

Storage Foundation and High Availability Solutions

7.4.2 HA and DR Solutions

Guide for Enterprise Vault - Windows

Last updated: 2020-06-05

Legal Notice

Copyright © 2020 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
2625 Augustine Drive
Santa Clara, CA 95054
<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

infoscaledocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing SFW HA for EV	7
	About clustering solutions with InfoScale products	7
	About high availability	8
	How a high availability solution works	8
	How VCS monitors storage components	9
	Shared storage—if you use NetApp filers	9
	Shared storage—if you use SFW to manage cluster dynamic disk groups	10
	Shared storage—if you use Windows LDM to manage shared disks	10
	Non-shared storage—if you use SFW to manage dynamic disk groups	11
	Non-shared storage—if you use Windows LDM to manage local disks	11
	Non-shared storage—if you use VMware storage	12
	About replication	12
	About disaster recovery	13
	What you can do with a disaster recovery solution	13
	Typical disaster recovery configuration	13
 Chapter 2	 Configuring high availability for Enterprise Vault with InfoScale Enterprise	 15
	Reviewing the HA configuration	16
	Active-Passive configuration	16
	Reviewing the disaster recovery configuration	18
	Sample disaster recovery configuration	19
	IP addresses for disaster recovery configuration	20
	Supported disaster recovery configurations for service group dependencies	21
	High availability (HA) configuration (New Server)	22
	Following the HA workflow in the Solutions Configuration Center	25
	Disaster recovery configuration	25
	DR configuration tasks: Primary site	26
	DR configuration tasks: Secondary site	29

Notes and recommendations for cluster and application configuration	31
IPv6 support	34
Configuring the storage hardware and network	35
Configuring cluster disk groups and volumes for Enterprise Vault	36
About cluster disk groups and volumes	37
Prerequisites for configuring cluster disk groups and volumes	37
Considerations for a fast failover configuration	38
Considerations for disks and volumes for campus clusters	39
Considerations for volumes for a Volume Replicator configuration	39
Sample disk group and volume configuration	40
Viewing the available disk storage	40
Creating a cluster disk group	41
Creating Volumes	41
About managing disk groups and volumes	41
Importing a disk group and mounting a volume	41
Unmounting a volume and deporting a disk group	42
Adding drive letters to mount the volumes	42
Deporting the cluster disk group	44
Configuring the cluster	44
Adding a node to an existing VCS cluster	45
Verifying your primary site configuration	49
Guidelines for installing InfoScale Enterprise and configuring the cluster	
on the secondary site	49
Setting up your replication environment	50
Setting up security for Volume Replicator	51
Assigning user privileges (secure clusters only)	53
Configuring disaster recovery with the DR wizard	54
Cloning the storage on the secondary site using the DR wizard (Volume	
Replicator replication option)	58
Installing and configuring Enterprise Vault on the secondary site	62
Configuring Volume Replicator replication and global clustering	64
Configuring global clustering only	72
Setting service group dependencies for disaster recovery	74
Verifying the disaster recovery configuration	74
Establishing secure communication within the global cluster (optional)	
.....	76
Adding multiple DR sites (optional)	78
Recovery procedures for service group dependencies	78

Chapter 3	Using the Solutions Configuration Center	81
	About the Solutions Configuration Center	81
	Starting the Solutions Configuration Center	82
	Options in the Solutions Configuration Center	82
	About launching wizards from the Solutions Configuration Center	83
	Remote and local access to Solutions wizards	84
	Solutions wizards and logs	85
	Workflows in the Solutions Configuration Center	86
 Chapter 4	 Installing and configuring Enterprise Vault for failover	 87
	Installing Enterprise Vault	87
	Configuring the Enterprise Vault service group	87
	Before you configure an EV service group	88
	Creating an EV service group	88
	Enabling fast failover for disk groups (optional)	92
	Configuring Enterprise Vault Server in a cluster environment	93
	Setting service group dependencies for high availability	94
	Verifying the Enterprise Vault cluster configuration	95
	Setting up Enterprise Vault	96
	Considerations when modifying an EV service group	96

Introducing SFW HA for EV

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 replication](#)
- [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 with Enterprise Vault:

- High availability failover cluster in an active/passive configuration on the same site
- 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.

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 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.

InfoScale Enterprise provides Volume Replicator for use in replication. Volume Replicator can be used for replication in either a replicated data cluster (RDC) or a wide area disaster recovery solution.

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

About disaster recovery

Wide area disaster recovery (DR) provides the ultimate protection for data and applications in the event of a disaster. If a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away. InfoScale Enterprise 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 1-1 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 application on System1 fails, the application comes online on node System2 and 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.

Configuring high availability for Enterprise Vault with InfoScale Enterprise

This chapter includes the following topics:

- [Reviewing the HA configuration](#)
- [Reviewing the disaster recovery configuration](#)
- [High availability \(HA\) configuration \(New Server\)](#)
- [Following the HA workflow in the Solutions Configuration Center](#)
- [Disaster recovery configuration](#)
- [Notes and recommendations for cluster and application configuration](#)
- [Configuring the storage hardware and network](#)
- [Configuring cluster disk groups and volumes for Enterprise Vault](#)
- [Configuring the cluster](#)
- [Adding a node to an existing VCS cluster](#)
- [Verifying your primary site configuration](#)
- [Guidelines for installing InfoScale Enterprise and configuring the cluster on the secondary site](#)

- [Setting up your replication environment](#)
- [Setting up security for Volume Replicator](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [Configuring disaster recovery with the DR wizard](#)
- [Cloning the storage on the secondary site using the DR wizard \(Volume Replicator replication option\)](#)
- [Installing and configuring Enterprise Vault on the secondary site](#)
- [Configuring Volume Replicator replication and global clustering](#)
- [Configuring global clustering only](#)
- [Setting service group dependencies for disaster recovery](#)
- [Verifying the disaster recovery configuration](#)
- [Establishing secure communication within the global cluster \(optional\)](#)
- [Adding multiple DR sites \(optional\)](#)
- [Recovery procedures for service group dependencies](#)

Reviewing the HA configuration

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


See [“Active-Passive configuration”](#) on page 16.

Active-Passive configuration

In a typical example of a high availability cluster, you create a virtual Enterprise Vault server in an Active-Passive configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node.

The following figure illustrates a typical Active-Passive configuration.

Figure 2-1 Active Passive configuration



Enterprise Vault Server is installed on both Node1 and Node2 and configured as a virtual server with a virtual IP address. Shared volumes are configured on shared storage for the following:

- MSMQ data
- Registry replication data
- Various EV services data (Indexing service, Shopping service, Vault store partitions, PST holding folders, etc.)

Veritas recommends as a best practice to configure SQL Server for high availability before configuring Enterprise Vault. You will specify the SQL virtual server name during EV configuration.

Configuring SQL Server for high availability is covered in the SQL Server solutions guides.

Sample Active-Passive configuration

A sample setup is used to illustrate the installation and configuration tasks for an Active-Passive configuration.

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

Table 2-1 Active-Passive configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	servers
EVDG	cluster disk group
EV_MSMQ_DATA	volume for MSMQ data
EV_MSMQ_LOG	volume for MSMQ log
EV_DATASTORE1	additional volume(s) for storing EV data, as appropriate for your needs
MSMQ_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the EV Server
CLUS1	EV cluster (if the cluster is not already created for SQL Server)
EV-VS	EV virtual server
EV_SG	EV service group

IP addresses for sample Active-Passive configuration

In addition to preparing the names you want to assign the Active-Passive configuration objects, for an IPv4 network, you should obtain all required IP addresses before beginning configuration. For an IPv6 network, IP addresses are generated during configuration.

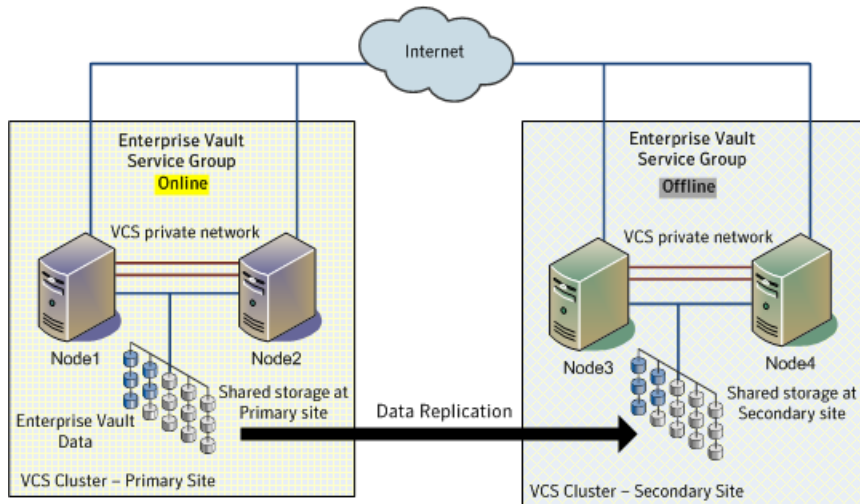
Each EV virtual server requires its own virtual IP address. In the sample configuration there is one EV virtual server. Therefore you would need one virtual server IP address. If you want to use the VCS notification service, you require a cluster IP address. The cluster IP address is also used by the Global Cluster Option for disaster recovery.

Reviewing the disaster recovery configuration

You may be preparing to configure both a primary site and a secondary site for disaster recovery.

The following table illustrates a typical Active-Passive disaster recovery configuration.

Figure 2-2 Typical DR configuration



In the example, the primary site consists of two nodes, Node1 and Node2. Similarly the secondary setup consists of two nodes, Node3 and Node4. Each site has a clustered setup with the nodes set up appropriately for failover within the site.

If the Enterprise Vault server on Node1 fails, Enterprise Vault comes online on node Node2 and 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. If there is a disaster at the primary site, Node3 at the secondary site takes over.

The cluster on the primary site has a shared disk group that is used to create the volumes required by Volume Replicator for setting up the Replicated Volume Group (RVG). The application data is stored on the volumes that are under the control of the RVG.

Sample disaster recovery configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

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

Table 2-2 Sample Disaster Recovery configuration objects

Object Name	Description
Primary site	
SYSTEM1 & SYSTEM2	first and second nodes of the primary site
EVDG	cluster disk group
EV_MSMTQ_DATA	volume for MSMQ data
EV_MSMTQ_LOG	volume for MSMQ log
EV_DATASTORE1	additional volume(s) for storing EV data, as appropriate for your needs
MSMQ_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the EV Server
CLUS1	EV cluster (if the cluster is not already created for SQL Server)
EV-VS	EV virtual server
EV_SG	EV service group
Secondary site	
SYSTEM3 & SYSTEM4	First and second nodes of the secondary site
	All the other parameters are the same as on the primary site.
DR Components (Volume Replicator only)	
EV_RDS	RDS Name
EV_RVG	RVG Name
EV_RVG_SG	Replication service group

IP addresses for disaster recovery configuration

In addition to preparing the names you want to assign configuration objects, for an IPv4 network, you should obtain all required IP addresses before beginning configuration. For an IPv6 network, IP addresses are generated during configuration.

You specify the following addresses during the replication process:

virtual server IP address	For a disaster recovery configuration, the virtual IP address for the virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.
Cluster IP address	You need one for the primary site cluster and one for the secondary site cluster.
Replication IP address	You need two IP addresses per application instance, one for the primary site and one for the secondary site.

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*.

High availability (HA) configuration (New Server)

The following table outlines the high-level objectives and the tasks to complete each objective for an Active-Passive configuration.

Note: Veritas recommends as a best practice to configure SQL Server for high availability before configuring Enterprise Vault for high availability. Configuring SQL Server for high availability is covered in the SQL Server solutions guides.

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>

Table 2-3 Enterprise Vault Server: Active-Passive configuration tasks

Action	Description
Review the HA configuration	<ol style="list-style-type: none">1 Understand active-passive configuration2 Review the sample configuration See “Reviewing the HA configuration” on page 16.
Configure the storage hardware and network	<ol style="list-style-type: none">1 Set up the storage hardware for a cluster environment2 Verify the DNS entries for the systems on which Enterprise Vault Server will be installed
Review pre-requisites and install InfoScale Enterprise	◆ Install InfoScale Enterprise on all the systems where you want to configure EV 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 31.

Table 2-3 Enterprise Vault Server: Active-Passive configuration tasks
(continued)

Action	Description
Configure disk groups and volumes for Enterprise Vault Server	<ol style="list-style-type: none"> 1 Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) 2 Create dynamic volumes for the MSMQ data, registry replication data, and EV services data <p>See “Configuring cluster disk groups and volumes for Enterprise Vault” on page 36.</p>
Configure VCS cluster	<p>If the cluster has not already been configured for SQL Server:</p> <ol style="list-style-type: none"> 1 Verify static IP addresses and name resolution configured for each node 2 Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster <p>See “Configuring the cluster” on page 44.</p>
Install Enterprise Vault on the cluster nodes	<p>◆ Ensure that the appropriate amount of local storage space is available on the node. This is required for storing temporary files during Enterprise Vault installation.</p> <p>Refer to the Enterprise Vault documentation for installation instructions</p>
Create an Enterprise Vault service group	<ol style="list-style-type: none"> 1 Ensure that you have met the prerequisites 2 Ensure that the disk group and volumes for the various Enterprise Vault components are mounted on the first node 3 Create a EV service group using the Enterprise Vault Cluster Setup Wizard 4 Bring the EV service group online on the first node <p>See “Configuring the Enterprise Vault service group” on page 87.</p>

Table 2-3 Enterprise Vault Server: Active-Passive configuration tasks
(continued)

Action	Description
Configure fast failover for disk groups (optional)	<ol style="list-style-type: none"> 1 Ensure that you have installed the Fast Failover option and met the prerequisites for storage 2 Use the Java Console to enable the FastFailover attribute for VMDg resources <p>See “Enabling fast failover for disk groups (optional)” on page 92.</p>
Configure Enterprise Vault for the cluster environment on the first node	<ol style="list-style-type: none"> 1 Launch the Enterprise Vault Configuration Wizard on the first node 2 Choose the option to create a new Enterprise Vault server with cluster support 3 Complete running the wizard on the first node <p>See “Configuring Enterprise Vault Server in a cluster environment” on page 93.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
Configure Enterprise Vault for the cluster environment on any additional nodes	<ol style="list-style-type: none"> 1 Bring the EV service group online on the first node 2 Launch the Enterprise Vault Configuration Wizard on the second node 3 Choose the option to add the node as a failover node for an existing clustered server 4 Complete running the wizard on the second node 5 Repeat these steps for any additional nodes in the EV cluster <p>See “Configuring Enterprise Vault Server in a cluster environment” on page 93.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
Perform additional configuration steps for Enterprise Vault	<p>See “Setting up Enterprise Vault” on page 96.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
(Optional) Configure the appropriate service group dependencies	<p>Configure the appropriate service group dependencies.</p> <p>See “Verifying the Enterprise Vault cluster configuration” on page 95.</p>

Table 2-3 Enterprise Vault Server: Active-Passive configuration tasks
(continued)

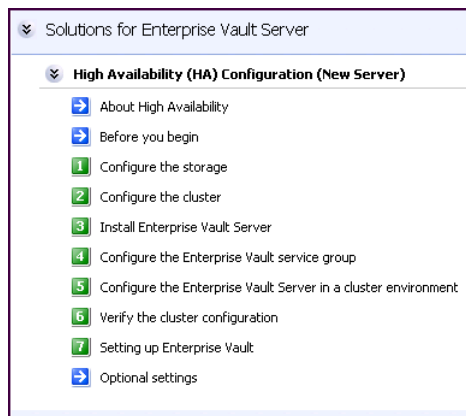
Action	Description
Verify the HA configuration	Test failover between nodes. See “Verifying the Enterprise Vault cluster configuration” on page 95.

Following the HA workflow in the Solutions Configuration Center

The Solutions Configuration Center helps you through the process of configuring HA for Enterprise Vault.

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

Figure 2-3 Configuration steps in the Solutions Configuration Center



See [“Workflows in the Solutions Configuration Center”](#) on page 86.

Disaster recovery configuration

For configuring disaster recovery, you first begin by configuring the primary site for high availability. After setting up an high availability environment for Enterprise Vault (EV) on a primary site, you can create a secondary or “failover” site for disaster recovery.

Note: Veritas recommends as a best practice to configure SQL Server for disaster recovery before configuring Enterprise Vault for disaster recovery. Configuring SQL Server for disaster recovery is covered in the SQL Server solutions guides.

The Disaster Recovery (DR) wizard helps you to configure the storage, Volume Replicator, and the global cluster on the secondary site.

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 81.

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 tasks for configuring the primary site for disaster recovery.

Table 2-4 Outlines the high-level tasks for configuring the primary site for disaster recovery.

Action	Description
Configure the storage hardware and network	For all nodes in the cluster: 1 Set up the storage hardware for a cluster environment 2 Verify the DNS entries for the systems on which EV will be installed See “Configuring the storage hardware and network” on page 35.
Review pre-requisites and install InfoScale Enterprise	Review prerequisites and install InfoScale Enterprise on all the systems where you want to configure high availability for EV. See <i>Veritas InfoScale Installation and Upgrade Guide</i> .
Review application-specific requirements	See “Notes and recommendations for cluster and application configuration” on page 31.

Table 2-4 Outlines the high-level tasks for configuring the primary site for disaster recovery. (*continued*)

Action	Description
Configure the cluster	<p>If the cluster has not already been configured for SQL Server:</p> <ol style="list-style-type: none"> 1 Verify static IP addresses and name resolution configured for each node 2 Configure cluster components using the Cluster Server Configuration Wizard (VCW) 3 Set up secure communication for the cluster <p>See “Configuring the cluster” on page 44.</p>
Configure cluster disk groups and volumes for Enterprise Vault	<ol style="list-style-type: none"> 1 Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) 2 Create dynamic volumes for the MSMQ data, registry replication data, and EV services data <p>See “Configuring cluster disk groups and volumes for Enterprise Vault” on page 36.</p>
Install Enterprise Vault on the cluster nodes	<p>◆ Ensure that the appropriate amount of local storage space is available on the first cluster node. This is required for storing temporary files during Enterprise Vault installation.</p> <p>Refer to the Enterprise Vault documentation for installation instructions</p>
Create an Enterprise Vault service group	<p>◆ Ensure that you have met the prerequisites</p> <ol style="list-style-type: none"> 1 Ensure that the disk group and volumes for the various Enterprise Vault components are mounted on the first node 2 Create a EV service group using the Enterprise Vault Cluster Setup Wizard 3 Bring the EV service group online on the first node <p>See “Configuring the Enterprise Vault service group” on page 87.</p>

Table 2-4 Outlines the high-level tasks for configuring the primary site for disaster recovery. (*continued*)

Action	Description
Configure Enterprise Vault for the cluster environment on the first node	<ul style="list-style-type: none"> ◆ Launch the Enterprise Vault Configuration Wizard on the first node ◆ Choose the option to create a new Enterprise Vault server with cluster support ◆ Complete running the wizard on the first node <p>See “Configuring Enterprise Vault Server in a cluster environment” on page 93.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
Configure Enterprise Vault for the cluster environment on any additional nodes	<ul style="list-style-type: none"> ◆ Bring the EV service group online on the first node <ol style="list-style-type: none"> 1 Launch the Enterprise Vault Configuration Wizard on the second node 2 Choose the option to add the node as a failover node for an existing clustered server 3 Complete running the wizard on the second node 4 Repeat these steps for any additional nodes in the EV cluster <p>See “Configuring Enterprise Vault Server in a cluster environment” on page 93.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
Perform additional configuration steps for Enterprise Vault	<p>See “Setting up Enterprise Vault” on page 96.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
(Optional) Configure the appropriate service group dependencies	<p>Configure the appropriate service group dependencies</p> <p>See “Verifying the Enterprise Vault cluster configuration” on page 95.</p>
Verify the primary site configuration	<p>Test failover between nodes on the primary site</p> <p>See “Verifying the Enterprise Vault cluster configuration” on page 95.</p>

DR configuration tasks: Secondary site

The following table outlines the high-level tasks for configuring the secondary site for disaster recovery.

Table 2-5 Configuring the secondary site for disaster recovery

Action	Description
Install InfoScale Enterprise and configure the cluster on the secondary site	Caution: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.
Verify that Enterprise Vault has been configured for high availability at the primary site	Verify that Enterprise Vault has been configured for high availability at the primary site and that the service group is online See “Verifying your primary site configuration” on page 49.
Set up security for Volume Replicator	Ensure that you have completed setting up Volume Replicator security before running the DR wizard See “Setting up security for Volume Replicator” on page 51.
(Secure cluster only) Assign user privileges	For a secure cluster only, assign user privileges See “Assigning user privileges (secure clusters only)” on page 53.
Start running the DR wizard	<ol style="list-style-type: none">1 Review prerequisites for the DR wizard2 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 54.
Clone the storage configuration (Volume Replicator only)	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 58.
Install Enterprise Vault on the cluster nodes	◆ Ensure that the appropriate amount of local storage space is available on the first cluster node. This is required for storing temporary files during Enterprise Vault installation. Refer to the Enterprise Vault documentation for installation instructions

Table 2-5 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
Create an Enterprise Vault service group	<ol style="list-style-type: none"> 1 Ensure that you have met the prerequisites 2 Ensure that the disk group and volumes for the various Enterprise Vault components are mounted on the first node 3 Create a EV service group for the secondary site using the same service group name, virtual server name, and configuration as on the primary site <p>See “Installing and configuring Enterprise Vault on the secondary site” on page 62.</p>
Configure Enterprise Vault for the cluster environment on the first node	<ol style="list-style-type: none"> 1 Launch the Enterprise Vault Configuration Wizard on the first node 2 Choose the option to create a new Enterprise Vault server with cluster support 3 Complete running the wizard on the first node <p>See “Installing and configuring Enterprise Vault on the secondary site” on page 62.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
Configure Enterprise Vault for the cluster environment on any additional nodes	<ol style="list-style-type: none"> 1 Bring the EV service group online on the first node 2 Launch the Enterprise Vault Configuration Wizard on the second node 3 Choose the option to add the node as a failover node for an existing clustered server 4 Complete running the wizard on the second node 5 Repeat these steps for any additional nodes in the EV cluster <p>See “Installing and configuring Enterprise Vault on the secondary site” on page 62.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>
Perform additional configuration steps for Enterprise Vault	<p>See “Setting up Enterprise Vault” on page 96.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>

Table 2-5 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
Configure replication and global clustering	Use the DR wizard to configure Volume Replicator replication and global clustering See “ Configuring Volume Replicator replication and global clustering ” on page 64.
(Optional) Configure the appropriate service group dependencies	Configure the appropriate service group dependencies See “ Setting service group dependencies for disaster recovery ” on page 74.
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 74.
(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 76.
(Optional) Add additional DR sites	Optionally, add additional DR sites to a Volume Replicator environment See “ Adding multiple DR sites (optional) ” on page 78.
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 78.

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:

- One static IP address per site for each Enterprise Vault virtual server.
- 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.
- 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.
- 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	<p>The following types of IPv6 addresses are supported:</p> <ul style="list-style-type: none"> ■ 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	<p>LLT over UDP is supported on both IPv4 and IPv6.</p> <p>You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.</p>
VCS agents, wizards, and other components	<p>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).</p> <p>The Veritas High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.</p>

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

Configuring the storage hardware and network

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

To configure the hardware

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

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel by clicking **Start > Control Panel**.
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, click **Adapter settings**.
- 4 Ensure 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 the public adapter is the first adapter in the **Connections** list. If necessary, use the arrow button to move the adapter to the top of the list.

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

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
 - 6 In the Public Status dialog box, on the General tab, click **Properties**.
 - 7 In the Public Properties dialog box, on the General tab, select the 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 that the **Register this connection's address in DNS** check box is selected.
 - 12 Make sure that the correct domain suffix is entered in the **DNS suffix for this connection** field.
 - 13 Click **OK**.

Configuring cluster disk groups and volumes for Enterprise Vault

Before configuring Enterprise Vault for high availability, you must create cluster disk groups and volumes using the Veritas Enterprise Administrator (VEA) console.

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

See [“About cluster disk groups and volumes”](#) on page 37.

See [“Prerequisites for configuring cluster disk groups and volumes”](#) on page 37.

See [“Considerations for a fast failover configuration”](#) on page 38.

See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 39.

See [“Sample disk group and volume configuration”](#) on page 40.

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

See [“Viewing the available disk storage”](#) on page 40.

See [“Creating a cluster disk group”](#) on page 41.

See [“Creating Volumes”](#) on page 41.

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.

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

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

Prerequisites for configuring cluster disk groups and volumes

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

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

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

- Determine the layout or configuration for each volume and the total number of disks needed.

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

For a fast failover configuration, See [“Considerations for a fast failover configuration”](#) on page 38.

For a disaster recovery configuration using Volume Replicator, See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 39.

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 disks and volumes for campus clusters

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

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

Consider the following when creating new volumes:

- For campus clusters, when creating a new volume, you must select the "mirrored across enclosures" option.
- Choosing "Mirrored" and the "mirrored across" option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.
- 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.
- You cannot selecting RAID-5 for mirroring.
- Selecting "stripe across enclosures" is not recommended because then you need four enclosures, instead of two.

Considerations for volumes for a 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)

- Data Change Object (DCO)
- Volumes with commas in the names
- 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 an SFW HA solution, you first create a cluster disk group (EVDG) on shared disks and then create volumes for the following:

- MSMQ data
- Registry replication data
- Various EV services data (Indexing service, Shopping service, Vault store partitions, PST holding folders, etc.)

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 Launch **Veritas Enterprise Administrator** from the **Apps** menu on the **Start** screen. Select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.

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

- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.

The internal names for the disks that the current system can access for available storage are displayed, with names Harddisk1, Harddisk2, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a cluster disk group

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

Creating Volumes

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

Before you begin, review the following topic if applicable to your environment:

- See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 39.

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 41.
- See [“Unmounting a volume and deporting a disk group”](#) on page 42.

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**.

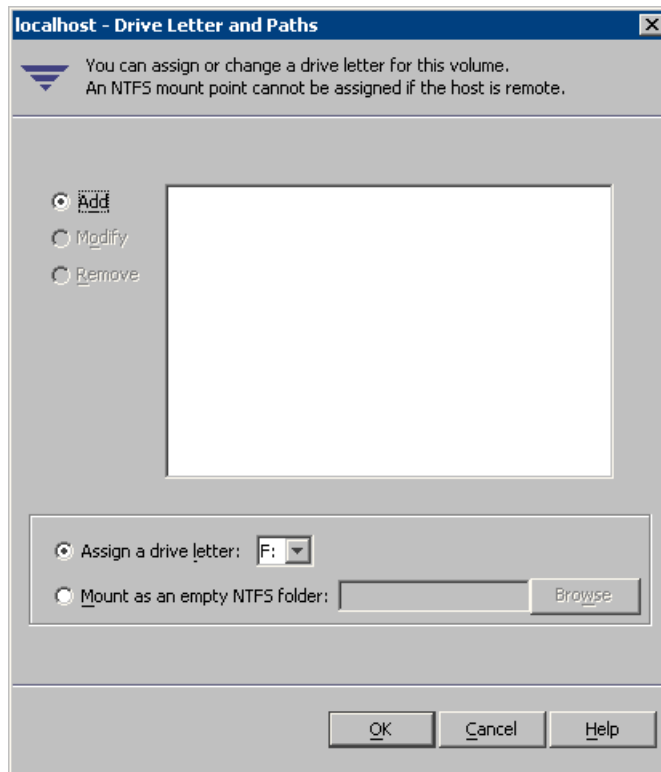
Adding drive letters to mount the volumes

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

To move ownership of the cluster disk group to another node, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node and then import it to the desired node.

To deport the cluster disk group

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

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses and that name resolution is configured for each node.
- Verify that you have the required privileges.
See [“Notes and recommendations for cluster and application configuration”](#) on page 31.

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

Adding a node to an existing VCS cluster

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

Prerequisites for adding a node to an existing cluster 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.

The VCS Cluster Configuration Wizard (VCW) configures VCS components and starts VCS services on the new node. The wizard does not configure any service groups on the new node.

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.
- InfoScale Enterprise is either not installed on the system or the version 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**.

Verifying your primary site configuration

Before you begin configuring disaster recovery, make sure that Enterprise Vault has been configured for high availability at the primary site.

If you have not yet configured Enterprise Vault 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 and that the service group is online.

Guidelines for installing InfoScale Enterprise and configuring the cluster on the secondary site

Use the following guidelines for installing InfoScale Enterprise and configuring the cluster on the secondary site.

Note: Veritas recommends as a best practice to configure SQL Server for disaster recovery before configuring Enterprise Vault for disaster recovery. If you have completed SQL Server DR configuration, the following steps may already be complete.

- Ensure that you have set up the components required to run a cluster.
- Ensure that when installing InfoScale Enterprise you install the appropriate disaster recovery options at both the primary and secondary sites, as follows:

Global Cluster Option	This InfoScale Availability option is required for a disaster recovery configuration.
Volume Replicator	This InfoScale Storage option is required for Volume Replicator replication.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

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](#)” on page 44.

Note: You do not need to configure the GCO option while configuring the cluster. This is done later using the Disaster Recovery wizard.

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

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 72.

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:

See “[Setting up security for Volume Replicator](#)” on page 51.

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 **VVR Security Service Configuration Wizard** from the **Apps** menu on 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	<p>The Available domains pane lists all the domains that are present in the Windows network neighborhood.</p> <p>Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.</p>
Adding a domain	<p>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.</p>

Click **Next**.

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

Selecting hosts	<p>The Available hosts pane lists the hosts that are present in the specified domain.</p> <p>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.</p>
Adding a host	<p>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.</p>

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.

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 Enterprise Vault 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 application service group.

```
hauser -add user [-priv <Administrator|Operator>  
[-group service_groups]]
```

- 4 Reset the configuration to read-only:

```
haconf -dump -makero
```

To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:

```
haconf -makerw
```

- 2 Add the user. Specify the name in the format `username@domain`.

```
hauser -add user [-priv <Administrator|Operator>]
```

- 3 Reset the configuration to read-only:

```
haconf -dump -makero
```

Configuring disaster recovery with the DR wizard

In an Enterprise Vault environment, the Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration
- Configure Volume Replicator replication and global clustering

You will need to exit the wizard after the storage cloning task to install the application, configure the Enterprise Vault service group, and configure Enterprise Vault for the cluster environment. Then you start the wizard again.

The DR Wizard list of service groups shows only those that 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.
- Enterprise Vault is configured for HA at the primary site and the EV service group is online on the primary site.
- Enough free disk space is available at the secondary site to duplicate the storage configuration at the primary site.
- One static IP address is available per application service group to be created.
- A minimum of one static IP address per site is available for each application instance running in the cluster.
- Global Cluster Option (GCO) is installed at the primary and secondary site, and one static IP address is available at each site for configuring GCO.
- The service group to be cloned can use either IPv4 IP addresses or IPv6 addresses but not a mixture of both.
- 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 prerequisite:

See [“Setting up security for Volume Replicator”](#) on page 51.

To start configuring disaster recovery with the DR wizard

- 1 Start the **Solutions Configuration Center** from the **Apps** menu on the **Start** screen. Expand the Solutions for Enterprise Vault tab and click **Disaster Recovery Configuration > Clone the storage on the secondary site> Disaster Recovery Configuration Wizard**.

Note: By design, the DR wizard requires specific settings for the Lanman attributes on the primary and secondary sites. Before beginning the DR configuration, the wizard checks for these values, and if they are not set as required, the wizard will automatically proceed with setting these values, both at the primary and secondary sites.

- 2 In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.
- 3 In the System Selection panel, complete the requested information:

System Name

Enter the IP address or Fully Qualified Host Name (FQHN) of the primary system where the application is online.

If you have launched the wizard on the system where the application is online at the primary site, you can also specify `localhost` to connect to the system.

- 4 In the Service Group Selection panel, select the Enterprise Vault service group for which you want to configure the storage and replication.

The panel lists only service groups that contain a MountV resource. The service group must not have a child service group, since the DR wizard does not support such a configuration for Volume Replicator replication.

Click **Next**.

- 5** In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.

Click **Next**.

- 6 In the Replication Options panel, select the replication method. Although you must select the replication method now, configuring replication and the global cluster option is done later, after service group creation. For Enterprise Vault, select the option to configure Volume Replicator and the Global Cluster option (GCO).

Configure 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.

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**.

7 Continue with cloning the storage.

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:

See [“Configuring disaster recovery with the DR wizard”](#) on page 54.

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	<div>Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary.<ul style="list-style-type: none">■ 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.</div>

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	<p>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.</p> <p>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.</p>
-----------------	--

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	<p>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.</p> <p>Select disks for each unavailable volume that you want to clone on to the secondary.</p>
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 When the Application Installation panel is displayed, click **Finish** to exit the wizard.
- 9 You must complete the following tasks for installing and configuring Enterprise Vault on the secondary site before you restart the Disaster Recovery Wizard:
 - Install Enterprise Vault on the secondary site nodes.
 - Run the Enterprise Vault Cluster Setup Wizard on the first node on the secondary site to configure the Enterprise Vault service group.
 - Run the Enterprise Vault Configuration Wizard on each node on the secondary site to configure Enterprise Vault for the cluster environment.

Installing and configuring Enterprise Vault on the secondary site

Perform the following steps when installing and configuring Enterprise Vault for the cluster on the secondary site.

Be sure to read these instructions before running the Enterprise Vault Cluster Setup wizard and the Enterprise Vault Configuration Wizard on the secondary site.

To install and configure Enterprise Vault on the secondary site:

- 1 Install Enterprise Vault on each node of the secondary site cluster. For installation and configuration instructions, see the Enterprise Vault documentation.
- 2 Before you launch the wizard to configure the EV service group, storage cloning must be complete. Verify that the cluster disk group is imported to the first node on the secondary site and the volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the DR Wizard does not mount the volumes on the secondary site and you must format the volumes and mount them manually.

See [“About managing disk groups and volumes”](#) on page 41.
- 3 Bring the EV service group offline on the primary site.

- 4 On the first node, launch the Enterprise Vault Cluster Setup Wizard to configure the EV service group.

Launch the **Solutions Configuration Center** from the **Apps** menu on the **Start** screen. Expand the Solutions for Enterprise Vault tab and click **Disaster Recovery Configuration > Configure the Enterprise Vault service group on the secondary site > Enterprise Vault Cluster Setup Wizard**.

Specify the same service group name and virtual server name on the secondary site as on the primary site.

For example, if the service group name on the primary site is EV_SG, use EV_SG for the service group name on the secondary site. If the virtual server name on the primary site is EV-VS, use EV-VS for the virtual server name on the secondary site.

Specify the same MSMQ and Replication Directory paths on the secondary site as on the primary site.

Use the same procedure as when configuring the service group on the primary site. Be sure to choose the wizard option to bring the service group online after creating it.

See [“Configuring the Enterprise Vault service group”](#) on page 87.

- 5 Use the Enterprise Vault Configuration Wizard to configure Enterprise Vault on each node on the secondary site, beginning with the node on which the service group is online (the first node).

Launch the Solutions Configuration Wizard. Expand the Solutions for Enterprise Vault tab and click **Disaster Recovery Configuration > Configure Enterprise Vault Server in a cluster environment > Enterprise Vault Configuration Wizard**.

Follow these guidelines on the first node on the secondary site:

- On the first node, select the option to create a new Enterprise Vault Server with cluster support, rather than the option to add the node as a failover node for an existing clustered server.
- Select the Enterprise Vault service group that you just created as the group in which to configure the resources for the Enterprise Vault services.
- When prompted to enter the SQL Server that you want to use for the Enterprise Vault database, specify the same database that you selected on the primary site. If SQL Server has been configured for high availability and disaster recovery, enter the name in the format `virtualservername\instancename`. For example, if the SQL virtual server name is virtualsvr and the instance name is EVSQL, enter virtualsvr\EVSQL. The wizard detects the existing database and updates it.

- When you finish running the wizard, bring the Enterprise Vault resources online and verify that the Enterprise Vault service group is online. The EV resources will be in an unknown state on the failover node because the Enterprise Vault Configuration Wizard has not yet been run on that node. For additional information on the wizard, see the Enterprise Vault documentation.
- 6 To configure Enterprise Vault on any additional node, keep the service group online on the first node. Launch the Enterprise Vault Configuration Wizard from the additional node. On the additional node, make sure that you select the option to add the node as a failover node for an existing clustered server.
 - 7 Once configuration is complete, verify the configuration on the secondary site by switching the service group from the first node to the failover node on the secondary site.
 - 8 Take the service group offline on the secondary site and bring it online on the primary site.

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 51.
- 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 DR wizard and proceed to the Replication Setup panel as follows:

- Start the **Solutions Configuration Center** from the **Apps** menu on the Start screen.
 - Expand the Solutions for Enterprise Vault tab and click **Disaster Recovery Configuration > Configure replication and the global cluster option (GCO) > 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**.

The wizard proceeds to the storage cloning panel. If it detects that the storage is identical on the secondary site, it proceeds to the next task. If it detects that the service group is created on the secondary site, it proceeds to the Internet Protocol panel.

- 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	<p>Displays the list of available volumes that have not been selected to be a part of the RVG.</p> <p>Either double-click on the volume name or use the > button to move the volumes into the Selected RVG Volumes pane.</p>
Selected RVG Volumes	<p>Displays the list of volumes that have been selected to be a part of the RVG.</p> <p>To remove a selected volume, either double-click the volume name or use the < button to move the volumes into the Available Volumes pane.</p>
Primary SRL	<p>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.</p> <p>Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>
Secondary SRL	<p>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.</p> <p>Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>
Start Replication after the wizard completes	<p>Select this check box to start replication automatically after the wizard completes the necessary configurations.</p> <p>Once replication is configured and running, deselecting the checkbox does not stop replication.</p>

Click **Advanced Settings** to specify some additional replication properties.

Advanced Replication Settings

Advanced Replication Settings for RVG_TESTFS_0

Replication Mode:	Synchronous Override	Protocol:	UDP
Log Protection:	AutoDCM	Packet Size (Bytes):	1400
Primary RLINK Name:	48326630361624	Latency Protection:	Fail
Secondary RLINK Name:	48326630361623	High Mark Value:	10000
Bandwidth:	Maximum	Low Mark Value:	9950
	<input type="checkbox"/> Mbps	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 | <p>Select the appropriate log protection from the list:</p> <ul style="list-style-type: none"> ■ 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	<p>By default, Volume Replicator replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit.</p> <p>The default unit is Megabits per second (Mbps) and the minimum allowed value is 1 Mbps.</p>
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	<p>By default, latency protection is set to Off.</p> <p>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.</p> <p>This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail option when the Secondary is connected.</p>
High Mark Value	<p>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.</p> <p>To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.</p>
Low Mark Value	<p>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.</p>

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 global clustering only

You have the option to use the DR wizard to configure global clustering (GCO) only.

Before configuring GCO:

- One static IP address must be available per site for configuring GCO.
- If you created secure clusters at the primary site and secondary site, ensure that you have configured a VCS user with the same name and privileges in each cluster, and the user must be added in the Administrator role.

The following procedure assumes that you have already configured the storage and installed and configured EV on the secondary site.

To configure GCO only

- 1 Launch the wizard and proceed to the GCO Setup panel as follows:
- 2 Launch the **Solutions Configuration Center** from the **Apps** menu on the **Start** screen. Expand the Solutions for Enterprise Vault tab and click **Disaster Recovery Configuration > Configure replication and the global cluster option (GCO) > 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 Global Cluster Option (GCO) only**. Click **Next** and continue to the GCO Setup panel.
- 5 In the GCO Setup panel, review the requirements. If you have met the requirements, 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	<p>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.</p> <p>Once GCO is configured and running, deselecting the checkbox does not stop GCO.</p>

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified. 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.

Setting service group dependencies for disaster recovery

Since Enterprise Vault requires the SQL database, if you cluster SQL Server using Cluster Server, Veritas recommends setting a service group dependency between the EV service group and the SQL service group.

In the disaster recovery environment with Volume Replicator replication, the EV Replicated Volume Group (RVG) should be linked as a parent to the SQL Server service group as a child. Set the same service group dependencies on the primary and secondary site clusters.

For the Volume Replicator environment, the DR wizard can configure DR only for a service group that has no child. Therefore, set service group dependencies only after running the DR wizard.

For more information on setting service group dependencies, see the Cluster Server Administrator's Guide.

Verifying the disaster recovery configuration

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- 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.
- 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.
- 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.
See [“Setting service group dependencies for disaster recovery”](#) on page 74.
- 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.
- 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 typically consists of
 - 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.

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 21.

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 2-6 Online, local, soft dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	<ul style="list-style-type: none"> ■ 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. 	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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. 	<ul style="list-style-type: none"> ■ 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.

Table 2-7 Online, local, firm dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	<ul style="list-style-type: none"> ■ 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. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Leave the RVG group online at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ 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. 	<ul style="list-style-type: none"> ■ 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 2-8 Online, local, hard dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	<ul style="list-style-type: none"> ■ 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. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Do not take the RVG group offline at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ 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. 	<ul style="list-style-type: none"> ■ 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 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
- Microsoft SQL Server
- Enterprise Vault Server (high availability and disaster recovery solutions)
- Microsoft SharePoint Server (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

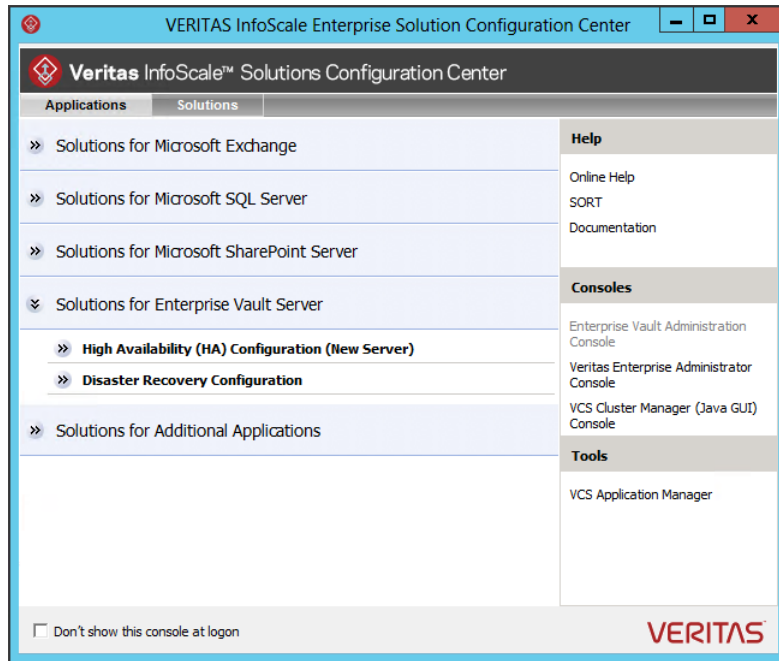
Navigate to the Apps menu and then click **SCC**.

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 Enterprise Vault Server.

Figure 3-1 Solutions Configuration Center for Enterprise Vault Server



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	<p>(For advanced users) Lists wizards by solution, without additional instructions, as follows:</p> <ul style="list-style-type: none"> ■ High Availability Configuration Wizards ■ Disaster Recovery Configuration Wizards ■ Quick Recovery Configuration Wizards ■ Fire Drill Configuration Wizards <p>You can go directly to a particular wizard.</p> <p>Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.</p> <p>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.</p> <p>Refer to the following Microsoft knowledge database article for more details:</p> <p>http://technet.microsoft.com/en-us/library/dd184075.aspx</p>
-----------	---

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	<p>Configures wide area disaster recovery, including cloning storage, cloning service groups, and configuring the global cluster</p> <p>Can also configure:</p> <ul style="list-style-type: none"> ■ Volume Replicator (Volume Replicator) replication ■ VCS resource for EMC SRDF and Hitachi TrueCopy array-based hardware replication <p>Note: Requires first configuring high availability on the primary site.</p> <p>To configure IPv6 settings, launch the wizard from a system on which the IPv6 stack is installed.</p>
Fire Drill Wizard	<p>Sets up a fire drill to test disaster recovery</p> <p>Note: Requires first configuring high availability on the primary site.</p> <p>To configure IPv6 settings, launch the wizard from a system on which the IPv6 stack is installed.</p>

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
Enterprise Vault Cluster Setup Wizard	Configures the service group for Enterprise Vault Server high availability
Enterprise Vault Cluster Wizard	Configures Enterprise Vault Server in a cluster environment
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

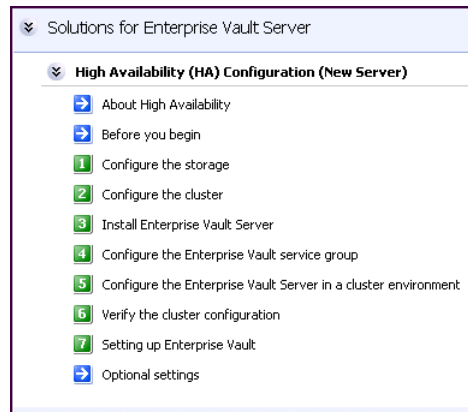
Workflows in the Solutions Configuration Center

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 Enterprise Vault Server from the Solutions Configuration Center.

Figure 3-2 Workflow for configuring high availability for Enterprise Vault Server



Installing and configuring Enterprise Vault for failover

This chapter includes the following topics:

- [Installing Enterprise Vault](#)
- [Configuring the Enterprise Vault service group](#)
- [Configuring Enterprise Vault Server in a cluster environment](#)
- [Setting service group dependencies for high availability](#)
- [Verifying the Enterprise Vault cluster configuration](#)
- [Setting up Enterprise Vault](#)
- [Considerations when modifying an EV service group](#)

Installing Enterprise Vault

Install Enterprise Vault on the cluster nodes.

For installation instructions, see the Enterprise Vault documentation.

Configuring the Enterprise Vault service group

Before you configure Enterprise Vault in an SFW HA cluster environment, you must configure a service group to represent the Enterprise Vault server.

This section describes how to configure an Enterprise Vault (EV) Server service group using the Enterprise Vault Cluster Setup Wizard.

The following topics describe how to configure an Enterprise Vault (EV) Server service group using the Enterprise Vault Cluster Setup Wizard:

- See [“Before you configure an EV service group”](#) on page 88.
- See [“Creating an EV service group”](#) on page 88.
- See [“Enabling fast failover for disk groups \(optional\)”](#) on page 92.

Before you configure an EV service group

Before you configure an Enterprise Vault service group, do the following:

- Verify that you have installed the Enterprise Vault Cluster Setup Wizard. If you have not installed it during InfoScale Enterprise installation, use Windows Add or Remove Programs to install it.
- Verify that you have configured a cluster using the VCS Cluster Configuration Wizard (VCW).
- Verify your DNS server settings. You must ensure that a static DNS entry maps the virtual IP address with the virtual server name. Refer to the appropriate DNS document for more information.
- Verify that the Veritas High Availability Engine (HAD) is running on the system from where you will run the Enterprise Vault Cluster Setup Wizard.
- Ensure that you have Cluster Administrator privileges. You must also be a Local Administrator on the node where you run the wizard.
- If you have configured a firewall, add the required ports and services to the Firewall Exception list.
For a detailed list of services and ports used by the InfoScale product, refer to the *Veritas InfoScale Installation and Upgrade Guide*.
- Verify that MSMQ is installed locally on each node that will be part of the EV service group.
- Ensure that you mount the shared volumes you created for EV on the node from which you will run the wizard and unmount the volumes from other nodes in the cluster.
See [“About managing disk groups and volumes”](#) on page 41.

Creating an EV service group

Complete the following steps to create a service group for EV.

To create the EV service group

- 1
- Start the **Enterprise Vault Cluster Setup Wizard** from the **Apps** menu on the **Start** screen.
- Alternatively, launch the **Solutions Configuration Center** from the **Apps** menu on the **Start** screen, expand the Solutions for Enterprise Vault Server tab, and click **High Availability (HA) Configuration (New Server) > Configure the Enterprise Vault Service Group > Enterprise Vault Cluster Setup Wizard**.
- 2
- Review the information in the Welcome panel and click **Next**.
- 3
- On the Wizard Options panel click **Create service group** and then click **Next**.
- 4
- On the Service Group Configuration panel, specify the service group details and then click **Next**

Service Group Name	Type a name for the EV service group.
Available Cluster Systems	<p>Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.</p> <p>To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.</p> <p>To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.</p> <p>System priority defines the order in which service groups are failed over to systems. 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.</p>
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox. For information about the AutoStartList attribute, see the Cluster Server Administrator's Guide.

- 5
- On the Virtual Server Configuration panel, specify information related to your network and then click **Next**.

- Select `IPv4` to configure an IPv4 address for the EV virtual server.
 - In the Virtual IP Address field, type a unique virtual IPv4 address for the EV virtual server.
 - In the Subnet Mask field, type the subnet to which the virtual IPv4 address belongs.
- Select `IPv6` to configure an IPv6 address for the EV virtual server. The IPv6 option is disabled if the network does not support IPv6.
 - Select the network from the drop-down list. The wizard uses the network prefix and automatically generates an IPv6 address that is valid and unique on the network.
- In the Virtual Server Name field, type a unique name for the EV virtual server. This is the name by which the EV server will be referenced by clients. The virtual name must not exceed 15 characters. You will need to specify the same name when running the Enterprise Vault Configuration Wizard.
- For each system in the cluster, select the public network adapter name. The Adapter Display Name field displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. Verify that you select the adapters assigned to the public network, not the private.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, specify the desired Organizational Unit in the domain and then click **OK**.

This sets the Lanman resource attributes `ADUpdateRequired` and `ADCriticalForOnline` to true. This allows the Lanman agent to update Active Directory with the SQL virtual server name.

You can type the OU details in the format

```
CN=Computers,DC=domainname,DC=com.
```

To search for the OU, click the ellipsis button and specify the search criteria in the Windows Find Organizational Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
- Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- 6 On the MSMQ and RegRep Directory Details panel, complete the following and then click **Next**:

MSMQ Directory	<p>Type a path or click ... (ellipsis button) to browse for the directory. All MSMQ data is stored at this location.</p> <p>Note: The wizard, by default, also stores the Indexing service and Shopping service data at this location. You can modify the location while configuring the EV Server, later.</p>
Replication Directory	<p>Type a path or click ... (ellipsis button) to browse for the registry replication directory. This directory contains the list of registry keys to replicate.</p> <p>Veritas recommends that you configure the MSMQ and registry replication directories on different volumes.</p>

Note: Make sure that this is a new or empty directory. This directory must not contain data pertaining to an MSMQ service group that was deleted.

- 7 On the Storage Location Details panel, select the volumes that you want to configure for EV data and then click **Next**.
- Select the volumes from the Available Volumes box and then click the right arrow button to move them to the Selected Volumes box. The volumes listed in the Available volumes box do not include the volumes you specified for MSMQ and registry replication.
- 8 On the Service Group Summary panel, review the service group configuration.

Resources	<p>Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.</p> <p>To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.</p>
-----------	--

Attributes	Displays the attributes and their configured values, for a resource selected in the Resources list.
Enable FastFailOver attribute for all the VMDg resources in the service group	<p>To enable all the VMDg resources in the service group for fast failover, select this checkbox.</p> <p>For information about the FastFailOver attribute, see the Cluster Server Administrator's Guide.</p>

■

Click **Next**.

A message appears informing you that the wizard will run commands to modify the service group configuration.

■

Click **Yes**.

The wizard starts running commands to create the service group. Various messages indicate the status of these commands.

9

On the completion dialog box, check **Bring the service group online** check box to bring the EV service group online on the local system, and then click **Finish**.

If you plan to enable fast failover for disk groups, see the following topic:

See [“Enabling fast failover for disk groups \(optional\)”](#) on page 92.

Otherwise, you can proceed to configuring EV using the EV Configuration Wizard.

See [“Configuring Enterprise Vault Server in a cluster environment”](#) on page 93.

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 38.

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.

Configuring Enterprise Vault Server in a cluster environment

The Enterprise Vault Configuration Wizard provides options for setting up Enterprise Vault in a cluster. You must run the wizard on each node of the cluster.

Before running the wizard, do the following:

- Ensure that the Enterprise Vault service group is configured and online on the first node where you run the Enterprise Vault Configuration Wizard.
- Ensure that Enterprise Vault is installed on all additional nodes where you run the wizard.

To configure Enterprise Vault Server in a cluster environment

- 1 Launch the **Solutions Configuration Center** from the **Apps** menu on the **Start** screen. Expand the Solutions for Enterprise Vault tab and click **High Availability (HA) Configuration (New Server) > Configure the Enterprise Vault Server in a cluster environment > Enterprise Vault Configuration Wizard**.
- 2 Use the Enterprise Vault Configuration Wizard to configure Enterprise Vault on the first node. Follow these guidelines for configuring in the EV cluster environment:
 - When configuring EV on the first node, select the option to create a new Enterprise Vault Server with cluster support.
 - Select the Enterprise Vault service group that you just created as the group in which to configure the resources for the Enterprise Vault services.
 - When prompted to enter the SQL Server that you want to use for the Enterprise Vault database, if SQL Server has been configured for high availability and disaster recovery, enter the name in the format `virtualservername\instancename`. For example, if the SQL virtual server name is `virtualsvr` and the instance name is `EVSQ`, enter `virtualsvr\EVSQ`.
 - For the computer alias, use the virtual server name that was specified when creating the Enterprise Vault service group.
 - When you finish running the wizard on the first node, bring the Enterprise Vault resources online and verify that the Enterprise Vault service group is online.
- 3 To configure Enterprise Vault on any additional node, keep the service group online on the first node. Launch the Enterprise Vault Configuration Wizard from the additional node. On the additional node, make sure that you select the option to add the node as a failover node for an existing clustered server.

Refer to the EV documentation for more information on the wizard.

Setting service group dependencies for high availability

Since Enterprise Vault requires the SQL database, if you cluster SQL Server using Cluster Server, Veritas recommends setting a service group dependency between the EV service group and the SQL service group.

For more information on setting service group dependencies, see the Cluster Server Administrator's Guide.

Verifying the Enterprise Vault cluster configuration

Simulating a failover is an important part of configuration testing. After completing the configuration, verify that failover occurs as desired.

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

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

To switch service groups

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

If there is more than one service group, you must repeat this step until all the service groups are switched.

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

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node, perform these steps sequentially:
 - Restart the node that you shut down in the first step.
 - 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.

Setting up Enterprise Vault

Use the Enterprise Vault Administration Console to set up Enterprise Vault for archiving.

Perform the following tasks depending on your environment:

- Add the EV services to the cluster nodes.
- Create retention categories or edit predefined categories to suit your environment.
- Create a default vault store and partition.
- Review the default settings for the site.
- Implement other types of archiving (for example, Exchange, Outlook Web Access, SharePoint Server) as per your requirements.

Make sure to configure EV Server data files on shared storage. Data file include Indexing service data, Shopping service data, Vault store partitions, PsT holding folders and EMC Centera staging areas.

For more information on how to perform these tasks, see the Enterprise Vault documentation.

Considerations when modifying an EV service group

The following table lists the items that you can modify in the EV service group. For more information on modifying service groups, see the Cluster Server Administrator's Guide.

Table 4-1 Items that can be modified in an EV service group

Item	Notes
System list	You can add or remove nodes from the service group SystemList. If you want to remove a node, ensure that you do not run the wizard to modify the service group from that node.

Table 4-1 Items that can be modified in an EV service group (*continued*)

Item	Notes
Volumes	You can add or remove volumes. If you remove a volume on which an Enterprise Vault service is configured, the service ceases to be highly available and is not monitored.
Virtual IP	You can change the virtual IP address if the service group is offline. You cannot change the virtual server name, which is fixed when you create the service group.

Note the following:

- You must run the wizard from a node on which the service group is online. You can then use the wizard to add resources to or remove them from the configuration.
- You must take the service group partially offline to change the resource attributes. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes. Mount all the volumes created to store Storage service data (vault stores), registry replication information, Shopping service data, Indexing data and MSMQ data.
- If you want to modify the system list or volumes, the service group must be online.

Note: If you add a system to an online service group, any resources with local attributes may briefly have a status of unknown. After you add the new node to the group, run the Enterprise Vault Configuration Wizard on this node to configure the Enterprise Vault services for it.
