

Veritas Storage Foundation™ and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange 2010

Windows Server 2008 (x64)

5.1 Service Pack 2



Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange 2010

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

Product version: 5.1. Service Pack 2

Document version: 5.1.SP2.1

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

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

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

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

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

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

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

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

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

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apac@symantec.com
------------------------	--

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

North America and Latin America	supportsolutions@symantec.com
---------------------------------	--

Contents

Section 1 Introduction and Concepts

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

About clustering solutions with SFW HA	18
About high availability	18
How a high availability solution works	19
About SFW HA support for Exchange Server 2010	19
How the Exchange 2010 HA solution differs from earlier	
Exchange HA solutions	20
About campus clusters	21
Differences between campus clusters and local clusters	22
Sample campus cluster configuration	22
What you can do with a campus cluster	24
About replication	25
About a replicated data cluster	25
How VCS replicated data clusters work	26
About disaster recovery	27
What you can do with a disaster recovery solution	28
Typical disaster recovery configuration	28
Where to get more information about Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange Server	30

Chapter 2 Introducing the VCS agent for Exchange 2010

About the VCS database agent for Microsoft Exchange 2010	32
Exchange 2010 database agent functions	33
Exchange 2010 database agent state definitions	33
Exchange 2010 database agent resource type definition	34
Exchange 2010 database agent attribute definitions	34
Exchange 2010 service group resource dependency graph	35
Exchange 2010 service group sample configuration	36

Chapter 3 Using the Solutions Configuration Center

About the Solutions Configuration Center	39
Available options from the Configuration Center	41
About running the Configuration Center wizards	43
Following the workflow in the Configuration Center	45
Solutions wizard logs	46

Section 2 Configuration Workflows

Chapter 4 Configuration workflows for Exchange Server

About using the workflow tables	50
High availability (HA) configuration (New Server)	51
High availability (HA) configuration (Existing Server)	53
VCS campus cluster configuration	55
VCS Replicated Data Cluster configuration	57
Disaster recovery configuration	61
DR configuration tasks: Primary site	61
DR configuration tasks: Secondary site	63

Section 3 Requirements and Planning

Chapter 5 Requirements and planning for your HA and DR configurations

Reviewing the requirements	70
Disk space requirements	70
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	70
Supported Exchange 2010 versions	71
System requirements for SFW HA	72
Network requirements for SFW HA	72
Permission requirements for SFW HA	73
Additional requirements for SFW HA	73
Best practices for SFW HA	74
Reviewing the HA configuration	75
Sample Exchange server HA configuration	76
IP addresses required	78
Following the HA workflow in the Solutions Configuration Center	79
Reviewing a standalone Exchange Server configuration	79
Sample standalone Exchange server configuration	81
Reviewing the campus cluster configuration	82
Campus cluster failover using the ForceImport attribute	83

Reviewing the Replicated Data Cluster configuration	84
Sample Exchange Server Replicated Data Cluster configuration	84
About setting up a Replicated Data Cluster configuration	85
About setting up replication	86
About configuring and migrating the service group	86
Reviewing the disaster recovery configuration	86

Section 4 Deployment

Chapter 6 Installing and configuring SFW HA

Configuring the storage hardware and network	92
Installing Veritas Storage Foundation HA for Windows	94
Installing Symantec Trusted certificate for unsigned drivers	94
Installing Storage Foundation HA for Windows	94
Configuring cluster disk groups and volumes for Exchange Server	98
About cluster disk groups and volumes	99
Prerequisites for configuring cluster disk groups and volumes	100
For standalone Exchange configurations	100
For Campus Cluster configurations	101
For configurations using VVR	101
Considerations for converting existing shared storage to cluster disk groups and volumes	101
Considerations for disks and volumes for campus clusters	101
Considerations for volumes for a VVR configuration	102
Sample disk group and volume configuration for Exchange 2010	103
Viewing the available disk storage	103
Creating a cluster disk group	104
Creating volumes	105
About managing disk groups and volumes	111
Importing a disk group and mounting a volume	111
Unmounting a volume and deporting a disk group	111
Adding drive letters to mount the volumes	112
Deporting the cluster disk group	113
Configuring the cluster	114
Configuring notification	121
Adding a node to an existing VCS cluster	122
Prerequisites for adding a node to an existing cluster	122
Adding a node to a VCS cluster	123

Chapter 7	Installing Exchange Server	
	About installing Exchange Server 2010	130
	Before you install Exchange Server 2010	130
	Privileges required for installing Exchange 2010	130
	Installing Exchange Server 2010	131
	Creating mailbox databases on shared storage	131
	Moving mailbox databases to shared storage	132
	Adding new Exchange servers to an existing cluster	133
Chapter 8	Configuring Exchange Server for failover	
	About configuring the VCS Exchange Server service group	140
	Prerequisites for configuring the Exchange Server service group	140
	Creating the Exchange Server 2010 service group	141
	Verifying the Exchange Server cluster configuration	146
	Determining additional steps needed	147
Chapter 9	Configuring campus clusters for Exchange Server	
	Tasks for configuring campus clusters	150
	Verifying the campus cluster: Switching the service group	150
	Setting the ForceImport attribute to 1 after a site failure	151
Chapter 10	Configuring Replicated Data Clusters for Exchange Server	
	Tasks for configuring Replicated Data Clusters for Exchange Server	153
	Creating the primary system zone	156
	Creating a parallel environment in the secondary zone	156
	Adding the systems in the secondary zone to the cluster	157
	Setting up security for VVR	161
	Setting up the Replicated Data Sets (RDS)	164
	Prerequisites for setting up the RDS for the primary and secondary zones	164
	Creating the Replicated Data Sets with the wizard	164
	Configuring a hybrid RVG service group for replication	176
	Creating the RVG service group	176
	Configuring the RVG service group for RDC replication	178
	Configuring the IP and NIC resources	178
	Configuring the VMDg resources for the disk groups	179
	Adding the VVR RVG resources for the disk groups	181
	Linking the VVR RVG resources to establish dependencies	182
	Deleting the VMDg resource from the Exchange Server service group	183

Configuring the RVG Primary resources	183
Creating the RVG Primary resources	184
Linking the RVG Primary resources to establish dependencies	184
Bringing the RVG Primary resources online	185
Configuring the primary system zone for the RVG	185
Setting a dependency between the service groups	186
Adding the nodes from the secondary zone to the RDC	187
Adding the nodes from the secondary zone to the RVG service group	187
Configuring secondary zone nodes in the RVG service group	188
Configuring the IP resources for fail over	188
Adding the nodes from the secondary zone to the Exchange Server service group	190
Configuring the zones in the Exchange Server service group	191
Verifying the RDC configuration	191
Bringing the service group online	191
Switching online nodes	192
Additional instructions for GCO disaster recovery	193

Chapter 11 Deploying disaster recovery for Exchange Server

Tasks for deploying a disaster recovery configuration of Microsoft Exchange	197
Reviewing the disaster recovery configuration	200
Supported disaster recovery configurations for service group dependencies	201
Setting up the secondary site: Installing SFW HA and configuring a cluster	202
Verifying your primary site configuration	203
Setting up your replication environment	203
Setting up security for VVR	204
Requirements for EMC SRDF array-based hardware replication	206
Software requirements for configuring EMC SRDF	206
Replication requirements for EMC SRDF	207
Requirements for Hitachi TrueCopy array-based hardware replication	208
Software requirements for Hitachi TrueCopy	208
Replication requirements for Hitachi TrueCopy	209
Assigning user privileges (secure clusters only)	211
Configuring disaster recovery with the DR wizard	212
Cloning the storage on the secondary site using the DR wizard (VVR replication option)	216
Creating temporary storage on the secondary site using	

the DR wizard (array-based replication)	220
Installing Exchange 2010	224
Cloning the service group configuration on to the secondary site using the DR wizard	224
Configuring replication and global clustering	227
Configuring VVR replication and global clustering	227
Configuring EMC SRDF replication and global clustering	235
Optional settings for EMC SRDF	237
Configuring Hitachi TrueCopy replication and global clustering	238
Optional settings for HTC	241
Configuring global clustering only	241
Verifying the disaster recovery configuration	243
Establishing secure communication within the global cluster (optional)	245
Adding multiple DR sites (optional)	247
Recovery procedures for service group dependencies	247
Possible task after creating the DR environment:	
Adding a new failover node to a VVR environment	251
Preparing the new node	251
Preparing the existing DR environment	251
Installing Exchange on the new node	252
Modifying the replication and Exchange service groups	252
Reversing replication direction	253

Chapter 12 Testing fault readiness by running a fire drill

About disaster recovery fire drills	255
About the Fire Drill Wizard	256
About Fire Drill Wizard general operations	256
About Fire Drill Wizard operations in a VVR environment	257
About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment	258
About post-fire drill scripts	260
Exchange 2007 scripts or cmdlets	261
Exchange 2010 scripts or cmdlets	261
Tasks for configuring and running fire drills	262
Prerequisites for a fire drill	264
Prerequisites for a fire drill in a VVR environment	264
Prerequisites for a fire drill in a Hitachi TrueCopy environment	265
Prerequisites for a fire drill in an EMC SRDF environment	266
Preparing the fire drill configuration	266
System Selection panel details	269
Service Group Selection panel details	269

Secondary System Selection panel details	269
Disk Selection panel details	269
Hitachi TrueCopy Path Information panel details	270
HTCSnap Resource Configuration panel details	271
SRDFSnap Resource Configuration panel details	271
Fire Drill Preparation panel details	272
Running a fire drill	272
Post fire drill operations panel details	274
Recreating a fire drill configuration that has changed	274
Restoring the fire drill system to a prepared state	277
Deleting the fire drill configuration	278
.....	279

Section 5 Reference

Appendix A VCS Logging

Index	285
-------------	-----

Introduction and Concepts

This section contains the following chapters:

- [Chapter 1, “Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange Server”](#)
- [Chapter 2, “Introducing the VCS agent for Exchange 2010”](#)
- [Chapter 3, “Using the Solutions Configuration Center”](#)

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

This chapter contains the following topics:

- [“About clustering solutions with SFW HA”](#) on page 18
- [“About high availability”](#) on page 18
- [“How a high availability solution works”](#) on page 19
- [“About SFW HA support for Exchange Server 2010”](#) on page 19
- [“How the Exchange 2010 HA solution differs from earlier Exchange HA solutions”](#) on page 20
- [“About campus clusters”](#) on page 21
- [“Differences between campus clusters and local clusters”](#) on page 22
- [“Sample campus cluster configuration”](#) on page 22
- [“What you can do with a campus cluster”](#) on page 24
- [“About replication”](#) on page 25
- [“About a replicated data cluster”](#) on page 25
- [“How VCS replicated data clusters work”](#) on page 26

- [“About disaster recovery”](#) on page 27
- [“What you can do with a disaster recovery solution”](#) on page 28
- [“Typical disaster recovery configuration”](#) on page 28
- [“Where to get more information about Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange Server”](#) on page 30

About clustering solutions with SFW HA

Storage Foundation HA for Windows (SFW HA) provides the following clustering solutions for high availability and disaster recovery:

- High availability failover cluster on the same site
- Campus cluster, in a two-node configuration with each node on a separate site
- Replicated data cluster, with a primary zone and a secondary zone existing within a single cluster, which can stretch over two buildings or data centers connected with Ethernet
- Wide area disaster recovery, with a separate cluster on a secondary site, with replication support using Veritas Volume Replicator or hardware replication

About high availability

The term high availability (HA) refers to a state where data and applications are highly available because software or hardware is in place to maintain the continued functioning in the event of computer failure. High availability can refer to any software or hardware that provides fault tolerance, but generally the term has become associated with clustering.

A cluster is a group of independent computers working together to ensure that mission-critical applications and resources are as highly available as possible. The group is managed as a single system, shares a common namespace, and is specifically designed to tolerate component failures and to support the addition or removal of components in a way that is transparent to users.

Local clustering provides high availability through database failover. This solution provides local recovery in the event of application, operating system, or hardware failure, and minimizes planned and unplanned application downtime.

The high availability solution includes procedures for installing and configuring clustered Exchange Server environments using Veritas Storage Foundation HA for Windows (SFW HA). SFW HA includes Veritas Storage Foundation for Windows and Veritas Cluster Server.

Setting up the clustered environment is also the first step in creating a wide-area disaster recovery solution using a secondary site.

How a high availability solution works

Keeping data and applications functioning 24 hours a day and seven days a week is the desired norm for critical applications today. Clustered systems have several advantages over standalone servers, including fault tolerance, high availability, scalability, simplified management, and support for rolling upgrades.

Using Veritas Storage Foundation HA for Windows as a local high availability solution paves the way for a wide-area disaster recovery solution in the future.

A high availability solution is built on top of a backup strategy and provides the following benefits:

- Reduces planned and unplanned downtime.
- Serves as a local and wide-area failover (rather than load-balancing) solution. Enables failover between sites or between clusters.
- Manages applications and provides an orderly way to bring processes online and take them offline.
- Consolidates hardware in larger clusters. The HA environment accommodates flexible fail over policies, active-active configurations, and shared standby servers for Exchange Server.

About SFW HA support for Exchange Server 2010

High availability support for Exchange 2010 is available 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.

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.

HA support is not available for public folders. VCS also does not provide HA 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.

How the Exchange 2010 HA solution differs from earlier Exchange HA solutions

The HA solution for Exchange 2010 differs from the HA solution for Exchange 2003 and Exchange 2007.

The differences are as follows:

- In Exchange 2003 and Exchange 2007, you configure a virtual server name that becomes the name of the Exchange server in a clustered environment. The active cluster node hosts the Exchange virtual server (EVS). The other nodes act as redundant servers ready to host the EVS if the active node fails. In case of a server failure, VCS moves the EVS from the active node to an alternate system in the service group.

The concept of Exchange virtual server (EVS) is not applicable to Exchange 2010. For Exchange 2010, you are not required to configure a virtual server name for the Exchange server. The unit of failover is the Exchange mailbox database. A cluster node hosts the Exchange mailbox databases. In case of a server failure, VCS moves the mailbox databases from the affected node to an alternate node in the service group. The databases then become active on that node.

- In Exchange 2003 and Exchange 2007, you run the Exchange Setup Wizard for VCS before and after installing Exchange. The wizard performs certain pre-installation and post-installation tasks that are required for making Exchange server highly available.

For Exchange 2010, you do not need to run the Exchange Setup Wizard for VCS before or after installing Exchange. There are no pre-install or post-install tasks. You can directly install Exchange 2010 on the required systems.

- In Exchange 2007, VCS requires that you install only the Exchange Mailbox Server role on the systems that are going to be part of the Exchange service group. You cannot install the other Exchange server roles on those systems. Therefore, other server roles such as Client Access server or Hub Transport server are to be configured on separate systems outside the Exchange service group.

For Exchange 2010, there is no such requirement. You can install the Mailbox Server role and the other Exchange server roles on the same

systems. However, to make the Exchange mailbox databases highly available, you must install at least the Mailbox Server role on the systems that are going to be part of the Exchange service group.

- In Exchange 2003 and Exchange 2007, the Exchange virtual server name is made highly available. In case of a server failure, VCS moves the EVS from one node to the other. A cluster node can host only one EVS at a given time. Even in an Any-to-Any configuration where you can configure two or more Exchange service groups sharing the same set of nodes, a single node cannot host more than one EVS at the same time.

For Exchange 2010, the unit of failover is the Exchange mailbox database. The databases are not bound to the Exchange server name. Each of the cluster nodes where you install the Mailbox Server role can host a separate set of mailbox databases that can be active at the same time. You can create multiple Exchange service groups sharing the same set of cluster nodes. Each cluster node hosts a separate Exchange service group and also serves as a failover target for the other service groups. VCS can move the mailbox databases across any of the configured cluster nodes within a service group. Thus, it is possible to fail over all the Exchange service groups to a single node, if required. The mailbox server on that node handles client requests related to databases that are active on that node.

This reduces server redundancy and helps optimize the resource requirements for making Exchange 2010 highly available.

- For Exchange 2010, you can use the Exchange 2010 Database Configuration Wizard to configure multiple Exchange 2010 service groups in a single workflow. Service groups are not bound to the Exchange server name. The unit of failover is the mailbox database. The wizard creates a service group per cluster disk group. If the selected mailbox databases reside on different cluster disk groups, the wizard automatically configures those databases in separate service groups.
- For Exchange 2010, there is no Exchange virtual server name and hence the Lanman and virtual IP resources are not configured in the service group. The VCS RegRep resource is also not required in the Exchange 2010 service group.

About campus clusters

Campus clusters are clusters in separate buildings (or sites) with mirrored SAN-attached storage located in each building. Typical campus clusters involve two sites; you can use more than two sites for additional redundancy. In a typical configuration, each node has its own storage array and contains mirrored data of the storage on the other array.

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

This solution provides a level of high availability that is above mirroring or clustering at a single site but is not as complex as disaster recovery with replication.


Differences between campus clusters and local clusters

The procedures for setting up a campus cluster are nearly the same as those for local clusters, except that a campus cluster has the nodes located in separate buildings, so the hardware setup requires SAN interconnects that allows these connections. Also, in a campus cluster, each node has its own storage array rather than having a shared storage array between the two clusters. Both local clusters and campus clusters have SFW dynamic disk groups and volumes, but the volumes on each campus cluster node are mirrors of one another.

Sample campus cluster configuration

[Figure 1-1](#) shows a sample configuration that represents a campus cluster with two sites, Site A and Site B.

Figure 1-1 Campus cluster: active-passive configuration



With SFW, a campus cluster can be set up using a Veritas Cluster Server (VCS) configuration. Both configurations involve setting up a single cluster with two nodes that are in separate buildings and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number of disks and contains mirrored data of the storage on the other array. SFW provides the mirrored storage and the disk groups that make it possible to fail over the storage by deporting the disk groups on one node and importing them on the other.

If a site failure occurs in a two-node campus cluster, the remaining cluster node will not be able to bring the cluster disk groups online because it cannot reserve a majority of disks in the disk groups. To allow for failover to the other site, a procedure forces the import to the other node, allowing a cluster disk group to be brought online on another node when that node has a minority of the cluster disks.

Implementing these force import procedures should be done with care. The primary site may appear to have failed but what really has happened is that both the storage interconnect between sites and the heartbeats have been lost. In that

case, cluster disk groups can still be online on the primary node. If a force import is done so that the data can be accessed on the secondary site, the cluster disks will be online on both sites, risking data corruption.

What you can do with a campus cluster

Administrators can use campus clusters to protect data from natural disasters, such as floods and hurricanes, and unpredictable power outages. Campus clusters provide disaster protection when an entire site goes down.

In the event of a site disaster, such as power failure in a building, campus clusters offer a level of high availability that surpasses mirroring or clustering at a single site by dispersing the clustered servers into different buildings or sites. This environment also provides a simpler solution for disaster recovery than a more elaborate SFW HA DR environment with replication software; however, a campus cluster generally stretches a shorter distance than a replication-based solution depending on the hardware.

About replication

The term replication refers to the use of a tool or service to automate the process of maintaining a consistent copy of data from a designated source (primary site) on one or more remote locations (secondary sites).

In the event that the primary site data center is destroyed, the application data is readily available at the remote site, and the application can be restarted at the remote site.

SFW HA provides Veritas Volume Replicator (VVR) for use in replication. VVR can be used for replication in either a replicated data cluster (RDC) or a wide-area disaster recovery solution.

The SFW HA disaster recovery solution also supports hardware replication.

For more information on VVR refer to the *Veritas Volume Replicator Administrator's Guide*.

About a replicated data cluster

A Replicated Data Cluster (RDC) uses data replication, instead of shared storage, to assure data access to all the nodes in a cluster. The Replicated Data Cluster configuration provides both local high availability and disaster recovery functionality in a single VCS cluster. You can set up RDC in a VCS environment using Veritas Volume Replicator (VVR).

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

For VVR replication to occur, the disk groups containing the Replicated Volume Group (RVG) must be imported at the primary and secondary zones. The replication service group must be online at both zones simultaneously, and must be configured as a hybrid VCS service group.

The Exchange service group is configured as a failover service group. The Exchange service group must be configured with an online local hard dependency on the replication service group.

Note: VVR supports multiple replication secondary targets for any given primary. However, RDC for VCS supports only one replication secondary for a primary.

An RDC configuration is appropriate in situations where dual dedicated LLT links are available between the primary zone and the secondary zone but lacks shared storage or SAN interconnect between the primary and secondary data centers. In an RDC, data replication technology is employed to provide node access to data in a remote zone. You must use dual dedicated LLT links between the replicated nodes.

How VCS replicated data clusters work


To understand how a RDC configuration works, let us take the example of Microsoft Exchange configured in a VCS replicated data cluster. The configuration has two system zones:

- Primary zone (zone 0) comprising nodes located at the primary site and attached to the primary storage
- Secondary zone (zone 1) comprising nodes located at the secondary site and attached to the secondary storage

Exchange is installed and configured on all nodes in the cluster. The Exchange mailbox databases are located on shared disks within each RDC zone and are replicated across RDC zones to ensure data concurrency. The Exchange service group is online on a system in the current primary zone and is configured to fail over in the cluster.

[Figure 1-2](#) shows failover in a replicated data cluster.

Figure 1-2 Failover in a replicated data cluster



In the event of a system or Exchange failure, VCS attempts to fail over the Exchange service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone (zone 1). VCS also redirects clients after the databases are online on the new location.

About disaster recovery

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

A disaster recovery (DR) solution is a series of procedures which you can use to safely and efficiently restore application user data and services in the event of a catastrophic failure. A typical DR solution requires that you have a source or *primary site* and a destination or *secondary site*. The user application data on the primary site is replicated to the secondary site. The cluster on the primary site

provides data and services during normal operations. In the event of a disaster at the primary site and failure of the cluster, the secondary site provides the data and services.

What you can do with a disaster recovery solution

A DR solution is vital for businesses that rely on the availability of data.

A well-designed DR solution prepares a business for unexpected disasters and provides the following benefits in a DR situation:

- Minimizes economic loss due to the unavailability or loss of data.
- Provides a plan for the safe and orderly recovery of data in the event of a disaster.
- Ensures safe and efficient recovery of data and services.
- Minimizes any decision making during DR.
- Reduces the reliance on key individuals.

Strategically planning a DR solution provides businesses with affordable ways to meet their service level agreements, comply with government regulations, and minimize their business risks.


Note: A DR solution requires a well-defined backup strategy. Refer to Veritas NetBackup or Backup Exec product documentation for information on configuring backup.

Typical disaster recovery configuration

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

[Figure 1-3](#) illustrates a typical DR configuration.

Figure 1-3 Typical DR configuration in a VCS cluster



The illustration displays an environment with a DR solution that is prepared for a disaster. In this case, the primary site consists of two nodes, System1 and System2. Similarly the secondary setup consists of two nodes, System3 and System4. Each site has a clustered setup with the nodes set up appropriately for failover within the site.

Data is replicated from the primary site to the secondary site. Replication between the storage is set up using a replication software. If the Microsoft Exchange Server on System1 fails, the database comes online on node System2 and System2 begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal.

When a failure occurs, such as an earthquake that destroys the data center in which the primary site resides, the DR solution is activated. System3 at the secondary site takes over, and the data that was replicated to the secondary site is used to restore the application services to clients.

Where to get more information about Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange Server

Table 1-1 shows the available Veritas Storage Foundation and High Availability Solutions solutions guides for Microsoft Exchange Server 2010. Separate guides are available for earlier versions of Exchange, for Microsoft SQL, for Enterprise Vault, for Microsoft SharePoint Server, and for additional application solutions.

Table 1-1 SFW HA solutions guides for Exchange Server 2010

Title	Description
<i>Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange 2010</i>	Solutions for Exchange Server 2010 and Veritas Cluster Server clustering with Veritas Storage Foundation HA for Windows ■ High availability (HA) ■ Campus clusters ■ Replicated data clusters ■ Disaster recovery (DR) with Veritas Volume Replicator or hardware array replication
<i>Veritas Storage Foundation and High Availability Solutions Quick Recovery Solutions Guide for Microsoft Exchange 2010</i>	Quick Recovery solutions for Exchange Server 2010 using either Veritas Storage Foundation for Windows or Veritas Storage Foundation HA for Windows.

Introducing the VCS agent for Exchange 2010

This chapter contains the following topics:

- [“About the VCS database agent for Microsoft Exchange 2010”](#) on page 32
- [“Exchange 2010 database agent functions”](#) on page 33
- [“Exchange 2010 database agent state definitions”](#) on page 33
- [“Exchange 2010 database agent resource type definition”](#) on page 34
- [“Exchange 2010 database agent attribute definitions”](#) on page 34
- [“Exchange 2010 service group resource dependency graph”](#) on page 35
- [“Exchange 2010 service group sample configuration”](#) on page 36

This chapter describes the VCS agent for Exchange Server and lists the resource type definition and attribute definitions of the agent. The resource type represents the VCS configuration definition of the agent and specifies how the agent is defined in the cluster configuration file, `main.cf`. The Attribute Definitions lists the attributes associated with the agent. The Required Attributes table lists the attributes that must be configured for the agent to function properly.

About the VCS database agent for Microsoft Exchange 2010

The VCS database agent for Microsoft Exchange 2010 provides high availability for Exchange 2010 databases in a VCS cluster. The agent monitors Exchange 2010 mailbox databases and critical Exchange services to ensure high availability. The agent detects a failure if any of the configured Exchange databases or critical Exchange services become unavailable. When this occurs, the agent moves the Exchange databases to the next available Exchange server in the service group's system list and if required, starts the critical Exchange services on that node. The databases then become active on that system.

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.

Exchange 2010 database agent functions

Agent functions include the following:

Online	The agent performs the following actions as part of its online function: ■ 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	The agent performs the following actions as part of its monitor function: ■ 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.

Exchange 2010 database agent state definitions

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.

Exchange 2010 database agent resource type definition

The VCS Database agent for Exchange 2010 is represented by the Exch2010DB resource type.

The resource definition is as follows:

```
type Exch2010DB (  
    static i18nstr ArgList[] = { DBName, MonitorService }  
    i18nstr DBName  
    boolean MonitorService = 1  
)
```

Exchange 2010 database agent attribute definitions

Review the following information to familiarize yourself with the agent attributes for a Exch2010DB resource type. This information will assist you during the agent configuration.

Table 2-1 Exchange 2010 database agent required attributes

Required Attributes	Type and Dimension	Definition
DBName	string-scalar	Name of the Exchange 2010 mailbox database to be made highly available.
MonitorService	boolean-scalar	Defines whether the agent should monitor the critical Exchange 2010 services. The value 1 (True) indicates that the agent monitors the critical services. The value 0 (False) indicates that it does not. Default is 1 (True). If this attribute is set to 1 (True), the agent monitors the following Exchange 2010 services internally: ■ Microsoft Exchange System Attendant (MSExchangeSA) ■ Microsoft Exchange Mail Submission (MSExchangeMailSubmission) Note: You cannot define which Exchange 2010 services should be monitored by the agent.


Exchange 2010 service group resource dependency graph

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 shared disk group is configured using the Volume Manager Disk Group agent (VMDg) resource. The MountV mount points are created using the MountV agent resource. The Exchange 2010 databases (db1, db2, db3) are configured using the Exch2010DB resource. Exchange databases db1 and db2 are using the same volume (ExchSG_vol1) while db3 is created on a separate volume (ExchSG_vol2) on the same disk group. The Exchange database resources come online after each of the underlying storage resources are brought online.

Figure 2-1 shows the dependencies in the Exchange 2010 database service group.

Figure 2-1 Exchange 2010 resource dependency



Exchange 2010 service group sample configuration

A sample VCS Exchange 2010 service group configuration file (main.cf) is included for reference.

```
include "types.cf"
cluster exchcluster (
    SecureClus = 1
)
    system System1 (
    )
    system System2 (
    )
    system System3 (
    )
    system System4 (
    )

group EXCHSG_SG1 (
    SystemList = { System1 = 0, System2 = 1, System3 = 2 }
)

    Exch2010DB db1-Exch2010DB (
        DBName = db1
    )
    Exch2010DB db2-Exch2010DB (
        DBName = db2
    )
    Exch2010DB db3-Exch2010DB (
        DBName = db3
    )
    Exch2010DB db4-Exch2010DB (
        DBName = db4
    )
    MountV EXCHSG_SG1-MountV (
        MountPath = "I:"
        VolumeName = vol1
        VMDGResName = EXCHSG_SG1-VMdG
    )
    MountV EXCHSG_SG1-MountV-1 (
        MountPath = "J:"
        VolumeName = vol2
        VMDGResName = EXCHSG_SG1-VMdG
    )
    NIC EXCHSG_SG1-NIC (
        MACAddress @System1 = 00-14-C2-3A-1F-70
        MACAddress @System2 = 00-19-BB-30-A1-D6
        MACAddress @System3 = 00-15-60-0F-20-BA
    )
    VMDg EXCHSG_SG1-VMdG (
        DiskGroupName = SG1DG3
```

```
DGGuid = 282aa5eb-f1fc-44c8-b62b-072924f6fa47
)
```

```
EXCHSG_SG1-MountV requires EXCHSG_SG1-VMDg
db1-Exch2010DB requires EXCHSG_SG1-NIC
db1-Exch2010DB requires EXCHSG_SG1-MountV
db2-Exch2010DB requires EXCHSG_SG1-NIC
db2-Exch2010DB requires EXCHSG_SG1-MountV
EXCHSG_SG1-MountV-1 requires EXCHSG_SG1-VMDg
db3-Exch2010DB requires EXCHSG_SG1-MountV-1
db3-Exch2010DB requires EXCHSG_SG1-NIC
db4-Exch2010DB requires EXCHSG_SG1-MountV-1
db4-Exch2010DB requires EXCHSG_SG1-NIC
```

```
// resource dependency tree
// group EXCHSG_SG1
// {
//   Exch2010DB db1-Exch2010DB
//   {
//     NIC EXCHSG_SG1-NIC
//     MountV EXCHSG_SG1-MountV
//     {
//       VMDg EXCHSG_SG1-VMDg
//     }
//   }
//   Exch2010DB db2-Exch2010DB
//   {
//     NIC EXCHSG_SG1-NIC
//     MountV EXCHSG_SG1-MountV
//     {
//       VMDg EXCHSG_SG1-VMDg
//     }
//   }
//   Exch2010DB db3-Exch2010DB
//   {
//     MountV EXCHSG_SG1-MountV-1
//     {
//       VMDg EXCHSG_SG1-VMDg
//     }
//     NIC EXCHSG_SG1-NIC
//   }
//   Exch2010DB db4-Exch2010DB
//   {
//     MountV EXCHSG_SG1-MountV-1
//     {
//       VMDg EXCHSG_SG1-VMDg
//     }
//     NIC EXCHSG_SG1-NIC
//   }
// }
```


Using the Solutions Configuration Center

This chapter covers the following topics:

- [About the Solutions Configuration Center](#)
- [Available options from the Configuration Center](#)
- [About running the Configuration Center wizards](#)
- [Following the workflow in the Configuration Center](#)
- [Solutions wizard logs](#)

About the Solutions Configuration Center

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

- Microsoft Exchange Server 2003, 2007, and 2010
- Microsoft SQL Server 2000, 2005, 2008, and 2008 R2
- Enterprise Vault Server (high availability and disaster recovery solutions)
- Microsoft SharePoint Server 2010 (high availability and disaster recovery solutions)
- Additional applications

Depending on the application, the following solutions may be available:

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

- Campus cluster disaster recovery, including the following:
 - Campus cluster using Veritas Cluster Server (SFW HA)
 - Campus cluster using Microsoft clustering
- Wide area disaster recovery involving multiple sites
- Quick Recovery for on-host recovery from logical errors in application data (available for Microsoft Exchange 2003, 2007, 2010, and for Microsoft SQL Server 2005, 2008, and 2008 R2)
- Fire drill to test the fault readiness of a disaster recovery environment

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

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


Available options from the Configuration Center

On the Applications tab, the Solutions Configuration Center is context-sensitive to the application. For example, the solutions displayed when you click the application name are those available for that application. The Solution Guides listed in the right pane match the selected application.

In addition, some choices can vary depending on the operating system of the node on which you launch the wizard. For example, since Microsoft Exchange 2003 runs only on 32-bit operating systems, on a 64-bit system only the Exchange 2007 and Exchange 2010 configuration wizards are shown.

Figure 3-1 shows the choices available on a 64-bit system when you click Solutions for Microsoft Exchange.

Figure 3-1 Solutions Configuration Center for Microsoft Exchange



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

Figure 3-2 shows one of the steps for implementing high availability for SQL Server, showing the SQL Server configuration wizards.

Figure 3-2 Context-sensitive step for SQL Server



Figure 3-3 shows one of the steps for implementing high availability for Exchange, showing the configuration wizards for Exchange.

Figure 3-3 Context-sensitive step for Exchange



About running the Configuration Center wizards

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

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

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

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

New Dynamic Disk Group Wizard	Launched from the Veritas Enterprise Administrator console
New Volume Wizard	Launched from the Veritas Enterprise Administrator console
Exchange Setup Wizard	Installs and configures Microsoft Exchange Server 2003 and 2007 for the high availability environment If Microsoft Exchange is already installed, refer to the documentation for further instructions. This wizard is not required for Exchange Server 2010.
Exchange Configuration Wizard	Configures the service group for Microsoft Exchange Server 2003 and 2007 high availability
Exchange 2010 Configuration Wizard	Configures the service group for Microsoft Exchange Server 2010 high availability

SQL Server Configuration Wizard	Configures the service group for SQL Server 2000 or SQL Server 2005 high availability You must first install SQL Server on each node according to the instructions in the documentation.
SQL Server 2008 Configuration Wizard	Configures the service group for SQL Server 2008 or SQL Server 2008 R2 high availability You must first install SQL server on each node according to the instructions in the documentation.
Enterprise Vault Cluster Setup Wizard	Configures the service group for Enterprise Vault Server high availability
MSDTC Wizard	Configures an MSDTC server service group for SQL Server 2000, 2005, or 2008 environments
MSMQ Configuration Wizard	Configures a Microsoft Message Queuing (MSMQ) service group
SharePoint 2010 Configuration Wizard	Configures SharePoint Server 2010 service groups You can run the wizard from any SFW HA cluster node where SharePoint Server is installed and configured.

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

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


Following the workflow in the Configuration Center

During the multi-step High Availability Configuration workflow, you may find it helpful to run an SFW HA client on another system and leave the Configuration Center open on that system. In this way, you can see what step comes next, drill down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format.

When setting up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration.

Figure 3-4 shows the high-level overview of the workflow steps for configuring high availability for Exchange from the Solutions Configuration Center.

Figure 3-4 Workflow for configuring Exchange high availability



Solutions wizard logs

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

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

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

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

For Windows Server 2003:

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

For Windows Server 2008:

C:\ProgramData\Veritas\winsolutions\log

Configuration Workflows

This section contains the following chapters:

- [Chapter 4, “Configuration workflows for Exchange Server”](#)

Configuration workflows for Exchange Server

This chapter contains the following topics:

- [“About using the workflow tables”](#) on page 50
- [“High availability \(HA\) configuration \(New Server\)”](#) on page 51
- [“High availability \(HA\) configuration \(Existing Server\)”](#) on page 53
- [“VCS campus cluster configuration”](#) on page 55
- [“VCS Replicated Data Cluster configuration”](#) on page 57
- [“Disaster recovery configuration”](#) on page 61

About using the workflow tables

Configuring a high availability or a disaster recovery environment involves a series of tasks such as evaluating the requirements, configuring the storage, installing and configuring VCS, installing and configuring the application, and so on. A configuration workflow table provides high level description of all the required tasks, with links to the topics that describe these tasks in detail.

Separate workflow tables are provided for High Availability (HA), Campus Cluster, Replicated Data Cluster (RDC) and Disaster Recovery (DR) configurations. Depending on the required high availability configuration, use the appropriate workflow table as a guideline to perform the installation and configuration.

Symantec recommends using the Solutions Configuration Center (SCC) as a guide for installing and configuring SFW HA for Exchange Server.

See “[About the Solutions Configuration Center](#)” on page 39.

The workflow tables are organized to follow the workflows in the Solutions Configuration Center.

For example, in using the Solutions Configuration Center to set up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration. Likewise, in this guide, you first refer to the High Availability workflow to set up high availability. You then continue with the appropriate workflow, either Replicated Data Cluster, campus cluster, or disaster recovery, for any additional solution that you want to implement.

High availability (HA) configuration (New Server)

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

Table 4-1 Task list: Exchange Server HA configuration tasks

Action	Description
Verify hardware and software requirements	■ Review the requirements See “Reviewing the requirements” on page 70.
Review the HA configuration	■ Review the sample configuration See “Reviewing the HA configuration” on page 75.
Configure the storage hardware and network	■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which Exchange will be installed See “Configuring the storage hardware and network” on page 92.
Install SFW HA	■ Install Veritas Storage Foundation HA for Windows See “Installing Veritas Storage Foundation HA for Windows” on page 94.
Configure disk groups and volumes for Exchange Server	■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Create dynamic volumes for the mailbox databases and logs See “Configuring cluster disk groups and volumes for Exchange Server” on page 98.
Configure VCS cluster	■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster See “Configuring the cluster” on page 114.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “About installing Exchange Server 2010” on page 130.

Table 4-1 Task list: Exchange Server HA configuration tasks (Continued)

Action	Description
Create Exchange databases on shared storage	■ Use the Exchange Management Console or the Exchange Management Shell to create Exchange mailbox databases on shared storage See “Creating mailbox databases on shared storage” on page 131.
Create an Exchange service group	■ Create a Exchange Server 2010 database service group using the Exchange 2010 Service Group Configuration Wizard See “About configuring the VCS Exchange Server service group” on page 140.
Verify the HA configuration	■ Test failover between nodes See “Verifying the Exchange Server cluster configuration” on page 146.
Proceed to the additional steps depending on the desired HA configuration.	See “Determining additional steps needed” on page 147.

High availability (HA) configuration (Existing 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 new Veritas Storage Foundation HA environment.

[Table 4-2](#) outlines the high-level objectives and the tasks to complete each objective for converting an existing standalone Exchange Server for high availability.

Table 4-2

Task list: Standalone Exchange Server HA configuration tasks

Action	Description
Verify hardware and software requirements	■ Review the requirements See “Reviewing the requirements” on page 70.
Review the HA configuration	■ Review the sample configuration See “Reviewing a standalone Exchange Server configuration” on page 79.
Configure the storage hardware and network	■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which Exchange will be installed See “Configuring the storage hardware and network” on page 92.
Install SFW HA	■ Install Veritas Storage Foundation HA for Windows See “Installing Veritas Storage Foundation HA for Windows” on page 94.
Configure disk groups and volumes for Exchange Server	■ Plan the storage layout ■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Create dynamic volumes for the mailbox databases and logs See “Configuring cluster disk groups and volumes for Exchange Server” on page 98.

Table 4-2 Task list: Standalone Exchange Server HA configuration tasks

Action	Description
Configure the VCS cluster	■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster ■ If the cluster is already configured, run VCW to add the Exchange server systems to the cluster See “Configuring the cluster” on page 114. See “Adding new Exchange servers to an existing cluster” on page 133.
Move the Exchange mailbox databases to shared storage	■ Use the Exchange Management Console or the Exchange Management Shell to move the Exchange databases to shared storage See “Moving mailbox databases to shared storage” on page 132.
Install and configure Exchange on the additional nodes	■ After moving the Exchange database, install Exchange server on additional nodes, if required See “About installing Exchange Server 2010” on page 130.
Create an Exchange service group	■ Create a Exchange Server 2010 service group using the Exchange 2010 Service Group Configuration Wizard See “About configuring the VCS Exchange Server service group” on page 140.
Verify the HA configuration	■ Test fail over between nodes See “Verifying the Exchange Server cluster configuration” on page 146.

VCS campus cluster configuration

You can install and configure a new Veritas Storage Foundation HA environment for Exchange Server in a campus cluster configuration. A campus cluster environment provides high availability and disaster recovery that extends beyond local clustering and mirroring at a single site, but is not as complex as SFW HA DR solution with replication.

[Table 4-3](#) outlines the high-level objectives and the tasks to complete each objective for a campus cluster configuration for Exchange.

Table 4-3 Task list: Exchange Server campus cluster configuration

Action	Description
Verify hardware and software prerequisites	■ Review the requirements “Reviewing the requirements” on page 70
Review the campus cluster configuration	■ Review the sample configuration “Reviewing the campus cluster configuration” on page 82
Configure storage hardware and network	■ Set up the network and storage for a cluster environment ■ Verify the DNS entries for the systems on which Exchange will be installed “Configuring the storage hardware and network” on page 92
Install SFW HA	■ Install Veritas Storage Foundation HA for Windows “Installing Veritas Storage Foundation HA for Windows” on page 94
Configure disk groups and volumes for Exchange Server	■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Create dynamic volumes for the mailbox databases and logs See “Configuring cluster disk groups and volumes for Exchange Server” on page 98.

Table 4-3 Task list: Exchange Server campus cluster configuration

Action	Description
Configure the VCS cluster	■ Verify static IP addresses and name resolution configured for each node ■ Run the Veritas Cluster Server Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster See “Configuring the cluster” on page 114.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “About installing Exchange Server 2010” on page 130.
Create Exchange databases on shared storage	■ Use the Exchange Management Console or the Exchange Management Shell to create Exchange mailbox databases on shared storage See “Creating mailbox databases on shared storage” on page 131.
Create an Exchange service group	■ Create a Exchange Server 2010 service group using the Exchange 2010 Service Group Configuration Wizard See “About configuring the VCS Exchange Server service group” on page 140.
Set the ForceImport attribute of the VMDg resource as per the requirement	■ If a site failure occurs, set the ForceImport attribute of the VMDg resource to 1 to ensure proper failover See “Setting the ForceImport attribute to 1 after a site failure” on page 151.

VCS Replicated Data Cluster configuration

You can install and configure a new Veritas Storage Foundation HA environment for Exchange Server in a Replicated Data Cluster configuration. The configuration process for a Replicated Data Cluster configuration has the following three main stages:

- Configure the SFW HA and Exchange Server components for high availability on the cluster nodes at the primary zone.
- Install and configure SFW HA and Exchange Server components on the cluster nodes at the secondary zone.
- Configure the VVR components for both zones.
Refer to the *Veritas Volume Replicator Administrator's Guide* for additional details on VVR.

Table 4-4 outlines the high-level objectives and the tasks to complete each objective for a Replicated Data Cluster configuration for Exchange.

Table 4-4 Task list: Exchange server Replicated Data Cluster configuration

Action	Description
Verify hardware and software prerequisites	■ Review the requirements See “ Reviewing the requirements ” on page 70.
Review the RDC configuration	■ Review the sample configuration See “ Reviewing the Replicated Data Cluster configuration ” on page 84.
Configure the storage hardware and network	■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which Exchange will be installed See “ Configuring the storage hardware and network ” on page 92

Table 4-4 Task list: Exchange server Replicated Data Cluster configuration

Action	Description
Install SFW HA	■ Install Veritas Storage Foundation HA for Windows ■ Select the option to install VVR; this also automatically installs the Veritas Cluster Server Agent for VVR ■ If you plan to configure a disaster recovery site in addition to configuring RDC, install the Global Cluster Option (GCO) See “Installing Veritas Storage Foundation HA for Windows” on page 94.
Configure cluster disk groups and volumes for Exchange Server	■ Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) ■ Create dynamic volumes for the mailbox databases and logs See “Configuring cluster disk groups and volumes for Exchange Server” on page 98.
Configure the cluster	■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication in the cluster See “Configuring the cluster” on page 114.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “About installing Exchange Server 2010” on page 130.
Create Exchange databases on shared storage	■ Use the Exchange Management Console or the Exchange Management Shell to create Exchange mailbox databases on shared storage See “Creating mailbox databases on shared storage” on page 131.
Create an Exchange service group	■ Create a Exchange Server 2010 database service group using the Exchange 2010 Service Group Configuration Wizard See “About configuring the VCS Exchange Server service group” on page 140.

Table 4-4 Task list: Exchange server Replicated Data Cluster configuration

Action	Description
Create the primary system zone	■ Create the primary system zone ■ Add the nodes to the primary zone See “Creating the primary system zone” on page 156.
Verify failover within the primary zone	■ Test failover between the nodes in the primary zone See “Verifying the Exchange Server cluster configuration” on page 146.
Create a parallel environment in the secondary zone	■ Install SFW HA on the systems in the secondary zone ■ Configure disk groups and volumes using the same names as on the primary zone ■ Install Exchange Server following the prerequisites and guidelines for installing on the secondary zone See “Creating a parallel environment in the secondary zone” on page 156.
Add the secondary zone systems to the cluster	■ Add the secondary zone systems to the cluster See “Adding the systems in the secondary zone to the cluster” on page 157.
Set up security for VVR on all cluster nodes	■ Set up security for VVR on all nodes in both zones This step can be done at any time after installing SFW HA on all cluster nodes, but must be done before configuring VVR replication. See “Setting up security for VVR” on page 161.
Set up the Replicated Data Set (RDS)	■ Use the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones See “Setting up the Replicated Data Sets (RDS)” on page 164.

Table 4-4 Task list: Exchange server Replicated Data Cluster configuration

Action	Description
Configure a hybrid RVG service group	■ Create a hybrid Replicated Volume Group (RVG) service group ■ Configure the hybrid RVG service group See “ Configuring a hybrid RVG service group for replication ” on page 176.
Set a dependency between the service groups	■ Set up a dependency from the VVR RVG service group to the Exchange Server service group See “ Setting a dependency between the service groups ” on page 186.
Add the nodes from the secondary zone to the RDC	■ Add the nodes from the secondary zone to the RVG service group ■ Configure the IP resources for failover ■ Add the nodes from the secondary zone to the Exchange Server service group See “ Adding the nodes from the secondary zone to the RDC ” on page 187.
Verify the RDC configuration	■ Verify that failover occurs first within zones and then from the primary to the secondary zone See “ Verifying the RDC configuration ” on page 191.

Disaster recovery configuration

You begin by configuring the primary site for high availability. After setting up an SFW HA high availability environment for Exchange on a primary site, you can create a secondary or “failover” site for disaster recovery.

The Disaster Recovery (DR) wizard helps you to clone the storage and service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

The DR wizard also helps you set up replication and the global clustering (GCO option). You can choose to configure replication using Veritas Volume Replicator (VVR) or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See “[About the Solutions Configuration Center](#)” on page 39.

To follow the workflow in the Solutions Configuration Center, the disaster recovery workflow has been split into two tables, one covering the steps for configuring high availability at the primary site, and the other covering the steps for completing the disaster recovery configuration at the secondary site.

DR configuration tasks: Primary site

[Table 4-5](#) outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the primary site.

Table 4-5 Task list: Exchange DR configuration at primary site

Action	Description
Verify hardware and software prerequisites	Review the requirements See “ Reviewing the requirements ” on page 70.
Understand the configuration	Understand the DR configuration See “ Reviewing the disaster recovery configuration ” on page 86.

Table 4-5 Task list: Exchange DR configuration at primary site (Continued)

Action	Description
Configure the storage hardware and network	For all nodes in the cluster: ■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which Exchange will be installed See “Configuring the storage hardware and network” on page 92.
Install SFW HA	■ Install Veritas Storage Foundation HA for Windows on all nodes that will become part of the cluster ■ Select the option to install the Global Cluster Option (GCO). ■ Select the option to install the appropriate replication agents for your configuration. See “Installing Veritas Storage Foundation HA for Windows” on page 94.
Configure the cluster	■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster See “Configuring the cluster” on page 114.
Configure cluster disk groups and volumes for Exchange Server	■ Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) ■ Create dynamic volumes for the mailbox databases and logs See “Configuring cluster disk groups and volumes for Exchange Server” on page 98.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “About installing Exchange Server 2010” on page 130.
Create Exchange databases on shared storage	■ Use the Exchange Management Console or the Exchange Management Shell to create Exchange mailbox databases on shared storage See “Creating mailbox databases on shared storage” on page 131.

Table 4-5 Task list: Exchange DR configuration at primary site (Continued)

Action	Description
Create an Exchange service group	■ Create a Exchange Server 2010 database service group using the Exchange 2010 Service Group Configuration Wizard See “About configuring the VCS Exchange Server service group” on page 140.
Verify the primary site configuration	■ Test failover between nodes on the primary site See “Verifying the Exchange Server cluster configuration” on page 146.

DR configuration tasks: Secondary site

[Table 4-6](#) outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the secondary site.

Table 4-6 Task list: Exchange DR configuration at secondary site

Action	Description
Install SFW HA and configure the cluster on the secondary site	Caution: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster. See “Setting up the secondary site: Installing SFW HA and configuring a cluster” on page 202.
Verify that Exchange Server has been configured for high availability at the primary site	■ Verify that Exchange has been configured for high availability at the primary site and that the service groups are online See “Verifying your primary site configuration” on page 203.
Set up the replication prerequisites	■ Ensure that replication prerequisites for your selected method of replication are met before running the DR wizard See “Setting up security for VVR” on page 204. See “Configuring EMC SRDF replication and global clustering” on page 235. See “Configuring Hitachi TrueCopy replication and global clustering” on page 238.

Table 4-6 Task list: Exchange DR configuration at secondary site (Continued)

Action	Description
(Secure cluster only) Assign user privileges	■ For a secure cluster only, assign user privileges See “Assigning user privileges (secure clusters only)” on page 211.
Start running the DR wizard	■ Review prerequisites for the DR wizard ■ Start the DR wizard and make the initial selections required for each task: selecting a primary site system, the service group, the secondary site system, and the replication method See “Configuring disaster recovery with the DR wizard” on page 212.
Clone the storage configuration (VVR replication only)	(VVR replication option) Clone the storage configuration on the secondary site using the DR wizard See “Cloning the storage on the secondary site using the DR wizard (VVR replication option)” on page 216.
Create temporary storage for application installation (other replication methods)	(EMC SRDF, Hitachi TrueCopy, or GCO only replication option) Use the DR wizard to create temporary storage for installation on the secondary site See “Creating temporary storage on the secondary site using the DR wizard (array-based replication)” on page 220.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “Installing Exchange 2010” on page 224
Clone the service group configuration	■ Clone the service group configuration from the primary to the secondary site using the DR wizard See “Cloning the service group configuration on to the secondary site using the DR wizard” on page 224.

Table 4-6 Task list: Exchange DR configuration at secondary site (Continued)

Action	Description
Configure replication and global clustering, or configure global clustering only	■ (VVR replication) Use the wizard to configure replication and global clustering ■ (EMC SRDF replication) Set up replication and then use the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Set up replication and then use the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Use the wizard to configure global clustering, and then set up replication See “Configuring replication and global clustering” on page 227,
Verify the disaster recover configuration	■ Verify that the secondary site has been fully configured for disaster recovery See “Verifying the disaster recovery configuration” on page 243.
(Optional) Add secure communication	■ Add secure communication between local clusters within the global cluster (optional task) See “Establishing secure communication within the global cluster (optional)” on page 245.
(Optional) Add additional DR sites	■ Optionally, add additional DR sites to a VVR environment See “Adding multiple DR sites (optional)” on page 247.
Handling service group dependencies after failover	■ If your environment includes dependent service groups, review the considerations for bringing the service groups online after failover to the secondary site See “Recovery procedures for service group dependencies” on page 247.

Requirements and Planning

This section contains the following chapter:

- [Chapter 5, “Requirements and planning for your HA and DR configurations”](#)

Requirements and planning for your HA and DR configurations

This chapter contains the following topics:

- [“Reviewing the requirements”](#) on page 70
- [“Reviewing the HA configuration”](#) on page 75
- [“Following the HA workflow in the Solutions Configuration Center”](#) on page 79
- [“Reviewing a standalone Exchange Server configuration”](#) on page 79
- [“Reviewing the campus cluster configuration”](#) on page 82
- [“Campus cluster failover using the ForceImport attribute”](#) on page 83
- [“Reviewing the Replicated Data Cluster configuration”](#) on page 84
- [“Reviewing the disaster recovery configuration”](#) on page 86

Reviewing the requirements

Verify that the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table 5-1](#) estimates disk space requirements for SFW HA.

Table 5-1 Disk space requirements

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

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

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

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW HA 5.1 Service Pack 2 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/358407>
- Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:
<http://entsupport.symantec.com/docs/358406>
- Review the Exchange Server environments supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing Veritas Storage Foundation HA for Windows (SFW HA) Microsoft Exchange Server solutions, ensure that you select the option to install the Veritas Cluster Server Application Agents for Microsoft Exchange

- 2003 or 2007 or 2010, if you plan to configure the respective Exchange versions for high availability.
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
 - When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported Exchange 2010 versions

[Table 36-9](#) lists the Microsoft Exchange Server 2010 versions supported with SFW HA.

Table 5-2 Supported Microsoft Exchange Server 2010 versions

Exchange Server 2010	Windows Servers
Microsoft Exchange Server 2010, (SP1 supported) Standard Edition or Enterprise Edition (Mailbox server role required)	■ Windows Server 2008 with Service Pack 2: Standard x64 Edition, Enterprise x64 Edition ■ Windows Server 2008 R2 on Standard x64 Edition, Enterprise x64 Edition

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
See "[Best practices for SFW HA](#)" on page 74.
- NIC teaming is not supported for the VCS private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

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

- One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.

Permission requirements for SFW HA

The following permissions are required:

- You must be a domain user.
- You must be a member of the local Administrators group on all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements for SFW HA

Please review the following additional requirements:

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

the same drive letter is available on all nodes and that the system drive has adequate space for the installation.

- For a Replicated Data Cluster, install only in a single domain.

Best practices for SFW HA

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.


Reviewing the HA configuration

In a typical example of an Exchange high availability configuration, the Exchange Mailbox Server role is installed on one or more cluster nodes. The Exchange mailbox databases are created on shared storage that is accessible from all the mailbox servers in the cluster. The Exchange mailbox databases are managed by a service group configured with a set of cluster nodes.

The mailbox databases are active on the node where the service group is online. If the active node fails, the mailbox databases are moved to an alternate mailbox server configured in the service group.

Figure 5-1 illustrates a typical two-node Exchange 2010 failover configuration. System1 and System2 are the mailbox servers that are part of the Exchange database service group. When the service group is online on System1, the mailbox databases are active on System1. The Client Access server directs all client requests to the mailbox server on System1. System2 acts as a redundant mailbox server as well as an additional Client Access server at the site. If System1 fails, all the mailbox databases are moved to System2, and the mailbox server on System2 starts accepting client requests for those databases.


Figure 5-1 Exchange 2010 failover configuration: Single service group



In this configuration, System2 functions as a redundant failover target for the mailbox databases that are active on System1. However, you can also configure System2 to host a different set of mailbox databases. You create a separate service group for those databases and then bring the service group online on System2. Thus System1 and System2 can both host mailbox databases and at the same time act as failover targets for each other.

Figure 5-2 illustrates such a configuration where multiple database service groups are configured on multiple mailbox servers; each server hosts a different set of mailbox databases and also serves as a failover target for the other mailbox databases configured in the cluster.

Figure 5-2 Exchange 2010 failover configuration: Multiple service groups



System1, System2, and System3 are the three mailbox servers that host Exchange service groups. Exchange mailbox databases DB1 and DB2 are configured in Exch_SG1 and are active on System1, DB3 and DB4 are configured in Exch_SG2 and are active on System2, and DB5 and DB6 are configured in Exch_SG3 and are active on System3.

All the cluster nodes are part of each database service group which means that each service group can failover on any of the three cluster nodes. If System1 fails, Group 1 (DB1, DB2) is moved to System2. System2 then hosts Group1 and Group2 at the same time. Similarly, if System2 fails, DB3 and DB4 are moved to System3. All the mailbox servers in the cluster host separate mailbox databases while simultaneously act as failover targets for other databases configured in the cluster.

Databases DB1, DB2 and their respective log volumes reside on the same disk group. Similar, DB3, DB4, and DB5, DB6, also reside on the same disk group. When Exch_SG1 service group fails over from System1 to System2, both the databases, DB1 and DB2, are moved to System2. Thus this configuration achieves multiple database mobility.

If you want to control database mobility on a per database basis, you need to configure the database and its volume on an independent disk group.

Sample Exchange server HA configuration

A sample setup is used to illustrate the installation and configuration tasks for an HA configuration.

Table 5-3 shows a sample configuration that will enable failover at the level of a single database. To enable single database failover, you must create an independent disk group for each database, and the disk group must contain the

database volume and log volume for that database only. Then create a service group for that one database.

If you choose to include multiple databases in the same disk group, all the included databases will be part of the same service group and will fail over together to another system in the cluster.

Table 5-3 Sample Exchange 2010 HA configuration objects

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3	Servers
EXCH_SG1	Exchange service group for Exchange database 1
SG1_DG	Cluster disk group for Exchange database 1 The disk group must contain both the database volume and the log volume to ensure successful failover of the database.
SG1_DB1	Volume for storing Microsoft Exchange mailbox database 1
DB1_LOG	Volume for storing Microsoft Exchange mailbox database 1 log file
EXCH_SG2	Exchange service group for Exchange database 2
SG2_DG	Cluster disk group for Exchange database 2 The disk group must contain both the database volume and the log volume to ensure successful failover of the database.
SG2_DB2	Volume for storing Microsoft Exchange mailbox database 2
DB2_LOG	Volume for storing the Microsoft Exchange mailbox database 2 log file

IP addresses required

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


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

Following the HA workflow in the Solutions Configuration Center

The Solutions Configuration Center helps you through the process of installing and configuring a new Veritas Storage Foundation HA environment for Exchange.

[Figure 5-3](#) shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

Figure 5-3 Configuration steps in the Solutions Configuration Center




See [“About the Solutions Configuration Center”](#) on page 39.

Reviewing a standalone Exchange Server configuration

A “standalone” Exchange server is a server that is installed and deployed in a production environment but is not configured for high availability.

[Figure 5-4](#) illustrates a standalone Exchange configuration where System1 and System2 are mailbox servers hosting a set of mailbox databases. System3 is a spare server that may run other application services. The mailbox databases, DB1, DB2, DB3 and DB4 reside on storage that is accessible only to the respective local systems. DB3 and DB4 volumes are not accessible from System1; DB1 and DB2 volumes are not accessible from System2.


Figure 5-4 Exchange 2010 initial standalone configuration



When this standalone Exchange setup is configured for HA, each of the mailbox servers become part of a cluster, the mailbox databases are moved to a shared storage that is accessible to all the mailbox servers, the databases are then managed by a VCS service group that consists of a set of cluster nodes.

Figure 5-5 illustrates the standalone Exchange configuration after it is made highly available. Mailbox servers System1 and System2 each host an Exchange service group that contains a set of mailbox databases. The spare server, System3, now acts as a mailbox server.

Figure 5-5 Exchange 2010 completed standalone cluster configuration



System1 and System3 are part of service group Exch_SG1; System2 and System3 are part of service group Exch_SG2.

Thus, System3 now acts as target failover node for both System1 and System2. If System1 faults, the mailbox databases DB1 and DB2 are moved to System3 and all client requests are directed to the mailbox server on System3. Similarly, if System2 faults, service group Exch_SG2 is brought online on System3 enabling continuous access to mailbox database DB3 and DB4. If System1 and System2 fault at the same time, both the service groups are brought online on System3. All the mailbox databases are active on System3.

Sample standalone Exchange server configuration

A sample setup is used to illustrate the installation and configuration tasks for making a standalone Exchange server highly available.


[Table 5-6](#) describes the objects created and used during the installation and configuration.

Table 5-6 Sample standalone Exchange 2010 HA configuration objects

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3	Servers
EXCH_SG1, EXCH_SG2	Exchange database service groups
SG1_DG	Cluster disk group The disk group must contain both the database volume and the log volume to ensure successful failover of the database.
SG1_DB1	Volume for storing a Microsoft Exchange mailbox database
DB1_LOG	Volume for storing a Microsoft Exchange mailbox database log file

Reviewing the campus cluster configuration

A campus cluster solution allows for clustered systems with mirrored or synchronously replicated storage arrays to be implemented in separate data centers, located either within the same building or separate buildings. A sample campus cluster configuration is a two-node campus cluster with each node in a separate site (Site A or Site B). In this example, each node has its own storage array with the same number of disks and contains mirrored data of the storage on the other array.



The two nodes are located several miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number of disks and contains mirrored data of the storage on the other array.

Plan for an equal number and size of disks on the two sites, because each disk group should contain the same number of disks on each site for the mirrored volumes.

Campus cluster failover using the ForceImport attribute

To ensure proper failover in a VCS campus cluster, you must verify the value of the ForceImport attribute of the VMDg resource. The table below lists failure situations and the outcomes depending on the settings for the ForceImport attribute. You can set this attribute to 1 (forcing the import of the disk groups to the other node) or 0 (not forcing the import).

Use the VCS Java Console or command line to modify the ForceImport attribute.

Table 5-8 Failure situations

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
1) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.	No interruption of service. Remaining disks in mirror are still accessible from the other node.	The Service Group does not failover. 50% of the mirrored disk is still available at remaining site.
2) Zone failure Complete Site failure, all accessibility to the servers and storage is lost.	Manual intervention required to online the Service Group at remaining site. Can not automatically import 50% of mirrored disk.	Automatic failover of Service Group to online site. Force Import must be set to True before site failure to ensure VCS can import 50% of mirrored disk.
3) Split-brain (loss of both heartbeats) If the public network link serves as a low-priority heartbeat, the assumption is made that the link is also lost.	No interruption of service. Can't import disks because the original node still has the SCSI reservation.	No interruption of service. Failover does not occur due to Service Group resources remaining online on the original nodes. Example: Online node has SCSI reservation to own disk.
4) Storage interconnect lost Fibre interconnect severed.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	No interruption of service. Service Group resources remain online, but 50% of the mirror disk becomes detached.

Table 5-8 Failure situations (Continued)

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
5) Split-brain and storage interconnect lost If a single pipe is used between buildings for the Ethernet and storage, this situation can occur.	No interruption of service. Cannot import with only 50% of disks available. Disks on the same node are functioning. Mirroring is not working.	Automatically imports 50% of mirrored disk to the alternate node. Disks online for a short period in both locations but offlined again due to IP and other resources being online on original node. No interruption of service.

Reviewing the Replicated Data Cluster configuration

During the Replicated Data Cluster configuration process you will create virtual IP addresses for the following:

- Replication IP address for the primary zone
- Replication IP address for the secondary zone

You should have these IP addresses available before you start deploying the RDC environment.

Sample Exchange Server Replicated Data Cluster configuration

The sample setup for a Replicated Data Cluster has four servers, two for the primary zone and two for the secondary zone. The nodes form two separate clusters, one at the primary zone and one at the secondary zone.

[Table 5-9](#) describes the objects created and used during the installation and configuration tasks.

Table 5-9 Exchange 2010 sample RDC configuration objects

Name	Object
Primary zone	
SYSTEM1, SYSTEM2	Servers at the primary zone

Table 5-9 Exchange 2010 sample RDC configuration objects

Name	Object
EXCH_SG1	Exchange service group
SG1_DG	Cluster disk group The disk group must contain both the database volume and the log volume to ensure successful failover of the database.
SG1_DB1	Volume for storing a Microsoft Exchange mailbox database
DB1_LOG	Volume for storing a Microsoft Exchange mailbox database log file
SG1_REPLOG	Replicator log volume for VVR replication
Secondary zone	
SYSTEM3, SYSTEM4	Servers at the secondary zone
All the other parameters are the same as on the primary zone.	
RDS and VVR Components	
EXCH_DG1_RDS	Replicated Data Set (RDS) name for Exchange mailbox database
EXCH_DG1_RVG	Replicated Volume Group (RVG) name for Exchange mailbox database
EXCH_RVG_SG	Replication service group for Exchange mailbox database and log files

About setting up a Replicated Data Cluster configuration

In the example, Exchange database is configured as a VCS service group in a four-node cluster, with two nodes in the primary RDC zone and two in the secondary RDC zone. In the event of a failure on the primary node, VCS can fail over the Exchange database to the second node in the primary zone. If the entire primary zone fails, VCS fails over the service group to a node in the secondary zone. The secondary zone now becomes the new primary zone.

The process involves the steps described in the following topics:

- [About setting up replication](#)
- [About configuring and migrating the service group](#)

About setting up replication

You set up replication between the shared disk groups. You use VVR to group the shared data volumes into a Replicated Volume Group (RVG), and create the VVR Secondary on hosts in your secondary zone.

You create a Replicated Data Set (RDS) with the Primary RVG and the Secondary RVG. The Primary RVG consists of volumes shared between the cluster nodes in the first zone (primary zone) and the Secondary RVG consists of volumes shared between the cluster nodes in the second zone (secondary zone).

Use the same disk group name and RVG name in both the zones so that the MountV resources are able to mount the same block devices.

About configuring and migrating the service group

For a successful wide-area failover, the mount points and applications must fail over to the secondary RDC zone. Additionally, the VVR secondary disk group and RVG must be imported and started on the secondary RDC zone.

In the RDC configuration, consider a case where the primary RDC zone suffers a total failure of the shared storage. In this situation, none of the nodes in the primary zone see any device. The service group cannot fail over locally within the primary RDC zone as the shared volumes cannot be mounted on any node. The service group must therefore fail over to a node in the secondary RDC zone.

The RVGPrimary agent ensures that VVR volumes at the secondary RDC zone are made writable. The Exchange mailbox databases are active on a node in the secondary RDC zone until the problem with the local storage is corrected. If the storage problem is corrected, you can switch the service group back to the cluster node in the primary zone.

Before switching the service group back to the original primary RDC zone, you must resynchronize any changed data from the active secondary RDC zone, since the failover. After the resynchronization completes, switch the service group to the primary zone.

Reviewing the disaster recovery configuration

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

If you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM4 and SYSTEM5 on the secondary site), the Exchange database can fail over from SYSTEM1 to SYSTEM2 or vice versa on the primary site, and SYSTEM4 to SYSTEM5 or vice versa on the secondary site.

Figure 5-6 provides a view of a cluster configuration on the primary site.

Figure 5-6 Cluster configuration on the primary site





Figure 5-7 displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.


Figure 5-7 Disaster Recovery environment



In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. When a failure occurs at the primary site, the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients.

Figure 5-8 illustrates this type of failure.

Figure 5-8 Application services restored after primary site failure



You can choose to configure replication using VVR or an agent-supported array-based hardware replication. You can use the DR wizard to configure VVR replication or required options for the VCS agents for EMC SRDF or Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

Deployment

This section contains the following chapters:

- [Chapter 6, “Installing and configuring SFW HA”](#)
- [Chapter 7, “Installing Exchange Server”](#)
- [Chapter 8, “Configuring Exchange Server for failover”](#)
- [Chapter 9, “Configuring campus clusters for Exchange Server”](#)
- [Chapter 10, “Configuring Replicated Data Clusters for Exchange Server”](#)
- [Chapter 11, “Deploying disaster recovery for Exchange Server”](#)
- [Chapter 12, “Testing fault readiness by running a fire drill”](#)

Installing and configuring SFW HA

This chapter contains the following topics:

- [“Configuring the storage hardware and network”](#) on page 92
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 94
- [“Configuring cluster disk groups and volumes for Exchange Server”](#) on page 98
- [“About managing disk groups and volumes”](#) on page 111
- [“Configuring the cluster”](#) on page 114
- [“Adding a node to an existing VCS cluster”](#) on page 122

Configuring the storage hardware and network

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

To configure the hardware

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

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.

- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Application Agent for Exchange. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

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

Installing Symantec Trusted certificate for unsigned drivers

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

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

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

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

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

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.

3 Review the links on the DVD browser panel.

The panel provides the **Late Breaking News** link to access the latest information about updates, patches, and software issues regarding this release, and a link to run the Configuration Checker to verify that your configurations meet all pertinent software and hardware requirements. The panel provides links to install the software (Storage Foundation for Windows or Storage Foundation HA for Windows) and access the documentation (Getting Started Guide, Installation and Upgrade Guide, and Release Notes).

The panel also provides links to access the Veritas Operations Services (VOS) site (VOS provides you four types of detailed reports about your computer and Symantec enterprise products, a checklist of configuration recommendations, and system and patch requirements to install or upgrade your software), contact the Symantec Technical Support, and see the contents of the DVD.

4 Under Install Storage Foundation HA, do one of the following:

- Click the **Complete/Custom** link to install server or client or both the components.
- Click the **Administrative Console** link to install only the client components.

Note: With the Administrative Console option, you will not be prompted for a product license or presented with a list of product options for SFW or SFW HA.

Click the **Complete/Custom** link.

- 5 On the Welcome panel, review the Welcome message and the listed prerequisites. Ensure that the prerequisites are met prior to proceeding. Click **Next**.
- 6 On the License Agreement panel, read the license agreement. If you agree to the license terms, click **I accept the terms of the License Agreement**, and then click **Next**.
- 7 On the License panel, enter the product license key before adding license keys for features. Click **Enter license key(s)**, provide the license key in the field below it, and then click **Add**.
- If you do not have a license key, click **Use embedded evaluation license key** to use the default evaluation license key. This license key is valid only for a limited evaluation period.
- To remove a license key, click the key, and then click **Remove**. To see a license key's details, click the key to display its details in the License key details area.

Click **Next** to continue.

- 8
- On the Option Selection panel, select the appropriate SFW product options, and click **Next**.

Veritas Cluster Server Application Agent for Exchange	Required to configure high availability for Exchange Server.
Client Components	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used while configuring high availability. Required to install the Solutions Configuration Center which provides information and wizards to assist configuration.
Global Cluster Option	Required for a disaster recovery configuration only.
Volume Replicator (VVR)	If you plan to use Veritas Volume Replicator (VVR) for replication, select the option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

- 9
- On the System Selection panel, do the following, and then click **Next**:
- To add a computer for installation, provide the name of the computer in the System Name box.

OR

If you do not know the name of the computer, click **Browse** to search for the computers available in your domain. The Select Systems dialog box appears. Select a computer from the Available Systems area, move it to the Selected Systems area, and then click **OK** to add it for installation.
- To change the installation path of an added computer, click the folder icon for the computer, and then select the installation path in the Browse For Folder dialog box.
- To know the verification status and other information of the added computer, click the information icon.
- To remove an added computer, select it, and then click the recycle bin icon.

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

The installer checks the prerequisites for the added computers and displays the results in the Status column. If a computer fails validation, address the issue, and repeat the validation process by clicking **Re-verify**.

- 10** Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages. If applicable to your installation, perform the procedure mentioned in the messages.

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

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

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

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

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

When installing Veritas Storage Foundation for Windows (Server Components) with the Microsoft Cluster Server (MSCS) option, you may want to adjust the minimum and maximum time for quorum arbitration to ensure optimal functioning of Veritas Storage Foundation for Windows dynamic volumes with MSCS. For more information, see the *Veritas Storage Foundation Administrator's Guide*.

Review the messages, and then click **OK**.

- 11** On the Pre-install Summary panel, the Pre-install Report is displayed with summarized information about the installation. Review the Pre-install Report. Click **Back** to make changes, if necessary. Click **Save Report** to save the report as a web page or text file on your computer.

It is recommended that you select the **Automatically reboot systems after installer completes the operation** check box to restart the computer after the installation is complete.

Click **Install** to install the software.

- 12 The Installation panel displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report, and address the reason for failure. You may have to either repair the installation or uninstall and re-install the software.
- 13 On the Post-install Summary panel, the Post-install Report is displayed with summarized information about the installation results along with links to the log files and installation summary for the computer. Click **Save Report** to save the report as a web page or text file on your computer. Review the Post-install Report and log files, and then click **Next**.
- 14 On the Finish panel, click **Finish** to complete the installation.
- 15 Click **Yes** to restart the local node.

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

To reset the driver signing options

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

Configuring cluster disk groups and volumes for Exchange Server

Before installing Exchange Server, you can create cluster disk groups and volumes using the Veritas Enterprise Administrator (VEA) console installed with SFW.

Note: Cluster disk groups and volumes are used to store the Exchange mailbox databases and log files and are therefore not required for the actual Exchange installation. You can choose to configure the storage after the Exchange installation.

Planning cluster disk groups and volumes is covered in the following topics:

- [“About cluster disk groups and volumes”](#) on page 99
- [“Prerequisites for configuring cluster disk groups and volumes”](#) on page 100
- [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 101
- [Considerations for disks and volumes for campus clusters](#)
- [“Considerations for volumes for a VVR configuration”](#) on page 102
- [“Viewing the available disk storage”](#) on page 103

Configuring cluster disk groups and volumes is covered in the following topics:

- [“Viewing the available disk storage”](#) on page 103
- [“Creating a cluster disk group”](#) on page 104
- [“Creating volumes”](#) on page 105

About cluster disk groups and volumes

SFW uses disk groups to organize disks or LUNs for management purposes. A dynamic disk group is a collection of disks that is imported or deported as a single unit. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

You create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in a high-availability cluster by first deporting the cluster disk group from the current node and then importing it on the desired node. In a campus cluster, the volumes are mirrored across the storage arrays.

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

Prerequisites for configuring cluster disk groups and volumes

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

- The type of volume configurations that are required
- The number of volumes or LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs that depend on the traffic load
- The disk groups and number of disks on each site
- Types of volumes required and location of the plex of each volume in the storage array

Complete the following tasks before you create the cluster disk group and volumes for Exchange:

- Determine the layout or configuration for each volume and the total number of disks needed.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the Expand Volume command but you can not decrease the size.
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must rescan, and if necessary, use the Write Signature command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

For Exchange 2010, there are additional requirements if you want to configure the disk group and volumes so that each Exchange database can fail over independently of any others.

See “[Sample disk group and volume configuration for Exchange 2010](#)” on page 103.

For standalone Exchange configurations

If the existing databases and logs are already on shared storage, read the following topic:

- “[Considerations for converting existing shared storage to cluster disk groups and volumes](#)” on page 101

For Campus Cluster configurations

For campus clusters, each disk group must contain an equal number of disks on each site. Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

For more information on disk groups and volumes for campus clusters, read the following topic:

- [“Considerations for disks and volumes for campus clusters”](#) on page 101

For configurations using VVR

For a Replicated Data Cluster (RDC) configuration or a disaster recovery (DR) configuration using Veritas Volume Replicator (VVR), read the following topic:

- [“Considerations for volumes for a VVR configuration”](#) on page 102

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

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

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

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

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

For a disaster recovery configuration using Veritas Volume Replicator, you need to allow additional disk space for a Storage Replicator Log volume.

See [“Considerations for volumes for a VVR configuration”](#) on page 102.

Considerations for disks and volumes for campus clusters

Ensure that each disk group has the same number of disks on each site. Each volume must be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

Consider the following when creating new volumes:

- For campus clusters, when creating a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.
- Symantec recommends using either simple mirrored (concatenated) or striped mirrored options for the new volumes. Striped mirrored gives you better performance compared to concatenated.
When selecting striped mirrored, select two columns in order to stripe one enclosure that is mirrored to the second enclosure.
- You cannot select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

Considerations for volumes for a VVR configuration

For a configuration using Veritas Volume Replicator (VVR), either a disaster recovery (DR) configuration on a secondary site or a Replicated Data Cluster (RDC), note the following:

- VVR does not support the following types of volumes:
 - SFW (software) RAID 5 volumes
 - Volumes with the Dirty Region Log (DRL)
 - Data Change Object (DCO)
 - Volumes with commas in the names
- A configuration with VVR requires a Storage Replicator Log (SRL) volume for each disk group that contains volumes that are replicated. You can create the SRL volume when configuring the other volumes for the application or you can create it later when you set up replication. If you create it later, ensure that you allow sufficient disk space for this volume. For more about VVR planning, see the *Veritas Volume Replicator Administrator's Guide*.
- Do not assign a drive letter to the Storage Replicator Log volume. This will limit access to that volume and avoid potential data corruption.

Sample disk group and volume configuration for Exchange 2010

When configuring disk groups and volumes for Exchange 2010 mailbox databases, you need to consider whether you want each mailbox database to be able to fail over independently, or whether you want to set up service groups that include multiple databases that will fail over as a group.

To enable single mailbox database failover, you must create an independent disk group for each database, and the disk group must contain the database volume and log volume for that database only. Then create a service group for that one database.

Each disk group must contain both the database volume and the log volume to ensure successful failover of the database.

If you choose to include multiple databases in the same disk group, all the included databases will be part of the same service group and will fail over together to another system in the cluster.

The following would be an example configuration for the single database failover scenario in which there are two Exchange mailbox databases:

- SG1_DG is a disk group that contains the following:
 - SG1_DB1 is the volume for Exchange mailbox database 1.
 - DB1_LOG is the volume that contains the transaction log for the database.
- SG2_DG is a separate disk group that contains the following:
 - SG1_DB2 is the volume for Exchange mailbox database 2.
 - DB2_LOG is the volume that contains the transaction log for the database.

Viewing the available disk storage

Before creating disk groups and volumes you may want to view available disk storage.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.

To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.
The internal names for the disks that the current system can access for available storage are displayed, with names Harddisk1, Harddisk2, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a cluster disk group


Use the Veritas Enterprise Administrator (VEA) to create a cluster disk group on the node where Exchange is installed. Repeat the procedure if you want to create additional disk groups.

To create a dynamic (cluster) disk group

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

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

6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, SG1_DG).
- Check the **Create cluster group** check box.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Creating volumes

This procedure will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure to create additional volumes.

Before you begin, make sure you review the following topics, if applicable to your environment:


- [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 101
- [“Considerations for disks and volumes for campus clusters”](#) on page 101
- [“Considerations for volumes for a VVR configuration”](#) on page 102

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example SG1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.




- Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
- To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
- You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the volume attributes.


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume. If you click the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
- Select a layout type.
For campus clusters, select either **Concatenated** or **Striped**.
- If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
For campus clusters, if you select **Striped**, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
- To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
For campus clusters, you select the **Mirrored** checkbox for either layout type.
- In the Mirror Info area, select the appropriate mirroring options.
For campus clusters, in the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
- Verify that **Enable logging** is not selected.
- Click **Next**.

- 8 Assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
 - To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
 - To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
 - If creating a Replicator Log volume for Veritas Volume Replicator, select **Do not assign a drive letter**.



- 9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- For a VVR configuration, for the Replicator Log volume only, clear the **Format this volume** check box.
- Select an allocation size or accept the default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

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

12 Repeat these steps to create additional volumes.

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

About managing disk groups and volumes

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

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

Importing a disk group and mounting a volume

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

To import a disk group

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

To mount a volume

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

Unmounting a volume and deporting a disk group

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

To unmount a volume and deport the dynamic disk group


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

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.

- 4 Select one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Deporting the cluster disk group

You can manually move ownership of the cluster disk group from one node to another by first deporting the disk group and then importing it on the desired node.

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name where the disk group is current imported, expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses and that name resolution is configured for each node.
- Verify that you have the required privileges.
See [“Reviewing the requirements”](#) on page 70.

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.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS, including instructions on adding cluster nodes or removing or modifying cluster configurations.

To configure a VCS cluster

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

Do one of the following:

- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.Proceed to [step 8](#) on page 115.
- To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box.

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

- Click **Next**.

If you chose to retrieve the list of systems, proceed to [step 6](#) on page 115. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**.

Do not specify systems that are part of another cluster.

Proceed to [step 8](#) on page 115.

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

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

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

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

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

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

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

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

Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
--------------	--

Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255. Caution: If you chose to specify systems and users manually in step 4 or if you share a private network between more than one domain, make sure that the cluster ID is unique.
Operating System	From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.
Available Systems	Select the systems that you wish to configure in the cluster. Check the Select all systems check box to select all the systems simultaneously. The wizard discovers the network interface cards (NICs) on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click Yes to configure a private link heartbeat.

- 10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.
If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 117.
- 11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:
 - To configure the VCS private network over the ethernet, complete the following steps:
 - Select the check boxes next to the two NICs to be assigned to the private network.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Select **Configure LLT over UDP**.
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Symantec recommends reserving at least two NICs exclusively for the VCS private network.
 - 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, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
 - For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.
- The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD),

which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

This account does not require Domain Administrator privileges.

Specify a domain user as follows:

- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 114, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the Create New User field, and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

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

Do one of the following:

- To use the single sign-on feature, complete the following steps:
 - Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify an existing root broker or a node in this cluster to configure as root broker**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers.

Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities.

These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.
- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 6-1 contains some more examples of search criteria.

Table 6-1 Search criteria examples

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

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

- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

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

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

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

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

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

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

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

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

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

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on the Notification resource.

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

- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See [“Configuring notification”](#) on page 121.

Configuring notification

This section describes steps to configure notification.

To configure notification

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

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

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.
 - Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click ‘+’ to add a field; click ‘-’ to remove a field.
 - Enter an SNMP trap port. The default value is “162”.
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.

- Type the name of the SMTP server.
 - Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
 - Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
 - Click + to add fields; click - to remove a field.
- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.
 - If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
 - 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
 - 6 Click **Configure**.
 - 7 Click **Finish** to exit the wizard.

Adding a node to an existing VCS cluster

You use the VCS Cluster Configuration Wizard (VCW) to add one or more nodes to an existing VCS cluster.

Prerequisites for adding a node to an existing cluster

Prerequisites for adding a node to an existing cluster are as follows:


- Verify that the logged-on user has VCS Cluster Administrator privileges.
- The logged-on user must be a local Administrator on the system where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select Services on the Administrative Tools menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node on which you run the wizard. Select Services on the Administrative Tools menu and verify that the Veritas High Availability Daemon is running.

Adding a node to a VCS cluster

Use the VCS Cluster Configuration Wizard (VCW) to add a node to the existing cluster.

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 In the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



Do one of the following:


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

Proceed to [step 8](#) on page 126.

- To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.


If you chose to retrieve the list of systems, proceed to [step 6](#) on page 125. Otherwise proceed to the next step.

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



- Type the name of a node in the cluster and click **Add**.
 - Type the name of the system to be added to the cluster and click **Add**.
- If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.
- Proceed to [step 8](#) on page 126.

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



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

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.


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

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

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

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

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**.
If you chose to specify the systems manually in [step 4](#), only the clusters configured with the specified systems are displayed.
- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**.
In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.
The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.
- 11 On the Cluster Details panel, check the check boxes next to the systems to be added to the cluster and click **Next**.



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

- 12 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.
If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.
- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**.
How you configure the VCS private network communication depends on

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

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - Select the check boxes next to the two NICs to be assigned to the private network.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.
The wizard will configure the LLT service (over ethernet) on the selected network adapters.
- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Check the **Configure LLT over UDP** check box.
 - Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
 - Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

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

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

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

- 14 On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.

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

Installing Exchange Server

This chapter contains the following topics:

- [“About installing Exchange Server 2010”](#) on page 130
- [“Before you install Exchange Server 2010”](#) on page 130
- [“Installing Exchange Server 2010”](#) on page 131
- [“Creating mailbox databases on shared storage”](#) on page 131
- [“Moving mailbox databases to shared storage”](#) on page 132
- [“Adding new Exchange servers to an existing cluster”](#) on page 133

About installing Exchange Server 2010

Installing Exchange Server 2010 in a VCS environment is a simple and straightforward process. Unlike the case for earlier versions of Microsoft Exchange, you do not need to run the Exchange Setup Wizard for VCS to perform any of the pre-installation and post-installation tasks for Exchange 2010. You simply run Microsoft Exchange 2010 Setup to install Exchange on the required systems.

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.

Note: You do not need to run the Exchange Setup Wizard for VCS before and after installing Exchange 2010.

Before you install Exchange Server 2010

Verify the following prerequisites before you install Exchange in a VCS environment:

- Ensure that the systems meet the minimum requirements for installing and configuring Exchange 2010. Refer to the Microsoft documentation for details.
- Symantec recommends that the Dynamic Update option for the DNS server be set to "Secure Only."
- VCS requires Microsoft Exchange to be installed on the same local drive on all nodes. For example if you install Exchange on drive C of one node, installations on all other nodes must be on their respective C drives. Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- If your cluster has an SQL service group configured, make sure to install Exchange Server on a node that is not in the SystemList attribute for the SQL service group.

Privileges required for installing Exchange 2010

Verify that the following privileges are available to the user installing Exchange:

- The logged-on user must be a domain user.

- The logged-on user must be a member of the local Administrators group on all nodes where Exchange will be installed.

Installing Exchange Server 2010

Run Microsoft Exchange 2010 Setup to install Exchange on the required systems. 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.
- For a standalone Exchange configuration, install Exchange on any additional systems that you plan to add to the Exchange service group.
- 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. Do not use the Exchange Setup Wizard for VCS.

Verify the following before you create mailbox databases:

- Create the disk group and volumes required for the Exchange mailbox databases and log files and ensure that the disk group is imported and the required volumes mounted on the system from where you will create the databases.
 - See [“Viewing the available disk storage”](#) on page 103.
 - See [“Creating a cluster disk group”](#) on page 104.
 - See [“Creating volumes”](#) on page 105.
 - See [“Importing a disk group and mounting a volume”](#) on page 111.
 - See [“Adding drive letters to mount the volumes”](#) on page 112.
- 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.

Moving mailbox databases to shared storage

Note: This is applicable if you are planning to configure an existing standalone Exchange Server setup for high availability.

Verify the location where all the existing Exchange mailbox databases and logs reside. If they reside on shared disks that are accessible from all the systems where Exchange is installed, then convert the shared disks to cluster disk groups and volumes.

See [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 101.

If the databases and log files do not reside on shared storage, you must move the databases to the appropriate cluster disk groups and volumes on the shared storage. 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.

Note: Do not use the Exchange Setup Wizard for VCS to move Exchange mailbox databases to shared storage. Use the Exchange Management Console.

To move the existing mailbox databases to shared storage

- 1 Verify that you have backed up your existing data.
- 2 Create the disk group and volumes required for the Exchange mailbox databases and log files and ensure that the dynamic cluster disk group is imported and the required volumes mounted on the system where the original database and log files are located on the local drives.
 - See [“Viewing the available disk storage”](#) on page 103.
 - See [“Creating a cluster disk group”](#) on page 104.
 - See [“Creating volumes”](#) on page 105.
 - See [“Importing a disk group and mounting a volume”](#) on page 111.
 - See [“Adding drive letters to mount the volumes”](#) on page 112.
- 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.

Adding new Exchange servers to an existing cluster

Note: This is applicable if you are planning to configure an existing standalone Exchange Server setup for high availability.

This procedure is required only if you have an existing VCS cluster that is running other applications in your environment, and you want to bring your standalone Exchange servers into that cluster.

Use the VCS Cluster Configuration Wizard (VCW) to add the Exchange server to the existing cluster.

Note: You can run VCW from the standalone Exchange server or from a node in the cluster.

Prerequisites for adding a node to an existing cluster are as follows:

- Verify that the logged-on user has VCS Cluster Administrator privileges.
- The logged-on user must be a local Administrator on the system where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select Services on the Administrative Tools menu and verify that the Veritas Command Server shows that it is started.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node on which you run the wizard. Select Services on the Administrative Tools menu and verify that the Veritas High Availability Daemon is running.

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 In the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.

Do one of the following:

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

Proceed to [step 8](#) on page 135.

- To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.

If you chose to retrieve the list of systems, proceed to [step 6](#) on page 134. Otherwise proceed to the next step.

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

- Type the name of a node in the cluster and click **Add**.
- Type the name of the system to be added to the cluster and click **Add**.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

Proceed to [step 8](#) on page 135.

- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems. Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

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

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

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.

- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

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

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**. If you chose to specify the systems manually in [step 4](#), only the clusters configured with the specified systems are displayed.
- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**. In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.

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

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

- 12 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.

If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.

- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**. How you configure the VCS private network communication depends on how it is configured in the cluster. If LLT is configured over ethernet, you have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have to use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - Select the check boxes next to the two NICs to be assigned to the private network.

Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Select **Configure LLT over UDP**.
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Symantec recommends reserving at least two NICs exclusively for the VCS private network.
 - 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, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
 - For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 14 On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.

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

Configuring Exchange Server for failover

This chapter contains the following topics:

- [“About configuring the VCS Exchange Server service group”](#) on page 140
- [“Prerequisites for configuring the Exchange Server service group”](#) on page 140
- [“Creating the Exchange Server 2010 service group”](#) on page 141
- [“Verifying the Exchange Server cluster configuration”](#) on page 146
- [“Determining additional steps needed”](#) on page 147

About configuring the VCS Exchange Server 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.

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.

Read the following topics:

- [Prerequisites for configuring the Exchange Server service group](#)
- [Creating the Exchange Server 2010 service group](#)

Prerequisites for configuring the Exchange Server service group

Note the following prerequisites for configuring the Exchange service group:

- Verify that you have completed all the steps mentioned in the high availability, campus cluster, or RDC workflows, up through the step of installing Exchange and creating the mailbox database on shared storage. See the following topics as appropriate:
 - See [“High availability \(HA\) configuration \(New Server\)”](#) on page 51.
 - See [“High availability \(HA\) configuration \(Existing Server\)”](#) on page 53.
 - See [“VCS campus cluster configuration”](#) on page 55.
 - See [“VCS Replicated Data Cluster configuration”](#) on page 57.
- Verify that the logged-on user has VCS Cluster Administrator privileges. This user classification is required to create and configure VCS service groups.
- The logged-on user must be a local Administrator on the node where you run the wizard.
- 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.
- Verify that the Veritas High Availability Daemon (HAD) is running on the node on which you run the wizard. Select Services on the Administrative Tools menu and verify that the Veritas High Availability Daemon is running.

- Import the cluster disk group and mount the shared volumes created to store the following data:
 - Exchange database
 - Database log

Mount them on the node where you run the wizard; unmount the drives from other nodes in the cluster.
- Verify that Microsoft Exchange Mailbox Server role is installed on all nodes that will be part of the service group.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service, `%vcs_home%\bin\CmdServer.exe`.
 Here `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141

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


Creating the Exchange Server 2010 service group

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

To configure the Exchange service group

- 1 Start the Exchange 2010 Database Configuration Wizard.
 From the Solutions Configuration Center, expand the Solutions for Microsoft Exchange Server tab and click **High Availability (HA) Configuration (New Server) > Configure Exchange service group > Exchange 2010 Configuration Wizard**.
 or
 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange 2010 Configuration Wizard**.
- 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.
- Click to select a database and then click 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 disk group where the selected database resides. Even if you select just one database, the wizard configures all the other databases residing on that disk group. Thus, all databases that reside on the same disk group are part of the same Exchange service group.
 - The wizard configures one service group per cluster disk group. If you select databases that reside on two separate disk groups, the wizard configures them in two separate service groups.
 - If databases are created on a disk group 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 also 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 cluster disk groups and 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_<dbname>. Here, <dbname> is the name of the database on the respective cluster disk group.

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 On the System Selection panel, specify the systems that will be part of the service group and then click **Next**.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the 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 cluster disk groups, the wizard creates multiple service groups, one for each cluster disk group. 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.
- 7 On the Network Configuration 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.
- 8 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.

- 9 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.

- 10 In the Completing the Exchange Configuration panel, select the **Bring the service group online** check box to bring the service group online on the local system and then click **Finish**.

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 disk groups, then you must run the wizard again (in the configure or modify mode) to configure the newly added databases for high availability.

- Verify the Transaction Log Path for the selected Exchange virtual server.
- Click **Next**.

Verifying the Exchange Server cluster configuration

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

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

The service group you selected is taken offline and brought online on the node that you selected.

Determining additional steps needed

This completes the high availability configuration steps for Exchange Server. Depending on the configuration being deployed, there are additional steps that you must perform to set up and complete the configuration.

[Table 8-1](#) contains a list of references to the chapters that describe configuration specific tasks in detail. Proceed to the desired chapter depending on the desired configuration.

You must perform the configuration specific tasks only after you complete the high availability steps mentioned in this and the earlier chapters.

Table 8-1 Additional Exchange Server configuration steps

Tasks	Refer to
Setting up a campus cluster configuration for Exchange Server	“VCS campus cluster configuration” on page 55
Setting up a replicated data cluster configuration for Exchange Server	“VCS Replicated Data Cluster configuration” on page 57
Setting up a disaster recovery configuration for Exchange Server	“Disaster recovery configuration” on page 61
Configuring and running a fire drill for Exchange Server configuration	“About disaster recovery fire drills” on page 255

Configuring campus clusters for Exchange Server

This chapter contains the following topics:

- [“Tasks for configuring campus clusters”](#) on page 150
- [“Verifying the campus cluster: Switching the service group”](#) on page 150
- [“Setting the ForceImport attribute to 1 after a site failure”](#) on page 151

Tasks for configuring campus clusters

In campus clusters you begin by configuring a high availability cluster and then continue with the steps specific to the campus cluster configuration.

Refer to the campus cluster configuration workflow table for a complete list of configuration steps.

See “[VCS campus cluster configuration](#)” on page 55.

[Table 9-1](#) shows the steps specific to the campus cluster configuration that are done after configuring high availability on the nodes.

Table 9-1 Completing campus cluster configuration

Action	Description
Verify the campus cluster configuration	Verify that failover occurs between the nodes. See “ Verifying the campus cluster: Switching the service group ” on page 150.
Set the ForceImport attribute	In case of a site failure, you may have to set the ForceImport attribute to ensure proper failover. See “ Setting the ForceImport attribute to 1 after a site failure ” on page 151.

Verifying the campus cluster: Switching the service group

Failover simulation is an important part of configuration testing.

To verify the campus cluster is functioning properly

- 1 Bring the service group online on one node as follows:
 - In the Cluster Explorer configuration tree, right-click the service group.
 - Click **Online**, and click the appropriate system from the menu.
- 2 Switch the service group to the other node as follows:
 - In the Cluster Explorer configuration tree, right-click the service group.
 - Click **Switch To**, and click the appropriate system from the menu.

Setting the ForceImport attribute to 1 after a site failure

ForceImport is a flag that defines whether the agent forcibly imports the disk group when exactly half the disks are available. The value 1 indicates the agent imports the configured disk group when half the disks are available. The value 0 indicates it does not. Default is 0. This means that the disk group will be imported only when SFW acquires control over the majority of the disks.

Warning: Set this attribute to 1 only after verifying the integrity of your data. If due caution is not exercised before setting this attribute to 1, you risk potential data loss.

You must set the ForceImport attribute for the VMDg resource to 1 after a site failure to ensure proper failover.

To set the ForceImport attribute to 1 from the Java Console

- 1 From the Cluster Explorer configuration tree, select the VMDg resource in the Exchange Server service group.
- 2 In the Properties View, click the **Edit** icon for the **ForceImport** attribute.
- 3 In the Edit Attribute dialog box, make the following selections:
 - Select the **Per System** option.
 - Select the system in Site B.
 - Select the **ForceImport** check box.
 - Click **OK**.
- 4 From the **File** menu of Cluster Explorer, click **Close Configuration**.
- 5 After the failover takes place, revert the ForceImport attribute to its original value.

To set the ForceImport attribute to 1 from the command line

- ◆ Use the following command for implementing the force import setting in VCS:

```
hares -modify <vmdg_resource_name> ForceImport 1|0
```

Example:

```
hares -modify vmdg_Dg1 ForceImport 1
```

Import is forced on vmdg_Dg1.

Configuring Replicated Data Clusters for Exchange Server

This chapter contains the following topics:

- [“Tasks for configuring Replicated Data Clusters for Exchange Server”](#) on page 153
- [“Creating the primary system zone”](#) on page 156
- [“Creating a parallel environment in the secondary zone”](#) on page 156
- [“Adding the systems in the secondary zone to the cluster”](#) on page 157
- [“Setting up security for VVR”](#) on page 161
- [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 164
- [“Configuring a hybrid RVG service group for replication”](#) on page 176
- [“Setting a dependency between the service groups”](#) on page 186
- [“Adding the nodes from the secondary zone to the RDC”](#) on page 187
- [“Verifying the RDC configuration”](#) on page 191
- [“Additional instructions for GCO disaster recovery”](#) on page 193

Tasks for configuring Replicated Data Clusters for Exchange Server

For a Replicated Data Cluster (RDC) you begin by configuring a high availability cluster on the primary zone systems.

You then continue with the steps specific to the RDC configuration.

For the complete RDC configuration workflow see “[VCS Replicated Data Cluster configuration](#)” on page 57.

[Table 10-1](#) shows the steps specific to the RDC configuration that are done after configuring high availability on the primary zone.

Table 10-1 Completing the configuration of a Replicated Data Cluster

Action	Description
Create the primary system zone	■ Create the primary system zone ■ Add the nodes to the primary zone See “ Creating the primary system zone ” on page 156.
Verify failover within the primary zone	See “ Verifying the Exchange Server cluster configuration ” on page 146.
Create a parallel environment in the secondary zone	■ Install SFW HA on the systems in the secondary zone ■ Configure disk groups and volumes using the same names as on the primary zone ■ Install Exchange Server on the systems in the secondary zone See “ Creating a parallel environment in the secondary zone ” on page 156.
Add the secondary zone systems to the cluster	Add the secondary zone systems to the cluster See “ Adding the systems in the secondary zone to the cluster ” on page 157.
Set up security for VVR on all cluster nodes	Set up security for VVR on all nodes in both zones This step can be done at any time after installing SFW HA on all cluster nodes, but must be done before configuring VVR replication. See “ Setting up security for VVR ” on page 161.
Set up the Replicated Data Set	Use the Setup Replicated Data Set Wizard to create Replicated Data Sets and start replication for the primary and secondary zones See “ Setting up the Replicated Data Sets (RDS) ” on page 164.

Table 10-1 Completing the configuration of a Replicated Data Cluster


Action	Description
Configure a hybrid RVG service group	■ Create a hybrid Replicated Volume Group (RVG) service group ■ Configure the hybrid RVG service group See “ Configuring a hybrid RVG service group for replication ” on page 176.
Set a dependency between the service groups	Set up a dependency from the RVG service group to the Exchange Server service group See “ Setting a dependency between the service groups ” on page 186.
Add the nodes from the secondary zone to the RDC	■ Add the nodes from the secondary zone to the RVG service group ■ Configure the IP resources for failover ■ Add the nodes from the secondary zone to the Exchange Server service group See “ Adding the nodes from the secondary zone to the RDC ” on page 187.
Verify the RDC configuration	■ Verify that failover occurs first within zones and then from the primary to the secondary zone See “ Verifying the RDC configuration ” on page 191.

Creating the primary system zone

In the service group, set up systems in the primary zone (zone 0) to specify that initial fail over occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 Select the Exchange service group (EXCH_SG1) in the left pane and then click **Properties** tab in the right pane.
- 3 In the Properties pane, click **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the SystemZones attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.



- 7 Click **OK**.
- 8 After setting up the primary system zone, you can verify the service group failover on systems within the primary zone.
See “[Verifying the Exchange Server cluster configuration](#)” on page 146.

Creating a parallel environment in the secondary zone

After setting up a SFW HA environment in the primary zone, you set up a parallel environment in the secondary zone (zone1).
Before you begin to configure the secondary zone, ensure that the service group is online in the primary zone.

Then complete the following tasks to configure the secondary zone, using the guidelines shown:

- [“Configuring the storage hardware and network”](#) on page 92
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 94
- [“Configuring cluster disk groups and volumes for Exchange Server”](#) on page 98
 During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as that at the primary zone:
 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- [“Installing Exchange Server”](#) on page 129
 Install Exchange on all nodes in the secondary zone.

Note: After you install Exchange on the nodes in the secondary zone, make sure to use VEA to remove all the drive letters from the configured volumes to avoid conflicts during the configuration of the zones.

- See [“Adding the systems in the secondary zone to the cluster”](#) on page 157. You do not create another cluster in the secondary zone. Instead you add the systems to the existing cluster.
 Note that you do not create another Exchange service group in the secondary zone. You continue with the remaining VVR configuration tasks, during which the secondary zone nodes are added to the Exchange service group.

For the complete RDC workflow, see [“VCS Replicated Data Cluster configuration”](#) on page 57.

Adding the systems in the secondary zone to the cluster

Add the systems in the secondary zone (zone1) to the existing cluster with the following procedure.

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.

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

Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.

- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 In the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.

Do one of the following:

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

Proceed to [step 8](#) on page 159.

- To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.

If you chose to retrieve the list of systems, proceed to [step 6](#) on page 158. Otherwise proceed to the next step.

- 5 On the System Selection panel, complete the following and click **Next**.
 - Type the name of a node in the cluster and click **Add**.
 - Type the name of the system to be added to the cluster and click **Add**.
If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.
Proceed to [step 8](#) on page 159.
- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems.
Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

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

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

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

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

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**. If you chose to specify the systems manually in [step 4](#), only the clusters configured with the specified systems are displayed.
- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**. In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.

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

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

- 12 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**. If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.

- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**. How you configure the VCS private network communication depends on how it is configured in the cluster. If LLT is configured over ethernet, you have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have to use the same on the nodes being added.

Do one of the following:

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

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Select **Configure LLT over UDP**.
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Symantec recommends reserving at least two NICs exclusively for the VCS private network.
 - 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, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

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

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

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

- 14 On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.

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

Setting up security for VVR

If you are using Veritas Volume Replicator (VVR) replication, you must configure the VxSAS service on all cluster nodes. For a Replicated Data Cluster environment, you configure the service on all nodes in both the primary and secondary zones.

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

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.

- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

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

To configure the VxSAS service

- 1 Launch the VVR Security Service Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard**.
or
Type `vxsasscfg.exe` at the command prompt.
- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows and then click **Next**:


Account name (domain\account)	Enter the administrative account name.
Password	Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong and then click **Next**:

Selecting domains	The Available domains pane lists all the domains that are present in the Windows network neighborhood. Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.
Adding a domain	If the domain name that you require is not displayed, click Add domain . This displays a dialog that allows you to specify the domain name. Click Add to add the name to the Selected domains list.

- 5 On the Host Selection panel, select the required hosts:



Selecting hosts

The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a host

If the host name you require is not displayed, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 6 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.
Click **Back** to change any information you had provided earlier.
- 7 Click **Finish** to exit the wizard.

Setting up the Replicated Data Sets (RDS)

Set up the Replicated Data Sets (RDS) in the primary zone (zone0) and secondary zone (zone1). You can configure an RDS for both zones using the Setup Replicated Data Set Wizard.

Prerequisites for setting up the RDS for the primary and secondary zones

Before you run the Setup Replicated Data Set Wizard, verify the following:

- Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- Verify that the cluster disk group is imported and the volumes are mounted in the primary and secondary zone
- Verify that you have configured security for VVR
See “[Setting up security for VVR](#)” on page 161.

Creating the Replicated Data Sets with the wizard

To create the Replicated Data Set

- 1 Use the Veritas Enterprise Administrator (VEA) console to launch the Setup Replicated Data Set Wizard from the cluster node on the Primary where the cluster disk group is imported:
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Setup Replicated Data Set**.
- 3 Read the information on the Welcome page and then click **Next**.

- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG) and then click **Next**.

Setup Replicated Data Set Wizard

Enter names for Replicated Data Set and Replicated Volume Group

Select the desired Primary host from the list of connected hosts.

Replicated Data Set name :

Replicated Volume Group name :

Primary Host :


Veritas Enterprise Administrator (VEA) should be connected to the desired Primary host.

< Back Next > Cancel Help

By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.


- 5
- Select from the table the dynamic disk group and data volumes that will undergo replication and then click **Next**.



To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

6 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (SG1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name	Enter the name for the volume in the Name field.
Size	Enter a size for the volume in the Size field.
Layout	Select the desired volume layout.

Disk Selection

Enables you to specify the disk selection method.

- Enable the **Thin Provisioned Disks Only** check box to ensure that the Replicator Log volume is created only on Thin Provisioned (TP) disks.

Note: The check box will remain disabled if the diskgroup does not have any TP disk.

If this option is selected along with the **Select disks automatically** option, then the Replicator Log volume will be created only on TP disks. However, if you enable this check box along with **Select disks manually** option, then the user can select only TP disks from **Available Disks**.

For more information on Thin Provisioning refer to the *Veritas Storage Foundation Administrator's Guide*.

- Choose the **Select disks automatically** option if you want VVR to select the disks.
- Choose the **Select disks manually** option to use specific disks from the Available disks pane for creating the volume. Either double-click on it or select **Add** to move the disks into the Selected disks pane.

- Click **OK** to create the Replicator Log volume.
- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.

7 Review the information on the summary page and click **Create Primary RVG**.

8 After the Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?

Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.

Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.

9 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field and then click **Next**.

If the Secondary host is not connected to VEA, the wizard tries to connect it when you click Next. This wizard allows you to specify only one Secondary

host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.

- 10 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

- the same or larger amount of space as that on the Primary
- Enough space to create volumes with the same layout as on the Primary. Otherwise, the RDS setup wizard enables you to create the required volumes manually.
- Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
- Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.

- 11 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.

This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.

- If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.

Depending on the discrepancies between the volumes on the primary site and the volume, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.

When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.

- 12
- Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

Setup Replicated Data Set Wizard

Edit replication settings

Edit replication settings or click next.

Primary side IP

10.217.53.214

Secondary side IP

10.217.53.215

Replication Mode

Synchronous Override

Replicator Log Protection

AutoDCM

Primary RLINK Name

Pri_RLINK

Secondary RLINK Name

Sec_RLINK

Advanced

DHCP addresses are not supported by VVR.

< Back

Next >

Cancel

Help

-
- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP	Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Secondary side IP	Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode	Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous. The default is synchronous override. Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously. If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as MISSING.
Replicator Log Protection	The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows. The DCM option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them. The Off option disables Replicator Log Overflow protection. In the case of the Bunker node. Replicator Log protection is set to Off, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.


If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK Name	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

- Click **Next** to start replication with the default settings.

- 13 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

Off is the default option and disables latency protection.

Fail enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.

Override enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value	Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.
Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.	
Protocol	UDP/IP is the default protocol for replication.
Packet Size	Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.
Bandwidth	By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.
Enable Compression	Enable this checkbox if you want to enable Compression for the secondary host.

Click **OK** to close the dialog box and then click **Next**.

14 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically	If virtual IPs have been created, select the Synchronize Automatically option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately. If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online. When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization. Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.
---------------------------	---

Synchronize from
Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu.
- Click **Next** to display the Summary page.

15 Review the information.

Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

If additional Exchange mailbox databases exist in separate disk groups, repeat the procedure “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 164 for the disk group that contains the databases. Provide unique names for the Replicated Data Set and the Replicated Volume Group.

Configuring a hybrid RVG service group for replication

Create and configure a hybrid Replicated Volume Group (RVG) service group for replication. The RVG service group is hybrid because it behaves as a failover service group within a zone and as a parallel service group between zones. For additional information about service group types, see the *Veritas Cluster Server Administrator's Guide*.

Configure the RVG service group's resources manually by copying and modifying components of the Exchange Server service group. Then create new RVG resources and bring them online.

Table 10-2 shows the resources in the hybrid RVG service group for replication.

Table 10-2 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg for the disk group	Volume Manager disk group with Exchange database files
VvrRvg for the disk group	Replicated volume group with Exchange database files


Creating the RVG service group

Create a hybrid replicated volume (RVG) service group, to contain the resources for replication.

To create a hybrid RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.

3 In the Add Service Group window, specify the following:



- Enter a name for the service group. Make sure the service group name is in uppercase.
For example, enter EXCH_RVG_SG.
- Select the systems in the primary zone (Zone 0) and click the right arrow to add them to the service group.
- Select **Hybrid**.
- Click **OK**.

Configuring the RVG service group for RDC replication

Configure the RVG service group’s resources manually for RVG by completing the following tasks:

- **“Configuring the IP and NIC resources”**
Create an IP resource, and copy the NIC resource of the Exchange Server service group (EXCH_SG1), paste and modify them for the RVG service group (EXCH_RVG_SG).
- **“Configuring the VMDg resources for the disk groups”**
Copy the VMDg resources for all disk groups in the Exchange Server service group (EXCH_SG1), paste and modify them for the RVG service group (EXCH_RVG_SG).
- **“Adding the VVR RVG resources for the disk groups”**
Create the VVR RVG resources for all the disk groups and enter the attributes for each of the disk groups and the replication IP address.
- **“Linking the VVR RVG resources to establish dependencies”**
Link the VVR RVG resources to establish the dependencies between the VMDg resources, the IP resource for replication, and the VVR RVG resources for the disk groups. Configure the RVG service group’s VMDg resources to point to the disk groups that contain the RVGs.
- **“Deleting the VMDg resource from the Exchange Server service group”**
Delete the VMDg resources from the Exchange Server service group, because they depend on the replication and were configured in the RVG service group.

Configuring the IP and NIC resources

Configure the following resources and attributes for the IP and NIC:

Table 10-3 IP and NIC resources

Resource	Attributes to Modify
IP	Address
NIC	(none)

To create the IP resource

- 1 In the VCS Cluster Explorer window, right-click the RVG service group (EXCH_RVG_SG) and then click **Add Resource**.

- 2 Specify a name for the resource and choose the Resource Type as IP. Then, specify the IP address, subnet mask and mac address attributes. The IP address is the address for the replication service group.
- 3 Click **OK**.

To create the NIC resource

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EXCH_SG1) in the left pane.
- 2 On the Resources tab, right-click the IP resource (EXCH_SG1-IP), and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the RVG service group (EXCH_RVG_SG).
- 4 On the Resources tab, right-click in the blank resource display area and click **Paste**.
- 5 In the Name Clashes window, change the names of the IP and NIC resources for the RVG service group.
- 6 Click **OK**.

To modify the IP resource and NIC

- 1 In the Resources tab display area, right-click the IP resource (EXCH_RVG_SG-IP) and select **View > Properties View**.
- 2 In the Properties View window, for the **Address** attribute, click **Edit**.
- 3 In the Edit Attribute window, enter the VVR IP address for the Primary Zone as the scalar value.
- 4 Close the Properties View window.

To enable the IP resource and NIC

- 1 In the Resources tab display area, right-click the IP resource (EXCH_RVG_SG-IP) and select **Enabled**.
- 2 In the Resources tab display area, right-click the NIC resource (EXCH_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg resources for the disk groups

You create the VMDg resource in the RVG service group by copying it from the Exchange Server service group and renaming it. You then clear the DGGuid attribute for the new VMDg.

You repeat these procedures for any additional VMDg resources that you want to create for replication.

You modify the attributes of the MountV resources in the Exchange Server service group for the new VMDg in the RVG service group.

Note: The MountV resources correspond to the volumes that you are configuring for replication. The table shows an example configuration. You may have additional volumes you want to include for replication.

Table 10-4 MountV resources

Resource	Attributes to Modify
Resources for disk groups for the Exchange files:	
MountV (for the Exchange Server mailbox database volume)	VMDg Resource Name Volume Name
MountV (for the log volume)	VMDg Resource Name Volume Name

To create the VMDg resource for a disk group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EXCH_SG1) in the left pane.
- 2 On the Resources tab, right-click the VMDg resource for the disk group, with the Exchange files (EXCH_SG1-VMDg), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (EXCH_RVG_SG).
- 4 On the Resources tab, right-click in the blank resource display area and click **Paste**.
- 5 In the Name Clashes window, change the name of the VMDg resource for the RVG service group, for example to EXCH_RVG_SG-VMDg.
- 6 Click **OK**.

To clear the DGGuid attribute for the new VMDg

- 1 In the Resources tab display area, right-click the new VMDg resource.
- 2 In the same Properties View window, for the **DGGuid** attribute, click **Edit**.
- 3 In the Edit Attribute window, clear the scalar value for the **DGGuid** attribute.
- 4 Close the Properties View window.

To modify the MountV resources in the Exchange Server service group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EXCH_SG1) in the left pane.
- 2 In the Resources tab display area, right-click the MountV resource for the Exchange Server files (EXCH_SG1-MountV) and select **View > Properties View**.
- 3 In the Properties View window, verify that the Volume Name attribute is the Exchange Server mailbox database files (SG1_DB1).
- 4 In the same Properties View window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the Edit Attribute window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example EXCH_RVG_SG-VMDg.
- 6 Close the Properties View window.
- 7 Repeat these steps to modify the VMDGResName value for the additional MountV resources for the Exchange database log volume.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (EXCH_RVG_SG).
- 2 In the Resources tab display area, right-click the VMDg resource (EXCH_RVG_SG-VMDg) and select **Enabled**.

Adding the VVR RVG resources for the disk groups

Add VVR RVG resources for replication of the disk groups. If the application has multiple disk groups, create a separate VvrRvg resource for each disk group. Configure the following attributes in the RVG service group for the VvrRvg resource:

Table 10-5 VvrRvg resources

Resource	Attributes to Modify
Resources for the disk group for the Exchange files:	
VvrRvg	VMDgResName IPResName

To create the VVR RVG resource for a disk group

- 1 In the left pane, select the RVG service group (EXCH_RVG_SG) and then right-click on it and select **Add Resource**.

- 2
- In the Add Resource window, do the following:
- Enter the **Resource Name** for the VVR RVG resource, for example, EXCH_VvrRvg.
- Select the **Resource Type** as VvrRvg.
- 3
- In the Add Resource window the attributes appear. For the **RVG** attribute, click **Edit**.
- 4
- In the Edit Attribute window, enter the name of the RVG group that is being managed, for example EXCH_RVG_SG.
- 5
- Click **OK**.
- 6
- In the Add Resource window, for the **VMDGResName** attribute, click **Edit**.
- 7
- In the Edit Attribute window, enter the name of the disk group containing the RVG, for example EXCH_RVG_SG-VMDg.
- 8
- Click **OK**.
- 9
- In the Add Resource window, for the **IPResName** attribute, click **Edit**.
- 10
- In the Edit Attribute window, enter the name of the IP resource managing the IP address for replication, for example EXCH_RVG_SG-IP.
- 11
- Click **OK**.
- 12
- In the Add Resource window, verify that the attributes have been modified and then click **OK**.

Linking the VVR RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the RVG service group to establish the dependencies between the resources. Start from the top parent and link the following resources:

Table 10-6 Dependencies for VVR RVG resources for RDC

Parent	Child
Resources for the disk group for the Exchange files:	
EXCH_VvrRvg	The IP for replication, for example EXCH_RVG_SG-IP
EXCH_VvrRvg	The VMDg for the Exchange files, for example EXCH_RVG_SG-VMDg

To link the VVR RVG resources

- 1
- In the left pane, select the RVG service group (EXCH_RVG_SG).

- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example SG1_DB1_VvrRvg.
- 4 Click the child resource, for example EXCH_RVG_SG-IP.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG resources.
Notice that when you enable a resource and the state of the entity which it is monitoring is online, the corresponding VCS agent reports status for that resource as online. You do not have to bring the resource online manually.

Deleting the VMDg resource from the Exchange Server service group

The VMDg resources must now be manually deleted from the Exchange Server service group because they depend on replication and were configured in the RVG service group.

To delete the VMDg Resources from the Exchange Server service group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EXCH_SG1) from the left pane.
- 2 In the Resources tab display area, right-click the VMDg resource for the first disk group (EXCH_SG1-VMDg) and select **Delete**.
- 3 Click **Yes** to confirm that you want to delete it (even if it is online).
- 4 In the Resources tab display area, right-click the VMDg resource for any additional disk group, if configured, and select **Delete**.
- 5 Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

Add resources of type RVGPrimary to the Exchange Server service group for each of the Exchange Server disk groups and configure the attributes.
Set the value of the **RvgResourceName** attribute to the name of the RVG resource for the RVGPrimary agent.
Configure the following attributes in the Exchange service group for the RVG Primary resources:

Table 10-7 RVG Primary resources

Resource	Attributes to Modify
Resources for the disk group for the Exchange files:	

Table 10-7 RVG Primary resources

Resource	Attributes to Modify
RVGPrimary	RvgResourceName

Creating the RVG Primary resources

For each disk group created for this Exchange mailbox database, create a separate RVG Primary Resource for replication.

To create the RVG Primary resource for an Exchange Server disk group

- 1 In the VCS Cluster Explorer window, right-click the Exchange Server service group (EXCH_SG1) in the left pane, and select **Add Resource**.
- 2 In the Add Resource window, do the following:
 - Enter the **Resource Name** for the RVG Primary resource for the Exchange Server disk group, for example EXCH_RvgPrimary.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the Add Resource window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the Edit Attribute window, enter the name of the VVR RVG resource that corresponds to the disk group, for example EXCH_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Veritas Cluster Server Administrator's Guide* for more information about the RVG Primary agent.
- 6 Verify that **Critical** and **Enabled** are both checked.
- 7 Click **OK**.

Linking the RVG Primary resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the Exchange Server service group (EXCH_SG1) to establish the dependencies between the resources for replication.

Link each MountV resource to the appropriate RVGPrimary resource.

Table 10-8 Dependencies for the RVG Primary resources for RDC

Parent	Child
EXCH_SG1-MountV	EXCH_RvgPrimary
EXCH_SG1-MountV-1	EXCH_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the Exchange Server service group (EXCH_SG1).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example EXCH_SG1-MountV.
- 4 Click the child resource, for example EXCH_RvgPrimary.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the Exchange Server service group (EXCH_SG1) online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the Exchange Server service group (EXCH_SG1).
- 2 In the right pane, on the Resources tab, right-click the first RVG Primary resource (EXCH_RvgPrimary) and select **Online > SYSTEM1**.
- 3 In the right pane, on the Resources tab, right-click any additional RVG Primary resource and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG

In the RVG service group, set up systems in the primary zone (Zone 0) to specify that initial fail over occurs to systems within the primary zone for the RVG service group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (EXCH_RVG_SG).

- 2 In the right pane, select the Properties tab.
- 3 In the Properties pane, click **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the SystemZones attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (type 0 for Zone 0) for the primary zone.
- 7 Click **OK**.

Setting a dependency between the service groups

The RVG service group must be online on both the primary and secondary zones. However, if a failover occurs from one node to another within the same zone, the RVG service group must fail over along with the database service group.

To ensure that the Exchange Server service group and the RVG service group fail over and switch together, set up an online local hard dependency from the RVG service group to the Exchange Server service group.

The Exchange service group (for example, EXCH_SG1) is dependent on the replication service group (for example, EXCH_RVG_SG).

To set up an online local hard dependency

- 1 From VCS Cluster Explorer, in the left pane, select the cluster (MYCLUSTER).
- 2 In the right pane, select the Service Groups tab.
- 3 Click the **Link** button to create a dependency between service groups.
- 4 Click the Exchange Server service group (the parent service group), for example EXCH_SG1.
- 5 Click the RVG service group (the child resource), for example EXCH_RVG_SG.
- 6 In the Link Service Groups window, do the following:
 - Select the Relationship of **online local**.
 - Select the Dependency Type of **hard**.
 - Click **OK**.

Adding the nodes from the secondary zone to the RDC

Configuration of the systems in the Primary Zone (Zone 0) is complete. The systems in the Secondary Zone (Zone 1) can now be added to the RDC configuration.

Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

To add the nodes from the secondary zone to the RVG

- 1 From the active node of the cluster in the primary zone, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the Welcome page, and click **Next**.
- 3 In the Wizard Options dialog box, do the following:
 - Click **Modify an existing replication service group**.
The existing replication service group is selected, by default.
 - Click **Next**.
- 4 If a VCS notice message appears asking if you want to continue, click **Yes**.
- 5 Specify the system priority list as follows:
 - In the Available Cluster Systems box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group's system list.
To remove a node from the service group's system list, click the node in the Systems in Priority Order box, and click the left arrow icon.
 - To change the priority of a node in the system list, click the node in the Systems in Priority Order box, then click the up and down arrow icons.
The node at the top of the list has the highest failover priority.
 - Click **Next**.
- 6 If a message appears indicating that the configuration will be changed from Read Only to Read/Write, click **Yes**.
- 7 Review the Disk Group and Replicated Volume Group Configuration and click **Next**.

- 8 In the IP Resource Options dialog box, select **Modify IP resource** and click **Next**.
- 9 If a VCS error appears, click **OK**.
- 10 In the Network Configuration dialog box, verify that the selected adapters are correct and click **Next**.
- 11 Review the summary of the service group configuration.
The Resources box lists the configured resources. Click a resource to view its attributes and their configured values in the Attributes box.
- 12 Click **Next** to modify the replication service group. When prompted, click **Yes** to modify the service group.
- 13 Click **Finish**.

Configuring secondary zone nodes in the RVG service group

Specify Zone 1 as the zone for the nodes in RVG service group at the secondary zone.

To specify the secondary zone for the nodes in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (EXCH_RVG_SG).
- 2 In the right pane, select the Properties tab.
- 3 In the Properties pane, click **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the SystemZones attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the Edit Attribute dialog box, click the plus (+) sign and enter the systems and the zone number (type 1 for Zone 1) for the secondary zone.
- 8 Click **OK** and close the Attributes View window.

Configuring the IP resources for fail over

Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.

In the event of a system or Exchange failure, VCS attempts to fail over the Exchange service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the

primary zone, VCS switches the service group to a node in the current secondary system zone.

To modify the IP resources in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (EXCH_RVG_SG).
- 2 In the right pane, select the Resources tab.
- 3 Right-click the RVG IP resource (EVS1_RVG_SG-IP) and select **View > Properties View**.
- 4 In the Edit Attributes window, edit the Address attribute as follows:
 - Select **Per System**.
 - Select the first node (SYSTEM1) in the primary zone and enter the virtual IP address for the primary zone.
 - Select the second node in the primary zone and enter the virtual IP address for the primary zone (the same IP address as the first node).
 - Repeat for all nodes in the primary zone.
 - Select the first node (SYSTEM3) in the secondary zone and enter the virtual IP address for the secondary zone.
 - Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - Repeat for all nodes in the secondary zone.
 - Click **OK**.
- 5 In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone have the same IP address. The IP address at the primary zone and the secondary zone should be different.
- 6 Close the Properties View window.

As this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.

Adding the nodes from the secondary zone to the Exchange Server service group

Use the Exchange 2010 Database Configuration Wizard to add the nodes from the secondary zone (Zone 1) to the Exchange database service group.

To add the nodes from the secondary zone to the Exchange database service group

- 1 Start the Exchange 2010 Database Configuration Wizard.
Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server 2010 Configuration Wizard**.
- 2 Review the prerequisites on the Welcome panel and then click **Next**.
- 3 On the Service Group Options panel, click **Modify database service group**, select the Exchange service group, and then click **Next**.
- 4 On the Exchange Database Selection panel, click **Next**.
- 5 On the Exchange Service Group Configuration panel, click **Next**.
- 6 On the System Selection panel, select the required systems and then click **Next**.
 - The Available Cluster Systems box displays the systems from the secondary zone (Zone 1). Select a system and then click the right arrow icon to move the system 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 icon.
- 7 On the Network Configuration panel, click **Next**.
- 8 On the Service Group Summary panel, review the service group configuration and click **Next**.
- 9 Click **Yes** on the dialog box that prompts you that the wizard will modify the configuration.
- 10 In the Completing the Exchange Configuration panel, clear the **Bring the service group online** check box and 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.

Configuring the zones in the Exchange Server service group

Specify Zone 1 as the zone for the Exchange **database** service group nodes in the secondary zone.

To specify the secondary zone for the nodes in the Exchange database service group

- 1 From VCS Cluster Explorer, in the left pane, select the Exchange database service group (EXCH_SG1).
- 2 In the right pane, select the Properties tab.
- 3 In the Properties pane, click **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the SystemZones attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the Edit Attribute dialog box, click the plus (+) sign and enter the systems and the zone number (type 1 for Zone 1) for the secondary zone.
- 8 Click **OK** and close the Attributes View window.

Verifying the RDC configuration

After completing all the configuration tasks for the primary and secondary zones, you can bring the service group online, then verify the configuration.

Perform the following tasks:

- [Bringing the service group online](#)
- [Switching online nodes](#)

Bringing the service group online

After completing all configuration tasks, ensure that the RVG service group is online in both the primary and secondary zone. Then you can bring the Exchange **database** service group online in the primary zone.

To bring the Exchange service group online

- 1 From VCS Cluster Explorer, in the left pane, right-click the Exchange **database** service group (EXCH_SG1).
- 2 Click **Online**.

Switching online nodes

Failover simulation is an important part of configuration testing. Test the failover by switching online nodes.

Note: This should never be tested on systems with live data. A reliable and tested backup should be available. A tested backup means that it has been tested successfully by a restore.

The RVG service group is online in both the primary and secondary zone. However, within a zone, if more than one node is configured, the RVG service group should fail over with the database service group.

Switch the database service group between nodes using Veritas Cluster Manager (Java Console). When you complete the procedure, you will see the online system role shift from one system to the other.

If you enter the system name manually from the Java Console, specify the name in upper case.

To switch online nodes

- 1 Open the Veritas Cluster Manager (Java Console).
Click **Start > All Programs > Veritas > Veritas Cluster Manager (Java Console)**.
- 2 Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, in the File menu, click **New Cluster**. In the New Cluster - Connectivity Configuration window, enter the computer name in the **Host name** field and click **OK**.
- 3 In the Machinename - Login window, enter your user name and password in the respective fields and click **OK**.
- 4 Right-click the service group in the left pane, and select an alternate system name from the **Switch To** entry.
- 5 In the Question dialog box, click **Yes** to confirm you do want to switch the service group to the other node.

Additional instructions for GCO disaster recovery

After completing the tasks for setting up a replicated data cluster for Exchange, you can optionally create a secondary site for wide area disaster recovery using the SFW HA Global Cluster option (GCO).

With this option, if a disaster affects a local or metropolitan area, data is failed over to a site hundreds or thousands of miles away.

To configure disaster recovery using a secondary site, you must install the SFW HA Global Cluster Option on all nodes on the primary (replicated data cluster) site cluster, as well as the secondary (DR) site cluster. GCO configuration also requires a static IP address available for each site.

You can use the Disaster Recovery (DR) wizard when setting up the secondary site. The secondary site is not configured as a replicated data cluster. There can be only one replicated data cluster in the DR environment. The DR wizard does the following tasks:

- Clones the storage
- Clones the application service group
- Sets up VVR replication for the secondary site
- Configures the primary and secondary site clusters as global clusters

See “[Disaster recovery configuration](#)” on page 61.

Deploying disaster recovery for Exchange Server

This chapter contains the following topics:

- [“Tasks for deploying a disaster recovery configuration of Microsoft Exchange”](#) on page 197
- [“Reviewing the disaster recovery configuration”](#) on page 200
- [“Setting up the secondary site: Installing SFW HA and configuring a cluster”](#) on page 202
- [“Verifying your primary site configuration”](#) on page 203
- [“Setting up your replication environment”](#) on page 203
- [“Assigning user privileges \(secure clusters only\)”](#) on page 211
- [“Configuring disaster recovery with the DR wizard”](#) on page 212
- [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 216
- [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 220
- [“Installing Exchange 2010”](#) on page 224
- [“Cloning the service group configuration on to the secondary site using the DR wizard”](#) on page 224
- [“Configuring replication and global clustering”](#) on page 227
- [“Verifying the disaster recovery configuration”](#) on page 243

- [“Establishing secure communication within the global cluster \(optional\)”](#) on page 245
- [“Adding multiple DR sites \(optional\)”](#) on page 247
- [“Recovery procedures for service group dependencies”](#) on page 247
- [“Possible task after creating the DR environment: Adding a new failover node to a VVR environment”](#) on page 251

Tasks for deploying a disaster recovery configuration of Microsoft Exchange

Before setting up disaster recovery at the secondary site, you must complete the high availability configuration on the primary site.

See “[High availability \(HA\) configuration \(New Server\)](#)” on page 51.

You can also configure disaster recovery for a primary site that is configured as a replicated data cluster.

See “[VCS Replicated Data Cluster configuration](#)” on page 57.

After setting up the SFW HA environment for Exchange on a primary site, you can create a secondary or “failover” site for disaster recovery using the Disaster Recovery (DR) wizard.

The DR wizard helps you to clone the storage and service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

After service group configuration, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See “[About the Solutions Configuration Center](#)” on page 39.

This document covers how to use the DR wizard to configure disaster recovery. However, if necessary, you can configure disaster recovery manually. To do so, you first configure the high availability and Exchange components on both sites. You then configure the disaster recovery components: Veritas Volume Replicator (VVR) and the Global Cluster Option. You also have the choice of using array-based hardware replication for your disaster recovery solution. See the *Veritas Cluster Server Administrator's Guide* for information on configuring the Global Cluster Option. See the *Veritas Volume Replicator Administrator's Guide* for information on configuring VVR. For information on configuring array-based hardware replication, see the VCS hardware agent documentation for the particular array you want to configure.

Table 11-1 outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the secondary site using the Disaster Recovery wizard.

Table 11-1 Configuring the secondary site for disaster recovery

Action	Description
Install SFW HA and configure the cluster on the secondary site	Install SFW HA and configure the cluster at the secondary site. Caution: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster. See “Setting up the secondary site: Installing SFW HA and configuring a cluster” on page 202.
Verify that Exchange Server has been configured for high availability at the primary site	Verify that Exchange has been configured for high availability at the primary site and that the service groups are online See “Verifying your primary site configuration” on page 203.
Set up the replication prerequisites	Ensure that replication prerequisites for your selected method of replication are met before running the DR wizard See “Setting up security for VVR” on page 204. See “Configuring EMC SRDF replication and global clustering” on page 235. See “Configuring Hitachi TrueCopy replication and global clustering” on page 238.
(Secure cluster only) Assign user privileges	For a secure cluster only, assign user privileges See “Assigning user privileges (secure clusters only)” on page 211.
Start running the DR wizard	■ Review prerequisites for the DR wizard ■ Start the DR wizard and make the initial selections required for each task: selecting a primary site system, the service group, the secondary site system, and the replication method See “Configuring disaster recovery with the DR wizard” on page 212.

Table 11-1 Configuring the secondary site for disaster recovery (Continued)

Action	Description
Clone the storage configuration (VVR replication only)	(VVR replication option) Clone the storage configuration on the secondary site using the DR wizard See “Cloning the storage on the secondary site using the DR wizard (VVR replication option)” on page 216.
Create temporary storage for application installation (other replication methods)	(EMC SRDF, Hitachi TrueCopy, or GCO only replication option) Use the DR wizard to create temporary storage for installation on the secondary site See “Creating temporary storage on the secondary site using the DR wizard (array-based replication)” on page 220.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “Installing Exchange 2010” on page 224
Clone the service group configuration	Clone the service group configuration from the primary to the secondary site using the DR wizard See “Cloning the service group configuration on to the secondary site using the DR wizard” on page 224.
Configure replication and global clustering, or configure global clustering only	■ (VVR replication) Use the wizard to configure replication and global clustering ■ (EMC SRDF replication) Set up replication and then use the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Set up replication and then use the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Use the wizard to configure global clustering, and then set up replication See “Configuring replication and global clustering” on page 227,
Verify the disaster recover configuration	Verify that the secondary site has been fully configured for disaster recovery See “Verifying the disaster recovery configuration” on page 243.


Table 11-1 Configuring the secondary site for disaster recovery (Continued)

Action	Description
(Optional) Add secure communication	Add secure communication between local clusters within the global cluster (optional task) See “ Establishing secure communication within the global cluster (optional) ” on page 245.
(Optional) Add additional DR sites	Optionally, add additional DR sites to a VVR environment See “ Adding multiple DR sites (optional) ” on page 247.
Handling service group dependencies after failover	If your environment includes dependent service groups, review the considerations for bringing the service groups online after failover to the secondary site See “ Recovery procedures for service group dependencies ” on page 247.

Reviewing the disaster recovery configuration


In a disaster recovery environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. [Figure 11-1](#) displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 11-1 Disaster Recovery environment



When a failure occurs at the primary site, the DR solution is activated. The data that was replicated to the secondary site is used to restore the application services to clients. [Figure 11-2](#) illustrates this type of failure:

Figure 11-2 Application services restored after primary site failure



You can choose to configure replication using VVR or an agent-supported array-based hardware replication. You can use the DR wizard to configure VVR replication or required options for the VCS agents for EMC SRDF or Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

Supported disaster recovery configurations for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

Service group dependency configurations are described in detail in the VCS documentation.

See *Veritas Cluster Server Administrator's Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

The Disaster Recovery wizard supports only one level of dependency (one child). In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported. If you need to configure more

levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

The wizard clones dependent service groups as global groups.

Setting up the secondary site: Installing SFW HA and configuring a cluster

After completing the HA configuration on the primary site, repeat the appropriate tasks to complete the SFW HA installation and configure the cluster at the secondary site.

The storage configuration will be handled by the DR wizard.

Use the following guidelines for installing SFW HA and configuring the cluster on the secondary site.

- Ensure that you have set up the components required to run a cluster.
See “[Configuring the storage hardware and network](#)” on page 92.
- Ensure that when installing SFW HA you install the appropriate disaster recovery options at both the primary and secondary sites. Be sure to select the following installation options as appropriate for your environment:

Veritas Cluster Server Application Agent for Exchange	Required to configure high availability for Exchange
Client	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration
Global Cluster Option	Required for a disaster recovery configuration
Veritas Volume Replicator	Required if you plan to use VVR for replication
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent

For more information see the *SFW HA Installation and Upgrade Guide*.

- Configure the cluster with the VCS Cluster Configuration Wizard (VCW). Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.
See “[Configuring the cluster](#)” on page 114.

Note: You do not need to configure the GCO option while configuring the cluster. This is done later using the Disaster Recovery wizard.

Verifying your primary site configuration

Make sure that Exchange has been configured for high availability at the primary site. If you have not yet configured Exchange for high availability at the primary site, go to High Availability (HA) Configuration in the Solutions Configuration Center and follow the steps in the order shown.

To verify the configuration, use the Cluster Manager (Java console) on the primary site and check the status of the service group in the tree view. Verify that all the resources are online.

Note: If you are setting up a replicated data cluster at the primary site, use the replicated data cluster instructions rather than the high availability configuration steps in the Solutions Configuration Center. See [“VCS Replicated Data Cluster configuration”](#) on page 57.

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Veritas Volume Replicator (VVR)
- EMC SRDF
- Hitachi TrueCopy

For array-based hardware replication, you can use any replication agent supported by Veritas Cluster Server. The DR wizard can help with configuring the methods listed above. If you choose a different replication method, you must run the wizard first to complete configuring global clustering; then afterwards, you configure replication separately.

See [“Configuring global clustering only”](#) on page 241.

Before configuring replication with the wizard, ensure that you set up the replication environment prerequisites. Choose from the following topics, depending on which replication method you are using:

- [“Setting up security for VVR”](#) on page 204

- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 206
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 208

Setting up security for VVR

If you are using Veritas Volume Replicator (VVR) replication, you must configure the VVR Security Service (VxSAS) on all cluster nodes on both the primary and secondary sites.

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

The procedure has these prerequisites:

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

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

To configure the VxSAS service

- 1 Launch the VVR Security Service Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard**.
or
Type `vxscfg.exe` at the command prompt.
- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows and then click **Next**:

Account name	Enter the administrative account name.
(domain\account)	

Password Specify a password.


If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong and then click **Next**:

Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.
 Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

- 5 On the Host Selection panel, select the required hosts:



Selecting hosts	The Available hosts pane lists the hosts that are present in the specified domain. Move the appropriate host from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.
Adding a host	If the host name you require is not displayed, click Add host. In the Add Host dialog specify the required host name or IP in the Host Name field. Click Add to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 6 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.
Click **Back** to change any information you had provided earlier.
- 7 Click **Finish** to exit the wizard.

Requirements for EMC SRDF array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for EMC SRDF. The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also configures the Symm heartbeat. Optional resource settings are left in the default state.

For more information about the EMC SRDF agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for configuring EMC SRDF”](#) on page 206
- [“Replication requirements for EMC SRDF”](#) on page 207

Software requirements for configuring EMC SRDF

The EMC SRDF agent supports SYMCLI versions that EMC recommends for the firmware on the array. The agent supports SRDF on all microcode levels on all

Symmetrix arrays, provided that the host/HBA/array combination is in EMC's hardware compatibility list.

To use the DR wizard to configure the required agent settings for EMC SRDF, ensure that the following software requirements are met:

- The EMC Solutions Enabler is installed on all cluster nodes.
- The SYMCLI version that is installed supports the generation of XML output.
- The SYMCLI version and the microcode level support dynamic swapping.
- The VCS EMC SRDF agent is installed on all cluster nodes.

Replication requirements for EMC SRDF

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that no devices are RDF2.
- On the secondary site, the wizard verifies that no devices are RDF1.

Otherwise, the wizard displays an invalid configuration message and is unable to proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All disks in SFW disk groups must belong to the same device group.
- The device group must not span more than one array (no composite device groups).
- A device group can contain one or more disk groups.
- Dynamic swap must be enabled on both sites.
- On the primary site:
 - All devices must be RDF1 and part of an RDF1 device group.
 - Devices must have write access.
- On the secondary site:
 - All devices must be RDF2 and part of an RDF2 device group.
 - Write access must be disabled.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the SRDF resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the SRDF resource, not to the array configuration. However, the SRDF resource will be unable to come online in the service group until replication has been configured correctly.

In addition, note the following agent requirement:

- Device group configuration must be the same on all nodes of the cluster.

Requirements for Hitachi TrueCopy array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for Hitachi TrueCopy. The wizard configures the required settings for the HTC resource in the VCS application service group. Optional settings are left in the default state.

For more information about the Hitachi TrueCopy agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for Hitachi TrueCopy”](#) on page 208
- [“Replication requirements for Hitachi TrueCopy”](#) on page 209

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

For details, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

To use the DR wizard to configure the required agent settings for Hitachi TrueCopy, ensure that the following requirements are met:

- RAID Manager is installed in the same location on all nodes on a site.

- Enter the primary and secondary site file paths for the horcm files on the Hitachi TrueCopy Path Information panel in the wizard. The default location is:
System Driver\Windows
- The horcm files are named horcmnn.conf (where nn is a positive number without a leading zero, for example, horcm1.conf but not horcm01.conf).

Replication requirements for Hitachi TrueCopy

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that all devices are the same type, but not S-SWS or SSUS.
- On the secondary site, the wizard verifies that all devices are the same type, but not P-VOL or PSUS.

Otherwise, the wizard displays an invalid configuration message and does not proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All configured instances are running.
- No disks in the SFW disk group span across the Device Group.
- A device group can contain one or more disk groups.
- The device group does not span more than one array.
- At the primary site, all devices are of the type P-VOL.
- At the secondary site, all devices are of the type S-VOL.
- All device groups at the primary site are paired to an IP address which must be online on the secondary node.
- Device group and device names include only alphanumeric characters or the underscore character.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the HTC resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the HTC resource, not to the array configuration. However, the HTC resource will be unable to come online in the service group until replication has been configured correctly.

Assigning user privileges (secure clusters only)

In order to enable remote cluster operations you must configure a VCS user with the same name and privileges in each cluster.

When assigning privileges in secure clusters, you must specify fully-qualified user names, in the format `username@domain`. You cannot assign or change passwords for users when VCS is running in secure mode.

You must assign service group rights to the Exchange service group as well as any dependent service groups except for the RVG service group.

See the *Veritas Cluster Server Administrator's Guide*.

To assign user privileges at the primary site

- 1 Set the configuration to read/write mode:
haconf -makerw
- 2 Add the user. Specify the name in the format `username@domain`.
hauser -add user [-priv <Administrator|Operator>]
- 3 Modify the attribute of the service group to add the user. Specify the Exchange service group and any dependent service groups except for the RVG service group.
hauser -add user [-priv <Administrator|Operator> [-group service_groups]]
- 4 Reset the configuration to read-only:
haconf -dump -makero

To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:
haconf -makerw
- 2 Add the user. Specify the name in the format `username@domain`.
hauser -add user [-priv <Administrator|Operator>]
- 3 Reset the configuration to read-only:
haconf -dump -makero

Configuring disaster recovery with the DR wizard

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (VVR replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure VVR replication, or configure the VCS hardware replication agent settings for EMC SRDF or Hitachi TrueCopy
- Configure global clustering

Warning: To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment that is not configured by the wizard, you must first run the wizard to configure global clustering before configuring replication.

The wizard allows you to exit after the logical completion of each task. Each time you re-start the wizard, you specify the primary site system, service group, secondary site system, and replication method, as described in the following procedure. Clicking **Next** then takes you to the start page of the process following the one that you had last completed.

The DR Wizard list of service groups shows only those that contain a MountV resource. For a dependent service group to be listed, the parent service group must also contain a MountV resource.

Warning: Once you have completed configuring replication and global clustering with the DR wizard, you cannot use the wizard to change the method of replication.

Before running the DR wizard to configure disaster recovery, ensure that you meet the following prerequisites:

- SFW HA is installed and a cluster is configured at the secondary site. Ensure that the name assigned to the secondary site cluster is different than the name assigned to the primary site cluster.
- Your application or server role is configured for HA at the primary site and all required services are running at the primary site.
- The clusters taking part in the DR configuration should have distinct names.

- Enough free disk space is available at the secondary site to duplicate the storage configuration at the primary site.
- For Exchange 2003 or 2007, one static IP address is available per application service group to be cloned.
- If using VVR for replication, a minimum of one static IP address per site is available for each application instance running in the cluster.
- Global Cluster Option (GCO) is installed at the primary and secondary site, and one static IP address is available at each site for configuring GCO.
- A VCS user is configured with the same name and privileges in each cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall is set to allow both ingoing and outgoing TCP requests on port 7419.

Note: The DR wizard does not support VVR configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- [“Setting up security for VVR”](#) on page 204
- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 206
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 208

To start configuring disaster recovery with the DR wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

Note: By design, for Exchange 2003 or 2007, the DR wizard requires specific settings for the Lanman attributes on the primary and secondary sites. Before beginning the DR configuration, the wizard checks for these values, and if they are not set as required, the wizard will automatically proceed with setting these values, both at the primary and secondary sites.

- 2 In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.
- 3 In the System Selection panel, complete the requested information:

System Name	Enter the IP address or Fully Qualified Host Name (FQHN) of the primary system where Exchange is online. If you have launched the wizard on the system where Exchange is online at the primary site, you can also specify <code>localhost</code> to connect to the system.
-------------	---

Click **Next**.

- 4 In the Service Group Selection panel, select the service group that you want to clone to the secondary site.
You can choose to clone only the parent service group by not selecting the dependent service group. Only online and local dependencies are supported, in soft, firm, or hard configurations. The wizard can configure only one level of dependency. In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported.
The panel lists only service groups that contain a MountV resource.
Click **Next**.
- 5 In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.
Click **Next**.

- 6 In the Replication Options panel, select the replication method. Although you must select the replication method now, configuring replication and the global cluster option is done later, after service group cloning.

Configure Veritas Volume Replicator (VVR) and the Global Cluster Option (GCO)	Select this option if you want to configure VVR replication. Select this option even if you plan to configure VVR replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a VVR environment. The wizard verifies each configuration task and recognizes if a task has been completed successfully. You cannot mix replication methods. That is, if your primary site is using array-based replication, and you select the VVR option, the wizard will warn you that you cannot use VVR replication for the disaster recovery site.
Configure EMC SRDF and the Global Cluster Option (GCO)	Select this replication option if you want to configure the settings for the VCS EMC SRDF agent. All disks used for the service group on the primary site must belong to an EMC SRDF array. Select this option even if you plan to configure EMC SRDF replication or the GCO option manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.
Configure Hitachi TrueCopy and the Global Cluster Option (GCO)	Select this replication option if you want to configure the settings for the VCS Hitachi TrueCopy agent. All disks used for the service group on the primary site must belong to a Hitachi TrueCopy array. Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.

Configure the Global Cluster Option (GCO) only	If you select this option, the DR wizard does not configure any replication settings. It configures the global cluster option. Select this option if you want to use the wizard in an array-based replication environment that is not supported by this wizard. You must configure replication manually after you finish the wizard. If you select the GCO only option, the DR wizard sets up the storage and service group configuration on the secondary site for an array-based hardware replication environment. Therefore, you cannot use this option to clone the storage and service group for a VVR replication environment.
--	---

Click **Next**.

- 7 Continue with the next DR configuration task.
For VVR replication, see [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 216.
For array-based replication, see [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 220.

Cloning the storage on the secondary site using the DR wizard (VVR replication option)

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at the secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, the wizard can correct this and then complete the configuration.

If you have not yet started the wizard, see the following topic before continuing with the storage cloning procedure:

- [“Configuring disaster recovery with the DR wizard”](#) on page 212.

To clone the storage configuration from the primary site to the secondary site (VVR replication method)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system. In the Replication Options panel, select the VVR replication method and click **Next**.

- Review the information in the Storage Validation Results panel. This panel compares the configuration at the secondary site with that on the primary. If the storage is already configured identically on both sites, the panel shows that results are identical. Otherwise, the panel shows the differences and recommended actions. You can toggle between a summary and detailed view of information about the differences.
The detailed view shows the following:

Disk Group	Displays the disk group name that needs to be created on the secondary site.
Volume	Displays the list of volumes, if necessary, that need to be created at the secondary site.
Size	Displays the size of the volume that needs to be created on the secondary site.
Mount	Displays the mount to be assigned the volume on the secondary site.
Recommended Action	Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary. ■ If the volume does not exist, a new volume will be created. ■ If the volume exists but is of a smaller size than that on the primary, the volume will be expanded to the required size. ■ If the volume is of a greater size than that on the primary, the volume will be recreated using the appropriate size. ■ If the volume is the same as that on the primary, the message indicates that the volumes are identical and no action is required.

The summary view shows the following:

Disk groups that do not exist	Displays the names of any disk groups that exist on the primary but do not exist on the secondary.
Existing disk groups that need modification	Displays the names of any disk groups on the secondary that need to be modified to match the primary.
Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.

If the panel displays a message indicating that the available disks are inadequate to clone the primary site configuration on the secondary, you

can free some disks on the secondary or add more storage. Then click **Refresh/Validate** to have the wizard update its information about the secondary storage configuration.

You continue with the wizard to provide information for the recommended actions. Before proceeding to the service group configuration, the wizard ensures that the configuration of the disk groups and volumes for the service group is the same at the primary and secondary site.

Click **Next**.

- 3
- In the Disk Selection for Storage Cloning panel, for each of the disk groups that does not exist or is not same as the corresponding disk group at the primary site, select disks that the wizard can use to create the respective disk groups at the secondary site.

Selecting Disks

For each of the disk groups that needs to be created, select the required disks from the Available Disks pane. Either double-click on the host name or the >> option to move the hosts into the Selected disks pane.

Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures.

Click **Next**.

- 4
- In the Volume Layout for Secondary Site Storage panel, complete the requested information:

Disk Group

Displays the disk group name to which the volume belongs.

Volume (Volume Size)

Displays the name and the size of the volume, corresponding to that on the primary, that needs to be created on the secondary.

Available Disks

Select the disks on which you want the wizard to create the volumes. From the Available Disks pane, either double-click on the disk name or the >> option to move the disks into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group.

Select disks for each unavailable volume that you want to clone on to the secondary.

Layout

By default, the same layout as the one specified for the primary volume is selected. Click **Edit** to change the layout to suit your specific requirements.

Selected Disks

Displays the list of disks that have been moved in from the Available Disks pane.

Cloning the storage on the secondary site using the DR wizard (VVR replication option)

View Primary Layout Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 5 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.
- 6 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 7 In the Storage Cloning Configuration Result screen, view the results and click **Next**.
- 8 In the Exchange Installation panel, review the information. If Exchange is already installed on the required secondary site nodes, click **Next** to continue with service group cloning. Otherwise, proceed with installation as follows:
 - For Exchange 2003 or Exchange 2007, before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.
 - For Exchange 2003 or Exchange 2007, the system gets restarted when the Exchange installation is complete. Therefore, if you are running the DR wizard from a system where you need to install Exchange 2003 or 2007, click **Finish** to exit the wizard before proceeding with installation. If the DR Wizard is run from a remote node, you can keep the wizard running on that node while you install the application locally on each of the required nodes.
 - If you keep the wizard running during installation, once application installation is complete, click **Next** to proceed with service group cloning. Otherwise, restart the DR wizard and continue through the wizard from the Welcome panel.

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

To enable you to install applications that require volume mount points, the DR wizard can create a temporary disk group, DR_APP_INSTALL__DG, which contains the volumes and mount points for use in application installation. The temporary configuration uses 500 MB volumes or the volume size at the primary site, depending on which is smaller. The wizard deletes the temporary configuration after application installation.

Note: Although Exchange 2010 does not require volume mount points for installation, you still follow this procedure to proceed through the wizard. On the storage cloning panel, you can uncheck the option for storage cloning, as described in the procedure.

If you are starting the wizard for the first time, see the following topic before continuing with the storage cloning procedure:

- [“Configuring disaster recovery with the DR wizard”](#) on page 212.

To create temporary storage for application installation (array-based replication)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system.
- 2 In the Replication Options panel, select the array-based replication method you plan to use and click **Next**:
 - EMC SRDF
 - Hitachi TrueCopy
 - Global Cluster Option only (select if you are using another agent-supported array-based replication method)
- 3 If you selected Hitachi TrueCopy replication, the Hitachi TrueCopy File Paths panel is displayed. The wizard populates the fields if it locates the files in the default location. Otherwise, fill in the file path information for both the primary and secondary sites as follows:

RAID Manager bin path	Path to the RAID Manager Command Line interface
	Default: C:\HORCM\etc
	where C is the system drive.

HORCM files location	Path to the horcm configuration files (horcmnn.conf) Default: C:\Windows where C is the system drive An horcm configuration file is required by the RAID Manager on all nodes; however the wizard does not validate this.
----------------------	--

- 4
- In the Storage Cloning panel, you can choose whether or not to perform storage cloning, which creates a temporary storage disk group and volumes for application installation. The wizard will delete the temporary storage once you confirm application installation is complete. Exchange 2003 and 2007 require the temporary storage for installation. Exchange 2010 does not require the temporary storage since it does not require volume mount points for installation. Choose one of the following:

 - For Exchange 2003 or 2007, if you have not yet installed the application on all nodes, leave **Perform storage cloning** checked and click **Next**. Continue with the next step in this procedure.
 - For Exchange 2010, since temporary storage is not required for installation, you can uncheck **Perform storage cloning** and click **Next**. Proceed with the application installation. Once installation is complete, you can continue with the procedure for service group cloning.
 - If you have already installed the application on all nodes, uncheck **Perform storage cloning** and click **Next**. Continue with the procedure for service group cloning.
- 5
- The Storage Validation Results panel shows the temporary storage configuration that the wizard will configure at the secondary site. You can click **Show Summary** to toggle to a summary view and toggle back to a detailed view by clicking **Show Details**. The detailed view shows the following:

Disk Group	Displays the name of the single disk group required on the secondary site for temporary storage: DR_APP_INSTALL_DG
Volume	Displays the list of volumes required at the secondary site.
Size	Displays the size of the volumes required on the secondary site.
Mount	Displays the mounts required at the secondary site.

Recommended Action	Indicates the action that the wizard will take at the secondary site.
--------------------	---

The summary view shows the following:.

Existing configuration	Displays the existing secondary configuration.
Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.

If the panel displays a message indicating that the available disks on the secondary are inadequate, you can free some disks on the secondary or add more storage. Then click **Refresh/Validate** so that the wizard can update its information about the secondary storage configuration. Click **Next**.

- 6
- In the Disk Selection for Storage Cloning panel, a default disk selection is shown for the temporary storage at the secondary site. You can change the selection by moving disks to and from the Available Disks and Selected Disks pane. Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures. Click **Next**.
- 7
- The Volume Layout for Secondary Site Storage panel shows a default volume layout for the temporary storage based on the primary site volume layout. Optionally, you can change the default disk assignment and layout for any volume:

Disk Group	Displays the DR_APP_INSTALL__DG disk group.
Volume (Volume Size)	Displays the name and the size of the volume to be created on the secondary.
Available Disks	Displays the disks that are available for the volumes. To select a disk, either double-click on the host name or click the >> button to move the hosts into the Selected Disks pane.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been selected for the volume. To remove a disk from the list, select it and click the << button.

View Primary Layout Displays the volume layout at the primary site.

Click **Next**.

- 8 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the temporary storage configuration at the secondary site.
- 9 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 10 In the Storage Configuration Cloning Result screen, view the results and click **Next**.
- 11 In the Exchange Installation panel, review the information and do one of the following:
 - For Exchange 2003 or Exchange 2007, before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.
 - For Exchange 2003 or Exchange 2007, the system gets restarted when the Exchange installation is complete. Therefore, if you are running the DR wizard from a system where you need to install Exchange 2003 or 2007, click **Finish** to exit the wizard before proceeding with installation. If the DR Wizard is run from a remote node, you can keep the wizard running on that node while you install the application locally on each of the required nodes.
 - If you keep the wizard running during installation, once application installation is complete, click **Next** to proceed with service group cloning. Otherwise, restart the DR wizard and continue through the wizard from the Welcome panel.

Once the application is installed, the temporary cloned storage is no longer needed. Before beginning service group cloning, the wizard displays the Temporary Storage Deletion panel to confirm the deletion of the temporary storage.

Installing Exchange 2010

Installing Exchange on the secondary site uses the same procedure as when installing Exchange on the primary site.

See [Chapter 7, “Installing Exchange Server”](#) on page 129.

Cloning the service group configuration on to the secondary site using the DR wizard

Before cloning a service group on the secondary site, verify the following:

- On the secondary site, the application is installed.
- For Exchange 2003 or Exchange 2007, if the secondary site is in a different subnet, ensure that a static IP address is available on the secondary site to assign to the virtual server.

If you are launching the wizard for the first time, see the following topic for additional information:

- [“Configuring disaster recovery with the DR wizard”](#) on page 212.

Note: Although you can view the cloning progress in the VCS Java Console, do not save and close the configuration while cloning is in progress. Otherwise, the cloning fails and you have to delete the service group on the secondary site and run the wizard again.

To clone the service group configuration from the primary site to the secondary site

- 1 At the primary site, verify that you have brought the application service group online.
- 2 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 In the Welcome panel, click **Next** and continue through the wizard, providing the requested information for the primary site system, the service group, the secondary site system, and the replication method.

If you selected the VVR replication method, the wizard proceeds to the storage cloning task and notifies you if it detects that the storage is identical. Click **Next** until you reach the Service Group Analysis panel. If you selected an array-based replication method (EMC SRDF, HTC, or GCO only), the temporary storage is no longer needed once the application is installed and the wizard confirms whether or not to delete it.

- 4 (Array-based replication method only) In the Temporary Storage Deletion panel, confirm whether or not to delete the cloned storage:
- If the application is already installed on the required nodes, leave **Delete cloned storage** checked and click **Next**. When the wizard prompts you to confirm deleting the shared storage, click **Yes**.

■ If you want to delete the cloned storage manually later, uncheck **Delete cloned storage** and click **Next**.
- 5 (Array-based replication method only) If you selected to delete the cloned storage, the wizard shows the progress of the tasks in the Implementation panel. If the storage deletion fails, the wizard will show a failure summary page. Otherwise, when it shows the tasks are complete, click **Next**.
- 6 Review the following information displayed in the Service Group Analysis panel and click **Next** to continue with service group cloning.

Service Group Name	Displays the list of application-related service groups present on the cluster at the primary site.
Service Group Details on the Primary Cluster	Displays the resource attributes for the service group at the primary site. The NIC resource consists of the MAC address. For Exchange 2003 or 2007 service groups: The IP resource consists of the IP address and subnet mask.
Service Group Details on the Secondary Cluster	Displays a message to indicate whether the service group or the corresponding attributes have been configured at the secondary site.

- 7 In the Service Group Cloning panel, specify the requested system information for the secondary site.

Service Group Name	Depending on the application service group already created at the primary site, and subsequently selected on the Service Group Selection page, the wizard displays the names of the service groups that will be cloned at the secondary site.
--------------------	---

Available Systems	Displays a list of available systems on the secondary cluster that are not yet selected for service group cloning. Select any additional secondary systems on which you want the wizard to clone the application service group configuration. Either double-click on the system name or use the > option to move the hosts into the Selected Systems pane. Note: If you want to add systems to a service group after you finish cloning the service group configuration with the DR wizard, you cannot do so by running the DR wizard again. Instead, run the VCS configuration wizard and edit the system list of the existing service group.
Selected Systems	Displays the list of selected systems. The secondary system that you selected earlier in the wizard is listed by default.

Click **Next**.

- 8
- In the Service Group Attribute Selection panel, complete the requested information to create the required resources on the secondary site. The panel also displays the service group resource name and the attribute information at the primary site.

Resource Name	Displays the list of resources that exist on the primary cluster.
Attribute Name	Displays the attribute name associated with each of the resources displayed in the Resource Name column. If you need to edit additional attributes that are not shown, you must edit them manually on the secondary site service group once service group cloning is complete.
Primary Cluster	Displays the primary attribute values for each of the displayed attributes.
Secondary Cluster	The default is the same as the primary cluster. The following applies to Exchange 2003 or 2007 only: The same virtual IP address (for the Exchange virtual server) can be used if both sites exist on the same network segment. You can specify different attributes depending on your environment. For the MACAddress attribute select the appropriate public NIC from the drop-down list.

Click **Next**.

- 9 In the Service Group Summary, review the attribute information that will be cloned on to the secondary cluster. Click **Back** to change any of the secondary service group attributes. Otherwise, click **Next** to proceed with cloning the service group configuration on the secondary site.
- 10 In the Implementation panel, wait until all the tasks are completed. The progress bar indicates the status of the tasks. Successful tasks are marked with a check symbol. If some task could not be completed successfully, the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**
- 11 If the cloning failed, review the troubleshooting information. Otherwise, click **Next** to continue with the replication and GCO configuration, or with GCO only, depending on which option you selected.
Optionally, you can exit the wizard at this point and launch the wizard again later. When you launch the wizard again, continue through the wizard, specifying the primary site system, the service group, the secondary site system, and the replication method. Click **Next** to continue to the replication and/or GCO configuration task.

Configuring replication and global clustering

After creating the identical service group configuration on both sites, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication.

If you are using an array-based replication that is not supported by the wizard, you configure global clustering only. In this case, you must complete configuring global clustering before configuring replication.

The following topics cover the steps required for each replication method:

- [“Configuring VVR replication and global clustering”](#) on page 227
- [“Configuring EMC SRDF replication and global clustering”](#) on page 235
- [“Configuring Hitachi TrueCopy replication and global clustering”](#) on page 238
- [“Configuring global clustering only”](#) on page 241

Configuring VVR replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure VVR replication and global clustering.

Note: By default, in an Exchange or SQL Server environment, the DR wizard configures all the volumes in a disk group under one Replicated Volume Group (RVG). If you require a different organization, you should configure it using the Veritas Enterprise Administrator (VEA) rather than the DR wizard. For information on setting up VVR replication with the VEA, see *Veritas Volume Replicator, Administrator's Guide*.

You can then return to the wizard to configure global clustering.

Before you begin, ensure that you have met the following prerequisites:


- Ensure that Veritas Volume Replicator is installed at the primary and secondary site.
- Ensure that Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- Ensure that VVR Security Service (VxSAS) is configured at the primary and secondary site. See the following topic:
 - [“Setting up security for VVR”](#) on page 204
- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that, if using secure clusters, you configure a VCS user with the same name and privileges in each cluster.

Use the following procedure to configure VVR replication and global clustering with the DR wizard.

To configure VVR replication and GCO

- 1 Verify that the application server service group is online at the primary site and the appropriate disk groups are imported at the secondary site.
- 2 If the wizard is still open after the previous wizard task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
 - On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.

- On the Replication Methods panel, click **Configure VVR and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**. If not, click **Cancel** and restart the wizard after meeting the requirements.
- 4 In the Replication Settings for Replicated Volume Group panel, specify the requested information. If you are adding a DR site to an existing DR configuration, fields that must match the existing settings, such as the RVG or RDS name, are dimmed so that you cannot change them.




Disk Group	The left column lists the disk groups. By design, an RVG is created for each disk group.
RVG Name	Displays the default RVG name. If required, change this to a name of your choice.

RDS Name	Displays the default Replicated Data Set (RDS) name. If required, change this to a name of your choice.
Available Volumes	Displays the list of available volumes that have not been selected to be a part of the RVG. Either double-click on the volume name or use the > option to move the volumes into the Selected RVG Volumes pane.
Selected RVG Volumes	Displays the list of volumes that have been selected to be a part of the RVG. To remove a selected volume, either double-click the volume name or use the < option to move the volumes into the Available Volumes pane.
Primary SRL	If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the name, size, and disk. Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.
Secondary SRL	If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the same name and size as you specified for the primary SRL. Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.
Start Replication after the wizard completes	Select this check box to start replication automatically after the wizard completes the necessary configurations. Once replication is configured and running, deselecting the checkbox does not stop replication.

- Click **Advanced Settings** to specify some additional replication properties. The options on the dialog box are described column-wise, from left to right; refer to the *Veritas Volume Replicator*

Administrator's Guide for additional information on VVR replication options.



Replication Mode Select the required mode of replication; **Synchronous**, **Asynchronous**, or **Synchronous Override**. The default is synchronous override.

Log Protection Select the appropriate log protection from the list.

The **AutoDCM** is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **Off** option disables Replicator Log Overflow protection.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

If the Secondary becomes inactive due to disconnection or administrative action then Replicator log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between primary and secondary RVG is broken, then, any new writes to the primary RVG are failed.

Primary RLINK Name	Enter a name of your choice for the primary RLINK. If you do not specify any name then the wizard assigns a default name.
Secondary RLINK Name	Enter a name of your choice for the Secondary RLINK. If you do not specify any name then the wizard assigns a default name.
Bandwidth	By default, VVR replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit. The default unit is Mega bits per second (Mbps) and the minimum allowed value is 1 Mbps.
Protocol	Choose TCP or UDP. UDP/IP is the default replication protocol.
Packet Size (Bytes)	Default is 1400 Bytes. From the drop-down list, choose the required packet size for data transfer. The default unit for the packet size is Bytes. You can set the packet size only if the protocol is UDP/IP.
Latency Protection	By default, latency protection is set to Off. When this option is selected the High Mark Value and the Low Mark Value are disabled. Select the Fail or Override option to enable Latency protection. This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail option when the Secondary is connected.
High Mark Value	This option is enabled only when Latency Protection is set to Override or Fail. It specifies the maximum number of pending updates by which the secondary site can be behind the primary site. The default value is 10000. To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.
Low Mark Value	This option is enabled only when Latency Protection is set to Override or Fail . When the updates in the Replicator log reach the High Mark Value , then the writes to the system at the primary site continues to be stalled until the number of pending updates on the Replicator log falls back to the Low Mark Value . The default is 9950.

Initial Synchronization	If you are doing an initial setup, then use the Auto Synchronous option to synchronize the secondary site and start replication. This is the default. When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization. If you want to use the Synchronize from Checkpoint method then you must first create a checkpoint. If you have a considerable amount of data on the primary data volumes then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from the checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.
-------------------------	---

To apply changes to advanced settings, click **OK**. On the Replication Settings for Replicated Volume Group panel click **Next**.

- In the Replication Attribute Settings panel, specify required replication attribute information for the cluster at the primary and secondary site. Click the arrow icon to expand an RVG row and display the replication attribute fields. If you are configuring an additional secondary site (multiple DR sites), some fields are disabled.

Disk Group	Displays the list of disk groups that have been configured.
RVG Name	Displays the Replicated Volume Groups corresponding to the disk groups.
IP Address	Enter replication IPs that will be used for replication, one for the primary site and another for the secondary site.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC from the drop-down list for the system at the primary and secondary site.
Copy	Enables you to copy the above network settings to any additional RVGs that are listed on this screen. If there is only one RVG, this option does not apply.

After specifying the replication attributes for each of the RVGs, click **Next**.

- 6
- In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7
- In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters. If you have a printer installed, you can click the printer icon at the bottom of the scrollable list to print the settings.
Click **Next** to implement the settings.
- 8
- In the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for global clustering. If a task could not be completed successfully, it is marked with an (x) symbol. For any critical errors, the wizard displays an error message. For less critical errors, the Information column displays a brief description

about the task failure and the next screen displays additional information on what action you can take to remedy it. Click **Next**.

- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Configuring EMC SRDF replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the SRDF resource in the application service group.

Ensure that you have met the prerequisites for replication. See the following topic:

- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 206

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings. See the following topic:

- [“Optional settings for EMC SRDF”](#) on page 237

To configure EMC SRDF replication and GCO

- 1 Verify that you have brought the application service group online at the primary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.

- In the Replication Methods panel, click **Configure EMC SRDF and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the SRDF resource cannot come online in the service group.

- 4 In the SRDF Resource Configuration panel, the wizard populates the required resource fields if replication has been configured. Otherwise, you must enter the required resource settings manually.

Symmetrix Array ID (SID)	Specify the array ID for the primary site and for the secondary site.
Device Group name	Specify the name of the Symmetrix device group that contains the disks of the disk group for the selected instance.
Available VMDG Resources	Select the disk groups associated with the selected application instance.

- 5 If you want to configure an additional SRDF resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.

IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

Click **Next**.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.
- 10 Proceed with configuring additional optional settings for the SRDF resource if desired, and then verifying the disaster recovery configuration.

Optional settings for EMC SRDF

The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also detects and configures the SymHome attribute.

Other settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for EMC*

SRDF, Configuration Guide. If you change any settings, ensure that you edit the resource on both the primary and secondary sites.

The wizard also detects and configures the SymHome attribute.

The optional settings use the following defaults:

Option	Default setting
DevFOTime	2 seconds per device required for a device to fail over
AutoTakeover	The default is 1; the agent performs a read-write enable on partitioned devices in the write-disabled state during a failover, if devices are consistent.
SplitTakeover	The default is 1; the agent brings service groups online on the R2 side even if the devices are in the split state because they are read-write enabled.

Configuring Hitachi TrueCopy replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the HTC resource in the application service group.

Ensure that you have met the prerequisites. See the following topic:

- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 208

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings. It uses defaults for optional settings. See the following topic:

- [“Optional settings for HTC”](#)

To configure Hitachi TrueCopy replication and GCO

- 1 Verify that you have brought the application server service group online at the primary site
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server**

> **Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the HTC resource cannot come online in the service group.

- 4 In the HTC Resource Configuration panel, the wizard populates the required resource fields if the horcm file is configured properly. If not, you can configure the horcm file and click **Refresh** to populate the fields. Alternatively, enter the required resource settings manually:

Instance ID	Specify the instance number of the device group. Multiple device groups may have the same instance number.
Device Group name	Specify the name of the Hitachi device group that contains the disk group for the selected instance. The device group name must be the same on both the primary and secondary sites.
Available VMDG Resources	Select the disk groups associated with the selected application instance.
Add, Remove, Reset buttons	Click Add or Remove to display empty fields so that you can manually add or remove additional resources. Click Refresh to repopulate all fields from the current horcm file.

- 5 If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information

can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7
- In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8
- In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✕) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9
- In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

- 10 Proceed with configuring additional optional settings for the HTC resource if desired, and then verifying the disaster recovery configuration.

Optional settings for HTC

The wizard configures the required settings for the HTC resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

The optional settings use the following defaults:

Option	Default setting
LinkMonitor	The default is 0; the agent does not periodically attempt to resynchronize the S-VOL side if the replication link is disconnected. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the pairresync command.
SplitTakeover	The default is 0; the agent does not permit a failover to S-VOL devices if the replication link is disconnected; that is, if P-VOL devices are in the PSUE state.

Configuring global clustering only

If you are using a replication method that the DR wizard does not configure, you must select the replication option to configure global clustering only.

For the GCO only option, you use the wizard to complete all DR tasks except the replication configuration task. You must complete the final wizard task of configuring global clustering before configuring replication.

Before configuring GCO:

- Ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- If you created secure clusters at the primary site and secondary site, ensure that you have configured a VCS user with the same name and privileges in each cluster, and the user must be added in the Administrator role.

The following procedure assumes that you have completed the earlier wizard tasks through the service group cloning task and are continuing with the final step of configuring global clustering.

To configure GCO only

- 1
- If the wizard is still open after the service group cloning task, continue with the GCO Setup panel. Otherwise, launch the wizard and proceed to the GCO Setup panel as follows:
- Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
- In the Replication Methods panel, click **Configure Global Cluster Option (GCO) only**. Click **Next** and continue to the GCO Setup panel.
- 2
- In the GCO Setup panel, review the requirements. If you have met the requirements, click **Next**.
- 3
- In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.

Start GCO after configuration

Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes.

Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 4 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified. Click **Next**.
- 5 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (x) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 6 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Verifying the disaster recovery configuration

The steps you need to take to verify your DR configuration depend on the type of replication you are using.

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- For VVR replication, confirm that the configuration of disk groups and volumes at the DR site have been created by the DR wizard storage cloning.
- Confirm that the application VCS service group has been created in the DR cluster including the same service group name, same resources, and same dependency structure as the primary site's application VCS service group.
- Confirm that the application service group is online at the primary site. The application service group should remain offline at the DR site.
- For VVR replication:
 - Ensure VVR replication configuration. This includes ensuring that the RVGs have been created at primary and secondary with the correct

volume inclusion, replication mode, Replicator Log configuration, and any specified advanced options.

- Confirm that the replication state matches what was specified during configuration. If specified to start immediately, ensure that it is started. If specified to start later, ensure that it is stopped.
- Ensure that the VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.
- Confirm that the RVG service groups are online at the primary and secondary sites.
- Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- For array-based replication, verify that the required array resource is created in the primary and secondary cluster's application service group and that a dependency is set between the VMDg resource and the array resource.
- For EMC SRDF replication, verify that the SRDF resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, verify that the HTC resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, you must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the disk groups online. This must be performed only once, after which the failover works uninterrupted. For more information, see *Veritas™ Cluster Server Hardware Replication Agent for Hitachi TrueCopy Installation and Configuration Guide*.
- Ensure that the application service groups are configured as global.
- Check to ensure that the two clusters are communicating and that the status of communication between the two clusters has a state of Alive.
- If you are using VVR for replication and configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you are using VVR for replication and chose to start replication manually in the DR wizard, to avoid replicating large amounts of data over the network the first time, then you will need to start the process necessary to synchronize from checkpoint. This typically consists of
 - starting a VVR replication checkpoint

- performing a block level backup
- ending the VVR replication checkpoint
- restoring the block level backup at the DR site
- starting replication from the VVR replication checkpoint

To learn more about the process of starting replication from a checkpoint, refer to the *Veritas Volume Replicator Administrator's Guide*.

- Do not attempt a wide area failover until data has been replicated and the state is consistent and up to date. The Solutions Configuration Center provides a Fire Drill Wizard to test wide area failover for VVR-based replication.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.

- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

To add the -secure option to the StartProgram resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add -secure switch to the path of the executable Scalar Value.
For example:
"C:\Program Files\Veritas\Cluster Server\bin\wac.exe" -secure
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.
Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:
vssat setuptrust --broker <host:port> --securitylevel <low|medium|high> [--hashfile <filename> | --hash <root hash in hex>]
For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:
from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low  
from RB2, type:  
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the ClusterService-Proc (wac) resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

Adding multiple DR sites (optional)

In a Veritas Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. Veritas Cluster Server supports up to four DR sites. In other replication environments, additional DR sites require manual configuration.

Run the DR wizard and on the Secondary System selection panel, select the new site.

Before you start the wizard on the task of configuring replication and global clustering, ensure that the cluster service group is online at the existing primary and secondary sites. This enables the wizard to configure GCO not only between the selected primary site and the new secondary site but also between the new site and the earlier configured secondary site. Otherwise, the wizard displays a warning message after the global clustering task.

When configuring the VVR replication settings with the wizard for the additional site, fields that must match existing settings are dimmed so that you cannot change them. For example, you cannot change the RVG name or RVG layout on the Replication Settings panel. Similarly, on the Global Cluster Settings panel, GCO has already been configured at the primary site, so the primary site fields are dimmed.

Recovery procedures for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

See [“Supported disaster recovery configurations for service group dependencies”](#) on page 201.

The procedure and requirements for bringing service group dependencies online at the secondary site depends on their configuration: soft, firm, or hard. In general, if a child or parent remains online at the primary site, you take it offline before you bring the child and parent service groups online in the correct order on the secondary site.

An exception is the RVG service group, used for VVR replication, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, soft dependency link.

Table 11-2 Online, local, soft dependency link

Failure condition	Results	Action required	
The child service group fails	■ The parent remains online on the primary site.	1	Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the child service group only.	2	Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.		
The parent service group fails	■ The child remains online on the primary site.	1	Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2	Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.		

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, firm dependency link.

Table 11-3 Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent). Leave the RVG group online at the primary site.
The parent service group fails	■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, hard dependency link.

Table 11-4 Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent). Do not take the RVG group offline at the primary site.

Table 11-4 Online, local, hard dependency link (Continued)

Failure condition	Results	Action required
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.	

Possible task after creating the DR environment: Adding a new failover node to a VVR environment

The following procedures describe how to add an additional node to the cluster at either the primary or secondary site after your disaster recovery environment is in operation. The clusters at each site are not required to have the same number of nodes or the same failover configuration.

See the following topics:

- [“Preparing the new node”](#) on page 251
- [“Preparing the existing DR environment”](#) on page 251
- [“Installing Exchange on the new node”](#) on page 252
- [“Modifying the replication and Exchange service groups”](#) on page 252
- [“Reversing replication direction”](#) on page 253

Preparing the new node

Install SFW HA on the new system and then add the system to the cluster.

To install SFW HA and add the system to the cluster

- 1 Refer to [“Installing Veritas Storage Foundation HA for Windows”](#) on page 94 for installation instructions.
- 2 Use the **Cluster Operations** option of the VCS Configuration wizard (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**) to add the new system to the cluster. If necessary, refer to the *Veritas Cluster Server Administrator’s Guide* for information on this procedure.

Preparing the existing DR environment

If you plan to add a failover node to the secondary site, you must temporarily switch the roles of the primary and secondary sites so that the current site becomes the primary. This action reverses the direction of replication.

To prepare the existing DR environment

- 1 If you are adding the failover node to the cluster at the primary site, proceed directly to [step 2](#). If you are adding a failover node to the secondary site, you must switch the roles of the primary and secondary sites. This action reverses the direction of replication.

- In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
 - Click **Switch To**, and click **Remote switch**.
 - In the **Switch global group** dialog box:
 - Click the cluster at the secondary site you want to switch the group to.
 - Click the specific system where you want to bring the global Exchange service group online.
 - Click **OK**.
- 2 Take the global Exchange service group offline at the current primary site.
 - 3 Take the VVR replication service group offline.

Installing Exchange on the new node

Install Exchange on the new node, but do not add the node to the service group SystemList.

To prepare the node and install Exchange

- 1 Import the disk group on the new node. Follow the procedure described in [“Setting up the secondary site: Installing SFW HA and configuring a cluster”](#) on page 202.
- 2 From the VEA navigation tree, right-click the RVG for the primary site, and click **Enable Data Access**.
- 3 Install Exchange. See [Chapter 7, “Installing Exchange Server”](#) on page 129.

Modifying the replication and Exchange service groups

Add the new failover node to the system lists in the Replication and Exchange service groups.

To add the failover node to the system lists

- 1 Bring the replication service group online on an existing cluster node of the current primary site.
- 2 Bring the MountV resources of the corresponding Exchange service group online on the same node.
- 3 Use the **Modify an existing replication service group** option of the Volume Replicator Agent Configuration Wizard (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard**) to add the new node to the system list for the

replication service group. If necessary, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for information on this procedure.

- 4 Use the **Modify service group** option of the Exchange Server Configuration Wizard (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Exchange Server Configuration Wizard**) to add the new node to the system list for the Exchange service group. Check the check box to bring the service group online after the wizard completes. If necessary, refer to the *Veritas Cluster Server Administrator's Guide* for information on this procedure.
- 5 After bringing the Exchange service group online, you must use Exchange to configure all the database stores to automatically mount on start-up.

Reversing replication direction

If you added a failover node at the original secondary site and migrated the RVG in [“Preparing the existing DR environment”](#) on page 251, move the global Exchange service group back to the original primary site and reverse the direction of replication. These actions switch the Primary and Secondary sites back to their original roles.

To reverse the replication direction

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 In the **Switch global group** dialog box:
 - Click the cluster to switch the group to.
 - Click the specific system where you want to bring the global Exchange service group online.
 - Click **OK**.

Testing fault readiness by running a fire drill

Topics in this chapter include:

- [About disaster recovery fire drills](#)
- [About the Fire Drill Wizard](#)
- [About post-fire drill scripts](#)
- [Tasks for configuring and running fire drills](#)
- [Prerequisites for a fire drill](#)
- [Preparing the fire drill configuration](#)
- [Running a fire drill](#)
- [Recreating a fire drill configuration that has changed](#)
- [Restoring the fire drill system to a prepared state](#)
- [Deleting the fire drill configuration](#)

About disaster recovery fire drills

A disaster recovery plan should include regular testing of an environment to ensure that a DR solution is effective and ready should disaster strike. This testing is called a fire drill.

A fire drill is performed at the secondary site using a special service group for fire drills. The fire drill service group uses a copy of the data that is used by the application service group.

About the Fire Drill Wizard

Veritas Storage Foundation HA for Windows (SFW HA) provides a Fire Drill Wizard to help you set up and run a fire drill on a disaster recovery environment. You launch the Fire Drill Wizard from the Solutions Configuration Center.

The Fire Drill Wizard tests the fault readiness of a disaster recovery configuration by mimicking a failover from the primary site to the secondary site. The wizard does this without stopping the application at the primary site and disrupting user access.

The Fire Drill Wizard supports conducting a fire drill for a disaster recovery site that uses Veritas Volume Replicator (VVR) or that uses Hitachi TrueCopy or EMC SRDF hardware replication.

About Fire Drill Wizard general operations

The Fire Drill Wizard performs the following operations:

- Prepares for the fire drill by creating a fire drill service group on the secondary site
The fire drill service group is a copy of the application service group. When creating the fire drill service group, the wizard uses the application service group name, with the suffix `_fd`. The Exchange fire drill service group contains only the VMDg and mountV resources, and in the case of hardware replication, the HTC Snap or SRDF Snap resource. The wizard renames the fire drill service group resources with a prefix `FDnn` and changes attribute values as necessary to refer to the FD resources.
The wizard also supports fire drill service groups created under a different naming convention by an earlier version of the wizard.
- Runs the fire drill by bringing the fire drill service group online on the secondary site
Optionally the wizard runs Eseutil to check for data consistency as part of the fire drill test. The fire drill tests the replication and consistency of the data to verify that the data will be available if the Exchange service group fails over and comes online at the secondary site should the need arise.
Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online on the primary site.

Note: The fire drill service group for Exchange contains only the data resources, not the application, so that the Exchange application does not itself come online during a fire drill.

- Restores the fire drill configuration, taking the fire drill service group offline

After you complete the fire drill, you run the wizard to restore the fire drill configuration to a prepared state. Otherwise, the fire drill service group remains online. If you run a fire drill on one service group, restore that service group before you continue with a fire drill on another service group. You must also restore the fire drill configuration before you can delete it.

Warning: If the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting. Therefore, always use the wizard to restore the fire drill configuration to a prepared state as soon as possible after completing the fire drill testing for a service group.

See [“Restoring the fire drill system to a prepared state”](#) on page 277.

- Deletes the fire drill configuration

The details of some Fire Drill Wizard operations are different depending on the replication environment.

See [“About Fire Drill Wizard operations in a VVR environment”](#) on page 257.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 258.

About Fire Drill Wizard operations in a VVR environment

The general operations of the Fire Drill Wizard are the same in all replication environments.

- Prepares for the fire drill, creating a fire drill service group on the secondary site
- Runs the fire drill, bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration, taking the fire drill service group offline
- Deletes the fire drill configuration

See [“About the Fire Drill Wizard”](#) on page 256.

However, the following additional Fire Drill Wizard operations are specific to a Veritas Volume Replicator (VVR) environment.

In a VVR environment, when preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, replaces the RVGPrimary resources with VMDg resources
- Uses the SFW HA VxSnap feature to prepare snapshot mirrors for use during the fire drill
You assign one or more disks for the mirrored volumes while running the wizard. Mirror preparation can take some time, so you can exit the wizard once this step is started and let the preparation continue in the background.

In a VVR environment, when running the fire drill, the wizard does the following:

- Detaches the mirrors from the original volumes to create point-in-time snapshots of the production data
- Creates a fire drill disk group on the secondary site with a snapshot of the application data to use for testing purposes

In a VVR environment, when restoring the fire drill system to a prepared state, the wizard does the following:

- Takes the fire drill service group offline
- Disables the fire drill service group resources
- Imports the fire drill disk group
- Joins the fire drill disk group to the application service group disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

In a VVR environment, when deleting the fire drill configuration, the wizard does the following:

- Deletes the fire drill service group and any associated registry entry
- Performs the snap abort operation on the snapshot mirrors to free up the disk space

About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment

The general operations of the Fire Drill Wizard are the same in all replication environments.

- Prepares for the fire drill, creating a fire drill service group on the secondary site
- Runs the fire drill, bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration, taking the fire drill service group offline

- Deletes the fire drill configuration

See “[About the Fire Drill Wizard](#)” on page 256.

However, additional Fire Drill Wizard operations are specific to a Hitachi TrueCopy or EMC SRDF replication environment.

In a Hitachi TrueCopy or EMC SRDF replication environment, the wizard performs the following additional actions during preparation, running of the fire drill, restoring the configuration, and deleting the configuration. You must configure the ShadowImage (for Hitachi) or BCV (for SRDF) pairs before running the wizard.

When preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, the wizard creates HTCSnap or SRDFSnap resources for each HTC and SRDF resource in the application service group. The SRDFSnap and HTCSnap resources from the firedrill service group are linked to the respective resources configured in the main application service group.
- The wizard configures the Snap resource. The following Snap resource attributes are set to a value of 1:
 - UseSnapshot (take a local snapshot of the target array)
 - RequireSnapshot (require a successful snapshot for the Snap resource to come online)
 - MountSnapshot (use the snapshot to bring the fire drill service group online)

When running the fire drill, the wizard brings the HTCSnap or SRDFSnap agent online. The HTCSnap or SRDFSnap agent manage the replication and mirroring functionality according to the attribute settings. The Snap agents take a consistent snapshot of the replicating data using the mirroring technology provided by the array vendor. The Snap agents also import the disk group present on the snapshot devices with a different name.

In more detail, the Snap agent does the following:

- Suspends replication to get a consistent snapshot
- For HTCSnap, takes a snapshot of the replicating application data on a ShadowImage device
- For SRDFSnap, takes a snapshot of the replicating application data on a BCV device
- Resumes replication
- Modifies the disk group name in the snapshot

When restoring the fire drill configuration to a prepared state, the wizard does the following:

- Takes the fire drill service group offline, thus also taking offline the SRDF and HTC Snap agents
This action reattaches the hardware mirrors to the replicating secondary devices and resynchronizes them.

When deleting the fire drill configuration, the wizard does the following:

- Deletes the fire drill service group
- Deletes any associated registry entry

If you want to remove the hardware mirrors, you must do so manually.

For more information about the Hitachi TrueCopy Snap agent functions, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

For more information about the EMC SRDF Snap agent functions, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

About post-fire drill scripts

You can specify a script for the Fire Drill Wizard to run on the secondary site at the end of the fire drill.

For the wizard to run the script, the script must exist on the secondary system that you are specifying for the fire drill.

Note: The wizard does not support using script commands to launch a user interface window. In such a case, the process is created but the UI window does not display.

Optionally, you can specify to run a Windows PowerShell cmdlet by creating a .bat file.

To run a cmdlet, create a .bat file with the following entry:

```
%windir%\system32\WindowsPowerShell\v1.0\PowerShell.exe -command  
"$ScriptName"
```

Where

\$ScriptName = .ps1 script (fully qualified) / cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\PowerShell.exe -command  
C:\myTest.ps1
```

Specify the name of the .bat file as the script to run.

For Exchange Server 2007 or 2010, go to one of the following topics for more information on Exchange scripts or cmdlets:

- [“Exchange 2007 scripts or cmdlets”](#) on page 261
- [“Exchange 2010 scripts or cmdlets”](#) on page 261

Exchange 2007 scripts or cmdlets

For Exchange 2007, Symantec recommends using scripts rather than cmdlets as Exchange is not brought online at the secondary site, and the Exchange Management shell should be run under the virtual environment context, which is not available on a node where Exchange is offline.

To specify a cmdlet for Exchange 2007, use the following entry in the .bat file:

```
%windir%\system32\WindowsPowerShell\v1.0\PowerShell.exe -PSConsoleFile  
"$ExchDir"\bin\exshell.psc1 -command "$ScriptName"
```

Where

\$ExchDir equals HKLM -> SOFTWARE -> Microsoft -> Exchange -> Setup::
Services key.

\$ScriptName = .ps1 script (fully qualified) / cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\PowerShell.exe -  
PSConsoleFile "D:\Program Files\Microsoft\Exchange Server\bin\exshell.psc1"  
-command C:\myTest.ps1
```

Exchange 2010 scripts or cmdlets

To use Exchange cmdlets for Exchange 2010, use the following entry in the .bat file:

```
%windir%\System32\WindowsPowerShell\v1.0\powershell.exe -noexit  
-command ". '$ExchangeDir'\bin\RemoteExchange.ps1';  
Connect-ExchangeServer -auto; $ScriptName"
```

Where

\$ExchangeDir is the directory where Exchange 2010 is installed. For example:
C:\Program Files\Microsoft\Exchange Server\V14

\$ScriptName is the .ps1 script (with fully qualified path) / cmdlet specified by the user. For example: C:\ListMailboxes.ps1

The following would be an example of the .bat file entry:

```
C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe -noexit  
-command ". 'C:\Program Files\Microsoft\Exchange  
Server\V14\bin\RemoteExchange.ps1'; Connect-ExchangeServer -auto;  
C:\ListMailboxes.ps1"
```

Tasks for configuring and running fire drills

While running the Fire Drill Wizard, the following sequence of actions are available:

- Prepare the fire drill configuration
- Run the fire drill or delete the configuration
- Restore the fire drill configuration after running a fire drill
- Run another fire drill or delete the configuration

In addition, you have the option to recreate a fire drill configuration that has changed.

After an action is complete, the next action becomes available in the wizard. You can select the next action or exit the wizard and perform the next action later.

[Table 12-1](#) gives more details of the process of configuring and running fire drills with the wizard.

Table 12-1 Process for configuring and running fire drills

Action	Description
Verify the hardware and software prerequisites	Before running the wizard, review the prerequisites and make sure that they are met. See “Prerequisites for a fire drill” on page 264.
Prepare the fire drill configuration	Use the wizard to configure the fire drill. See “Preparing the fire drill configuration” on page 266.
Recreate a fire drill configuration that has changed	If a fire drill configuration exists for the selected service group, the wizard checks for differences between the fire drill service group and the application service group. If differences are found, the wizard can recreate the fire drill configuration before running the fire drill. See “Recreating a fire drill configuration that has changed” on page 274.

Table 12-1 Process for configuring and running fire drills

Action	Description
Run the fire drill	Use the wizard to run the fire drill. Running the fire drill brings the fire drill service group online. Optionally you can specify a script to be run once the fire drill is complete. See “Running a fire drill” on page 272. Confirm the resources are online and replicated data is available Note: After completing the fire drill testing, run the wizard again as soon as possible to restore the configuration. Otherwise the fire drill service group remain online. Be sure to restore one fire drill service group to a prepared state before running a fire drill on another service group.
Restore the fire drill configuration to a prepared state	Use the wizard to restore the fire drill system to a state of readiness for future fire drills or to prepare for removal of the fire drill configuration This is a required action after running the fire drill. See “Restoring the fire drill system to a prepared state” on page 277. This operation reattaches snapshot mirrors and takes the fire drill service group offline.
Delete the fire drill configuration	If a fire drill service group is no longer needed, or if you want to free up resources, use the wizard to remove the fire drill configuration See “Deleting the fire drill configuration” on page 278. The wizard deletes the service group on the secondary site. In a VVR environment, the wizard performs a snap abort to delete the snapshot mirrors created on the secondary site for use in the fire drill. In hardware replication environments, you can delete these manually. If a fire drill has been run, the wizard ensures that you first restore the fire drill configuration to a prepared state before this option becomes available. This ensures that mirrors are reattached and the fire drill service group is offline before the configuration is deleted.

Prerequisites for a fire drill

Before running the Fire Drill Wizard make sure that you meet the following general requirements:

- You can run the Fire Drill Wizard from any node in the domain of the cluster, as long as the SFW HA client is installed on that node.
- If the cluster is secured, the login you use to run the Fire Drill Wizard must have the appropriate permissions to make changes in the cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall must be set to allow both ingoing and outgoing TCP requests on port 7419.
- If you want the fire drill wizard to run a script that you supply, ensure that the script file is available on any secondary site nodes where you plan to run the fire drill.
- If you specify for the fire drill wizard to run Eseutil, the output files are placed by default in the system's TEMP environment variable folder (for example, C:\Windows\Temp). If you want the output files to go to another folder, use the WINSOL_ESEUTIL_OUT_DIR environment variable to define the output file location.

Additional requirements apply to specific replication environments.

See [“Prerequisites for a fire drill in a VVR environment”](#) on page 264.

See [“Prerequisites for a fire drill in a Hitachi TrueCopy environment”](#) on page 265.

See [“Prerequisites for a fire drill in an EMC SRDF environment”](#) on page 266.

Prerequisites for a fire drill in a VVR environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 264.

Make sure that the following additional prerequisites are met before configuring and running a fire drill in a Veritas Volume Replicator (VVR) environment:

- The primary and secondary sites must be fully configured with VVR replication and the global cluster option.
- The Veritas FlashSnap option must be installed on all nodes of the secondary site cluster.

- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- On the secondary site, empty disks must be available with enough disk space to create snapshot mirrors of the volumes. Snapshot mirrors take up the same amount of space as the original volumes. In addition, two disk change object (DCO) volumes are created for each snapshot mirror, one for the source volume and one for the snapshot volume. The two DCO volumes must be on different disks. Allow 2 MB additional space for each DCO volume.

The empty disks must be in the same disk group that contains the RVG. If the disk group does not have empty disks available, you must use the VEA to add the disks to the disk group before you run the wizard. The secondary system must have access to the disks or LUNs.
- All disk groups in the service group must be configured for replication. The Fire Drill wizard does not support a VVR configuration in which disk groups are excluded from replication. However, you can exclude individual volumes within a disk group from replication.

Prerequisites for a fire drill in a Hitachi TrueCopy environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 264.

Make sure that the following prerequisites are met before configuring and running a fire drill in a Hitachi TrueCopy environment:

- The primary and secondary sites must be fully configured with Hitachi TrueCopy replication and the global cluster option. The configuration must follow the applicable instructions in the Veritas Storage Foundation HA for Windows documentation for configuring disaster recovery with Hitachi TrueCopy.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- Make sure that Hitachi RAID Manager/Command Control Interface (CCI) is installed.
- ShadowImage for TrueCopy must be installed and configured for each LUN on the secondary site target array. ShadowImage pairs must be created to allow for mirroring at the secondary site.

- The name of the ShadowImage device group must be the same as the replicated device group for both replicated and non-replicated LUNs that are to be snapshot. The instance number should be different.
- Make sure the HORCM instance managing the S-VOLs runs continuously; the agent does not start this instance.

Prerequisites for a fire drill in an EMC SRDF environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 264.

Make sure that the following prerequisites are met before configuring and running a fire drill in an EMC SRDF environment:

- The primary and secondary sites must be fully configured with EMC SRDF replication and the global cluster option. The configuration must follow the applicable instructions in the Veritas Storage Foundation HA for Windows documentation for configuring disaster recovery with EMC SRDF.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- The infrastructure to take snapshots at the secondary site must be properly configured between the secondary site source and target arrays. This process involves associating Symmetric Business Continuance Volumes (BCVs) and synchronizing them with the secondary site source.
- If you plan to run a fire drill on SRDF/A devices, you must have a TimeFinder/CG license. Make sure TimeFinder for SRDF is installed and configured at the target array.
- To take snapshots of R2 devices, BCVs must be associated with the RDF2 device group and fully established with the devices.
- To take snapshots of non-replicated devices, create an EMC Symmetrix device group with the same name as the SFW disk group. The device group must contain the same devices as in the disk group and have the corresponding BCVs associated.

Preparing the fire drill configuration

Preparing the fire drill configuration creates a fire drill service group. You specify the application service group and the secondary system to use. Only one service group can be prepared for a fire drill at one time.

For a Veritas Volume Replicator (VVR) environment, the preparation step also prepares snapshot mirrors of production data at the specified node on the secondary site.

Note: Preparing the snapshot mirrors takes some time to complete.

Before you prepare the fire drill configuration with the Fire Drill Wizard, make sure that you meet the prerequisites.

See [“Prerequisites for a fire drill”](#) on page 264.

To prepare the fire drill configuration

- 1 Open the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**).
- 2 Start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, review the information and click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
 See [“System Selection panel details”](#) on page 269.
- 5 In the Service Group Selection panel, select the service group that you want to use for the fire drill and click **Next**.
 See [“Service Group Selection panel details”](#) on page 269.
- 6 In the Secondary System Selection panel, select the cluster and the system to be used for the fire drill at the secondary site, and then click **Next**.
 See [“Secondary System Selection panel details”](#) on page 269.
- 7 If the Fire Drill Prerequisites panel is displayed, review the information and ensure that all prerequisites are met. Click **Next**.
 See [“Prerequisites for a fire drill”](#) on page 264.
 Otherwise, if a fire drill service group already exists on this system for the specified service group, one of the following panels is displayed:

If the Run Fire Drill option or Delete Fire Drill options are shown, a fire drill service group has already been prepared.	You can run the fire drill with no further preparation. Click Run Fire Drill and follow the procedure for running a fire drill. See “Running a fire drill” on page 272.
--	---

If the Fire Drill Restoration panel is displayed, the fire drill service group remains online from a previous fire drill.	Follow the procedure for restoring the fire drill configuration to a prepared state. This must be done before running a new fire drill. See “Restoring the fire drill system to a prepared state” on page 277.
If the Recreate Fire Drill Service Group panel is displayed, a fire drill service group has already been prepared but is not up to date.	You can choose to recreate the fire drill configuration to bring it up to date. See “Recreating a fire drill configuration that has changed” on page 274. Or you can clear the check box to recreate the configuration and run the fire drill on the existing configuration.

8 The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

VVR replication	Disk Selection panel See “Disk Selection panel details” on page 269.
Hitachi TrueCopy replication	Horcm Files Path Selection panel See “Hitachi TrueCopy Path Information panel details” on page 270. HTCSnap Resource Configuration panel See “HTCSnap Resource Configuration panel details” on page 271.
EMC SRDF replication	SRDFSnap Resource Configuration panel See “SRDFSnap Resource Configuration panel details” on page 271.

Click **Next**.

- 9 In the Fire Drill Preparation panel, the wizard shows the status of the preparation tasks.
See [“Fire Drill Preparation panel details”](#) on page 272.
When preparation is complete, click **Next**.
- 10 The Summary panel displays the message that preparation is complete.
To run the fire drill now, click **Next**. Continue with the procedure to run the fire drill.
See [“Running a fire drill”](#) on page 272.

To run the fire drill later, click **Finish**. The fire drill preparation remains in place.

System Selection panel details

Use the System Selection panel of the wizard to specify a system in the primary site cluster.

All systems containing online global service groups are available to select. The default system is the node where you launched the wizard (localhost) if a global service group is online on that system. When selecting a system you can specify either a fully qualified host name or IP address.

Service Group Selection panel details

Use the Service Group Selection panel of the wizard to select the service group that you want to use for the fire drill. You can select only one service group at a time for a fire drill.

Secondary System Selection panel details

Use the Secondary System Selection panel of the wizard to select the cluster and the system to be used for the fire drill at the secondary site.

The selected system must have access to the replicated data.

The system must have access to disks for the snapshots that will be created for the fire drill.

Disk Selection panel details

During fire drill preparation in a VVR replication environment, you must ensure that information is available to the wizard for creating the fire drill snapshots.

Use the Disk Selection panel of the wizard to review the information on disks and volumes and make the selections for the fire drill snapshots, as follows:

Volume	Select the volumes for the fire drill snapshots. By default all volumes associated with the service group are selected. If you deselect a volume that might result in the fire drill service group failing to come online, the wizard displays a warning message. Note: The Disk Selection panel also appears if the wizard is recreating a fire drill service group to which volumes have been added. In that case, only the new volumes are shown for selection.
Disk Group	Shows the name of the disk group that contains the original volumes. This field is display only.
Fire Drill DG	Shows the name of the fire drill disk group that running the fire drill will create on the secondary system to contain the snapshots. This field is display only. For the fire drill disk group name, the wizard prefixes the original disk group name with <i>FDnn</i>.
Disk	Click the plus icon to the right of the Disk column and specify the disk to be used for the snapshot volume. Repeat for each row that contains a selected volume. You can store multiple snapshot volumes on the same disk, if the production volumes reside on disks in the same disk group. If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard.
Mount Details	Shows the mount details for the snapshot volumes on the secondary system, which match the mounts for the production volumes. This field is display only.

Hitachi TrueCopy Path Information panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the Hitachi TrueCopy Path Information panel is displayed.

The wizard populates the path field with the customary default location:

C:\Windows

where C is the system drive.

If the horcm configuration files are in a different location, edit the field to specify that location.

HTCSnap Resource Configuration panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the wizard discovers the HTC resources and non-replicating SFW disk groups in the application service group

This information is used to configure the HTCSnap resources.

The wizard lists each HTCSnap resource that will be configured. You can clear the HTCSnap resource name check box if you do not want to include its dependent disk groups in the fire drill.

You must specify the ShadowImage instance.

The HTCSnap Resource Configuration panel shows the following:

Target Resource Name	The panel shows the HTC resource name in the case of a Replication Device Group or the disk group resource name in the case of a non-replicating disk group.
ShadowImage Instance ID	For every HTC resource, specify the ID of the ShadowImage instance associated with the replicating secondary devices.
Refresh	If you click the Refresh button, the wizard rediscovers and validates the HTC configuration.

More information about HTCSnap resource configuration and operation is available.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 258.

SRDFSnap Resource Configuration panel details

During fire drill preparation in an EMC SRDF replication environment, the wizard validates whether at least one BCV device is attached to every device (RDF2) of the SRDF device group. If not, the wizard displays an informational message on this panel. The panel shows as the Target Resource Name the name of the resource that is managing the LUNs that you want to snapshot. For data being replicated from the primary site, the Target Resource Name is the name of the SRDF resource. For data that is not replicated, the Target Resource Name is the name of the disk group resource.

The wizard lists each SRDFSnap resource that will be configured. You can clear the SRDFSnap resource name check box if you do not want to include its dependent disk groups in the fire drill.

More information about SRDFSnap resource configuration and operation is available.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 258.

Fire Drill Preparation panel details

After you enter the information required to prepare a fire drill configuration, the Fire Drill Preparation panel is displayed. You wait while the wizard completes the preparation tasks.

The fire drill service group is created on the secondary site (but remains offline).

In addition, for a VVR replication environment, the snapshot mirrors for the volumes are prepared; this can take some time. You may want to minimize the wizard while the task runs in the background. You can also track the mirror preparation progress in the VEA. When done, the wizard displays a message that the fire drill preparation is complete. If the wizard is completing the preparation steps as part of recreating a fire drill configuration, snapshot mirrors are prepared only for new volumes.

See [“Recreating a fire drill configuration that has changed”](#) on page 274.

Running a fire drill

After you complete the initial fire drill preparation step using the Fire Drill Wizard, you can run the fire drill immediately without exiting the wizard or run the wizard later to run the fire drill.

Running the fire drill does the following:

- Creates the snapshots
- Enables the firedrill resources
- Brings the fire drill service group online
- Optionally runs Eseutil with the /g option
- Optionally, executes a specified command to run a script
See [“About post-fire drill scripts”](#) on page 260.

For details on the operations that occur when running a fire drill, see the following topics:

- [“About Fire Drill Wizard operations in a VVR environment”](#) on page 257
- [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 258

Warning: After running the fire drill, the fire drill service group remains online. After you verify the fire drill results, run the wizard again to restore the system to the prepared state. Otherwise, if the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting. See [“Restoring the fire drill system to a prepared state”](#) on page 277.

To run a fire drill

- 1 If you completed the initial preparation and have not exited the wizard, or if you are returning to this procedure after recreating a fire drill service group, go to [step 8](#). Otherwise, if you need to restart the wizard, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
- 5 In the Service Group Selection panel, select the service group and click **Next**.
- 6 In the Secondary System Selection panel, specify the system previously prepared for the fire drill at the secondary site and click **Next**.
If the fire drill configuration is in a prepared state, the wizard compares the resources of the fire drill service group with the resources of the application service group.
- 7 If the application service group changed since the fire drill configuration was prepared, the wizard displays the Recreate Fire Drill Service Group panel, showing the differences. Choose one of the following:
 - Leave the option checked to recreate the configuration before running the fire drill and click **Next**. You complete additional steps in the wizard before running the fire drill.
For more information, see [“Recreating a fire drill configuration that has changed”](#) on page 274.
 - To run the fire drill on the existing configuration, clear the option to recreate the fire drill service group and click **Next**.
- 8 In the Fire Drill Mode Selection panel, click **Run Fire Drill** and click **Next**.

- 9 In the Post Fire Drill Script panel, you have the option to specify the full path to a script for the wizard to run after the running the fire drill. In addition, you can specify to run the Eseutil consistency check.
See [“Post fire drill operations panel details”](#) on page 274.
- 10 In the Fire Drill Implementation screen, wait until all fire drill tasks are performed and click **Next**. The Summary panel displays the message that the fire drill is complete. You can leave the wizard running while you verify the results or exit the wizard. To exit the wizard, click **Finish**.
- 11 Run your own tests to verify the fire drill results.

Warning: You should always restore the fire drill system to a prepared state immediately after completing fire drill testing on a service group.

- 12 Restore the fire drill configuration to the prepared state.
See [“Restoring the fire drill system to a prepared state”](#) on page 277.

Post fire drill operations panel details

In the Post Fire Drill Script panel, the wizard displays options for the following actions that it can perform after bringing the fire drill service group online:

- Specify the full path to a script for the wizard to run on the secondary system right after running the fire drill. The script must already exist on the secondary system.
For more information, see [“About post-fire drill scripts”](#) on page 260.
- Check the **Run Eseutil** check box if you want the Eseutil consistency check run on the fire drill snapshots once they are created.
The Eseutil output files are placed by default in the system’s TEMP environment variable folder unless you defined another location by using the WINSOL_ESEUTIL_OUT_DIR environment variable.

Recreating a fire drill configuration that has changed

When you run the Fire Drill wizard, a fire drill service group may already exist for the selected application service group. However, the application service group may have changed since the fire drill service group was created. Therefore, the wizard compares the resource names of the two service groups. If differences are found, the wizard lists them on the Recreate Fire Drill Service Group panel.

You have the following choices from the Recreate Fire Drill Service Group panel:

- Leave the option checked to recreate the fire drill service group.
 Proceed with using the wizard to recreate the configuration to match the application service group.
 The wizard deletes the existing fire drill configuration first, before creating the new one.
 For a VVR replication environment, the wizard handles existing volumes as follows: It does not delete the mirrors for volumes that still exist. When it recreates the fire drill configuration, it prepares new mirrors only for new volumes. If volumes have been removed, the wizard displays an additional option to snap abort the obsolete snapshot volumes to free up disk space.
- Clear the option to recreate the fire drill service group. You can then proceed with using the wizard to do either of the following:
 - Run the fire drill, ignoring the differences.
 - Delete the entire fire drill configuration. Then start over with preparing the fire drill configuration.

Note: The wizard does not check for changes in volume attributes, such as the MountPath attribute. For example, if you have a MountV resource with an attribute that points to drive Y and you change that attribute to point to drive X, the wizard does not identify this change and does not give the option to recreate the fire drill service group.

You can choose whether to manually edit the fire drill service group for such changes and then run the fire drill, ignore the differences, or delete the configuration and start over.

The following procedure describes the choice of recreating the fire drill configuration.

To recreate the fire drill configuration if the service group has changed

- 1 In the Recreate Fire Drill Service Group panel, leave the option checked to recreate the configuration before running the fire drill.
 For a VVR replication environment, if volumes have been removed, optionally select to snap abort the volumes.
 Click **Next**.
- 2 In the Fire Drill Mode Selection panel, Delete Fire Drill Configuration is selected. Click **Next**, and click **Yes** to confirm the deletion.
- 3 The Fire Drill Deletion panel shows the progress of the deletion.
 For a VVR replication environment, the wizard leaves the existing fire drill snapshot volumes so that those snapshot mirrors do not have to be

prepared again. If volumes were removed and you selected the option to snap abort, the wizard snap aborts the snapshots of those volumes.

Warning: If you close the wizard after deleting the fire drill configuration without continuing on to the fire drill preparation step, the information of the existing snapshot volumes is lost.

When all tasks are complete, click **Next**.

- 4
- In the Fire Drill Prerequisites panel, review the information and ensure that all prerequisites are met. Click **Next**.
See “[Prerequisites for a fire drill](#)” on page 264.
- 5
- The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

VVR replication	If volumes have been added, the Disk Selection panel is displayed. Specify the information for the added volumes. If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard. See “Disk Selection panel details” on page 269.
Hitachi TrueCopy replication	Horcm Files Path Selection panel See “Hitachi TrueCopy Path Information panel details” on page 270. HTCSnap Resource Configuration panel See “HTCSnap Resource Configuration panel details” on page 271.
EMC SRDF replication	SRDFSnap Resource Configuration panel See “SRDFSnap Resource Configuration panel details” on page 271.

Click **Next**.

- 6
- The Fire Drill Preparation panel is displayed. Wait while the wizard recreates the fire drill service group.

For VVR replication environments, wait while the wizard starts mirror preparation.

Mirror creation can take some time. You may want to minimize the wizard while the task runs in the background. You can also close the wizard and track the mirror preparation progress in the VEA.

- 7 Once preparation is complete, click **Next**. The Summary page is displayed. To continue with running the fire drill, click **Next**. See [“Running a fire drill”](#) on page 272.

Restoring the fire drill system to a prepared state

After running a fire drill and verifying the results, use the Fire Drill Wizard as soon as possible to restore the fire drill system at the secondary site to a prepared state. A prepared state is the initial fire drill configuration created by the wizard, in which the fire drill service group has been prepared but is offline. Restoring the fire drill system to a prepared state is required for any of the following:

- Making the secondary system available for failover of the application service group at the primary site.
- Running another fire drill.
- Deleting the fire drill configuration after a fire drill has been run.

For details on the operations that occur when restoring a fire drill configuration, see the following topics:

- [“About Fire Drill Wizard operations in a VVR environment”](#) on page 257
- [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 258

To restore the fire drill system to a prepared state

- 1 If you completed running a fire drill and have not exited the wizard, go to [step 8](#). Otherwise, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.

- 7 In the Fire Drill Restoration Information panel, review the requirements for restoration and click **Next**.
- 8 In the Fire Drill Restoration screen, wait until the screen shows the restoration tasks are completed and click **Next**.
- 9 In the Summary screen, click **Next** if you want to delete the fire drill configuration. Otherwise click **Finish** to exit the wizard, leaving the fire drill configuration in a prepared state.

Deleting the fire drill configuration

If you no longer need a fire drill configuration you can delete it.

Deleting a fire drill configuration deletes the fire drill service group on the secondary site.

In a VVR replication environment, deleting a fire drill configuration also performs a snap abort of the snapshot mirrors created on the secondary site for use in the fire drill. It frees up the disk space used for the snapshot mirrors for other use.

In a Hitachi TrueCopy or EMC SRDF environment, you could manually remove mirrors after the deletion is complete.

To delete a fire drill configuration

- 1 If you have just used the wizard to prepare or restore a fire drill configuration and have not exited the wizard, go to [step 10](#). Otherwise continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft Exchange**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 If the wizard detects that the fire drill service group is different from the application service group, it displays the Recreate Fire Drill Service Group panel. Clear the option to recreate the fire drill service group and click **Next**.

- 8 If the wizard detects that the fire drill service group is still online, the Fire Drill Restoration panel is displayed. Review the requirements for restoration and click **Next**.
- 9 In the Restore Fire Drill screen, wait until the screen shows the restoration tasks are completed. Then click **Next**.
- 10 In the Fire Drill Mode Selection panel, click **Delete Fire Drill Configuration** and click **Next**, and click **Yes** to confirm the deletion.
- 11 The Fire Drill Deletion panel shows the progress of the deletion. Wait until all tasks are complete and then click **Next**.
If errors occur while deleting the fire drill configuration, the wizard will list any incomplete steps so that you can complete them manually.
- 12 The Summary panel is displayed. Click **Finish**.

Reference

This section contains the following chapter:

- [Appendix A, “VCS Logging”](#)



VCS Logging

VCS generates two error message logs: the engine logs and the agent logs. Log file names are appended by letters. Letter A indicates the first log file, B the second, C the third, and so on.

The agent log is located at %VCS_HOME%\log\agent_A.txt. The format of agent log messages is:

Timestamp (Year/MM/DD) | Mnemonic | Severity | UMI | Agent Type |
Resource Name | Entry Point | Message Text

A typical agent log resembles:

```
2010/04/12 17:48:39 VCS ERROR V-16-20079-7  
Exch2010DB:QRVCS3-Exch2010DB:attrchanged:Failed to parse the  
argument list.
```

Here,

- Timestamp denotes the date and time when the message was logged.
- Mnemonic denotes which Symantec product logs the message. For VCS application agent for Microsoft Exchange, mnemonic is 'VCS'.
- Severity denotes the seriousness of the message. Severity of the VCS error messages is classified into the following types:
 - CRITICAL indicates a critical error within a VCS process. Contact Technical Support immediately.
 - ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action.
 - WARNING indicates a warning or error, but not an actual fault.
 - NOTE informs that VCS has initiated an action.
 - INFO informs about various state messages or comments.Of these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.

- UMI or Unique Message ID is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the Exch2010DB agent would resemble: V-16-20079-7
Originator ID for all VCS products is 'V-16.' Category ID for Exch2010DB agent is 20079. Message ID is a unique number assigned to the message text.
- Message text denotes the actual message string.

You can view these message logs using Notepad or any text editor. All messages are logged to the engine and the agent logs. Messages of type CRITICAL and ERROR are also written to the Windows event log.

The following table lists the messages of type ERROR and WARNING.

Index

A

- agent functions
 - Exch2010DB agent 33
- agent state definition
 - Exch2010DB agent 33
- attributes
 - for Exch2010DB agent 34

C

- campus cluster
 - defined 21
 - failover using the forceimport attribute 83
 - forceimport 151
 - new installation 149
 - overview 24
 - site failure 151
 - verifying cluster configuration 150
- cloning for DR
 - Exchange service group 224
 - secondary storage (array-based replication) 220
 - secondary storage (VVR replication) 216
- cluster
 - configure LLT over ethernet 116
 - configure LLT over UDP 117
- clusters
 - assigning user privileges 211
 - configuring the cluster 114
 - switching online nodes 192
 - verifying campus cluster configuration 150
 - verifying the HA failover configuration 146
- cmdlets 260
- configuration overview
 - Exchange Server
 - disaster recovery 86
 - high availability 75
- configure
 - LLT over ethernet 116
 - LLT over UDP using VCW 117
- configure cluster
 - ethernet 116

UDP 117

configuring standalone Exchange Server 53

D

- disaster recovery (DR)
 - adding a new failover node
 - DR 251
 - cloning Exchange service group 224
 - cloning secondary storage (VVR replication) 216
 - configuring GCO with DR wizard 241
 - configuring replication with DR wizard 227
 - creating temporary storage (array-based replication) 220
 - DR wizard overview 212
 - DR wizard requirements 212
 - illustrated 28
 - multiple sites 247
 - typical configuration 28
- disk groups
 - cloning for secondary site (array-based replication) 220
 - cloning for secondary site (VVR replication) 216
 - creating 104
 - deporting 113
 - overview 99
 - preconditions 100
- DR wizard
 - cloning Exchange Server 2005 service group 224
 - cloning secondary storage (array-based replication) 220
 - cloning secondary storage (VVR replication) 216
 - configuring replication and GCO 227
 - overview 212
 - requirements 212
- drive letters
 - adding to mount volumes 112
- driver signing options

resetting 98

E

EMC SRDF

- configure SRDF replication with the DR wizard 235
- requirements for DR wizard 206

Exch2010DB agent

- attributes 34
- functions 33
- state definition 33
- type definition 34

Exchange

- tasks for setting up a secondary site 156

Exchange 2010 Database agent 32

Exchange high availability, VCS application agent 21

Exchange installation

- additional nodes
 - DR secondary site 224
- first node
 - DR secondary site 224

Exchange post-installation

- additional nodes
 - DR 224

Exchange Server

- configuring service group 140
- service group 140

Exchange Server 2010 agent

- supported services 32

Exchange Server Configuration Wizard 141

ExchDB2010 agent attributes

- DBName 34
- MonitorService 34

F

Fire Drill Wizard

- actions 266
- changing a fire drill configuration 274
- deleting the configuration 278
- overview 256
- preparing the configuration 266
- prerequisites for a fire drill 264, 265, 266
- restoring the prepared configuration 277
- running a fire drill 272

forceimport

- attribute for campus cluster 83
- defined 83

setting after a site failure 151

functions

- Exch2010DB agent 33

G

Global Cluster Option

- secure configuration 245

Global Cluster Option (GCO)

- configuring with the DR wizard 241
- prerequisites 241

H

HA

- Sample configuration 79

HA configuration overview 75

hardware configuration for a cluster 92

high availability (HA)

- defined 18
- verifying the failover 146

Hitachi TrueCopy

- configure replication with the DR wizard 238
- requirements for DR wizard 208

I

installing

- Exchange, secondary site tasks 156

installing SFW HA

- HA 94

L

LLT over ethernet

- configuring using VCW 116

LLT over UDP

- configuring using VCW 117

M

multiple DR sites 247

N

network and storage, configuring

- campus cluster 84

- DR (primary site) 202

network configuration for the cluster 92

P

- permissions requirements 73
- post-fire drill scripts 260
- prerequisites
 - service group for Exchange Server 140
 - SFW HA 70
- primary host, defined 27

R

- replicated data clusters
 - overview of setup 85
- replicated data clusters (RDC)
 - setting up a secondary Exchange site 156
- replication
 - configuring for EMC SRDF with the DR wizard 235
 - configuring for HTC with the DR wizard 238
 - configuring for VVR with DR wizard 227
 - defined 25
 - setting up a Replicated Data Set (RDS) 164
- requirements
 - Exchange 2010 software 71
 - permissions 73
- requirements, additional for SFW HA 73
- requirements, network 72
- requirements, system 72
- resetting
 - driver signing options 98
- resource type
 - Exch2010DB agent 34

S

- sample configurations
 - campus cluster 82
 - Exchange Server
 - RDC 84
- scripts 260
- secondary host, defined 27
- secondary site
 - setting up for disaster recovery 212
- secure clusters
 - assigning user privileges 211
- secure GCO, establishing 245
- Security Services
 - configuring 118
- service groups
 - configuring for Exchange Server 140
 - prerequisites for configuring for Exchange

- Server 140
- VCS Exchange Configuration wizard 141
- setting bandwidth
 - using RDS wizard 232
- SFW HA
 - additional requirements 73
 - best practices 74
 - network requirements 72
 - system requirements 72
- SFW HA installation 94
- site failure, forceimport attribute 151
- software requirements
 - Exchange 2010 71
- Solutions Configuration Center
 - context sensitivity 41
 - overview 39
 - running wizards remotely 43
 - wizard descriptions 43
 - workflow for active/active configuration 79
- SRDF
 - configuring replication with the DR wizard 235
- standalone Exchange server
 - converting to a “clustered” Exchange server
 - HA 122
- state definition
 - Exch2010DB agent 33
- storage cloning with the DR wizard
 - for array-based replication 220
 - for VVR replication 216
- storage hardware configuration 92
- supported services 32
- switching online nodes 192

T

- type definition
 - Exch2010DB agent 34

U

- user privileges
 - assigning 211

V

- VCS
 - configuring the cluster 114
 - configuring the Exchange Server service
 - group 140

- switching online nodes 192
- VCS Application Agent 21
- VCS Configuration Wizard 114
- volumes
 - adding drive letters 112
 - considerations for VVR and disaster recovery 102
 - creating on a cluster disk group 105
 - preconditions on a cluster disk group 100
 - sample configuration 103
- VVR
 - configuring replication with DR wizard 227
 - setting up RDS (VCS) 164
 - VxSAS service 204
- VVR security service configuration 161
- VxSAS service 161, 204