definitions	72
	About resource type definitions	72
	Exchange database agent	72
	Exchange 2010 database agent resource type definition	72
	Exchange 2010 database agent attribute definitions	73
	Dependency graph for an Exchange cluster	73
Appendix B	Sample Configurations	75
	About the sample configurations	75
	Sample service group configuration	75

Introducing the VCS agents for Exchange and NetApp

This chapter includes the following topics:

- [About VCS support for Exchange 2010 and NetApp](#)
- [About the VCS database agent for Exchange 2010](#)
- [About the VCS hardware replication agent for NetApp](#)
- [How the VCS agents make Exchange highly available](#)
- [Typical Exchange configurations in a VCS cluster](#)

About VCS support for Exchange 2010 and NetApp

VCS support for Exchange Server 2010 includes high availability for Exchange 2010 mailbox databases.

VCS provides a database agent for Exchange 2010 that monitors the mailbox databases configured on shared storage. You must install the Exchange 2010 Mailbox Server role to allow VCS to make the databases highly available. The agent internally monitors a critical set of Exchange 2010 services to verify the availability of the Mailbox Server and the configured databases.

VCS however does not provide high availability support for public folders. VCS also does not provide high availability support for mailbox databases that are configured

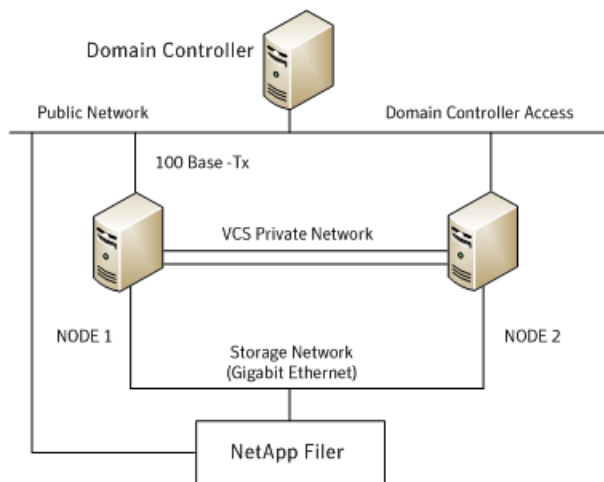
in an Exchange 2010 Data Availability Group (DAG). If you wish to make those databases highly available with VCS, you must first remove them from the DAG.

The VCS 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 databases 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 72.

About the VCS database agent for Exchange 2010

The agent provides “Active-Active” support for Exchange 2010 wherein a single cluster node can host multiple Exchange 2010 service groups at the same time, if required.

The VCS Exchange 2010 Database agent (Exch2010DB) monitors the Exchange mailbox databases, brings them online and takes them offline. The agent also starts the following Exchange services if they are not running already, and monitors their status:

- Microsoft Exchange System Attendant (MSEExchangeSA):
The Exchange component responsible for monitoring, maintenance, and Active Directory lookup services, and ensuring that operations run smoothly.
- Microsoft Exchange Information Store (MSEExchangeIS):
The Exchange storage used to hold messages in users' mailboxes and in public folders.
- Microsoft Exchange Mail Submission (MSEExchangeMailSubmission):
This service submits messages from the Mailbox Server to the Hub Transport Server.

The agent internally monitors these services; the Exchange 2010 service group does not contain separate resources for these services.

Agent functions

The agent functions include the following:

- | | |
|---------|--|
| Online | <ul style="list-style-type: none"> ■ Checks if the mailbox database file is available on the configured shared volume. ■ Checks the status of the Microsoft Exchange Information Store (MSEExchangeIS) service and starts the service if it is not running. ■ Updates the Windows Active Directory (AD) to bind the Exchange mailbox database to the Exchange Mailbox Server. ■ Mounts the Exchange mailbox database on the node. |
| Offline | Dismounts the Exchange mailbox database from the node. |
| Monitor | <ul style="list-style-type: none"> ■ Verifies the status of the mailbox database on the node. If the database is mounted, the agent reports the resource as ONLINE. If the database is dismounted, the agent resource is marked as OFFLINE. ■ If the agent is unable to retrieve the database status, the agent queries the Service Control Manager (SCM) for the status of the Microsoft Exchange Information Store (MSEExchangeIS) service. If the service is running, the agent reports the resource as UNKNOWN; otherwise the resource is marked as OFFLINE. |
| Clean | Forcibly dismounts the Exchange mailbox database from the node. |

Agent state definitions

The agent state definitions are as follows:

ONLINE	Indicates that the configured mailbox database is mounted and active on the cluster node.
OFFLINE	Indicates that the configured mailbox database is dismounted from the cluster node.
UNKNOWN	Indicates that the agent is unable to determine the status of the configured mailbox database on the cluster node.

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
- Performs the following tasks:
- 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	DNS-resolvable name or IP address of the locally attached filer. Type and dimension: string-scalar
StorageIP	The private storage IP address of the filer. Type and dimension: string-scalar

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

Attribute	Description
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 break-off 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	Resynchronizes an original source volume with a broken-off volume. After synchronization, the original source volume becomes the target volume. The broken-off volume was initially the target volume, but was broken off as a result of a take over.

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"> ■ async: Indicates that the mirror should be configured in the asynchronous mode. ■ semi-sync: Indicates that the mirror should be configured in the semi-synchronous mode. ■ sync: Indicates that the mirror should be configured in the synchronous mode. <p>The default is async (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 VCS agents make Exchange highly available

The VCS database agent for Exchange detects a failure in case of a system failure or if the Exchange Mailbox Server becomes unavailable on a node. VCS moves the mailbox databases configured on that node to the next available cluster node in the service group's system list. The agent also starts the critical Exchange 2010 services on that node, if required. The databases then become active on the new node. The client requests are then handled by the Mailbox Server on the new node, thus maintaining continuous availability of the Exchange 2010 mailbox databases.

The following topics describe 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 service groups 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)

Typical Exchange configurations in a VCS cluster


The VCS database agent for Exchange supports the Active-Active and Disaster Recovery type of configuration.

Active-Active failover configuration

In an Active-Active configuration, a single cluster node can host multiple Exchange 2010 database service groups.

The following figure illustrates a two node active-active configuration. The Exchange Server databases are configured on the shared storage on volumes or LUNs, such that each Exchange Server database is configured in a separate service group and each of this service group can failover to other node in the cluster.

Figure 1-2 Active-Active failover configuration




Disaster recovery configuration for Exchange

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 1-3 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:

<https://sort.veritas.com/documents>

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 (NICs), and SCSI controllers are installed and connected to each system.
Veritas recommends the following actions for network adapters:
 - Disable the ethernet auto-negotiation options on the private NICs to prevent:
 - Loss of heartbeats on the private networks
 - VCS from mistakenly declaring a system as offlineContact the NIC manufacturer for details on this process.
 - Remove TCP/IP from the 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.
- Veritas 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. Veritas 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.
- In an any-to-any configuration, you can add the systems for all the Exchange servers when creating the cluster, not only the first Exchange server.

To configure a VCS cluster using the wizard

- 1 Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > 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**.

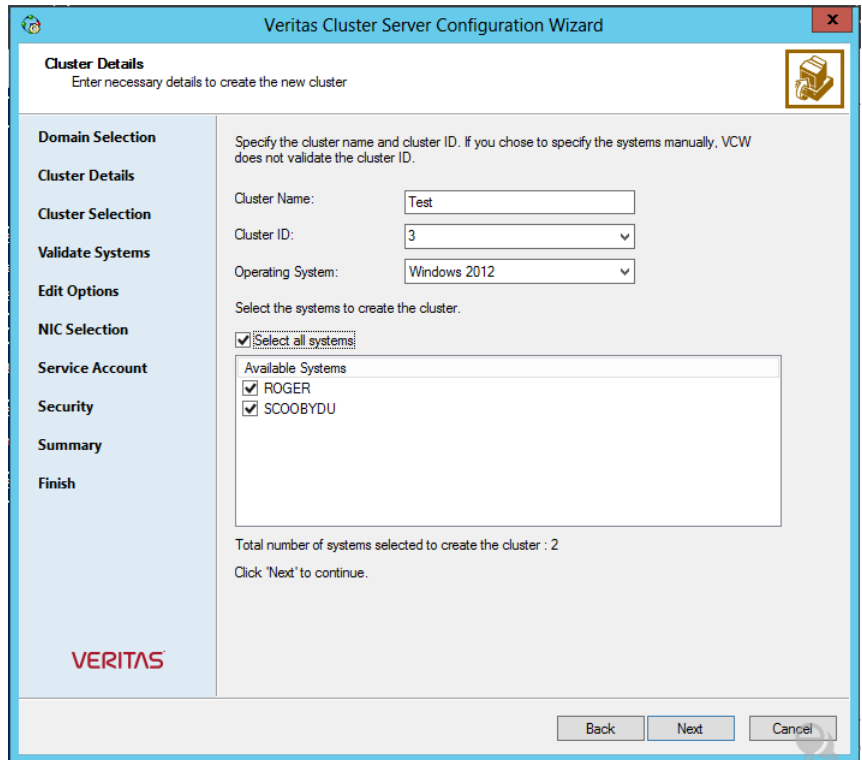
Select the system to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

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. Veritas 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. |

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. Veritas 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, Veritas 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. Veritas 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. Veritas 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.
Veritas 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 30.

- 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 33.

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.

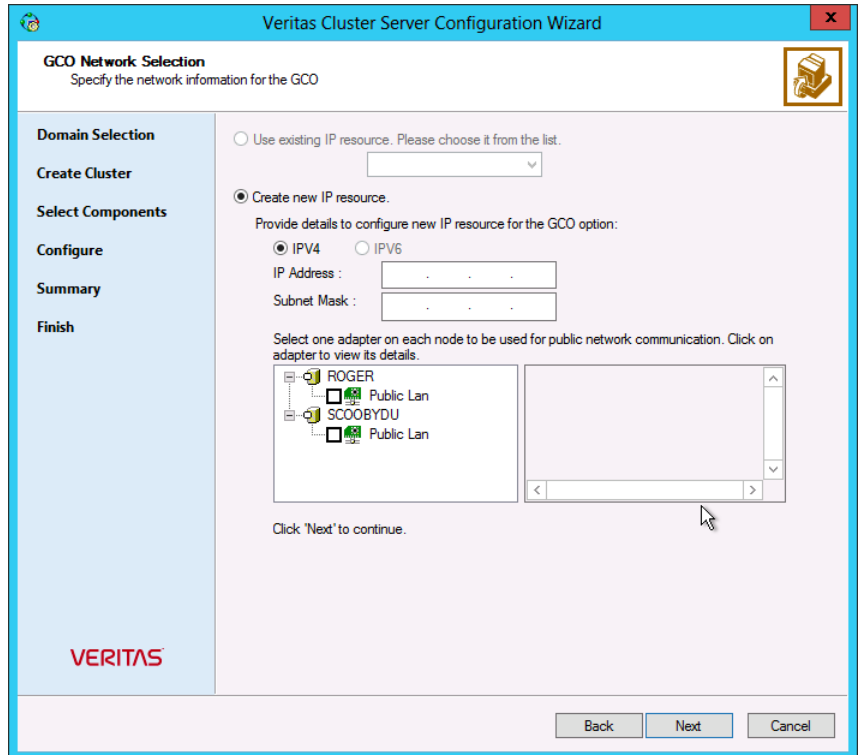
Configuring Wide-Area Connector process for global clusters

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 Server 2010

This chapter includes the following topics:

- [About installing Exchange 2010 in a VCS environment](#)
- [Before you install Exchange 2010](#)
- [Managing storage using NetApp filer](#)
- [Managing storage using Windows Logical Disk Manager](#)
- [Installing Exchange Server 2010](#)

About installing Exchange 2010 in a VCS environment

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

Note: In case of Exchange 2010, the unit of failover is the Exchange mailbox database. The VCS concept of Exchange virtual server (EVS) is not applicable. Hence, you do not need to install Exchange on a virtual server name to facilitate high availability.

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

See [“High availability configuration for a standalone server”](#) on page 55.

Before you install Exchange 2010

Verify the following prerequisites before you install Exchange in an InfoScale environment:

- Ensure that the appropriate InfoScale product is installed and the cluster is configured.
- Ensure that the systems meet the minimum requirements for installing and configuring Exchange.
For details, refer to the Microsoft documentation.
- Ensure that Exchange is installed on the same local drive on all the InfoScale cluster nodes. Make sure that the same drive letter is available on all the nodes and that the drive has adequate space for the installation.

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.

To configure Microsoft iSCSI initiator on a Windows Server system

- 1 Start the Microsoft iSCSI initiator.
- 2 On the Discovery tab, click **Discover Portal**.
- 3 On the Discover Target Portal dialog box, specify the DNS name for the NetApp filer and then click **OK**.
- 4 On the Target tab, click **Connect**.
- 5 On the Connect to Target dialog box, clear the **Add this connection to list of Favorite Targets** check box and then click **Ok**.
- 6 On the Targets tab, verify that the newly added portal is listed under the **Select a target** box and the status appears as 'connected.' Then, click **OK**.

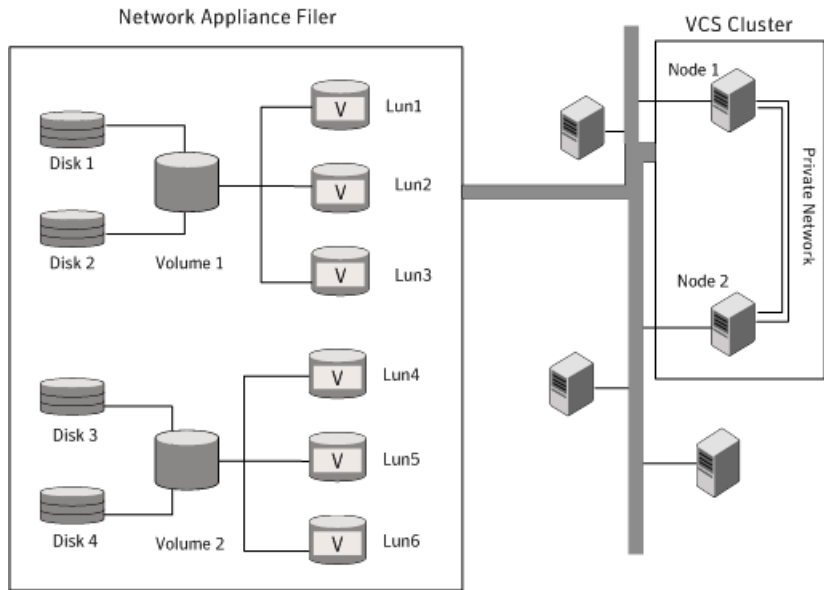
Managing storage using NetApp filer

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



Veritas 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.

Veritas 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 41.

- Create volumes
See “[Creating volumes \(if you use Windows LDM\)](#)” on page 41.
- Mount volumes
See “[Mounting volumes \(if you use Windows LDM\)](#)” on page 42.
- Unassign the drive letter
See “[Unassigning a drive letter](#)” on page 42.
- Release the disks
See “[Releasing disks \(if you use Windows LDM\)](#)” on page 43.

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 43.

- 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:\>havol -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:\>havol -scsitest /REL:<disk #>
```

For example, to release disk 4, type:

```
C:\>havol -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 Server 2010

Run Microsoft Exchange 2010 Setup to install Exchange on each cluster node where you want to configure the Exchange service group. Refer to the Microsoft documentation for instructions.

Note the following requirements while installing Exchange:

- Ensure that you select to install the Mailbox server role (Mailbox Role check box). You can also install the other Exchange server roles as per your requirement, but you must install the Exchange Mailbox server role.
- Install the Exchange Server program files to the local system disks. You do not need to install the Exchange Server program files to a shared location.
- Install the Exchange Management Tools on the systems where you install Exchange mailbox server role or on other systems if you wish to manage Exchange remotely.

Creating mailbox databases on shared storage

After installing Exchange on the required servers, create the Exchange mailbox databases on shared storage. Use the Exchange Management Console or the Exchange Management Shell to create the databases.

Verify the following before you create mailbox databases:

- Map the database LUNs and volumes required for the Exchange mailbox databases and log files on the system from where you will create the databases. See [“Managing storage using NetApp filer”](#) on page 37.
- Ensure that the database files and log files are on the separate luns but are a part of the same volume on the filer.
- Ensure that you have the permissions required to create mailbox databases. Refer to the Microsoft documentation for details.
- While creating the mailbox databases, ensure that the path you specify for the mailbox database and the log file is on the shared disk.

Configuring the Exchange database service group

This chapter includes the following topics:

- [About configuring the Exchange service group](#)
- [Prerequisites for configuring the Exchange database service group](#)
- [Creating the Exchange database service group](#)
- [About verifying the service group configuration](#)
- [About modifying the Exchange database service group configuration](#)
- [Deleting the Exchange service group](#)

About configuring the Exchange service group

Configuring the Exchange database service group involves creating the Exchange service group and its resources and then defining the attribute values for the configured resources.

See [“Creating the Exchange database service group”](#) on page 47.

See [“Exchange 2010 database agent attribute definitions”](#) on page 73.

A VCS Exchange database service group is used to move the configured Exchange mailbox databases between the mailbox servers configured in the service group. If an active node fails, the mailbox databases are moved to an alternate system in the service group thus ensuring continuous availability.

See [“About modifying the Exchange database service group configuration”](#) on page 53.

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 72.

Prerequisites for configuring the Exchange database 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 21.
- 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
See [“Managing storage using NetApp filer”](#) on page 37.
See [“Managing storage using Windows Logical Disk Manager”](#) on page 40.
- 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.

Creating the Exchange database service group

Use the Exchange 2010 Configuration Wizard to configure the Exchange database service groups.

To configure the Exchange service group

- 1 Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Exchange 2010 Configuration Wizard**.

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

- 2 Review the prerequisites on the Welcome panel and then click **Next**.
- 3 On the Service Group Options panel, click **Configure database service groups** and then click **Next**.
- 4 On the Exchange Database Selection panel, select the databases that you wish to configure for high availability.
 - The Databases box displays the databases discovered on the local system. Databases that reside on shared storage and are not part of another service group are available for selection.
 - Select a database and then select the **Make the database highly available** check box.
 - When you select a database, the wizard automatically selects all the other databases that reside on the same volume. Even if you select just one database, the wizard configures all the other databases residing on that volume. Thus, all databases that reside on the same volume are part of the same Exchange service group.
 - The wizard configures one service group per volume. If you select databases that reside on two separate volumes, the wizard configures them in two separate service groups.
 - If databases are created on a volume that is already a part of an existing Exchange service group, then the wizard automatically adds the newly created databases to that existing service group. This happens even if you do not explicitly select those new databases. Thus, even while creating a service group, the wizard modifies existing service groups automatically.

You can also run the wizard in the modify mode to add newly created databases to existing service groups.

Note that the wizard is able to configure the new databases in existing service groups only if the corresponding volumes are mounted on the node where the wizard is running.

- Click **Select All** if you wish to select all the available databases. The wizard configures all the eligible databases in the service group.
 - Click **Next**.
- 5** On the Exchange Service Group Configuration panel, specify the service group name and then click **Next**.

The Service Groups box displays the default name that the wizard assigns to the service group. The default service group name is of the format EXCHSG_<tablename>. Here, <tablename> is the name of the database on the respective volume. To specify another name, select the service group in the Service Groups list and then type a name in the **Service Group Name** field.

You can specify a name only to newly created service groups. You cannot edit names of Exchange service groups already configured in the cluster.

- 6** For disaster recovery, if you have configured replication using NetApp SnapMirror, select the **Configure NetApp SnapMirror Resource(s)** checkbox.
- 7** On the System Selection panel, specify the systems that will be part of the service group and then click **Next**.
- In the Available Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the Selected Systems box. The Selected Systems list represents the service group's system list.
 - To remove a system from the service group's system list, select the system in the Selected Systems box and click the left arrow.
 - To change a system's failover priority in the service group's system list, select the system in the Selected Systems list 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.
 - If you have selected databases that reside on different volumes, the wizard creates multiple service groups, one for each volume. In such a case, the systems that you select here are configured in the system list of all those service groups. You cannot choose systems on a per service group basis while creating the service groups.

To modify the service group system list, run the wizard again in the modify mode and then define the system list for each of the service groups.

- 8 On the Initiator Selection panel, select the desired initiator for each of the selected system. Click **Next**

The initiator enables to you to connect to the NetApp filer. The selected initiator is configured as an attribute of the NetApp Snapdrive resource.

- 9 On the Network Adapter Selection panel, select a public network adapter for each system in the service group and then click **Next**.

To select a public adapter, click the **Adapter Display Name** field and then select an adapter from the drop-down list.

The selected adapter is used to detect network failures on the configured system.

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

Note: If you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that share 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 51.

- 10 On the Service Group Summary panel, review the service group configuration summary, change the resource names, if desired, and then click **Next**.

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

The wizard assigns unique names to resources. Change names of resources, if desired. To edit a resource name, select the resource name and either click it or press the **F2** key. Press **Enter** after editing each resource name. To cancel editing a resource name, press the **Esc** key.

- 11 Click **Yes** on the dialog box that prompts you that the wizard will modify the configuration.

The wizard runs command to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

- 12 On the Completing the Exchange 2010 Database Configuration Wizard panel, select **Bring the service group online** check box to bring the service group online on the local system and then click **Finish**.

Sometimes the wizard may fail to bring the service group online. In such a case, you must probe the resources and bring the service group online manually. You can use the Cluster Manager (Java Console) to perform the tasks.

After creating service groups, if you create fresh mailbox databases on the same volume, then you must run the wizard again (in the configure or modify mode) to configure the newly added databases for high availability.

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 service group 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.
- Ensure that the database files and log files are on the separate luns but are a part of the same volume on the filer.
- If you are restoring database backup on a different exchange server, then make sure that you select **Restore backup created on a different server**
- Ensure that you bring the database service group online and freeze it, before you begin to restore the database. This prevents the service group to fault and failover to another node, during the restore process.

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

The following topics provide 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 database 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 database 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.
- 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 database 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 Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Exchange Server 2010 Configuration Wizard**

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

- 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 45.

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 Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Exchange Server 2010 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:

- [High availability configuration for a standalone server](#)
- [Moving mailbox databases to shared storage](#)

High availability configuration for a standalone server

A “standalone” Exchange server is an Exchange server that is already deployed in a messaging environment but is not configured for high availability. You can convert an existing standalone Exchange server into a clustered Exchange server in a VCS environment.

The following table outlines the high-level objectives and the tasks to complete each objective for converting an existing standalone Exchange Server for high availability.

Table 5-1 Task list: Standalone Exchange server HA configuration tasks.

Task	Description
Verify hardware and software requirements	See <i>Veritas InfoScale Installation and Upgrade Guide</i>
Configure storage hardware and network	See “Managing storage using NetApp filer” on page 37.

Table 5-1 Task list: Standalone Exchange server HA configuration tasks.
(continued)

Task	Description
Install InfoScale Availability	See <i>Veritas InfoScale Installation and Upgrade Guide</i>
Add the standalone server to the VCS cluster	Run the Veritas Cluster Server Configuration Wizard and use the Add Node option to add the standalone server to the VCS cluster. See <i>Cluster Server Administrator's Guide</i>
Move the Exchange mailbox databases to shared storage	See "Moving mailbox databases to shared storage" on page 56.
Install and configure Exchange on the additional nodes	See "About installing Exchange 2010 in a VCS environment " on page 36.
Create the Exchange database service group(s)	See "Creating the Exchange database service group" on page 47.

Moving mailbox databases to shared storage

If the databases and log files do not reside on shared storage, you must move them to the database volume and log LUN that is connected to the standalone Exchange server. This ensures proper failover operations in the cluster.

Use the Exchange Management Console to move the mailbox databases and log files to shared storage. Refer to the Microsoft documentation for instructions.

To move the existing mailbox databases to shared storage

- 1 Verify that you have backed up your existing data.
- 2 Create and mount the database and log LUNs on the nodes where the original Exchange mailbox databases and log files reside.
- 3 Use the Exchange Management Console to move the Exchange mailbox database and log files to the shared storage.

Refer to the Microsoft documentation for instructions.

Note: Ensure that you move the databases and log files to separate LUNs but on the same Volume.

Deploying Disaster Recovery for Exchange Server

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](#)
- [Configure NetAppSnapMirror resources at the secondary site](#)

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 6-1 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 2010 in a VCS environment " on page 36. ■ Configure the Exchange service group See "About configuring the Exchange service group" on page 45.
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 2010 in a VCS environment " on page 36.

Table 6-1 Tasks for configuring a disaster recovery setup (*continued*)

Tasks	Description
Back up and restore the Exchange data files	<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.
Configure the Exchange service group at the secondary site	<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 45.</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>
Configure replication using NetApp SnapMirror	<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 6-1 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 60.
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.

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.

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 Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Exchange 2010 Configuration Wizard**.
On Windows 2012 operating systems, use the **Apps** menu.
- 3 Review the information on the Welcome panel and click **Next**.
- 4 On the Service Group Options panel, click **Modify database service group**, select the service group to be modified, and click **Next**.
- 5 On the Exchange Database Selection, click **Next**.
- 6 On the Exchange Service Group Configuration panel, select **Configure the NetApp SnapMirror resource(s)** check box and then click **Next**.
Click **Yes** on the warning that appears.
- 7 Accept default values in the subsequent dialog boxes and click **Next** till you reach the wizard completion panel.
- 8 On the Completing the Exchange 2010 Database Configuration Wizard panel, make sure the **Bring the service group online** check box is selected. Click **Finish**.

Configure NetAppSnapMirror resources at the secondary site

You must configure the NetAppSnapMirror resources at the secondary site, to ensure the following:

- 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.

Note: Ensure that the resource is identical to the one on the primary site, including the case sensitive name and the dependencies.

Use the Java Console to configure the NetAppSnapMirror resource.

To configure the NetAppSnapMirror resource

- 1** In the tree view, select the service group to which you want to add the resource, right-click and select **Add Resource**.
- 2** On the Add Resource panel, enter the resource name in the **Resource Name** text box.
- 3** From the Resource Type drop-down list, select **NetAppSnapMirror**.
- 4** Edit resource attributes according to your configuration.
- 5** Select the **Critical** and **Enabled** check boxes, if applicable. The Critical option is selected by default.

A critical resource indicates the service group is faulted when the resource, or any resource it depends on, faults. An enabled resource indicates agents monitor the resource; you must specify the values of mandatory attributes before enabling a resource. If a resource is created dynamically while VCS is running, you must enable the resource before VCS monitors it. VCS will not bring a disabled resource nor its children online, even if the children are enabled.

- 6** Click **Ok**.
- 7** Save and close the configuration.

Removing the software components

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 64.
- 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 65.
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 Review the information in the Welcome panel and click **Next**.
- 2 In the Available Option panel, click **Configure/Remove highly available Exchange Server** and click **Next**.
- 3 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.

- 4 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.

- 5 Once all the commands are executed, click **Next**.
- 6 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 Review the information in the Welcome panel and click **Next**.
- 2 In the Available Option panel, click **Configure/Remove highly available Exchange Server** and click **Next**.
- 3 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.
- 4 Select the Exchange virtual server for which you are removing the failover node 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 running commands to set up the VCS environment for removing the node from the Exchange service group.
- 6 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.
- 7 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.

- 8** Restarting the node automatically launches the Exchange Server Setup Wizard for VCS. Review the information in the Welcome dialog box 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** In the Microsoft Exchange Installer Welcome panel, read the welcome information and click **Next**.
- 11** 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.

- 12** Once all the tasks are complete, click **Next**.
- 13** 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

This chapter includes the following topics:

- [About troubleshooting VCS agents for Exchange](#)
- [VCS logging](#)
- [VCS Cluster Configuration Wizard \(VCW\) logs](#)
- [VCWsilent logs](#)
- [NetApp agents error messages](#)

About troubleshooting VCS agents for Exchange

The following topics describe how to troubleshoot common problems in the VCS agents for NetApp and Microsoft Exchange. They list the error messages and describe the problems 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 8-1 Log message components and their description

Log message component	Description
Timestamp	Denotes the date and time when the message was logged.
Mnemonic	Denotes which Veritas 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.

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 8-2 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 8-2 NetApp agents error messages (*continued*)

Message	Description
<p>Failed to connect/disconnect virtual disk. (error_type: %s, error_code: 0x%s, error_message: %s)</p>	<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>
<p>Unable to create/delete online lock file %s. Error code %s,</p>	<p>Make sure you have write permissions on the specified directory.</p>

Resource type definitions

This appendix includes the following topics:

- [About resource type definitions](#)
- [Exchange database agent](#)

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 database agent

This section provides the resource type definition and attribute definitions for the VCS agent for Exchange Server 2010.

Exchange 2010 database agent resource type definition

```
type Exch2010DB (
    static int IMF{} = { Mode=2, MonitorFreq=1, RegisterRetryLimit=3 }
    static i18nstr ArgList[] = { DBName, MonitorService }
    i18nstr DBName
    boolean MonitorService = 1
)
```


Exchange 2010 database agent attribute definitions

The following table describes the Exchange 2010 database agent required attributes

Table A-1 Exchange 2010 database agent required attributes

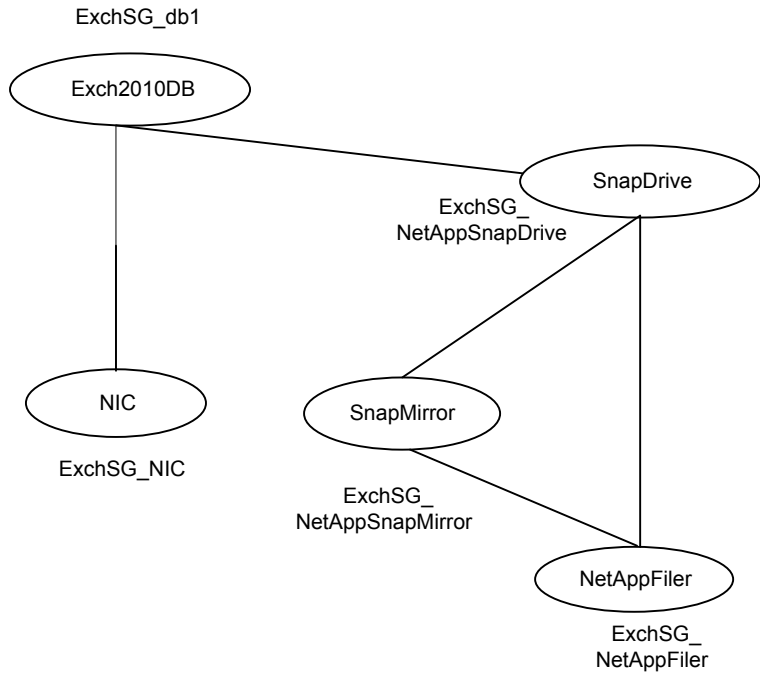
Required Attributes	Type-Dimension	Description
DBName	string-scalar	Name of the Exchange 2010 mailbox database to be made highly available.
MonitorService	boolean-scalar	<p>Defines whether the agent should monitor the critical Exchange 2010 services.</p> <p>The value 1 (True) indicates that the agent monitors the critical services. The value 0 (False) indicates that it does not.</p> <p>Default is 1 (True).</p> <p>If this attribute is set to 1 (True), the agent monitors the following Exchange 2010 services internally:</p> <ul style="list-style-type: none"> ■ Microsoft Exchange System Attendant (MSEExchangeSA) ■ Microsoft Exchange Mail Submission (MSEExchangeMailSubmission) <p>Note: You cannot define which Exchange 2010 services should be monitored by the agent.</p>

Dependency graph for an Exchange cluster

The following diagram illustrates a typical Exchange 2010 database service group configured to make Exchange 2010 mailbox databases highly available in a VCS cluster. The dependency graph depicts the resource types, resources, and resource dependencies within the service group. Review the dependency carefully before configuring the agent.

The following figure shows the dependencies in the Exchange 2010 database service group.

Figure A-1 shows the dependencies in the Exchange 2010 database service group.



Sample Configurations

This appendix includes the following topics:

- [About the sample configurations](#)
- [Sample service group configuration](#)

About the 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.

Sample service group configuration

A sample Exchange cluster service group configuration file (main.cf) is as follows:

```
include "types.cf"
cluster Exch2010 (
  UserNames = { a = aLLf }
  Administrators = { a }
)
system System1 (
)
group EXCHSG_SG1 (
  SystemList = { System1 = 0 }
)
```

```
Exch2010DB db1-Exch2010DB (
  DBName = db1
)
NIC EXCHSG_SG1-NIC (
  MACAddress @System1 = 00-11-43-DD-F4-80
)
NetAppFiler EXCHSG_SG1-NetAppFiler (
  FileName = vcsnetapp1
  StorageIP = "10.217.19.140"
)
NetAppSnapDrive EXCHSG_SG1-NetAppSnapDrive (
  FilerResName = EXCHSG_SG1-NetAppFiler
  VolumeName = Vol1
  ShareName = DEPRECIATED_VALUE
  LUN = "EXch.lun"
  MountPath = I
  Initiator @System1 = { "iqn.1991-05.com.microsoft:System1.e10dc.com" }
)
NetAppSnapMirror EXCHSG_SG1-NetAppSnapMirror (
  FilerResName = EXCHSG_SG1-NetAppFiler
  VolumeName = Vol1
)
EXCHSG_SG1-NetAppSnapDrive requires EXCHSG_SG1-NetAppFiler
EXCHSG_SG1-NetAppSnapDrive requires EXCHSG_SG1-NetAppSnapMirror
EXCHSG_SG1-NetAppSnapMirror requires EXCHSG_SG1-NetAppFiler
SG1-Exch2010DB requires EXCHSG_SG1-NIC
SG1-Exch2010DB requires EXCHSG_SG1-NetAppSnapDrive
// resource dependency tree
//
// group EXCHSG_SG1
// {
//   Exch2010DB SG1-Exch2010DB
//     {
//       NIC EXCHSG_SG1-NIC
//       NetAppSnapDrive EXCHSG_SG1-NetAppSnapDrive
//         {
//           NetAppFiler EXCHSG_SG1-NetAppFiler
//           NetAppSnapMirror EXCHSG_SG1-NetAppSnapMirror
//             {
//               NetAppFiler EXCHSG_SG1-NetAppFiler
//             }
//           }
//         }
//       }
//     }
// }
```

```
// }
```