

Storage Foundation and High Availability Solutions 7.4 HA and DR Solutions Guide for Microsoft Exchange 2010 - Windows

Last updated: 2018-05-31

Legal Notice

Copyright © 2018 Veritas Technologies LLC. All rights reserved.

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

This product may contain third-party software for which Veritas is required to provide attribution to the third-party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

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

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

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

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

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

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Section 1	Introduction and Concepts	12
Chapter 1	Introducing Storage Foundation and High Availability Solutions for Microsoft Exchange Server	13
	About clustering solutions with InfoScale products	14
	About high availability	14
	How a high availability solution works	15
	How VCS monitors storage components	15
	Shared storage—if you use NetApp filers	16
	Shared storage—if you use SFW to manage cluster dynamic disk groups	17
	Shared storage—if you use Windows LDM to manage shared disks	17
	Non-shared storage—if you use SFW to manage dynamic disk groups	17
	Non-shared storage—if you use Windows LDM to manage local disks	18
	Non-shared storage—if you use VMware storage	18
	About SFW HA support for Exchange Server 2010	19
	About campus clusters	19
	Differences between campus clusters and local clusters	20
	Sample campus cluster configuration	20
	What you can do with a campus cluster	21
	About replication	21
	About a replicated data cluster	22
	How VCS replicated data clusters work	23
	About disaster recovery	24
	What you can do with a disaster recovery solution	25
	Typical disaster recovery configuration	25

Chapter 2	Introducing the VCS agent for Exchange 2010	27
	About the VCS database agent for Microsoft Exchange 2010	27
	Exchange 2010 database agent functions	28
	Exchange 2010 database agent state definitions	29
	Exchange 2010 database agent resource type definition	29
	Exchange 2010 database agent attribute definitions	30
	Exchange 2010 service group resource dependency graph	30
	Exchange 2010 service group sample configuration	31
Section 2	Configuration Workflows	34
Chapter 3	Configuring high availability for Exchange Server with InfoScale Enterprise	35
	Reviewing the HA configuration	36
	Sample Exchange server HA configuration	39
	Reviewing a standalone Exchange Server configuration	40
	Sample standalone Exchange server configuration	42
	Reviewing the campus cluster configuration	43
	Reviewing the Replicated Data Cluster configuration	44
	Sample Exchange Server Replicated Data Cluster configuration	45
	About setting up a Replicated Data Cluster configuration	46
	Reviewing the disaster recovery configuration	47
	Active-passive DR configuration	47
	Following the HA workflow in the Solutions Configuration Center	49
	VCS campus cluster configuration	50
	VCS Replicated Data Cluster configuration	52
	Disaster recovery configuration	54
	DR configuration tasks: Primary site	55
	DR configuration tasks: Secondary site	56
	About installing the Veritas InfoScale products	58
	Notes and recommendations for cluster and application configuration	58
	IPv6 support	61
	Campus cluster failover using the ForceImport attribute	62
	Configuring the storage hardware and network	64
	Configuring disk groups and volumes for Exchange Server	66
	About cluster disk groups and volumes	66
	Prerequisites for configuring cluster disk groups and volumes	67

Considerations for a fast failover configuration	68
Considerations for converting existing shared storage to cluster disk groups and volumes	69
Considerations when creating disks and volumes for campus clusters	69
Considerations for volumes for a Volume Replicator configuration	71
Sample disk group and volume configuration for Exchange 2010	71
Viewing the available disk storage	72
Creating a dynamic disk group	72
Adding disks to campus cluster sites	74
Creating volumes for high availability clusters	74
Creating volumes for campus clusters	78
About managing disk groups and volumes	82
Importing a disk group and mounting a volume	83
Unmounting a volume and deporting a disk group	84
Adding drive letters to mount the volumes	84
Deporting the cluster disk group	86
Configuring the cluster using the Cluster Configuration Wizard	86
Configuring notification	95
Adding nodes to a cluster	98

Chapter 4 Using the Solutions Configuration Center 103

About the Solutions Configuration Center	103
Starting the Solutions Configuration Center	104
Options in the Solutions Configuration Center	104
About launching wizards from the Solutions Configuration Center	105
Remote and local access to Solutions wizards	106
Solutions wizards and logs	107
Workflows in the Solutions Configuration Center	108

Section 3 Deployment 109

Chapter 5 Installing Exchange Server 2010 110

About installing Exchange Server 2010	110
Before you install Exchange Server 2010	110
Privileges required for installing Exchange 2010	111
Installing Exchange Server 2010	111
Creating mailbox databases on shared storage	111
Moving mailbox databases to shared storage	112

	Adding new Exchange servers to an existing cluster	113
Chapter 6	Configuring Exchange Server for failover	115
	Tasks for configuring a new server for high availability	115
	Tasks for configuring an existing server for high availability	117
	About configuring the Exchange 2010 service group	119
	Prerequisites for configuring the Exchange Server service group	119
	Creating the Exchange Server 2010 service group	120
	Configuring the service group in a non-shared storage environment	124
	Enabling fast failover for disk groups (optional)	126
	Verifying the Exchange Server cluster configuration	127
	Determining additional steps needed	128
Chapter 7	Configuring campus clusters for Exchange Server	129
	Tasks for configuring campus clusters	129
	Verifying the campus cluster: Switching the service group	130
	Setting the ForceImport attribute to 1 after a site failure	130
Chapter 8	Configuring Replicated Data Clusters for Exchange Server	132
	Tasks for configuring Replicated Data Clusters for Exchange Server	132
	Creating the primary system zone for the application service group	134
	Creating a parallel environment in the secondary zone	135
	Setting up security for Volume Replicator	136
	Setting up the Replicated Data Sets (RDS)	138
	Prerequisites for setting up the RDS for the primary and secondary zones	138
	Creating the Replicated Data Sets with the wizard	139
	Configuring a RVG service group for replication	150
	Creating the RVG service group	151
	Configuring the resources in the RVG service group for RDC replication	152
	Configuring the RVG Primary resources	160
	Configuring the primary system zone for the RVG service group	162
	Setting a dependency between the service groups	163
	Adding the nodes from the secondary zone to the RDC	163

Adding the nodes from the secondary zone to the RVG service group	164
Configuring secondary zone nodes in the RVG service group	166
Configuring the RVG service group NIC resource for fail over (VMNSDg only)	166
Configuring the RVG service group IP resource for failover	167
Configuring the RVG service group VMNSDg resources for fail over	169
Adding the nodes from the secondary zone to the Exchange Server service group	170
Configuring the zones in the Exchange Server service group	172
Configuring the application service group IP resource for fail over (VMNSDg only)	172
Configuring the application service group NIC resource for fail over (VMNSDg only)	173
Verifying the RDC configuration	174
Bringing the service group online	175
Switching online nodes	175
Additional instructions for GCO disaster recovery	176

Chapter 9

Deploying disaster recovery for Exchange Server	177
Tasks for deploying a disaster recovery configuration of Microsoft Exchange	178
Tasks for setting up DR in a non-shared storage environment	182
Reviewing the disaster recovery configuration	185
Supported disaster recovery configurations for service group dependencies	186
Setting up the secondary site: Installing InfoScale Enterprise and configuring a cluster	186
Verifying your primary site configuration	187
Setting up your replication environment	188
Requirements for EMC SRDF array-based hardware replication	189
Requirements for Hitachi TrueCopy array-based hardware replication	190
Assigning user privileges (secure clusters only)	192
About configuring disaster recovery with the DR wizard	193
Configuring disaster recovery with the DR wizard	195
Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)	198

Creating temporary storage on the secondary site using the DR wizard (array-based replication)	202
Installing Exchange 2010	206
Cloning the service group configuration from the primary site to the secondary site	206
Configuring the Exchange service group in a non-shared storage environment	210
Configuring replication and global clustering	210
Configuring Volume Replicator replication and global clustering	211
Configuring EMC SRDF replication and global clustering	218
Configuring Hitachi TrueCopy replication and global clustering	221
Configuring global clustering only	225
Creating the replicated data sets (RDS) for Volume Replicator replication	228
Creating the Volume Replicator RVG service group for replication	228
Configuring the global cluster option for wide-area failover	229
Linking clusters: Adding a remote cluster to a local cluster	230
Converting a local service group to a global service group	231
Bringing a global service group online	233
Verifying the disaster recovery configuration	234
Establishing secure communication within the global cluster (optional)	236
Adding multiple DR sites (optional)	238
Recovery procedures for service group dependencies	238
Possible task after creating the DR environment: Adding a new failover node to a Volume Replicator environment	240
Preparing the new node	241
Preparing the existing DR environment	241
Installing Exchange on the new node	242
Modifying the replication and Exchange service groups	242
Reversing replication direction	243

Chapter 10	Testing fault readiness by running a fire drill	244
	About disaster recovery fire drills	244
	About the Fire Drill Wizard	245
	About Fire Drill Wizard general operations	245
	About Fire Drill Wizard operations in a Volume Replicator environment	247

About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment	249
About post-fire drill scripts	251
Exchange 2010 scripts or cmdlets	252
Tasks for configuring and running fire drills	252
Prerequisites for a fire drill	254
Prerequisites for a fire drill in a Volume Replicator environment	255
Prerequisites for a fire drill in a Hitachi TrueCopy environment	256
Prerequisites for a fire drill in an EMC SRDF environment	256
Preparing the fire drill configuration	257
System Selection panel details	260
Service Group Selection panel details	261
Secondary System Selection panel details	261
Disk Selection panel details	261
Hitachi TrueCopy Path Information panel details	262
HTCSnap Resource Configuration panel details	262
SRDFSnap Resource Configuration panel details	263
Fire Drill Preparation panel details	264
Running a fire drill	264
Post fire drill operations panel details	266
Re-creating a fire drill configuration that has changed	267
Restoring the fire drill system to a prepared state	269
Deleting the fire drill configuration	271
Fire Drill Deletion panel details	272
Considerations for switching over fire drill service groups	272

Section 4 Reference 273

Appendix A Using Veritas AppProtect for vSphere 274

About Just In Time Availability	275
Prerequisites	279
Setting up a plan	281
Deleting a plan	283
Managing a plan	283
Viewing the history tab	285
Limitations of Just In Time Availability	285
Getting started with Just In Time Availability	286
Supported operating systems and configurations	288
Viewing the properties	289
Log files	289

	Plan states	290
	Troubleshooting Just In Time Availability	292
Appendix B	Troubleshooting	293
	VCS logging	293
	Exchange Service agent error messages	294
	Troubleshooting Microsoft Exchange uninstallation	296
	Troubleshooting Exchange Setup Wizard issues	297

Introduction and Concepts

- [Chapter 1. Introducing Storage Foundation and High Availability Solutions for Microsoft Exchange Server](#)
- [Chapter 2. Introducing the VCS agent for Exchange 2010](#)

Introducing Storage Foundation and High Availability Solutions for Microsoft Exchange Server

This chapter includes the following topics:

- [About clustering solutions with InfoScale products](#)
- [About high availability](#)
- [How a high availability solution works](#)
- [How VCS monitors storage components](#)
- [About SFW HA support for Exchange Server 2010](#)
- [About campus clusters](#)
- [Differences between campus clusters and local clusters](#)
- [Sample campus cluster configuration](#)
- [What you can do with a campus cluster](#)
- [About replication](#)
- [About a replicated data cluster](#)

- [How VCS replicated data clusters work](#)
- [About disaster recovery](#)
- [What you can do with a disaster recovery solution](#)
- [Typical disaster recovery configuration](#)

About clustering solutions with InfoScale products

Veritas InfoScale products provide the following clustering solutions for high availability and disaster recovery:

- High availability failover cluster in an active-passive configuration 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 Volume Replicator or hardware replication

About high availability

The term high availability 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 and application 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 configuring clustered environments using InfoScale Enterprise. InfoScale Enterprise includes Storage Foundation for Windows and 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 InfoScale Enterprise 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 applications.

How VCS monitors storage components

VCS provides specific agents that monitor storage components and ensure that the shared disks, disk groups, LUNs, volumes, and mounts are accessible on the system where the application is running. Separate agents are available for shared and non-shared storage and for third-party storage arrays such as NetApp filers. Your storage configuration determines which agent should be used in the high availability configuration.

For details on the various VCS storage agents, refer to the *Cluster Server Bundled Agents Reference Guide*.

Shared storage—if you use NetApp filers


The VCS hardware replication agents for NetApp provide failover support and recovery in environments that employ NetApp filers for storage and NetApp SnapMirror for replication. The agents enable configuring NetApp filers over an iSCSI or Fibre Channel (FC) connection in a VCS cluster environment.

The VCS agents for NetApp are as follows:

- NetAppFiler
- NetAppSnapDrive
- NetAppSnapMirror

These agents monitor and manage the state of replicated filer devices and ensure that only one system has safe and exclusive access to the configured devices at a time. The agents can be used in local clusters, single VCS replicated data clusters, and multi-cluster environments that are set up using the VCS Global Cluster Option (GCO).

In a typical configuration, the agents are installed on each system in the cluster. The systems are connected to the NetApp filers through a dedicated (private) storage network. VCS cluster systems are physically attached to the NetApp filer via an ethernet cable supporting iSCSI or FC as the transport protocol.



VCS also provides agents for other third-party hardware arrays. For details on the supported arrays, refer to the product Software Compatibility List (SCL).

Shared storage—if you use SFW to manage cluster dynamic disk groups

The VCS MountV and VMDg agents are used to monitor shared storage that is managed using Storage Foundation (SFW). SFW manages storage by creating disk groups from physical disks. These disk groups are further divided into volumes that are mounted on the cluster systems.

The MountV agent monitors volumes residing on disk groups. The VMDg agent monitors cluster dynamic disk groups and is designed to work using SCSI reservations. Together the MountV and VMDg agents ensure that the shared cluster dynamic disk groups and volumes are available.

Shared storage—if you use Windows LDM to manage shared disks

The VCS Mount and DiskReservation (DiskRes) agents are used to monitor shared disks that are managed using Windows Logical Disk Management (LDM).

The Mount agent monitors basic disks and mount points and ensures that each system is able to access the volume or mount path in the same way. The DiskRes agent monitors shared disks and uses persistent reservation to ensure that only one system has exclusive access to the disks. During failovers, these agents ensure that the disks and volumes are deported and imported on the node where the application is running.

Non-shared storage—if you use SFW to manage dynamic disk groups

VCS introduces the Volume Manager Non-Shared Diskgroup (VMNSDg) agent to support local non-shared storage configurations that are managed using SFW. The VMNSDg agent works without SCSI reservations and is designed for locally attached storage devices that do not support SCSI.

The VMNSDg agent monitors and manages the import and deport of dynamic disk groups created on local storage. The only difference between the VMDg agent and the VMNSDg agent is that the VMDg agent is designed for shared cluster dynamic disk groups and uses SCSI reservations, whereas the VMNSDg agent supports only non-shared local dynamic disk groups and works without SCSI reservations.

The VMNSDg agent can be used to set up single node Replicated Data Clusters (RDC) or Disaster Recovery (DR) configurations with replication set up between the sites.

During a failover, the VCS MountV and VMNSDg agents deport the locally attached storage from the affected node and then import the locally attached storage of the target node. Replication ensures that the data is consistent and the application is up and running successfully.

Note: The VMNSDg agent does not support fast failover and Intelligent Monitoring Framework (IMF).

Non-shared storage—if you use Windows LDM to manage local disks

VCS introduces the NativeDisks agent to support local non-shared storage configurations managed using Windows LDM. The NativeDisks agent works without SCSI reservations and is designed for local storage that does not support SCSI.

Together with the Mount agent, the NativeDisks agent monitors and manages the import and deport of basic local disks on the system. The only difference between the DiskRes agent and the NativeDisks agent is that the DiskRes agent is designed for shared disks and uses SCSI reservations, whereas the NativeDisks agent supports only non-shared local disks and works without SCSI reservations.

Note: The NativeDisks agent does not support fast failover and Intelligent Monitoring Framework (IMF).

Non-shared storage—if you use VMware storage

VCS introduces the VMwareDisks agent to support storage configurations in a VMware virtual environment. The agent is platform independent and supports VMware Virtual Machine Disk (VMDK), Raw Device Mapping (RDM) disk files (virtual), and storage that is configured using Network File System (NFS). The VMwareDisks agent works without SCSI reservations and supports locally attached non-shared storage.

VMware features such as snapshots, vMotion, and DRS do not work when SCSI disks are shared between virtual machines. The VMwareDisks agent is designed to address this limitation. With this agent, the disks can now be attached to a single virtual machine at a time in the VCS cluster. On failover, along with the service group, the VMwareDisks agent moves the disks to the target virtual machine.

The VMwareDisks agent communicates with the host ESXi server to configure storage. This agent manages the disk attach and detach operations on a virtual machine in the VCS cluster. The agent is VMware HA aware. During failovers, the agent detaches the disk from one system and then attaches it to the system where the application is actively running. The VMwareDisks agent presents the virtual disks to the operating system. On Windows, the agent relies on the VMNSDg agent (in case of SFW-managed local storage) and the NativeDisks agent (in case of LDM-managed local storage) for initializing and managing the virtual disks. On Linux, the agent relies on the LVM and VxVM agents.

Note: The VMwareDisks agent does not support fast failover and Intelligent Monitoring Framework (IMF).

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.

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 Storage Foundation (SFW) dynamic disk groups and volumes, but the volumes on each campus cluster node are mirrors of one another.

Sample campus cluster configuration

The following figure 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 Storage Foundation for Windows (SFW), a campus cluster can be set up using a 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 Volume Replicator for use in replication. Volume Replicator can be used for replication in either a replicated data cluster or a wide area disaster recovery solution.

The SFW HA disaster recovery solution also supports hardware replication.

For more information on Volume Replicator refer to the *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 RDC 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 Volume Replicator.

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

For Volume Replicator 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.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG service group.

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

Note: Volume Replicator 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 an application configured in a VCS replicated data cluster.


The configuration has the following 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

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

The following figure shows failover in a replicated data cluster.

Figure 1-2 Failover in a replicated data cluster




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

While this example required using shared storage, you can also set up an RDC cluster that uses non-shared storage. This involves installing and configuring the application on a single system in each of the RDC zones. The application data is located on the local disks attached to the system within each RDC zone. The data is replicated between the systems across the RDC zones to ensure concurrency.

The application service group is online on the single node in the primary RDC zone (Zone 0). In the event of a failure, VCS switches the service group to the node in the secondary RDC zone (Zone 1). Data replication ensures that the application is able to successfully handle client requests from the new node.

The following figure shows failover in a replicated data cluster using non-shared storage.

Figure 1-3 Failover in a replicated data cluster using non-shared storage



Note: The VCS VMNSDg agent is used to monitor the non-shared storage.

About disaster recovery

Wide area disaster recovery 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. 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.

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.

The following figure illustrates a typical DR configuration.

Figure 1-4 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, Exchange 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.

Introducing the VCS agent for Exchange 2010

This chapter includes the following topics:

- [About the VCS database agent for Microsoft Exchange 2010](#)
- [Exchange 2010 database agent functions](#)
- [Exchange 2010 database agent state definitions](#)
- [Exchange 2010 database agent resource type definition](#)
- [Exchange 2010 database agent attribute definitions](#)
- [Exchange 2010 service group resource dependency graph](#)
- [Exchange 2010 service group sample configuration](#)

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 | <p>The agent performs the following actions as part of its online function:</p> <ul style="list-style-type: none">■ Checks if the mailbox database file is available on the configured shared volume.■ Checks the status of the Microsoft Exchange Information Store (MSEExchangeIS) service and starts the service if it is not running.■ Updates the Windows Active Directory (AD) to bind the Exchange mailbox database to the Exchange Mailbox Server.■ Mounts the Exchange mailbox database on the node. |
| Offline | <p>Dismounts the Exchange mailbox database from the node.</p> |

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 (MSExchangeIS) 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.
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 int IMF{} = { Mode=2, MonitorFreq=1, RegisterRetryLimit=3 }
    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	Definition
DBName	Name of the Exchange 2010 mailbox database to be made highly available. Type and Dimension: string-scalar
MonitorService	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 (MSEExchangeSA) ■ Microsoft Exchange Mail Submission (MSEExchangeMailSubmission) Note: You cannot define which Exchange 2010 services should be monitored by the agent. Type and Dimension: boolean-scalar


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.

The following figure 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
//   }
// }
```

Configuration Workflows

- [Chapter 3. Configuring high availability for Exchange Server with InfoScale Enterprise](#)
- [Chapter 4. Using the Solutions Configuration Center](#)

Configuring high availability for Exchange Server with InfoScale Enterprise

This chapter includes the following topics:

- [Reviewing the HA configuration](#)
- [Reviewing a standalone Exchange Server configuration](#)
- [Reviewing the campus cluster configuration](#)
- [Reviewing the Replicated Data Cluster configuration](#)
- [Reviewing the disaster recovery configuration](#)
- [Following the HA workflow in the Solutions Configuration Center](#)
- [VCS campus cluster configuration](#)
- [VCS Replicated Data Cluster configuration](#)
- [Disaster recovery configuration](#)
- [About installing the Veritas InfoScale products](#)
- [Notes and recommendations for cluster and application configuration](#)
- [Campus cluster failover using the ForceImport attribute](#)
- [Configuring the storage hardware and network](#)

- [Configuring disk groups and volumes for Exchange Server](#)
- [About managing disk groups and volumes](#)
- [Configuring the cluster using the Cluster Configuration Wizard](#)


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 shared storage is monitored by specific VCS storage agents. 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.

The following figure illustrates a typical two-node Exchange 2010 failover configuration that uses shared storage. 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 3-1 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.

The following figure 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 3-2 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. 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.

Sample Exchange server HA configuration

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

The following table 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 3-1 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

The configuration process requires the following IP addresses; for an IPv4 network, you will need to specify the addresses; for an IPv6 network, they are generated:


Cluster IP address	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 Volume Replicator only)	For a disaster recovery configuration using Volume Replicator, an IP address is required for each Replicated Data Set (RDS), one for the primary site and one for the secondary site. Two IP addresses are required per Replicated Volume Group (RVG).

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.

The following figure 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 3-3 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.

The following figure 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 3-4 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.

The following table describes the objects created and used during the installation and configuration.

Table 3-2 Sample standalone Exchange 2010 HA configuration object

Name	Object
SYSTEM1, SYSTEM2, SYSTEM3	Servers
EXCH_SG1, EXCH_SG2	Exchange database service groups

Table 3-2 Sample standalone Exchange 2010 HA configuration object
(continued)

Name	Object
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 campus cluster involves an active-passive configuration for Exchange with one to one failover capabilities. In an active-passive configuration, one or more Exchange virtual servers can exist in a cluster, but each server must be managed by a service group configured with a set of nodes in the cluster. In this case, EVS1 can fail over from SYSTEM1 to SYSTEM2 and vice versa.

The following figure illustrates an active-passive configuration with one to one failover capabilities.

Figure 3-5 Campus cluster: Active-Passive configuration



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

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

Reviewing the Replicated Data Cluster configuration

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

- The IP address should be the same on all nodes at the primary and secondary zones.
- Replication IP address for the primary zone
- Replication IP address for the secondary zone

You should have these IP addresses available before you start deploying the RDC environment (for an IPv4 network, you will need to specify the addresses; for an IPv6 network, they are generated).

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.

The following table describes the objects created and used during the installation and configuration tasks.

Table 3-3 Exchange 2010 sample RDC configuration objects

Name	Object
Primary zone	
SYSTEM1, SYSTEM2	Servers at the primary zone
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 Volume Replicator replication
Secondary zone	
SYSTEM3, SYSTEM4	Servers at the secondary zone

All the other parameters are the same as on the primary zone.

Table 3-3 Exchange 2010 sample RDC configuration objects (*continued*)

Name	Object
RDS and Volume Replicator 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

The process involves the steps described in the following topics:

- See [“About setting up replication”](#) on page 46.
- See [“About configuring and migrating the service group”](#) on page 47.

About setting up replication

Use Volume Replicator to set up replication between the disk groups in the RDC primary and secondary zones.

Note the following:

- Use Volume Replicator to group the data volumes into a Replicated Volume Group (RVG). Create a primary RVG on the hosts in the first zone (primary zone) and create a secondary RVG on hosts in the second zone (secondary zone).
- If using shared storage, the primary RVG consists of volumes shared between the cluster nodes in the primary zone and the secondary RVG consists of volumes shared between the cluster nodes in the secondary zone.
- If using non-shared storage, the primary RVG consists of volumes created on the local disks on the node in the primary zone and the secondary RVG consists of the volumes created on the local disks on the node in the secondary zone.
- Create a Replicated Data Set (RDS) with the primary RVG and the secondary RVG.
- 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 Volume Replicator 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 storage. In this situation, none of the nodes in the primary zone see any device. The service group cannot fail over locally within the primary RDC zone, because the volumes cannot be mounted on any node. The service group must therefore fail over to a node in the current secondary RDC zone.

The RVGPrimary agent ensures that Volume Replicator volumes at the secondary RDC zone are made writable. The application can be started at the secondary RDC zone and run there until the problem with the local storage is corrected. If the storage problem is corrected, you can switch the application back to the primary zone using VCS.

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

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.

Review the DR information for the configurations you have planned as follows:


- See [“Active-passive DR configuration”](#) on page 47.

Active-passive DR configuration

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.


The following figure provides a view of an active-passive cluster configuration on the primary site.

Figure 3-6 Cluster configuration on the primary site



The following figure displays an environment that is prepared for a disaster with a DR solution. In this case, the primary site is replicating its application data to the secondary site.


Figure 3-7 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.

The following figure illustrates this type of failure.

Figure 3-8 Application services restored after primary site failure




You can choose to configure replication using Volume Replicator or an agent-supported array-based hardware replication. You can use the DR wizard to configure Volume Replicator 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.

Following the HA workflow in the Solutions Configuration Center

The Solutions Configuration Center helps you through the process of installing InfoScale Enterprise and configuring a new Storage Foundation and High Availability Solutions environment for Exchange.

The following figure shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

Figure 3-9 Configuration steps in the Solutions Configuration Center



See [“Workflows in the Solutions Configuration Center”](#) on page 108.

VCS campus cluster configuration

You can configure a new Storage Foundation and High Availability Solutions environment for your application 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.

Veritas recommends using the Solutions Configuration Center as a guide for installing InfoScale Enterprise and configuring SFW HA for your application.

See [“Workflows in the Solutions Configuration Center”](#) on page 108.

The following table outlines the high-level tasks to complete each objective for a campus cluster configuration for your application.

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

Action	Description
Review the campus cluster configuration	Review the sample configuration See “Reviewing the campus cluster configuration” on page 43.

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

Action	Description
Configure storage hardware and network	■ Set up the network and storage for a cluster environment ■ Verify the DNS entries for the systems on which the application will be installed
Review pre-requisites and install InfoScale Enterprise	■ Install InfoScale Enterprise on all the systems where you want to configure Exchange Server for high availability. Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.
Review application-specific requirements	See “Notes and recommendations for cluster and application configuration” on page 58.
Configure disk groups and volumes for Exchange Server	■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Add disks to campus cluster sites to enable site allocation. ■ Create dynamic volumes for the mailbox databases and logs See “Configuring disk groups and volumes for Exchange Server” on page 66.
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 See “Configuring the cluster using the Cluster Configuration Wizard” on page 86.
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 110.
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 111.
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 Exchange 2010 service group” on page 119.
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 130.

VCS Replicated Data Cluster configuration

The configuration process for a Replicated Data Cluster configuration includes the following main stages:

- Configure the SFW HA and the application components for high availability on the primary zone nodes.
- Install InfoScale Enterprise and configure SFW HA and the application components on the secondary zone.
- Configure the Volume Replicator components for both zones.
Refer to the *Volume Replicator Administrator's Guide* for additional details on Volume Replicator.

The following table outlines the high-level tasks to complete each objective for a Replicated Data Cluster configuration for your application.

Table 3-5 Process for deploying a Replicated Data Cluster

Action	Description
Understand the configuration	■ Understand active-passive configuration and zone failover in a RDC environment ■ Review the sample configuration See "Reviewing the Replicated Data Cluster configuration" on page 44.
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 the application will be installed
Review pre-requisites and install InfoScale Enterprise	■ Install InfoScale Enterprise on all the systems where you want to configure Exchange Server for high availability. Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.
Review application-specific requirements	See "Notes and recommendations for cluster and application configuration" on page 58.
Configure disk groups and volumes for the Exchange Server	■ Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) (if using a shared storage configuration) ■ Create dynamic disk groups using VEA (for a single-node configuration using non-shared storage) ■ Create dynamic volumes for the mailbox databases and logs See "Configuring disk groups and volumes for Exchange Server" on page 66.

Table 3-5 Process for deploying a Replicated Data Cluster *(continued)*

Action	Description
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 using the Cluster Configuration Wizard” on page 86.
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 110.
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 111.
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 Exchange 2010 service group” on page 119.
Create the primary system zone	■ Create the primary system zone ■ Add the nodes to the primary zone See “Creating the primary system zone for the application service group” on page 134.
Verify failover within the primary zone	Test failover between the nodes in the primary zone See “Verifying the Exchange Server cluster configuration” on page 127.
Create a parallel environment in the secondary zone	■ Install InfoScale Enterprise on the systems in the secondary zone ■ Configure disk groups and volumes using the same names as on the primary zone ■ Install your application following the prerequisites and guidelines for installing on the second zone. See “Creating a parallel environment in the secondary zone” on page 135.
Add the secondary zone systems to the cluster	Add the secondary zone systems to the cluster.

Table 3-5 Process for deploying a Replicated Data Cluster *(continued)*

Action	Description
Set up security for Volume Replicator on all cluster nodes	Set up security for Volume Replicator on all nodes in both zones. This step can be done at any time after installing InfoScale Enterprise on all cluster nodes, but must be done before configuring Volume Replicator replication. See “Setting up security for Volume Replicator” on page 136.
Set up the Replicated Data Set	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 138.
Configure a RVG service group	■ Create a Replicated Volume Group (RVG) service group ■ Configure the RVG service group See “Configuring a RVG service group for replication” on page 150.
Set a dependency between the service groups	■ Set up a dependency from the Volume Replicator RVG service group to the Exchange Server service group See “Setting a dependency between the service groups” on page 163.
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 163.
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 174.

Disaster recovery configuration

You begin by configuring the primary site for high availability. After setting up an SFW HA high availability environment for your application 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 Volume Replicator (Volume

Replicator) 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. Veritas recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [“Workflows in the Solutions Configuration Center”](#) on page 108.

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

The following table outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the primary site.

Table 3-6 Configuring the primary site for disaster recovery

Action	Description
Understand the configuration	Understand the DR configuration See “Reviewing the disaster recovery configuration” on page 185.
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 the application will be installed
Review pre-requisites and install InfoScale Enterprise	■ Install InfoScale Enterprise on all the systems where you want to configure Exchange Server for high availability. Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.
Review application-specific requirements	See “Notes and recommendations for cluster and application configuration” on page 58.
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 See “Configuring the cluster using the Cluster Configuration Wizard” on page 86.

Table 3-6 Configuring the primary site for disaster recovery (*continued*)

Action	Description
Configure disk groups and volumes for the Exchange Server	■ Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) (if using a shared storage configuration) ■ Create dynamic disk groups using VEA (for a single-node configuration using non-shared storage) ■ Create dynamic volumes for the mailbox databases and logs See “Configuring disk groups and volumes for Exchange Server” on page 66.
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 110.
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 111.
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 Exchange 2010 service group” on page 119.
Verify the primary site configuration	Test failover between nodes on the primary site See “Verifying the Exchange Server cluster configuration” on page 127.

DR configuration tasks: Secondary site

The following table outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the secondary site.

Table 3-7 Configuring the secondary site for disaster recovery

Action	Description
Review pre-requisites, install InfoScale Enterprise, configure the cluster on the secondary site	Warning: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster. ■ Install InfoScale Enterprise on all the systems where you want to configure Exchange Server for high availability. Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.
Verify that Exchange Server has been configured for high availability at the primary site	Verify that your application 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 187.

Table 3-7 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
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 Volume Replicator” on page 136. See “Requirements for EMC SRDF array-based hardware replication” on page 189. See “Configuring Hitachi TrueCopy replication and global clustering” on page 221.
(Secure cluster only) Assign user privileges	For a secure cluster only, assign user privileges See “Assigning user privileges (secure clusters only)” on page 192.
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 195.
Clone the storage configuration (Volume Replicator replication only)	(Volume Replicator 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 (Volume Replicator replication option)” on page 198.
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 202.
Install and configure Exchange Server	Follow the guidelines for installing Exchange Server in the Storage Foundation and High Availability Solutions environment See “Installing Exchange 2010” on page 206.
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 from the primary site to the secondary site” on page 206.

Table 3-7 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
Configure replication and global clustering, or configure global clustering only	■ (Volume Replicator 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 210.
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 234.
(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 236.
(Optional) Add additional DR sites	Optionally, add additional DR sites to a Volume Replicator environment See “Adding multiple DR sites (optional)” on page 238.
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 238.

About installing the Veritas InfoScale products

For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.

You can use Veritas Operations Manager to monitor the status of the application. For more information, see the Veritas Operations Manager product documentation.

Notes and recommendations for cluster and application configuration

- Review the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:

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

Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.

The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems.

Refer to the following Microsoft knowledge database article for more details:

<http://technet.microsoft.com/en-us/library/dd184075.aspx>

- 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 *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. Veritas recommends three NICs.
- NIC teaming is not supported for the VCS private network.
- Static IP addresses are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated.

Static IP addresses are required 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 Volume Replicator replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.

- For Volume Replicator replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
 Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in SFW HA because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
 See the *Cluster Server Bundled Agents Reference Guide*.
- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.
- For a disaster recovery configuration, all sites must reside in the same Active Directory domain.
- If Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).
- 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.
- If User Access Control (UAC) is enabled on Windows systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, C:\Program Files\Veritas). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.
- For a Replicated Data Cluster, install only in a single domain.

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- 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)
- This is applicable for a Replicated Data Cluster configuration.
 This is applicable for a Replicated Data Cluster configuration. You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, 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.
- To configure a RDC cluster, you need to create virtual IP addresses for the following:
 - Application virtual server; this IP address should be the same on all nodes at the primary and secondary zones
 - Replication IP address for the primary zone
 - Replication IP address for the secondary zone
 Before you start deploying your environment, you should have these IP addresses available.

IPv6 support

For IPv6 networks, the following is supported:

Types of addresses	The following types of IPv6 addresses are supported: ■ Unicast addresses: Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration: Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	LLT over UDP is supported on both IPv4 and IPv6. You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.

VCS agents, wizards, and other components

VCS agents that require an IP address attribute and wizards that configure or discover IP addresses now support IPv6 addresses (of the type described above).

The Veritas High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.

Note: Pure IPv4, pure IPv6, and dual-stack (IPv4 and IPv6 on the same system) configurations are supported.

Campus cluster failover using the ForceImport attribute

Automated recovery is handled differently in a VCS campus cluster than with a VCS local cluster. The outcomes of failure situations depend on the settings for the ForceImport attribute of the VMDg resource. To ensure proper failover in a VCS campus cluster, you must verify the value of the ForceImport attribute.

You can set this attribute as follows:

- ForceImport set to 1 automatically forces the import of the disk groups to the other node
- ForceImport set to 0 does not force the import

The advantage of automatic failover in the event of site failure comes at the cost of potential data loss if all storage and network communication paths between the sites are severed. Choose an option that is suitable given your cluster infrastructure, uptime requirements, and administrative capabilities.

You can use the VCS Java Console or command line to modify the ForceImport attribute. For more information on modifying ForceImport:

See [“Setting the ForceImport attribute to 1 after a site failure”](#) on page 130.

The following table lists failure situations and the outcomes depending on the settings for the ForceImport attribute of the VMDg resource.

Table 3-8 Failure situations in a VCS campus cluster

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic forced import)
1) Application fault May mean the services stopped for an application, a NIC failed, or a database table went offline.	Application automatically moves to another node.	Service Group failover is automatic on the standby or preferred system or node. Service Group failover is automatic on the standby or preferred system or node.
2) Server failure May mean a power cord became unplugged or a failure caused the system to stop responding.	Application automatically moves to other node. 100% of the disks are still available.	Service Group failover is automatic on the standby or preferred system or node. 100% of the mirrored disks are still available.
3) 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.
4) 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.
5) 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.
6) 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 3-8 Failure situations in a VCS campus cluster (*continued*)

Failure Situation	ForcelImport set to 0 (import not forced)	ForcelImport set to 1 (automatic forced import)
7) 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 offline again due to IP and other resources being online on original node. No interruption of service.

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat these procedures 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 the Microsoft cluster from mistakenly declaring a system down, Veritas recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters. Verify that each system can access the storage devices.
- 4 Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel by clicking **Start > Control Panel**.

On Windows 2012 operating systems, use the **Settings** menu from the **Start** screen.
- 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 that the public network adapter is the first bound adapter by following these steps sequentially:
 - In the Network Connections window, click **Advanced > Advanced Settings**.
 - In the Adapters and Bindings tab, verify that 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 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 appropriate IP version and then 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.

To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.

- 13 Close the window.

Configuring 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.

You create cluster disk groups if you are using a shared storage environment and dynamic disk groups in case of a non-shared storage environment.

Note: Cluster disk groups and volumes are used to store the Exchange mailbox databases and log files only. 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](#)
- [Prerequisites for configuring cluster disk groups and volumes](#)
- [Considerations for a fast failover configuration](#)
- [Considerations for converting existing shared storage to cluster disk groups and volumes](#)
- [Considerations when creating disks and volumes for campus clusters](#)
- [Considerations for volumes for a Volume Replicator configuration](#)
- [Sample disk group and volume configuration for Exchange 2010](#)
- [Viewing the available disk storage](#)

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

- [Creating a dynamic disk group](#)
- [Adding disks to campus cluster sites](#)
- [Creating volumes for high availability clusters](#)
- [Creating volumes for campus clusters](#)

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.

If using a shared storage, 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.

If using non-shared storage, you create dynamic disk groups and volumes on the locally attached storage on each node separately. Replication is set up between the volumes to ensure data concurrency.

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

Prerequisites for configuring cluster disk groups and volumes

Before you create a disk group (cluster disk group or dynamic 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.
- If using shared storage, 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.

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 71.

For more information on disk group and volume requirements for specific configurations, see the following topics:

- For service groups with many disk groups, you may want to implement the fast failover feature. Read the following topic: See [“Considerations for a fast failover configuration”](#) on page 68.
- If the existing databases and logs are already on shared storage, read the following topic: See [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 69.
- For more information on disk groups and volumes for campus clusters, read the following topic: See [“Considerations when creating disks and volumes for campus clusters”](#) on page 69.
- For a Replicated Data Cluster (RDC) configuration or a disaster recovery (DR) configuration using Volume Replicator (Volume Replicator), read the following topic: See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 71.

Considerations for a fast failover configuration

For VCS service groups that contain many disk groups, you can greatly reduce failover time by implementing fast failover.

Fast failover speeds up the failover of storage resources in several ways:

- Fast failover provides a "read-only deported" mode for disk groups on inactive nodes. This mode speeds up the process of importing a disk group.
- Fast failover maintains the current disk group configuration in memory on the inactive nodes. Any changes are automatically synchronized so that all nodes maintain an identical disk group configuration.

For more details about fast failover, refer to the *Storage Foundation Administrator's Guide*.

Take the following storage-related requirements into account if you are planning to implement fast failover:

- Fast failover is currently not supported for the following:
 - RAID-5 volumes
 - SCSI-2
 - Active/Passive (A/P) arrays for DMP
- In synchronous mode of replication, if fast failover is set, then the RVG cannot be stopped and started when a disk group fails over to another node. If the RLINK is in hard synchronous mode, it may not be connected when the volume arrives, and the I/Os may fail. In such case, the Event Viewer displays NTFS and ReFS errors and file system reports the volume as RAW. Therefore, fast failover is not supported if the RLINK is in hard synchronous mode.
- The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click Properties. Disk group version upgrade is required after upgrading SFW HA on the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

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

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.

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

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

See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 71.

Considerations when creating disks and volumes for campus clusters

When you create the disk groups for a campus cluster, ensure that each disk group has the same number of disks on each physical site. You create each volume as 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.

Veritas recommends using the SFW site-aware allocation feature for campus cluster storage. Site-aware allocation can ensure that site boundary limits are maintained for operations like volume grow, subdisk move, and disk relocation.

Enabling site-aware allocation for campus clusters requires the following steps in the VEA:

- After creating the disk groups, you tag the disks with site names to enable site-aware allocation. This is a separate operation, referred to in the VEA as adding disks to a site.
As an example, say you had a disk group with four disks. Disk1 and Disk2 are physically located on Site A. Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to "site_a" and add Disk3 and Disk4 to "site_b".
- During volume creation, you specify the volume site type as Site Separated. This ensures that the volume is restricted to the disks on the selected site.

Note: The hot relocation operation does not adhere to site boundary restrictions. If hot relocation causes the site boundary to be crossed, then the Site Separated property of the volumes is changed to Siteless. This is done so as not to disable hot relocation. To restore site boundaries later, you can relocate the data that crossed the site boundary back to a disk on the original site and then change back the properties of the affected volumes.

For more information on site-aware allocation, refer to the *Storage Foundation Administrator's Guide*.

When you create the volumes for a campus cluster, consider the following:

- During disk selection, configure the volume as "Site Separated" and select the two sites of the campus cluster from the site list.
- For volume attributes, select the "mirrored" and "mirrored across enclosures" options.
- Veritas 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.
- During the volume creation procedure for Site Separated volumes, you can only create as many mirrors as there are sites. However, once volume creation is complete, you can add additional mirrors if desired.
- Choosing "Mirrored" and the "mirrored across" option without having two enclosures that meet requirements causes new volume creation to fail.

- You cannot selecting RAID-5 for mirroring.
- Selecting "stripe across enclosures" is not recommended because then you need four enclosures, instead of two.
- Logging can slow performance.

Considerations for volumes for a Volume Replicator configuration

For a configuration using Volume Replicator, either a disaster recovery configuration on a secondary site or a Replicated Data Cluster, note the following:

- Volume Replicator does not support the following types of volumes:
 - SFW (software) RAID 5 volumes
 - Volumes with the Dirty Region Log (DRL)
 - Volumes with commas in the names
 - Data Change Object (DCO)
- A configuration with Volume Replicator 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 Volume Replicator planning, see the *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 > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator**, or, on Windows 2012 operating systems, from the **Apps** menu, 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 dynamic disk group

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

Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

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

Follow the steps in this section to create one or more disk groups for your application.

To create a dynamic disk group

- 1 Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) or, on Windows 2012 operating systems, from the **Apps** menu, 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 as follows:
 - Enter the name of the disk group (for example, **SG1_DG**).
 - Check the **Create cluster group** check box if you wish to create cluster dynamic disk groups that are used in a shared storage environment.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.

Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier.

For example, entering **TestGroup** as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.

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

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

Adding disks to campus cluster sites

For campus cluster storage, Veritas recommends using Storage Foundation (SFW) site-aware allocation. To enable site-aware allocation, you assign a site name to disks after they are added to a disk group. In the VEA assigning a site name is referred to as adding disks to a site.

For example, Disk1 and Disk2 are physically located on Site A and Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to site_a and add Disk3 and Disk4 to site_b.

To add disks to a site

- 1 From the VEA console, right-click a disk that needs to be added to a site and select **Add Disk to Site**.

Disks must be part of a dynamic disk group in order to add them to a site.

- 2 In the Add Disk to a Site screen, choose one of the following:
 - Choose **Select a new site** and specify a new site name.
The site name can include any alphanumeric value and valid characters like the period (.), dash (-), and underscore (_). It cannot exceed 31 characters. Site names are case insensitive; all names are converted to lowercase.
 - Choose **Available Sites** and select a site from the list.
- 3 From the **Available Disks** column, select the disk or disks to add to the specified site.
- 4 Click **OK**.

Creating volumes for high availability clusters

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

You can use this procedure for volumes in a high availability cluster. For volumes in a campus cluster, see the following:

- See [“Creating volumes for campus clusters”](#) on page 78.

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

- See [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 69.
- See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 71.

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 > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu, 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 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 that the appropriate disk group name appears that in the **Group name** drop-down list.
 - For Site Preference, leave the setting as **Siteless** (the default).
 - Automatic disk selection is the default setting. To manually select the disks, click **Manually select disks** and use the **Add** and **Remove** buttons to move the appropriate disks to the **Selected disks** list. Manual selection of disks is recommended.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling track alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
 - Click **Next**.
- 7 Specify the volume attributes.

- Enter a 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 volume layout type. To select mirrored striped, select both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
In the Mirror Info area, select the appropriate mirroring options.
 - 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 disk.

- If creating a Replicator Log volume for Volume Replicator, select **Do not assign a drive letter**.
 - Click **Next**.
- 9** Create an NTFS file system.
- Make sure that the Format this volume checkbox is checked and click **NTFS**.
 - For a Volume Replicator 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**.
- 10** Click **Finish** to create the new volume.
- 11** Repeat these steps to create additional volumes.

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

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

Creating volumes for campus clusters

This section will guide you through the process of creating a volume on a dynamic disk group for a campus cluster.

For creating volumes for other types of clusters:

- See [“Creating volumes for high availability clusters”](#) on page 74.

Before you begin, review the following topics:


- See [“Considerations when creating disks and volumes for campus clusters”](#) on page 69.
- See [“Adding disks to campus cluster sites”](#) on page 74.

Use the following procedure to create dynamic volumes for a campus cluster.

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

To create dynamic volumes

- 1 Launch the VEA console from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
 You can right-click the disk group you have just created.
 For example, **SG1_DG**.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume as follows:



Group name	Make sure the appropriate disk group is selected.
Site preference	Select the Site Separated option.
Select site from	Select the campus cluster sites. Press CTRL to select multiple sites. Note: If no sites are listed, the disks have not yet been added to a site.
Auto select disks	Automatic disk selection 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 picks two disks (one from each array) with the most space.
Manually select disks	If you manually select disks, use the Add and Remove buttons to move the appropriate disks to the Selected disks list.
Disable Track Alignment	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 as follows:

localhost - New Volume Wizard

New Volume Wizard

Select the attributes for this volume.

Volume name: EVS1_SG1_D81

Size: 1 GB Max Size

Layout

☒ Concatenated

☐ Striped

☐ RAID-5

Columns: 2

Stripe unit size (Sectors): 128

☐ Stripe across: Port

Mirror Info

☒ Mirrored

Total mirrors: 2

☐ Mirror across: Port

☐ Enable logging

Concatenated: A simple volume with a single copy of data on one or more disks.

< Back

Next >

Cancel

Help

Volume name	Specify a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
Size	Specify a size for the volume. If you click Max Size , the Size box shows the maximum possible volume size for that layout in the dynamic disk group.
Layout	Ensure that the Mirrored checkbox is selected. Select either the Concatenated or Striped layout type. If you are creating a striped volume, the Columns and Stripe unit size boxes need to have entries. Defaults are provided. In addition, click the Stripe across checkbox and select Ports from the drop-down list.
Mirror Info	Click Mirror across and select Enclosures from the drop-down list. When creating a site separated volume, as required for campus clusters, the number of mirrors must correspond to the number of sites. If needed, you can add more mirrors after creating the volume.
Enable logging	Verify that this option is not selected.

Click **Next**.

- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
 - To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
 - To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.

Click **Next**.

- 9 Create an NTFS file system.
 - Make sure the **Format this volume** checkbox is checked and select **NTFS**.
 - Select an allocation size or accept the default.
 - The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space.
Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 10 Click **Finish** to create the new volume.
- 11 Repeat these steps to create additional volumes as needed.

Note: 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 cluster dynamic 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.

Managing disk groups and volumes involves the following:

- See [“Importing a disk group and mounting a volume”](#) on page 83.
- See [“Unmounting a volume and deporting a disk group”](#) on page 84.

Note: (Disaster recovery configurations only) 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 *Storage Foundation Administrator's Guide* for more information.

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 **Assign a Drive Letter** and select a drive letter from the drop-down list.
 - 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.

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

Before installing Exchange on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

Note: This is not applicable if you are using non-shared storage (dynamic disk groups).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Launch the VEA console from **Start > All Programs > Veritas > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu, 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 using the Cluster Configuration Wizard

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

and global clusters (GCO). You can also use VCW to modify or delete cluster configurations.

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

Note the following prerequisites before you proceed:

- The required network adapters (NICs), and SCSI controllers are installed and connected to each system.

Veritas recommends the following actions for network adapters:

- Disable the ethernet auto-negotiation options on the private NICs to prevent:
 - Loss of heartbeats on the private networks
 - VCS from mistakenly declaring a system as offlineContact the NIC manufacturer for details on this process.
- Remove TCP/IP from the private NICs to lower system overhead.
- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

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

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list. When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

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

To configure a VCS cluster using the wizard

- 1 Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.

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

- Clear **Specify systems and users manually**.

- Click **Next**.

Proceed to step 8.

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

- Select **Specify systems and users manually**.

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

- Click **Next**.

If you chose to retrieve the list of systems, proceed to step 6. Otherwise, proceed to the next step.

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

Do not specify systems that are part of another cluster.

Proceed to step 8.

- 6 On the System Selection panel, specify the systems for the cluster and then click **Next**.

Do not select systems that are part of another cluster.

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

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

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

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

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- Product is either not installed or there is a version mismatch.

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

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

Veritas Cluster Server Configuration Wizard

Cluster Details
Enter necessary details to create the new cluster

Specify the cluster name and cluster ID. If you chose to specify the systems manually, VCV does not validate the cluster ID.

Cluster Name:

Cluster ID:

Operating System:

Select the systems to create the cluster.

☒ **Select all systems**

Available Systems

- ☒ ROGER
- ☒ SCOOPYDU

Total number of systems selected to create the cluster : 2

Click 'Next' to continue.

VERITAS

Back Next Cancel

Specify the cluster details as follows:

- | | |
|------------------|---|
| Cluster Name | Type a name for the new cluster. Veritas recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | <p>Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 65535.</p> <p>Note: If you chose to specify systems and users manually in step 4 or if you share a private network between more than one domain, make sure that the cluster ID is unique.</p> |
| Operating System | <p>From the drop-down list, select the operating system.</p> <p>All the systems in the cluster must have the same operating system and architecture.</p> |

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

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

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

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


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

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

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

Do one of the following:

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



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

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

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

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

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

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

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

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

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

Specify the domain user details as follows:

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

- To specify a new user, click **New user** and type a valid user name in the Create New User field and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

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

Do one of the following:

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

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

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

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

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

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

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

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

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
See [“Configuring notification”](#) on page 95.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.
You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.

Configuring notification


This section describes steps to configure notification.

To configure notification

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

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

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



The screenshot shows the 'Notifier SNMP Configuration' window of the Veritas Cluster Server Configuration Wizard. The window has a blue title bar and a sidebar on the left with navigation links: Domain Selection, Create Cluster, Select Components, Configure, Summary, and Finish. The main area is titled 'Notifier SNMP Configuration' with the subtitle 'Specify information about SNMP console.' Below this, it says 'Enter the name or the IP address of the SNMP console and then select the desired severity level.' There is a table with two columns: 'SNMP Console' and 'Severity Information'. The first row has a text input field with the placeholder 'Click here to change the text..' and a dropdown menu with 'Information' selected. Below the table, there are instructions: 'Click on '+' button to add more consoles.' and 'Click '-' to remove a console.' with corresponding '+' and '-' buttons. The 'SNMP Trap Port' is set to '162'. A note states: 'Note: SNMP console must be MIB 2.0 compliant.' and a prompt says 'Click 'Next' to continue.' At the bottom right are 'Back', 'Next', and 'Cancel' buttons. The Veritas logo is in the bottom left corner.

SNMP Console	Severity Information
Click here to change the text..	Information

SNMP Trap Port: 162

Note: SNMP console must be MIB 2.0 compliant.

Click 'Next' to continue.

Do the following:

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

- Enter an SNMP trap port. The default value is 162.
- 3** If you chose to configure SMTP, specify information about SMTP recipients and then click **Next**.

The screenshot shows the 'Notifier SMTP Configuration' window of the Veritas Cluster Server Configuration Wizard. The window has a blue title bar and a sidebar on the left with navigation links: Domain Selection, Create Cluster, Select Components, Configure, Summary, and Finish. The main area is titled 'Notifier SMTP Configuration' with the subtitle 'Specify information about SMTP recipients.' It contains a text box for 'SMTP Server Name / IP', a table for 'Enter SMTP recipients and select a severity level for each recipient.', and buttons for '+', '-', and 'Next'. The table has two columns: 'Recipients' and 'Severity'. The 'Recipients' column has a placeholder text 'Click here to change the text..'. The 'Severity' column has a placeholder text 'Information'. Below the table, there are instructions: 'Click '+' to add a recipient. Click '-' to remove a recipient.' and 'Click 'Next' to continue.' At the bottom right, there are three buttons: 'Back', 'Next', and 'Cancel'. The Veritas logo is visible in the bottom left corner of the main area.

Recipients	Severity
Click here to change the text..	Information

Do the following:

- Type the name of the SMTP server.
 - Click a field in the **Recipients** column and enter a recipient for notification. Enter recipients as admin@example.com.
 - Click the corresponding field in the **Severity** column and select a severity level for the recipient.
VCS sends messages of an equal or higher severity to the recipient.
 - Click the + icon to add fields; click the - icon to remove a field.
- 4** On the Notifier Network Card Selection panel, specify the network information and then click **Next**.

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster.
 The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS starts and click **Configure**.
 - 6 Click **Finish** to exit the wizard.

Adding nodes to a cluster

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.

Click **Start > All Programs > Veritas > 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.

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

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

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

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

If you chose to retrieve the list of systems, proceed to step 6. Otherwise proceed to the next step.

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

- Type the name of an existing 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.

- 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.

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

- The system does not respond to a ping request.
- WMI access is disabled on the system.
- The wizard is unable to retrieve information about the system's 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.

Click **Next** to proceed.

- 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 (a cluster that is not a secure cluster).

- 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, do the following:
 - Select the check boxes next to the two NICs to be assigned to the private network.

Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
 - If you have only two NICs on a selected system, 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. Veritas recommends that you do not select teamed NICs for the private network. The wizard configures the LLT service (over Ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, do the following:
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Veritas recommends reserving at least two NICs exclusively for the VCS private network. 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.
 - 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. In case of IPv4, each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
 - For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 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 credentials 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**.

Using the Solutions Configuration Center

This chapter includes the following topics:

- [About the Solutions Configuration Center](#)
- [Starting the Solutions Configuration Center](#)
- [Options in the Solutions Configuration Center](#)
- [About launching wizards from the Solutions Configuration Center](#)
- [Remote and local access to Solutions wizards](#)
- [Solutions wizards and logs](#)
- [Workflows in the Solutions Configuration Center](#)

About the Solutions Configuration Center

The Storage Foundation and High Availability Solutions Configuration Center guides you through setting up your Storage Foundation (SFW) or Storage Foundation and High Availability Solutions (SFW HA) environment.

The Solutions Configuration Center provides solutions for the following applications:

- Microsoft Exchange Server 2010
- Microsoft SQL Server 2012, 2014, and 2016
- Enterprise Vault Server (high availability and disaster recovery solutions)
- Microsoft SharePoint Server 2010 and 2013 (high availability, disaster recovery, and Quick 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 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
- Fire drill to test the fault readiness of a disaster recovery environment

Starting the Solutions Configuration Center

Depending on the operating system, you can start the Solutions Configuration Center from the **All Programs** menu, the **Run** menu, or from the **Apps** menu.

To start the Solutions Configuration Center

- ◆ Click **Start > All Programs > Veritas > Veritas Storage Foundation > Solutions Configuration Center**.

or

Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**.

or

Click **Start > Run**, type **scc**, and press Enter.

or


Navigate to the Apps menu and then click **SCC**.

Options in the Solutions 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 steps that are shown when you click on a solution are customized for that application.

The following figure shows the solutions available when you click Solutions for Microsoft Exchange.

Figure 4-1 Solutions Configuration Center for Microsoft Exchange



About launching wizards from the Solutions Configuration Center

The Solutions Configuration Center provides two ways to access wizards:

Applications	Lists solutions by application. Provides step-by-step configuration instructions that include buttons to launch the appropriate wizard for each step.
--------------	---

Solutions	(For advanced users) Lists wizards by solution, without additional instructions, as follows: ■ High Availability Configuration Wizards ■ Disaster Recovery Configuration Wizards ■ Quick Recovery Configuration Wizards ■ Fire Drill Configuration Wizards You can go directly to a particular wizard. Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems. The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems. Refer to the following Microsoft knowledge database article for more details: http://technet.microsoft.com/en-us/library/dd184075.aspx
-----------	--

Remote and local access to Solutions wizards

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

Disaster Recovery Configuration Wizard	Configures wide area disaster recovery, including cloning storage, cloning service groups, and configuring the global cluster Can also configure: ■ Volume Replicator (Volume Replicator) replication ■ VCS resource for EMC SRDF and Hitachi TrueCopy array-based hardware replication Note: Requires first configuring high availability on the primary site. To configure IPv6 settings, launch the wizard from a system on which the IPv6 stack is installed.
Fire Drill Wizard	Sets up a fire drill to test disaster recovery Note: Requires first configuring high availability on the primary site. To configure IPv6 settings, launch the wizard from a system on which the IPv6 stack is installed.

Quick Recovery Configuration Wizard	Schedules preparation of snapshot mirrors and schedules the Quick Recovery snapshots
VCS Configuration Wizard	Sets up the VCS cluster
Volume Replicator Security Service Configuration Wizard	Configures the Volume Replicator security service

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 2010 Configuration Wizard	Configures the service group for Microsoft Exchange Server 2010 high availability
MSMQ Configuration Wizard	Configures a Microsoft Message Queuing (MSMQ) service group
SFW Configuration Utility for Hyper-V Live Migration Support	Configures SFW for Microsoft Hyper-V Live Migration support on the selected systems

Solutions wizards and 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 folder:

C:\ProgramData\Veritas\winsolutions\log


Workflows in the Solutions 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 Solutions 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.

The following figure shows the high-level overview of the workflow steps for configuring high availability for Exchange from the Solutions Configuration Center.

Figure 4-2 Workflow for configuring high availability for Exchange



Deployment

- [Chapter 5. Installing Exchange Server 2010](#)
- [Chapter 6. Configuring Exchange Server for failover](#)
- [Chapter 7. Configuring campus clusters for Exchange Server](#)
- [Chapter 8. Configuring Replicated Data Clusters for Exchange Server](#)
- [Chapter 9. Deploying disaster recovery for Exchange Server](#)
- [Chapter 10. Testing fault readiness by running a fire drill](#)

Installing Exchange Server 2010

This chapter includes the following topics:

- [About installing Exchange Server 2010](#)
- [Creating mailbox databases on shared storage](#)
- [Moving mailbox databases to shared storage](#)
- [Adding new Exchange servers to an existing cluster](#)

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.
- Veritas 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.

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.
 The user must have write permissions for objects corresponding to these nodes in the Active Directory.

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 72.
 - See [“Creating a dynamic disk group”](#) on page 72.
 - See [“Adding disks to campus cluster sites”](#) on page 74.
 - See [“Importing a disk group and mounting a volume”](#) on page 83.
 - See [“Adding drive letters to mount the volumes”](#) on page 84.
- 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 69.

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 72.

See [“Creating a dynamic disk group”](#) on page 72.

See [“Adding disks to campus cluster sites”](#) on page 74.

See [“Importing a disk group and mounting a volume”](#) on page 83.

See [“Adding drive letters to mount the volumes”](#) on page 84.
- 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.
- On the node on which you run the wizard, select Services on the Administrative Tools menu and verify that the Veritas High Availability Engine service is running.

Configuring Exchange Server for failover

This chapter includes the following topics:

- [Tasks for configuring a new server for high availability](#)
- [Tasks for configuring an existing server for high availability](#)
- [About configuring the Exchange 2010 service group](#)
- [Prerequisites for configuring the Exchange Server service group](#)
- [Creating the Exchange Server 2010 service group](#)
- [Configuring the service group in a non-shared storage environment](#)
- [Verifying the Exchange Server cluster configuration](#)
- [Determining additional steps needed](#)

Tasks for configuring a new server for high availability

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

Table 6-1 Task list: Exchange Server HA configuration tasks

Action	Description
Review the HA configuration	■ Understand active-passive and any-to-any configuration ■ Review the sample configuration See “Reviewing the HA configuration” on page 36.
Configure the storage hardware and network	■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which the application will be installed
Review pre-requisites and install InfoScale Enterprise	Install InfoScale Enterprise on all the systems where you want to configure Exchange Server for high availability. For more information, see the <i>Veritas InfoScale Installation and Upgrade Guide</i>.
Review application-specific requirements	See “Notes and recommendations for cluster and application configuration” on page 58.
Configure disk groups and volumes for the Exchange Server	■ Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) (if using a shared storage configuration) ■ Create dynamic disk groups using VEA (for a single-node configuration using non-shared storage) ■ Create dynamic volumes for the mailbox databases and logs See “Configuring disk groups and volumes for Exchange Server” on page 66.
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 See “Configuring the cluster using the Cluster Configuration Wizard” on page 86.
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 110.
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 111.
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 Exchange 2010 service group” on page 119.

Table 6-1 Task list: Exchange Server HA configuration tasks *(continued)*

Action	Description
Configure fast failover for disk groups (optional)	■ Ensure that you have installed the Fast Failover option and met the prerequisites for storage ■ Use the Java Console to enable the FastFailover attribute for VMDg resources. See “Enabling fast failover for disk groups (optional)” on page 126.
Verify the HA configuration	Test failover between nodes See “Verifying the Exchange Server cluster configuration” on page 127.
Proceed to the additional steps depending on the desired HA configuration.	See “Determining additional steps needed” on page 128.

Tasks for configuring an existing server for high availability

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 Storage Foundation and High Availability Solutions environment.

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

Table 6-2 Task list: Standalone Exchange Server HA configuration tasks

Action	Description
Review the HA configuration	■ Review the sample configuration See “Reviewing a standalone Exchange Server configuration” on page 40.
Configure the storage hardware and network	■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which the application will be installed
Review pre-requisites and install InfoScale Enterprise	Install InfoScale Enterprise on all the systems where you want to configure Exchange Server for high availability. See the <i>Veritas InfoScale Installation and Upgrade Guide</i>.

Table 6-2 Task list: Standalone Exchange Server HA configuration tasks
(continued)

Action	Description
Review application-specific requirements	See “Notes and recommendations for cluster and application configuration” on page 58.
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 disk groups and volumes for Exchange Server” on page 66.
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 application server systems to the cluster See “Configuring the cluster using the Cluster Configuration Wizard” on page 86.
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 112.
Install and configure Exchange on the additional nodes	Install Exchange Server on additional nodes if required See “About installing Exchange Server 2010” on page 110.
Create an Exchange service group	■ Create a Exchange Server 2010 service group using the Exchange 2010 Service Group Configuration Wizard See “About configuring the Exchange 2010 service group” on page 119.
Configure fast failover for disk groups (optional)	■ Ensure that you have installed the Fast Failover option and met the prerequisites for storage ■ Use the Java Console to enable the FastFailover attribute for VMDg resources. See “Enabling fast failover for disk groups (optional)” on page 126.
Verify the HA configuration	Test fail over between nodes See “Verifying the Exchange Server cluster configuration” on page 127.

About configuring the Exchange 2010 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.

Refer to the following topics:

- See [“Prerequisites for configuring the Exchange Server service group”](#) on page 119.
- See [“Creating the Exchange Server 2010 service group”](#) on page 120.

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 [“Tasks for configuring a new server for high availability”](#) on page 115.
- See [“Tasks for configuring an existing server for high availability”](#) on page 117.
- See [“VCS campus cluster configuration”](#) on page 50.
- See [“VCS Replicated Data Cluster configuration”](#) on page 52.
- 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.

- On the node on which you run the wizard, select Services on the Administrative Tools menu and verify that the Veritas High Availability Engine service is running.
- Import the disk group (cluster disk group in case of shared storage and dynamic disk group in case of non-shared storage) and mount the volumes created to store the following data:
 - Exchange database
 - Database logMount 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 required ports and services to the Firewall Exception list.

Creating the Exchange Server 2010 service group

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

If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

See [“Configuring the service group in a non-shared storage environment”](#) on page 124.

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 > Veritas > Veritas Cluster Server > Configuration Tools > Exchange 2010 Configuration Wizard**.

or

On Windows Server 2012 operating systems, launch the wizard from the **Apps** menu on the Start screen.

- 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.

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_<databasename>. Here, <databasename> 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.

Configuring the service group in a non-shared storage environment

If you are using a non-shared storage configuration, you have to use the VCS MountV – VMNSDg agents to monitor your local storage. Currently, the service group configuration wizards do not support configuring these agents in the service group. You have to configure these agents manually by using the Cluster Manager (Java Console) or the VCS commands.

VCS provides templates for configuring service groups that use non-shared storage agent resources.

The Java Console templates are located in the following directory:

```
%VCS_HOME%\Templates
```

Here, `%VCS_HOME%` is the default product installation directory, typically, `C:\Program Files\Veritas\Cluster Server`.

For information about adding a service group using templates from the Java Console, refer to the *Cluster Server Administrator's Guide*.

The following steps describe how to create a service group using the Cluster Manager (Java Console).

To configure the service group in a non-shared storage environment

- 1 Open the Cluster Manager (Java Console) from **Start > All Programs > Veritas > Veritas Cluster Server** and then click **Veritas Cluster Manager - Java Console** or, on Windows Server 2012 operating systems, from the **Apps** menu.
- 2 Log on to the cluster. On the Cluster Monitor window click **File > New Cluster**, then on the New Cluster window type **localhost** in the Host name field, and then click **OK**.
- 3 Launch the service group configuration wizard. From the Cluster Explorer window menu, click **Tools > Configuration Wizard**.
- 4 On the Service Group Configuration Wizard Welcome panel, click **Next**.
- 5 Fill in the following information and then click **Next**:
 - Specify a name for the service group.
 - Select the systems for the service group. Click a system in the Available Systems box and then click the right arrow to move the systems to Systems for Service Group.
 - Leave the service group type as the default, Failover.
- 6 Click **Next** again.

- 7 In the Templates list, select the desired service group template depending on the configuration and then click **Next**.

Template name	Description
Exchange2010SG-VMNSGroup	Use this template to create a single node high availability service group that uses non-shared storage. This template includes resources for configuring MountV and VMNSDg agents.
Exchange2010SG-VirtVMNSGroup	Use this template to create a single node high availability service group in a VMware virtual environment. This template includes resources for configuring MountV, VMwareDisks, and VMNSDg agents.
VvrRvgVMNSRVGGroup	Use this template to create a Volume Replicator replication service group on a single node that uses non-shared storage.

The Templates box lists the templates available on the system to which Cluster Manager is connected. The resource dependency graph of the templates, the number of resources, and the resource types are also displayed.

- 8 Click **Next**. The wizard starts creating the service group.
- 9 After the service group is successfully created, click **Next** to edit attributes using the wizard.
- 10 The wizard lists the resources and their attributes. You must specify values for the mandatory attributes that appear in bold. The remaining resources listed in the window are preconfigured by the template and do not require editing.

To modify an attribute, do the following:

- Click the resource.
- Click the attribute to be modified.
- Click the **Edit** icon at the end of the table row.
- In the Edit Attribute dialog box, enter the attribute values.
- Click **OK**.

For details on application-specific agent attributes, refer to the application-specific agent or solutions guide.

For details on the storage and network agent attributes, refer to the *Cluster Server Bundled Agents Reference Guide*.

- 11 Click **Finish**.
- 12 Right-click the newly created service group and select **Enable Resources**.
- 13 Right-click the newly created service group, select **Online** from the context menu, and then select a system on which to bring the service group online.

If you are configuring the service group on a node at the secondary site in a DR environment, bring the service group online only after completing all the DR configuration steps.

Enabling fast failover for disk groups (optional)

For service groups that contain many disk groups, you can greatly reduce failover time by implementing the SFW fast failover feature for disk groups.

More information is available about fast failover benefits and requirements.

See [“Considerations for a fast failover configuration”](#) on page 68.

For implementing the fast failover feature, VCS provides a new attribute, `FastFailOver`, for the Volume Manager Diskgroup (VMDg) resource. This attribute determines whether or not a disk group is enabled for fast failover.

Note: The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click **Properties**. Disk group version upgrade is required after upgrading SFW HA on the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

You can enable fast failover for all the VMDg resources while configuring the service group using the configuration wizard. The service group configuration wizard provides a checkbox to enable fast failover.

Perform these steps if you did not enable fast failover using the wizard or if you have configured the service group manually.

The following procedure describes how to enable the `FastFailOver` attribute using the VCS Java Console.

To enable the FastFailover attribute for a VMDg resource

- 1 In Cluster Manager (Java Console), select a service group with a VMDg resource configured for it.

Select the Properties tab from the right pane.
- 2 Scroll down to choose the **FastFailOver** attribute and click to edit the attribute value.
- 3 In the Edit Attribute dialog box, check the **FastFailOver** check box and then click **OK**.
- 4 Repeat these steps for every VMDg resource for which you want to enable fast failover.

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 From 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 From 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. Depending on the configuration being deployed, there are additional steps that you must perform to set up and complete the configuration.

The following table 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 6-3 Additional Exchange Server configuration steps

Tasks	Refer to
Setting up a campus cluster configuration for Exchange Server	See “Tasks for configuring campus clusters” on page 129.
Setting up a replicated data cluster configuration for Exchange Server	See “Tasks for configuring Replicated Data Clusters for Exchange Server” on page 132.
Setting up a disaster recovery configuration for Exchange Server	See “Tasks for deploying a disaster recovery configuration of Microsoft Exchange” on page 178.
Configuring and running a fire drill for Exchange Server configuration	See “Tasks for configuring and running fire drills” on page 252.

Configuring campus clusters for Exchange Server

This chapter includes the following topics:

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

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 50.

The following table shows the steps specific to the campus cluster configuration that are done after configuring high availability on the nodes.

Table 7-1 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 130.

Table 7-1 Completing campus cluster configuration (*continued*)

Action	Description
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 130.

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 application 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 vmdgResourceName 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 includes the following topics:

- [Tasks for configuring Replicated Data Clusters for Exchange Server](#)
- [Creating the primary system zone for the application service group](#)
- [Creating a parallel environment in the secondary zone](#)
- [Setting up security for Volume Replicator](#)
- [Setting up the Replicated Data Sets \(RDS\)](#)
- [Configuring a RVG service group for replication](#)
- [Setting a dependency between the service groups](#)
- [Adding the nodes from the secondary zone to the RDC](#)
- [Verifying the RDC configuration](#)
- [Additional instructions for GCO disaster recovery](#)

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 52.

The following table shows the steps specific to the RDC configuration that are done after configuring high availability on the primary zone.

Table 8-1 Completing the configuration of a Replicated Data Cluster

Action	Description
Create the primary system zone and then verify failover within the primary zone	■ Create the primary system zone ■ Add the nodes to the primary zone See “Creating the primary system zone for the application service group” on page 134.
Create a parallel environment in the secondary zone	■ Install InfoScale Enterprise 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 135.
Add the secondary zone systems to the cluster	Add the secondary zone systems to the cluster
Set up security for Volume Replicator on all cluster nodes	Set up security for Volume Replicator on all nodes in both zones This step can be done at any time after installing InfoScale Enterprise on all cluster nodes, but must be done before configuring Volume Replicator replication. See “Setting up security for Volume Replicator” on page 136.
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 138.

Table 8-1 Completing the configuration of a Replicated Data Cluster
(continued)

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 RVG service group for replication” on page 150.
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 163.
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 163.
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 174.

Creating the primary system zone for the application service group

In the service group, set up systems in the primary zone (zone 0) to specify that initial failover 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 the **Properties** tab in the right pane.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.

- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone. Make sure you specify the systems in uppercase.

In case of a non-shared storage configuration, add only the single node to 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 RDC configuration”](#) on page 174.

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 (zone 1).

Before you begin to configure the secondary zone, do the following:

- Offline the following resources in the Exchange service group in the primary zone:
 - Exchange Server resource
 - Exchange virtual server name resource
 - Exchange virtual IP resourceThe remaining resources should be online, including the storage resources.
- In VEA, make sure to remove all the drive letters from the configured volumes, to avoid conflicts when configuring the zones.

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

-
- Installing the Veritas InfoScale product.
For more information, see the *Veritas InfoScale Installation and Upgrade Guide*.
- See [“Configuring disk groups and volumes for Exchange Server”](#) on page 66.
During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as the cluster at the primary zone:
 - Disk group name
 - Volume sizes

- Volume names
- Drive letters
- Installing Exchange Server on additional nodes
Follow the pre-installation, installation, and post-installation steps in the section on installing on additional nodes to 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.

- Adding a system to the existing cluster
You do not create another cluster in the secondary zone. Instead you add the systems to the existing cluster.

You do not create another Exchange Server service group in the secondary zone. You continue with the remaining Volume Replicator configuration tasks, during which the secondary zone nodes will be added to the Exchange Server service group.

Setting up security for Volume Replicator

If you use Volume Replicator for replication, you must configure the Veritas Volume Replicator Security Service (VxSAS) on all the cluster nodes.

For details on this required service, see the *Volume Replicator Administrator's Guide*.

After you install InfoScale Storage or InfoScale Enterprise, launch the Veritas Volume Replicator Security Service Configuration Wizard. This wizard lets you complete the Volume Replicator security service configuration.

Prerequisites for configuring VxSAS

- The wizard requires you to be logged on with administrative privileges.
- The account that 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.
- The systems on which you want to configure VxSAS must be accessible from the local system.

To configure VxSAS

- 1 Launch the Veritas Volume replicator Security Service Configuration Wizard from **Start > All Programs > Veritas > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

Optionally, run `vxsascfg.exe` from the command prompt of the required machine.

- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows:

Account name Enter the administrative account name.
 (domain\account)

Password Specify a password

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

Click **Next**.

- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that lets you specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

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 VxSAS.

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.

When configuring VxSAS in a firewall setup, the VxSAS wizard may not be able to configure the machines that are across the firewall, although the Host Selection dialog may list these nodes. In this case, configure VxSAS locally on the machines that are across the firewall.

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 intended Primary host is connected to VEA, if you are configuring the RDS from a remote client or from a host that is not the Primary.
- Verify whether the IP version preference is set before you configure replication.

If you specify host names when you configure replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP version Volume Replicator uses to resolve the host names.

Use one of the following methods to set the IP preference:

- Veritas Enterprise Administrator (VEA) GUI—select the appropriate options on the Control Panel > VVR Configuration > IP Settings tab.
- Run the `vxtune ip_mode [ipv4 | ipv6]` command at the primary site as well as the secondary site.
- Verify that the data volumes are not of the following types as Volume Replicator does not support these types of volumes:
 - Storage Foundation (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
For the Replicator Log volume, in addition to the above types, make sure that the volume does not have a DCM
 - Volumes names containing a comma
 - Secondary volume of a size smaller or greater than that on the Primary
- Verify that the disk group is imported and the volumes are mounted in the primary and secondary zone.
- Verify that you have configured security for Volume Replicator.
Verify that the VxSAS account has been configured with the same username and password for all the hosts, which are intended to be a part of the same RDS.
See [“Setting up security for Volume Replicator”](#) on page 136.

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.

Start VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator**.

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

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 appears:

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 *Storage Foundation Administrator's Guide*.

- Choose the **Select disks automatically** option if you want Volume Replicator 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, Volume Replicator 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 Volume Replicator 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 secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log. When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.
- 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** (default) 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 Volume Replicator updates in the Replicator Log. From there, Volume Replicator 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

- **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 Volume Replicator 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 Volume Replicator assigns a default name.

- If you want to specify advanced replication settings, click **Advanced**. Edit the replication settings for a secondary host as needed.



The image shows a Windows-style dialog box titled "Advanced Replication Settings". It contains several configuration options:

- Latency Protection:** A dropdown menu currently set to "Off".
- High Mark Value:** An empty text input field.
- Low Mark Value:** An empty text input field.
- Protocol:** A dropdown menu currently set to "UDP/IP".
- Packet Size(Bytes):** A dropdown menu currently set to "1400".
- Bandwidth:** A dropdown menu set to "Maximum Available" followed by a text input field and the unit "Mbps".
- Enable Compression:** An unchecked checkbox.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom.

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.

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, Volume Replicator 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.
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, Volume Replicator 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**.

- 13** On the Start Replication page, choose the appropriate option as follows:
- To add the Secondary and start replication immediately, select **Start Replication** with one of the following options:

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, Volume Replicator 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 and ReFS 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 *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.

14 Review the information.

Click **Back** to change any information you had specified.

Otherwise, click **Finish** to add the secondary host to the RDS and exit the wizard.

Configuring a RVG service group for replication

If you are setting up a RDC configuration, 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.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG service group.

For additional information about service group types, see the *Cluster Server Administrator's Guide*.

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

The following table shows the resources in the RVG service group for replication.

Table 8-2 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg (shared storage) or VMNSDg (non-shared storage) for the disk group	Volume Manager disk group with Exchange database files
VvrVrg for the system files 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.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG service group.

To create a 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.
- Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- Select **Hybrid**.
If you creating the RVG service group for a DR configuration in a non-shared storage environment, select **Failover**.
- Click **OK**.

Note: If you are setting up replication in a non-shared storage environment, you can use the replication service group template, **VvrRvgVMNSRVGGroup**, available in the Java Console. For an RDC configuration, ensure that you select the service group type as **Hybrid** while creating the service group using the Configuration Wizard from Java Console.

Configuring the resources in 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).

- [Modifying the DGGuid attribute for the new disk group resource in the RVG service group](#)
Copy the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) 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 Volume Replicator RVG resources for the disk groups](#)
Create the Volume Replicator RVG resources for all the disk groups and enter the attributes for each of the disk groups and the replication IP address.
- [Linking the Volume Replicator RVG resources to establish dependencies](#)
Link the Volume Replicator RVG resources to establish the dependencies between the VMDg or VMNSDg resources, the IP resource for replication, and the Volume Replicator RVG resources for the disk groups. Configure the RVG service group's VMDg or VMNSDg resources to point to the disk groups that contain the RVGs.
- [Deleting the VMDg or VMNSDg resource from the Exchange Server service group](#)
Delete the VMDg or VMNSDg 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 agents.

The following table shows the resource attributes to modify.

Table 8-3 IP and NIC resources

Resource	Attributes to modify
IP	Address
NIC	(none)

Note: In a non-shared storage environment, if you use the Java Console template, **VvrRvgVMNSRVGGroup**, to create the RVG service group, then do not recreate these resources; you modify the attributes of the existing IP and NIC resources in the service group.

To create the IP resource and NIC resource

- 1 In the **VCS Cluster Explorer** window, select the **Exchange Server service group (EVS1_GRP)** in the left pane.
- 2 On the **Resources** tab, right-click the **IP resource (EVS1_SG1-IP)**, and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the **RVG service group (EVS1_RVG_SG)**.
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the names of the IP and NIC resources for the RVG service group.
- 6 Click **OK**.

To modify the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (EVS1_RVG_SG-IP) and select **View > Properties View**.
- 2 In the **Properties View** window, for the **Address** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, enter the Volume Replicator IP address for the Primary Zone as the scalar value. This is the IP address you specified as the Primary side IP address while configuring the Replicated Data Set (RDS) earlier using the RDS wizard.
- 4 Close the **Properties View** window.

To enable the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (EVS1_RVG_SG-IP) and select **Enabled**.
- 2 In the **Resources** tab display area, right-click the NIC resource (EVS1_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg or VMNSDg resources for the disk groups

Configuration involves the following tasks:

- Create the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resource in the RVG service group by copying it from the Exchange Server service group and renaming it.
- Modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service group to ensure the desired failover behavior.

- Modify the attributes of the MountV resources in the Exchange Server service group for the new VMDg or VMNSDg in the RVG service group.
- Repeat these procedures for any additional VMDg or VMNSDg resources that you want to create for replication.
- If you are creating a DR configuration in a non-shared storage environment, modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service groups at both the sites (primary site and the disaster recovery site) separately.

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.

The following table shows the MountV resources and attributes to configure for the example configuration.

Table 8-4 MountV resources and attributes to modify

Resource	Attributes to modify
MountV (for the Exchange Server database volume)	VMDg Resource Name Volume Name
MountV (for the log volume)	VMDg Resource Name Volume Name
MountV (for the registry volume)	VMDg Resource Name Volume Name

To create the VMDg or VMNSDg resource in the RVG service group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EVS1_GRP) in the left pane.
- 2 On the Resources tab, right-click the VMDg or VMNSDg resource for the disk group that you want to configure for the RVG and click **Copy > Self**.
For example, right-click EVS1_SG-VMDg or EVS1_SG-VMNSDg.
- 3 In the left pane, select the RVG service group (EVS1_RVG_SG).
- 4 On the Resources tab, right-click in the blank resource display area and click **Paste**.

- 5 In the Name Clashes window, change the name of the VMDg or VMNSDg resource for the RVG service group.

For example change EVS1_SG-VMDg to EVS1_RVG_SG-VMDg or EVS1_RVG_SG-VMNSDg.

- 6 Click **OK**.

Note: Modify the DGGuid attribute of the new VMDg or VMNSDg resource before you perform this next procedure.

See [“Modifying the DGGuid attribute for the new disk group resource in the RVG service group”](#) on page 157.

To modify the MountV resources in the Exchange Server service group

- 1 In the VCS Cluster Explorer window, select the Exchange Server service group (EVS1_GRP) in the left pane.
- 2 In the Resources tab display area, right-click the MountV resource for the Exchange Server files (EVS1_SG1-MountV) and select **View > Properties View**.
- 3 In the Properties View window, verify that the **Volume Name** attribute is the Exchange Server database files (EVS1_SG1_DATA).
- 4 In the same Properties View window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the Edit Attribute window, modify the **VMDGResName** scalar value to be the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resource that was just created in the RVG service group.

For example, EVS1_RVG_SG-VMDg or EVS1_RVG_SG-VMNSDg.
- 6 Close the Properties View window.
- 7 Repeat these steps to modify the VMDGResName value for the additional MountV resources for the Exchange log volume and the Exchange registry replication volume.

To enable the VMDg or VMNSDg resource in the RVG service group

- 1 In the left pane, select the RVG service group (EVS1_RVG_SG).
- 2 In the Resources tab display area, right-click the VMDg or VMNSDg resource (EVS1_RVG_SG-VMDg or EVS1_RVG_SG-VMNSDg) and select **Enabled**.

Modifying the DGGuid attribute for the new disk group resource in the RVG service group

To modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service group

- 1 From the VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the Resources tab.
- 3 In the Resources tab display area, right-click the new VMDg or VMNSDg resource and click **View > Properties View**.
- 4 In the Properties View window, locate the DGGuid attribute and click the edit icon.
- 5 In the Edit Attribute window, edit the attribute by performing the following actions sequentially:
 - Select **Per System**.
 - From the dropdown list select the first node in the primary zone (Zone 0).
 - In the **Scalar Value** field specify the GUID of the disk group that is imported on the node.
Run the `VMGetDrive` utility at the command prompt to retrieve the GUID.
 - Repeat the previous two steps, and select a different node from the dropdown list each time. You must specify the GUID separately for each node displayed in the dropdown list.
In case of a shared storage environment (VMDg resource), if there are multiple nodes in the primary zone, then the disk group GUID will be the same for all systems within the zone. However, the GUID will always be different across zones.
- 6 In the Properties View window, verify that all nodes in the RDC primary zone have DGGuid values specified.

Note: If you are creating a DR configuration manually for a non-shared storage environment, you have to modify the DGGuid attribute of the VMNSDg resource in the RVG service groups at both the sites (primary site and the disaster recovery site) separately.

- 7 Close the Properties View window.

Adding the Volume Replicator RVG resources for the disk groups

Add Volume Replicator RVG resources (VvrRvg) for replication of the disk groups. If the application has multiple disk groups, create a separate VvrRvg resource for each disk group.

The following table lists the attributes that you must configure in the RVG service group for the VvrRvg resource.

Table 8-5 VvrRvg resource and attributes to modify

Resource	Attributes to Modify
VvrRvg	VMDgResName IPResName

To create the Volume Replicator RVG resource for a disk group containing the system files

- 1 In the left pane, select the RVG service group (EVS1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the Add Resource window, specify the following:
 - Enter a resource name for the Volume Replicator RVG resource. For example, enter EVS1_VvrRvg
 - In the Resource Type list, select **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, enter EVS1_RVG_SG.

The RVG name is the name you specified when you created the Replicated Data Set (RDS) earlier using the RDS wizard. You can retrieve the RVG name by running the command `vxprint -VPl`.
- 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 disk group containing the RVG.

For example, for the system files disk group, enter EVS1_RVG_SG-VMDg or EVS1_RVG_SG-VMNSDg.
- 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, enter EVS1_RVG_SG-IP.
- 11 Click **OK**.
- 12 In the Add Resource window, verify that the attributes have been modified:
- 13 Click **OK**.

Linking the Volume Replicator RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the Volume Replicator RVG service group to establish the dependencies between the resources.

You start from the top parent and link the parent and child resources as shown in the following table\.

Table 8-6 Dependencies for Volume Replicator RVG resources for RDC

Parent	Child
EVS1_ VvrRvg	The IP for replication, for example EVS1_RVG_SG-IP.
EVS1_ VvrRvg	The VMDg or VMNSDg for the Exchange files. For example EVS1_RVG_SG-VMDg or EVS1_RVG_SG-VMNSDg.

To link the Volume Replicator RVG resources

- 1 In the left pane, select the RVG service group (EVS1_RVG_SG).
- 2 Click the **Link** button in the right pane.
- 3 To link the VvrRvg resource to the IP resource, click the parent resource, for example EVS1_DB1_VvrRvg, and then click the child resource, for example EVS1_RVG_SG-IP.
- 4 When prompted to confirm, click **OK**.
- 5 To link the VvrRvg resource to the VMDg or VMNSDg resource, click the parent resource, for example EVS1_DB1_VvrRvg, and then click the child resource, for example EVS1_RVG_SG-VMDg or EVS1_RVG_SG-VMNSDg.

- 6
- When prompted to confirm, click **OK**.
- 7
- 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 or VMNSDg resource from the Exchange Server service group

The VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) 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 or VMNSDg resources from the Exchange service group

- 1
- In the VCS Cluster Explorer window, select the Exchange Server service group (EVS1_GRP) from the left pane.
- 2
- In the Resources tab display area, right-click the VMDg or VMNSDg resource for the first disk group (EXCH_SG1-VMDg or EXCH_SG1-VMNSDg) 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 or VMNSDg 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. This is the name of the VvrRvg resource in the RVG replication service group.

Configure the following attributes in the Exchange service group for the RVG Primary resources:

Table 8-7 RVG Primary resources

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

Creating the RVG Primary resources

For each disk group created for the Exchange Server, create a separate RVG Primary Resource for replication.

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

- 1 In the VCS Cluster Explorer window, right-click the Exchange Server service group (EVS1_GRP) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window, specify the following:
 - Enter a resource name for the RVG Primary resource for the Exchange Server system files disk group. For example, enter EVS1_RvgPrimary.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the Add Resource window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the Edit Attribute window, enter the name of the Volume Replicator RVG resource, for example INST1_VvrRvg and click **OK**. This is the name of the VvrRvg resource in the RVG replication service group.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Volume Replicator 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_GRP) to establish the dependencies between the resources for replication. Link each MountV resource to the appropriate RVGPrimary resource.

You start from the top parent and link the parent and child resources as shown in the following table.

Table 8-8 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_GRP).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource.
For example EXCH_SG-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_GRP) 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.
- 2 In the right pane on the Resources tab, right-click the first RVG Primary resource and select **Online > SYSTEM1**.
- 3 In the right pane on the Resources tab, right click the second RVG Primary resource and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG service group

In the RVG service group, set up systems in the primary zone (Zone 0) to specify that initial failover 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 the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the SystemZones attribute.

- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (type 0 for Zone 0) for the primary zone.

In case of a non-shared storage configuration, add only the single node to the primary zone.

- 7 Click **OK**.

Setting a dependency between the service groups

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

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

The Exchange service group (for example, EXCH_GRP) 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 application service group (parent service group).
For example, click the Exchange Server service group EXCH_SG.
- 5 Click the RVG service group (the child resource).
For example, click the RVG service group EXCH_RVG_SG.
- 6 In the **Link Service Groups** window, specify 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 nodes in the Secondary Zone (zone 1) can now be added to the RDC configuration.

See the following topics:

- See [“Configuring secondary zone nodes in the RVG service group”](#) on page 166.
- See [“Configuring the RVG service group IP resource for failover”](#) on page 167.

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.

Note: In case of a non-shared storage environment, perform this task manually using the Java Console. You cannot use the wizard to modify the RVG service group.

Use the following procedure if the RVG service group contains a VMDg resource (shared storage environment).

To add the nodes from the secondary zone to the RVG

- 1 From the active node of the cluster in the primary zone, launch the configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard**, or, on Windows 2012 operating systems, from the **Apps** menu.
- 2 Read and verify the requirements on the Welcome page, and click **Next**.
- 3 In the Wizard Options panel, specify 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 priority.

- To enable the service group to automatically come online on one of the systems, select the Include selected systems in the service group's AutoStartList attribute checkbox.
For information about the AutoStartList attribute, see the *Cluster Server Administrator's Guide*.
 - Click **Next**.
- 6** If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
 - 7** Review the Disk Group and Replicated Volume Group Configuration and click **Next**.
 - 8** In the IP Resource Options panel, select **Modify IP resource** and click **Next**.
 - 9** If a VCS error appears, click **OK**.
 - 10** In the Network Configuration panel, verify that the selected adapters are correct and click **Next**.
 - 11** Review the summary of the service group configuration as follows:
 - The Resources box lists the configured resources. Click a resource to view its attributes and their configured values in the Attributes box.
 - Click **Next** to modify the replication service group.
 - 12** When prompted, click **Yes** to modify the service group.
 - 13** Click **Finish**.

Use the following procedure if the RVG service group contains a VMNSDg resource (non-shared storage environment).

To add nodes from the secondary zone to the RVG service group using Java Console

- 1** From VCS Cluster Explorer, in the left pane, right-click the RVG service group and select **View > Properties View**.
- 2** In the Attributes window, click **Show all attributes**.
- 3** From the attributes list, select the attribute SystemList and click the edit icon.
- 4** In the Edit Attribute window, edit the SystemList attribute as follows:
 - Click the + button to add an empty row.
 - In the System field type the cluster node name from the secondary zone.
 - In the Priority field type **1**.

- Click **OK**.
- 5 Close the Attributes window.

Configuring secondary zone nodes in the RVG service group

Specify zone 1 for the nodes in the secondary zone.

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

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

In case of a non-shared storage configuration, add only the single node to the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Configuring the RVG service group NIC resource for fail over (VMNSDg only)


This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment where the RVG service group contains a VMNSDg resource.

Modify the MACAddress attribute of the NIC resource in the RVG service group to ensure desired fail over behavior in the RDC.

To modify the NIC resource in the RVG service group

- 1 From the VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the Resources tab.
- 3 In the Resources tab display area, right-click the NIC resource and click **View > Properties View**.
- 4 In the Properties View window, locate the MACAddress attribute and click the edit icon.

- 5 In the Edit Attribute window, edit the attribute by performing the following actions sequentially:



- Select **Per System**.
 - From the dropdown list, select the node in the RDC primary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - Click **OK**.
- 6 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
 - 7 Close the Properties View window.

Configuring the RVG service group IP resource for failover

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.

Use the following procedure to modify the IP resources.

Note: For IPv6 networks, modify the IPv6 resources.

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 and select **View > Properties View**.
- 4** In the Edit Attributes window, edit the Address attribute.
 - Select **Per System**.
 - Select the first node in the primary zone and enter the virtual IP address for the primary zone.
 - Select the second node in the primary zone and enter the virtual IP address for the primary zone (the same IP address as the first node).
 - Repeat for all nodes in the primary zone.
 - Select the first node in the secondary zone (SYSTEM3) and enter the virtual IP address for the secondary zone.
 - Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - Repeat for all nodes in the secondary zone.
 - Click **OK**.
- 5** In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone have the same IP address. The IP address at the primary zone and the secondary zone should be different
- 6** This step is applicable only if you are using a non-shared storage environment (VMNSDg agent).
 In the Edit Attributes window, edit the MACAddress attribute as follows:
 - Select **Per System**.
 - From the dropdown list, select the node in the RDC primary zone.
 - In the Scalar Value field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.

- In the Scalar Value field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - Click **OK**.
- 7 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
 - 8 Close the Properties View window.
- Since this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.


Configuring the RVG service group VMNSDg resources for fail over

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment where the RVG service group contains a VMNSDg resource. Modify the DGGuid attribute of the VMNSDg resources in the RVG service group to ensure the desired failover behavior in the RDC.

To modify the VMNSDg resources in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the Resources tab.
- 3 Right-click the RVG VMNSDg resource and select **View > Properties View**.
- 4 In the Edit Attributes window, edit the DGGuid attribute by performing the following actions sequentially:
 - Select **Per System**.
 - Select the node in the RDC secondary zone.
 - In the Scalar Value field, enter the GUID of the dynamic disk group that is imported on the single node in the RDC secondary zone.
 - You can retrieve the disk group details running the `VMGetDrive` utility from the command prompt.
 - Click **OK**.

- 5 In the Properties View window, verify that the DGGuid for the nodes in the primary and secondary zone are 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 2010 Database Configuration Wizard to add the nodes from the secondary zone (Zone 1) to the database service group.

Note: In case of a non-shared storage environment, perform this task manually using the Java Console. You cannot use the wizard to modify the application service group.

Use the following procedure if the service group contains a VMDg resource (shared storage environment).

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

- 1 Start the Exchange 2010 Database Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Exchange Server 2010 Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu.
- 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.
 - To enable all the VMDg resources in the service group for fast failover, select the **Enable FastFailOver attribute for all the VMDg resources in the service group** checkbox. For information about the FastFailOver attribute, see the Storage Foundation Administrator's Guide.
 - 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.

Use the following procedure if the service group contains a VMNSDg resource (non-shared storage environment).

To add nodes from the secondary zone to the application service group using Java Console

- 1 From VCS Cluster Explorer, in the left pane, right-click the Exchange service group (EXCH) and select **View > Properties View**.
- 2 In the Attributes window, click **Show all attributes**.
- 3 From the attributes list, select the attribute SystemList and click the edit icon.
- 4 In the Edit Attribute window, edit the SystemList attribute as follows:
 - Click the + button to add an empty row.

- In the System field type the cluster node name from the secondary zone.
 - In the Priority field type 1.
 - Click **OK**.
- 5** Close the Attributes windows.

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 Server database service group

- 1** From VCS Cluster Explorer, in the left pane, select the Exchange database service group (EXCH).
- 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.

In case of a non-shared storage configuration, add only the single node to the secondary zone.
- 8** Click **OK** and close the Attributes View window.

Configuring the application service group IP resource for fail over (VMNSDg only)

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment (VMNSDg agent).

Modify the IP resource in the application service group to ensure the desired failover behavior in the RDC.

Note: For IPv6 networks, modify the IPv6 resources.

To modify the IP resource in the application service group

- 1 From VCS Cluster Explorer, in the left pane, select the application service group.
- 2 In the right pane, select the Resources tab.
- 3 Right-click the IP resource and select **View > Properties View**.
- 4 In the Edit Attributes window, edit the MACAddress attribute by performing these actions sequentially:
 - Select **Per System**.
 - From the dropdown list, select the node in the RDC primary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - Click **OK**.
- 5 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
- 6 Close the Properties View window.

Configuring the application service group NIC resource for fail over (VMNSDg only)


This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment (VMNSDg agent).

Modify the MACAddress attribute of the NIC resource in the application service group to ensure desired fail over behavior in the RDC.

To modify the NIC resource in the application service group

- 1 From the VCS Cluster Explorer, in the left pane, select the application service group.
- 2 In the right pane, select the Resources tab.
- 3 In the Resources tab display area, right-click the NIC resource and click **View > Properties View**.

- 4 In the Properties View window, locate the MACAddress attribute and click the edit icon.
- 5 In the Edit Attribute window, edit the attribute by performing the following actions sequentially:



- Select **Per System**.
 - From the dropdown list, select the node in the RDC primary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - Click **OK**.
- 6 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
 - 7 Close the Properties 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:

- See [“Bringing the service group online”](#) on page 175.

- See “Switching online nodes” on page 175.

Bringing the service group online

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

To bring the application service group online

- 1 From VCS Cluster Explorer, in the left pane, right-click the application service group.
- 2 Click **Online**.

Switching online nodes

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

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

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.

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

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

To switch online nodes

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, and click the Service Groups tab.
- 2 Switch the service group as follows:
 - 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.

- 3 Verify that the service group is online on the node you selected.
- 4 To move all the resources back to the original node, repeat step 2 for each of the service groups.

Additional instructions for GCO disaster recovery

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

With this option, if a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away.

To configure disaster recovery using a secondary site, you must install the SFW HA Global Cluster Option on all nodes on the primary (replicated data cluster) site cluster, as well as the secondary (DR) site cluster. GCO configuration also requires a static IP address available for each site.

You can use the Disaster Recovery (DR) wizard when setting up the secondary site. The secondary site is not configured as a replicated data cluster. There can be only one replicated data cluster in the DR environment.

The DR wizard does the following tasks:

- Clones the storage
 - Clones the application service group
 - Sets up Volume Replicator replication for the secondary site
 - Configures the primary and secondary site clusters as global clusters
- When cloning the service group, the wizard does not clone the settings that specify primary and secondary zones, because the secondary site cluster is not divided into zones.

Deploying disaster recovery for Exchange Server

This chapter includes the following topics:

- [Tasks for deploying a disaster recovery configuration of Microsoft Exchange](#)
- [Tasks for setting up DR in a non-shared storage environment](#)
- [Reviewing the disaster recovery configuration](#)
- [Setting up the secondary site: Installing InfoScale Enterprise and configuring a cluster](#)
- [Verifying your primary site configuration](#)
- [Setting up your replication environment](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [About configuring disaster recovery with the DR wizard](#)
- [Configuring disaster recovery with the DR wizard](#)
- [Cloning the storage on the secondary site using the DR wizard \(Volume Replicator replication option\)](#)
- [Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)](#)
- [Installing Exchange 2010](#)

- [Cloning the service group configuration from the primary site to the secondary site](#)
- [Configuring the Exchange service group in a non-shared storage environment](#)
- [Configuring replication and global clustering](#)
- [Creating the replicated data sets \(RDS\) for Volume Replicator replication](#)
- [Creating the Volume Replicator RVG service group for replication](#)
- [Configuring the global cluster option for wide-area failover](#)
- [Verifying the disaster recovery configuration](#)
- [Establishing secure communication within the global cluster \(optional\)](#)
- [Adding multiple DR sites \(optional\)](#)
- [Recovery procedures for service group dependencies](#)
- [Possible task after creating the DR environment: Adding a new failover node to a Volume Replicator environment](#)

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 [“Tasks for configuring a new server for high availability”](#) on page 115.

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 52.

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 Volume Replicator 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. Veritas recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [“About the Solutions Configuration Center”](#) on page 103.

Note: If you are using non-shared storage (dynamic disk groups monitored using VMNSDg agent), you cannot use the DR wizard to configure disaster recovery. You have to set up DR manually. Refer to the following workflow table available for configuring DR manually.

The following table 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 9-1 Configuring the secondary site for disaster recovery

Action	Description
Install InfoScale Enterprise and configure the cluster on the secondary site	Install InfoScale Enterprise and configure the cluster at the secondary site. 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 InfoScale Enterprise and configuring a cluster” on page 186.
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 187.

Table 9-1 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
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 Volume Replicator” on page 136. See “Configuring EMC SRDF replication and global clustering” on page 218. See “Configuring Hitachi TrueCopy replication and global clustering” on page 221.
(Secure cluster only) Assign user privileges	For a secure cluster only, assign user privileges See “Assigning user privileges (secure clusters only)” on page 192.
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 195.
Clone the storage configuration (Volume Replicator replication only)	(Volume Replicator 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 (Volume Replicator replication option)” on page 198.
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 202.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “Installing Exchange 2010” on page 206.

Table 9-1 Configuring the secondary site for disaster recovery *(continued)*

Action	Description
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 from the primary site to the secondary site” on page 206.
Configure replication and global clustering, or configure global clustering only	■ (Volume Replicator 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 210.
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 234.
(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 236.
(Optional) Add additional DR sites	Optionally, add additional DR sites to a Volume Replicator environment See “Adding multiple DR sites (optional)” on page 238.
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 238.

Tasks for setting up DR in a non-shared storage environment

The following table outlines the high-level objectives and tasks for a creating a single-node DR configuration at the secondary site. Refer to this table if you are setting up DR in a non-shared storage environment (dynamic disk groups configured using VMNSDg agent).

You cannot use the DR wizard to configure disaster recovery in a non-shared storage environment. You have to configure DR manually.

Note: Some procedures (for example, configuring Volume Replicator replication) are common if you are setting up a DR or an RDC configuration. To avoid duplication, the topics referenced in this table point to the procedures described in the RDC chapter covered earlier.

Table 9-2 Non-shared storage: Configuring Disaster Recovery

Action	Description
Install SFW HA and configure the cluster on the secondary site	■ Verify the software and hardware prerequisites ■ Set up the network and storage ■ Install the product ■ Configure the cluster at the secondary site Ensure that you select the GCO option to configure the Global Cluster Option resource for the cluster. Caution: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster. ■ Configure disk groups and volumes See “Setting up the secondary site: Installing InfoScale Enterprise and configuring a cluster” on page 186.
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 187.

Table 9-2 Non-shared storage: Configuring Disaster Recovery (*continued*)

Action	Description
Set up the replication prerequisites	■ Ensure that Volume Replicator replication prerequisites are met ■ Configure the VxSAS service for Volume Replicator, specifying the cluster nodes at both primary and secondary sites See “Setting up security for Volume Replicator” on page 136.
(Secure cluster only) Assign user privileges	For a secure cluster only, assign user privileges. See “Assigning user privileges (secure clusters only)” on page 192.
Install and configure Exchange Server	■ Follow the guidelines for installing Exchange Server in the SFW HA environment See “Installing Exchange 2010” on page 206.
Configure the Exchange service group for VCS (secondary site)	■ Configure the application service group manually using the Cluster Manager (Java Console) ■ Ensure that the name of the service group is the same as that on the primary site See “Configuring the Exchange service group in a non-shared storage environment” on page 210.
Set up the replicated data sets (RDS) for Volume Replicator replication	■ Use the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites ■ Use the Setup Replicated Data Set Wizard to create Replicator Log volumes for the primary and secondary sites See “Creating the replicated data sets (RDS) for Volume Replicator replication” on page 228.


Table 9-2 Non-shared storage: Configuring Disaster Recovery (*continued*)

Action	Description
Create the Volume Replicator RVG service group (repeat steps on primary and secondary site separately)	■ Create the Volume Replicator RVG service group for the replicated volume group ■ Use the Cluster Manager (Java Console) to manually create the service group ■ Create the RVG service group at the primary site and the secondary site separately ■ Bring the RVG service group online on the primary site See “Creating the Volume Replicator RVG service group for replication” on page 228.
Configure the global cluster option for wide-area failover	■ Link clusters (adding a remote cluster to a local cluster) ■ Converting the application service group that is common to all the clusters to a global service group ■ Converting the local service group to a global group ■ Bringing the global service group online See “Configuring the global cluster option for wide-area failover” on page 229.
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 234.
(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 236.
(Optional) Add additional DR sites	Optionally, add additional DR sites to a Volume Replicator environment See “Adding multiple DR sites (optional)” on page 238.
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 238.

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. The following figure 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 9-1 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. The following figure illustrates this type of failure:

Figure 9-2 Application services restored after primary site failure



You can choose to configure replication using Volume Replicator or an agent-supported array-based hardware replication. You can use the DR wizard to configure Volume Replicator 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 the *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.

In a hardware replication environment, the Disaster Recovery wizard supports one level of dependency (one child). If you need to configure more levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

In a Volume Replicator environment, the wizard cannot configure DR for a service group that has a child and you will need to configure the secondary site manually. For more information on configuring Volume Replicator, see the *Volume Replicator Administrator's Guide*. For more information on configuring GCO, see the *Cluster Server Administrator's Guide*.

Setting up the secondary site: Installing InfoScale Enterprise and configuring a cluster

After completing the HA configuration on the primary site, repeat the appropriate tasks to complete the product installation and configure the cluster at the secondary site.

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 64.

- Ensure that when installing InfoScale Enterprise 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:

Global Cluster Option	Required for a disaster recovery configuration
Volume Replicator	Required if you plan to use Volume Replicator for replication

For more information see the *Veritas InfoScale 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 using the Cluster Configuration Wizard”](#) on page 86.

Note: You do not need to configure the GCO option while configuring the cluster. This is done later using the Disaster Recovery wizard. However, in case you are configuring DR manually in a non-shared storage environment, you must configure GCO using VCW.

- The storage configuration will be handled by the DR wizard (in case of shared storage only). If you are configuring a DR setup that uses non-shared storage (VMNSDg agent), you have to configure the storage using VEA console. While creating disk groups and volumes for the secondary site, make sure to use the same names of volumes as those on the primary site. The size of the volumes on the secondary site must be equal to or larger than the size of the volumes on the primary site.

Ensure that you allow sufficient disk space to create a volume for the Volume Replicator Storage Replicator Log for each storage group. You can create the volume now, or later, when running the wizard to create replicated data sets (RDS).

See [“Configuring disk groups and volumes for Exchange Server”](#) on page 66.

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.

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Volume Replicator
- EMC SRDF
- Hitachi TrueCopy

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

For array-based hardware replication, you can use any replication agent supported by 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 225.

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:

- Volume Replicator
See [“Setting up security for Volume Replicator”](#) on page 136.
- EMC SRDF
See [“Requirements for EMC SRDF array-based hardware replication”](#) on page 189.
- Hitachi TrueCopy
See [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 190.

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 *Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide*.

Before using the DR wizard, review the following topics:

- See [“Software requirements for configuring EMC SRDF”](#) on page 189.
- See [“Replication requirements for EMC SRDF”](#) on page 189.

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.

Note: In addition, the agent requires that the 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 *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*.

Before using the DR wizard, review the following topics:

- See [“Software requirements for Hitachi TrueCopy”](#) on page 191.
- See [“Replication requirements for Hitachi TrueCopy”](#) on page 191.

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

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:

`systemDriver\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 *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
```

About configuring disaster recovery with the DR wizard

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (Volume Replicator replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure Volume Replicator 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:

- InfoScale Enterprise 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.
- 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.
- For IPv4 networks, static IP addresses are available to enter for the following (for IPv6, they are generated during configuration):
 - One static IP address at each site for configuring GCO.
 - If using Volume Replicator for replication, a minimum of one static IP address per site for each application instance running in the cluster.
- The service group to be cloned can use either IPv4 IP addresses or IPv6 addresses but not a mixture of both.
- To configure IPv6 settings, the wizard must be launched from a system on which the IPv6 stack is installed.

- For Volume Replicator replication, the service group to be cloned cannot contain a child service group.
- 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 Volume Replicator configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- See [“Setting up security for Volume Replicator”](#) on page 136.
- See [“Requirements for EMC SRDF array-based hardware replication”](#) on page 189.
- See [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 190.

Configuring disaster recovery with the DR wizard

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

This procedure describes how to configure disaster recovery using the wizard.

To start configuring disaster recovery with the DR wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**, or, on Windows 2012 operating systems, from the **Apps** menu.
- 2 Expand the Solutions for Microsoft Exchange Server tab.
Click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.

- 4** In the System Selection panel, provide information in the **System Name** field:

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 **localhost** to connect to the system.

Click **Next**.

- 5** In the Service Group Selection panel, select the service group that you want to clone to the secondary site.

For a hardware replication environment, 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 Volume Replicator environment, the DR wizard does not support configuring DR for a service group that has a child. If you select a service group that has a child, you will receive an error message when you select the Volume Replicator replication method later in the wizard.

The panel lists only service groups that contain a MountV resource.

Click **Next**.

- 6** 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**.

- 7 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 Volume Replicator (Volume Replicator) and the Global Cluster Option (GCO)	Select this option if you want to configure Volume Replicator replication. Select this option even if you plan to configure Volume Replicator replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a Volume Replicator 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 Volume Replicator option, the wizard will warn you that you cannot use Volume Replicator 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.

Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)

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 Volume Replicator replication environment.

Click **Next**.

8 Continue with the next DR configuration task.

For Volume Replicator replication:

See [“Cloning the storage on the secondary site using the DR wizard \(Volume Replicator replication option\)”](#) on page 198.

For array-based replication:

See [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 202.

Cloning the storage on the secondary site using the DR wizard (Volume Replicator 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.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you have not yet started the wizard, refer to the following topic before continuing with the storage cloning procedure:

To clone the storage configuration from the primary site to the secondary site (Volume Replicator 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 Volume Replicator replication method and click **Next**.
- 2
- Review the information in the Storage Validation Results panel. This panel compares the configuration at the secondary site with that on the primary. If the storage is already configured identically on both sites, the panel shows that results are identical. Otherwise, the panel shows the differences and recommended actions. You can toggle between a summary and detailed view of information about the differences.

The detailed view shows the following:

Disk Group	Displays the disk group name that needs to be created on the secondary site.
Volume	Displays the list of volumes, if necessary, that need to be created at the secondary site.
Size	Displays the size of the volume that needs to be created on the secondary site.
Mount	Displays the mount to be assigned the volume on the secondary site.
Recommended Action	Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary. ■ 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 >> button 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 >> button to move the disks into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group. Select disks for each unavailable volume that you want to clone on to the secondary.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been moved in from the Available Disks pane.
View Primary Layout	Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 5 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.

Note: On the VEA GUI of the secondary site, a Windows dialog box might appear prompting you to format a disk. Click **Cancel** to close the dialog.

The appearance of this dialog box has no impact on the operations being performed by the DR wizard. You can safely ignore it.

- 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, 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, 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: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you have already installed the application on all nodes, you can skip this storage cloning step by unchecking the Perform storage cloning check box on the Storage Cloning panel.

If you are starting the wizard for the first time, refer to the following topic before continuing with the storage cloning procedure:

See [“About configuring disaster recovery with the DR wizard”](#) on page 193.

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 The default path is <code>C:\HORCM\etc</code> , where <code>C</code> is the system drive.
HORCM files location	Path to the horcm configuration files (horcmnn.conf) The default path is: <code>C:\Windows</code> , where <code>C</code> is the system drive. The <code>horcm</code> configuration file is required by the RAID Manager on all nodes; however, the wizard does not validate its presence.

- 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 2010 does not require the temporary storage since it does not require volume mount points for installation.

Choose one of the following:

- 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**.

Note: If SCSI-3 support is enabled for using Persistent Group Reservations (PGR), and if one of the selected disks is not SCSI-3 compliant, the following error is displayed: "Unable to reserve a majority of dynamic disk group members. Failed to start SCSI reservation thread."

Recommended action: Click **Finish** to exit the wizard. Either replace the non-compliant disk with a SCSI-3 compliant disk, or enable SCSI-2 support, and then run the wizard again.

- 10 In the Storage Configuration Cloning Result screen, view the results and click **Next**.
- 11 In the Exchange Installation panel, review the information and once the 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 [“Installing Exchange Server 2010”](#) on page 111.

Cloning the service group configuration from the primary site to the secondary site

Before cloning a service group on the secondary site, verify if the application is installed on the secondary site.

Ensure that the SQL Server Full-Text Search service on the secondary site is configured to start in the manual mode and is initially in the stopped state.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you are launching the wizard for the first time, refer to the following topic for additional information:

See [“About configuring disaster recovery with the DR wizard”](#) on page 193.

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 by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.

Expand the Solutions for Microsoft Exchange Server tab.

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 Volume Replicator 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**.

- 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.
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.

- 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	For IPv6, select the network from the dropdown list. If you select the same subnet as the primary site, the primary site IP address will be used. Otherwise the IP address will be generated from the network. For the MACAddress attribute select the appropriate public NIC from the drop-down list. For IPv6 available NICs are those belonging to the selected IPv6 network.

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 the Exchange service group in a non-shared storage environment

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to clone the application service group created at the primary site if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure the service group manually using the Cluster Manager (Java Console).

Note the following before configuring the service group at the secondary site:

- Ensure that the application agent resources, the Lanman resource (if configured), and the IP resource is offline in the service group on the primary site. The remaining resources, including the storage resources, must be online.
- Ensure that the name of the service group is the same as that on the primary site.
- After configuring the service group do not bring it online on the secondary site at this time. You can bring it online later after completing all the DR configuration steps.

See [“Configuring the service group in a non-shared storage environment”](#) on page 124.

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 Volume Replicator or an agent-supported array-based hardware replication.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

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:

- See [“Configuring Volume Replicator replication and global clustering”](#) on page 211.
- See [“Configuring EMC SRDF replication and global clustering”](#) on page 218.

- See “[Configuring Hitachi TrueCopy replication and global clustering](#)” on page 221.
- See “[Configuring global clustering only](#)” on page 225.

Configuring Volume Replicator replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure Volume Replicator replication and global clustering.

Before you begin, ensure that you have met the following prerequisites:

- Ensure that Volume Replicator Security Service (VxSAS) is configured at the primary and secondary site.
See “[Setting up security for Volume Replicator](#)” on page 136.
- Verify whether the IP version preference is set before you configure replication. If you specify host names when you configure replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP version Volume Replicator uses to resolve the host names.

Use one of the following methods to set the IP preference:

- Veritas Enterprise Administrator (VEA) GUI—select the appropriate options on the Control Panel > VVR Configuration > IP Settings tab.
- Run the `vxtune ip_mode [ipv4 | ipv6]` command at the primary site as well as the secondary site.
- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that, for remote cluster configuration, you configure a VCS user with the same name and privileges in each cluster.

Use the following procedure to configure Volume Replicator replication and global clustering with the DR wizard.

To configure Volume Replicator 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 DR 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:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.

- Expand the Solutions for Microsoft Exchange Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - 4 On the Replication Methods panel, click **Configure Volume Replicator and the Global Cluster Option (GCO)**. Click **Next**.
 - 5 In the Internet Protocol panel, select IPv4 or IPv6 depending on which type of network you are using. (You must use the same on primary and secondary sites.) Click **Next**.
 - 6 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.
 - 7 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 > button 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 < button 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.

Advanced Replication Settings

Advanced Replication Settings for RVG_TESTFS_0

Replication Mode:

Synchronous Override

Log Protection:

AutoDCM

Primary RLINK Name:

48326630361624

Secondary RLINK Name:

48326630361623

Bandwidth:

Maximum

Mbps

Protocol:

UDP

Packet Size (Bytes):

1400

Latency Protection:

Fail

High Mark Value:

10000

Low Mark Value:

9950

Initial Synchronization:

Auto Synchronous

OK

Cancel

The options on the dialog box are described column-wise, from left to right:

- Replication Mode

Select the required mode of replication; **Synchronous**, **Asynchronous**, or **Synchronous Override** (default).

Log Protection	Select the appropriate log protection from the list: ■ 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, Volume Replicator replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit. The default unit is Megabits 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, Volume Replicator 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**.

For additional information on Volume Replicator replication options, refer to the *Volume Replicator Administrator's Guide*.

Click **Next**.

- 8 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	For IPv4 networks, enter replication IPs that will be used for replication, one for the primary site and another for the secondary site. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
Public NIC	Select the public NIC from the drop-down list for the system at the primary and secondary site. For IPv6, available NICs are those belonging to the selected network.
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**.

- 9** 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	For IPv4, enter a virtual IP for the WAC resource. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
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.

- 10** 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.

Otherwise, click **Next** to implement the settings.

- 11 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**.
- 12 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 one static address is available per site for configuring GCO.

See [“Requirements for EMC SRDF array-based hardware replication”](#) on page 189.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings.

See [“Optional settings for EMC SRDF”](#) on page 221.

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 by following these steps sequentially:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
- Expand the Solutions for Microsoft Exchange Server tab. 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.

Otherwise, 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 (x) 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 *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 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.

See [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 190.

Ensure that you have one static address is available per site for configuring GCO.

The wizard configures the required agent settings. It uses defaults for optional settings.

See [“Optional settings for HTC”](#) on page 225.

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 by following these steps sequentially:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
Expand the Solutions for Microsoft Exchange Server tab.
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.
- 4 In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 5 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.

- 6** 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 <code>horcm</code> file.

- 7** If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.

- 8 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.

- 9 In the Settings Summary panel, review the displayed information.

If you want to change any of the parameters specified for the replication resource settings or the global cluster settings, click **Back**.

Otherwise, click **Next**.

- 10 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**.

- 11 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.
- 12 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 *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 you have one static address is 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 by following these steps sequentially:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
- Expand the Solutions for Microsoft Exchange Server tab.
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	For IPv4, enter a virtual IP for the WAC resource. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
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.
- If you want to change any of the parameters specified, click **Back**.
- Otherwise, 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.

Creating the replicated data sets (RDS) for Volume Replicator replication

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to configure Volume Replicator replication if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure replication using the Setup Replicated Data Set Wizard.

Configuring Volume Replicator involves setting up the replicated data sets (RDS) on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

See [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 138.

Creating the Volume Replicator RVG service group for replication

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to configure Volume Replicator replication if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure the replication service group manually using the Cluster Manager (Java Console).

Complete the following procedures first on the node in the primary site. Then repeat all the steps on the node in the secondary site. You must follow the order of the procedures as mentioned.

Refer to the following topics:

- See [“Configuring a RVG service group for replication”](#) on page 150.

- See [“Creating the RVG service group”](#) on page 151.
- See [“Configuring the IP and NIC resources”](#) on page 153.
- See [“Configuring the VMDg or VMNSDg resources for the disk groups”](#) on page 154.
- See [“Adding the Volume Replicator RVG resources for the disk groups”](#) on page 158.
- See [“Linking the Volume Replicator RVG resources to establish dependencies”](#) on page 159.
- See [“Deleting the VMDg or VMNSDg resource from the Exchange Server service group”](#) on page 160.
- See [“Configuring the RVG Primary resources”](#) on page 160.
- See [“Creating the RVG Primary resources”](#) on page 161.
- See [“Linking the RVG Primary resources to establish dependencies”](#) on page 161.
- See [“Bringing the RVG Primary resources online”](#) on page 162.
- See [“Setting a dependency between the service groups”](#) on page 163.

Configuring the global cluster option for wide-area failover

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

The Global Cluster option is required to manage global clustering for wide-area disaster recovery. The process of creating a global cluster environment involves the following tasks:

- Connecting standalone clusters by adding a remote cluster to a local cluster
- Converting the local service group that is common to all the clusters to a global service group

Use the VCS Java Console and perform the following global cluster operations:

- See [“Linking clusters: Adding a remote cluster to a local cluster”](#) on page 230.
- See [“Converting a local service group to a global service group”](#) on page 231.
- See [“Bringing a global service group online”](#) on page 233.

Linking clusters: Adding a remote cluster to a local cluster

This is applicable only if you are setting up DR manually in a non-shared storage environment.

The VCS Cluster Manager (Java Console) provides a wizard to create global clusters by linking standalone clusters or bringing a standalone cluster into an existing global environment.

Note the following uses of the wizard:

- If you are creating a global cluster environment for the first time with two standalone clusters, run the wizard from either the cluster on the primary site or the cluster on the secondary site.
- If you are adding a standalone cluster to an existing global cluster environment, run the wizard from a cluster already in the global cluster environment.

The following information is required for the Remote Cluster Configuration Wizard in VCS Cluster Manager:

- The active host name or IP address of each cluster in the global configuration and of the cluster being added to the configuration.
- The user name and password of the administrator for each cluster in the configuration.
- The user name and password of the administrator for the cluster being added to the configuration.

Veritas InfoScale products do not support adding a cluster that is already part of a global cluster environment. To merge the clusters of one global cluster environment (for example, cluster A and cluster B) with the clusters of another global environment (for example, cluster C and cluster D), separate cluster C and cluster D into standalone clusters and add them one by one to the environment containing cluster A and cluster B.

To add a remote cluster in Cluster Explorer

- 1 From Cluster Explorer, click **Edit > Add/Delete Remote Cluster**.

or

From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.

- 2 Review the required information for the Remote Cluster Configuration Wizard and then click **Next**.
- 3 In the Wizard Options panel, click **Add Cluster**, then click **Next**.
- 4 In the New Cluster Details panel, enter the details of the new cluster.

If the cluster is not running in secure mode, specify the following:

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- If necessary, change the default port number.
- Enter the user name and the password.
- Click **Next**.

If the cluster is running in secure mode, specify the following:

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- Verify the port number.
- Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.
 If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
- Click **Next**.

5 Click **Finish**.

After running the wizard, the configurations on all the relevant clusters are in read-write mode; the wizard does not close the configurations.

6 Verify that the heartbeat connection between clusters is alive by entering `hahb -display` in the command window.

The state attribute in the output should show "alive". If the state is unknown, then take the ClusterService group offline and bring it online again.

Converting a local service group to a global service group

This is applicable only if you are setting up DR manually in a non-shared storage environment.

To convert a local service group to a global group

- 1 From Cluster Explorer, click **Edit > Configure Global Groups**.
 or
 From the Cluster Explorer configuration tree, right-click the cluster, and click **Configure Global Groups**.
 or
 From the Cluster Explorer configuration tree, right-click the service group, click **Configure As Global**, and proceed to step 3.
- 2 Review the information required for the Global Group Configuration wizard and click **Next**.
- 3 Enter the details of the service group to modify, as follows:
 - Click the name of the service group that will be converted from a local group to a global group, or vice versa.
 - From the Available Clusters box, click the clusters on which the group can come online. Click the right arrow to move the cluster name to the Clusters for Service Group box; for global to local cluster conversion, click the left arrow to move the cluster name back to the Available Clusters box. A priority number (starting with 0) indicates the cluster on which the group will attempt to come online. If necessary, double-click the entry in the Priority column and enter the new value.
 - Select the policy for cluster failover as follows:

Manual	Prevents a group from automatically failing over to another cluster.
Auto	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster fails.
Connected	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster.
 - Click **Next**.
- 4 Enter or review the connection details for each cluster. Click the Configure icon to review the remote cluster information for each cluster, as follows:

Cluster not in
secure mode

Follow these steps sequentially:

- Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
- Verify the port number.
- Enter the user name.
- Enter the password.
- Click **OK**.

Repeat these steps for each cluster in the global environment.

Cluster in secure
mode

Follow these steps sequentially:

- Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
- Verify the port number.
- Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and domain.
If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
- Click **OK**.

Repeat these steps for each cluster in the global environment.

5 Click **Next**, then click **Finish**.

At this point, you must bring the global service group online from Cluster Explorer.

Bringing a global service group online

This is applicable only if you are setting up DR manually in a non-shared storage environment.

To bring a remote global service group online from Cluster Explorer

- 1 In the Service Groups tab of the configuration tree, right-click the service group.
or

Click a cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.

- 2 Click **Online**, and click **Remote online**.
- 3 In the Online global group dialog box, specify the following:
 - Click the remote cluster to bring the group online.
 - Click the specific system, or click **Any System**, to bring the group online.

- Click **OK**.

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 Volume Replicator 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 Volume Replicator replication:
 - Ensure Volume Replicator 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 Volume Replicator 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 *Cluster Server Hardware Replication Agent for Hitachi TrueCopy 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 Volume Replicator for replication and configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you are using Volume Replicator 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 process typically consists of the following tasks:

- Starting a Volume Replicator replication checkpoint
 - Performing a block level backup
 - Ending the Volume Replicator replication checkpoint
 - Restoring the block level backup at the DR site
 - Starting replication from the Volume Replicator replication checkpoint
- To learn more about the process of starting replication from a checkpoint, refer to the *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 Volume Replicator-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 all the previous steps 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 > 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 the previous step 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 all the previous steps 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 rootHashInHex]
```

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 use the following commands:

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 all the previous steps for the additional clusters in the global cluster.

Adding multiple DR sites (optional)

In a Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. 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 Volume Replicator 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 186.

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 Volume Replicator 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 9-3 Online, local, soft dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	■ The parent remains online on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online.	■ Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
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.	■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, firm dependency link.

Possible task after creating the DR environment: Adding a new failover node to a Volume Replicator environment**Table 9-4** Online, local, firm dependency link

Failure condition	Result	Action required (sequentially)
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.	■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ 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 9-5 Online, local, hard dependency link

Failure condition	Result	Action required (sequentially)
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.
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.	■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Possible task after creating the DR environment: Adding a new failover node to a Volume Replicator 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](#)
- [Preparing the existing DR environment](#)
- [Installing Exchange on the new node](#)
- [Modifying the replication and Exchange service groups](#)
- [Reversing replication direction](#)

Preparing the new node

Install InfoScale Enterprise on the new system and then add the system to the cluster.

To install InfoScale Enterprise and add the system to the cluster

- 1 See [“About installing the Veritas InfoScale products”](#) on page 58.
- 2 Use the Cluster Operations option of the VCS Cluster Configuration wizard to add the new system to the cluster. If necessary, refer to the *Cluster Server Administrator's Guide* for information on this procedure.

Preparing the existing DR environment

If you plan to add a failover node to the secondary site, you must temporarily switch the roles of the primary and secondary sites so that the current site becomes the primary. This action reverses the direction of replication.

To prepare the existing DR environment

- 1 If you adding the failover node to the cluster at the primary site, proceed directly to step 2. If you are adding a failover node to the secondary site, you must switch the roles of the primary and secondary sites. This action reverses the direction of replication.
 - In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
 - Click **Switch To**, and click **Remote switch**.
 - In the Switch global group dialog box:
 - Click the cluster at the secondary site you want to switch the group to.
 - Click the specific system where you want to bring the global Exchange service group online.

- Click **OK**.
- 2 Take the global Exchange service group offline at the current primary site.
 - 3 Take the Volume Replicator 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.
See [“Setting up the secondary site: Installing InfoScale Enterprise and configuring a cluster”](#) on page 186.
- 2 From the VEA navigation tree, right-click the RVG for the primary site, and click **Enable Data Access**.
- 3 Install Exchange.
See [“Installing Exchange Server 2010”](#) on page 111.

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 to add the new node to the system list for the replication service group. If necessary, refer to the *Volume Replicator Administrator's Guide* for information on this procedure.
- 4 Use the Modify service group option of the 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 *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](#) 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

This chapter includes the following topics:

- [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](#)
- [Re-creating a fire drill configuration that has changed](#)
- [Restoring the fire drill system to a prepared state](#)
- [Deleting the fire drill configuration](#)
- [Considerations for switching over fire drill service groups](#)

About disaster recovery fire drills

A disaster recovery (DR) plan should include regular testing of an environment to ensure that a DR solution is effective and ready if a disaster strikes. 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

Storage Foundation and High Availability Solutions (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 Volume Replicator (Volume Replicator) or that uses Hitachi TrueCopy or EMC SRDF hardware replication.

In the Hitachi TrueCopy or EMC SRDF environments, the Fire Drill Wizard supports only the Gold configuration. For the Silver or Bronze configuration, you must manage (create, restore, delete) the fire drill configurations and run the fire drills manually. For further information about the Gold, Silver, and Bronze configurations, refer to the following documents:

Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide

Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide

Note: After upgrading to 6.0.1 or later, the existing fire drill service groups will not be usable. In a Hitachi TrueCopy or EMC SRDF environment, you must manually edit the existing fire drill service groups. In a Volume Replicator environment, you must use the Fire Drill Wizard to re-create them. For more information, see the *Veritas InfoScale Installation and Upgrade Guide*.

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 VVRSnap and mountV resources, and in the case of hardware replication,

the HTCSnap or SRDFSnap resource. The wizard renames the fire drill service group resources with a prefix FDnn and changes attribute values as necessary to refer to the FD resources.

The wizard also supports fire drill service groups created under a different naming convention by an earlier version of the wizard.

- Runs the fire drill by bringing the fire drill service group online on the secondary site

Optionally the wizard runs Eseutil to check for data consistency as part of the fire drill test. The fire drill tests the replication and consistency of the data to verify that the data will be available if the Exchange service group fails over and comes online at the secondary site should the need arise.

Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online on the primary site.

Note: The fire drill service group for Exchange contains only the data resources, not the application, so that the Exchange application does not itself come online during a fire drill.

- Restores the fire drill configuration, taking the fire drill service group offline
After you complete a fire drill, 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, after completing the fire drill testing for a service group, always use the wizard to restore the fire drill configuration to a prepared state as soon as possible.

See [“Restoring the fire drill system to a prepared state”](#) on page 269.

- 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 Volume Replicator environment”](#) on page 247.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 249.

About Fire Drill Wizard operations in a Volume Replicator 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 245.

However, the following additional Fire Drill Wizard operations are specific to a Volume Replicator environment.

Preparing the fire drill configuration

In a Volume Replicator environment, when preparing the fire drill configuration, the wizard does the following:

- Replaces the RVGPrimary resources with VVRSnap resources in the fire drill service group
- Uses the SFW HA VxSnap feature to prepare snapshot mirrors for use during the fire drill
While running the wizard, you assign one or more disks for the mirrored volumes. Mirror preparation can take some time, so you can exit the wizard after this step is started and let the preparation continue in the background.
- Sets the `offline-local-firm` dependency between the service groups, where the fire drill service group is the parent and the application service group is the child
- Configures the VVRSnap resource by setting the following attributes to the appropriate values:
 - RVG
 - AppDiskGroupName
 - DiskGroupName
- Sets the FireDrill attribute of the following resources to true:
 - IP

- Lanman
- RegRep
- Exch2010DB
- Sets the ForFireDrill attribute of the following resources to `true` in the fire drill service group:
 - MountV
 - VMDg

This indicates that the volume being monitored by the VVRSnap agent belongs to the fire drill disk group.

About running the fire drill

The Fire Drill Wizard brings the fire drill service group online. Optionally, you can also run the fire drill using the Veritas Operations Manager console.

In a Volume Replicator environment, when running the fire drill, the VVRSnap agent 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

About restoring the fire drill configuration

The Fire Drill Wizard takes the fire drill service group offline. Optionally, you can also restore the fire drill using the Veritas Operations Manager console.

In a Volume Replicator environment, restoring the fire drill system to a prepared state, the VVRSnap agent does the following:

- Joins the fire drill disk group to the application service group disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

About deleting the fire drill configuration

In a Volume Replicator environment, when deleting the fire drill configuration, the wizard does the following:

- Sets the FireDrill attribute of the following resources to `false`:
 - IP
 - Lanman

- RegRep
- Exch2010DB
- Unlinks the fire drill service group
- 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 Fire Drill Wizard performs the following basic operations in all replication environments:

- Prepares for the fire drill by creating a fire drill service group on the secondary site
- Runs the fire drill by bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration by taking the fire drill service group offline
- Deletes the fire drill service group and any associated registry entries

See [“About the Fire Drill Wizard”](#) on page 245.

In Hitachi TrueCopy or EMC SRDF replication environments, the Fire Drill 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.

About preparing the fire drill configuration

When preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, the wizard creates HTCSnap or SRDFSnap resources for each HTC and SRDF resource in the application service group. It links the fire drill service group to the corresponding application service group.
- In an HTC or SRDF environment, the wizard configures the Snap resource and sets the following attributes to the value 1, which indicates:
 - **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.

- In an EMC SRDF environment, the wizard sets the following attribute values:
 - It sets **CopyMode** to one of the following, which indicates:
 - **Mirror**: Use the TimeFinder Mirror technology to create snapshots.
 - **Clone**: Use the TimeFinder Clone technology to create snapshots.
 - **Snap**: Use the TimeFinder Snap technology to create snapshots.
 - When the TimeFinder Clone technology is used, it sets **UseTgt** to one of the following, which indicates:
 - **0**: Use BCV devices to create snapshots.
 - **1**: Use STD devices to create snapshots.
 - When the TimeFinder Snap technology is used, if a custom save pool area name is specified, it sets **SavePoolName** accordingly. The specified save pool area is used to create snapshots.
If no value is specified on the SRDFSnap Resource Configuration panel, the default save pool area is used.

For information about the actual procedure:

See [“Preparing the fire drill configuration”](#) on page 257.

About running the fire drill

When running the fire drill, the wizard brings the HTCSnap or SRDFSnap agent online. The HTCSnap or SRDFSnap agent manage the replication and mirroring functionality according to the attribute settings. The Snap agents take a consistent snapshot of the replicating data using the snapshot or 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, STD, or VDEV device
- Resumes replication
- Modifies the disk group name in the snapshot

For information about the actual procedure:

See [“Running a fire drill”](#) on page 264.

About restoring the fire drill configuration

When restoring the fire drill configuration to a prepared state, the wizard takes the fire drill service group offline, thereby taking the SRDF and HTC Snap agents offline.

This action reattaches the hardware mirrors to the replicating secondary devices and resynchronizes them.

For information about the actual procedure:

See [“Restoring the fire drill system to a prepared state”](#) on page 269.

About deleting the fire drill configuration

When deleting the fire drill configuration, the wizard does the following:

- Delinks the fire drill service group from the corresponding application service group.
- Deletes the fire drill service group
- Deletes any associated registry entries

If you want to remove the hardware mirrors, you must do so manually.

For information about the actual procedure:

See [“Deleting the fire drill configuration”](#) on page 271.

For more information about the Hitachi TrueCopy Snap agent functions, see *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*.

For more information about the EMC SRDF Snap agent functions, see *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.

For more information about scripts or cmdlets related to Exchange Server:

-
- See [“Exchange 2010 scripts or cmdlets”](#) on page 252.

To run a cmdlet

- 1 Create a .bat file with the following entry:

```
%windir%\system32\WindowsPowerShell\v1.0\
PowerShell.exe -command "$ScriptName"
```

In this entry, `$ScriptName` is either the fully qualified .ps1 script, or the cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\
PowerShell.exe -command C:\myTest.ps1
```

- 2 Specify the name of the .bat file as the script to run.

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"
```

Here, `$ExchangeDir` is the directory where Exchange 2010 is installed, and `$ScriptName` is the .ps1 script (with fully qualified path) or cmdlet specified by the user.

For example:

```
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 re-create 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.

The following table gives more details of the process of configuring and running fire drills with the wizard.

Table 10-1 Tasks 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 254.
Prepare the fire drill configuration	Use the wizard to configure the fire drill. See "Preparing the fire drill configuration" on page 257.
Re-create 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 re-create the fire drill configuration before running the fire drill. See "Re-creating a fire drill configuration that has changed" on page 267.
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 264. 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 groups remain online. It is recommended that you restore a fire drill service group to a prepared state before running a fire drill on another service group.

Table 10-1 Tasks for configuring and running fire drills (*continued*)

Action	Description
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 269. This operation takes the fire drill service group offline and reattaches snapshot mirrors.
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 271. The wizard deletes the service group on the secondary site. In a Volume Replicator 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 Volume Replicator environment”](#) on page 255.

See [“Prerequisites for a fire drill in a Hitachi TrueCopy environment”](#) on page 256.

See [“Prerequisites for a fire drill in an EMC SRDF environment”](#) on page 256.

Prerequisites for a fire drill in a Volume Replicator 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 254.

Make sure that the following additional prerequisites are met before configuring and running a fire drill in a Volume Replicator environment:

- The primary and secondary sites must be fully configured with Volume Replicator replication and the global cluster option.
- 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 Volume Replicator 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 254.

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. Make sure that you have configured 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 254.

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. Make sure that you have configured disaster recovery with EMC SRDF.

- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- To take snapshots of R2 devices, appropriate additional devices must be associated with the RDF2 device group and fully established with the devices.
- The infrastructure to take snapshots at the secondary site must be properly configured between the secondary site source and target arrays. Depending on the snapshot technology in use, this process involves the following tasks:
 - Mirror: Associate Symmetric Business Continuance Volumes (BCVs) and synchronize them with the secondary site source (STD devices).
 - Clone: Make sure that no clone session is in progress.
The source and target devices must be of the exact same size.
 - Snap: Make sure that sufficient save pool area is configured and that no snap session is in progress.
The source and target devices must be of the exact same size.
- 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 non-replicated devices, create a EMC Symmetrix device group with the same name as the SFW disk group. The device group must contain the same devices as in the disk group and have the corresponding BCV, STD, or VDEV devices 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 Volume Replicator 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 254.

To prepare the fire drill configuration

- 1 Open the Solutions Configuration Center (From **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu).
- 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 Start the Fire Drill Wizard (expand Solutions for Additional Applications, expand Fire Drill, expand Configure or run a fire drill, and click Fire Drill Wizard).
- 4 In the Welcome panel, review the information and click **Next**.
- 5 In the System Selection panel, specify a system in the primary site cluster and click **Next**.

See [“System Selection panel details”](#) on page 260.

- 6 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 261.

- 7 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 261.

- 8 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 254.

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 264.
---	---

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 269.
--	--

If the Re-create 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 re-create the fire drill configuration to bring it up to date. See “Re-creating a fire drill configuration that has changed” on page 267. Or you can clear the check box to re-create the configuration and run the fire drill on the existing configuration.
--	--

- 9 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.

Volume Replicator replication	Disk Selection panel See “Disk Selection panel details” on page 261.
Hitachi TrueCopy replication	Horcm Files Path Selection panel See “Hitachi TrueCopy Path Information panel details” on page 262. HTCSnap Resource Configuration panel See “HTCSnap Resource Configuration panel details” on page 262.
EMC SRDF replication	SRDFSnap Resource Configuration panel See “SRDFSnap Resource Configuration panel details” on page 263.

Click **Next**.

- 10 In the Fire Drill Preparation panel, the wizard shows the status of the preparation tasks.

See [“Fire Drill Preparation panel details”](#) on page 264.

When preparation is complete, click Next.

- 11 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 264.

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.

See [“Preparing the fire drill configuration”](#) on page 257.

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.

See [“Preparing the fire drill configuration”](#) on page 257.

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.

See [“Preparing the fire drill configuration”](#) on page 257.

Disk Selection panel details

During fire drill preparation in a Volume Replicator 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 re-creating 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 FDnn.
---------------	---

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. If the production volumes reside on disks in the same disk group, you can store multiple snapshot volumes on a single disk. If the volumes in a disk group are configured on multiple RVG resources, provide a separate disk for each RVG. Note: The Fire Drill Wizard does not allow creating mirrors of multiple RVGs from a single disk group on the same disk. You must select a different disk for each RVG in a 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 is a display-only field.

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.

See [“Preparing the fire drill configuration”](#) on page 257.

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 249.

SRDFSnap Resource Configuration panel details

Depending on the snapshot technology in use, the wizard validates the following when preparing for fire drill in an EMC SRDF replication environment:

- Mirror:
 - The number of BCV devices matches that of the STD devices.
 - The BCV devices are associated and synchronized with the STD devices.
- Clone:
 - The number of BCV devices (or STD devices in case of Targets) matches that of the STD devices.
 - No clone session is in progress.
- Snap:
 - The number of VDEV devices matches that of the STD devices.
 - No snap session is in progress.

If these criteria are not satisfied, the wizard displays a warning on this panel. The wizard does not check whether the sizes of the source and target devices match, and therefore does not display a warning. The following figure depicts such a warning.

However, you can proceed with the configuration. The wizards configures the fire drill service group, but is unable to bring the service group online.

This panel lists all the SRDFSnap resources that will be configured. If you do not want to include the dependent disk groups of a SRDFSnap resource in the fire drill, clear the check box against its name.

The name of the resource that is managing the LUNs that you want to snapshot appears as the Target Resource Name. 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.

You can specify the TimeFinder snapshot technology to be used for configuring fire drill for the SRDFSnap resources:

- **Mirror**
BCV devices are used to create snapshots.
- **Clone**
BCV devices are used to create snapshots. Optionally, you can specify that Target devices be used. If you select the **Use Target Devices** check box, STD devices are used to create snapshots.
- **Snap**
VDEV devices are used to create snapshots. The default SavePoolArea is used. Optionally, to use a different SavePoolArea, specify its name.

To discover the most recent SRDF configuration information, click **Refresh**.

See [“Preparing the fire drill configuration”](#) on page 257.

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 249.

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 Volume Replicator 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 re-creating a fire drill configuration, snapshot mirrors are prepared only for new volumes.

See [“Re-creating a fire drill configuration that has changed”](#) on page 267.

See [“Preparing the fire drill configuration”](#) on page 257.

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 251.

For details on the operations that occur when running a fire drill, refer to the following topics:

- See [“About Fire Drill Wizard operations in a Volume Replicator environment”](#) on page 247.
- See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 249.

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 269.

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 re-creating a fire drill service group, go to step 6.

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 Re-create Fire Drill Service Group panel, showing the differences.

Choose one of the following:

 - Leave the option checked to re-create the configuration before running the fire drill and click **Next**. You complete additional steps in the wizard before running the fire drill.

See [“Re-creating a fire drill configuration that has changed”](#) on page 267.
 - To run the fire drill on the existing configuration, clear the option to re-create 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 [“About post-fire drill scripts”](#) on page 251.
- 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 269.

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.

See [“About post-fire drill scripts”](#) on page 251.

- 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.

Re-creating 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. The wizard compares the resource names of the two service groups. If differences are found, the wizard lists them on the Re-create Fire Drill Service Group panel.

The wizard also checks the RVGs configured for disk groups. If a single RVG is configured per disk, the wizard allows you to re-create the service group; the existing snapshots are retained. If multiple RVGs are configured on a disk, the wizard only allows you to delete the service group; the existing snapshots are deleted. To create a corresponding new one, you need to launch the wizard again and perform the fire drill preparation steps.

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 re-create 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.

You have the following choices from the Re-create Fire Drill Service Group panel:

- Leave the option checked to re-create the fire drill service group. Proceed with using the wizard to re-create the configuration to match the application service group. The wizard deletes the existing fire drill configuration first, before creating the new one.

For a Volume Replicator replication environment, the wizard handles existing volumes as follows: It does not delete the mirrors for volumes that still exist. When it re-creates 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 re-create 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.

The following procedure describes the choice of re-creating the fire drill configuration.

Note: Veritas recommends that you do not use this procedure to re-create any existing fire drill service groups after performing an upgrade. Instead, use the Fire Drill Wizard to delete the existing service groups and create corresponding new ones.

To re-create the fire drill configuration if the service group has changed

- 1** In the Re-create Fire Drill Service Group panel, leave the option checked to re-create the configuration before running the fire drill.

For a Volume Replicator 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 Volume Replicator 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 254.

- 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.

Volume Replicator 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 261.

Hitachi TrueCopy replication	Horcm Files Path Selection panel See “Hitachi TrueCopy Path Information panel details” on page 262. HTCSnap Resource Configuration panel See “HTCSnap Resource Configuration panel details” on page 262.
------------------------------	---

EMC SRDF replication	SRDFSnap Resource Configuration panel See “SRDFSnap Resource Configuration panel details” on page 263.
----------------------	---

Click **Next**.

- 6 The Fire Drill Preparation panel is displayed. Wait while the wizard re-creates the fire drill service group.

For Volume Replicator 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 264.

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:

- See [“About Fire Drill Wizard operations in a Volume Replicator environment”](#) on page 247.
- See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 249.

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 [6](#).

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 Volume Replicator 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 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 If the wizard detects that the fire drill service group is different from the application service group, it displays the Re-create Fire Drill Service Group panel. Clear the option to re-create 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**.

Fire Drill Deletion panel details

The Fire Drill Deletion panel shows the progress of the deletion tasks. Wait until all tasks are complete before you click **Next** to proceed to the Summary panel and exit the wizard.

See [“Deleting the fire drill configuration”](#) on page 271.

The Fire Drill Deletion panel also appears if you selected the option to re-create a fire drill configuration. During a re-create operation, wait until all deletion tasks are complete and then click **Next** to continue with the fire drill preparation.

When the wizard re-creates a configuration in a Volume Replicator environment, it deletes the service group but maintains the required snapshot volumes. If you close the wizard without continuing to the preparation step, the snapshot volume information is lost and the fire drill configuration is fully deleted.

Considerations for switching over fire drill service groups

In a Volume Replicator environment, if you directly switch the fire drill service group from one node to another, the VVRSnap resource fails to come online on the target node. The fire drill service group depends on the RVG service group. To make the switch successfully, you must first switch the RVG service group to the intended node and then switch the fire drill service group.

Reference

- [Appendix A. Using Veritas AppProtect for vSphere](#)
- [Appendix B. Troubleshooting](#)

Using Veritas AppProtect for vSphere

This appendix includes the following topics:

- [About Just In Time Availability](#)
- [Prerequisites](#)
- [Setting up a plan](#)
- [Deleting a plan](#)
- [Managing a plan](#)
- [Viewing the history tab](#)
- [Limitations of Just In Time Availability](#)
- [Getting started with Just In Time Availability](#)
- [Supported operating systems and configurations](#)
- [Viewing the properties](#)
- [Log files](#)
- [Plan states](#)
- [Troubleshooting Just In Time Availability](#)

About Just In Time Availability

The Just In Time Availability solution provides increased availability to the applications on a single node InfoScale Availability cluster in VMware virtual environments.

Using the Just In Time Availability solution, you can create plans for:

- Planned Maintenance
- Unplanned Recovery

Planned Maintenance

In the event of planned maintenance, the Just In Time Availability solution enables you to clone a virtual machine, bring it online, and fail over the applications running on that virtual machine to the clone on the same ESX host. After the maintenance procedure is complete, you can fail back the applications to the original virtual machine. Besides failover and failback operations, you can delete a virtual machine clone, view the properties of the virtual machine and its clone, and so on.

Unplanned Recovery

When an application encounters an unexpected or unplanned failure on the original or primary virtual machine on the primary ESX host, the Just In Time Availability solution enables you to recover the application and bring it online using the unplanned recovery feature.

With **Unplanned Recovery Policies**, the Just In Time Availability solution enables you to set up recovery policies to mitigate unplanned failures that are encountered by an application. Just In Time Availability solution provides the following recovery policies; you may select one or all the recovery policies as per your need:

Unplanned Recovery Policies	Description
-----------------------------	-------------

Restart Application	
---------------------	--

	Just In Time Availability (JIT) solution attempts to restart the service group (SG), and bring the application online on the original virtual machine on primary ESX.
--	---

	Maximum three retry attempts are permitted under this policy.
--	---

	Note: If all the three attempts fail, application continues to remain in faulted state or continues with the next policy as selected while creating a plan.
--	--

Unplanned Recovery Policies	Description
Restart virtual machine (VM)	Just In Time Availability (JIT) solution performs the following subsequent tasks: ■ take the service group offline ■ shut down the virtual machine ■ power on the virtual machine ■ bring the service group online on the original virtual machine on primary ESX You are provided with Last attempt will be VM reset option to reset the virtual machine. By default, this checkbox is selected and the default retry attempt value is one. If you retain the default settings, then VM reset operation is performed on the virtual machine at the first attempt itself. Maximum three retry attempts are permitted for this operation. If you deselect the checkbox, then the virtual machine reset (VM Reset) operation is not performed.
Restart VM on target ESX	Using this policy, you can recover the faulted application on the virtual machine. In this policy, the original virtual machine is unregistered from the primary ESX; registered on the target ESX; and the faulted application is brought online on the target ESX.
Restore VM on target ESX	Using this policy, you can recover the faulted application on the virtual machine using a boot disk backup copy of the original virtual machine. In this policy, the original virtual machine is unregistered from the ESX and the boot disk backup copy of the original virtual machine is registered on target ESX. The faulted application is then brought online on the virtual machine.

Unplanned Recovery Policies

Unplanned Failback

The **Unplanned Failback** operation lets you fail back the application from the boot disk backup copy of virtual machine on the target ESX to the original virtual machine on primary ESX.

If you have selected either **Restart VM on target ESX** or **Restore VM on target ESX** or both the recovery policies, you can perform the **Unplanned Failback** operation.

On the **Plans** tab, in the plans table list, right-click the virtual machine and click **Unplanned Failback**.

Note: **Unplanned Failback operation** operation is disabled and not available for the plans and the virtual machines which have **Restart Application** and **Restart VM** policies as the only selected options.

Based on the selected recovery policy for a plan, Just In Time Availability (JIT) solution performs the necessary operations in the sequential order.

For example, if you have selected **Restart Application** and **Restart VM** as the recovery policy, then in the event of unplanned application failure, first it performs tasks for **Restart Application** policy and if that fails, it moves to the next policy.

You may select one or all the recovery policies based on your requirement.

[Table A-1](#) lists the sequence of tasks that are performed for each Unplanned Recovery policy.

Table A-1 Tasks performed for each Unplanned Recovery policy

Unplanned Recovery Policy	Tasks Performed
Restart Application	◆ Make an attempt to restart the application.
Restart virtual machine (VM)	1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Power on the virtual machine 4 Brings the service group(s) online

Table A-1 Tasks performed for each Unplanned Recovery policy (*continued*)

Unplanned Recovery Policy	Tasks Performed
Restart VM on target ESX	1 Takes the service group(s) offline 2 Shuts down the original virtual machine 3 Detaches the data disks from the original virtual machine 4 Unregisters the virtual machine from the primary ESX 5 Registers the original virtual machine on target ESX 6 Attaches the data disks back to the virtual machine 7 Power on the virtual machine 8 Brings the service group(s) online
Restore VM on target ESX	1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Detaches the data disks from the virtual machine 4 Unregisters the original virtual machine from the target ESX 5 Registers the boot disk backup copy of the original virtual machine to the target ESX 6 Attaches the data disks back to the virtual machine 7 Power on the virtual machine 8 Brings the service group(s) online
Unplanned Failback	1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Detaches the data disks from the virtual machine 4 Unregisters the virtual machine from the target ESX 5 Registers the virtual machine using the original boot disk backup copy to the primary ESX 6 Attaches the data disks to the virtual machine 7 Power on the virtual machine on primary ESX 8 Brings the service group(s) online on the virtual machine

Scheduler Settings

While creating a plan for unplanned recovery, with **Scheduler Settings**, you can set up a schedule for taking a back up of boot disk of all the virtual machines that are a part of the plan.

To use the Just In Time Availability solution, go to **vSphere Web Client > Home view > Veritas AppProtect**.

See [“Setting up a plan”](#) on page 281.

Prerequisites

Before getting started with Just In Time Availability, ensure that the following prerequisites are met:

- The Just In Time (JIT) solution feature cannot co-exist with VMware HA, VMware FT, and VMware DRS. This pre-requisite is applicable for **Unplanned Recovery** only.
- VIOM 7.2 version must be installed and configured using fully qualified domain name (FQDN) or IP.
- Make sure that you have the admin privileges for vCenter.
- VMware Tools must be installed and running on the guest virtual machine.
- VIOM Control Host add-on must be installed on VIOM server or machine.
- The virtual machines must be added in VIOM. The virtual machines, vSphere ESX servers, and VIOM must have the same Network Time Protocol (NTP) server configured.
- Make sure to specify VIOM Central Server FQDN or IP in the SNMP Settings of the vCenter Server.
- vCenter Server and VIOM must be configured using the same FQDN or IP address. Make sure that if FQDN is used to configure vCenter in VIOM Server that is used during the configuration.
- If raw disk mapping (RDM) disks are added to the virtual machine, then make sure that the virtual machine is in the physical compatibility mode. Veritas AppProtect does not support the virtual compatibility mode for RDM disks.
- For Microsoft Windows operating system, make sure that you have the Microsoft Windows product license key. The key is required to run the Sysprep utility, which enables customization of the Windows operating system for a clone operation.

- For RHEL7 and SUSE12 operating system, install the deployPkg plug-in file on the virtual machine.
For more information on installing the plug-in, see <https://kb.vmware.com/kb/2075048>
- Make sure that the InfoScale Availability service group is configured with one of the storage agents such as Mount, DiskGroup, LVMVolumeGroup, VMNSDg (for Windows), or DiskRes (for Windows), for the data disks. This configuration enables Veritas AppProtect to discover data disks for the applications. Also, ensure that the service group is online to determine data disk mapping.
- Virtual machines which have snapshots associated with them are not supported.
- Virtual machines with SCSI Bus Sharing are not supported.
- Make sure that the SNMP Traps are configured for the following from vCenter server to VIOM:
 - Registered virtual machine
 - Reconfigured virtual machine
 - Virtual machine which is getting cloned
- Make sure that the boot disk of VM's (vmdk) does not have spaces.
- For HA console add on upgrade from VIOM 7.1 to VIOM 7.2, refer *Veritas InfoScale Operations Manager 7.2 Add-ons User's Guide* for more details.
- Make sure to set the vSphere DRS Automation Level to manual, if you want to configure **Restart VM on target ESX** or **Restore VM on target ESX** policies for your plan.
- Ensure to update or edit the plan, when a virtual machine is migrated or if there are any modifications made to the settings of the virtual machines which are configured for that plan.
- Ensure to increase the tolerance limit of DiskRes resource to two, if you want to create a plan for unplanned recovery with **Restore VM on target ESX** as the unplanned recovery policy.

Note: This prerequisite is applicable for Windows operating system.

Setting up a plan

Plan is a template which involves a logical grouping of virtual machines so as to increase the availability of the application in the event of a planned failover and recovery of the application in the event of an unexpected application failure.

To set up a plan

- 1 Launch Veritas AppProtect from the **VMware vSphere Web Client > Home view > Veritas AppProtect** icon.

- 2 Click **Configure Plan**.

The **Plan Configuration** wizard appears.

- 3 Specify a unique **Plan Name** and **Description**, and then click **Next**.

The wizard validates the system details to ensure that all prerequisites are met.

- 4 Select the virtual machines that you want to include in the plan, review the host and operating system details, and then click **Next**.

The **Unplanned Recovery Settings** page appears.

- 5 On the **Unplanned Recovery Settings** page, you can configure the selected virtual machines for **Unplanned Recovery** as well.

Deselect the **Configure selected VMs for Unplanned Recovery as well** check box, if you do not want to include the selected virtual machines for unplanned recovery.

If you have selected the virtual machines for unplanned recovery, then set up the unplanned recovery policies as appropriate from the available options. You can set up policies to restart applications, restart virtual machines, restart virtual machine on target ESX, and restore a virtual machine on target ESX.

If you have selected **Restore VM on target ESX** as the unplanned recovery policy, then you can set up a schedule to create a boot disk back up copy of the virtual machine within the configured plan. You can set the frequency as daily, weekly, monthly, or manual as per your requirement.

After you have finished making necessary settings for Unplanned Recovery, Click **Next**.

- 6 The wizard validates the prerequisite attributes of the virtual machine and the ESX host, and adds the qualified virtual machines to the plan.

Click **Next** after the validation process completes.

- 7 In the **Disks** tab, you can view the selected application data disks. Just In Time Availability solution uses the selected data disks to perform detach-attach operation during a planned failover and unplanned recovery.

Note: If the disks are not auto-marked as selected to perform detach-attach operation, then first refresh the VIOM server and then the VCenter server in VIOM and then create a plan.

- 8 In the **Network Configuration** tab, specify the network interface configuration details for the cloned virtual machine. Make sure to specify at least one public interface and valid IP details.
- 9 In the **Unplanned Recovery Target** tab, specify the target ESX server to restore the virtual machine, and the target ESX port details.

Note: The **Unplanned Recovery Target** tab is visible only when **Restart VM on target ESX** or **Restore VM on target ESX** is selected.

- 10 In the **Windows Settings** tab, specify the domain name, Microsoft Windows product license key, domain user name, domain password, admin password, and time zone index.

Note: The **Windows Settings** tab is visible only when a Windows virtual machine is selected in the plan.

- 11 Click **Next**. The **Summary** wizard appears.
- 12 In the **Summary** wizard, review the plan details such as the plan name, unplanned recovery policies, schedule, and so on.

Deselect the **Start backup process on finish** checkbox if you do not want to initiate a backup process when the plan creation procedure is finished. This checkbox is selected by default.

Click **Create**. The plan is created and saved.

- 13 Click **Finish** to return to the plans tab and view the created plans.

See [“Managing a plan”](#) on page 283.

See [“Deleting a plan”](#) on page 283.

Deleting a plan

After you have finished performing failback operations from the clone to the primary virtual machine in case of planned maintenance and recovery operations in case of unplanned recovery, you may want to delete the plan.

To delete a plan

- 1 Launch **Veritas AppProtect** from the VMware vSphere Web Client Home view.
- 2 In the **Plans** tab, select the plan that you want to delete.
- 3 Click **Delete Plan**.

Note: The **Delete plan** icon is enabled only when the selected plan is in **Ready For Failover**, **Failed to Revert**, and **Failed to Failback** state.

Managing a plan

Planned Maintenance

After the maintenance plan is created, you can fail over the applications to the clone virtual machine and fail back the applications from the clone to the virtual machine. When the scheduled maintenance is complete, you can delete the cloned virtual machine or retain it for future use.

To perform failover, failback, revert, or delete clone operations, go to **Plans**, and select a plan. Based on the enabled operation, perform the following tasks:

To fail over the applications to the cloned virtual machine

- ◆ Click the **Failover** icon.

Just In Time Availability (JIT) performs the sequence of failover tasks, which includes taking the application offline, detaching the disks, cloning the virtual machine, attaching the disks, and so on.

To fail back the applications from the clone to the primary virtual machine

- ◆ Click the **Failback** icon.

Just In Time Availability (JIT) performs the sequence of failback tasks, which includes taking the application offline, detaching the disks, attaching the disks, and so on.

To revert a failover or a failback operation

- ◆ Click the **Revert** icon.

If the failover or a failback operation fails, the revert operation restores the applications on the virtual machine, and deletes the clone if created.

To delete a clone

- ◆ Click the **Delete Clone** icon.

After the failback operation is complete, you can delete the clone. By default, the revert operation deletes the clone.

Note: Alternatively, right-click **Plan** in the **Plans** table on the **Plans** wizard to perform failover, failback, revert, delete plan, and delete clone operations.

Unplanned Recovery

Once you have set up a plan for unplanned recovery during **Configure Plan** operation, based on the recovery policies selected for the plan, the application is recovered accordingly.

You can manage unplanned recovery policies settings by performing the following operations on the plan and its associated virtual machines.

Managing unplanned recovery settings

On the **Plans** tab, in the plans table which lists all the existing plans, navigate to the required plan and use the right-click option on the selected plan.

- **Edit:** Use this option to modify the configured plans settings such as adding or removing a virtual machine from the plan, and so on.
The same **Configuration Plan** wizard using which you had set up or configured a plan is displayed with pre-populated details.
See [“Setting up a plan”](#) on page 281.
- **Disable Unplanned Recovery:** Use this option to disable the Unplanned Recovery settings.
- **Enable Unplanned Recovery:** Use this option to enable the Unplanned Recovery settings.
- **Disable Scheduler:** Use this option to disable the scheduler settings.
- **Enable Scheduler:** Use this option to enable the scheduler settings.
- **Delete Plan:** Use this option to delete the created plan.
- **Properties:** Use this option to view the properties for unplanned recovery. It displays details such as the selected unplanned recovery policies and the

associated operations for the selected policies. It also provides information about the selected scheduler mode for performing boot disk back up operation for the selected virtual machines.

Managing virtual machines settings

On the **Plans** tab, in the plans table which lists all the existing plans and its associated virtual machines, navigate to the required virtual machine. Select the required virtual machine and use the right-click option on the selected virtual machine.

- **Remove VM From Plan:** Use this option to delete the virtual machine from the selected plan.
- **Create Clone Backup:** Use this option to create a boot disk backup copy of the virtual machine.
- **Unplanned Failback:** Use this option to fail back the application from the boot disk backup copy of the virtual machine on target ESX to the original virtual machine on primary ESX.

Note: This option is available only if you have set unplanned recovery policies as **Restart VM on target ESX** or **Restore VM on target ESX**.

- **Properties:** Use this option to view properties such as the last run time for backup operation, last successful backup attempt time and the target ESX details.

See [“Plan states”](#) on page 290.

Viewing the history tab

On the **History** tab, you can view the detailed summary of the operations that are performed on the virtual machine. The details include the plan name, virtual machine name, operation, the status of the operation, the start and the end time of the operation, and the description of the operation status.

To view the summary

- 1 Launch **Veritas AppProtect** from the VMware vSphere Web Client Home view.
- 2 Click the **History** tab.

Limitations of Just In Time Availability

The following limitations are applicable to Just In Time Availability:

- On a single ESX host only ten concurrent failover operations are supported. Across ESX hosts, twenty concurrent failover operations are supported.
- Linked mode vCenter is not supported.
- Only three backup operations per data store are active, the rest will be queued. Only five backup operations per ESX host are active, the rest will be queued.

See [“Supported operating systems and configurations”](#) on page 288.

Getting started with Just In Time Availability

You can access the Just In Time Availability solution from the **vSphere Web Client > Veritas AppProtect** interface.

The **Veritas AppProtect** is registered with Veritas InfoScale Operations Manager (VIOM), and is accessed from the **vSphere Web Client > Home** view.

[Figure A-1](#) describes the Veritas AppProtect interface in detail.

Figure A-1 Elements of the Veritas AppProtect interface




Table A-2 Elements of the Veritas AppProtect interface and the description

Label	Element	Description
1	Plans tab	Enables setting up a plan for a planned failover and unplanned recovery. Displays the plan attributes, and the virtual machines that are added to the plan. Displays the status of virtual machines for unplanned recovery and schedule for virtual machine back up operation based on the criteria set while configuring or editing the plan. Shows the enabled or disabled failover, failback, delete clone, revert, delete plan, and properties operations icons based on the state of the selected plan for planned failover.
2	History tab	Displays the status and the start and the end time of the specific operation performed on the created plans.
3	Configure Plan link	Opens the Plan Configuration wizard.
4	Plans table	Displays the attributes of the plan.
5	Failover icon	Fails over the applications from the original virtual machine to the clone.
6	Failback icon	Fails back the applications from the clone to the original virtual machine.
7	Delete Clone icon	Deletes the cloned virtual machine.
8	Revert State icon	Reverts the failed operation, restores the applications to the original virtual machines, and delete the clone virtual machines.
9	Delete Plan icon	Deletes the plan.
10	Properties icon	Displays the attributes of each virtual machine and the clone.

Table A-2 Elements of the Veritas AppProtect interface and the description
(continued)

Label	Element	Description
11	Operation-specific tabs	Displays the sequence of the tasks that are performed for the selected operation. Based on the operation that is executed, the associate tab opens. For Planned Maintenance 1 Failover 2 Failback 3 Revert 4 Delete Clone For Unplanned Recovery ◆ Unplanned Recovery Summary
12	Diagnostic information	Displays the logs that are reported for the Veritas AppProtect interface.

See “[Plan states](#)” on page 290.

Supported operating systems and configurations

Just In Time Availability supports the following operating systems:

- On Windows: Windows 2012, and Windows 2012 R2.
- On Linux: RHEL5.5, RHEL6, RHEL7, SUSE11, SUSE12.

Just In Time Availability supports the following configurations:

- Veritas Cluster Server (VCS) 6.0 or later, or InfoScale Availability 7.1 and later.
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) 7.1 and 7.2 version on the virtual machines.
For more information about VRTSsfmh, see the *Veritas InfoScale Operations Manager 7.2 User Guide*.
- Veritas InfoScale Operations Manager (VIOM) 7.2 as a central or managed server.
- VMware vSphere 5.5 Update 2, Update 3, or 6.0 and 6.0 Update 1 version.

Viewing the properties

Virtual Machine Properties

The **Virtual Machine Properties** window displays information about the virtual machine and its clone such as name, operating system, cluster name, service groups, DNS server, domain, IP addresses, and data disks.

To view the properties

- 1 On the **Plans** tab, select the virtual machine.
- 2 Click the **Properties** icon or right-click the virtual machine.

The **Virtual Machine Properties** window opens and displays the attributes of the virtual machine and its clone.

Plan Properties

The **Plan Properties** window displays information about the unplanned recovery policies selected; scheduler mode set; and the time when the last backup operation was run and was successful for a virtual machine.

To view properties for the plan

- 1 In the Plan Name table, select the plan.
- 2 Right-click the selected plan. A window with a list of options is displayed.
- 3 Click **Properties**

The **Plan Properties** window opens and displays the unplanned recovery policies selected and the schedule mode for virtual machine backup operation.

Log files

The following log files are helpful for resolving the issues that you may encounter while using Veritas AppProtect:

- Console related logs:

```
/var/opt/VRTSsfmcs/logs/*
```

These log files show console messages and are useful for debugging console issues.

- Operations logs:

```
/var/opt/VRTSsfmh/logs/vm_operations.log
```

This log file shows the messages pertinent to the Veritas AppProtect interface.

- VMware vSphere 6.0 logs:

C:\ProgramData\VMware\vCenterServer\logs\vsphere-client\logs*

These log files show the messages that are reported for the VMware vSphere Web Client version 6.0.

- VMware vSphere 5.5 U2 and U3 logs:

C:\ProgramData\VMware\VSphere Web Client\serviceability\logs*

These log files show the messages that are reported for the VMware vSphere Web Client version 5.5 U2 and U3.

- Veritas AppProtect interface logs:

The log file shows the logs that are reported for the Veritas AppProtect interface. To view the log files, on the **Planned Maintenance** tab or the **History** tab > **Diagnostic Information**.

Plan states

Based on the state of the plan, the operation icons are enabled and disabled on the **Plans** tab.

Table A-3 List of plan and operation states

Plan state	Failover	Failback	Revert	Delete clone	Delete Plan	Unplanned Failback	Create Clone backup	Properties
Ready For Failover	✓	–	–	✓ Note: Enabled when the selected maintenance plan has an associate clone.	✓ Note: Enabled when the selected maintenance plan does not have an associate clone.	–	✓	✓
Failed Over	–	✓	–	–	–	–	–	✓
Failed To Failover	–	–	✓	–	–	–	–	✓

Table A-3 List of plan and operation states (*continued*)

Plan state	Failover	Failback	Revert	Delete clone	Delete Plan	Unplanned Failback	Create Clone backup	Properties
Failed To Failback	–	–	✓	–	–	–	–	✓
Failed To Revert	–	–	✓	–	✓	–	–	✓
Unknown	–	–	✓	–	–	✓	–	✓
Failed To Delete Clone	–	–	–	✓	–	–	–	✓
Failover In Progress	–	–	–	–	–	–	–	✓
Failback In Progress	–	–	–	–	–	–	–	✓
Revert In Progress	–	–	–	–	–	–	–	✓
Delete Clone In Progress	–	–	–	–	–	–	–	✓
Application Faulted	–	–	–	–	–	–	–	✓
Failed To Restart VM	–	–	–	–	–	–	–	✓
Failed To Move VM	–	–	–	–	–	✓	–	✓
Failed To Restore VM	–	–	–	–	–	✓	–	✓
Unplanned	–	–	–	–	–	✓	✓	–
Unplanned Restored VM	–	–	–	–	–	✓	–	✓
Unplanned Failed to Failback	–	–	–	–	✓	–	–	–

Troubleshooting Just In Time Availability

Table A-4 lists the issues and the recommended solutions.

Table A-4 Issues and the corresponding resolutions

Issue	Recommended Solution
When setting up a maintenance plan, the registered virtual machine is not listed on the wizard.	To troubleshoot the issue, make sure the following: ■ ESX host on which the virtual machine resides, is connected to the vCenter. ■ The virtual machine is added as a managed host to Management Server. ■ On the virtual machine, at least one application is configured for monitoring, along with VCS. ■ The virtual machine is registered in VIOM. ■ VCS is configured on the virtual machine. ■ The virtual machine does not contain RHEL7 and SUSE 12, which are not supported. Note: Windows 2012 R2 and 2008 R2 are supported. ■ VCS is configured with the service groups.
When setting up a maintenance plan, the listed virtual machine is not available for selection.	To troubleshoot the issue, make sure the following: ■ The virtual machine is not configured for Global Cluster option (GCO). ■ Agents that support SAN are configured.
When Veritas AppProtect executes an operation, the timeout message is reported.	To troubleshoot the issue, perform the following: ■ If the failover or the failback operation fails, then click Planned Maintenance > Revert icon. Retry the operation. ■ If the delete plan or the delete clone operation fails, then retry the operation.
The revert operation failed.	Manually revert the virtual machine to its original state.

Troubleshooting

This appendix includes the following topics:

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

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

Here,

- Timestamp denotes the date and time when the message was logged.
- Mnemonic denotes which Veritas 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.

- ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action.
- WARNING indicates a warning or error, but not an actual fault.
- NOTE informs that VCS has initiated an action.
- INFO informs about various state messages or comments.
Of these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.
- UMI or Unique Message ID is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the ExchService agent would resemble: V-16-20024-13
Originator ID for all VCS products is 'V-16.' Category ID for ExchService agent is 20024. Message ID is a unique number assigned to the message text.
- Message text denotes the actual message string.

You can view these message logs using Notepad or any text editor. All messages are logged to the engine and the agent logs. Messages of type CRITICAL and ERROR are also written to the Windows event log.

Exchange Service agent error messages

The following table lists the Exchange Service agent error messages and their descriptions.

Table B-1 Exchange Service agent error messages

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

Table B-1 Exchange Service agent error messages *(continued)*

Message	Description
Failed to start the service (service name) Error Type, Error Code.	The agent failed to start the specified service. Solution: See the associated Windows error type and error code for more information.
Failed to stop the service (service name). Error Type, Error Code.	The agent failed to stop the service. Solution: See the associated Windows error type and error code for more information.
Failed to kill the service (service name) Error Type, Error Code.	The agent failed to terminate the service. Solution: See the associated Windows error type and error code for more information.
Configuration error. 'Service' attribute is not configured.	No value specified for the "Service" attribute. Solution: Specify a valid value for the attribute.
Configuration error. 'LanmanResName' attribute is not configured.	No value specified for the "LanManResName" attribute. Solution: Specify a valid value for the attribute.
Failed to set the virtual environment for service: (service name). Error Type, Error Code.	The agent failed to set the environment block for the service. The agent needs to set the environment block for starting the service in the context of the virtual server name. Solution: See the associated Windows error type and error code for more information.
Failed to remove virtual environment for Service = (service name). Error Type, Error Code.	The agent failed to remove the environment block for the service. While taking the resource offline, the agent stops the service and removes the environment block. Solution: See the associated Windows error type and error code for more information.
Configuration error. \"LanmanResName\" attribute is not configured.	No value specified for the "LanmanResName" attribute. Solution: Specify a valid value for the attribute.

Table B-1 Exchange Service agent error messages *(continued)*

Message	Description
Configuration error. DetailMonitoringInterval attribute is greater than zero but DBList is empty. No database is specified for detail monitoring.	Detail monitoring for databases is enabled and the monitoring interval (DetailMonitor attribute) is also specified. But there are no databases selected. The DBList attribute is empty. Solution: Select the databases for the detail monitoring.
FaultOnMountFailure flag is true.\"Auto Mount\" on database: (database names) is enabled but database is dismounted. Agent will return status as offline."	The attribute FaultOnMountFailure is set to True for databases that are set to mount automatically on startup. But these databases are dismounted. So the agent will fault the service group.
\"Auto Mount\" on database: (database names) is enabled but database is dismounted. Agent will return status as Unknown.	Databases that are set to mount automatically on startup are dismounted. If these databases are selected for detail monitoring, the agent will return an Unknown status and appropriate administrative action is required.
Failed to add computer account to 'Exchange Servers' group Error Type, Error Code.	Unable to add the computer account to the Exchange Servers group. Solution: Make sure that the user has permissions to add computer accounts to the Exchange Servers group. If the user has those permissions, see the error type and error code for further information.

Troubleshooting Microsoft Exchange uninstallation

You might encounter errors while removing Microsoft Exchange if any of the following requirements are not adhered to:

- User mailboxes exist.
- The Exchange Server to be uninstalled has routing group connectors configured.
- Public folder databases exist.

In any of the above scenarios, carry out the following steps to resolve the error.

- 1** Start the following Exchange services manually using the Service Control Manager:
 - MExchangeSA
 - MExchangeIS
- 2** Move or delete user mailboxes. See the Exchange documentation for instructions.
- 3** Move or delete public folder. See the Exchange documentation for instructions.
- 4** Stop all Exchange services started in Step 1.
- 5** Start the Exchange Setup Wizard for VCS and select the Remove Exchange option. Note that you must uninstall Exchange only by using the Exchange Setup Wizard for VCS.

Troubleshooting Exchange Setup Wizard issues

When adding a failover node to an existing Exchange cluster, the Exchange Setup Wizard may fail to rename the node during the pre-installation phase, and report the following error message:

`Failed to rename the node. Refer to the log file for further details.`

This can happen if the Exchange Setup Wizard is unable to delete the Exchange Virtual Server computer object in the Active Directory.

To resolve this issue, you must manually delete the Exchange Virtual Server computer object from the AD, and run the wizard again.